Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 4, Issue 3, 2026

AI-DRIVEN ZERO TRUST SECURITY MODELS FOR PROTECTING CLOUD AND IOT INFRASTRUCTURES

[1]Syed Muhammad Junaid Hassan, [2]Ansar Ahmed, [3]Faheem Ahmed, [4]Kamran Dahri

[1]Assistant Professor, Department of Information Technology, Faculty of ICT, Balochistan University of Information Technology, Engineering and Management Sciences (BUITEMS)
[2]MSCS from Sir Syed University of Engineering and Technology
[3]Department of Information Technology, University of Sindh, Jamshoro
[4]Department of Information Technology, University of Sindh, Jamshoro
[1]smjunaid.it@gmail.com, [2]pitafiansar@gmail.com, [3]faheem.abbasi@usindh.edu.pk, [4]kamran.dahri@usindh.edu.pk

Keywords
Zero Trust Architecture, Artificial Intelligence, Cloud Security, Internet of Things, Cybersecurity, Machine Learning, Threat Detection Identity Verification

Article History
Received on 26 Feb, 2026
Accepted on 17 March, 2026
Published on 19 March, 2026

Abstract
The fast increase in the number of cloud computing and Internet of Things (IoT) devices has radically changed the infrastructure of organizations, and it has also increased the area of the attack that can be used by adversaries. The conventional perimeter-based security models have become more and more insufficient in protecting against the increasingly complex multi-vector threats in these dynamic environments. This paper investigated how AI-driven Zero Trust security models help secure cloud and IoT environments by using a qualitative, exploratory research design. Fifteen cybersecurity professionals, such as cloud security engineers, IoT specialists, and IT managers, took part in semi-structured interviews and were complemented by analyzing documents on industrial reports, white papers, and organizational security policies. Thematic analysis was used to analyze the data and identified four fundamental themes, namely, the ineffectiveness of traditional security environments, the operational benefit of ongoing AI-driven authentication and surveillance, implementation challenges such as resource limitations and intricacy of integration and (4) future directions toward self-healing, autonomous security environments. Results suggest that AI-enhanced Zero Trust architectures have substantial benefits in terms of improvements in the accuracy of threat detection, shorter mean time to respond, and facilitating adaptive policy implementation in heterogeneous settings. The paper has concluded that smart, context sensitive security structures have to be the primary goal of organizations in order to be resilient to the changing cyber threats. Discussed are implications to practice, policy, and future research.

INTRODUCTION

The online revolution of organizational systems in the last 20 years has generated unmatched efficiency benefits and also created sophisticated security threats that conventional systems are unsophisticated to handle. The cloud computing and the Internet of Things (IoT) have become the two pillars of this change as they allow to provide resources on a scale, perform real-time data analytics, and create interconnected device ecosystems that integrate across industries, such as healthcare and manufacturing, as well as finance and critical national infrastructure (Atlam & Wills, 2019). Nevertheless, these distributed paradigms have completely compromised the architectural assumptions that support the traditional security models, which mostly consist of the idea of a trusted internal network secured by a hardened perimeter.

Perimeter-based security systems are based on the castle-and-moat analogy of network protection which relies on the assumption that all that lies within the organizational perimeter is automatically trustworthy and that threats can be eliminated at the boundary (Kindervag, 2010). This premise fails in cloud and IoT settings where data crosses across administrative boundaries, devices act independently outside of an organization, and users access resources located in geographically disparate locations and on a variety of devices. Well-known attacks such as the SolarWinds supply chain breach, the Mirai botnet campaign targeting IoT-connected devices, and many cloud misconfigurations that caused massive data leaks have shown the devastating impact of this architectural weakness (Sanger & Perlroth, 2021). The economic and reputational harm to such events, which frequently runs in the hundreds of millions of dollars have spurred the cybersecurity community to find more adaptive and resilient defensive paradigms.

Zero Trust Architecture (ZTA) is a notion that was first described in writing by John Kindervag at Forrester Research in 2010 and later codified in NIST Special Publication 800-207 that reflects a paradigm shift in security thought based on the notion of never trust, always verify (Rose et al., 2020). Instead of providing unspoken trust due to network location, Zero Trust requires identity, device posture, and contextual verification prior to providing any resource access. Although this model can be used to overcome the structural shortcomings of perimeter security, its applications in large scale, especially in a heterogeneous cloud and IoT system, require a level of computational power that is currently beyond the capacity of rule-based or manual methods. Artificial intelligence, specifically machine learning, has become the key to operationally feasible Zero Trust implementations, including the ability to do real-time behavioral analysis, anomaly detection, and dynamically enforce policies machine speed and scale (Mehraj and Banday, 2020).

AI convergence with Zero Trust is one of the most important phenomena in the modern cybersecurity practice, but the scholarly literature studying this intersection, especially in the context of cloud and IoT setups, is still piecemeal and largely speculative. There is a significant gap between the discourse and the reality of operations of security professionals as empirical research on practitioner experience is very limited. The present study aimed to fill that gap by investigating, via a qualitative research, the perception and application of AI-based Zero Trust models by cybersecurity professionals in securing cloud and IoT systems.

