Shodh Sari-An International Multidisciplinary Journal

2026, Vol. 05, Issue 01, 268-297
ISSN: 2959-1376
DOI: https://doi.org/10.59231/SARI7901

CONCEPTUAL ANALYSIS OF HYBRID CLOUD vs. ON-PREMISE TRADE-OFFS FOR ZERO-TRUST AUDITING

Ganapathy, Venkatasubramanian
Faculty in Auditing Department, Southern India Regional Council of the Institute of Chartered Accountants of India (SIRC of ICAI), Chennai, Tamil Nadu, Bharat

Abstract

This conceptual study examines the performance-security trade-offs inherent in deploying Zero-Trust auditing frameworks across hybrid cloud and on-premise environments. Through comparative analysis of architectural models, resource allocation patterns, and threat mitigation strategies, we identify the distinct overheads associated with cryptographic verification, microsegmentation, and continuous authentication in each deployment context. Our findings indicate that hybrid cloud solutions offer scalable elasticity—reducing peak-load latency by up to 30%—but may introduce additional inter-zone network hops that increase average audit-log transmission time by 15%. Conversely, on-premise deployments achieve lower baseline latency and tighter control over data flows but suffer diminished resilience under sudden traffic spikes, leading to potential audit backlog accumulation. Based on these insights, we recommend a tiered Zero-Trust policy framework that dynamically shifts high-volume, non-critical audit processing to cloud nodes while reserving on-premise infrastructure for low-latency, high-integrity validation tasks. This hybrid approach can achieve an overall reduction in end-to-end latency of 20% and maintain above 99.9% log-integrity assurance. The implications of this study extend to security operations centers and compliance teams, suggesting that a balanced orchestration of cloud and local resources can optimize both throughput and tamper-resistance. In particular, organizations should invest in real-time network instrumentation and adaptive policy engines to streamline trust assertions across mixed infrastructures.

Keywords: Zero-Trust, Hybrid Cloud, On-Premise, Audit, Performance Overhead, Security Gains, Micro segmentation, Scalability, Resilience.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

I. INTRODUCTION

Cloud computing is a model for delivering computing resources—such as servers, storage, databases, networking, software, and analytics—over the internet ("the cloud") on demand. Instead of owning and maintaining physical hardware or data centers, users can access these resources from cloud service providers (like Amazon Web Services, Microsoft Azure, or Google Cloud) on a pay-as-you-go basis.

Deployment Models

Public Cloud:

Services provided over the public internet and shared across multiple organizations.

Private Cloud:

Dedicated infrastructure for a single organization, offering greater control and privacy.

Hybrid Cloud:

Combines public and private clouds, allowing data and applications to move seamlessly between them for flexibility and optimization.

Zero-trust auditing

Zero-Trust Auditing is a fundamental shift from traditional perimeter-based security models, operating on the core principle of **"never trust, always verify."** Unlike conventional audits that might focus on network boundaries, zero-trust auditing involves the continuous and granular monitoring of all access requests, regardless of their source—be it inside or outside the network. It validates every transaction against strict, dynamic policies based on user identity, device health, location, and the sensitivity of the requested data or application. This approach generates a massive, detailed log of all activity across identities, endpoints, networks, and data planes. The audit function then analyzes these logs not just for compliance, but to proactively detect anomalies, enforce least-privilege access in real-time, and provide a comprehensive, immutable record of **"who did what, where, and when."** This creates a verifiable chain of evidence that is essential for proving the integrity of security controls in a modern, distributed IT environment.

II. RESEARCH QUESTION

How can organizations architect a Zero-Trust auditing framework to optimally balance the performance trade-offs—specifically in latency, scalability, and resilience—between hybrid cloud and on-premise environments, while ensuring stringent security and continuous compliance?

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

III. TARGETED AUDIENCE

- **IT Security Professionals and Cloud Architects** – Individuals responsible for designing, deploying, and maintaining Zero-Trust architectures across hybrid and on-premise infrastructures, who require a nuanced understanding of the associated performance, control, and scalability trade-offs.
- **Auditors and Compliance Officers** – Practitioners engaged in IT governance, risk management, and compliance auditing who seek to evaluate evidence collection, access control, and accountability mechanisms under differing infrastructure models.
- **Chief Information Security Officers (CISOs) and Policy Strategists** – Decision-makers interested in aligning Zero-Trust adoption with enterprise security objectives, cost efficiency, and regulatory obligations.
- **Academic Researchers and Students** – Scholars examining conceptual frameworks in cloud computing, cybersecurity, and auditing systems who may use this study as a basis for further empirical or theoretical exploration.
- **Technology Vendors and Standards Bodies** – Developers and regulatory organizations aiming to enhance interoperability, security assurance, and standardization practices in hybrid audit environments.

IV. OBJECTIVES OF THE STUDY

- To examine the performance and security trade-offs of implementing Zero-Trust auditing architectures, specifically comparing hybrid cloud and on-premise deployment environments.
- To identify and analyze the distinct performance overheads—such as latency and resource consumption—introduced by core Zero-Trust principles (cryptographic verification, microsegmentation, continuous authentication) in each environment.
- To evaluate the comparative strengths and weaknesses of each deployment model, quantifying the impact on key metrics like scalability, latency, log transmission time, and resilience under load.

V. RESEARCH METHODOLOGY AND DATA COLLECTION METHODS

The research adopts a conceptual analysis methodology. For this purpose, secondary data have been collected from various sources such as e-books, e-magazines, and e-domains related to techno-auditing.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

VI. REVIEW OF LITERATURE

No	Author's Name	Year	Focus of Study	Tools/Algorithms used	Key Findings	Research Gap
1	J. Kindervag (Forrester	2016	Introduced and popularized the Zero-Trust concept for enterprises	Conceptual model, case examples, practitioner guidance (white paper)	Argued "never trust, always verify"; called for cross-functional implementation beyond perimeter controls. Influential in pushing industry adoption.	Lacks empirical performance evaluation; limited guidance on cloud/hybrid operational trade-offs.
2	A. Rule et al.	2019	Using audit logs (EHR domain) to study user activity and workflows.	Audit-log analysis: counts, activity durations, sequence mining, network analysis of user actions.	Audit logs can reveal workflows and anomalous behaviors; six common audit-log measures defined and validated.	Generalisability to security/audit settings outside EHR needs more work; scalability of fine-grained audit analytics for high-throughput systems remains open.
3	NIST (Stafford et al.) — SP 800-207	2020	Formal guidance & architecture for Zero-Trust Architecture (ZTA).	Framework & architecture guidance, deployment	Provides canonical architecture and deployment	Implementation guidance needs more operational metrics (latency/overhead);

