




International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: <https://www.svedbergopen.com/>



Research Paper

Open Access

Zero Trust Security Framework for ePDG in 5G Untrusted Wi-Fi Access Networks

Saiprasad Charudatta Shrikhande

Independent Researcher, India. E-mail: saiprasadc1@gmail.com

Abstract

The integration of untrusted Wi-Fi networks into 5G architecture introduces significant security challenges because the access layer is, by definition, beyond the operator's trust boundary. The Evolved Packet Data Gateway (ePDG) provides secure connectivity from Wi-Fi to the 5G core using IKEv2/IPSec tunnels and EAP-AKA authentication, but its trust model remains essentially perimeter-based: once a device completes EAP-AKA, it is implicitly trusted for the lifetime of the session. Modern attack vectors — rogue access points, man-in-the-middle downgrades, replayed EAP-AKA nonces, credential theft, session hijacking, and device spoofing — are increasingly able to defeat that one-time, static model.

This paper proposes a Zero Trust Security Framework for ePDG that replaces perimeter trust with continuous verification, dynamic policy enforcement, and context-aware authentication. The framework introduces five concentric layers — Identity & Trust, Behavioral Analytics, Policy Decision Point (PDP), Micro-segmentation Policy Enforcement Point (PEP), and a Crypto Plane — wrapped around the ePDG core. A composite risk model $R = \alpha(1 - U) + \beta(1 - D) + \gamma N + \delta B$ combines user trust, device posture, network condition, and behavioral anomaly to drive per-session decisions in real time. Re-authentication, step-up MFA, and micro-segmentation isolation are triggered when the score crosses configurable thresholds. Across a six-attack benchmark (rogue AP, MITM, replay, credential theft, session hijack, device spoofing), the proposed framework lifts the mean detection rate from 46% (traditional ePDG) to 94% while adding only ≈ 34 ms of mean authentication latency. The design is standards-compliant (3GPP TS 23.402, NIST SP 800-207) and incrementally deployable on existing ePDG infrastructure.

Keywords— ePDG, Zero Trust Architecture, NIST SP 800-207, 5G non-3GPP access, untrusted Wi-Fi, IPSec, IKEv2, EAP-AKA, continuous authentication, micro-segmentation, behavioral analytics, dynamic risk scoring, policy enforcement point, 3GPP TS 23.402.

I. Introduction

The rapid growth of mobile data and the operational economics of Wi-Fi offload have made non-3GPP access an essential building block of 5G. In this architecture, the Evolved Packet Data Gateway (ePDG) is the secure anchor between an untrusted Wi-Fi network and the 5G core, terminating IKEv2/IPSec tunnels on the SWu interface and authenticating subscribers via EAP-AKA over SWm with the AAA/HSS [1-5]. The current security model rests on three implicit assumptions: that there is a stable trust boundary (the IPSec tunnel), that authentication is a one-time event (EAP-AKA at attach), and that a single point of verification — the AAA — is sufficient.

These assumptions no longer hold. Untrusted access networks are populated with rogue access points that impersonate operator SSIDs, with MITM proxies that attempt to downgrade IKEv2 ciphers, and with high-mobility devices whose identity context can change minute by minute. The 2020 publication of NIST SP 800-207 [2, 6-10] formalized an alternative model — Zero Trust Architecture (ZTA) — built on the maxim “never trust, always verify.” In a Zero Trust model, every access decision is treated as if the request originated from an open network: identity, device posture, environmental context, and behavior are continuously re-evaluated, and authorization is granular, time-bounded, and revocable [3, 11-13].

This paper makes three contributions. First, it formalizes the security gap in current ePDG deployments by mapping six concrete attack vectors against the implicit assumptions of perimeter trust. Second, it proposes a Zero Trust Security Framework specifically tailored to ePDG, organized as five concentric protection layers around the ePDG core (Identity, Analytics, Policy, Micro-segmentation, and Crypto Plane), with explicit Policy Decision Point (PDP) and Policy Enforcement Point (PEP) roles. Third, it presents a dynamic risk model and a

control-plane / data-plane separation that together enable continuous, per-session trust evaluation while remaining compatible with existing 3GPP reference points (SWu, SWm, S2b).




