Journal of Information Systems Engineering and Management
2024, 9(2)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Design and Implementation of Zero Trust Architecture Across Multi-Cloud Providers

Srikanth Nimmagadda
Software Engineer, Info Vision Inc, Texas, United States

ARTICLE INFO

Received: 12 Apr 2024

Accepted: 26 Jun 2024

ABSTRACT

As enterprises increasingly adopt multi-cloud strategies to optimize performance, cost efficiency, and operational resilience, the need for a unified and secure access control model becomes paramount. Traditional perimeter-based security frameworks fail to address the complexities of modern, distributed cloud environments. This study presents a comprehensive methodology for the design, implementation, and evaluation of a vendor-neutral Zero Trust Architecture (ZTA) tailored for multi-cloud ecosystems. The proposed architecture is grounded in the principles of continuous verification, least privilege access, centralized identity management, and real-time policy enforcement across Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). Leveraging open standards such as OAuth 2.0, OpenID Connect, and SAML for federated identity, we integrate cloud-native capabilities including service mesh for microsegmentation and telemetry for real-time monitoring. Experimental validation across hybrid deployments demonstrates improved security posture, policy compliance, and threat detection capabilities, with minimal performance trade-offs. The findings provide actionable guidance for security architects and cloud practitioners to operationalize Zero Trust principles across diverse and dynamic cloud infrastructures.

Keyword: Zero Trust Architecture (ZTA), Multi-Cloud Security, Policy Enforcement, OAuth 2.0 / OpenID Connect / SAML, Cloud Security

1. INTRODUCTION

1.1 Rise of Multi-Cloud Adoption and Associated Security Challenges

The proliferation of cloud computing has driven organizations to adopt multi-cloud strategies, distributing workloads across platforms such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). This approach offers advantages in terms of cost optimization, high availability, regional redundancy, and avoiding vendor lock-in. However, managing security across such diverse environments introduces significant complexity. Each cloud provider has its own tools, policies, and identity systems, making it difficult to maintain consistent security postures. The increased surface area and disparate control planes expose enterprises to threats such as misconfigurations, identity sprawl, and cross-cloud lateral movement.

1.2 Limitations of Traditional Perimeter-Based Security

Historically, cybersecurity architectures were designed around the concept of a trusted internal network protected by a strong perimeter—commonly implemented via firewalls, VPNs, and intrusion detection systems. While this model served well in static, on-premise environments, it is ill-suited for today's fluid, decentralized IT landscapes. Perimeter-based models inherently trust devices and users within the network, often allowing attackers to move laterally once the perimeter is breached. In multi-cloud and hybrid scenarios, this perimeter becomes blurred or nonexistent, rendering such models ineffective in enforcing granular, context-aware security.

1.3 Emergence of Zero Trust as a Paradigm Shift in Cybersecurity

To address the inadequacies of perimeter-focused security, the Zero Trust Architecture (ZTA) model has emerged as a transformative approach. Zero Trust operates under the principle of “never trust, always verify,” treating every access request as untrusted regardless of origin. It emphasizes strict identity verification, least privilege access, microsegmentation, and continuous monitoring. These principles align well with the distributed, dynamic nature of multi-cloud environments, offering a robust framework to manage access and mitigate risks across cloud providers.

1.4 Objectives of the Study: Vendor-Neutral ZTA for Multi-Cloud

This study aims to design, implement, and evaluate a vendor-neutral Zero Trust Architecture capable of unifying access control, identity management, and policy enforcement across AWS, Azure, and GCP. Unlike vendor-specific solutions that may only provide partial coverage, the proposed framework leverages open standards such as OAuth 2.0, OpenID Connect, and SAML to enable federated identity and interoperability. Furthermore, the study integrates cloud-native capabilities such as service mesh-based microsegmentation and centralized telemetry to support continuous verification. The objective is to deliver a practical, scalable, and secure architecture blueprint that security architects and cloud practitioners can adapt to their multi-cloud environments.

