

AI-DRIVEN ZERO TRUST SECURITY FRAMEWORK FOR DETECTING AND PREVENTING CYBER THREATS IN MULTI-CLOUD ENVIRONMENTS

Asiya Ali Khan^{*1}, Danish Hameed²

^{*1}National University of Sciences and Technology (NUST), Pakistan

²CTO, Arhamsoft Pvt. Ltd.

aalikhan.bee22seecs@seecs.edu.pk^{*1}, danish.hameed@arhamsoft.com²

DOI: <https://doi.org/10.5281/zenodo.23310945>

Keywords

Artificial Intelligence, Zero Trust Architecture, Multi-Cloud Security, Cyber Threat Detection, Adaptive Security, Identity-Centric Security, Cybersecurity Governance.

Article History

Received: 25 August 2026

Accepted: 09 October 2026

Published: 12 October 2026

Copyright @Author

Corresponding Author: *

Asiya Ali Khan

Abstract

Multi-cloud computing disperses applications, data, identities, workloads, and services throughout multiple cloud providers/administrative domains. While this architecture can be flexible, scalable, and specialize services, it can present security issues with different infrastructures, different controls, widely dispersed identities, communication between clouds, and policy enforcement complexity. This conceptual study explores the potential of combining Artificial Intelligence (AI) capabilities with Zero Trust Architecture (ZTA) to enhance adaptive threat detection and prevention in multi-cloud environments. The study is a qualitative exploratory study using documentary and secondary sources, only scholarly literature and cyber security authorities' guidance. Thematic analysis and synthesis of documentary evidence were conducted around five themes: multi-cloud security complexity, Zero Trust and identity-centric security, AI-supported threat detection and adaptive response, AI-Zero Trust integration, governance, explainability, privacy, human oversight and resilience. AI can be used to deliver analytical capabilities for detecting anomalous patterns, interpreting security telemetry, and assisting with contextual risk assessment, while Zero Trust offers principles and mechanisms to ensure continuous verification, least-privilege access, policy enforcement, and resource protection. The resulting framework is a continuous loop that connects multi-cloud telemetry, AI powered analysis, contextual verification, risk-adaptive access, least-privilege enforcement, micro-segmentation, adaptive response and continuous monitoring. The cross cutting safeguards are governance, privacy, explainability, human oversight and resilience. The framework is theoretical and has not been empirically applied or statistically proven, and there is no claim of security improvements being measured. The study offers a conceptual framework to support future experimental and organizational studies in the area of AI-powered Zero Trust security for multi-cloud environments.

INTRODUCTION

Conventional cyber security controls were built on the premise of a single central enterprise network, which is no longer the case. Applications, workloads, data, identities and services can now be spread out among geographically separated infrastructures, and users, devices and services can now access from anywhere. In the Zero Trust approach, the network perimeter disappears and users, assets, and resources are protected without any implicit trust being assumed because of network location or ownership (Rose et al., 2020). This architecture change is especially important for multi-cloud solutions, where security policies need to be effective and relevant to the diversity of infrastructure and administration.

Multi-cloud deployments could include multiple cloud service providers, application programming interfaces (APIs), identity and access management systems, virtual machines, containers, service meshes, remote connections, and third-party services. Data and workloads can also be transferred from one environment to another as needed. In a recent systematic literature review conducted by Diningrat et al. (2025), some common issues in multi-cloud security were identified: Data Protection, Cryptography, Authentication and Authorization, Service Availability, Confidential Computing, and Virtualization. These are just some of the reasons behind the fact that multi-cloud security cannot only be achieved by applying perimeter security.

