applied sciences

MDPI

Systematic Review

Systematic Literature Review of AI-Driven Multi-Cloud Anomaly Detection in Zero-Trust Frameworks

Ziad Almulla * and Abdullah Albuali *

Department of Computer Networks and Communications, College of Computer Sciences and Information Technology, King Faisal University, Al-Ahsa 31982, Saudi Arabia
* Correspondence: 225000830@student.kfu.edu.sa (Z.A.); aabuali@kfu.edu.sa (A.A.)

Abstract

Multi-cloud is becoming more challenging to secure as traditional perimeter-based security models have a hard time protecting workloads running across multiple cloud platforms, identities, and services. To address this challenge, organizations are shifting to Zero-Trust Architecture (ZTA), which focuses on constant verification and stringent access control, coupled with anomaly detection methodologies to gain better visibility and threat detection in the distributed cloud environment. This paper presents a Systematic Literature Review (SLR) of anomaly detection approaches in multi-cloud environments and how these are applied in zero-trust security models. The review is conducted according to the guidelines of the 2020 Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020), and is based on studies published between 2020 and 2025 selected from the databases of the following journals: Institute of Electrical and Electronics (IEEE) Xplore, Science Direct, MDPI, Google Scholar, and the Saudi Digital Library. Studies found on benchmark datasets such as CICIDS-2017 and UNSW-NB15 are not evaluated, as none addressed real multi-cloud environments. Although zero trust is highlighted in general, very few studies have implemented basics of zero trust such as micro-segmentation, identity federation, and enforcement through policy. Overall, this review identifies gaps around cross-cloud validation, explainability, and compliance-aware security design, including lack of attention to regulations such as the GDPR and HIPAA. These findings provide helpful recommendations for future research and development on practical and security solutions for multi-cloud environments.

Keywords: zero-trust architecture; federated learning; self-supervised learning; anomaly detection; multi-cloud security; cloud compliance; policy enforcement

Academic Editors: Emil Pricop and Khalil El-Khatib

Received: 7 January 2026
Revised: 23 May 2026
Accepted: 2 June 2026
Published: 12 June 2026

Correction Statement: This article has been republished with a minor change. The change does not affect the scientific content of the article and further details are available within the backmatter of the website version of this article.

1. Introduction

Enterprise cybersecurity has been fundamentally altered by the accelerated movement of mission-critical digital workloads to multi-cloud environments, forcing organizations to reinvent perimeter-centric models of security in favor of context-aware, adaptive, and zero-trust frameworks [1,2].

Zero trust is a security paradigm which enforces strict least-privilege access and advocates radical rethinking about of network protection. Zero trust involves never implicitly trusting anything or anyone ("never trust, always verify") by always verifying users, devices, and services [3].

As artificial intelligence technologies continue to advance faster than traditional security controls, machine learning-based anomaly detection is now a critical technology for

securing dynamic multi-cloud environments. AI-powered anomaly detection technology uses advanced analytics (federated learning, deep neural networks, and self-supervised models) to automatically identify bad behavior, deduce intricate attack patterns, and continuously adapt to new threats in real time while maintaining privacy and regulatory compliance across heterogeneous environments. The combination of AI and zero-trust principles allows for ongoing monitoring, risk-based decision-making, and automated policy enforcement, leading to a proactive defense architecture that is beyond the limitations of perimeter-based security and static access controls [4,5].

This paper explores the convergence of AI-based anomaly detection and zero-trust architectures in multi-cloud environments. The aim is to assess security, performance, and compliance in today's cloud ecosystems with respect to key technical challenges such as secure identity management, segmentation, continuous policy optimization, and scalable incident response.

This review highlights the urgency of organizations implementing adaptive and intelligence-driven security architectures that can reduce advanced cyber threats and improve enterprise resiliency while ensuring regulatory compliance during digital transformation [5].

1.1. Problem Statement

The security challenges in multi-cloud environments are significant, as the systems and data are spread across various cloud providers, each employing different technologies, configurations, and access control models. Such distribution makes it challenging for traditional anomaly detection and perimeter-based security solutions to operate, particularly against evolving real-time cyber threats. Although zero trust provides more comprehensive assurances with rigorous and ongoing verification, very few studies focus on detecting anomalies in multiple cloud platforms at the same time, and none of the existing frameworks yet combines zero trust with AI-based anomaly detection in a multi-cloud environment. This gap reduces the ability to proactively and comprehensively secure multi-cloud infrastructures.

As an architectural method that requires constant verification and tight access control, zero-trust architectures have become a new standard in the cybersecurity world. However, most existing research focuses on zero trust from a policy or architectural point of view, with some limited attention on threats such as anomaly detection techniques. Although some studies have identified the potential of integrating artificial intelligence and zero trust, these efforts are often conceptual, rely on isolated datasets, or provide no practical coordination between different cloud providers.

As a result, there is still no clear picture of how AI-enabled anomaly detection is actually applied in the context of zero-trust frameworks in multi-cloud environments or what practical challenges and limitations exist. To fill this gap, a systematic literature review is required to explore the existing research on AI-based anomaly detection, zero-trust frameworks, and multi-cloud security in order to offer structured insights that can contribute to future research and system development.

1.2. Research Questions

RQ1: Which AI methods and algorithms have been implemented to detect anomalies in multi-cloud environments, and how do they vary in terms of effectiveness and scalability?

RQ2: What are the most prevalent datasets and evaluation measures for evaluating AI-based anomaly detection techniques in multi-cloud environments?

RQ3: What is the conceptualization and integration level of zero-trust principles with AI-based anomaly detection solutions in existing multi-cloud security solutions?

RQ4: What are the challenges and limitations identified in the implementation of AI-enabled anomaly detection in zero-trust environments across various multi-cloud environments?

RQ5: What are the latest trends and future research opportunities to improve AI-based anomaly detection in accordance with zero-trust architecture in multi-cloud environments?

To answer these research questions, we provide an in-depth overview of cloud computing, multi-cloud environments, and the zero-trust framework. Next, we conduct a comprehensive systematic literature review for paper selection as per PRISMA 2020 guidelines.

1.3. Contributions

This SLR study provides the following contributions:

- An in-depth systematic taxonomy of AI-based techniques for anomaly detection techniques in cloud and multi-cloud environments, including supervised, unsupervised, self-supervised, federated learning, graph-based, and hybrid approaches.
- A mapping that connects AI methods, deployed models (single cloud, hybrid, multi-cloud), and zero-trust capabilities such as continuous monitoring, micro-segmentation, identity federation, and policy-based enforcement.
- A comparative analysis of widely used datasets and evaluation metrics, including their strengths, limitations, and suitability in multi-cloud and zero trust scenarios.
- Conducing comprehensive research gap, future work direction in multi cloud environments with aligned with privacy and compliance requirements.

2. Background

2.1. Cloud Computing and the Move Toward Multi-Cloud

Cloud computing is fundamentally changing the way that organizations manage and deliver IT resources by enabling scalable, elastic, and on-demand IT resources through various different deployment models, including Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). IaaS provides full on-demand service (for example virtual machine, operating system, networking, and storage), allowing for customer deployment without the need to manage infrastructure such as physical security cabinet. PaaS offers development services such as building, customizing, and deploying apps, while SaaS offers software applications to the customer over the internet without requiring any installation or maintenance [6].

In the last few years, organizations have been moving more towards multi-cloud environments due to strategic or operational needs. One major motivating factor is the need for vendor flexibility, which is the ability to choose services from different cloud providers and avoid relying on a single vendor. A second factor is about the cost based on service and performance. In addition, deployment of services over geographically and architecturally different clouds makes them more resilient and helps to achieve business continuity by minimizing the effect of service outages. As a result, multi-cloud adoption is especially attractive to organizations that are dealing with sensitive data and operating in environments with strict regulatory requirements such as the GDPR or NIST [7,8].

Despite these advantages, the use of multi-cloud environments brings a certain level of complexity to issues such as identity management, policies consistency, and threat detection. Differences in cloud-native security controls and visibility models make it hard to enforce uniform policies and centralized monitoring [9].

To keep up with these challenges, organizations are looking to zero-trust approaches to access and intelligent analytics in order to enable continuous verification and orchestrated security enforcement across the cloud platforms. As adoption of multi-cloud continues to grow, security solutions that integrate automation, adaptability, and centralized manage-

ment of dynamic and distributed environments are a must-have to manage cloud-based dynamic environments [10].

In the end, multi-cloud environments have their advantages—improved flexibility and operational resilience among others—but come with higher security complexity. The distributed nature and lack of end-to-end visibility of multi-cloud platforms ruin the effectiveness of traditional perimeter-based security approaches. This level of complexity makes an interesting argument for more advanced solutions; examples include anomaly detection with AI that runs across a number of different cloud providers with no issues and zero-trust architecture principles for continuous verification. These emerging challenges are a direct motivation for RQ3 and RQ4, which focus on the design of integrated approaches to the detection of anomalies and zero-trust practices in multi-cloud environments [11,12].

2.2. Cloud Services Models

There are three main service models of cloud computing, shown in Figure 1. Each has a different level of abstraction, control, and security responsibility between cloud provider and cloud user.

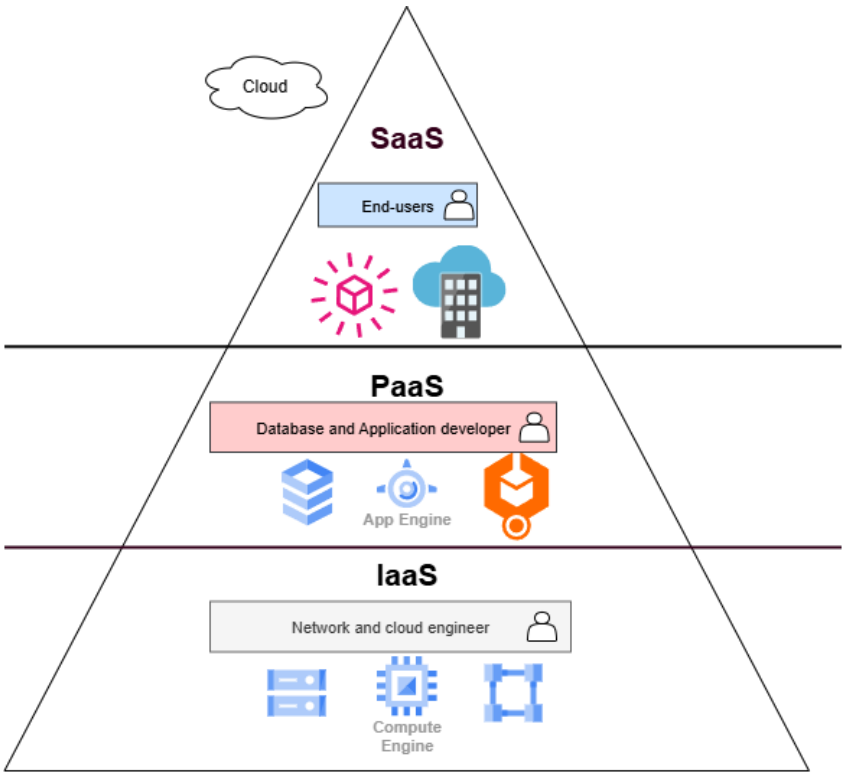

Figure 1. Cloud services models.