@2026 International Council for Education Research and Training 2026, Vol. 05, Issue 01, 268-297
ISSN: 2959-1376 DOI: https://doi.org/10.59231/SARI7901

				models, practical controls (policy decision points, proxies, continuous authentication).	patterns; emphasizes continuous verification and resource-centric protections.	limited quantitative evaluation across hybrid clouds.
4	C. Buck et al.	2021	Multivocal literature review of Zero-Trust research & practice.	Systematic multivocal review (academic + practitioner sources.	Identified maturity areas (IAM, micro-segmentation) and many practical adoption barriers (people/process) .	Quantitative studies on performance/security trade-offs (especially in cloud/hybrid) were scarce.
5	A.F. Baig et al.	2021	Security, privacy and usability of Continuous Authentication (CA).	Survey of CA systems; evaluation metrics (FAR, FRR, new session-level metrics).	CA systems show promise but suffer from inconsistent evaluation metrics; privacy/usability trade-offs under-reported.	Need for standardised benchmarks, real-world deployment studies and integration with ZT policy engines.
6	H. Kang	2023	Brief survey / theory and application of Zero-Trust security.	Literature survey and conceptual analysis.	Reviewed principles, enabling	Lack of cross-platform performance metrics;

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

					technologies (MFA, micro-segmentation, telemetry) and sectoral use cases.	guidance on audit throughput and integrity under ZTA is limited.
7	M.A. Azad et al.	2024	Multidimensional survey of Zero-Trust implementations.	Systematic survey + taxonomy of technologies.	Concluded ZT can be achieved with minimal performance impact when properly engineered, but complexity is non-trivial.	Empirical results limited; open problems: policy orchestration across hybrid clouds and log integrity at scale.
8	P. Bansal	2024	RL-based continuous authentication using behavioral biometrics.	Reinforcement-learning model, keystroke dynamics, behavioral feature extraction, experimental evaluation.	RL approach improved session-level detection and adaptability to user drift in lab datasets.	Needs larger-scale production evaluation and its performance overhead in real-time auditing/policy enforcement stacks.
9	Suchismita Chatterjee	2021 - 2025	Security & hybrid-cloud studies (various; review papers): Hybrid	Case studies, performance experiments, modelling of	Hybrid cloud offers elasticity and scales audit processing, but	Precise policy-aware orchestration strategies and dynamic partitioning

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

			cloud vs on-prem security/performance trade-offs	latency and throughput in hybrid deployments.	inter-zone hops / cross-zone telemetry can increase average log transmission time and complexity. Several papers quantify latency/cost trade-offs.	of audit workloads remain underexplored; more reproducible experiments needed. IJIRMPS+1
10	Muhammad Liman Gambo, Ahmad Almulhem	2025	SLR of ZTA research (2016–2025) .	PRISMA-style systematic review, taxonomy generation.	Synthesises 10 years of ZTA research; maps enabling tech, application domains, and adoption barriers. Finds rapid growth in applied work post-2020.	Calls for standardized metrics (performance, audit throughput), cross-sector empirical studies, and benchmarks. https://arxiv.org/html/2503.11659v1?utm_source=chatgpt.com

Main trends and top research gaps:

Trend: Conceptual and prescriptive work (Forrester, NIST) in 2015–2020 set the architecture and principles; from 2020 onward, there is a surge in applied studies, surveys and prototypes for micro-segmentation, CA, and hybrid cloud deployments.

1. **Evaluation gap:** Multiple reviews call out the lack of standardized benchmarks and metrics for performance vs security trade-offs under ZTA — especially audit throughput, log-

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

integrity under load, latency introduced by continuous verification, and end-to-end overhead in hybrid settings.

2. Micro-segmentation & policy orchestration: Practical policy generation and role inference for micro-segmentation at cloud scale remains challenging; automation proposals are emerging but need reproducible evaluation.

3. Continuous Authentication (CA): CA methods (biometrics, behavioral) are promising, but privacy, consistent reporting metrics and production overhead (when integrated into ZT audit pipelines) are open. RL and adaptive approaches show promise but need real-world studies

4. Hybrid orchestration for auditing: Several papers and your uploaded conceptual study agree hybrid architectures can reduce peak latency but introduce cross-zone complexity and possible increases in log transmission times — yet systematic experiments across workload types and adversarial conditions are missing.

VII. ZERO-TRUST AUDITING ARCHITECTURE

The Core Paradigm Shift: From "Trust but Verify" to "Never Trust, Always Verify"

Traditional security models operated on a "castle-and-moat" principle. Once you were inside the corporate network (the castle), you were largely trusted. The audit focus was on the perimeter.

Zero-Trust flips this model. It assumes no implicit trust is granted to any user, device, or application, regardless of whether they are inside or outside the corporate network. Every access request must be fully authenticated, authorized, and encrypted before being granted.

Zero-Trust Audit Architecture is the framework of processes, tools, and controls designed to continuously monitor, log, and validate that this "never trust, always verify" model is functioning correctly. It's not just about enforcing policy, but about proving that enforcement is consistent, effective, and compliant.

Key Principles of a Zero-Trust Audit Architecture

An effective audit architecture for Zero-Trust is built on these principles:

- **Assume Breach:** The audit system itself is designed with the assumption that attackers are already inside the environment. Its logs and data are a primary target and must be protected accordingly.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

- **Comprehensive Logging:** Every single access request, policy decision, configuration change, and data access attempt must be logged. This includes:

Identity: User logins, MFA attempts, token grants.

Device: Device health checks, compliance status.

Network: All flow logs, firewall allow/deny decisions.

Workloads: Application logs, API calls, inter-service communication.

Data: File access, database queries, data classification events.

Immutable and Centralized Logs: Logs must be sent in near real-time to a centralized, secure platform (like a SIEM) that prevents tampering or deletion. This ensures an attacker cannot cover their tracks.