There's an architectural and governance view of this distributed environment, and that's Zero Trust. The core of its concept is that trust is not something that should be assumed based on the network's structure. In contrast, the evaluation of access to resources should be done through explicit authentication and authorization, contextual information, least-privilege principles, and continuous monitoring (Rose et al., 2020). The cloud-native extension of the model adds to the original model by focusing on application/service identities, API gateways, service meshes, sidecar proxys, and granular

policy enforcement in multi-cloud and hybrid configurations (Chandramouli & Butcher, 2023). An AI tool offers a potentially complementary analytical capability. These capabilities include anomaly detection, behavioral analysis, threat-intelligence processing, and interpretation of vast amounts of security telemetry, which all can benefit from machine-learning and other AI techniques. Yet, AI systems themselves carry a risk of data quality, model behaviour, adversarial manipulation, transparency, privacy, and reliability issues. The NIST AI Risk Management Framework focuses on trustworthy and responsible AI and includes qualities of validity and reliability, security and resilience, transparency, explainability and interpretability, privacy and accountability (Tabassi, 2023). In addition, NIST's adversarial machine-learning taxonomy also captures the security vulnerabilities that can result from an attack or manipulation of a machine-learning system (Vassilev et al., 2024).

This study's research problem is thus conceptual, not experimental, and asks how analytical capabilities powered by AI and Zero Trust can be structured into a coherent security framework for multi-cloud, heterogeneous environments without devoting an entire component to AI, treating AI as an autonomous alternative to governance or human decision-making? Previous literature has dealt with topics such as multi-cloud security, Zero Trust, detection, identity, privacy or governance using an AI approach as distinct areas of concern. This study unites the above in one conceptual framework.

1.1 Research Questions

- 1) Which of the following security issues are frequently mentioned in the literature about multi-cloud environments?
- 2) How do capabilities powered by AI fit into the Zero Trust paradigm for cyber-threat detection and prevention?
- 3) How does AI and Zero Trust fit together in multi-cloud environments and what technical, organizational, governance, privacy and trustworthiness challenges does that pose?

4) What are the principles of a conceptual AI enabled Zero Trust security framework for multi-cloud environments?

1.2 Contribution and Scope

This study is not meant to present a new detection algorithm, software implementation or statistically validated security model but it is meant to present a conceptual framework that is supported by theory. The framework establishes the interrelationships among multi-cloud telemetry, AI-driven SI, continuous verification, risk-adaptive access control, least privilege, micro-segmentation, adaptive response, and governance protection. As the evidence base is documentary and secondary, it is not possible to estimate the accuracy of detection, the false-positive rate, the attack prevention rate, nor the return on security investment.

This module covers a literature search and the conceptual foundations.

2.1 Multi-Cloud Security as a Distributed Security Problem

There is more to multi-cloud security than just securing individual cloud platforms. The security context includes the relationships among identities, APIs, workloads, applications, devices, data flows, services, and administrative domains. The systematic review conducted by Diningrat et al. (2025) shows that multi-cloud security issues can be broken down into multiple technical domains, such as data security, authentication and authorization, availability, confidential computing, and virtualization. This distribution makes it even more significant to have consistent security policies and visibility between providers. Cloud security taxonomies also highlight the importance of looking at cloud security in a comprehensive manner as opposed to a set of individual controls (Mthunzi et al., 2020). In a multi-cloud scenario, authentication and data integrity are also critical aspects, since identities and data may span across multiple services and infrastructure domains (Megouache et al., 2020). The observations support the notion that the multi-cloud is not just a set of “different” clouds

but rather the context in which the proposed framework is to be realized.

This section covers Zero Trust Architecture and Identity-Centric Security. In this section, students will learn about Zero Trust Architecture and Identity-Centric Security.

NIST SP 800-207 defines Zero Trust Architecture as a method that shifts security from the static network perimeter to the protection of users, assets, services, workflows and resources. There is no implicit trust based on network location or ownership and authentication and authorization are distinguished from assumptions about network position (Rose et al., 2020).

NIST SP 800-207A also adds application and service identities to user identities and outlines API gateways, service meshes, and granular identity-tier policies as examples of mechanisms that can be used to achieve the objectives of the guidelines (Chandramouli & Butcher, 2023). This is directly relevant to multi-cloud environments as machine-to-machine interactions, microservices or distributed workloads can require a policy decision that goes beyond the user or network location.

Understanding Zero Trust as an architectural and governance approach, rather than a single product, is key. It must be implemented with identity management, policy definition, authentication, authorization, monitoring, segmentation and organizational governance. The implementation guidance by NIST also outlines Zero Trust as Secure Access, Risk-based Controls, Continuous Inspection, Monitoring and Logging (Kerman et al., 2020).