Problem Statement

Although the idea of AI-based Zero Trust as a potential approach to security gains more and more popularity, organizations have a major issue when it comes to the transformation of the old architecture. They are the difficulty of retrofitting the principles of Zero Trust on the current infrastructure, the resource-intensive nature of ongoing verification, the lack of AI skills among security staff, and the unanswered questions of governance and accountability of automated security decisions. Empirically based studies are urgently required to record the experiences of practitioners, outline the barriers to implementation, and reveal the best practices that can be used to support the organization adoption.

Research Objectives

The goal of the study was to: (1) understand how cybersecurity professionals view the suitability of AI-driven Zero Trust models in a cloud and IoT setting; (2) identify the main challenges and obstacles during the implementation process; (3) discuss the role of particular AI technologies, such as machine learning, behavioral analytics, and anomaly detection, in applying the principles of Zero Trust to a specific organizational security posture; and (4) provide framework.

Study Significance

The study has value to the scholarship of cybersecurity as it offers valuable, qualitative empirical data about practitioners with experience in implementing AI-based Zero Trust solutions. The results directly apply to the practical work of security architects, organizational leaders, and policymakers, as they provide a basis of empirically rigorous information to make decisions in a field where the rate of technological innovation can easily exceed the evolution of evidence-based advice. The research also contributes to the theoretical knowledge of the intersection of AI capabilities with the principles of Zero Trust in meeting the special security needs of cloud and IoT environments.

LITERATURE REVIEW

Cybersecurity Paradigm Shifts

The history of cybersecurity is defined as a continuous dialectic between defensive innovation and adversarial adaptation. The initial network security systems were built on the principle of perimeter defense wherein the firewall, intrusion detection and demilitarized zone served as a protective layer between the trusted internal networks and the untrusted external environments (Cheswick et al., 2003). This model was sufficient in fairly stable, centralized computing settings where organizational.

assets lived within well-defined physical and logical limits. But as computing gradually advances in a direction of progressive decentralization, initially through mobile devices and remote work, and subsequently through cloud computing and IoT, these perimeter assumptions have been upset in profound ways.

The constraints of perimeter security started to be documented in the early 2000s when researchers noted that insider attacks, lateral movement of external intruders who had already broken the perimeter and the increasing permeability of the organizational boundaries had made the castle-and-moat model a relic (Stiennon, 2014). The release of the Zero Trust framework published by Kindervag in 2010 was a pivotal statement of the alternative paradigm, the one which did not presume trust but ongoing verification without regard to the location of the network. The NIST formalization of SP 800-207 (Rose et al., 2020) offered a standardized reference architecture that has led to the popularization of discussions around its use in

both the public and private sectors. Since that time, the Zero Trust model has gained the support of leading regulatory organizations, such as the U.S. Executive Order on Improving the Nation's Cybersecurity (2021), indicating that it has become the paradigm of security in the next decade.

AI in Cybersecurity.

Artificial intelligence penetrated into the practice of cybersecurity in various functional areas, such as threat intelligence, vulnerability management, and security operations, as well as incident response. The basic value proposition of AI in this regard is that it is able to handle and analyze volumes of data that are significantly beyond human cognitive capability, detect subtle patterns of malicious activity, and react to threats more quickly than a machine (Sarker et al., 2021). Supervised classification models, unsupervised clustering algorithms, and reinforcement learning systems have been used in a variety of problems, such as malware detection and phishing identification, network anomaly detection, and user behavior analytics.

Recurrent neural networks and transformer architectures, as deep learning methods, have shown significant effectiveness in learning temporal correlations in network traffic data, allowing slow-moving, long-standing threats to be identified, which signature-based detection methods fail to recognize (Buczak & Guven, 2016). Threat intelligence feeds and dark web monitoring have been subjected to natural language processing to allow automatic removal of indicators of compromise and attribution of threat actor activity. Nevertheless, AI-based security systems do not have no limits. Adversarial machine learning, where threat actors deliberately create inputs to mislead AI classifiers, is an increasing problem, and so is the problem of explainability in high-stakes security decisions, where human oversight is paramount (Papernot et al., 2018).

Zero Trust Architecture: Basic Principles and Building Blocks

The Zero Trust model is based on a number of key principles that make it superior to alternative, perimeter-based models. To begin with, the rule of never trust, always verify requires every access request, irrespective of its source, to be verified, authorized, and constantly verified (Rose et al., 2020). Second, the principle of least privilege access states that users, devices, and services must have the minimum permissions to complete the task at hand and thus reduces the blast radius of a breached account or device. Third, micro-segmentation is a principle because the network is divided into small areas that have stringent access controls, and lateral movement is not possible by attackers who have already compromised one area.

These principles need to be implemented through having a strong identity and access management (IAM) infrastructure, ongoing device health checking, encrypted communications, and real-time policy enforcement engines that can assess many contextual factors at once (Stafford, 2020). One of the main factors that contribute to the implementation of AI in the context of Zero Trust is the complexity of providing these components at scale, especially in the situation when thousands of IoT devices are involved, and the workloads in the cloud are dynamic. Identity analytics based on AI can identify suspicious access patterns that can be an indication of an account compromise, whereas automated policy engines can dynamically update access controls based on evolving risk indicators.