Infrastructure as a Service (IaaS): is an on-demand service that provides simple computing elements such as virtual machines, storage, and networks. The operating system, applications, and data are controlled by the user, whereas the physical infrastructure is managed by the provider. Since users can directly access the software stack and data stored on it, security in IaaS is largely aimed at configuration security, patch management, and Identity and Access Management (IAM) [6].

Platform as a Service (PaaS): provides a development platform that includes hosting of applications, middleware, and database. Data management and applications are handled by the users, whereas the provider handles runtime, OS, and hardware security. This reduction in the extent of the management transfers security issues to application weaknesses, secure coding, and data protection within the platform ecosystem [13].

Software as a Service (SaaS) refers to the provision of complete controlled software programs that are conveyed using web interfaces. In this case, the provider takes care of practically everything, including application security as well as data storage and infrastructure. Control of access and authentication is primarily the role of end users; security concerns include identity management, data privacy, and compliance adherence [6].

These differences are important factors to take into account when establishing efficient security practice in multi-cloud environments where service models are applied. Since each model uses its own responsibility matrix, the relevant security controls and risk management practices should be adapted so as to offer sufficient protection to the various levels of control and common shared responsibility in multi-cloud deployments [13].

2.3. Security Challenges in Multi-Cloud Environments

As shown in Table 1, multi-cloud environments are characterized by complexity and heterogeneity, since organizations are consuming services across numerous different cloud providers simultaneously and each of these cloud providers has their own different architecture, security protocols, and configurations. It is this diversity that provides centralized governance and security management as a challenge, and that can also contribute to the possibility of configuration errors and inconsistency in security policies that may result in exploitable vulnerabilities. The decentralizing of workloads on the various platforms also makes systems more complex, such that it may become impossible to apply security policies and the proper monitoring of activities [7].

A more open attack surface is one of the key consequences of multi-cloud adoption. Cloud platforms not only provide more access to resources via APIs, management consoles, and network interfaces, they also open up opportunities to cyber attackers. Very often, the complexity of securing the disparate environments is the cause of misconfigurations, which is a major factor in data breaches and unauthorized access. This threat is further escalated by the fact that it is difficult to manage similar Identity and Access Management (IAM) practices in heterogeneous clouds, which can be used to trigger privilege escalations or lateral movement by attackers [7].

Moreover, data privacy, and regulatory compliance are more difficult in multi-cloud environments. Businesses are forced to deal with compliance requirements such as GDPR, HIPAA, and NCA ECC (the Saudi national cybersecurity authority) controls [14].

All of these require a high degree of control in the processing of data as well as in the encryption and auditing of data. The presence of effective cross-cloud compliance means that high-end and automated compliance monitoring tools and actionable analytics is required. Without these measures, there is always a possibility of loopholes in the data protection policies that can put organizations at the risk of lawsuits and financial consequences [8].

Table 1. List of challenges and potential threats in multi-cloud environments.

Multi-Cloud Challenge Area	Security Challenges	Potential Threats
Complexity and Heterogeneity	Different architectures and security models hinder unified governance and policy consistency.	Misconfigurations, policy gaps, exploitable vulnerabilities.
Expanded Attack Surface	Multiple access points (APIs, consoles, interfaces) and inconsistent IAM increase exposure.	Unauthorized access, lateral movement, privilege escalation, data breaches.
Data Privacy and Compliance	Varied regulations and encryption requirements complicate unified compliance and data protection.	Regulatory violations, data leaks, legal and financial penalties.

2.4. Evolution of the Threat Landscape in Cloud Computing

The threat landscape has transformed significantly as companies have ceased to use conventional IT infrastructure in favor of multi-cloud systems. Early cloud threats were primarily in the form of insider attacks, through which dangerous inside users exploited privileged access in order to compromise systems. These threats have become more advanced in the years, and today there are also Advanced Persistent Threats (APTs) in which the attacker will perform covert and persistent attacks with the purpose of stealing the critical data during a prolonged period of time [15].

Lateral movement within the cloud has emerged as a serious strategy that attackers have adopted to gain privileges and move freely across cloud resources. Lateral movement is particularly difficult to identify and address in multi-cloud contexts where the security controls and access controls apply to various environments. Weaknesses and misconfigurations of identity federation can offer another point of access on top of the compromised account, and may be exploited to make further attacks [8,16].

Account takeover attacks are becoming both more sophisticated and more common. These attacks target cloud authentication systems using stolen credentials, phishing, and bots designed for tasks such as credential stuffing. The scale and automation possible with cloud services exacerbates the effects of these attacks, requiring constant authentication and sophisticated anomaly detection engines. Depending on the nature of the cloud architecture, attackers are increasingly able to use AI technologies and machine learning to attack while evading detection [15].

Dynamic models of security such as zero trust with Artificial Intelligence (AI)-enhanced anomaly detection can be used to observe and adapt to these new risks in cloud ecosystems. The continuous development of attacks necessitates preemptive context-sensitive security measures to protect sensitive workloads that operate with a variety of cloud providers [17,18].

2.5. Zero Trust Security Paradigm

Zero trust is a security paradigm developed as a reaction to the shortcomings of older and more perimeter-based security models, particularly when it comes to new more distributed and dynamic IT environments. The essence of zero trust can be summarized in the statement "never trust, always verify", meaning that no user, device, or system is implicitly trusted, regardless of whether or not it is within the network perimeter. Rather, identity and context verification must be performed continuously in order to access any resource [19].

The critical elements of zero trust are: (1) microsegmentation, in which the network is divided into small segments in order to limit the movement of lateral attackers and enforce access control between various segments; (2) continuous monitoring, in which the user behavior, network traffic, and system events are examined in real time to detect abnormalities that could be an indication of a compromise; (3) least privilege, which is based on the need to ensure that users and systems have access to the lowest possible set of access they require to conduct their operations, thereby minimizing the risk exposure that compromised credentials represent; and (4) identity-centrism in the approach to security, meaning that zero trust applies strong authentication, multi-factor authentication, and identity governance to control and audit access with greater precision [20].

Multi-cloud environments are extremely dynamic, heterogeneous, and complex, which makes zero trust especially important. The disparity between the security model and policies of the various cloud providers tends to create security loopholes which are not addressed by traditional security models. Zero-trust architecture in multi-cloud environments offers a central security platform that enforces similar policies around security access, risk assessment, and dynamic controls for all cloud platforms, thereby addressing propagation of the attack surface and providing more security for distributed workloads [21,22].

2.6. AI-Driven Anomaly Detection for Cloud Security

Anomaly detection in cloud security refers to the ability to identify unauthorized access, insider threats, and new types of cyberattacks, many of which emerge in the form of anomalous system or network behavior. Traditional rule-based detection methods are not effective in modern cloud environments because of their dynamic nature, large amounts of data, and the constantly evolving threat environment. As a result, organizations are relying more and more on machine learning and artificial intelligence to detect known and unknown threats [23]. Machine learning methods for anomaly detection can be grouped into two main categories. Supervised learning techniques are trained on labeled datasets to identify the difference between normal and malicious behavior, and are highly accurate provided that labeled data are available. On the other hand, unsupervised learning techniques operate by identifying anomalies without know any information about attack patterns. Deep learning algorithms—specifically neural networks, LSTM models, and autoencoders—are very good at detecting complex relationships in cloud telemetry data as well as at detecting advanced threats and zero-day attacks that can be missed by simpler methods [18].

AI-based anomaly detection has a number of advantages compared to legacy approaches. AI systems can dynamically adapt to changing threats via continuous learning, allowing them to detect new attacks that are not in historical signatures and produce actionable alerts with fewer false positives. When implemented into cloud security architectures, AI anomaly detection provides real-time monitoring and analysis which can detect and respond to the incidents quickly, making the entire cloud more resilient with high effectiveness in the threat mitigation [8,15,24].

To sum up, cloud anomaly detection research is highly dependent on individual cloud datasets, and does not include the integration of zero-trust controls. This gap is the motivation for the RQ1, RQ2, and RQ3, which explore the AI techniques and datasets applied in multi-cloud environments and how they interact with zero-trust principles.

2.7. Integration of Zero Trust and AI in Multi-Cloud Environments

The dynamic data-driven security mechanisms generated by the AI techniques that traditional static policies cannot offer are the best way of enforcing zero-trust principles in multi-cloud environments (shown in Figure 2). One AI-based technique is behavioral analytics, which enables continuous profiling of user and entity behavior along with detection of anomalies as opposed normal behavior. Behavioral analytics can be used to detect malicious intent or compromised accounts, in turn facilitating access control by risk, where permissions are granted or denied according to real-time risk analysis as compared to fixed policies. This approach is exactly in keeping with the principle of "never trust, always verify" that is the core of zero trust [1].

AI can also be used to detect new and advanced threats in real time through the analysis of massive streams of heterogeneous data generated in distributed multi-cloud environments. Combined with automated policy enforcement, AI-driven systems have the ability to respond quickly when anomalies are identified by isolating corrupted resources, altering access control, or initiating incident response processes with a small number of human interventions. This requirement of continuous change is necessary in multi-cloud infrastructure contexts, which by definition lack a centralized security perimeter and include multiple service providers with dissimilar security postures [19].

Although AI and zero trust are complementary, issues as data privacy, explainability of models, and the threat of adversarial attacks on the AI system should be expected. Research, development, and maturity of standards-based frameworks is needed to support AI-enabled policy automation, scalable cross-cloud identity federation, and AI architectures

that are capable of scaling and operating within compliance and performance limits. More and more sophisticated multi-cloud environments will demand the development of these spheres to achieve the maximum potential of AI-enhanced zero trust [3,19].

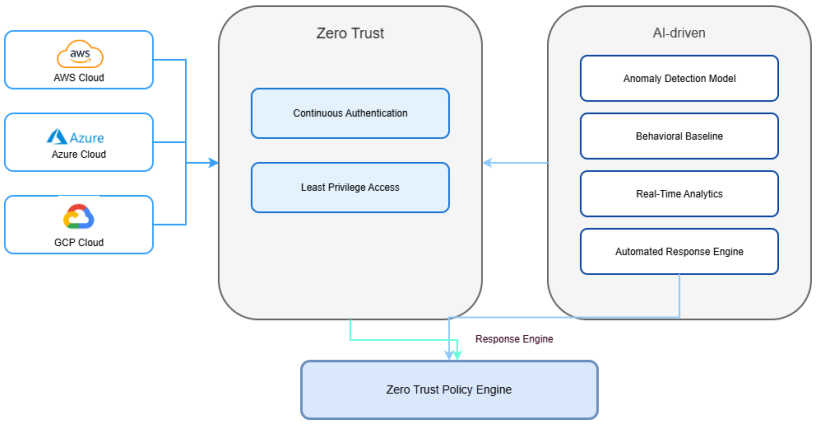

Figure 2. Levels of zero-trust integration among the selected studies.

2.8. Regulatory and Governance Considerations

Security mechanisms that are deployed in multi-cloud environments must operate within the context of data protection and sector-specific regulations, which play a critical role in the way cloud systems are designed and secured. Regulations such as the General Data Protection Regulation (GDPR), Health Insurance Portability and Accountability Act (HIPAA), and California Consumer Privacy Act (CCPA) have introduced stringent mandates in relation to data collection, processing, access control, auditing, and accountability [25–27].