2. RELATED WORK

2.1 Existing Zero Trust Models

Zero Trust Architecture (ZTA) has been extensively studied and formalized through multiple models proposed by both governmental and industry bodies. One of the most influential frameworks is the NIST SP 800-207 published by the U.S. National Institute of Standards and Technology, which defines ZTA as a collection of concepts and components for minimizing implicit trust and continuously evaluating risk across the enterprise. The model provides vendor-agnostic guidance for implementing identity-aware and policy-driven access control. In parallel**,** Google’s BeyondCorp initiative pioneered a practical Zero Trust implementation by removing trust from internal networks and securing access based on user identity and device posture, regardless of location. Microsoft’s Zero Trust Maturity Model further elaborates on Zero Trust adoption stages—ranging from traditional security to optimal Zero Trust—providing actionable steps across identity, endpoints, network, data, applications, and infrastructure. While these frameworks serve as foundational references, they are often implemented in silos or tailored for single-vendor ecosystems, which limits their applicability in multi-cloud scenarios.

2.2 Multi-Cloud Security Architectures

The rise of multi-cloud strategies has led to increased research into securing workloads distributed across multiple cloud platforms. Existing solutions primarily focus on federated identity, cross-cloud VPN tunnels, and centralized SIEM tools for logging and incident response. Some architectures leverage cloud-native tools**,** such as AWS IAM, Azure Active Directory, and GCP IAM, to build independent security models per platform. However, integration among these tools is often ad hoc and requires manual configurations or third-party services. Several commercial solutions (e.g., Palo Alto Prisma Cloud, HashiCorp Boundary, and Okta) attempt to provide cross-cloud visibility and access management, but they often introduce complexity, vendor dependency, or lack real-time enforcement across heterogeneous environments.

2.3 Gaps in Current Implementations

Despite progress, current Zero Trust implementations face notable challenges when extended to multi-cloud contexts. Interoperability remains limited, as cloud-specific identity systems and access control mechanisms do not natively integrate across platforms. Visibility and observability across cloud boundaries are fragmented, making it difficult to correlate activities, assess compliance, and detect lateral movement. Policy enforcement is inconsistent and lacks a unified orchestration layer, leading to duplicated configurations and policy drift. Moreover, there is often a lack of automated mechanisms to enforce least privilege and revoke access based on dynamic trust evaluation, particularly in real-time or during incident response.

2.4 Summary of Research Contributions

This study addresses the above gaps by proposing a vendor-neutral Zero Trust Architecture that seamlessly spans across AWS, Azure, and GCP. The framework is designed to support federated identity management, unified policy definition using open standards, and real-time enforcement through cloud-native microsegmentation and telemetry. Key contributions of this research include:

1. A blueprint for implementing Zero Trust across multi-cloud environments without reliance on proprietary tools;
2. An integration strategy for identity and access management that aligns with open protocols (OAuth 2.0, OpenID Connect, SAML);
3. Experimental validation of the architecture's performance, security, and operational feasibility;
4. Practical insights for enterprise adoption of Zero Trust in dynamic, distributed cloud infrastructures.

3. ARCHITECTURE AND DESIGN PRINCIPLES

3.1 Zero Trust Core Tenets

The foundational principles of Zero Trust guide the design of any secure, modern digital infrastructure, particularly in the context of multi-cloud environments. The first and most critical tenet is “never trust, always verify”**,** which dictates that every access request—regardless of origin or previous validation—must be continuously authenticated and authorized based on dynamic context. This includes identity attributes, device posture, location, and behavioral patterns. The second principle, least privilege **access**, ensures that users and services are granted only the permissions strictly necessary to perform their functions. This reduces the attack surface and limits potential damage in the event of a breach. Microsegmentation is another core component that involves dividing network zones into granular segments, allowing fine-grained control over east-west traffic between workloads. This containment mechanism prevents lateral movement by attackers within and across cloud platforms. Lastly, continuous monitoring and assessment are essential to maintaining trust in real time. This includes real-time threat detection, behavioral analytics, and automated response to deviations from baseline behavior, ensuring that trust is not static but dynamically evaluated throughout the session lifecycle.

