pISSN:2349-5707 Volume:12, Issue 10, Octobar-2025 eISSN:2349-5715

Advancing Zero-Trust Architectures in Multi-Tenant Cloud Environments: Integrating Human Factors, Protocol Frameworks, and Endpoint Augmentation

John A. Davenport
Global Institute for Cybersecurity Studies, University of Edinburgh, United Kingdom

ABSTRACT

Background: The rapid adoption of cloud computing and multi-tenant platforms has reoriented traditional perimeter-based security approaches toward architectural paradigms that assume compromise and distrust implicit trust relationships. Zero-trust architecture (ZTA) offers a principled response by asserting continuous verification, least privilege, and microsegmentation. However, theoretical design and practical deployments must negotiate a complex interplay of technical protocols, human factors such as security fatigue and usability, and emergent augmentations including blockchain-assisted endpoints and software defined perimeters. This article synthesizes foundational standards, empirical studies on authentication usability, and recent scholarly advances to propose a holistic view of ZTA in multi-tenant clouds. (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014; Stanton et al., 2016).

Objectives: To produce an integrative, publication-ready treatment that (1) explicates the rationale for zero-trust adoption in multi-tenant clouds, (2) maps architectural components and protocol choices to threat models and tenant isolation requirements, (3) examines human factors that influence adoption and effectiveness, (4) surveys endpoint augmentation approaches and blockchain integrations, and (5) propose a research and deployment agenda addressing measurement, governance, and continuous adaptation. Each major claim is grounded in the provided literature.

Methods: This article undertakes an analytical synthesis of the provided references, drawing inference via cross-comparative analysis, threat modeling thought experiments, and rigorous conceptual elaboration. Rather than reporting primary empirical measurements, the methodology organizes extant evidence into a theoretically consistent framework suitable for informing future empirical studies and operational translations (Weidman & Grossklags, 2017; Strouble et al., 2009).

Results: Integrated analysis reveals (a) zero-trust principles align tightly with multi-tenant cloud requirements but require careful translation into identity, device, and network controls; (b) usability and security fatigue are high-impact constraints that degrade real-world assurance unless mitigated through design choices in authentication modalities and organizational policy; (c) protocol stacks such as XMPP for security information exchange and Software Defined Perimeter specifications provide operational primitives but must be reconciled with dynamic access control and telemetry; and (d) endpoint augmentation via blockchain and advanced authorization models shows promise for distributed trust but introduces new complexity and governance questions (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014; Alevizos et al., 2022).

Conclusions: Zero-trust architecture is a necessary and promising reorientation for securing multi-tenant cloud environments, but its success depends on integrating human-centred design, automated continuous validation, and transparent governance. The future research agenda must include rigorous field experiments, standardized metrics for assurance and usability, and exploration of hybrid technical patterns that combine centralized control with provable decentralized attestations. (Khan, 2023; Hariharan, 2025; He et al., 2022).

Keywords: Zero trust, multi-tenant cloud, security fatigue, software defined perimeter, identity and access management, blockchain endpoints

INTRODUCTION

The velocity of digital transformation has pushed organizations of every scale to adopt cloud computing and multi-tenant platforms as foundational infrastructure. Multi-tenant clouds support economies of scale, elasticity, and rapid deployment but simultaneously introduce complex security challenges that arise from shared resources, heterogeneous tenants, and dynamic service boundaries. Traditional perimeter-centric security models — which rely on a hardened network boundary and implicit trust for internal entities — falter in environments where assets, identities, and workloads traverse untrusted networks and where tenants coexist on shared physical and virtual infrastructure. Zero-trust architecture (ZTA) emerges in this context as an alternative security philosophy that rejects inherent trust, instead relying on continuous verification of identity, device posture, and authorization decisions (Khan, 2023; He et al., 2022).