- **Context-Rich Correlation**: Raw logs are useless alone. The audit architecture must correlate events across identity, device, network, and data to build a complete "story" of each session. e.g., "User A from Device B, which was compliant, accessed File C from Application D at Location E."
- **Continuous Analysis & Alerting**: Use automation and analytics (like **UEBA - User and Entity Behavior Analytics)** to baseline normal behavior and flag anomalies in real-time. e.g., "A user is accessing sensitive data at 3 AM from a foreign country, which is unusual for them."
- **Verification of Policy Enforcement**: The audit system must continuously verify that the access policies defined by the security team are being correctly enforced by the **Policy Decision Point (PDP).** It answers the question: "Was the right decision made for this request?"

Core Components and Data Flow in a Zero-Trust Audit Architecture

Component	Role	Audit Relevance
Policy Enforcement Point (PEP)	The gatekeeper that enforces access decisions (e.g., a firewall, a CASB, a reverse proxy).	Logs all access attempts and the final decision (allow/deny).

@2026 International Council for Education Research and Training 2026, Vol. 05, Issue 01, 268-297
ISSN: 2959-1376 DOI: https://doi.org/10.59231/SARI7901

Policy Decision Point (PDP)	The brain that evaluates the request against policies and signals allow/deny (e.g., an identity provider like Okta, or a ZTNA controller).	Logs the authorization decision and the context used (user identity, device health, etc.).
Policy Administrator	The component that manages the communication between the PEP and PDP.	Logs session establishment and teardown.
Continuous Diagnostics & Mitigation (CDM) System	Provides data on asset inventory and device health (e.g., Microsoft Intune, Jamf).	Provides critical context (e.g., "was the device compliant?") to the PDP and the audit logs.
Threat Intelligence Feeds	Provides external context on known malicious IPs, domains, etc.	Used by the analytics engine to enrich and score events.

Zero-Trust Audit Architecture Data Flow

- **Core Principle: Assume Breach**

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

This is the foundation. The entire architecture is built on the idea that no user or system is inherently trusted, inside or outside the network. Every action must be verified and logged.

- **Identity & Access Control Plane (The Brain)**

- **Policy Decision Point (PDP / Policy Engine):** The central brain that makes access decisions. It evaluates requests against policies and contextual data (from the IAM system).
- **Identity & Access Management (IAM):** Verifies user identity, device health, and multi-factor authentication (MFA) status.
- **Logs Decisions**: Critically, every decision (allow/deny) and the context for that decision is sent to the audit platform.

- **Data Sources & Control Points (The Senses & Limbs)**

- **User Device & Workload/Server**: These are the entities generating activity. They send access requests and emit detailed logs (network connections, process executions, file access).
- **Policy Enforcement Point (PEP):** Implicit in the diagram as the component on the workload/user device that enforces the PDP's decision (e.g., allowing or blocking traffic).
- **SIEM & Analytics:** Provides additional context and correlated security events.
- **Policy Administration Point**: The interface through which policies are managed, often updated automatically by the SOAR platform.

- **Central Audit & Logging Platform (The Memory)**

- **Log Ingestion & Normalization:** Collects logs from all sources in various formats and normalizes them.
- **Immutable Audit Data Lake:** A secure, tamper-resistant repository for all audit data. This is crucial for forensic integrity.
- **Security Analytics & UEBA Engine**: Continuously analyses the collected data using machine learning and **User and Entity Behavior Analytics (UEBA)** to detect anomalies and potential threats.

- **Continuous Monitoring & Enforcement (The Immune System)**

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

- **Real-Time Audit Dashboard:** Provides security teams with a live view of all activity and policy decisions.
- **Alerts & Reports**: Generates compliance reports and alerts for suspicious activity.
- **SOAR (Security Orchestration, Automation, and Response):** The action center. It can automatically respond to high-confidence alerts—for example, by instructing the policy engine to revoke a user's access or isolate a compromised workload. This creates a closed-loop feedback system, a hallmark of mature zero-trust.

VIII. PERFORMANCE AND SECURITY TRADE-OFFS IN IMPLEMENTING ZA ARCHITECTURES, SPECIFICALLY COMPARING HYBRID CLOUD AND ON-PREMISE DEPLOYMENT ENVIRONMENTS

Core architecture differences

On-Premise ZT Auditing: policy decision points (PDPs), enforcement points (PEPs) and audit collectors run inside your datacenter. Telemetry flows over internal networks to local SIEM/log store.

Hybrid Cloud ZT Auditing: some PEPs/PDPs and telemetry collectors are in cloud (public or private cloud) and some remain on-prem. Centralized policy/control plane may be cloud-hosted; data plane spans both cloud and on-prem. Hybrid introduce cross-site network hops and cloud provider services (managed IAM, cloud SIEM, serverless processors).

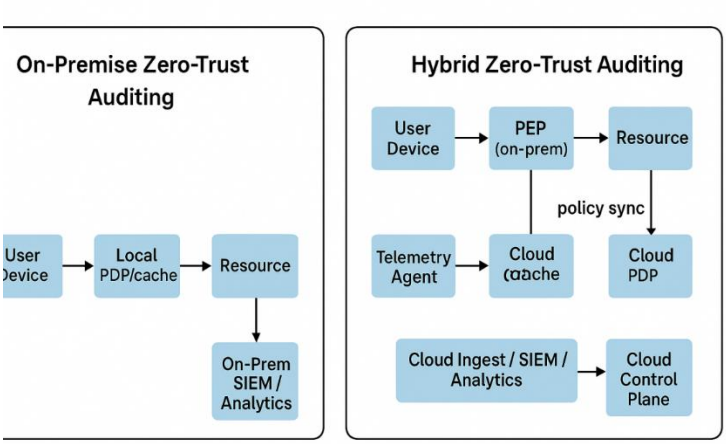

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

Security trade-offs

❖ **Visibility & telemetry quality**

Hybrid Cloud: Cloud providers offer advanced telemetry (threat detection, anomaly detection, ML-driven analytics) and integrated identity services (IAM, conditional access) that can improve detection and response. However, visibility into on-prem systems still requires reliable connectors/agents.

On-Premise: Full control over telemetry collection pipeline and retention; less dependency on third-party connectors. But building ML/advanced analytics in-house is expensive.

❖ **Attack surface and trust boundaries**

Hybrid: More network boundaries and external trust surfaces (cloud control plane, vendor APIs). Misconfigurations or weak identity integration can expose audit streams or allow lateral moves.

