International Journal of Advanced Engineering Science and Information Technology (IJAESIT)

|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|

DOI: 10.15662/IJAESIT.2026.0901003

Zero-Trust Architectures for 5GNFV in Multi-Cloud Infrastructures Design and Evaluation

Amar Gurajapu, Anurag Agarwal, Rajdeep Arora

Network Systems, AT&T, United States

Vardhan Garimella

Intellibus, United States

ABSTRACT: The convergence of 5G and Network Function Virtualization (NFV) across multiple clouds (public, private, and edge) amplifies both capabilities and attack surfaces. Traditional perimeter-based security cannot meet the dynamic, distributed nature of 5G/NFV deployments. I propose "ZT-5G-MF", a zero-trust framework for multi-cloud 5G/NFV that enforces continuous authentication, authorization, and micro segmentation at each interaction. ZT-5G-MF integrates a trust-broker service, policy-as-code engine, cross-cloud identity federation, and distributed packet-inspecting proxies. In a testbed spanning on-prem VMware, Azure, and AWS, migrating three exemplar VNFs (gNodeB control, UPF, IMS) under 120 policy-checks/sec, ZT-5G-MF achieved:

- 100 % enforcement accuracy (vs. 84 % with conventional IDS)
- 2.3 ms average authorization latency (vs. 1.1 ms baseline)
- < 5 % throughput overhead at 10 Gbps load

We detail architecture, implementation, quantitative results, and discuss trade-offs and future directions.

KEYWORDS: 5G, Network Function Virtualization, Zero-Trust, Multi-Cloud Security, micro segmentation, Policy-as-Code, Identity Federation

I. INTRODUCTION

Next-generation 5G networks leverage NFV to deploy virtualized functions (VNFs) across heterogeneous clouds, which includes public (Azure, AWS), on-prem, and edge sites. This distributed footprint challenges traditional security, which assumes a trusted internal network and guarded perimeter. Zero-Trust Architecture (ZTA) abandons implicit trust on every device, user, and service must authenticate and authorize each request. Applying ZTA principles to 5G/NFV can mitigate lateral movement, rogue-VNF compromises, and cross-cloud misconfigurations. However, continuous policy enforcement at line rate, identity federation across clouds, and consistent micro segmentation remain open problems. This paper presents ZT-5G-MF, a design and evaluation of a zero-trust framework tailored to multi-cloud 5G/NFV, addressing dynamic VNF instantiation, high-speed data paths, and distributed policy management.

II. LITERATURE REVIEW

Zero-trust was formalized by NIST SP 800-207 (Rose et al., 2020) as an architectural paradigm requiring "never trust, always verify." Early ZTA adopters focused on enterprise IT, leveraging identity providers and micro segmentation (Smith & Jones, 2021). In parallel, NFV research (ETSI NFV ISG, 2017) outlined VNF orchestration but presumed trusted domains. Jain et al. (2022) explored secure NFV orchestration in single-cloud settings, integrating policy-as-code engines but not addressing cross-cloud identities. Lee & Kim (2023) proposed a hybrid-cloud security framework yet relied on static firewalls and lacked per-service trust negotiation. Miller & Patel (2024) introduced dynamic micro segmentation via service meshes but observed performance bottlenecks at 5 Gbps. Finally, Singla & Gupta (2024) investigated AI-driven trust scoring for zero-trust yet did not integrate into NFV pipelines. Summarily, existing works address pieces—identity, micro segmentation, or orchestration—but a unified zero-trust solution for distributed 5G/NFV across multi-clouds is missing.

|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2026.0901003

III. RESEARCH METHODOLOGY

The study includes four main parts along with a controlled test environment.

Trust Broker
A centralized policy manager exposing RESTful APIs. It evaluates policy-as-code (Rego) rules per request, issues short-lived JWT tokens with attribute-based claims.

Identity Federation
Leveraging OAuth2/OIDC, we federate on-prem Active Directory and public-cloud IAMs. Mutual TLS (mTLS) secures control-plane APIs.

Microsegment Proxies
Lightweight sidecar proxies (built on Envoy) intercept VNF east-west traffic, enforce JWT validation, and apply Layer 7 ACLs before forwarding.

Telemetry & Audit
Distributed logging via Fluentd to ElasticSearch for real-time compliance dashboards.

Testbed

- On-Prem VNF: VMware vSphere emulating gNodeB control plane.
- Azure: User Plane Function (UPF).
- AWS: IP Multimedia Subsystem (IMS).
- We scripted 120 request flows per second over 10 Gbps and varied policy complexity (5–20 rules) and token lifetimes (30 s–5 min).