The conceptual roots of zero trust are clear: trust must be earned and continuously validated, not assumed by network location or by virtue of being "inside" a perimeter (Tyler & Viana, 2021). In multi-tenant clouds the imperative for zero trust is heightened: misconfigurations, lateral movement, and tenant compromise can propagate across shared substrates absent strict isolation and least-privilege controls. However, translating zero-trust principles into deployable mechanisms demands confronting several interlocking issues. First, architectural choices must reconcile identity systems, device attestations, and fine-grained authorization to support dynamic access patterns typical in cloud workloads (Yao et al., 2020). Second, human factors — notably security fatigue, usability trade-offs associated with multi-factor authentication, and organizational behavioral responses to new security regimes — significantly condition whether technical controls achieve their intended protective outcomes (Stanton et al., 2016; Weidman & Grossklags, 2017; Strouble et al., 2009). Third, novel augmentations such as blockchain-assisted endpoint attestations or software-defined perimeters (SDPs) promise improvements in trust distribution but introduce complexity and governance questions that must be resolved before they can be widely adopted (Alevizos et al., 2022; Software Defined Perimeter Working Group, 2014).

The literature supplied for synthesis spans protocol standards and RFCs that cover messaging for security information exchange (Cam-Winget et al., 2019), practical SDP specifications, studies of authentication usability and security fatigue, and a set of focused zero-trust research outputs addressing healthcare, IoT, and cloud environments (Chen et al., 2020; Tyler & Viana, 2021; Mehraj & Banday, 2020). These documents collectively provide a comprehensive corpus to ground a rigorous exploration of how zero-trust approaches can be realized, measured, and institutionalized in multi-tenant cloud contexts. This article thus seeks to fill an important literature gap by weaving standards, human factors research, and emerging technical augmentations into a detailed, actionable conceptual framework for zero-trust in multi-tenant clouds. It explicitly articulates the tensions, trade-offs, and research questions that must be addressed for practical, scalable, and usable implementations.

METHODOLOGY

This article employs a structured analytical synthesis methodology designed to transform the supplied references into a coherent, publication-ready argument and framework. The methodologies are entirely text-based and theoretical, emphasizing conceptual rigour rather than experimental data collection. The approach consists of five overlapping steps: corpus familiarization, theme extraction, architectural mapping, threat-driven design reasoning, and generation of a research and deployment agenda.

Corpus Familiarization: Each provided reference was examined for its primary claims, methodological approach, and implications for zero-trust deployment in cloud environments. Standards and RFCs were analyzed for operational primitives (message patterns, telemetry requirements), empirical studies for human

factors evidence (usability, fatigue, authentication acceptance), and zero-trust research articles for design patterns and formalized frameworks. Textual cues (for example, those highlighting recommended controls, constraints, or potential failure modes) were cross-referenced to establish a set of recurring themes (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014; Stanton et al., 2016).

Theme Extraction: From the corpus, recurring themes were extracted, categorized, and prioritized. Themes included: continuous verification and identity centrality; device attestation and endpoint augmentation; usability trade-offs in authentication; telemetry and security information sharing; microsegmentation and dynamic authorization; blockchain as a provenance and attestation layer; and governance and auditable policy enforcement. Each theme was linked to primary source claims to ensure that subsequent elaboration remained traceable to the provided literature (Khan, 2023; He et al., 2022; Alevizos et al., 2022).

Architectural Mapping: The themes were mapped onto a reference architecture for multi-tenant clouds. This step entailed describing component responsibilities (identity and access management, policy decision points, policy enforcement points, telemetry collectors, and endpoint attestation subsystems), and how these components interact in high-level flows. Protocol choices were considered where standards like XMPP for security information exchange or the SDP specification offered operational guidance (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014).

Threat-driven Design Reasoning: The architecture and associated controls were evaluated against representative threat scenarios typical of multi-tenant clouds: tenant compromise enabling lateral movement, cross-tenant data leakage, insider threat exacerbated by BYOD policies, and sophisticated supply-chain attacks. Design reasoning emphasized how zero-trust controls mitigate each threat and where residual risk remains, while ensuring that each mitigation claim is tied back to literature that supports the mechanism or demonstrates constraints (Kim et al., 2020; Weidman & Grossklags, 2017).

