



International Journal of Engineering Research and Science & Technology

www.ijerst.org

ISSN : 2319-5991

Vol. 18 No. 2 (2022)



ijerst.editor@gmail.com

editor@ijerst.com

Research Paper

ZERO TRUST NETWORK ARCHITECTURE IMPLEMENTATION IN HYBRID ENTERPRISES: A PRACTICAL EVALUATION

Amina Yusuf

Department of Information Security, Ahmadu Bello University, Nigeria

Received: 07-5-2022

Accepted: 16-6-2022

Published: 25-6-2022

ABSTRACT

Zero Trust Network Architecture (ZTNA) offers a transformative shift from perimeter-based security models to identity- and context-driven enforcement policies. This paper presents a practical evaluation of ZTNA implementation across hybrid enterprises that combine on-premise infrastructure with public cloud services. Using Azure Active Directory, AWS IAM, and open-source tools such as Istio and HashiCorp Boundary, we deploy a ZTNA model that enforces least-privilege access, mutual TLS authentication, and microsegmentation. Our test environment spans 3 cloud VPCs and 2 on-premise data centers connected via VPN and SD-WAN. We assess improvements in lateral movement prevention, authentication reliability, and policy enforcement latency. Results show a 60% reduction in lateral exposure (as measured by attacker path simulations using BloodHound) and a 32% improvement in incident detection times due to contextual access logs and identity-aware proxying. Integration challenges include policy synchronization across disparate platforms, legacy system incompatibility, and increased administrative overhead. We propose a phased implementation roadmap consisting of asset classification, access segmentation, identity federation, and continuous verification. Policy misconfigurations emerged as a primary concern, requiring rigorous testing and simulation. Overall, this study demonstrates that ZTNA is feasible and effective in hybrid enterprises when implemented with strong identity governance and proper tooling. Recommendations and lessons learned are provided to guide security architects through operationalization and performance tuning of Zero Trust frameworks.

1. INTRODUCTION

Enterprise networks are increasingly evolving toward hybrid architectures that combine traditional on-premise infrastructure with cloud services from providers like AWS and Azure. This transition, driven by the need for scalability, operational flexibility, and cost efficiency, has introduced significant complexity in managing secure access. Conventional perimeter-based security

models, which operate under the assumption that internal networks are inherently trusted, are ill-suited to defend against modern cyber threats, particularly those involving credential compromise, insider abuse, or lateral movement within a network.

The Zero Trust security model challenges this outdated assumption by mandating strict identity verification and access controls at every interaction point—

regardless of the network location of the user or asset. By adopting the “never trust, always verify” principle, Zero Trust Network Architecture (ZTNA) enforces continuous evaluation of trust based on identity, device health, location, and behavior. This represents a paradigm shift, especially in hybrid environments where a consistent and scalable enforcement mechanism is required across cloud and on-prem platforms.

The growing threat landscape, combined with compliance mandates such as HIPAA, PCI-DSS, and GDPR, further underscores the importance of adopting Zero Trust principles. Implementing ZTNA, however, is not a purely technical exercise—it involves aligning policies across identity providers, cloud-native access controls, legacy infrastructure, and operational teams. Enterprises must consider not just technology selection, but also organizational readiness, training, and cultural shifts required to enable least-privilege principles across the board.

This paper presents a practical implementation and evaluation of ZTNA in a hybrid enterprise using tools available as of 2022. We integrate Microsoft Azure AD for identity federation, AWS IAM for cloud-native access control, Istio for workload segmentation, and HashiCorp Boundary for session-based access brokering. Performance metrics, architectural design, and integration challenges are documented, along with a phased roadmap for gradual Zero Trust adoption.

2. Related Work

Zero Trust, as a formalized security concept, gained early momentum with Google’s BeyondCorp initiative (2014), which aimed to enable secure enterprise access without reliance on VPNs. This was

followed by Forrester’s Zero Trust eXtended (ZTX) framework, which emphasized seven pillars: users, devices, networks, applications, data, automation, and visibility. These foundational works prompted both academia and industry to reassess network security design.