These regulations are especially relevant in multi-cloud environments where data can be transferred between providers, regions, and locations and where the risk of misconfiguration and non-compliance is high. In addition to global regulations, there are national regulations such as the Saudi National Cybersecurity Authority Essential Cybersecurity Controls (NCA ECC/CCC) and Personal Data Protection Law (PDPL). These further define the expectations regarding governance and security, especially for organizations operating in regulated or critical sectors [28–30].

Despite the significant impact of these regulatory frameworks on cloud deployments in the wild, most existing studies of AI-based anomaly detection are mainly concerned with detection accuracy and the performance of various detection methods, with little regard for explicit compliance or governance requirements. In many cases, regulatory aspects are either implicitly assumed or are handled indirectly by means of techniques such as privacy-preserving learning, data minimization, or decentralized training approaches. Explicit mechanisms that provide a linkage between the results of anomaly detection and zero-trust enforcement, auditing, and regulatory accountability are rarely in place [31].

This gap demonstrates a need in the research to improve the integration of anomaly detection techniques with the principles of zero trust and the needs of regulation, especially in the case of complex multi-cloud environments, where compliance, transparency, and trust can be as important as detection performance.

3. Research Methodology

This paper study follows the 2020 Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines to ensure a transparent, rigorous, and reproducible research process. The adopted methodology is based on four main stages: first, the eligibility criteria used to select which studies are included or excluded; second,

the information sources (databases and relevant digital libraries) used for identification and documentation; third, a systematic search strategy with selected keywords and Boolean operators to capture the most relevant literature; and fourth, structured data selection, collection, and synthesis. This approach brings clarity, improves reliability, and contributes to the replicability of the review.

This review complies with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) [32]. None of the authors registered this review in any systematic review registry.

3.1. Inclusion and Exclusion Criteria

Clear inclusion and exclusion criteria were defined to ensure that only relevant high-quality studies were included in this review.

3.1.1. Inclusion Criteria

Studies were included if they: focused on security concerns on multi-cloud environments; identified issues with anomaly detection in the context of zero-trust frameworks; or discussed challenges in cloud environments or integration with the zero-trust framework. Peer-reviewed studies published between 2020 and 2025 were considered for inclusion.

3.1.2. Exclusion Criteria

Studies were excluded if they: were unrelated to multi-cloud security or zero trust; did not deal with or were outside the scope of the concept of anomaly detection, as specified by the topic of anomaly detection using artificial intelligence; or were published prior to 2020.

3.2. Information Sources

A systematic literature search was carried out in several scholarly databases and digital libraries in order to ensure completeness. Recent relevant publications on AI-driven anomaly detection in multi-cloud environments within zero-trust frameworks were retrieved from IEEE Xplore, ScienceDirect (Elsevier), MDPI, Google Scholar, and the Saudi Digital Library.

3.3. Search Strategy

A systematic search strategy was employed to systematically identify studies pertaining to anomaly detection, zero-trust principles and multi-cloud security. The Boolean operators (`AND`, `OR`, and `NOT`) were used to refine the results and maximize relevance. The primary search string was:

```
(‘‘multi-cloud’’ OR ‘‘hybrid cloud’’) AND (‘‘anomaly detection’’ OR ‘‘AI-
driven security’’) AND (‘‘Zero Trust’’ OR ‘‘zero-trust architecture’’).
```

To ensure comprehensive coverage, keywords such as multi-cloud anomaly detection, zero-trust access control, cloud workload monitoring, AI-driven cloud threat detection, micro-segmentation, and continuous verification were added to fit within the syntax of each database.

To avoid duplication, the following AI sub-techniques were excluded from the primary search string as separate keywords: federated learning, graph neural networks, self-supervised learning, and deep learning, as these topics are already covered by the more general terms AI-driven security and anomaly detection. They were added manually to the tests and resulted in many duplicates and irrelevant results. The selected formulation was broad enough to capture studies that utilized any of these types of AI, yet specific enough to identify those that addressed multi-cloud, anomaly detection, and zero-trust contexts.

Search Filters Applied

The results were further filtered to include only publications from 2020 to 2025 involving machine learning, anomaly detection, and zero-trust architectures.

3.4. Selection Process

This study follows the PRISMA 2020 guidelines, including the three phases of identification screening, and inclusion.

First, records were collected from the following sources: IEEE Xplore, MDPI, Saudi Digital Library, and Google Scholar. Initial screening of the articles was performed by excluding them on the basis of title and abstract relevance. Full text eligibility assessment was then performed based on the predefined inclusion and exclusion criteria.

3.5. Quality Assessment

To ensure the inclusion of high-quality and methodologically sound studies, a structured Quality Assessment (QA) process was applied to all candidate papers that passed the screening phase. Each study was scored using a checklist that comprised seven criteria suited to the scope of AI-enabled anomaly detection and zero-trust security in multi-cloud environments.

The quality criteria included the following:

- Clarity of the research objectives.
- Description of the AI technique and system architecture.
- Clear definition of what is a cloud or multi-cloud environment.
- Evaluation and reporting of performance measures.
- Clear consideration of zero-trust principles.
- Relevance to anomaly detection and cloud security.
- Level of the methodology used to indicate reproducibility.

Each aspect was rated on a basic scale of 0 to 2, where 0 indicates that the problem was not covered, 1 that it was partially covered, and 2 that it was clearly explained. The scores for all the criteria were then combined to form an overall quality score. Studies with a total score lower than 8 were excluded from the final review. These scores were used solely to support the study selection process, and did not form part of the study results.

3.6. Data Extraction and Coding

To ensure consistency and transparency of the analysis, we followed a structured process for extracting data from all of the studies in the final review. We followed a set template in order to gather the same key information from each study, ensuring that they fit our research questions.

For each selected paper, the following information was extracted:

- Author(s) and year of publication.
- Application domain (e.g., enterprise systems, IoT, healthcare, telecommunication).
- Model of cloud deployment (e.g., single cloud, hybrid, multi-cloud, edge/fog).
- Type of AI learning technique (e.g., supervised, unsupervised, self-supervised, federated, graph-based, hybrid).
- Datasets used for evaluation.
- Reported evaluation metrics (e.g., accuracy, precision, recall, F1 score).
- Zero-trust principles (e.g., continuous monitoring, identity management, micro-segmentation, policy enforcement).
- Key findings and limitations.

The data gathered in this way were used to create landscape tables providing a comparison of zero trust and legacy models, which are presented in the results and discussion section.

3.7. Identification Phase

In the identification phase, we identified 245 records from the IEEE Xplore, Saudi Digital Library, MDPI, and Google Scholar databases. Of these, 120 duplicate records and 15 ineligible records were indicated by automation tools and eliminated. Ten records were also eliminated due to other reasons, leaving 100 records.

3.8. Screening Phase

The remaining 100 records were screened in terms of title and abstract to relevance to AI-driven anomaly detection and zero trust in multi-cloud environments. According to the inclusion criteria, 60 records were eliminated based due to irrelevance, leaving the total number of studies as 40.

3.9. Eligibility and Inclusion Phase

Of the 40 identified studies, five could not be retrieved. Of the other 35 studies evaluated as eligible, 13 were excluded for the following reasons: eight were unrelated to the abstracts, and five were not relevant enough to the area of research.

As shown in Figure 3, the 22 studies that passed all inclusion criteria were included the final in-depth systematic review.

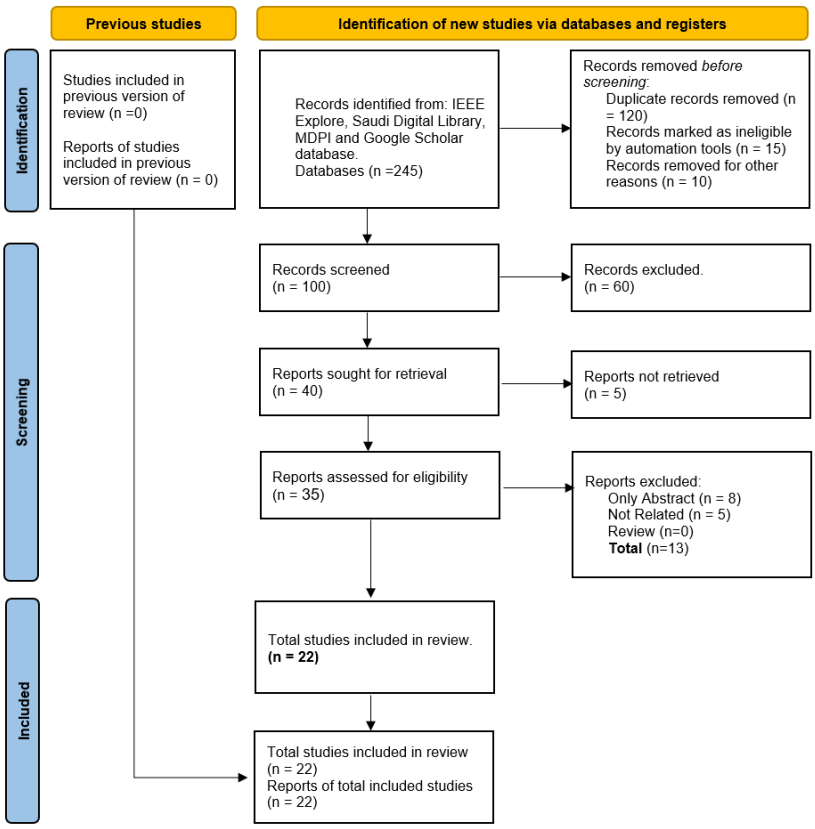

Figure 3. PRISMA 2020 flow diagram for study selection.

Figure 4 presents a line chart that shows the distribution of the 22 chosen papers by year of publication. The figure indicates that research activity has been increasing over the years, reaching a peak in 2025.

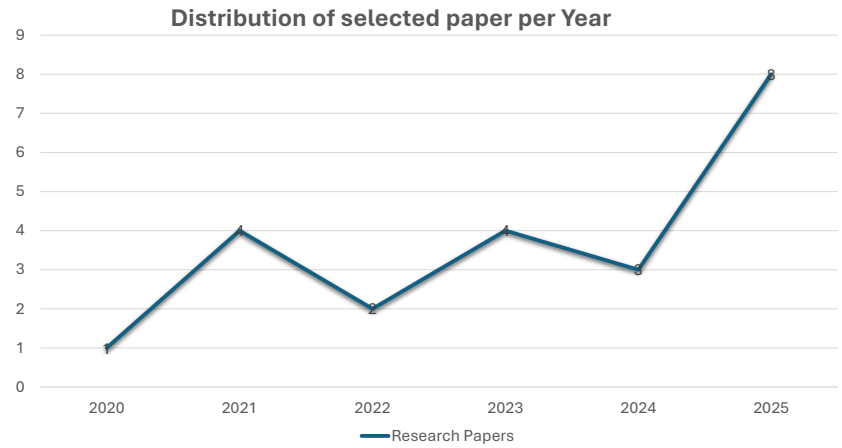

Figure 4. Paper distribution by year of publication.