2.3 AI for Cybersecurity

Cyber-security applications such as anomaly detection, intrusion detection, behavioral analysis, malware analysis, and security-data processing have been explored using AI and machine-learning techniques. In terms of the efficacy of using machine learning and AI in cyber security, the literature reviews suggest that these methods can be applied to the analysis of complex security information, but their performance is dependent on the data, the way

the model is designed, how it is deployed, and the evaluation protocols (Dasgupta et al., 2020; Naik et al., 2022; Nassif et al., 2021).

The security of the AI system should also be taken into account. In adversarial machine-learning research, an attack class is defined by the specific adversarial attack scenario and the considerations for mitigating that type of attack are identified for the machine-learning system lifecycle (Vassilev et al., 2024). Another challenge for cybersecurity is explainability, that is, the security analysts can want to know the reasons behind an alert or the security analyst's recommendation made by the model (Charmet et al., 2022; Rjoub et al., 2023). As such, the framework views an AI capability as a tool to aid analysis, not supplant it, within the context of security governance and human oversight.

2.4 AI-Zero Trust Complementarity

AI and Zero Trust are complementary concepts. AI can analyze the telemetry and detect patterns that could be relevant to a contextual risk assessment, and Zero Trust provides the architectural principles that are used to evaluate and enforce access. This model does not equate an AI output to an authorization decision. However it becomes one factor of a policy guided process that integrates policy, contextual info, resource sensitivity, device and identity.

This distinction is significant since the conceptual framework should not become a self-serving claim of autonomous security. AI has been already adopted for analytical and decision-support roles (Tabassi, 2023), and the concept of trustworthy-AI includes risk management, transparency, security, privacy, and accountability.

3. Conceptual Framework

The proposed framework brings together the three areas identified in the study: multi-cloud infrastructure as the operating context, AI as an adaptive analytical layer and Zero Trust as the foundation for decisions and enforcement based on identity. The process is cyclical instead of linear because information about security

gathered by enforcement and monitoring can be added to the process for further analysis.

The conceptual cycle looks like this: Multi-Cloud Environment → Security Telemetry and Signals → AI-Supported Security Intelligence → Continuous Identity and Contextual Verification → Risk-Adaptive Access Decision → Least-Privilege Enforcement and Micro-Segmentation → Adaptive Prevention and Response → Continuous Monitoring → Feedback to Security Intelligence. The entire cycle is governed by principles of governance, privacy, explainability, human oversight, interoperability and resilience.

The desired connectivity between the elements of continuous verification, least privilege, explicit authorization, AI-assisted behavioural assessments, contextual risk, adaptive access control, micro-segmentation, adaptive prevention and response, and cross-cutting principles of governance, privacy and explainability, resilience, and human oversight is shown in Figure 1. The figure is only a conceptual figure and not an implemented or tested system.

3.1 Framework Layers

The first layer is Multi-Cloud Infrastructure and Data Environment (MCI&D). The first layer is Multi-Cloud Infrastructure and Data Environment (MCI&D).

The first layer is the operational environment which includes cloud platforms, workloads, applications, APIs, identities, devices, services and data flows. These elements can be spread across multiple providers and administrative domains, leading to the production of diverse security telemetry which needs to be read within its context.

Layer 2: AI-Supported Security Intelligence

The second layer has analytical capabilities like anomaly detection, behavior analysis, threat-intelligence processing, context-based analysis and risk prioritization. It is used to structure the details of security information and to try to locate any potentially significant patterns. It is not a stand-alone security authority.

Eliminate trust in your network by using Layer 3 of the Zero Trust Security Model.

Layer 3: The third layer is based on Identity-centric security principles

How access is limited is determined by continuous verification, explicit authorization, least privilege, contextual access decisions, policy evaluation, micro-segmentation. The decisions can be informational based on AI, but the framework still provides for policy and governance mechanisms between the analytical output and enforcement.