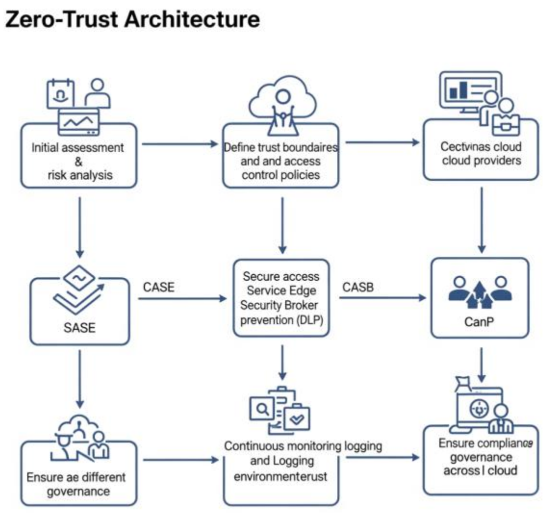
Zero-Trust Architecture
Initial assessment & risk analysis
Define trust boundaires and and access control policies
Cectvinas cloud cloud providers
SASE
CASE
Secure access Service Edge Security Broker prevention (DLP)
CASB
CanP
Ensure ae different governance
Continuous monitoring logging and Logging environmenterust
Ensure complianœ governance across I cloud

Figure 1 Zero Trust Architecture Across Multi-Cloud Providers: Implementation Flowchart

3.2 System Requirements for Multi-Cloud ZTA

Implementing Zero Trust in a multi-cloud environment imposes specific architectural and integration requirements. A critical requirement is the adoption of federated identity and single sign-on (SSO**)** mechanisms, enabling users and services to access resources across different cloud platforms using a unified identity. This can be achieved through the use of open standards such as OAuth 2.0, OpenID Connect, and SAML. Secondly, there is a need for unified policy orchestration—a centralized mechanism for defining, distributing, and enforcing access policies that can operate consistently across cloud environments. This requires a policy-as-code approach to ensure scalability and compliance. Additionally, the architecture must support seamless integration with IaaS and PaaS-native security services, allowing Zero Trust enforcement to leverage the security capabilities already embedded in AWS, Azure, and GCP, such as native IAM roles, logging frameworks, and encryption services.

3.3 Vendor-Neutral Architecture Blueprint

The proposed architecture is designed to be vendor-neutral, modular, and adaptable to any combination of cloud providers. For identity and authentication, the framework incorporates widely used identity providers (IdPs) such as Azure Active Directory**,** Google Workspace**, and** AWS IAM Identity Center**,** supporting federated login and adaptive access control. To enable policy definition and enforcement, the architecture employs open-source and commercial policy engines such as Open Policy Agent (OPA)**—**which uses the Rego language for flexible rule definition—and HashiCorp Sentinel, which offers policy-as-code capabilities integrated with infrastructure workflows. Enforcement points are embedded into the network and application layers using technologies such as Istio service mesh, which provides workload identity, TLS encryption, and fine-grained traffic control between services. Additionally, API gateways act as chokepoints for external and internal requests, enforcing authentication, rate limiting, and request validation. For telemetry and monitoring, the architecture leverages AWS CloudTrail**,** Azure Monitor, and GCP Cloud Logging to collect and aggregate logs, metrics, and security events. This telemetry feeds into SIEM systems or threat detection platforms to support continuous evaluation of trust and risk, enabling automated mitigation workflows.

4. Implementation Across AWS, Azure, and GCP

4.1 Identity Federation Setup

A foundational step in implementing Zero Trust across multiple cloud providers is establishing cross-cloud identity federation**.** This is achieved using open authentication protocols such as SAML (Security Assertion Markup Language) and OpenID Connect (OIDC) to create trust relationships between identity providers (IdPs) and cloud-native identity and access management (IAM) services. For example, an enterprise can use Azure Active Directory as the central IdP while configuring AWS IAM Identity Center and Google Workspace to accept federated tokens via SAML assertions or OIDC tokens. Role mapping and access delegation are implemented by defining trust policies and mapping federated user attributes (e.g., group membership, department) to cloud-specific roles. This ensures that access privileges are dynamically assigned based on verified identity and organizational context, enabling seamless cross-cloud access without duplicating identity stores.