Fig. 1. Zero Trust layered security architecture — five concentric defence layers wrapped around the ePDG secure core, deflecting untrusted Wi-Fi threats at multiple, independent control points.

2. Background and System Model

2.1 The ePDG in 5G Non-3GPP Access

Defined in 3GPP TS 23.402 [1], the ePDG provides secure connectivity for User Equipment (UE) attaching to the 5G core via untrusted, non-3GPP access — typically public, enterprise, or carrier-managed Wi-Fi. The UE establishes an IKEv2/IPSec tunnel to the ePDG over the SWu interface, mutually authenticates with the AAA via EAP-AKA on SWm (HSS-backed via SWx), and the ePDG forwards subscriber traffic to the UPF (S/PGW-U) over the S2b interface using GTP-U. The control-plane counterpart S2b-c uses GTPv2-C per TS 29.274 [3, 14-16].

2.2 Limitations of the Current Security Model

Although IPSec and EAP-AKA are individually strong primitives, the way they are composed in current ePDG deployments creates three structural weaknesses. First, authentication is a one-shot event at attach: there is no requirement for continuous re-verification of the UE's identity or device posture during the session. Second, the trust boundary is implicit and binary — once the IPSec tunnel is up, all subsequent flows inside it are treated as trustworthy. Third, the access network itself is not attested: a rogue Wi-Fi AP can present a valid-looking SSID and still relay legitimate IKEv2 traffic, while passively logging or actively manipulating side channels. These limitations are the technical roots of the threat landscape examined next.

2.3 Threat Model

Six attack classes drive the design of the proposed framework. They are listed here, and the countermeasure mapping is illustrated in Fig. 3.

- T1 — Rogue Access Point: an adversary stands up a Wi-Fi SSID that impersonates the operator's offload SSID or a partner hotspot and proxies legitimate IKEv2 traffic to harvest metadata or to inject downgrades.
- T2 — Man-in-the-Middle (MITM): an attacker between UE and ePDG attempts to downgrade IKEv2 cipher suites or substitute certificates that the UE may accept by default.
- T3 — Replay Attack: previously captured EAP-AKA messages or IKE_AUTH sequences are replayed to confuse session state on the ePDG.
- T4 — Credential Theft: phished SIM credentials, extracted IMSI/Ki pairs, or stolen long-term keys are used to impersonate a legitimate subscriber.

- T5 — Session Hijack: a legitimate tunnel is taken over after attach, typically by exploiting weak path-binding or by spoofing a roaming event.
- T6 — Device Spoofing: a cloned IMEI or a forged device fingerprint is used to bypass posture checks.

3. Proposed Zero Trust Security Framework

3.1 Design Principles

The framework adopts three NIST SP 800-207 [2] principles and adapts them to the ePDG environment. (i) No implicit trust: every access decision — including those inside an established IPSec tunnel — is re-evaluated against current context. (ii) Continuous verification: identity, device posture, and behavior are re-checked at configurable intervals and on event triggers, not only at attach. (iii) Granular, dynamic authorization: the Policy Decision Point (PDP) issues short-lived, scope-limited authorizations that the Policy Enforcement Point (PEP) — inside the ePDG — enforces per session and per resource.

3.2 Five Concentric Defense Layers

Fig. 1 illustrates the layered structure. L1 — Identity & Trust performs multi-factor authentication, device fingerprinting, and certificate posture verification at the outer boundary. L2 — Threat Intel & Behavioral Analytics continuously scores the session, drawing on operator threat feeds and on per-UE behavioral baselines. L3 — Policy Decision Point implements dynamic, context-aware access control using the risk model described in §3.4. L4 — Micro-segmentation enforces per-session isolation so a compromised tunnel cannot pivot laterally; this is the PEP and lives inside the ePDG. L5 — Crypto Plane is the existing IPSec / IKEv2 / EAP-AKA layer, but hardened with Perfect Forward Secrecy and cipher pinning to defeat downgrade attempts.




Fig. 2. Authentication workflow in ePDG — a six-phase Zero Trust sequence: connect, verify identity, EAP-AKA, evaluate risk, continuous monitoring, and re-authentication.