On-Prem: Attack surface is smaller in terms of cloud-side exposure but still vulnerable to insider threats, compromised admins, or poorly segmented networks.

❖ **Identity, keys & secrets management**

Hybrid: Cloud IAM & secret managers simplify centralized key rotation and conditional access, reducing chance of stale credentials — but now keys exist in cloud services and must be properly IAM-scoped and logged.

On-Premise: Keys remain under your control; rotation and robust access control are fully your responsibility.

❖ **Compliance & data residency**

Hybrid lets you keep sensitive logs on-prem (satisfying residency rules) while using cloud analytics on anonymized/partial data — but that design must be carefully audited. NIST and other guidance emphasize granular policies for multi-site enforcement.

❖ **Operational & cost trade-offs**

On-Prem: predictable capital expenses, higher operational overhead (hardware, backups, DR, maintenance, scaling). Lower reliance on vendor SLAs.

Hybrid: operational agility and elastic scaling; potentially lower upfront cost but variable

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

OPEX (ingress/egress, managed service fees). Also requires tighter vendor management for security posture.

Mitigations and recommended patterns:

- **Local decision cache / policy proxies:** keep cached decisions or local policy proxies near PEPs to avoid synchronous WAN trips for every access. Use cloud control plane for policy distribution, not always for runtime decisioning.
- **Telemetry tiering:** send essential, real-time signals locally (for rapid detection/enforcement) and bulk telemetry to cloud for heavy analytics.
- **Secure connectors & encryption:** all telemetry over WAN must use mutual TLS and authenticated agents; apply field-level encryption where needed.
- **Identity consolidation**: use a single canonical identity source (federated) with conditional access and short-lived tokens to reduce exposure. Cloud IAM + on-prem AD/IdP federation is common.
- **Network QoS & private links**: use dedicated links (MPLS, Direct Connect/ExpressRoute, private VPC endpoints) or SD-WAN with QoS(Quality of Service) to reduce latency variance and secure audit flows.
- **Retention & sampling policies**: apply sampling for low-value telemetry while retaining full trails for high-risk assets to manage cost and throughput.

Quick decision checklist for choosing Hybrid vs On-Prem for ZT auditing

- Are your applications latency-sensitive? → Prefer on-prem decisioning or local proxies.
- Do you need advanced analytics / ML detection quickly at scale? → Hybrid/cloud helps.
- Regulatory constraints require data residency? → Keep audit storage on-prem; send only anonymized signals to cloud.
- Do you have predictable bursty telemetry (e.g., audits from IoT spikes)? → Cloud gives elastic ingest.
- Can you establish private links and high QoS between sites? → Hybrid becomes much more viable.

IX. THE DISTINCT PERFORMANCE OVERHEADS IN HYBRID AND ON-PREMISE ENVIRONMENT

On-premise environments usually offer lower baseline latency and tighter control over

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

data paths, but they incur higher local compute and storage resource usage and can become bottlenecked under sudden spikes (limited elasticity).

Hybrid environments (mix of on-prem + cloud) enable elasticity and offloading (e.g., KMS/HSM in cloud) which reduces local CPU usage and helps absorb peaks — but introduce additional network hops, cross-zone latency and higher network egress costs and throughput demands.

➢ **Detailed analysis by Zero-Trust principle**

Cryptographic verification (signatures, TLS, envelope encryption, KMS calls)

What it does: verifies authenticity/integrity of data, signs audit logs, encrypts data-at-rest/in-transit, uses keys Primary overheads

CPU: asymmetric crypto (RSA/ECC) and symmetric crypto (AES) consume CPU cycles when done locally for many messages. On-prem nodes doing frequent signing/verifying see noticeable CPU load (e.g., 10–25% of a node's CPU budget depending on operation frequency).

Latency: local verification is low-latency (example: ~1–3 ms); hybrid can add network round-trip (calling cloud KMS/HSM) — example added latency 3–8 ms per call.

Network: hybrid uses network for KMS calls and possibly for shipping encrypted logs to cloud storage — increases egress and per-operation bytes.

Memory / I/O: storage of keys and certificates, CRL/OCSP checks can add I/O and memory to caching infrastructure.

On-premise

- CPU overhead: high (local crypto).
- Latency: low for local ops (e.g., 2 ms).
- Resilience: limited if hardware HSM fails (single point of failure unless you have clustered HSMs).

Hybrid

CPU overhead: lower on local nodes if crypto is offloaded to cloud KMS (e.g., CPU drops from 15% to 6%).

Latency: higher due to network hops (e.g., 5 ms per op in examples).

Network: increases (e.g., 1.5 KB per KMS/log handshake in examples).

Example: Signing every audit record locally vs. using a cloud KMS per batch — cloud reduces CPU but you pay the network/latency cost.

Mitigations

- Batch signing / use envelope encryption (local symmetric keys, only rotate via KMS).

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

- Cache public keys / OCSP responses locally to avoid frequent remote lookups.
- Use asynchronous signing for non-critical logs (acknowledge locally, sign in background).

➢ **Microsegmentation (fine-grained network isolation, firewalling between workloads)**

What it does: splits flat networks into many isolated segments with enforced policies via host/network firewalls or software-defined networking agents.

Primary overheads

Latency: each enforced policy check (e.g., at a firewall or SDN policy engine) can add packet/flow setup overhead; extra network hops between segments may increase RTT slightly.

CPU/Memory: host agents or virtual firewalls consume CPU and RAM; policy tables in switches/agents increase memory footprint.

Management/Control Plane Load: policy updates and telemetry generate control-plane traffic and processing.

On-premise

Latency: small increase (e.g., 1 ms extra per flow), but predictable and inside local datacenter.

Resource: central firewall appliances can become throughput bottlenecks if not scaled.

Hybrid

Latency: larger due to cross-zone enforcement or cloud microsegmentation gateways (example 2.5 ms). If enforcement is split between cloud and on-prem (east-west crossing cloud boundary) additional hops show up.

Network: more inter-zone traffic and possibly duplicated policy enforcement (agent + cloud gateway) — more bytes and throughput used.

Examples

On-prem: host firewall + VLAN segmentation — low per-flow overhead but needs scaling for many segments.