Layer 4: Adaptive Prevention and Response

The fourth layer maps security decisions to controlled actions. Actions can range from access restriction, isolation, altering policy conditions,

prioritizing alerts, launching incident response processes or increasing monitoring based on organizational policy and context. These actions create extra analytes information for the analytical layer.

Layer 5: Governance, Explainability, Privacy, Human Oversight, and Resilience

The fifth layer provides protection for the whole cycle. It incorporates accountability, transparency, explainability, protection of privacy, considerations for compliance, human decision authority, and resilience. This is in line with NIST's AI RMF, which considers trustworthiness a socio-technical issue rather than just a technical characteristic (Tabassi, 2023).




Figure 2. Infrastructure enforcement architecture for heterogeneous multi-cloud security.

Figure 2 provides an architectural view of decentralized workloads, endpoint and workload security, inter-cloud communication and access management, data privacy and regulatory controls, cloud-native security integration, and centralized governance. It should be interpreted as a conceptual mapping of control domains rather than a specification for a particular vendor implementation.

4. Research Methodology

4.1 Research Philosophy

An interpretivist research philosophy is used in the study because it focuses on the linkages between concepts across technical and

organizational literature, including concepts related to trust, identity, risk, explainability, governance, and adaptive security. These concepts are not reduced to a single measurable

variable nor tested for a causal relationship in the study.

4.2 Research Design

A qualitative, exploratory documentary research design was used. The design is suitable for a conceptual synthesis, where the main purpose is to synthesize arguments and principles from established scholarly and authoritative sources. The manuscript does not include a systematic review of literature because the original evidence record does not include a reproducible protocol for searching a database, a complete screening flow and a quantified final study sample is not available. This distinction is of key importance for methodological accuracy.

4.3 Sources of Evidence

The evidence base is comprised of peer-reviewed research, scholarly reviews, cybersecurity standards, government guidance, and authoritative technical publications. NIST publications related to Zero Trust, AI Risk Management, Cybersecurity Risk Management and Cloud-native access control are given special attention. Contextualisation of the use of AI in cybersecurity, multi-cloud security, explainability and related technical issues by scholarly sources. Relevance to the research questions, conceptual contribution, methodological/institutional authority, and relevance to today's cybersecurity were used to assess the evidence. Established concepts were supported by the use of foundational sources, while the topic is evolving more rapidly, more recent sources were used. As this is a conceptual study, a numeric claim for the number of items studied is not made.

4.4 Documentary Analysis

The documents were taken to be a source of concepts, arguments, principles, risks and proposed controls. The literature review focused on convergence and divergence within the field, specifically with the multi-cloud security literature, the Zero Trust model, the detection capabilities enabled by AI, identity and access management, privacy, explainability, governance

and resilience. The goal was to find relationships that could be the basis for a coherent conceptual framework, and not to estimate effect sizes or predictive performance.

4.5 Thematic Analysis

Thematic analysis approach was used in the analysis based on Braun and Clarke's work on thematic analysis. The process was characterized by the familiarization with the evidence, meaning-making, coding, development and review of candidate themes, refining the thematic structure, and interpretive synthesis. Themes were chosen if they were conceptually relevant to the research questions, not because of the number of times they appeared. (Braun & Clarke, 2021).

4.6 Analytical Themes

Multi-cloud security complexity: distributed resources, diverse infrastructures, APIs, identities, data flows, dependent services and cross-provider risks.

Zero Trust and Identity centric Security: Continuous verification, explicit authorization, least privilege, contextual access, policy enforcement, segmentation and monitoring.

3.AI-driven detection and adaptive response: anomaly detection, behavioral analysis, threat intelligence processing, risk analysis, adaptive monitoring and the limitations of AI.

4.AI-Zero Trust integration: How security telemetry, AI-enhanced contextual interpretation, access decisions, access enforcement and feedback all work together.

Governance and trustworthiness: explainability, privacy, accountability, human oversight, interoperability, resilience and responsible governance.

4.7 Trustworthiness and Rigor

Triangulation between scholarly and authoritative sources, conceptual consistency, clear links between research questions and themes and clear distinctions between documented evidence and the researcher's conceptual synthesis provided support for rigor.