3.3 Authentication Workflow

The six-phase workflow in Fig. 2 governs every session. Phase 1 establishes the IKEv2/IPSec tunnel on SWu. Phase 2 sends the UE's identity claim to the Identity Engine, which performs MFA and device fingerprinting and returns a user-trust score U and a device-trust score D . Phase 3 runs the standard EAP-AKA exchange via the AAA on SWm. Phase 4 calls the Policy Engine (PDP) with the full context (U, D, N, B), which returns an access decision and a current risk score R . Phase 5 streams session events to the Analytics Engine, which continuously updates the behavioral anomaly score B . Phase 6 triggers re-authentication or session termination whenever R crosses the configured threshold R_{th} .



Fig. 3. Threat model for untrusted Wi-Fi access. Each attack vector T_i (left) is mitigated by a specific Zero Trust countermeasure C_i (right) implemented by the ePDG.

3.4 Dynamic Risk Model

Trust is computed as a continuous, weighted composite of four signals. Let $U \in [0,1]$ denote user trust (from MFA and historical reputation), $D \in [0,1]$ device trust (from posture / attestation), $N \in [0,1]$ network risk (from AP reputation, geofence, and ToD), and $B \in [0,1]$ behavioral anomaly (from the analytics engine). The instantaneous risk score is:

$$R(t) = \alpha(1 - U) + \beta(1 - D) + \gamma N + \delta B, \quad \alpha + \beta + \gamma + \delta = 1$$

Two thresholds drive policy: R_{th} (default 0.55) triggers re-authentication or step-up MFA, and R_{block} (default 0.80) terminates the session and quarantines the device. Weights $\alpha, \beta, \gamma, \delta$ are tuned per operator policy and per service class (e.g., a corporate VoWiFi subscriber may give greater weight to D than a consumer voice subscriber).

3.5 Control Plane / Data Plane Separation

To keep performance acceptable in production, all trust computation runs out-of-band on the control plane while user traffic continues to flow on the IPSec/GTP-U data plane. The control plane consists of Identity Engine, AAA/EAP-AKA, PDP, Risk Engine, Analytics, and Audit/SIEM, connected by a signed Trust & Policy Bus. The PDP, Risk Engine, and Analytics each emit policy and risk signals into the PEP hooks inside the ePDG; the ePDG enforces them on the data plane without breaking the IPSec session. This separation is what makes continuous evaluation feasible without inflating per-packet latency.




Fig. 4. Control plane and data plane separation under Zero Trust. The control plane carries identity, policy, risk, and analytics; the data plane carries encrypted user traffic that the ePDG PEP enforces against policy in real time.

3.6 Algorithm — Per-Session Trust Loop

Algorithm 1: Continuous Trust Evaluation Loop at the ePDG

Input : New SWu/IPSec session s

Context vector $ctx = (U, D, N, B)$

Thresholds R_{th}, R_{block} ; eval interval Δt

Output : Per-session policy decision

1. Establish IKEv2/IPSec(s) // Phase 1
2. $(U, D) \leftarrow \text{IdentityEngine.verify}(s)$ // Phase 2 — MFA + fingerprint
3. AAA.EAP_AKA(s) // Phase 3
4. $R \leftarrow \alpha(1-U) + \beta(1-D) + \gamma N + \delta B$
5. if $R \geq R_{block}$ then Terminate(s); return
6. PDP.grant(s , scope, ttl= Δt) // short-lived authorization
7. repeat every Δt :
8. $ctx \leftarrow \text{TelemetryCollect}(s)$
9. $B \leftarrow \text{AnalyticsEngine.score}(s)$
10. $R \leftarrow \alpha(1-U) + \beta(1-D) + \gamma N + \delta B$
11. if $R \geq R_{block}$ then Terminate(s); break
12. elif $R \geq R_{th}$ then StepUpMFA(s); ReEvaluate()
13. else PDP.renew(s , scope, ttl= Δt)
14. end repeat

4. Evaluation Methodology

4.1 Setup

The framework was evaluated in a virtualized testbed reproducing a CMG-style ePDG deployment over an NFVI substrate. Six attack scenarios (T1–T6) were generated using a combination of in-house traffic generators and open-source red-team tools (e.g., a rogue-AP rig, an IKEv2 downgrade proxy, and a replay engine). Each scenario was executed 1,000 times against both a baseline ePDG (one-shot EAP-AKA, no continuous verification) and the proposed Zero Trust ePDG (Algorithm 1, $\Delta t = 30$ s, $R_{th} = 0.55$, $R_{block} = 0.80$).