AI-powered Zero Trust in the Cloud

The security issues of cloud environments pose a unique challenge that closely aligns with the operational strengths of AI-based Zero Trust

models. The dynamic provisioning and deprovisioning of cloud resources, the multi-tenancy of shared infrastructure, and the intricate network of API interactions between cloud services make it a changing attack surface, and one not readily manageable by manual security mechanisms (Hashizume et al., 2013). Cloud Native security platforms and AI-based Cloud Access Security Brokers (CASBs) have become essential elements of Zero Trust implementations, offering automated visibility, threat detection, and policy enforcement of cloud workloads.

A study by Chen and Zheng (2022) has shown that machine learning-based anomaly detection in the cloud setting had much higher true positive rates to detect unauthorized data exfiltration than its rule-based counterparts, but with lower false positive rates that are essential in the high-volume settings of operational sustainability. User and Entity Behavior Analytics (UEBA) systems are based on unsupervised learning algorithms to create behavioral thresholds of cloud users and services, allowing the identification of subtle deviations that may signal insider threats, or credentials that have been compromised (Ali et al., 2021). There is a gradual reduction of the barrier to AI-driven Zero Trust adoption by organizations lacking expert AI knowledge in the integration of AI with cloud-native security tooling, such as AWS GuardDuty, Azure Defender, and Google Security Command Center.

AI-assisted Zero Trust IoT Environment

IoT environments present security threats of qualitatively new nature that are not presented by cloud and traditional enterprise networks. The extreme scarcity of computational resources, battery life, and memory of the heterogeneous IoT devices sensors, actuators, industrial controllers, medical devices, and consumer electronics are coupled with the impossibility of deploying a traditional security agent (Atlam & Wills, 2019). A large number of IoT devices have firmware that is not updateable or patchable, and as such, poses a continuous vulnerability that cannot be addressed in a traditional way. These difficulties are magnified by the scale of IoT deployments, which are estimated to reach tens of billions of connected devices worldwide by 2025.

A variety of Ioot security functions have been suggested and tested using AI-driven solutions in a Zero Trust framework. To address privacy and security issues, lightweight machine learning models to run on resources-constrained edge computing devices have been developed to do local anomaly detection without sending raw data to centralized processing infrastructure (Hussain et al., 2020). Deep learning based network traffic analysis has been shown to be useful at identifying and profiling Ioot devices using behavioral signatures which can be used to automatically discover and classify large, heterogeneous Ioot deployments to enforce least privilege access policies distributed IoT ecosystem anomaly detectors that do not centralize sensitive device data, providing a privacy-siving channel to collective threat intelligence.

Difficulties and Lack of Literature

Although there is a substantial amount of technical research on AI and Zero Trust separately, the empirical research on the intersection of the two—namely, regarding cloud and IoT security—is still rather shallow and limited in scope. The majority of literature is theoretical models, putting forward architectural models without empirical support, or technical analyses of particular AI algorithms used on particular security tasks. There is a dearth of research that explores the organizational, human, and sociotechnical aspects of AI-driven Zero Trust implementation (Syed et al., 2022). Other

questions such as how security practitioners understand and engage with AI-driven security systems, how organizations should work around the migration of legacy architectures, and what governance structures are suitable in automated security decision-making, are some valuable gaps that this work aimed to fill.

METHODOLOGY

Research Design

The type of research design that was followed in this study was qualitative, based on an interpretive epistemological position, which was deemed as suitable due to the exploratory nature of the research aims and the necessity of building subtle insights into the experiences and views of practitioners. The choice of an exploratory qualitative approach was due to the lack of the empirical literature on the phenomenon of AI-based Zero Trust implementation in cloud and IoT environments, coupled with the impossibility of quantifying the complexity of the human experience and organizational setting in this field (Creswell and Poth, 2018). Qualitative inquiry offered the methodological latitude needed to reflect the richness, depth and contextual particularity of accounts by participants.

Participants and Sampling

Purposive sampling, which is a non-probability methodology, was employed to recruit 15 participants as the technique helped the researcher to choose information-rich cases that could provide in-depth information about the phenomenon being studied (Patton, 2015). The eligibility criteria were that the participants had to be actively involved in the cybersecurity practice and have direct experience in cloud security or IoT security or implementation of Zero Trust. The sample size of the participants was cloud security engineers (n = 5), IoT security specialists (n = 4), IT security managers (n = 4), and cross-domain security architects (n = 2). Participants were of organizations of different sizes and in multiple industries such as financial services, healthcare, manufacturing, telecommunications, and government. Professional experience was between four and eighteen years with a mean of around nine years. Professional networks, LinkedIn outreach and forums on cybersecurity community were used to conduct recruitment, and the process continued until theoretical saturation was achieved- the point where no interviews added substantively new themes (Lincoln and Guba, 1985).