This study does not assert that outputs generated by AI will necessarily be trusted and does not claim to be empirically tested with the proposed framework.

4.8 Scope and Limitations

This study uses only secondary documentary sources. Not a single observation of primary organization, not a controlled experiment, not an operational deployment, not a statistical validation, not a performance benchmarking are included. As a result, the framework will not be able to assert detection accuracy, prevent false positives, preventiveness, latency, cost savings or causal improvements in security. Also, because there is no reproducible systematic-review protocol, the study should not be interpreted as a comprehensive review of all available literature.

5. Findings and Discussion

Thematic synthesis resulted in five related findings. As a whole, they enable a conceptual model where multi-cloud infrastructure becomes the security context, AI becomes the adaptive analytical support, and Zero Trust becomes the set of principles and enforcement mechanisms – the identity-centric ones. Findings are understood as a conceptual linkage of the evidence reviewed and not as empirical measurements.

This section includes the following topics: Multi-Cloud Complexity, Distributed Cyber Risk.

The literature suggests that a multi-cloud configuration spreads security relevant assets and interactions over providers, locations, services, identities, and technologies. Diningrat et al. (2025) list data, cryptography, authentication and authorization, availability, confidential computing, and virtualization as some of the areas that are of repeated concern in multi-cloud security. In addition, NIST's cloud-native Zero Trust guidance acknowledges the need to have granular access control in multi-cloud and hybrid environments (Chandramouli & Butcher, 2023).

The implication in concept is that the security controls are not provider specific security measures. Infrastructure boundaries must have consistent identity, policy, monitoring and

response capabilities. The security decisions, in this case, do not need to be the same for all providers, but must be interpretable and enforceable for different providers and heterogeneous environments.

5.1. Zero Trust as the Identity-Centric Security Foundation

The other theme has to do with the applicability of Zero Trust to distributed environments. The ZTA model developed by NIST does not rely on implicit trust based on the network location, but focuses on protection of users, assets, services, workflows and resources (Rose et al., 2020). Similarly, this applies to cloud-native application and service identities, API gateways, and service meshes (Chandramouli & Butcher, 2023).

The synthesis places Zero Trust as the foundation for decision and enforcement of the framework. Continuous verification and least privilege aren't put forward as technologies that can ensure a system's security – they are a set of principles which dictate how authentication, authorization, monitoring, segmentation, and policy enforcement needs to work.

5.2. AI-Enabled Threat Detection and Adaptive Security

The third theme relates to the analytical function of AI. The literature reviewed shows that artificial intelligence and machine learning tools can help process complex security data and assist with various security-related tasks, including anomaly detection and behavioural analysis, threat intelligence, etc. But the red flags are there and not using AI is not the same as believing it to be infallible or as if it works on its own. Data quality, model design, deployment environment, and the security of the AI lifecycle affect the performance of AI (Dasgupta et al., 2020; Nassif et al., 2021; Naik et al., 2022).

The definition of AI in the framework is thus a layer of analysis. It can support contextual risk interpretation and prioritization, but policy, identity controls and governance are still needed. This will also take care of threats like adversarial

manipulation and unclear model conduct (Vassilev et al., 2024; Charmet et al., 2022).

The introduction of AI into the Zero Trust landscape is an ongoing process that can be viewed as a risk cycle.

The fourth theme relates the analytical capabilities of AI and the decision and enforcement rationale of Zero Trust. Multi-cloud activity generates security signals, AI can help organize and interpret those signals, identity and context can be reassessed and Zero Trust policies decide whether access is allowed, limited, or put under further controls. Monitoring then provides additional information, a feedback loop.

Conceptual link between detection and prevention is provided in this relationship. Importantly, the framework does not assert any claim that the decisions made using AI are inherently more secure. The suggested value is architectural; when analyzing data is not isolated in a detection system, but instead is made part of a controlled Zero Trust decision process.