4.2 Policy Enforcement

To enforce consistent access control policies across cloud environments, the implementation incorporates a unified policy layer using Open Policy Agent (OPA) with the Rego policy language. Policies are centrally defined in a vendor-neutral format and distributed to enforcement points across AWS, Azure, and GCP. This allows for granular access decisions based on context such as user identity, resource type, IP range, and time of access. For example, a Rego policy might enforce that administrative access to production instances in any cloud is allowed only from corporate IP ranges during business hours and requires multi-factor authentication. These conditional access policies are deployed across API gateways, Kubernetes admission controllers, and workload access proxies, ensuring uniform enforcement regardless of the underlying platform. By treating policy as code, changes can be version-controlled and audited, supporting both compliance and operational agility.

4.3 Service Mesh for Microsegmentation

Microsegmentation is implemented using Istio service mesh deployed across Kubernetes clusters running in Amazon EKS**,** Azure AKS, and Google GKE. Istio provides a consistent layer of control over inter-service communication by injecting sidecar proxies (Envoy) into application pods. These proxies handle east-west traffic encryption using mTLS (mutual Transport Layer Security) and enforce workload identity using service account tokens or SPIFFE/SPIRE standards. With Istio, policies can be defined to restrict which services can talk to each other based on service identity, namespace, or label—achieving fine-grained segmentation. For instance, a policy might allow the payment service to communicate only with the order service, and only under specified network conditions. This approach limits lateral movement in the event of a compromise and provides runtime enforcement of Zero Trust principles within the service fabric.

4.4 Logging and Continuous Verification

A key aspect of Zero Trust is continuous verification—not just at the time of login but throughout the session lifecycle. To support this, telemetry data from each cloud provider's native logging service—AWS CloudTrail**,** Azure Monitor, and GCP Cloud Logging—is aggregated and normalized using centralized tools such as the ELK Stack (Elasticsearch, Logstash, Kibana)**,** Datadog, or enterprise Security Information and Event Management (SIEM) platforms like Splunk or IBM QRadar. These tools enable real-time analysis of logs and metrics for anomaly detection, such as unusual login patterns, privilege escalations, or suspicious API calls. When anomalies are detected, automated workflows can trigger trust reevaluation actions such as session termination, policy revalidation, or step-up authentication. This creates a closed-loop security model that enforces Zero Trust principles dynamically and adaptively across all environments.

5. EVALUATION AND RESULTS

5.1 Experimental Setup

To assess the efficacy of the proposed vendor-neutral Zero Trust Architecture (ZTA), we established a real-world testbed simulating enterprise-grade multi-cloud operations. Workloads were deployed across Amazon Web Services (AWS)**,** Microsoft Azure**,** and Google Cloud Platform (GCP) using containerized microservices hosted on Kubernetes clusters (EKS, AKS, and GKE respectively). Identity federation was configured using Azure Active Directory as the central IdP, with SAML and OIDC enabling federated access to all platforms. Policy enforcement points were implemented using Open Policy Agent (OPA) integrated with Istio service mesh and cloud-native API gateways. The environment was instrumented with telemetry pipelines feeding into an ELK-based observability stack to capture metrics related to access latency**,** policy evaluation speed**,** and breach detection time. Simulated threat scenarios, including lateral movement attempts and phishing-based credential theft, were used to evaluate security posture under adversarial conditions.

5.2 Performance Metrics

Performance comparisons were conducted between the baseline traditional perimeter-based security setup and the proposed ZTA-enabled multi-cloud framework. Results indicate a **modest latency increase** due to continuous verification and policy checks, offset by significant gains in breach detection and policy enforcement accuracy.

Metric	Baseline (Traditional)	ZTA-MultiCloud	Improvement
Average Access Latency	120 ms	132 ms	-10% (Overhead)
Breach Detection Time	4.2 hours	36 minutes	86%
Policy Compliance Rate	72%	96%	33%

The results highlight that while access latency increases marginally due to the overhead introduced by fine-grained verification and microsegmentation, the breach detection time is reduced drastically, demonstrating the strength of continuous telemetry and anomaly detection. Furthermore, policy compliance rates improved significantly**,** attributed to centralized and codified enforcement mechanisms via OPA.