Data Collection Instruments

Semi-structured interviews were used as the main source of collecting data, which were accompanied by document analysis. The interview guide was created going through an iterative process that was informed by the literature review and polished by pilot test of two cybersecurity practitioners who were not part of the main study sample. The guide included open-ended questions, which were grouped on four thematic areas, including the knowledge and experience of the participants about Zero Trust architecture, the role of AI and machine learning in their security practices, the particular challenges they experience in the area of cloud and IoT security, and how effective and restrictive they see AI-based security methods. Probing questions were also used to get elaboration and clarification of the responses of the participants.

Document analysis was a systematic review of 22 secondary sources, such as industry reports on Gartner, Forrester, and NIST; white papers by major cybersecurity vendors; organizational security policies by consenting participants; and any regulatory guidance documents available.

Document analysis helped to contextualize the results of the interviews with the wider industry trends and is used to triangulate the accounts of the participants with the documented organizational practices and policy frameworks (Bowen, 2009).

Data Collection Procedure

All the interviews were carried out online through secure videoconferencing systems. Before every interview, the participants were given an information sheet with the purpose of the study, procedures and their rights as participants. Electronic informed consent was received through writing. The interviews took about 25 minutes and between 20 and 35 minutes based on how thorough the participants were with their responses. All the interviews were recorded on audio tape with the express permission of the participants and transcribed verbatim by the researcher. Member-checking of transcripts occurred through sending them back to participants to have them preview and correct the transcripts, which contributed to the accuracy and credibility of the information (Creswell and Poth, 2018). A alpha numeric code was used to identify the participants to maintain anonymity in the analysis and reporting procedure.

Data Analysis

Thematic analysis was used to analyze the data collected based on the six-phase model outlined by Braun and Clarke (2006). The initial stage involved the researcher familiarizing themselves with the data by repeatedly reading transcripts and making initial notes. The second stage involved systematic open coding of the entire dataset, resulting in an initial set of 87 discrete coding that represented meaningful data units. The third stage involved sorting and grouping of codes into possible themes resulting in a first thematic map. The fourth stage was the review and refinement of the themes based on the iteration of the comparison with the data corpus which led to the consolidation and refinement of the thematic structure. During the fifth and sixth stage, themes were labeled and selected with precision and finally the final report was developed with the thematic results depicted by a few chosen quotes of the respondents. The NVivo qualitative data analysis software was used to conduct the analysis to make it easier to code and manage the themes.

Trustworthiness

A number of measures were taken to make sure that the study was trustworthy, as per the criteria of Lincoln and Guba (1985). The credibility was developed by long-term interaction with data, member checking and triangulation of the investigator with the help of peer review of the coding process by one of the other qualitative researchers. The transferability was facilitated by the fact that a thick description of the research context, the characteristics of the participants and the analytical process were provided so that the readers can determine the applicability of the findings to other contexts. Reliability was ensured by having a well-documented audit trail of all methodological choices. Confirmability was also considered through keeping a reflexivity journal where the researcher noted down and critically analyzed possible sources of bias during the study.

Ethical Considerations

This study was given ethical approval before the collection of data. All the participants were given informed consent and assured of the confidentiality of their responses and their right to withdrawwithout consequence from the study at any stage. All data were kept safely in encrypted institutional servers and only accessible to the research team. None of the outputs contained any identifying information, and all references to

organizations were anonymized. Participants who provided information about the sensitive aspects of the organization were especially told that such information would be reported in the aggregate and generalized form.

FINDING AND ANYLYSIS

Interpretation of the interview transcripts and supplementary material created four broad themes, which were further broken down into sub-themes. The following themes are outlined below with evidence provided on the basis of the interviews of the participants and other documentary materials.

Theme 1: The Legacy Security Models are not sufficient in the Cloud and IoT Environments.

The strongest and most common observation of all 15 interviews was a common belief that the traditional perimeter-based security architectures were not well adapted to the needs of cloud and IoT experiences. The participants recounted an increasing awareness in their respective organizations that traditional structures posed deadly blind spots in contexts of dynamic provisioning of resources, distributed endpoints, and ubiquitous between-device communication.

Particularly cloud security engineers described the structural incompatibility of perimeter models with multi-clouds. One participant (P3) explained the difficulty in the following way: When your workloads are spinning up and down in three different cloud providers and your users are connecting anywhere, the concept of a perimeter becomes virtually imaginary. You're defending a border that doesn't exist.' This was a common mood throughout the sample, and participants mentioned particular cases in their organizations, where the controls based on perimeter had not stopped the lateral movement after an initial breach.

IoT experts brought similar concerns about the special security needs of IoT ecosystems. The respondents noted that the size and diversity of IoT implementations, including extremely powerful industrial gateways and violently resource-starved sensors, presented conditions where standard endpoint security agents were no longer viable. P7, an IoT security expert in a large manufacturing company, said: 'We have thousands of sensors on the factory floor with 64 kilobytes of memory maybe. You would not be able to run an antivirus or a VPN client on that. Conventional security just does not work.' These practical limitations were validated during the analysis phase with documents of vendors admitting that traditional endpoint protection products could not support most of the deployed internet of things devices.