6. Explainability, Privacy, Governance, Human Oversight and Resilience

The fifth theme sets the parameters for the operation of AI-powered Zero Trust. Continuous monitoring can include context attributes, device attributes, behavioral attributes, and context signals. The data that are collected and processed can have security implications, but privacy and accountability are also issues. NIST's AI RMF focuses on the trustworthy features of AI systems which include validity and reliability, security and resilience, transparency, explainability and interpretability, privacy, and accountability (Tabassi, 2023).

Explainability has a special significance when AI results affect security operations. An analyst might want to know why a particular event was considered suspicious, or why an access recommendation was made. There is hence still an element of human supervision for unusual, critical and unclear circumstances. Resilience is also addressed as a cross-cutting requirement: A multi-cloud security architecture must continue to deliver critical security services in the event of

compromise or failure of components, services, or identities.

The next question concerns the integrated framework and practical interpretation. The next is related to integrated framework and practical interpretation.

The five themes can be combined into one socio-technical security cycle. Security telemetry is generated by multi-cloud infrastructure, potentially meaningful patterns identified by AI-supported intelligence, evaluation of identity and context by Zero Trust, risk-adaptive access control applied to access, bounded exposure via least privilege and micro-segmentation, application of authorized protective actions via adaptive response, and continuous monitoring that provides new information to the cycle. This is constrained and guided by governance, privacy, explainability, human oversight, interoperability and resilience.

6.1 Conceptual Process

7. Distribute and standardize pertinent security telemetry sourced from multi-cloud resources, identities, applications, workloads, APIs, devices and services.

7. Apply analytical methods (with the support of AI) to detect abnormal behavior, relationships between events and contextual risk indicators.

Before allowing continued access, re-evaluate identity, device, resource and contextual information.

Implement risk-based and least-privilege authorisation based on organisational policy.

Implement micro-segmentation and other enforcement measures to restrict unauthorized lateral movement.

Implement preventative and/or responsive policy actions as needed when conditions warrant them.

12. Keep monitoring and provide analytical and decision layer with new evidence.

Ensure governance, privacy, explainability, accountability, human oversight, interoperability and resilience throughout.

6.2 Practical Design Principles

Access control should be based on identity, and not just network location.

Unless there is a valid policy that is specifically approved for automated action, the output of AI systems should not be considered security intelligence or decision support.

The permissions that can be accessed by users, services, workloads and applications should be limited in scope by adhering to the principle of least privilege.

Cross-cloud policies should be formulated with enough level of granularity to enable coherent enforcement across the cloud and flexibility in implementation details.

Monitoring should feedback to security operations and risk assessment processes.

High impact or ambiguous decisions involving AI should be reviewed by humans.

Privacy and data-governance requirements should be taken into account when handling behavioral and identity telemetry.

Resilience needs to be able to continue to provide security services if one or more cloud components, or cloud services, are degraded or compromised.

7. Research implications and future work.

The framework offers a number of directions for empirical research. First, future studies can be conducted in controlled multi-cloud testbeds to evaluate the precision of the detection, recall, false positive rate, response latency and resource overhead of the proposed architecture. Second, experiments can be conducted to benchmark static policy enforcement against AI-driven access decisions that use risk assessment. Third, researchers can research techniques to explain security alerts and access recommendations. Finally, research can study the implications of privacy-preserving methods on the usefulness of behavioral telemetry. Lastly, governance, human oversight, interoperability, and resilience can be assessed in real deployments.

Future studies are required since the current framework is conceptual. Current evidence indicates that there are connections between the

concepts, but fail to prove a numerical amount of increased security as a result of the proposed integration.

8. Conclusion

This research created a conceptual framework of Artificial Intelligence in the context of Zero Trust Security Model to solve the issue of Multi-cloud environments using a qualitative synthesis of scholarly and authoritative documentary evidence. The analysis highlighted five interrelated themes: Multi-Cloud Complexity, Zero Trust and identity-centric security, AI-powered threat detection, AI-ZT integration, Governance, explainability and privacy, and human oversight and resilience.