In general, the research methodology provides a clear and understandable systematic review protocol, allowing for a direct connection between the research questions and research selection process, quality evaluation, and systematic data extraction. Following the selection process in accordance with PRISMA helps to ensure that all literature published in the field is covered, and the quality check excludes research with weak methodologies. The analysis based on the RQs is directly supported by the extracted data, which allows us to conduct an accurate synthesis of AI methods, datasets, zero-trust integration, and research gaps in the field of multi-cloud anomaly detection.

4. Literature Review

In this section, we review the 22 selected studies that discuss the topics around AI-driven anomaly detection, zero-trust architectures, and zero-trust approaches, with a strong focus on architecture and policy design. These studies consist of survey and conceptual papers, and are grouped into thematic groups to highlight common patterns, differences, and research gaps across the reviewed works.

4.1. AI-Only Anomaly Detection

Almajed et al. [33] showed the effectiveness of machine learning (supervised, unsupervised, hybrid, ensemble) in detecting anomaly in cloud computing in relation to security threats, including data breaches and insider attack. The article combines the findings of 32 recent papers that have evaluation measures such as the rate of detection, false positives, and F1-score; datasets such as KDD-CUP, CICIDS-2017, and UNSW-NB15 are examined. The review concludes that hybrid and ensemble ways occasionally reach detection accuracy > 99%. Nonetheless, there are still issues with complexity in implementation, real-time adjustment, and variety of the datasets. The review points out that there is no considered solution that deals with distributed learning, federated learning, or dynamic policy implementation that is needed by zero trust in a multi-cloud environment. These models are typically not inter-provider transferable and fail to provide zero trust with constant context validation, and as such are inadequate to detect anomalies in AI-based multi-cloud context detection in zero-trust environments.

Bakro et al. [34] introduced a better cloud Intrusion Detection System (IDS) with a determination of bio-inspired features using a random forest classifier. Based on datasets

including UNSW-NB15, CIC-DDoS2019, and Kyoto, the work has a high accuracy (up to 99) and evaluates performance through the levels of true/false positives and ROC. Though this approach enhances efficiency and detection, it was designed and implemented in single-cloud environments only and cannot support policy harmonization, federated learning of anomalies, or AI explainability in multi-cloud entities with zero trust. Therefore, it cannot be conveniently applied to zero-trust highly-regulated cross-provider issues of detection.

Jaber et al. [35] aimed to improve NSL-KDD data. Their hybrid intrusion detection system is based on the capabilities of a fuzzy C-means clustering and SVM with large classification rates of attacks of different types. The proposal is excellent at reducing false positives; however, this was proven only on an artificial single-cloud grid, not on federated or multi-cloud zero-trust models. The approach excludes dynamic policy implementation, real-time cross-cloud integration, and AI privacy-preservation, which are required for zero trust in heterogeneous environments.

Thapa et al. [36] provided a survey review of machine learning-based anomaly detection techniques in cloud computing designed to combat cyberattacks, insider threats, and the weaknesses of traditional manual analysis. In connection to public, private, hybrid, and multi-cloud industry, the author thoroughly evaluated supervised, unsupervised, and hybrid methods (neural networks, SVMs, clustering, ensemble models) for cloud anomaly detection. Accuracy, precision, recall, F1-score, rates of false positive/negative, and complexity of computation were the performance measures. Results from an experiment revealed that such techniques as RNN-LSTM and CNN demonstrate more than 98 percent accuracy in both structured network traffic and anomaly cases. High requirements for labeled data, interpretability, alert fatigue, and model transferability were identified as limitations. The reviewed datasets were CICIDS, NSL-KDD, and workflow logs on clouds. Although this paper offers useful suggestions on algorithm choice, it is not a complete fit for the topic of AI-based anomaly detection in multi-cloud and zero-trust contexts, since it concentrates primarily on detection methods rather than complete zero-trust architectures, there is no integration of automated policy enforcement, and there is no coverage of privacy, explainability, or cross-cloud orchestration in regulated environments.

Megouache et al. [37] studied a hybrid clustering/classification scheme. The framework was tested with K-means and Extreme Learning Machines (ELM) for detecting cloud intrusion, and the results reached above 99% on the KDD99/NSL-KDD datasets. The presence of computational efficiency with large datasets is emphasized, although the techniques were not evaluated regarding multi-cloud setup and federated setups or zero-trust preparedness. Absence of context-aware automatic policy coordination and inter-provider outlier identification is a hindrance to applicability in next-generation AI-enabled multi-cloud zero-trust security.

Moreira et al. [38] suggested an intelligent system for anomaly detection in smart and fog/cloud environments. Machine learning was used to identify network anomalies for both cloud and fog infrastructure. Substantiating the quality in the dataset CICIDS using models such as random forest, their findings indicated that the model has high value; however, there were concerns in relation to particular infrastructure and decreased recall in some conditions. The multi-layer architecture lacks federated learning, adaptive zero-trust policy enforcement, and explainable AI. These are needed to operate in regulated multi-cloud systems, restricting this model's generalization for AI-based anomaly detection in zero-trust architectures.

To sum up, these studies demonstrate how machine learning models can obtain high detection in cloud-based anomaly detection, especially when tested using benchmark data such as CICIDS and NSL-KDD. However, most of the methods are detected individually and tested in one cloud. This lack of integration with zero-trust mechanisms, coordination

across clouds, and federated or privacy-preserving learning makes them inapplicable to real-world multi-cloud and zero-trust scenarios.

4.2. AI-Enabled Zero-Trust Architectures

Talati et al. [15] introduced a general AI-based anomaly detection platform for multi-cloud environments. The platform is based on federated learning, graph neural networks, and self-supervised learning. It improves threat detection without infringing data privacy and attempts to solve the weakness of legacy signature-based systems in identifying advanced and dynamic threats in cloud environments. Specifically, multi-cloud security (AWS, Azure, GCP) is approached and the methodology includes federated intelligence, graph analytics, hierarchical federated learning, post-quantum cryptography, and automated incident response at distributed infrastructures. Important measures are detection accuracy (F1-score, recall, precision), false positive rate, scalability, and adversarial robustness. Experiments indicate good detection rates (i.e., 96.8 percent with self-supervised learning, 92.5 percent with federated models) and large decrease in the average time to detect/respond to incidents. The drawbacks are the resource overhead of federated learning, complexity of real-time enforcement of policies, and difficulties in integrating quantum-resistant cryptography. Datasets used included logs from AWS, Azure, Google Chronicle, CICIDS, DARPA, and more. Tools included Apache Spark, Kafka, Kubernetes, PyTorch, etc. Although this work contributes to the privacy-preserving, quantum-resilient, and scalable detection of anomalies, it is not completely applicable to multi-cloud zero trust because of inabilities to administer real-time cross-cloud policy enforcement and edge case scalability as well as persistent tradeoffs between security efficacy and computational overhead. Lack of dynamic and context-aware zero-trust implementation and insufficient attention to explainable AI models in compliance-related settings are the key gaps.

Micheal et al. [39] suggested combining AI-enabled anomaly detection with zero trust risk-adaptive authorization for the security of cloud APIs, with an emphasis on defense-in-depth and continuous monitoring of user transactions. They aimed to identify API-specific attacks (e.g., abuse, credential misuse, sophisticated threats) on multi-tenant setups (e.g., SaaS, financial and healthcare APIs). Using an architectural methodology that involves layers (API gateways, unsupervised detection, policy enforcement), they tested models such as isolation forest, autoencoders, LSTM and random forests based on features extracted from request patterns. Measures such as accuracy, recall, precision, risk-adjusted access control, and false positive/negative rates were used. The findings reflected better advanced and subtle API attacks, dynamic access control, and regulatory compliance; however, there were still issues of model drift, explainability, false positives, and API performance implications. No specific datasets were revealed, but simulated real world API transaction data were incorporated with the use of standard tools such as Kafka, Spark, and explainable AI frameworks. This solution, while sturdy in terms of API level anomaly detection, is not sufficient to address the full scope of AI-driven anomaly detection for general multi-cloud environments or specifically in terms of automated zero-trust enforcement at scale, federated cross-cloud learning, and adaptive policy harmonization. The approach by the authors does not put emphasis on integration of the zero-trust context, explainability, or real-time multi-cloud requirements for regulated sectors.

Parisa et al. [4] suggested zero-trust security models based on AI in the context of retail cloud infrastructure with continuous authentication, anomaly detection, and threat mitigation of customer data. They covered cloud threats in retail contexts such as insider fraud, ransomware, and supply chain attacks. The evaluated the accuracy of threat detection along with the response time, false positive rate, and overhead of system performance, which they tested against several attack scenarios with the help of

machine learning algorithms (random forest, XGBoost, LSTM) and SOAR automation in multi-cloud testbeds (AWS, Azure, GCP). Findings indicated that the suggested AI-based zero trust model beat conventional methods and methods across all three aspects, i.e., detection improved by up to 18.8, false positives minimized, and compliance with the GDPR and PCI DSS requirements. Limitations such as computational overhead, alert fatigue, ethical AI considerations, and integration issues were apparent in the hybrid legacy cloud environment. The datasets consisted of retail data logs in the real world (UNSW-NB15 and CICIDS-2017). The tools were typical cloud security automation/orchestration. This solution is not as fully suitable for AI-based anomaly detection in multi-cloud zero trust, as it is limited to the vertical scope (only retail), does not support federated learning to cooperate with other clouds, and is not context-aware for managing real-time adaptive zero trust enforcement, which is critical in dynamic and multi-provider environments.

Narang et al. [40] studied cloud-based IoT and IIoT. The authors introduced a zero-trust intrusion detection system that operates on ensemble AI models (particularly XGBoost) to identify DDoS, enumeration, and malware threats with high accuracy (94.55%), recall (98.38%), and F1-score (94.22%). The range is IIoT and smart grid cybersecurity at the edges, and incorporates AI predictive analytics and zero-trust continuous validation. As an evaluation criterion, the Edge-IIoT dataset was considered, which evaluates performance in the classification of multi-class attacks, alongside a comparison with other algorithms (RNN, KNN, AdaBoost) and the ability to detect threats. Findings indicated excellent performance, risk reduction, and response time in comparison to traditional models, but emphasized computational complexity and scalability. The authors suggested the adoption of federated learning in the future in terms of privacy and robustness, yet it has not been implemented. The proposal, though state-of-the-art, does not suffice in general AI-driven detection of anomalies in multi-cloud zero trust environments because it remains vertical (IoT/edge) deployment, includes limited federated work, and lacks generalized and privacy-sensitive anomaly detection and adaptiveness for real-time zero trust coordination of activities across a wide range of clouds.