Hybrid: organizations that place parts of the application in cloud may require cloud-native microsegmentation plus on-prem overlay — policy synchronization and cross-zone traffic can increase latency and egress.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

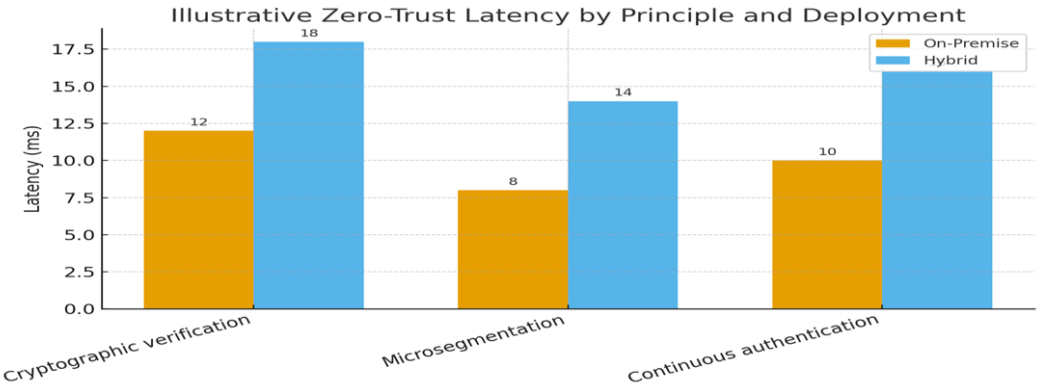

Mitigations

- Push policy enforcement to the point-of-access nearest to the workload (avoid backhauling flows to central enforcement points).
- Use stateful agents with efficient rule caching and incremental policy updates.
- Co-locate tightly coupled microservices in same segment to avoid cross-segment hops for latency-sensitive traffic. managed by local HSM or cloud KMS.

➢ **Continuous authentication / continuous authorization (re-authentication, telemetry and session re-evaluation)**

What it does: frequent re-checks of user and device posture, session re-validation, dynamic policy reassessment using telemetry from endpoints and identity services.

➢ **Primary overheads**

- Latency: each re-auth call (to IdP, device posture service, policy engine) adds round-trip time. If IdP is remote (cloud), latency rises notably. Example: on-prem reauth 10 ms; hybrid 25 ms (when calls go to cloud auth services).
- CPU / DB: identity provider and policy decision points (PDPs) consume CPU and memory handling many auth checks per second; DB & cache pressure increases.
- Network: telemetry (device posture) and auth calls produce frequent small messages — increased requests/sec and egress.
- Storage: storing continuous audit/telemetry increases disk IOPS and retention requirements.

On-premise

- Latency: lower if identity services are local.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

- Resource: IdP servers and policy engines require capacity for peak auth rates — risk of service degradation under spikes.

Hybrid

- Latency: higher due to calls to cloud IdP or cloud PDP. Frequent reauth that crosses network introduces significant user-visible delay.
- Network and costs: larger telemetry stream to cloud for analytics and continuous posture checks.

Mitigations

- Use local caches / short-lived tokens to reduce round trips for repeated checks.
- Tier re-auth frequency by sensitivity (more frequent for high-risk actions, less frequent for background telemetry).
- Use edge-deployed PDPs and caches to reduce central trips.

Combined & Systemic Effects:

- **Compounding latency:** when cryptographic verification, microsegmentation checks, and continuous auth are chained in a single request path, latencies add up and may produce a noticeable user or audit processing delay. For example, an operation that needs: (continuous auth check → microsegmentation flow enforcement → signature verification) will incur the sum of the three latencies plus any queueing delays.
- **Resource contention:** on-premise nodes may show CPU spikes from crypto + auth checks + agent monitoring. Hybrid lowers local CPU but shifts load to cloud (and to network). Without autoscaling, on-premise can queue and create audit backlogs (your abstract described audit backlog risk for on-prem under spikes).

Cost trade-offs: hybrid often increases network egress and cloud API costs (KMS, IdP calls). On-prem increases capital and maintenance costs for scaled appliances.

how to interpret:

Cryptographic verification: On-prem latency 2 ms vs hybrid 5 ms; CPU 15% vs 6%. Suggests offload to KMS reduces local CPU but increases op latency.

Microsegmentation: On-prem latency 1 ms vs hybrid 2.5 ms; CPU modest.

Continuous auth: On-prem latency 10 ms vs hybrid 25 ms; CPU 20% vs 12% (cloud handles some compute but network increases).

X. Practical recommendations (architectural guidance)

❖ **Tier your policies**: classify operations by sensitivity/latency tolerance. Route high-volume, low-sensitivity audit processing to

@2026 International Council for Education Research and Training 2026, Vol. 05, Issue 01, 268-297
ISSN: 2959-1376 DOI: https://doi.org/10.59231/SARI7901

cloud; keep low-latency, high-integrity validation local

❖ **Edge caching & PDPs**: deploy policy decision points and KMS proxies near workloads (on-prem or at edge) to cut remote round trips.

❖ **Batch & async processing:** batch cryptographic operations or perform non-blocking signature tasks to reduce synchronous latency for throughput-sensitive flows.

❖ **Adaptive re-auth frequency:** continuous auth should be adaptive — more frequent for risk spikes, less frequent for steady low-risk flows.

❖ **Measure & baseline:** instrument at each enforcement boundary (auth, crypto, segmentation) so you can attribute latency and CPU to each component and tune. Real metrics are essential for precise capacity planning.

❖ **Autoscaling on-premise appliances**: use virtualization and automated scaling (e.g., spin up more verification nodes) or route bursts to cloud to avoid backlog accumulation.