Research and Deployment Agenda Generation: Finally, the analysis distilled gaps and proposed a prioritized agenda encompassing measurement approaches (usability metrics, assurance metrics), governance structures (policy lifecycle, tenant consent), and recommended future studies (field experiments to quantify security–usability trade-offs, longitudinal studies of fatigue, and pilot deployments of blockchain endpoint attestations). This agenda is motivated directly by literature-identified gaps, for example the complexities of combining blockchain with endpoint attestation and the demonstration that security fatigue reduces compliance (Alevizos et al., 2022; Stanton et al., 2016; Strouble et al., 2009).

Throughout the methodology, claims made about effectiveness, trade-offs, or operational constraints were accompanied by explicit citations from the provided corpus to ensure rigorous traceability.

RESULTS

This section presents descriptive findings that emerge from synthesizing the provided literature into an integrated account of zero-trust architectures for multi-tenant cloud environments. The results are organized into thematic clusters: architectural alignment and primitives; identity, authentication, and human factors; endpoint augmentations and blockchain; telemetry and security information exchange; and governance, policy, and operationalization.

Architectural Alignment and Primitives

Zero-trust principles — continuous verification, least privilege, microsegmentation, and assuming breach — map directly onto the security needs of multi-tenant clouds. Multi-tenant clouds require fine-grained access control pathways that isolate tenant resources while permitting controlled cross-tenant services such as

federation or shared platform services. The architectural primitives necessary to implement zero trust include robust identity and access management (IAM), policy decision points (PDP) and policy enforcement points (PEP) distributed near resources, device and workload attestation mechanisms, and a telemetry fabric capable of continuous posture assessment (Khan, 2023; He et al., 2022; Yao et al., 2020).

Identity becomes the central control plane: when location and network perimeter cannot reliably denote trust, identity assertions — both for human users and workloads — become the primary signals used for authorization. This necessitates identity systems that support multi-factor authentication (MFA), federated identity across tenants and providers, and dynamic attribute-based access control (ABAC) models capable of reflecting current context (Mehraj & Banday, 2020; Yao et al., 2020). The PDP must synthesize contextual signals — identity attributes, device posture, network telemetry, behavioral indicators — and compute authorization decisions in real time, while PEPs enforce decisions at the resource boundary. Microsegmentation is operationalized by continuously managed policies that restrict lateral movement and apply least privilege, reducing the blast radius of any compromise (Khan, 2023).

Identity, Authentication, and Human Factors

The literature demonstrates that while MFA and stronger authentication schemes materially increase assurance, they are not panaceas; usability, fatigue, and employee perceptions determine long-term effectiveness. Strouble et al. (2009) and Weidman & Grossklags (2017) provide empirical evidence that adding two-factor security systems impacts productivity and acceptance. Strouble et al. (2009) observe measurable productivity and usability effects when two-factor systems are introduced, while Weidman & Grossklags (2017) report mixed employee perceptions during institutional transitions to BYOD second-factor authentication, indicating that technical strength must be balanced against workflow integration and user burden.

Security fatigue — the cognitive and motivational depletion that results from persistent security demands — is a recurrent barrier to secure behavior (Stanton et al., 2016). When users are asked to manage multiple authentication tokens, respond to frequent prompts, or navigate complex recovery processes, adherence drops and workaround behaviors proliferate (Stanton et al., 2016). Therefore, deployment of zero-trust controls must be guided by human-centred design: adaptive authentication mechanisms that escalate only when risk indicators justify friction, single sign-on (SSO) systems that reduce repeated credential entry, and careful attention to provisioning and recovery flows that minimize user disruption (Strouble et al., 2009; Weidman & Grossklags, 2017).

The trade-off between security and usability yields complex governance questions: mandating strict MFA with punitive noncompliance can increase security but may also incentivize shadow IT and noncompliant device usage. Conversely, lax enforcement preserves usability but increases attack surface. The literature suggests mitigation strategies such as context-aware step-up authentication (granting low-risk access with minimal friction while requiring stronger factors for sensitive actions), organizational training tailored to reduce fatigue by emphasizing clear, task-relevant prompts, and automation to minimize repetitive security decisions (Stanton et al., 2016; Yao et al., 2020).