Banerjee et al. [4] aimed to move forward to the next generation of AI-based zero trust of retail cloud. The authors suggested an orchestrated model of anomaly detection, continuous authentication, and SOAR automation to improve the security of distributed retail infrastructure. The idea is to prevent insider, ransomware, and API abuse threats through the use of real-time threat analytics. They used a mixed-method experimental approach, simulated experiments involving the deployment of multi-clouds (AWS, GCP and Azure), and registry benchmarking (old and new zero-trust AI architectures). The value of qualitative data is performance in detection/mitigation rates, false positives, latency, and alignment with regulations. Findings pointed to 16–19% detection, 60% faster mitigation, and great compliance, but they observed high complexity of integration, high computing cost, and high scalability barriers to large-scale retail. Data collection consisted of generated and open-source logs, and the engines were popular cloud-native security and coordination frameworks. The model is good in the retail context but has poor cross-sector and federated multi-cloud suitability because of its poor contextual adaptation, absence of federated AI learning, and insufficient integration of explainable AI, all of which are needed to provide scalable transparent zero-trust anomaly detection in heterogeneous clouds.

Jonnakuti et al. [6] introduced a federated learning AI pipeline of secure multi-cloud (AWS, Azure, GCP) AI workloads based on the principles of zero-trust identity federation, per-service authentication, end-to-end encryption, and real-time anomaly detection. The study attempted to strike a balance between distributed training privacy, compliance (GDPR, HIPAA), and operational trust over heterogeneous clouds. They described a modular architecture that uses Kubernetes orchestration, SIEM centralization, identity

federation between OIDC/SAML, and single-session cryptographic validation. Assessment was based on latency, exposure to breaches, policy enforcement faithfulness, and resilience, while the measures to be evaluated were simulation efficiency, security alerting in real time, and integrity of the audit logs. Healthcare, banking, and defense are examples of case studies, and the cloud-native deployment tools (Docker, Celery, policy-as-code). Critical results revealed that federated AI pipelines have better security but indicated a lack of simplicity in cross-cloud orchestration, key management and observability. This solution does not have empirical large scale attack validation and is limited to support dynamic explainable anomaly detection and smooth context adaptation across various clouds. It is an innovative architecture that is not complete to support generalized and AI-enhanced multi-cloud zero-trust anomaly detection, in particular contextualized real-time policy and alert customization across multiple sectors.

In general, the above studies demonstrate the possibility of integrating AI-based anomaly detection with the concepts of zero trust to enhance security and compliance, although with limits in supporting federated learning, real-time cross-cloud enforcement, explainability, and scalable deployment in heterogeneous multi-cloud environments.

4.3. Zero Trust-Heavy Studies Focusing on Architecture and Policy

Ofili et al. [41] explored the topic of federal cloud security. The paper examined the synergy of the AI, zero trust, and CISA compliance frameworks in government IT modernization, addressing such topics as insider threats, ransomware, supply chain attacks, and compliance enforcement. The application areas were federal government and critical infrastructure, with the methodology comprising literature review, real-life case studies analysis, regulatory document analysis, and AI model analysis. The paper examined machine learning-driven threat detection systems, automated compliance auditing, and security orchestration (SIEM, UEBA) in the context of public clouds. The discussed measures included detection accuracy, access failure rates, policy violation frequency, incident response time, and compliance audit results. The results suggested that AI-enhanced zero trust policies (i.e., continuous authentication, micro-segmentation, adaptive analytics) might drastically lower the rate of unauthorized access and response time; however, they create concerns around AI bias, explainability, and federal application. The paper discusses tools such as FedRAMP, CISA, and Trusted Internet Connections (TIC) as well as platforms such as AWS, Azure, and GCP, but does not provide extensive experimental datasets. Although the article provides information on regulatory compliance and AI-enhanced orchestrations, it is not specialized to multi-cloud zero-trust anomaly detection, and lacks attention to dynamic federated AI learning, suitability to heterogeneous cloud/data flows, and explainability in highly regulated multi-cloud situations. The main gaps are insufficiency of progressive and cross-cloud privacy-sharing AI integration and scalable and automated enforcement of zero-trust policies across boundaries.

Manne et al. [8] systematically focused on the implementation strategies of zero-trust architectures in multi-clouds, and identified several challenges that have been flexible in implementing the same, including interoperability, identity federation, policy enforcement, network segmentation, and constant checkups. The paper is targeted at general enterprise, government, finance, and healthcare usage of multi-cloud (AWS, Azure, GCP) infrastructures. The NIST SP 800-207 presented in the paper is analyzed with references to existing industry frameworks based on real-world case studies, the tools surveys (e.g., CSPM, SASE, SIEM), and comparative analysis. Measures that are assessed are architecture interoperability, detection coverage/precision, policy consistency, and compliance. Major discoveries about the successful use of AI-based security analytics and software-defined perimeters in improving adaptive security, but indicate ongoing challenges of policy fragmentation, IAM

integration, and heterogeneous compliance on clouds. There are no presented empirical datasets; instead, there is a reference to tools and standards (AWS Security Hub, Azure Policy, OPA, Sentinel). The work does not entirely fit within multi-cloud AI-based zero-trust anomaly detection because it is rather generic; it does not focus on scalable federated AI anomaly learning or on the context-related and explainable real-time response across cloud providers. The biggest void is lack of depth in application of adaptive cross-cloud AI-based anomaly detection and real-time zero-trust orchestration in complex settings.

Kumar et al. [13] provided a review focusing on the topic of zero-trust architectures as a method of advanced cybersecurity in small and medium-sized enterprises that undergo digital transformation and cloud adoption. They discussed limitations in resources and expertise along with benefits of zero-trust architectures in minimizing attack surfaces, perpetuating authentication, and facilitating compliance in hybrid, cloud, and remote work settings. Methodologically, the paper used a comparative study of zero-trust architectures and legacy access models (perimeter-based and RBAC) with case studies, simulation findings, industry surveys, and performance benchmarking. Security efficacy was determined through the detection rates, false positives, and time-to-response. The results validate zero-trust architecture as an effective tool to enhance the resilience, regulatory compliance, and incident response of small and medium-sized enterprises; however, they cite high costs, complexity, integrability with the legacy system, and user acceptance as challenges. No specific datasets or tools were outlined; the evidence was obtained through the literature, reports made by practitioners, and standards such as NIST SP 800-207. An AI-based multi-cloud anomaly detection is not entirely appropriate with the paper because it focuses on SME resource and change management rather than multi-cloud orchestration, AI model integration, or automated context-aware zero-trust enforcement. Limitations in the area of focus on AI-automated detection, federated learning, and real-time cross-cloud policy adaptation are the weaknesses.

Sharma et al. [7] studied security threats in multi-cloud architectures, namely, authentication, policy inconsistency, and data breaches. They offered a risk mitigation framework based on zero trust that uses AI as its driving force along with centralized orchestration, AI-based threat intelligence, and automated compliance enforcement. The methodology encompassed the quantification of risks via the STRIDE/DREAD models and employed the evaluation metrics of risk score reduction, mean time to detection, and rate of compliance. They reported a 40–60% increase in simulated security posture. The focus of the approach was on risk taxonomy and macro-level controls, with little support of cross-cloud anomaly detection in real time, continuous implementation of policies related to zero trust, or federated AI learning. The main weaknesses are scalability and granularity, which do not entirely make it usable as a powerful AI-driven multi-cloud zero trust anomaly detector.

KumarManne et al. [8] reviewed architectural frameworks applying zero-trust security in multi-cloud microservices platforms, covering network segmentation, least-privilege access, identity federation, and automated threat response. The industries covered consisted of distributed SaaS, fintech, and applications based on microservices. The methodology included analysis of architectures, comparison of best practices, and reviews of real implementations in the area of cloud-native security, continuous monitoring, and automated enforcement of AI-based policies. Performance was measured with the security incidence reduction, enforcement latency, and consistency of cross-cloud policy, but no specific empirical measurements were provided. The main conclusions are that cross-cloud policy enforcement and attack containment are beneficial to architectures, but there are still challenges associated with the alignment of IAM policy, scaling AI-based anomaly detection, and adaptive response on a provider-wide basis. None of the datasets/tools were mentioned explicitly; the case-based insights were with respect to familiar industry security tools and standards. This review is not

perfectly applicable to the field of AI-powered multi-cloud anomaly detection in the context of zero trust because of the lack of attention to federated learning, explainable AI, and real-time context-based orchestration. Disadvantages are the absence of in-depth experimental validation and scalable automated detection results.

The main focus of these studies is architectural design, policy implementation, and identity management in zero-trust systems for multi-cloud environments. Even though they can deliver useful information about governance, interoperability, and compliance, AI-based anomaly detection is not always integrated at a high level. Therefore, these solutions do not have the scalable, adaptive, and explainable AI processes that are needed to detect anomalies and coordinate actions in the real time among multiple cloud providers.

4.4. Survey/Conceptual Works

The review by Sarkar et al. [3] provides an in-depth comparison of zero-trust network models, frameworks, and proof-of-concepts in cloud computing as applied to trust-based authorization and access control and network visibility. It covers the overall threats of cloud computing in general APTs, ransomware, insider threats, and proactive intelligent orchestration of zero-trust security. These application areas are hybrid and public clouds, and their scope is expanded to enterprise and government services. The methodology includes taxonomy of existing methods, tables on comparative features, survey of mainstream vendor products, and critical discussion of design issues. The metrics that are put into consideration are the solution adaptability, ability to automate, effectiveness in detecting attacks, and complexity of management. They found that state-of-the art models yield better detection and resilience through AI and behavioral analytics, although most are still in their infancy, with no global standards or interoperability. The drawbacks are resource requirements, the integration constraints of legacy systems, and lack of uniformity in enforcement of zero-trust. Specification of datasets is not mentioned, since the study is a review; commercial and academic prototypes are mentioned among the tools and frameworks. The key issue that makes multi-clouds and AI-based zero-trust anomaly detection less suitable is the lack of unified and federated AI models implementation and real-time harmonization of policies and a lack of comprehensive explainability of audit-compliant environments. Earlier suggestions do not have enough strength to support diversified, real-time, and multi-provider cloud infrastructure.

Ghasemshirazi et al. [1] presented a questionnaire investigating theories, applications, and challenges of zero trust, including integration with cloud, IoT, and future digital infrastructures. The authors consulted case studies and best practices that can be used to implement zero trust, concentrating on micro-segmentation, continuous authentication, and policy management. There is a strong focus on how AI and machine learning can supplement zero trust to be more context-aware, particularly in anomaly detection. Aspects of performance metrics are presented in a theoretical format; nonetheless, the article is largely abstract and does not provide any quantitative analysis or testing using multi-cloud data. The suggested integrations are mostly idealistic and do not cover automated cross-cloud compatibility of policies or learning between anomalies; as such, it cannot sustain dynamic and AI-directed multi-cloud zero-trust security assessments.

Nasiruzzaman et al. [19] traced the development of zero-trust architectures and how they are moving away from the old technical style of building perimeters into the current strategy of always verifying. The examples of Google BeyondCorp and NIST frameworks provide case studies of the problem of integration, performance overhead, and micro-segmentation. The hypothetical reconsideration focuses on the future of AI and ML to implement unbiased threats and dynamic policy integration. However, AI- and federated learning-based real-time detection of anomalies in multi-cloud zero-trust environments and

is not practically evaluated. The majority of the solutions are identified as conceptual, and lack validated cross-cloud AI models, continuous policy adaptation, and data protection provisions that are needed to achieve regulated multi-cloud zero-trust deployments.