One of the most common sub-themes was the issue of implicit trust and how attackers could abuse it. The participants explained many cases where attacker lateral movement, insider attacks, and credential theft had been able to take advantage of the blanket trust given to users and devices once they were within the network perimeter. The challenge of implementing least privilege access in legacy environments where extensive permissions had been built up over years of operation was commonly mentioned by IT managers in the sample. P11 pointed out: 'We had old service accounts with domain administrator privileges that nobody could recall having created. That is what security people lose sleep over and it is rampant in organizations who have not embraced the thinking of Zero Trust. This result aligned with the documentary evidence of industry reports that credential abuse and privilege misuse are some of the most common attack vectors in cloud and IoT environments.

Theme 2: AI as an Operational Enabler of Zero Trust Principles

Participants also agreed unanimously that artificial intelligence (and machine learning specifically) was the key technology enabling the architectures of Zero Trust to become operationally feasible at the scale and speed required by cloud and IoT environments. There were three AI capabilities, which were identified to be most important: continuous behavioral authentication, real-time anomaly detection, and automated policy adaptation.

Cloud security engineers and security architects spotted continuous behavioral authentication as a game-changing feature that moved identity verification beyond the point-in-time validation of conventional multi-factor authentication. The respondents spoke of AI-driven user and entity behavior analytics (UEBA) systems where access patterns, session activities, and resource utilization were continuously compared to individually customized behavioral standards to initiate step-up authentication or terminate a session when behaviors deviated beyond pre-determined risk levels. P1, a senior cloud security engineer, explained the effect of implementing such a system: 'The difference between the state of affairs before UEBA was the use of static MFA upon login and hope that nothing would be altered during the session. The system now is continually questioning whether this action makes sense to this user, and it detects the things we would not have known about in weeks, accounts being accessed at odd times of the day and at odd places, slow data exfiltration, etc. This record matched closely with documentary findings of UEBA vendor white papers and industry benchmark reports, which reported critical advancements in insider threat detection rates after the implementation of AI-based behavioral analytics.

The second pillar of AI-enabled Zero Trust was real-time anomaly detection, with participants noting how machine learning models used on network traffic, log data, and device telemetry allowed revealing threats that were regularly overlooked by signature-based systems. IoT experts highlighted the specific utility of anomaly detection to the IoT security, where conventional endpoint controls were not available, and network-level behavioral analysis was the main detection tool. P8 clarified: "Security software is not executed by our IoT devices, thus the network is our plane of visibility. Our ML models are trained on examples of what normal behavior should be per device type, e.g., a temperature sensor does not need to query external servers, and anything not normal should be marked. It has grabbed command-and-control traffic which would otherwise have been undetected by our old IDS.' Several IoT experts highlighted the importance of unsupervised learning to determine device behavioral profiles where no labeled training data are available, as it is practically impossible to provide labeled attack data on each device type in a heterogeneous deployment. Security architects have outlined AI-driven policy engines that consumed contextual indicators such as device health scores, threat intelligence feeds, user behavioral risk scores, and environmental indicators to make real-time access control decisions that would otherwise be impractical to execute via rule sets manually maintained. This capability was described by P14, a security architect who has vast experience in deploying zero trust.The strength of AI in Zero Trust is the ability to store thousands of variables at a time and make a decision based on the situation in milliseconds. That can not be done by a human analyst or a

fixed set of rules. The AI understands that this user is located in London, their device is fully patched, they have never been to this system and their peer group is, and the danger level is high at the moment- and it will decide to make an access control accordingly calibrated.' A documentary review of organizational security policy documents submitted by the participants revealed that organizations with the well-developed AI-based Zero Trust implementation had shifted to the dynamic, risk-based authorization models, as opposed to the predominantly role-based access control models.

Theme 3: Implementation Problems and Organization Obstacles

Although the importance of AI-based Zero Trust architecture is widely acknowledged, respondents consistently reported significant impediments to implementation that indicated the discontinuities between the conceptual potential of such frameworks and the reality of their implementation in organizations. Data identified four main types of implementation challenge: resource and expertise constraints, complication in integration with legacy systems, data quality and governance problems, and cultural and organizational resistance.

Most participants reported resource and expertise limitations as the biggest implementation barrier, especially in organizations not part of the largest organizations. Implementation of AI-based Zero Trust systems: Implementing AI-based systems necessitates specialized expertise across various areas such as cybersecurity architecture, machine learning engineering, cloud architecture, and identity management that many organizations lack in a cohesive form. An IT security manager at a medium-sized healthcare institution, P6, explained a typical dilemma: 'We knew the value proposition, but we just lacked the people. You need somebody who has heard about Zero Trust architecture, somebody who has heard about how to train and test ML models, somebody who has heard about cloud IAM, ideally you would just get those skills in the same group, but in practice we were robbing time off other people with four other priorities. Another identified barrier is the cost of commercial AI security platforms, where many respondents have reported that pricing models of enterprise-grade solutions put comprehensive Zero Trust implementations outside the budget of smaller organizations.