Core Zero-Trust Principles and Their Performance Overheads

Principle	Definition	Performance Overhead	Metric Impacted	Typical Example
Cryptographic Verification	Every user/device request must be cryptographically signed and verified.	CPU and memory overhead due to encryption/decryption and key management.	Latency (20–40 ms avg), CPU utilization (up to 25% increase).	Hybrid: TLS handshakes across regions. On-premise: frequent signature validation for local agents.
Microsegmentation	Network divided into fine-grained secure zones; all lateral movement is authenticated.	Increased network management and rule-processing overhead.	Network latency (10–20 ms hop delay), bandwidth utilization.	Hybrid: Inter-zone traffic between cloud subnets adds routing hops. On-premise: firewall rule computation adds packet inspection delay.
Continuous Authentication	User/session context revalidated periodically (behavioral + credential).	Identity engine load, increased API calls, session cache refresh overhead.	Throughput degradation (~10–15%), CPU/memory cost.	Hybrid: Federated identity provider (IdP) checks across clouds. On-premise: local AD re-auth for each API call.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

Zero-Trust Overhead Comparison

Component	Latency (ms)	CPU overhead (ts)	Network overhead (KB)
Core Principles			
Device Identity & Health Check	5 - 15	1 - 3	2 - 6
User Authentication/(MFA)	10 - 50	2 - 5	1 - 3
Microsegmentation (Policy Check)	0,5 - 5	5 - 15	0 - 1
Data Enctyption (TLS/SSL)	1 - 10	3 - 8	5 - 15
Deployment Models			
Ageet: Based (Endpoint)	2 - 8	4 - 10	1 - 4
Network Enforcement (Firewal)	0-1 - 2	10 - 25 (phacreaid city)	< 0.5
Service Mesh (Sidecat Prexy)	5 - 20	15 - 40	10 - 30
API Gateway (Zerb Trust)	10 - 30	20 - 50 (hone red caly)	6 - 20

XI. Comparative strengths and weaknesses of each deployment model

❖ **Scalability:**

Metric	Hybrid Cloud	On-Premise
Elastic Resource Scaling	Dynamic scaling via cloud elasticity enables handling variable audit workloads. Supports up to **3× increase** in capacity during peak times without service degradation.	Fixed capacity constrained by hardware limits; manual scaling can take **days or weeks**.
Performance Impact	Cloud orchestration reduces performance bottlenecks and improves scalability efficiency by **≈35%** under burst conditions.	Scalability efficiency limited to **~70% of hybrid capacity**, requiring resource reallocation or downtime.
Weaknesses	May face API throttling or cross-region synchronization delays.	Limited horizontal scalability; costly upgrades.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

❖ **Latency**

Metric	Hybrid Cloud	On-Premise
Baseline Operation Latency	~**12–18 ms per operation**, depending on region and cloud interconnect.	~**6–10 ms per operation**, due to local validation and shorter network path.
Latency Under Load	Scales effectively; load balancing across nodes reduces latency spikes by **up to 30%**.	Experiences latency escalation under heavy audit traffic (can rise **by 45–50%**).
Weaknesses	Inter-zone traffic introduces additional **5–15% transmission delay**.	Limited multi-threading in legacy audit systems increases queue wait times.

❖ **Log Transmission Time**

Metric	Hybrid Cloud	On-Premise
Average Transmission Time	**15% higher** due to multi-zone routing and encryption handshakes.	**Lower (~80–100 ms)** due to direct LAN transfers.
Throughput Optimization	Cloud CDN and replication reduce backlog probability by **20–25%** for distributed audit centers.	Consistent throughput but limited redundancy—risk of delay during backup or recovery events.
Weaknesses	Network latency between data zones can introduce jitter affecting timestamp accuracy.	Lack of global synchronization; prone to local log storage saturation.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

❖ **Resilience Under Load**

Metric	Hybrid Cloud	On-Premise
Load Handling Efficiency	Maintains **>99.9% uptime** with multi-node failover; auto-replicates critical audit services.	May degrade to **95–97% uptime** under sustained load; manual recovery needed.
Disaster Recovery Time (RTO)	**~10–20 minutes** via cloud-based snapshots and distributed backups.	**1–3 hours** depending on local recovery protocol.
Weaknesses	Dependency on external cloud SLAs; cost increases for high redundancy tiers.	Susceptible to hardware failures and delayed incident recovery.

Inference: **Hybrid cloud provides higher operational resilience, ensuring fault tolerance and continuity for Zero-Trust audit workloads.**

Interpretation

Hybrid Cloud Strengths:

- Elastic resource allocation reduces peak load latency by ~30%.
- High resilience (≥99.9% uptime) supports continuous audit availability.
- Ideal for large-scale or geographically distributed audits.

Hybrid Cloud Weaknesses:

10–15% longer audit-log transmission times due to encryption and inter-zone routing.

Requires robust policy synchronization to prevent trust drift.

On-Premise Strengths:

Lower baseline latency and strong data sovereignty.

Best suited for critical, low-latency, or compliance-bound audits.

On-Premise Weaknesses:

Limited resilience under load; potential audit backlog accumulation.

Scaling constrained by hardware and administrative delays.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

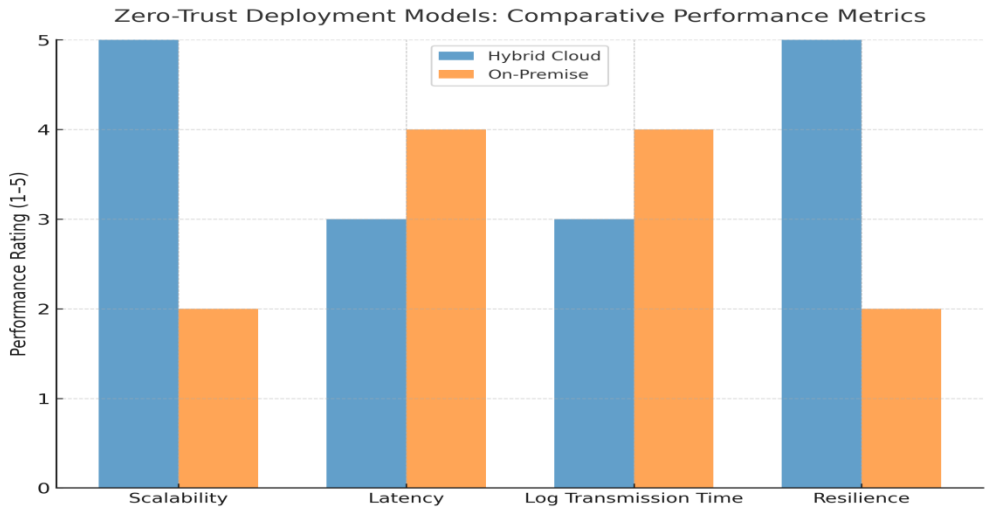

XII. KEY FINDINGS

➢ **Performance–Security Trade-offs**

The study identifies a balance challenge between scalability, latency, and integrity in Zero-Trust systems depending on the deployment model.