Shnakra et al. [42] discussed the topics of zero-trust implementation in contemporary enterprise networks. Their paper explores NIST SP 800-207 and the Forrester ZTX model along with industrial examples (Google BeyondCorp, Microsoft, US DoD). The approach is qualitative based on secondary data and case comparisons, with the evaluation of threats reduced, latency, and coverage of zero-trust pillars as metrics. In this case, AI is used primarily to respond to threats and identify users, not for federated multi-cloud anomaly detection. The gaps associated with practice consist of dangerous natural support of real-time and adaptive detection and no empirical validation of dynamic and multi-normative cloud zero trust.

Westrebreg et al. [9] reflected on how to adapt the concept of zero trust in the context of enterprise security for a network. They focused on the migration of the model of traditional edge security to continuous authentication, micro-segmentation, and risk-adaptive policies. The paper relies on a positive qualitative approach and imposes custodial systems of NIST and ISO/IEC Gunnery of examinatory security controls. Although the research suggests integration points of AI analytics, the same is not widely validated. This is not how federated AI learning, cross-cloud orchestration, or explainable AI is used in the implementation, as these factors are critical to multi-cloud zero trust anomaly detection. The granularity of the solution and automation is small.

The survey and conceptual studies provide useful frameworks and design insights on zero trust and AI-enabled security. However, most are still theoretical, with little multi-cloud experimentation and few concrete implementations of anomaly detection. This gap is the difference between zero-trust concepts at the high level and the practical requirements of the AI-driven security solutions in real regulated multi-cloud environments.

5. Discussion of Related Work

The reviewed studies are summarized in Tables 2 and 3, highlighting their research focus, methods, key findings, and limitations.

Table 2. Related work summary (Part 1).

Author(s)	Research Focus	Methods and Metrics	Key Findings and Limitations
[15]	AI-based anomaly detection in multi-cloud using federated and self-supervised learning	Federated intelligence, graph neural networks, post-quantum cryptography; F1-score, recall, precision, false positive rate, scalability	96.8% detection and faster response; but resource overhead, complex policy enforcement, limited real-time zero trust context
[39]	AI anomaly detection with zero-trust risk-adaptive authorization for APIs	Isolation forest, autoencoders, LSTM, random forest; accuracy, recall, precision, FP/FN rates	Improved API attack detection and dynamic access control; but model drift, explainability gaps, limited general multi-cloud zero trust integration
[13]	Zero-trust architecture for SMEs adopting cloud	Comparative analysis with case studies; detection rates, false positives, time-to-response	Enhanced SME resilience and compliance; but cost/complexity high, minimal AI automation, lacks multi-cloud orchestration
[41]	AI and zero trust for federal cloud modernization	Literature and case studies; detection accuracy, policy violations, response time, compliance audits	Strong regulatory insight and adaptive analytics; but lacks federated AI learning, real-time cross-cloud enforcement, and explainable detection

Table 2. *Cont.*

Author(s)	Research Focus	Methods and Metrics	Key Findings and Limitations
[4]	AI-driven zero trust for retail cloud security	ML models (RF, XGBoost, LSTM), SOAR automation; detection accuracy, false positives, response time	Better detection and compliance for retail but heavy computational overhead, alert fatigue, and sector-specific scope
[8]	Implementing ZTA in multi-cloud with AI analytics	Standards (NIST 800-207), tool survey; coverage, policy consistency, compliance	Effective adaptive security; but generic approach, lacks scalable federated AI anomaly detection and real-time zero-trust orchestration
[3]	Review of zero-trust network models for cloud	Comparative taxonomy of frameworks and vendor solutions; adaptability, automation, detection	Improved resilience with AI and behavioral analytics; but no universal standards, integration barriers, and weak cross-cloud support
[36]	Survey of ML anomaly detection methods in cloud	Review of supervised/unsupervised/hybrid; accuracy, F1, FP/FN, complexity	High accuracy (>98%) but needs labeled data, lacks privacy, automated zero-trust policy, and cross-cloud adaptation
[6]	Federated AI pipeline for secure multi-cloud under zero trust	Modular federated architecture, OIDC/SAML, Kubernetes; latency, breach exposure, policy fidelity	Better privacy and orchestration; but no large-scale attack validation, limited explainable detection, and context adaptation
[40]	Zero-trust IDS for IoT/IIoT edge security	Ensemble AI (XGBoost vs RNN/KNN); accuracy 94.55%, recall 98.38%, F1 94.22%	Excellent IoT detection but vertical focus only; lacks federated, multi-cloud generalization and adaptive zero trust enforcement
[4]	AI and zero trust with SOAR for retail multi-cloud	Simulated multi-cloud deployments, comparative benchmarking; detection/mitigation, latency	Detection improvement of 16–19%, faster mitigation; but costly integration, complex scaling, and limited cross-sector applicability

Table 3. Related Work Summary (Part 2).

Author(s)	Research Focus	Methods and Metrics	Key Findings and Limitations
[8]	Zero-trust for multi-cloud microservices	Architectural analysis and best-practice comparison; incident reduction, enforcement latency	Improved cross-cloud policy enforcement but lacks federated AI anomaly detection and strong experimental validation
[33]	Systematic review of ML anomaly detection in cloud	Synthesis of 32 studies; detection rates, false positives, F1-score	Ensemble methods >99% accuracy but complex, limited real-time adaptation and no federated zero trust support
[7]	AI-driven zero-trust risk mitigation for multi-cloud	STRIDE/DREAD risk models; risk score reduction, MTTR, compliance	40–60% posture improvement but macro-level only; lacks real-time anomaly detection and dynamic zero-trust policy
[1]	Conceptual zero trust and AI integration	Theoretical review of ZT models; potential metrics discussed conceptually	Highlights AI potential but no empirical evaluation or cross-cloud zero trust practical application

Table 3. *Cont.*

Author(s)	Research Focus	Methods and Metrics	Key Findings and Limitations
[19]	Evolution of zero-trust architecture with AI role	Case studies on BeyondCorp and NIST frameworks; qualitative evaluation	Shows AI promise but no tested AI-driven anomaly detection or federated learning for multi-cloud zero trust
[42]	Zero-trust models (Forrester, NIST) in enterprise networks	Qualitative case comparisons; threat reduction, latency, ZT coverage	Mainly threat response and auth; lacks adaptive real-time anomaly detection and empirical multi-cloud validation
[9]	Adapting zero trust to enterprise networks	Constructive qualitative study with NIST/ISO frameworks	Suggests AI analytics but limited automation; no federated learning or cross-cloud zero trust orchestration
[34]	IDS with bio-inspired feature selection	Random forest + feature selection; ROC, accuracy up to 99%	Strong single-cloud detection but no policy harmonization, federated anomaly learning, or multi-cloud zero trust integration
[35]	Hybrid IDS using Fuzzy C-Means + SVM	Clustering + SVM; classification accuracy, FP reduction	High accuracy but validated only on single-cloud; no dynamic zero trust or privacy-preserving AI
[37]	Hybrid K-means + Extreme Learning Machine IDS	K-means + ELM; >99% accuracy, computational efficiency	No evaluation in federated multi-cloud; lacks context-aware zero trust orchestration
[38]	Intelligent anomaly detection for fog/cloud	Random forest on CICIDS; accuracy and recall focus	Good detection but tied to specific infra; no federated learning or adaptive zero trust for multi-cloud

6. Results and Discussion

This section reports and analyzes the findings reported by the 22 reviewed studies, which are organized according to the defined research questions (RQ1–RQ5).

6.1. RQ1—AI Methods and Algorithms

To address RQ1, which explores the AI methods and algorithms applied for anomaly detection in multi-cloud environments, ref. [15] discusses the significant increase in detection accuracy and adaptability when using state-of-the-art AI techniques, in particular self-supervised learning, federated learning, and Graph Neural Networks (GNNs). Self-supervised models such as SimCLR and LogBERT were able to achieve more than 96% accuracy on the detection of anomalies without being given labeled logs, outperforming traditional rule-based systems, which did not exceed about 86% [15]. As illustrated in Figure 5, the most frequently adopted methods in the reviewed literature are supervised learning and deep learning approaches; self-supervised, federated, and graph-based learning approaches are less common, but are emerging in the field of multi-cloud environments.

These methods not only help to provide better detection of zero-day and unknown threats, they also limit human involvement, which is critical for the scalability of multi-cloud infrastructure. GNNs result in better detection for a number of complex attack behaviors, including lateral movement and privilege elevation. Enterprise-scale deployments of GCNs and GATs have been able to reduce undetected access misconfiguration errors by 43% [15], providing significantly higher detection accuracy than traditional static log-based detection methods.

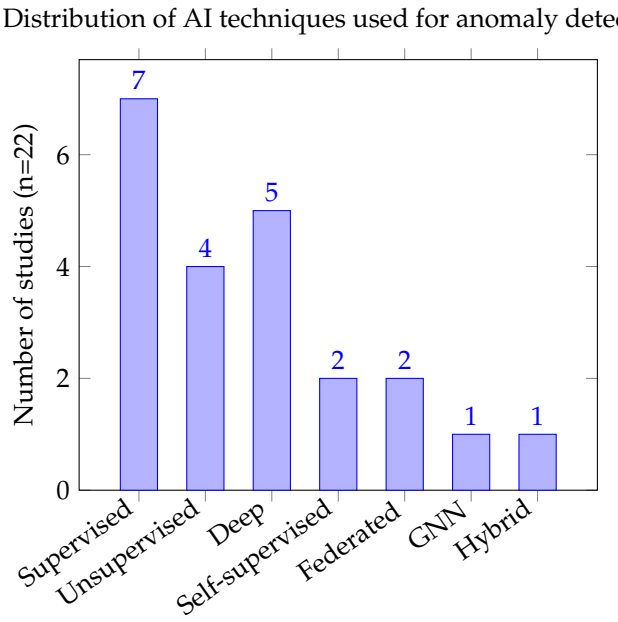

Figure 5. Counts of reviewed studies by primary AI technique category (synthesized from the 22 selected papers).

6.2. RQ2—Datasets and Evaluation Metrics

In relation to RQ2, which examines the use of datasets and evaluation metrics for the incorporation of zero-trust paradigms, one of the biggest challenges in multi-cloud security is identifying the perfect balance between privacy, compliance, and distributed learning.

Distributed anomaly detection through federated learning has proved effective for implementing GDPR- and HIPAA-compliant anomaly detection, achieving 92.5 percent accuracy [15] without having to store sensitive logs centrally across AWS, Azure, and Google Cloud. Hierarchical federated learning frameworks save even more communication overhead, almost 37% [15] over traditional federated learning. This enables continuous learning and synchronization of models in heterogeneous multi-cloud setups, thereby ensuring confidentiality—an essential element in zero-trust environments.

6.3. RQ3—Integration of Zero-Trust Principles

RQ3 [8] addresses how to combine AI-based access control and zero trust concepts to stop credential abuse and insider threats, which is an especially severe problem in sensitive sectors such as finance and healthcare. Through continuous authentication, real-time behavioral analytics, and micro-segmentation, ref. [8] reduced the number of incidences of insider attacks by 41%. In addition, AI-based dynamic risk scoring and policy enforcement can provide for faster isolation of compromised workloads, significantly reducing the attack blast radius—something that is hard to realize using static Role-Based Access Control (RBAC) mechanisms.