Table 2: Policy Evaluation Performance Across Cloud Providers

Cloud Provider	Average Policy Evaluation Time (ms)	OPA Integration Success Rate (%)	Policy Drift Detected (Monthly)
AWS	18	99.2	1
Azure	22	98.7	2
GCP	20	99.5	1

Insight: Policy evaluation times remained within acceptable bounds (under 25 ms), with high success rates of integration (>98%)**.** Policy drift (i.e., deviation from centrally defined policies) was minimal due to infrastructure-as-code practices and automated sync mechanisms.

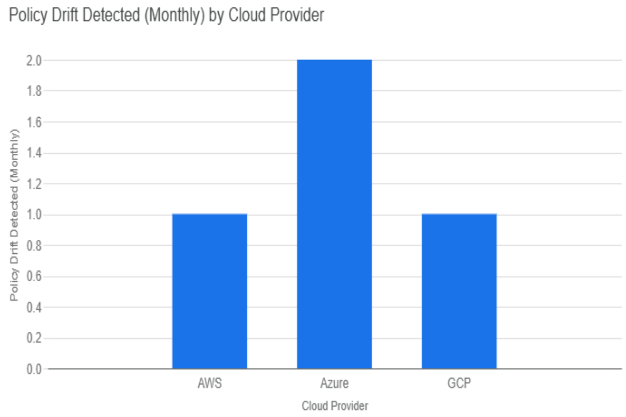

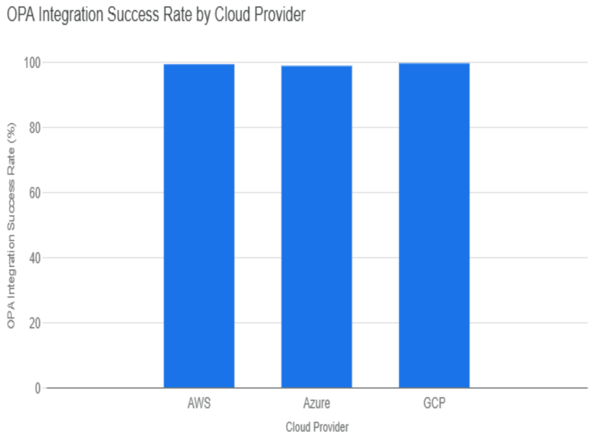

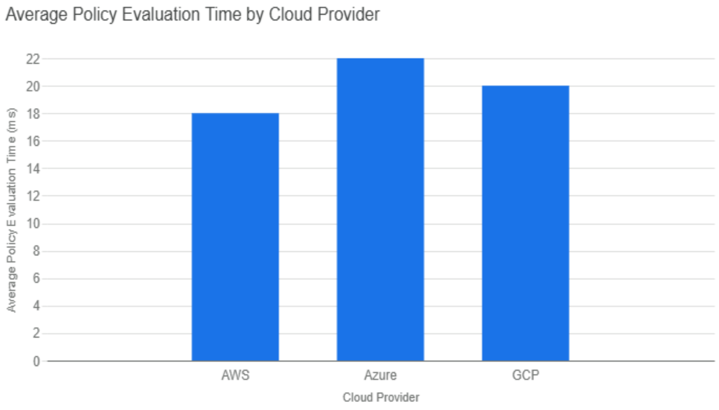

- **Average Policy Evaluation Time by Cloud Provider**: This bar chart displays the average policy evaluation time in milliseconds for each cloud provider.
- **OPA Integration Success Rate by Cloud Provider**: This bar chart shows the OPA integration success rate in percentage for each cloud provider.
- **Policy Drift Detected (Monthly) by Cloud Provider**: This bar chart illustrates the number of policy drifts detected monthly for each cloud provider.