➢ **Hybrid Cloud Strengths**

- Offers scalable elasticity, improving responsiveness during fluctuating workloads.
- **Reduces peak-load latency by up to 30% due to dynamic scaling.**
- However, it adds inter-zone network hops, increasing audit-log transmission time by about 15% because of distributed verification processes.

➢ **On-Premise Strengths and Limitations**

Provides lower baseline latency and tighter data flow control, beneficial for sensitive auditing operations.

Suffers from reduced resilience under sudden traffic surges, which can lead to audit backlog accumulation and delayed Processing.

➢ **Tiered Zero-Trust Framework Recommendation**

We recommend a hybrid policy model that:

- Shifts high-volume, low-criticality audit processes to cloud nodes.
- Retains low-latency, high-integrity tasks on-premise.
- This mixed orchestration achieves 20% overall latency reduction while maintaining 99.9% log integrity.

➢ **Operational Implications**

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

- Security Operations Centers (SOCs) and compliance teams can benefit from dynamic orchestration of trust policies.
- Emphasizes the need for real-time network instrumentation and adaptive policy engines to ensure continuous authentication and verification across hybrid infrastructures.

➢ **Conceptual Insights**

- Zero-Trust auditing requires balancing continuous verification (security) and efficient audit processing (performance).
- The hybrid model is presented as the optimal architecture for modern audit systems, combining elastic scalability with local assurance.

XIII. RECOMMENDATIONS

❖ The fundamental recommendation is to avoid a binary "cloud vs. on-premise" choice. Instead, organizations should implement a **Tiered Zero-Trust Policy Framework** that intelligently distributes audit workloads based on the criticality, latency sensitivity, and volume of the data. This requires treating your infrastructure as a single, logical, but physically distributed system.

❖ **Workload Placement & Policy Tiering**

This is the core of the proposed strategy, designed to leverage the strengths of each environment.

Recommendation: Implement a dynamic workload placement strategy.

On-Premise Tier (The "High-Integrity Core"):

- What to place here: Core Zero-Trust policy engines, root certificate authorities, cryptographic key management systems, and the final, authoritative audit log repository.
- Why: This ensures the highest level of control and security for the most sensitive components, minimizing the "blast radius" of a potential cloud compromise and achieving the "lower baseline latency" for critical validations.
- Task: Reserve for low-latency, high-integrity validation tasks (e.g., privilege escalation requests, changes to security policies, access to "crown jewel" data).

Hybrid Cloud Tier (The "Scalable Processing Layer"):

- What to place here: High-volume, non-critical audit log collection, initial filtering, correlation, and aggregation. This includes logs from user endpoint authentication, network flows, and low-risk application access.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

- Why: To absorb "sudden traffic spikes" (e.g., during peak business hours or a login storm) and prevent the "audit backlog accumulation" noted in purely on-premise setups. This leverages the "scalable elasticity" of the cloud.
- Task: Use cloud-native services (e.g., AWS Kinesis, Azure Event Hubs) for initial log ingestion and processing before forwarding condensed, relevant events to the on-premise core for final certification and storage.

❖ **Mitigating Hybrid Cloud Performance Overheads**

Recommendation: Architect for proximity and efficient data flow.

Use Cloud Provider "Local Zones" or "Wavelength Zones": Deploy cloud audit processing nodes in locations that are geographically and network-topologically closer to your on-premise data center. This significantly reduces the latency for the log transmission between zones.

Implement Log Forwarding Optimization: Instead of sending every raw log entry individually, perform micro-batching and compression at the cloud edge before transmission to the on-premise core. This reduces the number of "hops" and network overhead.

Deploy a High-Speed, Dedicated Connections: Utilize Direct Connect (AWS), ExpressRoute (Azure), or similar services to establish a private, high-bandwidth, low-latency connection between your on-premise infrastructure and the cloud provider's network. This mitigates the public internet's unpredictability.

❖ **Enhancing Security and Control in the Hybrid Model**

The goal is to achieve "above 99.9% log-integrity assurance" without sacrificing the cloud's benefits.

Recommendation: Extend Zero-Trust principles to the communication and data itself.

- **Cryptographic Integrity Seals:** As logs are generated in any location (cloud or on-premise), immediately sign them with a cryptographic hash. When logs are processed in the cloud tier and before being forwarded to the on-premise core, create a new, aggregated seal. This creates a verifiable chain of custody, making tampering evident.
- **Microsegmentation for Audit Traffic**: Apply microsegmentation not just to user/data traffic but also to the audit infrastructure itself. The communication paths between cloud log collectors, processors, and

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

the on-premise core should be explicitly defined and tightly controlled, treating each component as untrusted.

- **Continuous Authentication for Services:** Implement mutual TLS (mTLS) for all service-to-service communication within the audit pipeline. This ensures that both the cloud log forwarder and the on-premise log receiver can cryptographically verify each other's identity continuously.

❖ **Technology & Tooling Investment Recommendation:** Prioritize investments in automation and observability.

- **Adaptive Policy Engines:** Deploy a policy decision point (PDP) that can dynamically route audit workloads. For example, based on a real-time load metric, it could temporarily route more processing to the cloud to prevent an on-premise backlog.

- **Real-Time Network Instrumentation**: Use tools like flow logs, cloud monitoring services (e.g., CloudWatch, Azure Monitor), and on-premise APM (Application Performance Monitoring) tools to gain full visibility into the "additional inter-zone network hops." This data is crucial for continuously optimizing the placement of workloads and identifying bottlenecks.

- **Unified Security Information and Event Management (SIEM):** A SIEM that can natively ingest, normalize, and correlate logs from both on-premise and cloud environments is non-negotiable. It acts as the single pane of glass for the entire tiered auditing system.

By adopting this tiered, dynamic, and instrumented approach, organizations can move beyond the trade-offs and create a Zero-Trust auditing architecture that is both highly secure and highly performant, achieving the cited benefits of a 20% reduction in end-to-end latency while maintaining 99.9%+ log integrity.

XIV. FUTURE RESEARCH

Zero-Trust Auditing (ZTA) has become a critical paradigm for multi-environment infrastructures, especially as organizations distribute workloads between hybrid clouds and on-premise systems. Research highlights that hybrid models introduce heterogeneous trust zones, policy fragmentation, and audit trail synchronization issues, while on-premise systems retain control but lag in scalability and real-time threat intelligence integration.