Academic research has since explored various aspects of Zero Trust, from policy enforcement using SDN to behavioral authentication for dynamic trust scoring. For example, studies have evaluated microsegmentation using service mesh tools such as Istio, as well as continuous authentication using biometric and contextual signals. However, many implementations focus either on cloud-only deployments or hypothetical scenarios that lack real-world hybrid complexity.

Commercial ZTNA offerings such as Zscaler Private Access, Palo Alto Prisma Access, and Cisco Duo provide managed services for endpoint-to-application access. However, these are often black-box systems with limited customizability or integration flexibility, making them unsuitable for hybrid or multicloud deployments that demand fine-tuned control and transparency.

Recent technical evaluations of service meshes like Istio and Linkerd, and access proxies like HashiCorp Boundary or Teleport, have provided valuable insights into performance trade-offs and scalability limitations. However, few studies have combined these tools into a full-fledged ZTNA implementation across multiple cloud and on-prem domains.

This paper fills that gap by providing an end-to-end, tool-agnostic implementation of ZTNA, benchmarked in a realistic hybrid enterprise environment. We contribute not only architectural patterns

and performance benchmarks, but also operational lessons learned that are critical for practical deployment.

3. Methodology and System Design

Our ZTNA implementation is designed around the core principle of least-privilege access, enforced through federated identity, dynamic session controls, and workload-aware network segmentation. The system is deployed across:

- **Three cloud VPCs:** two hosted on AWS and one on Azure, each running microservice-based applications with independent IAM and networking layers.
- **Two on-premise data centers:** connected via site-to-site IPsec VPN and SD-WAN for resilient interconnectivity and policy testing across network layers.

The architecture includes the following components:

- **Identity Management:** Azure AD serves as the central identity provider, federated with AWS IAM for cross-cloud authentication. Conditional access policies are enforced based on device posture, location, and user risk scores.
- **Access Brokering:** HashiCorp Boundary brokers access to sensitive infrastructure and applications using short-lived credentials and session policies, integrating with Azure AD groups.
- **Network Segmentation:** Istio service mesh (v1.11) handles workload-to-workload mTLS, authorization policies, and observability. Sidecars are injected into services in all cloud VPCs.
- **Audit and Monitoring:** All authentication and access logs are aggregated into a centralized ELK

(Elasticsearch, Logstash, Kibana) stack, enriched with AWS CloudTrail and Azure Monitor for complete telemetry.

Design goals include minimizing lateral movement, improving incident detection time, and reducing policy drift across platforms. The system is subjected to red team testing and automated attack simulations using tools like BloodHound and Caldera.

4. ZTNA Components and Tooling

Each component was carefully selected to ensure open standards, integration flexibility, and scalability. Their combined use enabled a comprehensive and transparent Zero Trust framework:

- **Azure Active Directory (AD):** Serves as the core identity federation layer, enabling MFA, user risk scoring, and device compliance policies. Integration with Microsoft Endpoint Manager allows for conditional access based on device health.
- **AWS IAM:** Provides cloud-native role-based access policies for EC2, Lambda, S3, and other AWS services. Trust relationships are established with Azure AD for SAML-based federation and role assumption.
- **Istio Service Mesh:** Used for service-level authentication and policy enforcement. Enables mutual TLS, JWT token validation, and RBAC for east-west traffic within cloud environments. It also provides traffic telemetry and fault injection capabilities.
- **HashiCorp Boundary:** Acts as a dynamic session-based access proxy. Policies define who can access what, when, and for how

long. Sessions are ephemeral and logged, with optional integration into Vault for credential management.

- **ELK + CloudWatch:** Provide layered telemetry. ELK enables deep search and visualization of audit trails, while CloudWatch provides operational metrics and alerting for AWS-hosted resources.

Together, these tools support ZTNA implementation across identity, network, and session dimensions. APIs and policy-as-code templates were used to synchronize configuration across environments.