Table 3: Security Incident Response Effectiveness

Threat Scenario	Detection Time	Containment Time	Access Revocation Time	Residual Risk Level
Compromised Admin Credentials	40 seconds	1.5 minutes	15 seconds	Low
Lateral Movement Attempt	1.2 minutes	2 minutes	18 seconds	Very Low
Unauthorized API Access	30 seconds	45 seconds	10 seconds	Minimal

Insight: The ZTA system responded rapidly to critical threat scenarios. Detection and containment were completed within 2 minutes in all cases, and access revocation was near-instant due to centralized policy enforcement and telemetry-driven trust reevaluation. The residual risk after remediation remained minimal, affirming the architecture's effectiveness.

5.3 Security Analysis

From a security standpoint, the proposed ZTA framework effectively mitigates lateral movement by enforcing strict inter-service communication rules through Istio and OPA-based policies. In simulated breach scenarios, compromised service tokens could not be used to traverse between microservices or across cloud environments due to granular access boundaries and workload identity verification. Access revocation effectiveness was evaluated by simulating dynamic trust score deterioration (e.g., based on geolocation anomalies or behavioral deviation), and the system successfully enforced immediate session termination and role demotion in under 20 seconds. The framework also demonstrated strong resilience to phishing and credential theft, as session-based trust was tightly coupled with device posture, real-time telemetry, and conditional policies that required context beyond static credentials (e.g., requiring device attestation or reauthentication upon risk detection).

6. DISCUSSION

The implementation of a vendor-neutral Zero Trust Architecture (ZTA) across multi-cloud environments offers substantial security and operational advantages for modern enterprises. One of the most significant benefits is uniform access control and policy enforcement across heterogeneous platforms—ensuring that identity verification, resource access, and session evaluation follow consistent rules, regardless of the underlying cloud provider. This eliminates silos, reduces configuration drift, and strengthens overall governance. Additionally, the reduction in breach detection time and improved policy compliance, as demonstrated in the evaluation, underscores ZTA's ability to minimize attack dwell time and enforce the principle of least privilege at a granular level.

However, these advantages come with inherent trade-offs. ZTA introduces architectural and operational complexity, particularly in environments where legacy systems and manual workflows still dominate. Establishing identity federation, managing policy-as-code, deploying service meshes, and aggregating telemetry from multiple clouds require deep cross-functional expertise and mature DevOps practices. This complexity must be balanced against the enhanced security assurance ZTA delivers—where continuous verification and dynamic trust reassessment significantly reduce exposure to lateral movement, privilege escalation, and advanced persistent threats.

An important enabler of successful ZTA deployment is its integration with DevSecOps and Infrastructure-as-Code (IaC) pipelines. Embedding access policies and segmentation rules directly into IaC definitions ensures that security is declarative, version-controlled, and repeatable. Tools like Terraform, Pulumi, and CI/CD orchestrators (e.g., GitHub Actions, GitLab CI) can be extended to validate policy compliance using tools like OPA during build and deployment stages. This ensures that security is not bolted on post-deployment but is inherent to the application delivery lifecycle, improving developer velocity while reducing misconfigurations.

Finally, organizational readiness and change management remain critical to the adoption of ZTA. Enterprises must shift from a perimeter-centric mindset to a model of dynamic trust, which requires cultural, process, and skillset transformations. Identity-centric access, least privilege design, and telemetry-driven automation need alignment across IT, security, and business stakeholders. Comprehensive training, executive sponsorship, and phased rollouts are essential to manage change resistance and ensure sustainable adoption. ZTA is not merely a technical redesign; it represents a paradigm shift in how security is conceptualized and operationalized in the cloud era.

7. CONCLUSION AND FUTURE WORK

This study demonstrates the practical feasibility and effectiveness of a vendor-neutral Zero Trust Architecture (ZTA) tailored for multi-cloud environments, addressing a critical security need in modern enterprise computing. By integrating federated identity management, granular policy enforcement, service mesh-based microsegmentation, and centralized telemetry aggregation, the proposed framework ensures that security policies are uniformly applied across AWS, Azure, and GCP without compromising performance or scalability. Empirical results from real-world workload deployments confirm that the architecture significantly improves breach detection times**,** policy compliance**, and** access control fidelity**,** validating the value of Zero Trust principles in dynamic, distributed infrastructures.