6.4. RQ4—Operational Challenges and Limitations

With regard to RQ4, studies such as [15] are based on the analysis of operational challenges and limitations. Studies show that AI-powered SOAR can lead to a reductions in MTTD and MTTR of 53% and 47% compared to the standard practices found in traditional SOC. This is achieved by automating honeypot deployment, integrating real-time threat intelligence such as MITRE ATT&CK into the system, and the ability to continuously tune an adaptive model for faster and more autonomous response loops, which aligns with the continuous verification principle of zero trust.

6.5. RQ5—Emerging Trends and Research Opportunities

Regarding RQ5, the reviewed literature addresses some of the growing trends in the field, such as self-supervised learning, federated architectures, graph-based models, and the new interest in quantum-resilient cryptography for cloud security. A few recent publications have been made on the feasibility of performing encrypted log analysis in (near)-real time without full decryption, with a low computation overhead reported at approximately 4.8% [15]. This means that privacy-preserving anomaly detection will be potentially more practical in regulated multi-cloud environments.

To conclude the discussion section, studies such as [15,36] addressed both RQ1 and RQ4 by showing near-real-time detection of anomalies with low latency along with small extra computational overhead in large-scale distributed environments. Containerized implementations (Docker, Kubernetes) will offer consistent deployments on multiple platforms as well as elastically scalable implementations of AI-driven zero-trust frameworks, enabling AI-driven zero-trust frameworks to exceed other approaches (rule-based, non-adaptive hybrid). In addition, as shown in Table 4, the reviewed studies show that AI-based zero-trust systems are clearly superior to traditional machine learning and static models in terms of detection accuracy, false positive rate, compliance, and scalability. These results based on the reviewed studies directly respond to RQ1 and RQ4 by illustrating the practical efficiency and flexibility of zero-trust frameworks that are built on AI in multi-cloud environments.

Table 4. Qualitative synthesis of AI-driven multi-cloud anomaly detection studies compared to legacy ML/security models.

Metric	AI-Driven Zero Trust (Synthesized from Reviewed Works)	Legacy ML/Security Models
Anomaly Detection Accuracy	96–99%	85–92%
False Positive Rate	1.5–2.5%	4–10%
Mean Time to Detect/Respond	−53%/−47%	Baseline
Compliance (GDPR, HIPAA, CCPA)	Full	Partial
Adaptive Access Control Effectiveness	High (dynamic, per session)	Moderate (policy-based)
Computational Overhead (Multi-cloud)	<10% resource impact	15–25% (or more)
Scalability (Events/s)	Millions with low latency	Often bottlenecked

To sum up, these findings validate the idea that AI-driven approaches can strengthen zero trust for multi-cloud environments, especially federated and graph-based/edge-deployed architectures. However, there are still challenges around model explainability, cross-cloud identity federation, and standardized benchmarking, highlighting the need for continued research into fully autonomous privacy-preserving zero trust models for multi-cloud infrastructure.

7. Conceptual Reference Architecture for AI-Driven Anomaly Detection in Multi-Cloud Zero-Trust Environments

Based on a systematic analysis of the 22 selected studies, a conceptual reference architecture for AI-based anomaly detection in multi-cloud environments under a zero-trust security model is presented here. We stress that this is a synthesized reference model and not an original design contribution; it is neither implemented nor empirically validated, and is presented only to organize the architectural patterns, design principles, and integration strategies observed across the reviewed literature (RQ1–RQ3). Its usefulness

is in summarizing how the existing components fit together, not the proposal of a new system. Figure 6 shows the layered structure of the architecture.

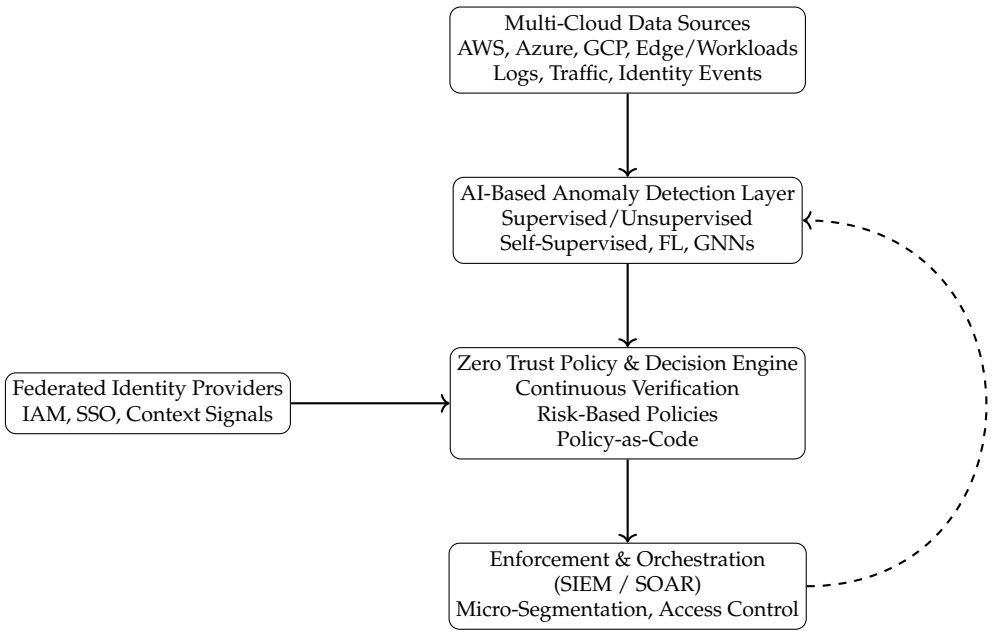

Figure 6. Conceptual reference architecture for AI-driven anomaly detection in multi-cloud zero-trust environments.

The architecture is structured in five integrated layers that work together to help with anomaly detection and zero-trust enforcement across multi-cloud environments. At the foundation, a layer of data collection pools various types of telemetry (including network traffic, application logs, identity and access events and system metrics) from multiple cloud service providers (i.e., AWS, Azure, Google Cloud). Importantly, this process of collection maintains provider autonomy and data sovereignty, which is consistent with the privacy and regulatory issues raised in RQ4.

Building on this telemetry, the anomaly detection layer utilizes a range of AI techniques discussed in the literature, including supervised and unsupervised learning methods, deep learning, self-supervised methods, federated learning, and graph based models. Among these, federated learning is of particular importance, since it makes it possible to perform cross-cloud behavioral analysis without having to share raw data, in turn facilitating compliance with regulations such as the GDPR and HIPAA.

The policy and decision layer then translates the AI-generated risk signals into actionable zero-trust decisions. This is made possible by enabling continuous verification, risk-aware access control, and policies such as code enforcement across cloud environments. Federated identity management can be supported by standards such as OIDC and SAML, which support uniform authentication and authorization with contextual risk information. Finally, the enforcement and orchestration layer uses SIEM and SOAR capabilities to automate response actions such as session revocation, workload isolation, and micro-segmentation across distributed multi-cloud infrastructures.

Illustrative Use Case: Cross-Cloud Identity-Based Anomaly Detection

To provide practical context, an illustrative use case is presented to demonstrate how the proposed architecture would operate in a realistic multi-cloud scenario.

Scenario: An enterprise has workloads that are running in AWS, Azure, and Google Cloud using federated identity management. A developer account is hacked, which leads to an abnormal series of access attempts and lateral movements over cloud providers.

This use case directly supports RQ3 in terms of showing the integration between AI-based anomaly detection and zero-trust enforcement mechanisms, and also supports RQ4 in showing how privacy, identity federation, and compliance issues are handled in multi-cloud environments. While individual elements of this architecture have been assessed in the past, the literature can be seen to have a gap on end-to-end implementation and large-scale validation, stimulating future research directions (RQ5). Table 5 presents the mapping between the illustrative use case and the reference architecture.

Table 5. Mapping between the illustrative use case and the reference architecture.

Architecture Layer	Use Case Description
Data Collection Layer	Identity logs, API access logs, and network telemetry are collected across the cloud providers, revealing abnormal locations and patterns of login and access.
Anomaly Detection Layer	Federated and self-supervised AI model detects deviation in behavior on a cross-cloud basis without sharing raw identity or workload data.
Policy and Decision Layer	The zero-trust policy engine considers the anomaly scores together with the contextual risk and classifies the session as high risk based on the continuous verification.
Identity Provider Layer	Federated identity services enforce step-up authentication and limit access to sensitive resources.
Enforcement and Orchestration Layer	SIEM is used to correlate the alerts while SOAR playbooks revoke the sessions, isolate affected workloads, and trigger automated incident response.

8. Conclusions and Future Work

This paper presented a systematic literature review of the latest research on AI-based anomaly detection in multi-cloud environments under the zero-trust framework. In total, 22 primary studies published between 2020 and 2025 are analyzed. The review is aimed at identifying the level of zero-trust integration along with relevant artificial intelligence techniques, datasets and evaluation practices, and major challenges limiting deployment in real-world scenarios. The results show that techniques based on supervised learning and deep learning are the most widely used in anomaly detection, while emerging approaches such as federated learning, graph neural networks, and self-supervised learning have also attracted interest because of their potential benefits in terms of scalability, privacy preservation, and adaptability in distributed cloud environments. However, it is also apparent from the review that there is strong reliance on a small number of benchmark datasets as well as on controlled or single-cloud experimental configurations, which limits generalization of the reported results to real multi-cloud deployments. From the perspective of zero trust, only a few studies have reviewed zero-trust architectures that incorporate continuous verification, identity federation, micro-segmentation, and policy enforcement. Most works have dealt with anomaly detection as an isolated part of the overall system, not as part of an integrated zero-trust system. In addition, explicit integration of regulatory and compliance requirements such as GDPR, HIPAA, CCPA, PDPL, and NCA ECC/CCC into AI-driven anomaly detection pipelines is largely missing, even though there has been plenty of discussion about compliance at a conceptual level. Based on architectural patterns across the reviewed literature, this study also synthesizes a conceptual reference architecture for AI-driven anomaly detection in multi-cloud zero-trust environments. We emphasize that this reference model for zero-trust policies emphasizes is a logical combination of AI-based

detection layers, identity providers, a zero-trust policy engine, and enforcement points. It should be noted that this reference architecture is derived from the literature review, and is not implemented or experimentally evaluated within the present work.

Future Work and Directions

A number of research directions open up from the gaps identified in this review.

In architectural and systems research, future research could aim at implementing and evaluating the proposed reference architecture in a real multi-cloud environment. This would highlight the integration between AI, zero-trust policies, and operation by leveraging security controls system such as Security Event Information Management (SIEM) and Security Orchestration Automation and Response (SOAR) systems.

In the context of privacy, compliance, and governance, there is a clear need for privacy-preserving and federated anomaly detection mechanisms that explicitly meet a set of regulatory frameworks such as GDPR, HIPAA, CCPA, PDPL, and NCA ECC/CCC.

For benchmarks and datasets, future research should focus on creating standard multi-cloud and zero-trust environments which include datasets that capture traffic for identifying events, policy decisions, and network packets. Availability of open and reusable data would improve comparison and benchmarking between different research studies.