Endpoint Augmentations and Blockchain

A growing body of work explores augmenting ZTA with blockchain primitives for endpoint attestation and provenance. Alevizos et al. (2022) review approaches that leverage blockchain to secure endpoint attestations and maintain immutable logs of device posture and firmware state. Blockchain can provide tamper-evident provenance records asserting prior attestations, software versions, and integrity measurements; these

immutable attestations can then be referenced by PDPs to support authorization decisions without requiring centralized trust anchors.

However, integrating blockchain into zero-trust endpoint verification presents trade-offs. First, blockchain systems introduce latency and scalability considerations that can conflict with real-time authorization needs. Second, the governance of a blockchain overlay — who operates nodes, how consensus is achieved, and how privacy of tenant attestations is preserved — creates complex organizational and legal questions (Alevizos et al., 2022). Third, immutable records that capture endpoint state may inadvertently reveal sensitive configuration or compliance data, raising confidentiality concerns. Therefore, blockchain augmentations appear most promising in limited, high-assurance scenarios where provable audit trails are essential and the consortium governance model can be clearly specified (Alevizos et al., 2022).

Beyond blockchain, endpoint augmentations include hardware attestation (for example, TPM and secure enclave attestations), continuous telemetry agents reporting posture, and lightweight measurement protocols to provide rapid evidence of device integrity. The combination of hardware-rooted attestation plus transient telemetry provides stronger signals than either alone — hardware roots offer stronger initial trust anchors, while telemetry captures runtime deviations indicative of compromise (Kim et al., 2020; Yao et al., 2020).

Telemetry and Security Information Exchange

Operational zero-trust architectures depend fundamentally on continuous telemetry: identity logs, device posture reports, network flow data, application audit trails, and behavioral analytics. Sharing and correlating these signals across control planes enable PDPs to compute contextually accurate decisions, and enable security teams to detect anomalies and escalate responses. Standards and protocols that facilitate security information exchange can therefore be central to scalable zero-trust deployments.

The RFC on using Extensible Messaging and Presence Protocol (XMPP) for security information exchange presents an operationally viable message bus for distributing security-relevant messages with low latency and flexible routing semantics (Cam-Winget et al., 2019). XMPP's extensibility and publish/subscribe patterns map well to use cases such as distributing device posture updates, broadcasting revocation notices, and enabling cross-service telemetry. The Software Defined Perimeter (SDP) specification articulates a model that hides services from unauthenticated discovery and establishes ephemeral, authenticated connections only after identity and posture are verified; SDP thus operationalizes zero-trust principles at the network connection level (Software Defined Perimeter Working Group, 2014).

A critical insight is that telemetry pipelines must support both real-time enforcement (fast decision loops for session establishment and step-up authentication) and forensic/audit capabilities (for after-the-fact investigations). The balancing act is nontrivial: high-throughput telemetry can be expensive and privacy-sensitive, while sparse telemetry reduces detection fidelity. Moreover, telemetry quality can be degraded by encryption, network segmentation, or privacy preservation choices. Therefore, architectures must specify tiered telemetry policies that align with tenant expectations and regulatory requirements; for instance, coarse-grained signals for operational decisions and more granular telemetry retained under stricter governance for audit purposes (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014).

Governance, Policy, and Operationalization

Translating zero-trust architecture into sustained practice requires governance frameworks that cover policy lifecycle management, tenant consent models, audit and compliance, and alignment with organizational risk appetites. ZTA shifts control emphasis from network boundaries to identity and policy; consequently,

governance must manage policy definitions (what actions are allowed under which contexts), policy distribution (how PDPs and PEPs receive and reconcile policies), and conflict resolution (how overlapping tenant policies or provider policies interact).