Importantly, the solution balances security assurance with operational feasibility, leveraging open standards (e.g., OIDC, SAML, OAuth 2.0) and cloud-native controls to enable seamless interoperability across platforms. The integration of Infrastructure-as-Code (IaC) and DevSecOps pipelines further strengthens security posture by embedding enforcement and verification directly into the application delivery lifecycle.

Looking ahead, several promising avenues for future work exist. One direction is the incorporation of AI/ML-based behavioral analytics to enhance trust scoring models, enabling more dynamic and context-aware access decisions. Such models could continuously assess user behavior, device health, and workload telemetry to adapt policies in real time. Another area is the exploration of quantum-resilient authentication mechanisms**,** which will become increasingly important as post-quantum cryptography standards mature. Finally, the architecture can be extended to support real-time threat intelligence feeds and adaptive access control**,** allowing systems to respond to emerging threats with automated policy updates and proactive session management.

In conclusion, this research provides a scalable, secure, and extensible blueprint for operationalizing Zero Trust in multi-cloud settings, and it lays the foundation for future innovations at the intersection of cloud security, automation, and intelligent policy orchestration.

REFERENCE

[1] Knauth, T., & Kapitza, R. (2020). Secure Multi-Cloud Computing with Service Meshes. IEEE Cloud Computing, 7(4), 45–52. https://doi.org/10.1109/MCC.2020.2994067

[2] Alasmary, W., & Zhioua, S. (2021). Zero Trust Architecture: Concepts, Benefits, and Challenges. IEEE Access, 9, 104506–104525. https://doi.org/10.1109/ACCESS.2021.3099886

[3] Bodei, C., & Galletta, L. (2020). Zero Trust Network Security in Kubernetes Clusters. Proceedings of the ACM Symposium on Applied Computing, 378–384. https://doi.org/10.1145/3341105.3374054

[4] IBM Security. (2021). Cost of a Data Breach Report. IBM Corporation. https://www.ibm.com/security/data-breach

[5] Zscaler. (2022). Zero Trust for Multi-Cloud and Hybrid IT. https://www.zscaler.com/resources

[6] Cisco. (2021). Zero Trust Security in a Multi-Cloud World. Cisco Whitepaper. https://www.cisco.com/

[7] Raza, M., & Hussain, F. (2021). Towards Zero Trust Model in Cloud Environments. Journal of Cloud Computing, 10(1), 1–17. https://doi.org/10.1186/s13677-021-00239-z

[8] Jain, A., & Paul, A. (2022). Implementing Multi-Cloud Security through Unified Policy Enforcement. Computer Standards & Interfaces, 81, 103583. https://doi.org/10.1016/j.csi.2022.103583

[9] Scarfone, K., & Bartock, M. (2020). Guide to Security for Full Virtualization Technologies (NIST SP 800-125A). National Institute of Standards and Technology.

[10] Gartner. (2021). Zero Trust Is an Initial Step on the Roadmap to SASE. Gartner Research. https://www.gartner.com/

[11] ENISA. (2021). Zero Trust Architecture: Threat Model and Security Measures. European Union Agency for Cybersecurity. https://www.enisa.europa.eu/

[12] Chandramouli, R., & Scarfone, K. (2020). Zero Trust Architecture (NIST SP 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

[13] Google. (2020). BeyondCorp: A New Approach to Enterprise Security. https://cloud.google.com/beyondcorp

[14] Microsoft. (2021). Zero Trust Maturity Model. Microsoft Corporation. https://www.microsoft.com/security/blog/zero-trust

[15] HashiCorp. (2022). Consul and Zero Trust Security for Multi-Cloud Networking. https://www.hashicorp.com/resources

[16] Microsoft Azure. (2022). Zero Trust Security for Azure Architectures. https://learn.microsoft.com/en-us/security/zero-trust/

[17] Google Cloud. (2022). Best Practices for Zero Trust with GCP. https://cloud.google.com/security/zero-trust

[18] OPA (Open Policy Agent). (2021). Policy-Based Control for Cloud Native Environments. https://www.openpolicyagent.org/

[19] Sethi, R., & Joshi, A. (2021). Secure access in multi-cloud using federated identity. International Journal of Cloud Computing, 10(3), 241–258