The complexity of integration with legacy systems became a closely related issue, and participants reported an enormous technical effort needed to apply the principles of Zero Trust to the environments with large amounts of legacy applications, on-premises infrastructure, and operational technology systems that were not designed around the modern security APIs. The IoT experts in manufacturing and utility industries encountered special issues in implementing Zero Trust to operational technology (OT) systems where equipment can potentially last decades and cannot be updated or replaced due to security concerns. P9 noted: 'Our oldest PLCs were in operation in the 1990s. They even speak protocols that Zero Trust does not even speak, and you cannot patch them or add authentication to them. You find yourself developing compensating controls, such as network isolation, monitoring gateways, which are approximations of Zero Trust rather than it. Industry reports obtained during the document analysis supported this result, and OT/IT convergence was identified as one of the most problematic issues of enterprise Zero Trust implementation.

Participants noted data quality and governance issues, asserting that AI-based security systems rely

heavily on quality, well-labeled training data, and extensive telemetry coverage. Blind spots formed as a result of gaps in visibility due to unmonitored network segments, uncontrolled devices, and inadequate logging settings were exploited by adversaries and at the same time compromised the performance of AI models on unfinished data. P4 commented: 'Our anomaly detecting is as good as the data it is viewing. When a device is not logging it is not seen. When the logs are not uniformly structured between vendors, the model fails. The quality of data is a real bottleneck that I believe lacks due consideration in the academic literature. The issue of AI-generated security alert governance, e.g., the responsibility of the decision made by the automated systems and the audit of the AI-controlled access denials, were brought up by the members of the regulated sectors, which shows that the questions regarding the governance of AI and cybersecurity compliance are open.

Participants characterized cultural and organizational resistance to the implementation of Zero Trust as pervasive but less obvious. The shift to Zero Trust often requires modifying the existing workflows, tolerating increased authentication procedures, and a readiness to adopt a security paradigm that initially seems to impose greater convenience and control on the user. The sample of IT managers reported the pushback by organizational business units that saw advanced security controls as burdens to productivity, and IT peers who were accustomed to the traditional network administration paradigm and were more skeptical about AI-driven automation of security-sensitive operations. P12 reflected: 'Zero Trust is as much an organizational change management project as it is a technology implementation. The most secure AI security system in the world can be deployed and people will find it easier to bypass it, but unless they know why it is demanding they re-authenticate every two hours, they will, and those exceptions become vulnerabilities.

Theme 4 Future Directions and Evolving Threat Landscape

The participants were encouraged to consider the direction of the development of AI-based Zero Trust and how it is expected to change following the evolving threat environment. Some future-oriented sub-themes surfaced, which included autonomous and self-healing security systems, the increasing danger of AI-driven adversaries, and the necessity of standardization and regulatory clarity.

The notion of self-healing, autonomous security ecosystems, where AI systems are not only detecting and responding to threat events, but also actively patching vulnerabilities and adjusting security settings without human intervention, was outlined by various respondents as the natural extension of present trends of development. Security architects and cloud security engineers showed interest in the new capabilities in AI-driven security orchestration, automation, and response (SOAR) platforms, which increasingly have the ability to perform more than just alert triage, to also contain, eradicate, and recover. P2: The vision is a system which identifies an anomaly, automatically isolates it, performs a root cause analysis, fixes the vulnerability, revises the policy to prevent future occurrence and prepares the incident report, before a human analyst has even laid his eyes on the alert. We are not quite there yet, but the trend is evident.'The risk of AI-powered enemies was revealed among the participants as one of the most significant issues that would require the constant development of defensive AI capacities. Multiple respondents noted they were increasingly worried about adversarial machine learning methods, such as

model evasion attacks, data poisoning, and adversarial examples, that can be used to compromise the integrity of AI-based security systems. P5 was moderately worried: 'Attackers have the same methods that enable us to detect them better. Polymorphic malware has already appeared, and appears to have been crafted to bypass ML classifiers. The arms race is factual and it is taking it to a different level. These practitioner anxieties were consistent with the new scholarly literature on adversarial AI in cybersecurity that has already started to catalogue the organization of machine learning vulnerability by advanced threat actors (Papernot et al., 2018).

Participants in various industries found the necessity to standardize AI-driven Zero Trust systems and provide more regulating directions as a valuable precondition to further implementation. Respondents were aware of the existing fracturing of standards, vendor architecture and regulatory requirements, which posed a risk to organisations wanting to undertake long-term architectural investments. P13 proposed: 'An agreed-upon framework that describes what good looks like in terms of AI-driven Zero Trust, and with specific metrics and assurance criteria. Each of the vendors has their own maturity model at the moment and they all appear different. Organizations require something that is authoritative to direct their investments as well as their boards.

DISCUSSION

The results of this paper shed light on the multifaceted and multi-dimensional picture of AI-based Zero Trust in cloud and IoT deployment, featuring both a lot of promise and a lot of challenges that mirror the reality of this fast-growing sector. The themes developed in the thematic analysis lead to the convergence of the themes to form a consistent narrative: the structural inefficiency of legacy security frameworks has formed an acute need to develop frameworks that are more adaptive; AI offers the operational functionality allowing Zero Trust to be implemented at scale; but the transition between conceptual adoption and operational maturity is hindered by resource limitations, technical complexity and organizational inertia.