Examples from the literature highlight sectoral considerations. Tyler and Viana (2021) emphasize a framework tailored to healthcare organizations that accounts for safety-critical access patterns and regulatory constraints; healthcare requires balancing immediate clinical access needs with privacy and auditability. Chen et al. (2020) describe an approach for 5G smart healthcare where zero-trust principles are applied to IoT device interactions, necessitating specialized controls for constrained devices and latency-sensitive clinical services. These domain examples illustrate that governance must be contextually adapted: a one-size-fits-all policy model is impractical across tenants and sectors.

Operationalization also raises questions about responsibility and function location. Should PDPs be centralized within the cloud provider control plane, or should tenants operate their own PDPs? Centralization enables consistent policy enforcement and economies of scale, but raises trust concerns for tenants who may not wish to cede policy control. Distributed or federated PDPs empower tenant autonomy but complicate cross-tenant interoperability and increase the risk of inconsistent enforcement. Hybrid models — where providers offer policy-as-a-service with tenant-scoped policy domains, accompanied by verifiable attestation logs — may offer pragmatic trade-offs (Khan, 2023; Hariharan, 2025).

Finally, governance must address change management and policy evolution. Zero trust is fundamentally dynamic: device state, threat indicators, and tenant requirements evolve. Governance must therefore institute processes for continuous policy testing, safe deployment of policy updates, rollback mechanisms, and observable impact analysis. Absent disciplined governance, dynamically changing policies risk unintended availability impacts or privilege creep.

DISCUSSION

The integration of zero-trust architecture into multi-tenant cloud environments is both necessary and challenging. The evidence synthesized from the provided corpus supports several key interpretive conclusions, identifies limitations of current approaches, and proposes prioritized directions for future research and operational practice.

Interpretation of Findings

First, zero trust addresses the fundamental structural weaknesses of perimeter-based security in cloud settings by reorienting enforcement around identity and continuous verification (Khan, 2023; He et al., 2022). This conceptual shift reduces reliance on brittle network boundaries and focuses defensive investments on controlling who and what can access resources at the time of access. The literature suggests that when identity is accurately modeled and dynamically evaluated against contextual signals such as device posture and behavioral analytics, authorization decisions become both more secure and more granular (Yao et al., 2020; Mehraj & Banday, 2020).

Second, human factors emerge as a decisive operational constraint. The strongest technical controls are only as effective as the people who operate and interact with them. Evidence that two-factor systems affect productivity and that transition to BYOD second-factor authentication provokes mixed perceptions demonstrates that usability directly mediates adoption and adherence (Strouble et al., 2009; Weidman & Grossklags, 2017). Security fatigue can erode compliance and encourage risky workarounds if systems are not designed to minimize unnecessary friction (Stanton et al., 2016). Therefore, deployment strategies that rely

solely on enforcement without addressing cognitive load and workflow integration are likely to underperform.

Third, the promise of decentralized attestations and blockchain augmentation is significant but conditional. Provable, immutable attestations can enhance auditability and cross-tenant trust, particularly in consortium or regulatory contexts where shared provenance matters. However, latency, scalability, privacy, and governance introduce substantial implementation burdens that must be reconciled. Blockchain is not a universal remedy; rather, it is a complementary tool that offers unique properties when applied judiciously to specific assurance needs (Alevizos et al., 2022).

Fourth, operational primitives such as XMPP for security information exchange and Software Defined Perimeter specifications provide technical building blocks for implementing zero-trust patterns. XMPP's message addressing and extensibility support flexible telemetry distribution, and the SDP specification enforces the "deny until proven safe" posture that ZTA requires (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014). Yet, connecting these primitives into a cohesive, scalable control plane that meets tenant privacy and performance expectations demands careful engineering and policy work.

LIMITATIONS

The analytical synthesis has inherent limitations arising from (1) the absence of primary empirical data collection within this work, (2) potential selection bias induced by relying solely on the provided corpus, and (3) the rapid evolution of threat landscapes and cloud platforms beyond the dates of cited works. The absence of field measurements — for example, quantifying latency impacts of blockchain attestations on authorization decision times — limits the ability to assert operational feasibility across diverse scale regimes. Additionally, because some cited works focus on domain-specific cases such as healthcare or IoT, extrapolations to general multi-tenant cloud contexts require careful validation (Chen et al., 2020; Tyler & Viana, 2021).

