



LJMU Research Online

Hashem Eiza, M, Akwirry, B, Raschella, A, Mackay, M and Maheshwari (Postdoc), M

A Hybrid Zero Trust Deployment Model for Securing O-RAN Architecture in 6G Networks

<https://researchonline.ljmu.ac.uk/id/eprint/26906/>

Article

Citation (please note it is advisable to refer to the publisher's version if you intend to cite from this work)

Hashem Eiza, M ORCID logoORCID: <https://orcid.org/0000-0001-9114-8577>, Akwirry, B, Raschella, A, Mackay, M and Maheshwari (Postdoc), M A Hybrid Zero Trust Deployment Model for Securing O-RAN Architecture in 6G Networks. Future Internet. (Accepted)

LJMU has developed **LJMU Research Online** for users to access the research output of the University more effectively. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LJMU Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain.

The version presented here may differ from the published version or from the version of the record. Please see the repository URL above for details on accessing the published version and note that access may require a subscription.

For more information please contact researchonline@ljmu.ac.uk

<http://researchonline.ljmu.ac.uk/>

A Hybrid Zero Trust Deployment Model for Securing O-RAN Architecture in 6G Networks

Max Hashem Eiza^{1,*}, Brian Akwirry², Alessandro Raschella¹, Michael Mackay¹ and Mukesh Kumar Maheshwari^{1,3}

¹ School of Computer Science and Mathematics, Liverpool John Moores University, Liverpool L3 3AF, UK; [M.HashemEiza, A.Raschella, M.I.Mackay, M.K.Maheshwari]@ljmu.ac.uk

² School of Engineering and Computing, University of Lancashire, Preston PR1 2HE, UK; BO-Aboakwirry@uclan.ac.uk

³ Department of Electrical Engineering Bahria University, Karachi Campus, 75260 Pakistan. Email: mukeshkumar.bukc@bahria.edu.pk

* Correspondence: M.HashemEiza@ljmu.ac.uk

Abstract: The evolution toward sixth generation (6G) wireless networks promises higher performance, greater flexibility, and enhanced intelligence. However, it also introduces a substantially enlarged attack surface driven by open, disaggregated, and multi-vendor Open RAN (O-RAN) architectures that will be utilised in 6G networks. This paper addresses the urgent need for a practical Zero Trust (ZT) deployment model tailored to O-RAN specification. To do so, we introduce a novel hybrid ZT deployment model that establishes the trusted foundation for AI/ML-driven security in O-RAN, integrating macro-level enclave segmentation with micro-level application sandboxing for xApps/rApps. In our model, the Policy Decision Point (PDP) centrally manages dynamic policies, while distributed Policy Enforcement Points (PEPs) reside in logical enclaves, agents, and gateways to enable per-session, least-privilege access control across all O-RAN interfaces. We demonstrate feasibility via a Proof of Concept (PoC) implemented with Kubernetes and Istio and based on the NIST Policy Machine (PM). The PoC illustrates how pods can represent enclaves and sidecar proxies can embody combined agent/gateway functions. Performance discussion indicates that enclave-based deployment adds 1–10 ms of additional per-connection latency while CPU/memory overhead from running a sidecar proxy per enclave is approximately 5–10 % extra utilisation, with each proxy consuming roughly 100–200 MB of RAM.

Keywords: 6G, Security, Next-G Networks, O-RAN, Open RAN, Zero Trust, NIST PM

1. Introduction

Research and development activities towards the sixth generation (6G) of wireless communications are gaining pace. The first commercial rollout of 6G networks is expected around 2030 while the standardisation process will start this year (2025) according to the 3GPP development roadmap [1]. 6G networks are envisaged to revolutionise the wireless networking ecosystem providing faster, more responsive, and increasingly ubiquitous connectivity than ever before. Moreover, the number and variety of potential 6G applications (e.g., Internet of Things (IoT)/Industry 4.0, Augmented Reality (AR)/Virtual Reality (VR), vehicle-to-everything) across diverse industries and services introduces unprecedented challenges in meeting network requirements. These requirements will include wider coverage, higher data rates, ultra-low latency, ultra-high location accuracy, integration of communications and sensing, more intelligence, stronger security, and better sustainability in comparison with fifth generation (5G) networks.

To accommodate the diverse needs of 6G networks and its applications, new approaches to network architecture are required. This includes openness, disaggregation, and multi-vendor interoperability of the 6G Radio Access Network (RAN). In this context, Open RAN, aka O-RAN as referred to by the O-RAN Alliance [2], are expected to offer several benefits to 6G networks [3]. O-RAN allows for a greater flexibility in selecting and optimising network functions, components, and services. It can also contribute to a smoother transition to 6G by allowing multi-vendor network components to be reused to support 6G scenarios. However, this openness and disaggregation dramatically increases the attack surface and insider threat potential. There is a wide consensus that 6G networks will include O-RAN components and 3GPP as well as legacy deployments [4]. Therefore, there are a wide range of ongoing research activities within organisations such as O-RAN Alliance, 3GPP, and 6G Smart Networks and Services Industry Association (6G-IA) to identify key architecture principles relevant to the standardisation of a future 6G RAN. These principles are guiding the development of different aspects of future O-RAN standards [5].