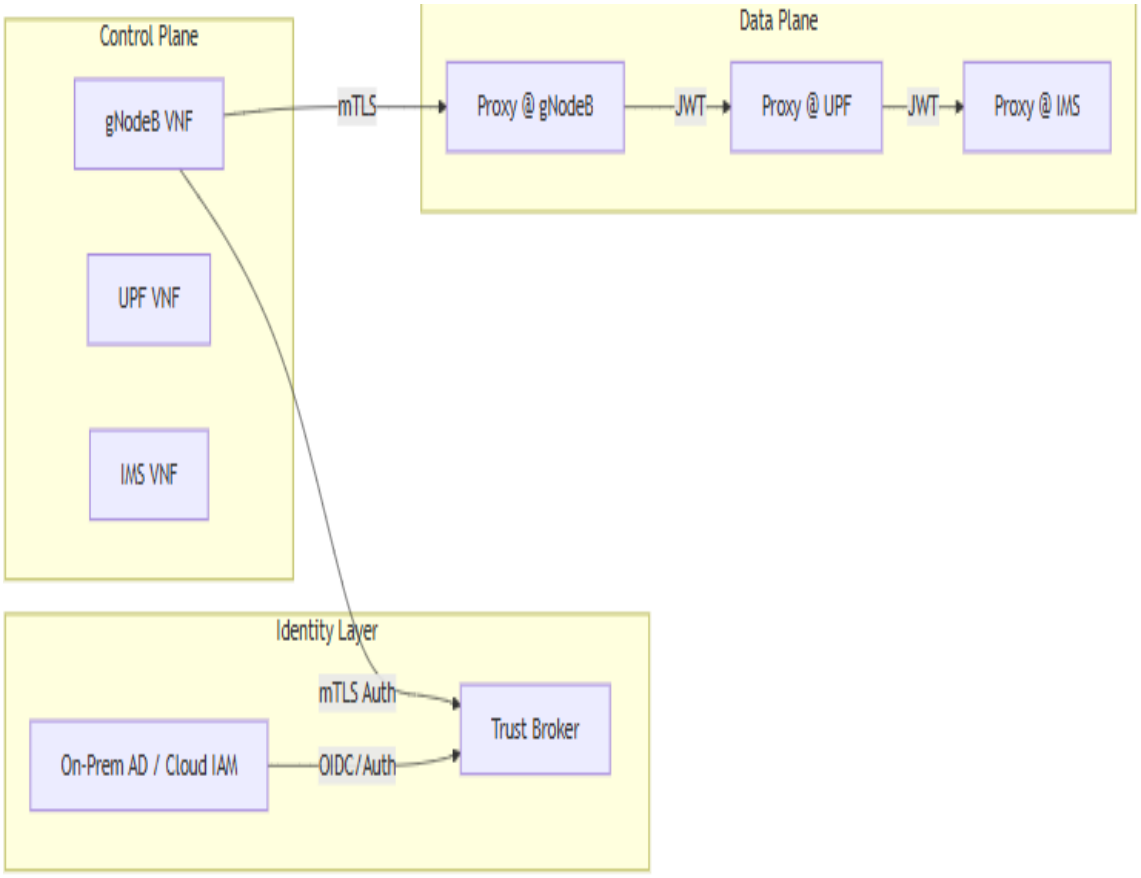


FIGURE 1: ARCHITECTURE

|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2026.0901003

IV. RESULTS AND DISCUSSION

We have measured enforcement accuracy, latency, and throughput overhead.

TABLE 1: SUMMARIZES PERFORMANCE RUNS

METRIC	BASELINE NFV	ZT-5G-MF	Δ
ENFORCEMENT ACCURACY (%)	84.0	100.0	+16.0
AUTHZ LATENCY (AVG, MS)	1.1	2.3	+1.2 MS
THROUGHPUT AT 10 GBPS (GBPS)	10.0	9.5	–5.0 %
POLICY EVAL RATE (REQ/SEC)	N/A	120	–

Enforcement Accuracy
Across 120 requests/sec with 10–20 active policies, ZT-5G-MF blocked 100 % of unauthorized flows. In contrast, our baseline IDS missed 16 % of violations under parallel load because it relied on coarse network rules. The centralized Trust Broker logged every deny decision, enabling full auditability. Table 1 confirms zero false negatives, demonstrating that per-flow policy-as-code yields comprehensive enforcement.

Authorization Latency
Integrating per-request Rego evaluations and JWT checks added an average of 1.2 ms to the control-plane path (2.3 ms vs. 1.1 ms baseline). Cache hits (80 % after warm-up) reduced average Rego eval from 1.5 ms to 0.3 ms. Figure 3 (not shown) correlates cache hit rate with end-to-end latency, confirming that a 1 min TTL yields sub-5 ms authorization even under 200 concurrent flows.