Future Research Directions

To bridge the gap between conceptual promise and operational reality, the literature synthesis suggests a research agenda organized into measurement, design, governance, and socio-technical experimentation.

Measurement: Develop standardized metrics for zero-trust assurance and usability. Metrics should include decision latency (time from access request to authorization decision), false acceptance and false rejection rates for dynamic policy engines, user task completion times under different authentication regimes, and measurable indicators of security fatigue over time. These metrics must be validated via field studies across representative tenant profiles to ensure ecological validity (Strouble et al., 2009; Stanton et al., 2016).

Design: Investigate hybrid attestations that combine hardware roots of trust, ephemeral telemetry, and selective immutable provenance. Experimental work should quantify trade-offs between latency and assurance for blockchain-backed attestations, and explore layered caching or off-chain approaches to mitigate performance impacts while preserving integrity guarantees (Alevizos et al., 2022; Kim et al., 2020).

Governance: Study policy distribution models that balance provider consistency with tenant autonomy. Comparative studies of centralized PDPs (policy-as-a-service) versus federated PDPs are needed, focusing on operational complexity, auditability, and tenant perceptions of control. Investigate legal and contractual frameworks for cross-tenant telemetry sharing that reconcile privacy regulations with security needs (Khan, 2023; Hariharan, 2025).

Socio-technical Experimentation: Conduct longitudinal deployments to quantify how security fatigue evolves under different authentication policies, step-up models, and UX interventions. Experiment with adaptive

authentication strategies that minimize friction while preserving assurance, and measure their impact on both compliance and operational security incidents (Weidman & Grossklags, 2017; Stanton et al., 2016).

Practical Recommendations for Deployment

For organizations and cloud providers seeking to adopt zero trust in multi-tenant contexts, the analysis suggests several pragmatic recommendations grounded in the literature.

Adopt Identity as Primary Control Plane: Prioritize investment in robust IAM, including MFA, attribute management, and federated identity models. Enable ABAC to support dynamic, context-sensitive authorization. Where feasible, implement SSO to reduce repeated authentication friction (Yao et al., 2020; Mehraj & Banday, 2020).

Design for Usability: Implement adaptive authentication that escalates only upon anomalous signals; simplify recovery and provisioning flows; conduct user-centred testing to identify and reduce security fatigue drivers. Provide clear communication and training to improve perceptions and adherence (Strouble et al., 2009; Stanton et al., 2016).

Instrument Telemetry Carefully: Build a telemetry fabric that supports both real-time enforcement and audit needs. Use standards-based messaging for security events and posture updates, and align telemetry granularity with privacy and regulatory constraints (Cam-Winget et al., 2019; Software Defined Perimeter Working Group, 2014).

Pilot Endpoint Augmentations Conservatively: When experimenting with blockchain-backed attestations, start with controlled pilots where consortium governance, latency budgets, and privacy encodings are well-defined. Evaluate alternative architectures (off-chain registries, hashed anchors) to balance performance and provable integrity (Alevizos et al., 2022).

Establish Clear Governance: Create policy lifecycle processes with testing, safe deployment, rollback, and impact analysis. Define tenant roles and responsibilities for policy control, and provide transparency regarding telemetry usage and audit rights to foster trust (Khan, 2023; Tyler & Viana, 2021).

CONCLUSION

Zero-trust architecture offers a principled and necessary approach for securing multi-tenant cloud environments. The synthesis of standards, protocol frameworks, human-factors studies, and experimental augmentations presented in this article demonstrates both the potential and the complexity inherent in operationalizing zero trust. Identity-centered control planes, continuous device and workload attestation, real-time telemetry, and microsegmentation are the core technical scaffolding. Yet, human factors such as security fatigue, usability trade-offs of stronger authentication, and tenant perceptions of control critically determine whether these architectures produce intended security outcomes.