The synthesis takes into account a complementary division of roles. AI can help with the context interpretation of security telemetry, help find potentially anomalous patterns, and help with context interpretation of risk. Zero Trust offers the principles and enforcement mechanisms that allow identities, devices, services and resources to be constantly checked and limited. The framework thus views AI as an adaptive analytical capability within a larger Zero Trust security framework instead of it being an independent security governance solution.

The main constraint is methodological; the model has not been applied, tested and statistically tested. It is thus to be taken as a conceptual contribution. Empirical studies are needed to measure the impact of the proposed integration on measurable security outcomes and to understand how explainability, privacy, interoperability, human oversight, governance, and resilience impact on operational performance.

REFERENCES

- Braun, V., & Clarke, V. (2021). Can I use TA? Should I use TA? Should I not use TA? Comparing reflexive thematic analysis and other pattern-based qualitative analytic approaches. *Counselling and*

- Psychotherapy Research.
<https://doi.org/10.1002/capr.12360>
- Charmet, F., Tanuwidjaja, H. C., Ayoubi, S., Heydari, S. S., & others. (2022). Explainable artificial intelligence for cybersecurity: A literature survey. *Annals of Telecommunications*, 77, 789–812. <https://doi.org/10.1007/s12243-022-00926-7>
- Chandramouli, R., & Butcher, Z. (2023). A Zero Trust Architecture model for access control in cloud-native applications in multi-location environments (NIST Special Publication 800-207A). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207A>
- Dasgupta, D., Akhtar, Z., & Sen, S. (2020). Machine learning in cybersecurity: A comprehensive survey. *The Journal of Defense Modeling and Simulation*, 19(1), 57–106. <https://doi.org/10.1177/1548512920951275>
- Diningrat, D. C., Suhardi, & Rahardjo, B. (2025). Security issues in multi-cloud: A systematic literature review. *IEEE Access*, 13, 70006–70017. <https://doi.org/10.1109/ACCESS.2025.3561352>
- Kerman, A., Borchert, O., Rose, S., & Division, E. (2020). Implementing a Zero Trust Architecture (NIST Cybersecurity White Paper 20). National Institute of Standards and Technology.
- Megouache, L., Zitouni, A., & Djoudi, M. (2020). Ensuring user authentication and data integrity in multi-cloud environment. *Human-Centric Computing and Information Sciences*, 10, Article 15. <https://doi.org/10.1186/s13673-020-00224-y>
- Mthunzi, S., Benkhelifa, E., Bosakowski, T., Ghedira Guegan, C., & Barhamgi, M. (2020). Cloud computing security taxonomy: From an atomistic to a holistic view. *Future Generation Computer Systems*, 107, 620–644. <https://doi.org/10.1016/j.future.2019.11.013>
- Naik, B., Mehta, A., Yagnik, H., & others. (2022). The impacts of artificial intelligence techniques in augmentation of cybersecurity: A comprehensive review. *Complex & Intelligent Systems*, 8, 1763–1780. <https://doi.org/10.1007/s40747-021-00494-8>
- Nassif, A. B., Talib, M. A., Nasir, Q., Albadani, H., & Dakalbab, F. M. (2021). Machine learning for cloud security: A systematic review. *IEEE Access*, 9, 20717–20735. <https://doi.org/10.1109/ACCESS.2021.3054129>
- National Institute of Standards and Technology. (2020). Security and privacy controls for information systems and organizations (NIST Special Publication 800-53, Revision 5). U.S. Department of Commerce.
- National Institute of Standards and Technology. (2024). NIST Cybersecurity Framework 2.0. U.S. Department of Commerce. <https://doi.org/10.6028/NIST.CSWP.29>
- Rjoub, G., Bentahar, J., Abdel Wahab, O., Mizouni, R., et al. (2023). A survey on explainable artificial intelligence for cybersecurity. *IEEE Transactions on Network and Service Management*, 20(4), 5115–5140. <https://doi.org/10.1109/TNSM.2023.3282740>
- Rose, S. W., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207>

Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>

Vassilev, A., Oprea, A., Fordyce, A., & Andersen, H. (2024). Adversarial machine learning: A taxonomy and terminology of attacks and mitigations (NIST AI 100-2e2023). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-2e2023>.