5. Experimental Setup

To simulate a realistic enterprise environment, we created a hybrid infrastructure composed of five network zones: three cloud-based VPCs (two in AWS, one in Azure) and two physically separated on-premise data centers. Each cloud VPC hosted containerized microservices behind Istio-enabled Kubernetes clusters, while on-prem servers hosted legacy apps and a private GitLab CI/CD pipeline.

Workload simulation involved 50 synthetic user profiles spanning three roles—administrators, developers, and contractors—each assigned differentiated access scopes via Azure AD groups and AWS IAM policies. We created detailed

identity and access control matrices mapping users to applications and services, enabling granular testing of policy enforcement.

Traffic and attack simulations were conducted using:

- **Apache JMeter** for simulating legitimate workloads (e.g., logins, file uploads, API requests).
- **BloodHound** for mapping Active Directory relationships and simulating lateral movement pathways.
- **MITRE Caldera** to launch automated attack chains, including credential reuse, port scans, and privilege escalation.

Experiments ran over 30 days with the following goals:

- Evaluate authentication latency, access denial accuracy, and policy evaluation times.
- Measure lateral movement feasibility pre- and post-ZTNA deployment.
- Quantify log generation volume and effectiveness in incident detection.
- Stress test policy engines under concurrent user load.

All sessions were logged and enriched with contextual metadata such as device state, geolocation, and behavioral baselines, stored in ELK for analysis.


6. EVALUATION METRICS AND RESULTS

ZTNA implementation led to notable improvements across key security and operational metrics:

Metric	Pre-ZTNA	Post-ZTNA	Delta
Avg. Lateral Movement Path (BloodHound)	3.7 hops	1.5 hops	↓ 60%
Unauthorized Session Block Rate	N/A	100%	Secure enforcement
Incident Detection Time	5.2 min	3.5 min	↓ 32%
Session Initiation Latency	220 ms	263 ms	↑ 19%
False Positives (Policy Violations)	N/A	5/1000	Acceptable margin
Log Events per Day	~9,000	~37,000	↑ 4.1x

The reduction in lateral exposure was confirmed through repeated attack graph generation and red team validation. While there was a slight increase in user authentication latency due to access proxying and mTLS handshakes, the trade-off was deemed acceptable given the security gains.

Figure 1: Effect of ZTNA on Security and Performance Metrics



The centralized log aggregation also proved critical in reducing detection times. Analysts were able to correlate abnormal behaviors across identity, session, and network layers with greater precision using Kibana dashboards.

7. CHALLENGES AND MITIGATION STRATEGIES

Implementing ZTNA in a hybrid environment presented multiple integration and operational challenges:

- **Policy Sprawl:** IAM policies differed across Azure and AWS,

causing drift. We mitigated this with policy-as-code using Terraform and Open Policy Agent (OPA) templates.

- **Legacy Compatibility:** Some internal services lacked support for federated SSO or modern TLS. These were wrapped with reverse proxies or migrated to containerized equivalents.
- **Tooling Fragmentation:** Managing logs, metrics, and policies across Istio, Boundary, and IAM increased overhead. This was addressed by adopting unified dashboards and automated policy testing pipelines.
- **Initial Latency and Complexity:** The introduction of proxies and identity federation caused delays in early phases. Gradual rollout with A/B testing and optimization of TLS handshakes helped reduce performance impact.

Security training for IT staff was also necessary, especially around conditional access, RBAC modeling, and interpreting new types of audit logs. We observed that early stakeholder buy-in significantly accelerated adoption and troubleshooting.

8. PHASED IMPLEMENTATION ROADMAP

To ensure scalable and low-risk deployment, we developed a four-phase ZTNA adoption strategy:

Phase 1: Asset Discovery and Classification

- Conducted an inventory of all users, devices, services, and data flows.
- Tagged assets with sensitivity and regulatory labels (e.g., HIPAA, internal-only).

Phase 2: Access Segmentation and Isolation

- Introduced Istio service mesh for internal traffic.
- Created logical access zones with minimal trust across boundaries (e.g., dev ↔ staging ↔ prod).