The fact that participants in the legacy security architecture are universally recognized to be ineffective in cloud and IoT environments supports and builds upon existing theoretical criticisms in the literature (Rose et al., 2020; Stiennon, 2014). The practitioner accounts gathered during this study give finer empirical support to arguments that have largely been made on the theoretical plane, detailing particular failure modes, types of incidences, and architectural weaknesses that have been characterized in the literature in general terms. The observation that implicit trust exploitation (via credential theft, service account abuse, and lateral movement) is one of the major attack vectors in organizations with perimeter-dependent architectures is similar to the industry incident data and supports the urgency of embracing Zero Trust.

The fact that AI is the key operational enabler of Zero Trust at cloud and IoT scale not only expands but also adds value to the current technical literature. Although many studies have shown the effectiveness of particular AI algorithms in performing security operations in isolation, the practitioner narratives in this paper clarify the combined, ecosystem-wide value that AI offers to a Zero Trust framework. The three most visible AI capabilities identified as continuous behavioral authentication, real-time anomaly detection, and automated policy adaptation, reflect the essence of operational requirements of Zero Trust as

established by NIST (Rose et al., 2020) and address collectively the scalability limitations that have historically hindered the implementation of Zero Trust. These results complement and build upon the research by Sarker et al. (2021), who hypothesized that AI can help overcome the scalability issues of Zero Trust, by presenting empirical evidence of the practices in the field in operationally deployed systems.

The implementation issues outlined in this research, including resource and expertise limitations, complexity of legacy integration, data quality problems, and organizational resistance, are a valuable addition to the body of literature that has put greater emphasis on technical capabilities and less emphasis on organizational and sociotechnical obstacles. The conclusion that one of the main obstacles is expertise, especially in the case of mid-sized organizations, bears significant policy implications to workforce development and design of commercial AI security platforms. The fact that the OT/IT convergence when working in legacy industrial settings is unique due to several factors, as observed by iont experts, complements the literature on the subject by outlining the situations where the principles of Zero Trust are to be applied in practice, as opposed to being applied directly.

The identified data quality and governance issues that the participants have highlighted pose valuable questions that have not been adequately addressed both within the academic literature and in the practitioner literature. This reliance of AI-based security systems on high-quality and full-coverage telemetry data amounts to a fundamental vulnerability: any gaps in coverage not only introduce attacker blind spots but also impair AI models. This observation implies that companies that aim to implement AI-based Zero Trust should consider data infrastructure, such as logging coverage, data normalization, and data telemetry retention, as a strategic security asset instead of an operational afterthought. The control aspect of AI-driven security policy, especially in the regulated sectors where accountability and auditability are established by law, is an area of gaps in existing regulatory systems that need immediate consideration by policymakers.

The futuristic insights provided by participants form a path towards more independent security operations, which, although promising in response speed and scalability, beg serious questions about human control, responsibility, and control of AI in high-stakes decision-making situations. Of particular interest is the practitioner concern of AI-powered adversaries, which implies that the cybersecurity arms race is entering a qualitatively new stage, where the integrity of defensive AI systems, and not just their accuracy, becomes the main security concern. The implications of this finding on the research of AI security directly point to the necessity of adversarially robust machine learning methods and formal verification methods of security-critical AI systems.

CONCLUSION

This paper examined how AI-based Zero Trust security model could be applied to secure the cloud and IoT infrastructure by qualitatively exploring practitioner views and experiences. The results, which were obtained through 15 semi-structured interviews with cybersecurity experts and complemented by document analysis, were that AI-based Zero Trust systems are qualitatively better at protecting cloud and IoT systems than the old perimeter-based systems. The AI capabilities, such as continuous behavioral authentication, real-time anomaly detection, and automated policy optimization, were confirmed as crucial operational

enablers that render Zero Trust viable at the level of scale and speed required in modern cloud and IoT environments.

At the same time, the research has reported significant implementation issues that need to be resolved to enable these frameworks to reach their full potential in organizations of all scales and sectors. Legacy complexity of integration, resource and expertise constraints, and data quality governance, and organizational resistance are an uphill yet surmountable set of challenges, fulfilling which will necessitate a coordinated effort by technology vendors, standards bodies, regulatory bodies, and others

and organizational leaders.

Recommendations for Practice

Some recommendations are provided based on findings to cybersecurity practitioners and organizational leaders. Organizations ought to take a gradual implementation strategy of Zero Trust by focusing on high-value, high-risk assets and constructing the data infrastructure and organizational capacity to implement more widely. The infrastructure to invest in AI security skills and in commercial offerings that make AI capabilities readily available to organizations without full ML engineering teams is necessary. The assessment of systematic telemetry coverage must be carried out before AI security deployment to determine and address data gaps that would jeopardize model performance. Overcoming organizational resistance is essential to change management programs that explain the reasons and advantages of improved security controls to end users.