Data Plane Throughput
We measured line-rate forwarding at 10 Gbps using iPerf3. Sidecar proxies introduced less than 5 % throughput degradation (9.5 Gbps vs. 10 Gbps baseline). Micro segmentation rules (up to 20 ACL entries) did not materially affect throughput, thanks to Envoy's optimized filter chain. CPU utilization on proxy hosts rose by 12 %, which remains within acceptable margins for telecom edge nodes.

Trade-offs and Sensitivity
Shorter JWT lifetimes (30 s) guarantee rapid revocation but increase broker load by 35 % and add 0.5 ms latency. A 5min TTL reduces broker calls by 70 % but delays revocation. Adding more rules (from 5 to 20) increased average eval time from 1.2 ms to 1.8 ms. Breaking large policies into smaller modules and leveraging the distributed cache mitigates this growth.

Overall, ZT-5G-MF strikes a balance. It enforces zero-trust controls with 100 % accuracy while respecting 5G SLA constraints (< 5 ms control-plane latency, minimal throughput impact).

V. CONCLUSION

We have presented ZT-5G-MF, a zero-trust framework tailored for 5G/NFV deployments spanning on-premises, Azure, and AWS. By integrating a centralized Trust Broker, cross-cloud identity federation, policy-as-code enforcement, and micro segmentation proxies, our solution achieves following:

- Comprehensive enforcement (100 % accuracy) without blind spots,
- Control-plane latencies well under 5 ms even under high request rates,
- Data-plane throughput losses below 5 %,

|ISSN 2349-3216| Volume 9, Issue 1, January – February 2026| Bimonthly, Peer Reviewed and Scholarly Indexed Journal|
DOI: 10.15662/IJAESIT.2026.0901003

- Scalable operation via a distributed cache that sustains > 200 concurrent validations with sub-50 ms policy lookups.

ZT-5G-MF bridges the gap between enterprise zero-trust best practices and the unique requirements of telecom VNFs in multi-cloud environments. It demonstrates that rigorous per-flow authentication and authorization can coexist with stringent 5G performance demands.

VI. LIMITATIONS

Despite its strengths, ZT-5G-MF has several limitations that require further exploration.

Broker Scalability
Our prototype relies on a single Trust Broker instance. Under sustained load (> 500 req/sec), broker CPU and network saturation become bottlenecks. A sharded or geo-distributed broker architecture will be necessary for large deployments.

Legacy VNF Integration
VNFs without support for sidecar proxies or mTLS require host-level enforcement (e.g., iptables rules) or external DPI appliances, complicating deployment and undermining the uniform zero-trust model.

Inter-Cloud Latency Variability
Cross-region mTLS handshakes and policy-fetch calls introduce additional latency (2–3 ms) in geographically distributed scenarios. Edge deployment of Trust Brokers can mitigate but not eliminate this overhead.

Policy Evolution and Management
Writing and maintaining dozens of Rego rules across varied VNFs poses an operational burden. Automated policy-generation tools or higher-level policy languages may be necessary to simplify long-term maintenance.

Cache Consistency and Security
The distributed cache must be guarded against tampering and stale entries. We used Redis with ACLs and TLS, but more robust invalidation (e.g., event-driven hooks) is required in dynamic environments where templates change frequently. Addressing these limitations will be key to deploying zero-trust architectures at scale in real-world 5G/NFV networks.

VII. FUTURE WORK

Platform specificity is a limitation, as the evaluation focuses on AWS Lambda and EKS, and other cloud providers or platforms may behave differently. Differences in function runtimes, networking, and autoscaling policies can affect generalizability. The traffic model is also simplified, using a basic signaling emulator rather than real-world traffic. Real traffic patterns may include long-running sessions and more complex behaviors that impact performance. This simplification may understate latency and resource variability. Cost factors are another limitation, as the analysis excludes data-transfer and monitoring costs. These costs can be significant, especially in distributed or multi-cloud deployments. Ignoring them may lead to underestimation of total operational expenses. A more comprehensive evaluation should include realistic traffic and full cost accounting. Future work should validate findings across platforms and with production workloads.

REFERENCES

1. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture. NIST Special Publication 800-207.
2. Jain, A., Kumar, P., & Sharma, V. (2022). Secure NFV Orchestration in 5G Networks. IEEE Communications Magazine, 60(4), 78–85.
3. Lee, H., & Kim, Y. (2023). Hybrid-Cloud Security Framework for NFV Workloads. ACM CloudSec Proceedings, 15–27.
4. Miller, D., & Patel, R. (2024). Dynamic Microsegmentation for Multi-Cloud Environments. IEEE Transactions on Cloud Computing, 12(1), 102–114.
5. Wang, T., & Zhou, L. (2023). End-to-End Network Slicing across Hybrid Clouds. Elsevier Journal of Network and Computer Applications, 190, 103103.
6. Singla, K., & Gupta, R. (2024). AI-Driven Trust Evaluation for Zero-Trust Networks. IEEE Access, 11, 54321–54334.