Phase 3: Identity Federation and Policy Unification

- Federated Azure AD with AWS IAM and on-prem AD.
- Implemented role-based policies and device health checks for all users.

Phase 4: Continuous Verification and Monitoring

- Enforced session-based access via HashiCorp Boundary.
- Enabled anomaly detection through ELK alerts and automated policy re-evaluation.

Each phase included rollback plans and validation through penetration testing and compliance checks. The roadmap proved adaptable to enterprises at different maturity levels, offering a modular path to full ZTNA deployment.

9. CONCLUSION

This paper presented a real-world implementation and evaluation of Zero Trust Network Architecture in a hybrid

enterprise setting using publicly available tools as of 2022. By integrating identity federation, dynamic access brokering, and service mesh-based segmentation, we were able to achieve measurable improvements in security posture, incident detection, and access governance.

Key findings include:

- A 60% reduction in lateral movement paths across hybrid assets
- Complete mitigation of unauthorized access attempts through session policies
- Improved observability and telemetry volume with minimal impact on performance

Despite challenges with legacy system integration and policy alignment, the phased approach and tool integration strategies we outlined enabled smooth ZTNA operationalization. This research provides a replicable blueprint for enterprises transitioning toward Zero Trust, with actionable insights and performance benchmarks.

Future work will explore integration with behavioral analytics engines, zero-touch device onboarding, and real-time risk scoring to further automate and enhance trust decisions.

REFERENCES

1. Chandramouli, R., & Scarfone, K. (2020). *Zero Trust Architecture (SP 800-207)*. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207>
2. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture*. NIST Special Publication 800-207. <https://csrc.nist.gov/publications/detail/sp/800-207/final>

3. Kindervag, J. (2010). *Build Security Into Your Network's DNA: The Zero Trust Network Architecture*. Forrester Research.
4. Google. (2014). *BeyondCorp: A New Approach to Enterprise Security*.
<https://research.google/pubs/archive/43231.pdf>
5. Microsoft. (2021). *Zero Trust Deployment Guide*. Microsoft Corporation.
<https://www.microsoft.com/security/blog/zero-trust>
6. Amazon Web Services. (2022). *Security Pillar – AWS Well-Architected Framework*.
<https://docs.aws.amazon.com/wellarchitected/latest/security-pillar>
7. HashiCorp. (2021). *HashiCorp Boundary: Identity-Based Access Management*.
<https://www.hashicorp.com/products/boundary>
8. IBM. (2020). *Zero Trust Security: A Guide to Risk-Based Protection*.
<https://www.ibm.com/security/zero-trust>
9. Bellamkonda, Srikanth. "The Role of Network Engineers in Securing Cloud-based Applications and Data Storage."
10. Cisco Systems. (2021). *Zero Trust Security Framework*.
<https://www.cisco.com/c/en/us/products/security/zero-trust/index.html>
11. Microsoft Azure. (2021). *Azure Active Directory Conditional Access Documentation*.
<https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/>
12. Red Hat. (2021). *Service Mesh With Istio: Architecture, Patterns, and Best Practices*.
<https://www.redhat.com/en/topics/service-mesh/what-is-istio>
13. AWS. (2021). *IAM Identity Federation*.
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_provider_s.html
14. Suci, G., Vulpe, A., & Dobre, C. (2020). Implementation of Zero Trust Architecture in Hybrid Cloud Environments. *Informatica Economică*, 24(2), 21–30.
<https://doi.org/10.24818/issn14531305/24.2.2020.02>
15. Newman, M. (2020). *Istio: Up and Running: Using a Service Mesh to Connect, Secure, Control, and Observe*. O'Reilly Media.
16. Fernandes, E., Rahmati, A., & Prakash, A. (2020). Security Implications of Converged IT/OT Networks in Hybrid Clouds. *IEEE Cloud Computing*, 7(3), 70–78.
<https://doi.org/10.1109/MCC.2020.2988481>