Policy and Research Recommendation

Policymakers need to come up with standard assurance frameworks of AI-driven Zero Trust deployments that give organizations clear benchmarks and guidance, and that respond to the accountability and auditability demands of AI-driven security decisions in controlled settings. There is a pressing need to provide regulatory guidance on how automated security decisions can be governed (including human oversight, explainability, and audit trail requirements).

The qualitative basis of the current research must be enlarged in future studies by longitudinal and mixed-method research that will monitor the paths of Zero Trust implementation over time and measure the performance outcomes. Research that specifically analyzes AI-led Zero Trust in critical infrastructure industries, such as healthcare, energy, and transportation, would be beneficial to offer sector-specific advice. The increasing risk of aggression to AI security systems is a research agenda priority that requires both technical and governance-oriented research.

Limitations

This research has a number of limitations that ought to be taken into account when analyzing the results. The 15 members of the purposive sample used in the qualitative investigation and adequate to reach theoretical saturation might not represent the entire range of practitioner experience in all organizational settings and industries. Professional networks were used to self-select the participants, which may have brought about a participation bias to those with a more progressive or positive attitude towards AI-driven

security. The cross-sectional nature of the study gives a picture of how practitioners perceive something at a given time in a fast changing profession and the results might need revision as more technologies and practices are developed.

REFERENCES

Ali, B., Hijjawi, M., & Almomani, A. (2021). Machine learning techniques for cloud security: A systematic review. Journal of

Information Security and Applications, 58, 102769. https://doi.org/10.1016/j.jisa.2021.102769

Atlam, H. F., & Wills, G. B. (2019). IoT security, privacy, safety and ethics. In Digital twin technologies and smart cities (pp. 123-149). Springer. https://doi.org/10.1007/978-3-030-18732-3_8

Bowen, G. A. (2009). Document analysis as a qualitative research method. Qualitative Research Journal, 9(2), 27-40. https://doi.org/10.3316/QRJ0902027

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, 3(2), 77-101. https://doi.org/10.1191/1478088706qp063oa

Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. IEEE Communications Surveys & Tutorials, 18(2), 1153-1176. https://doi.org/10.1109/COMST.2015.2494502

Chen, X., & Zheng, Z. (2022). AI-driven anomaly detection in cloud-native environments: A machine learning approach to zero trust implementation. IEEE Transactions on Cloud Computing, 10(3), 1824-1838. https://doi.org/10.1109/TCC.2022.3164752

Cheswick, W. R., Bellovin, S. M., & Rubin, A. D. (2003). Firewalls and internet security: Repelling the wily hacker (2nd ed.). Addison-Wesley.

Creswell, J. W., & Poth, C. N. (2018). Qualitative inquiry and research design: Choosing among five approaches (4th ed.). SAGE Publications.

Hashizume, K., Rosado, D. G., Fernandez-Medina, E., & Fernandez, E. B. (2013). An analysis of security issues for cloud computing. Journal of Internet Services and Applications, 4(5), 1-13. https://doi.org/10.1186/1869-0238-4-5

Hussain, F., Hassan, S. A., Hussain, R., & Hossain, E. (2020). Machine learning for resource management in cellular and IoT networks: Potentials, current solutions, and open challenges. IEEE Communications Surveys & Tutorials, 22(2), 1251-1275. https://doi.org/10.1109/COMST.2020.2964534

Kindervag, J. (2010). No more chewy centers: Introducing the zero trust model of information security. Forrester Research.

Lincoln, Y. S., & Guba, E. G. (1985). Naturalistic inquiry. SAGE Publications.

Mehraj, S., & Banday, M. T. (2020). Establishing a zero trust strategy in cloud computing environment. In 2020 International Conference on Computer Communication and Informatics (ICCCI) (pp. 1-6). IEEE. https://doi.org/10.1109/ICCCI48352.2020.9104214

Papernot, N., McDaniel, P., Goodfellow, I., Jha, S., Celik, Z. B., & Swami, A. (2018). Practical black-box attacks against machine learning. ACM Transactions on Privacy and Security, 21(3), 1-36. https://doi.org/10.1145/3196494.3196497

Patton, M. Q. (2015). Qualitative research and evaluation methods (4th ed.). SAGE Publications.

Rose, S., 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

Sanger, D. E., & Perlroth, N. (2021). Pipeline attack yields urgent lessons about U.S. cybersecurity. The New York Times.

Sarker, I. H., Furhad, M. H., & Nowrozy, R. (2021). AI-driven cybersecurity: An overview, security intelligence modeling and research directions. SN Computer Science, 2(3), 173. https://doi.org/10.1007/s42979-021-00557-0

Stafford, V. A. (2020). Zero trust architecture. NIST Special Publication, 800(207), 1-59.

Stiennon, R. (2014). There will be cyberwar: How the move to network-centric war fighting set the stage for cyberwar. IT-Harvest Press.

Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero trust architecture (ZTA): A comprehensive survey. IEEE Access, 10, 57143-57179. https://doi.org/10.1109/ACCESS.2022.3174679