Blockchain and other decentralized technologies augment zero trust by providing immutable provenance and new attestation models, but they are not universally applicable and demand careful governance and performance engineering. Protocol primitives such as XMPP for security information exchange and the Software Defined Perimeter specification provide practical paths for implementing essential capabilities, but they must be embedded within policy and governance frameworks that articulate tenant rights, auditability, and change management.

The path forward requires interdisciplinary work: rigorous field studies to quantify the interplay of usability

and assurance; engineering research to reduce latency and scalability barriers for augmented attestations; policy work to reconcile tenant autonomy with provider consistency; and organizational research to craft sustainable governance models. By explicitly integrating technical, human, and governance considerations, the research and practitioner communities can transform zero trust from a guiding philosophy into a matured, operational security paradigm for multi-tenant clouds. The literature synthesized here provides a robust foundation for this transition, and the prioritized agenda offers a roadmap for future empirical inquiry and responsible deployment (Khan, 2023; Hariharan, 2025; He et al., 2022).

REFERENCES

1. Cam-Winget N (ed.), Appala S, Pope S, Saint-Andre P (2019) Using Extensible Messaging and Presence Protocol (XMPP) for Security Information Exchange. (Internet Engineering Task Force (IETF)), IETF Request for Comments (RFC) 8600. https://doi.org/10.17487/RFC8600

2. Software Defined Perimeter Working Group "SDP Specification 1.0" Cloud Security Alliance. April 2014.

3. Stanton B, Theofanos MF, Spickard Prettyman S, Furman S (2016) Security Fatigue. IT Professional 18(5):26-32. https://doi.org/10.1109/MITP.2016.84

4. Strouble D, Shechtman GM, Alsop AS (2009) Productivity and Usability Effects of Using a Two-Factor Security System. SAIS 2009 Proceedings (AIS, Charleston, SC), p 37. Available at http://aisel.aisnet.org/sais2009/37

5. Weidman J, Grossklags J (2017) I Like It but I Hate It: Employee Perceptions Towards an Institutional Transition to BYOD Second-Factor Authentication. Proceedings of the 33rd Annual Computer Security Applications Conference (ACSAC 2017) (ACM, Orlando, FL), pp 212-224. https://doi.org/10.1145/3134600.3134629

6. M. J. Khan, "Zero trust architecture: Redefining network security paradigms in the digital age," World Journal of Advanced Research and Reviews, pp. 105-116, 2023.

7. Hariharan, R. (2025). Zero trust security in multi-tenant cloud environments. Journal of Information Systems Engineering and Management, 10.

8. D. Tyler and T. Viana, "Trust no one? a framework for assisting healthcare organisations in transitioning to a zero-trust network architecture," Applied Sciences, p. 7499, 2021.

9. B. Chen, S. Qiao, J. Zhao, D. Liu, X. Shi, M. Lyu and Y. Zhai, "A security awareness and protection system for 5G smart healthcare based on zero-trust architecture," IEEE Internet of Things Journal, pp. 10248-10263, 2020.

10. S. Mehraj and M. T. Banday, "Establishing a zero trust strategy in cloud computing environment," International Conference on Computer Communication and Informatics, pp. 1-6, 2020.

11. L. Alevizos, V. T. Ta and M. Hashem Eiza, "Augmenting zero trust architecture to endpoints using blockchain: A state- of- the- art review," Security and Privacy, p. 191, 2022.

12. Kim, J. Oh, J. Ryu and K. Lee, "A review of insider threat detection approaches with IoT perspective," IEEE Access, pp. 78847-78867, 2020.

13. Q. Yao, Q. Wang, X. Zhang and J. Fei, "Dynamic access control and authorization system based on zero-

trust architecture," Proceedings of the 2020 1st International Conference on Control, Robotics and Intelligent System, pp. 123-127, 2020.

14. Y. He, D. Huang, L. Chen, Y. Ni and X. Ma, "A survey on zero trust architecture: Challenges and future trends," Wireless Communications and Mobile Computing, 2022.