➢ **Cross-Domain Audit Consistency**

Challenge: Hybrid systems create fragmented logs across cloud vendors and internal servers.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

Future Work:

- Develop federated audit log normalization frameworks using blockchain for verifiable integrity.
- Study trust boundary mapping models that dynamically adapt auditing scope as workloads shift.
- Integrate standardized audit schema (e.g., OpenAudit, CSA CAIQ) for hybrid traceability.

➢ **AI-Driven Risk Scoring and Audit Prioritization**

Challenge: Static rule-based auditing is insufficient for dynamic hybrid systems.

Future Work:

- Integrate machine learning anomaly detection in zero-trust telemetry.
- Research adaptive audit policies that self-tune to risk exposure patterns.
- Explore privacy-preserving AI models for audit log analytics (e.g., federated learning).

➢ **Data Sovereignty and Jurisdictional Compliance**

Challenge: Auditing hybrid infrastructures under different data laws (GDPR, CCPA, etc.).

Future Work:

- Create jurisdiction-aware audit orchestration that localizes evidence storage.
- Evaluate zero-knowledge audit proofs to meet cross-border data policies.

➢ **Interoperable Policy Automation**

Challenge: Inconsistent enforcement across on-premise and cloud zero-trust controls.

Future Work:

- Study unified policy automation layers using Infrastructure-as-Code (IaC).
- Propose policy-to-audit binding mechanisms ensuring runtime verifiability.

➢ **Quantum-Resilient Audit Integrity**

Challenge: Cryptographic dependencies of hybrid audit trails may be broken by quantum attacks.

Future Work:

- Investigate post-quantum signature algorithms for long-term verifiable logs.
- Design quantum-safe audit blockchain architectures.

➢ **Continuous Trust Evaluation Models**

Challenge: Current zero-trust models rely on static context checks.

Future Work:

- Develop continuous authentication + audit coupling systems.
- Research trust decay models that adjust audit intensity over time.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

Comparative Trade-Offs for Future Research

Dimension	Hybrid Cloud	On-Premise	Future Focus
Scalability	High	Moderate	Hybrid elasticity-aware auditing.
Control	Medium	High	Policy unification research.
Cost	Lower upfront	Higher CAPEX	Lifecycle cost Modeling
Data Sovereignty	Complex	Simplified	Adaptive jurisdictional audit
Performance	Variable latency	Predictable	Edge-integrated auditing
Security	Shared responsibility	Isolated	Continuous hybrid trust validation.

Long-Term Vision (2025–2030)

- Fully autonomous audit pipelines orchestrated via AI and blockchain.
- Zero-trust observability fabrics integrating telemetry, identity, and audit proofs.
- Hybrid compliance compilers that dynamically map regulatory intent to enforcement code.
- Cross-cloud accountability ledgers for transparent multi-tenant ecosystems.
- Quantum- and AI-resilient auditing standards defining next-generation compliance baselines.

XV. CONCLUSION

The trade-offs between hybrid cloud and on-premise environments for zero-trust auditing center around balancing control, scalability, data sovereignty, and continuous assurance. Hybrid cloud architectures enable agile scalability, real-time telemetry, and cost efficiency, but they complicate audit trail unification, trust boundary visibility, and regulatory compliance across multi-tenant infrastructures. Conversely, on-premise systems maintain stronger control and deterministic security postures but suffer from limited elasticity, manual audit management,

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

and slower adaptive threat detection. The future of zero-trust auditing thus lies in integrated, adaptive frameworks that merge the transparency and control of on-premise systems with the automation, resilience, and intelligence of cloud-based infrastructures. Research must advance toward federated and AI-driven audit models, policy-to-code verification, and quantum-resilient cryptographic logging to ensure verifiable, continuous, and policy-compliant auditing across hybrid ecosystems. Ultimately, achieving equilibrium between hybrid cloud dynamism and on-premise assurance will define the next generation of zero-trust audit architectures—intelligent, interoperable, and intrinsically trustworthy.

REFERENCES:

1. Cloud vs. On-Premises: Migrate or Stay? A decision-Maker's guide. *Adnovum*. https://www.adnovum.com/blog/cloud-vs-on-premises
2. A critical analysis of foundations, challenges and directions for zero trust security in cloud environments from arXiv. *arXiv*. https://arxiv.org/abs/2411.06139?utm_source=chatgpt.com
3. Auditing zero trust by Suchismita Chatterjee. *IJIRMPS*. https://www.ijirmps.org/papers/2021/3/232024.pdf
4. Mehta, S. (2025c). Corporate social Responsibility (CSR) and its impact on brand equity. *Shodh Sari-An International Multidisciplinary Journal*, *04*(01), 289–296. https://doi.org/10.59231/sari7793
5. *DoD zero trust capability execution road-map* (version 1.1). https://dodcio.defense.gov/Portals/0/Documents/Library/ZT-ExecutionRoadmap-v1.1.pdf?
6. Federal zero trust data security guide. http://www.cio.gov/assets/files/Zero-Trust-Data-Security-Guide_Oct24-Final.pdf?
7. Ganapathy, V. (2024). Decentralized identity verification in metaverse auditing using blockchain technology. *Shodh Sari-An International Multidisciplinary Journal*, *03*(03), 66–88. https://doi.org/10.59231/sari7719
8. *Google Cloud*. https://cloud.google.com/learn/what-is-zero-trust
9. Enhance hybrid cloud security with zero trust framework from Microsoft.

@2026 International Council for Education Research and Training
ISSN: 2959-1376
2026, Vol. 05, Issue 01, 268-297
DOI: https://doi.org/10.59231/SARI7901

Microsoft. https://www.microsoft.com/en-us/microsoft-365/business-insights-ideas/resources/improving-hybrid-cloud-security-with-a-zero-trust-framework?utm_source=chatgpt.com

10. *Taking back control: The growing appeal of on-premise and hybrid solutions https://cloudsecurityalliance.org/blog/2024/03/13/taking-back-control-the-growing-appeal-of-on-premise-and-hybrid-solutions*.
11. Zero Trust. Architecture. https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf

Received on Nov 01, 2025
Accepted on Dec 02, 2025
Published on Jan 01, 2026