4.2 Metrics

- Attack detection rate: fraction of attack instances that the ePDG correctly identifies and blocks before service impact.
- False-positive rate: fraction of legitimate sessions terminated as suspicious.
- Authentication latency: end-to-end time from IKE_SA_INIT to PDP grant.
- Risk-score responsiveness: time from event onset to R crossing R_{th} .

5. Results and Discussion




Fig. 5. Attack detection rate comparison across six attack classes. The Zero Trust ePDG framework lifts detection from a 46% mean baseline to 94% mean across the benchmark.

5.1 Attack Detection

Fig. 5 reports per-attack detection rates. The baseline ePDG detects only 31–64% of attacks depending on type, because its only signal is EAP-AKA correctness; once the cryptographic exchange succeeds the session is treated as trustworthy. The proposed framework detects 89–99% of attacks, with the largest gains on Rogue AP (+52 pp), Credential Theft (+53 pp), and Device Spoofing (+58 pp). These three are precisely the cases where one-time cryptographic verification is least informative and where continuous behavioral / posture signals dominate.

Fig. 6. Authentication Latency Evaluation




Fig. 6. Authentication latency evaluation: (a) cumulative distribution of end-to-end auth latency; (b) latency broken down by phase. Zero Trust adds ≈ 30 ms of identity + risk overhead but keeps tail latency bounded.

5.2 Authentication Latency

The CDF in Fig. 6(a) shows that the proposed framework adds approximately 30 ms to the median authentication path, dominated by the Identity Engine (MFA + fingerprint) and Risk Engine evaluation. The tail (95th percentile) remains under 220 ms, which is well inside the typical attach-time budget for VoWiFi. Fig. 6(b) breaks the latency down by phase: IPsec setup and EAP-AKA are essentially unchanged, while the new Identity (18 ms) and Risk (12 ms) stages account for the additional cost.



Fig. 7. Dynamic risk scoring trend during a 300-second session. Two anomalies are detected; the score crosses R_{th} , triggers step-up MFA, and recovers to baseline after re-authentication.

5.3 Continuous Trust Evaluation

Fig. 7 illustrates the dynamic risk score $R(t)$ over a 300-second session under two injected anomalies. Anomaly A — a MITM downgrade attempt at $t \approx 80$ s — drives R briefly above 0.55, immediately triggering step-up MFA. Anomaly B — a slow lateral movement attempt from $t \approx 150$ s — produces a gradual upward drift in R that the policy engine captures; a forced re-authentication at $t \approx 240$ s restores the score to its baseline. Critically, the IPsec session itself is not torn down for transient excursions: the operator's policy only requires step-up authentication, which preserves continuity of service while still neutralizing the threat.

5.4 Summary

Metric	Traditional ePDG	Zero Trust ePDG (Proposed)
Mean Attack Detection Rate	46 %	94 % (+48 pp)
Rogue-AP Detection	42 %	94 %
MITM Detection	55 %	97 %
Credential-Theft Detection	38 %	91 %
Median Auth Latency	100 ms	134 ms (+34 ms)
Continuous Re-evaluation	None	Every 30 s + on event
Trust Model	Perimeter (one-shot)	Continuous, context-aware

Table I. Security and performance comparison: Traditional ePDG vs. Zero Trust ePDG.

6. Deployment Considerations

The framework is designed for incremental adoption on existing ePDG installations. The Identity Engine, Risk Engine, Analytics, and PDP run as microservices on the operator's control-plane VMs (or CNFs) and communicate with the ePDG over the existing OAM and SWm signaling fabric. The PEP is implemented as a hook inside the ePDG PCEF/ADC pipeline, so no change is required on the SWu, S2b, or SGi reference points. Operator-side data sources — UE profile, AP reputation, threat intel feeds — are integrated via the Trust & Policy Bus. The framework adds negligible per-packet overhead because trust computation runs on the control plane out-of-band.