Overall, this review offers a structured and comprehensive survey for understanding the state of AI-based anomaly detection in multi-cloud zero-trust environments. In addition, it identifies key research opportunities for developments in safe, scalable, and compliant cloud security architectures towards the construction of robust cybersecurity frameworks in light of the constantly increasing complexity of the threat landscape.

Author Contributions: Conceptualization, Z.A., A.A. methodology, Z.A.; formal analysis, Z.A.; investigation, Z.A.; writing—original draft preparation, Z.A.; writing—review and editing, Z.A., A.A.; visualization, Z.A.; supervision, A.A.; project administration, A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia [Grant No. KFU263044].

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Acknowledgments: The authors acknowledge the support of the Deanship of Scientific Research, Vice Presidency for Graduate Studies and Scientific Research, King Faisal University, Saudi Arabia [Grant No. KFU263044].

Conflicts of Interest: All authors declare no conflicts of interest. The authors disclose that reference [33] includes co-author A. Albuali, who is also an author of the present manuscript. This prior work is cited transparently as part of the reviewed literature.

References

1. Ghasemshirazi, S.; Shirvani, G.; Alipour, M.A. Zero Trust Applications, Challenges, and Opportunities. *arXiv* **2024**, arXiv:2309.03582.
2. Rose, S.; Borchert, O.; Mitchell, S.; Connelly, S. *Zero Trust Architecture*; Technical Report NIST SP 800-207; National Institute of Standards and Technology: Gaithersburg, MD, USA, 2020. [CrossRef]
3. Sarkar, S.; Choudhary, G.; Shandilya, S.K.; Hussain, A.; Kim, H. Security of Zero Trust Networks in Cloud Computing: A Comparative Review. *Sustainability* **2022**, *14*, 11213. [CrossRef]
4. Parisa, S.K.; Banerjee, S.; Whig, P. AI-Driven Zero Trust Security Models for Retail Cloud Infrastructure: A Next-Generation Approach. *Int. J. Sustain. Dev. Field IT* **2023**, *15*, 15.

5. Lian, H.; Wang, X.; Zhang, C. AI-Powered Anomaly Detection in Cloud Environments: A Lightweight Security Framework under Zero Trust Architecture. *Acad. Nexus J.* **2025**, *4*. Available online: https://ruj.uj.edu.pl/entities/publication/a3906326-ccbe-4584-9510-e8eccc5a51e9 (accessed on 1 June 2026).
6. Jonnakuti, S. Zero-Trust Architectures for Secure Multi-Cloud AI Workloads. *Int. J. Lead. Res. Publ.* **2021**, *2*, 88–97. [CrossRef]
7. Sharma, S. Multi-Cloud Environments: Reducing Security Risks in Distributed Architectures. *J. Artif. Intell. Cyber Secur.* **2021**, *5*, 1–5.
8. Manne, T.A.K. Implementing Zero Trust Architecture in Multi-Cloud Environments. *Int. J. Comput. Eng.* **2023**, *4*, 1–9. [CrossRef]
9. Westerberg, J. Network Security in Zero Trust: Adapting Enterprise Network to Zero Trust. Master's Thesis, Tampere University of Applied Sciences, Tampere, Finland, 2021.
10. Alavizadeh, H.; Alavizadeh, H.; Jang-Jaccard, J. Automation and Orchestration of Zero Trust Architecture: Potential Solutions and Challenges. *Mach. Intell. Res.* **2024**, *21*, 294–317. [CrossRef]
11. Gambo, M.L.; Almulhem, A. Zero Trust Architecture: A Systematic Literature Review. *arXiv* **2025**, arXiv:2503.11659. [CrossRef]
12. Suleiman, B.; Zhang, Y.; Alibasa, M.J.; Youssef, J.; Bao, Y. A Systematic Literature Review on the Implementation and Challenges of Zero Trust Architecture Across Domains. *Sensors* **2025**, *25*, 6118. [CrossRef]
13. Kumar, P. Zero Trust Architecture for Sme Cybersecurity Enhancing Resilience in the Digital Transformation Era. *Int. J. Progress. Res. Eng. Manag. Sci.* **2025**, *5*, 2791–2819.
14. Mosli, R. Optimal Job Role Allocation for Compliance with NCA-ECC Controls Using the Saudi Cybersecurity Workforce Framework. *IEEE Access* **2024**, *12*, 128235–128245. [CrossRef]
15. Talati, D.V. Enhancing Multi-Cloud Security with Quantum-Resilient AI for Anomaly Detection. *World J. Adv. Res. Rev.* **2022**, *13*, 629–638. [CrossRef]
16. Zanasi, C.; Russo, S.; Colajanni, M. Flexible Zero Trust Architecture for the Cybersecurity of Industrial IoT Infrastructures. *Adv. Hoct. Netw.* **2024**, *156*, 103414. [CrossRef]
17. Sun, P.J. Security and Privacy Protection in Cloud Computing: Discussions and Challenges. *J. Netw. Comput. Appl.* **2020**, *160*, 102642. [CrossRef]
18. Zhang, Y.; Suleiman, B.; Alibasa, M.J.; Zheng, X.; Moustafa, N.; Bao, Y. Privacy-Aware Anomaly Detection in IoT Environments Using FedGroup: A Group-Based Federated Learning Approach. *J. Netw. Syst. Manag.* **2024**, *32*, 20. [CrossRef]
19. Nasiruzzaman, M.; Ali, M.; Salam, I.; Miraz, M.H. The Evolution of Zero Trust Architecture (ZTA) from Concept to Implementation. *Int. J. Cyber Secur.* **2025**, *5*, 1–8.
20. Almulla, Z.T.; Rahman, H. The Role of Network Segmentation and Micro-Segmentation in Operational Technology Security. In *Proceedings of the 2025 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)*; IEEE: New York, NY, USA, 2025; pp. 0342–0347. [CrossRef]
21. Ward, R.; Beyer, B. BeyondCorp: A New Approach to Enterprise Security. *Login USENIX Mag.* **2014**, *39*, 6–11.
22. Laghari, A.A.; Khan, A.A.; Ksibi, A.; Hajjej, F.; Kryvinska, N.; Almadhor, A.; Mohamed, M.A.; Alsubai, S. A Novel and Secure Artificial Intelligence Enabled Zero Trust Intrusion Detection in Industrial Internet of Things Architecture. *Sci. Rep.* **2025**, *15*, 26843. [CrossRef] [PubMed]
23. Cui, L.; Qu, Y.; Xie, G.; Zeng, D.; Li, R.; Shen, S.; Yu, S. Security and Privacy-Enhanced Federated Learning for Anomaly Detection in IoT Infrastructures. *IEEE Trans. Ind. Inform.* **2022**, *18*, 3492–3500. [CrossRef]
24. Prabu, K.; Vanitha, A.; Radhakrishnan, R.; Rajakumar, P.; Thenmozhi, E.; Ramesh, S.P. Enhancing Cloud Security with AI-Driven Anomaly Detection for Zero-Day Threats. In *Proceedings of the 2024 1st International Conference on Advances in Computing, Communication and Networking (ICAC2N)*; IEEE: New York, NY, USA, 2024; pp. 1087–1092. [CrossRef]
25. Nass, S.J.; Levit, L.A.; Gostin, L.O. *Beyond the HIPAA Privacy Rule: Enhancing Privacy, Improving Health Through Research*; National Academies Press: Washington, DC, USA, 2009. [CrossRef]
26. European Parliament and Council of the European Union. Regulation (EU) 2016/679 (General Data Protection Regulation). *Off. J. Eur. Union* **2016**, *L119*, 1–88. . . [CrossRef]
27. State of California. *California Consumer Privacy Act of 2018*; California Civil Code §1798.100–1798.199; California State Legislature: Sacramento, CA, USA, 2018.
28. National Cybersecurity Authority (NCA). *Essential Cybersecurity Controls (ECC-1:2018)*; National Cybersecurity Baseline; National Cybersecurity Authority (NCA): Riyadh, Saudi Arabia, 2018.
29. National Cybersecurity Authority (NCA). *Cloud Cybersecurity Controls (CCC-1:2020)*; Official National Cybersecurity Framework; National Cybersecurity Authority (NCA): Riyadh, Saudi Arabia, 2020.
30. Kingdom of Saudi Arabia. Personal Data Protection Law (PDPL), 2021. Royal Decree No. M/19. Available online: https://sdaia.gov.sa/en/Research/Pages/DataProtection.aspx (accessed on 1 June 2026).
31. Nwachukwu, C.; Durodola-Tunde, K.; Akwiwu-Uzoma, C. AI-driven anomaly detection in cloud computing environments. *Int. J. Sci. Res. Arch.* **2024**, *13*, 692–710. [CrossRef]

32. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. [CrossRef]
33. Almajed, H.; Alsaqer, A.; Albuali, A. Towards Effective Anomaly Detection Machine Learning Solutions in Cloud Computing. *IJACSA Int. J. Adv. Comput. Sci. Appl.* **2025**, *16*, 1335–1351. [CrossRef]
34. Bakro, M.; Kumar, R.R.; Husain, M.; Ashraf, Z.; Ali, A.; Yaqoob, S.I.; Ahmed, M.N.; Parveen, N. Building a cloud-IDS by hybrid bio-inspired feature selection algorithms along with random forest model. *IEEE Access* **2024**, *12*, 8846–8874. [CrossRef]
35. Jaber, A.N.; Rehman, S.U. FCMSVM based intrusion detection system for cloud computing environment. *Clust. Comput.* **2020**, *23*, 3221–3231. [CrossRef]
36. Thapa, P.; Arjunan, T. AI-Enhanced Cybersecurity: Machine Learning for Anomaly Detection in Cloud Computing. *J. Emerg. Technol. Innov.* **2025**, *9*, 25–37.
37. Megouache, L.; Zitouni, A.; Sadouni, S.; Djoudi, M. Machine Learning for Cloud Data Classification and Anomaly Intrusion Detection. *Rev. Des Sci. Et Technol. De L'Information-Ser. ISI* **2024**, *29*, 1809–1819. [CrossRef]
38. Moreira, D.A.; Marques, H.P.; Costa, W.L.; Celestino, J.; Gomes, R.L.; Nogueira, M. Anomaly detection in smart environments using AI over fog and cloud computing. In *Proceedings of the IEEE Consumer Communications & Networking Conference (CCNC)*; IEEE: New York, NY, USA, 2021.
39. Micheal, L. Integrating AI-Powered Anomaly Detection with Zero-Trust Authorization for Cloud APIs. *ResearchGate* 2025, *unpublished manuscript*.
40. Narang, A.; Gogineni, L. Zero-Trust Security in Intrusion Detection Networks: An AI-Powered Threat Detection in Cloud Environments. *Int. J. Sci. Res. Mod. Technol.* **2025**, *4*, 60–70. [CrossRef]
41. Ofili, B.T.; Erhabor, E.O.; Obasuyi, O.T. Enhancing Federal Cloud Security with AI Zero Trust, Threat Intelligence and CISA Compliance. *World J. Adv. Res. Rev.* **2025**, *25*, 2377–2400. [CrossRef]
42. Sunkara, G. Implementing Zero Trust Architecture in Modern Enterprise Networks. *SAMRIDDHI A J. Phys. Sci. Eng. Technol.* **2025**, *17*, 1. [CrossRef]