7. Future Work

Three directions are being pursued. First, the Risk Engine is being extended from a hand-weighted linear score to a learned model — initially a gradient-boosted classifier and then a small reinforcement-learning policy that adapts the weights α , β , γ , δ per service class. Second, the framework is being integrated with confidential computing primitives (Intel TDX / AMD SEV-SNP) so that the PDP and Risk Engine can run inside attestable enclaves, removing the operator itself from the trust base for high-assurance customers. Third, federated learning is being explored for the Analytics Engine so that per-UE behavioral baselines can be improved without centralizing raw behavioral data.

8. Conclusion

This paper presented a Zero Trust Security Framework specifically tailored to the ePDG in 5G untrusted Wi-Fi access networks. Five concentric defence layers, a six-phase continuous authentication workflow, an explicit threat-to-countermeasure mapping, and a dynamic risk model based on $R = \alpha(1 - U) + \beta(1 - D) + \gamma N + \delta B$ together replace the implicit, one-shot trust of current deployments with explicit, continuous verification. Across a six-attack benchmark, the framework raises mean detection from 46 % to 94 % while adding only ≈ 34 ms of authentication latency. The design is standards-compliant with 3GPP TS 23.402 and NIST SP 800-207, is incrementally deployable on existing infrastructure, and provides a credible upgrade path for operators carrying VoWiFi, enterprise Wi-Fi offload, and roaming traffic over untrusted access.

References

- [1] 3GPP TS 23.402, "Architecture enhancements for non-3GPP accesses," Release 13.
- [2] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST Special Publication 800-207, Aug. 2020.
- [3] 3GPP TS 29.274, "Evolved General Packet Radio Service (GPRS) Tunnelling Protocol for Control plane (GTPv2-C); Stage 3," Release 13.
- [4] 3GPP TS 33.501, "Security architecture and procedures for 5G System," Release 17.
- [5] J. Kindervag, "No More Chewy Centers: Introducing the Zero Trust Model of Information Security," Forrester Research, 2010.

- [6] ETSI GS NFV-SEC 001, "Network Functions Virtualisation (NFV) Security; Problem Statement," ETSI, 2014.
- [7] M. Shirode, J. Doshi, and T. Kothari, "Zero Trust Security Framework for Network Functions Virtualisation Environments," Proc. IEEE GLOBECOM Workshops, 2021.
- [8] H. Ning, H. Liu, J. Ma, L. T. Yang, and R. Huang, "Security in 5G Networks: A Survey," IEEE Communications Surveys & Tutorials, 2019.
- [9] S. Sikdar, S. Chen, and K. Ren, "Secure 5G Architectures for IoT and Mobile Edge Computing," IEEE Internet of Things Journal, 2019.
- [10] C. Kolias, A. Stavrou, J. Voas, I. Bojanova, and R. Kuhn, "Learning Internet-of-Things Security 'Hands-On'," IEEE Security & Privacy, 2016.
- [11] Cisco Systems, "ePDG Security and Wi-Fi Offload — Architecture White Paper," Cisco, 2020.
- [12] IETF RFC 7296, "Internet Key Exchange Protocol Version 2 (IKEv2)," 2014.
- [13] IETF RFC 5448, "Improved Extensible Authentication Protocol Method for 3rd Generation Authentication and Key Agreement (EAP-AKA')," 2009.
- [14] M. Liyanage, A. Braeken, A. Jurcut, M. Ylianttila, and A. Gurtov, "A Comprehensive Survey of 5G Security," IEEE Communications Surveys & Tutorials, vol. 22, no. 1, 2020.
- [15] B. Zhang and J. McCune, "Continuous Authentication for Cellular Networks," Proc. ACM CCS, 2022.
- [16] Malik, N., & Singh, K. (2025). Enhancing agro-industrial development in Haryana through crop diversification and value chain integration. *Journal of Artificial Intelligence in Agriculture and Applied Research*, 1(1).