Resilience, security, and privacy of 6G networks are among the top priorities identified by those architecture principles. The adoption of open architectures and interfaces, cloud deployments, and multi-vendors' networks will expand the attack surface especially to internal attacks due to Advanced Persistent Threats (APTs) and lateral movement. Hence, in 6G networks, it is critical to pursue a Zero Trust Architecture (ZTA) during the standardisation and design phases to achieve a strong security posture to protect against evolving threats [3]. ZTA is based on the Zero Trust (ZT) security paradigm where threats can originate from both external and internal sources [5]. Hence, ZT emphasises verification and validation of every access request to minimise the risk of security breach. All assets and resources are secured as micro-perimeters and no entity, whether internal or external, is assumed to be trusted for access to applications and data regardless of its location and ownership. Therefore, the pursuit of ZT implementation, especially in critical infrastructure such as 5G/6G networks, is now a strategic goal set by governments, standardisation organisations, and industry (e.g., [6], [7]). Examples of the most recent activities and reports for implementing ZT in O-RAN and 5G/6G networks are O-RAN Alliance Working Group (WG11) ZT for secure O-RAN white paper [8] and the National Institute of Standards and Technology (NIST) collaboration with O-RAN Alliance and Alliance for Telecommunications Industry Solutions (ATIS) to fully incorporate ZT into emerging standards for 5G and 6G wireless standards [9].

On the other hand, there are multiple guidance documents that define ZTA, its tenets, different deployment models, and maturity models for full implementation of ZT. Examples of these documents that are relevant to the mobile industry are NIST SP 800-207 ZTA [10], National Security Agency (NSA) and Cybersecurity and Infrastructure Security Agency (CISA) Enduring Security Framework (ESF) Security Guidance for 5G Cloud Infrastructures [11], and CISA ZT Maturity Model (ZTMM) [12]. All these documents explain how to incrementally strive towards a full implementation of ZT through different phases and across different pillars in critical infrastructure such as mobile networks. However, there is a lack of details regarding the deployment and implementation of different ZTA components such as Policy Decision Point (PDP) and Policy Enforcement Point (PEP).

1.1. Motivation & Contributions

Despite the growing consensus around ZT as a foundational principle for securing critical infrastructure, current efforts within the O-RAN domain (e.g., O-RAN Alliance WG11 and their next-Generation Research Group (nGRG) Security) remain conceptual, lacking concrete deployment models tailored to its layered and interface-heavy design.

Most existing works in the literature either focus on securing specific O-RAN interfaces or propose high-level ZT frameworks without addressing practical enforcement mechanisms, dynamic policy management, or the complexities of microservice-based components such as xApps and rApps. This gap leaves operators and vendors with limited technical guidance on how to implement ZT within O-RAN in a scalable, enforceable, and standards-aligned manner. There is little to no discussion about how various security controls operate in practice and how they should be implemented to achieve a mature ZTA in O-RAN.

This paper addresses this gap by proposing and validating a novel hybrid ZT deployment model for securing O-RAN architecture based on O-RAN Alliance technical specification in [13], as to be explained later in Section 2.1. To the best of our knowledge, this is the first work to propose a ZT deployment model purpose-built for O-RAN with full coverage of all O-RAN interfaces and their security controls. The proposed model uniquely integrates:

1. Enclave-based segmentation to enforce macro-level isolation between logically grouped O-RAN functions,
2. Device application sandboxing to isolate xApps and rApps, thereby limiting intra-enclave lateral movement at the micro-level,
3. Distributed PEPs and a centralised PDP, aligned with NIST ZTA tenets and O-RAN Alliance specifications,
4. A Kubernetes and Istio-based Proof of Concept (PoC) implementation, leveraging the NIST Policy Machine (PM) to demonstrate technical feasibility,
5. A comprehensive mapping of ZT deployment approaches (i.e., agent/gateway-based, enclave, resource portal, and sandboxing) to their applicability in O-RAN contexts, and
6. A performance analysis based on the PoC implementation that considers latency, CPU/memory impact, scalability, and trade-offs of the proposed hybrid ZT model.

Together, these contributions lay a foundation for both research and industry adoption of ZT in O-RAN for the future 6G networks. Note that when we make references to O-RAN security controls and requirements, the reader can refer to [13] and [14] for full details. Furthermore, for more details on ZT logical components including PDP, which is composed of the Policy Engine (PE) and the Policy Administrator (PA), and PEP, we refer the reader to NIST SP 800-207 publication [10]. Unlike AI-centric works, our hybrid ZT model provides the security fabric (i.e., enclaves, sandboxing, dynamic PDP/PEP) necessary to deploy and protect AI-driven xApps/rApps. Without ZTA's per-session verification, AI models themselves could become attack surfaces [49]. Note that prior AI/ML works focus on specific interface and/or component in O-RAN while our model is the first to secure all O-RAN interfaces with ZT principles, creating a trusted foundation for network-wide AI/ML deployment.

The rest of the paper is organised as follows. In Section II, the latest works related to securing O-RAN interfaces and ZTA are presented. Section III presents the hybrid ZT deployment model for O-RAN and explains the rationale behind it in detail. The strategy for implementing the proposed model, a PoC implementation, and its potential impact on O-RAN performance are explained in Section IV. Further analysis of the proposed deployment and implementation strategies, and discussion on future work are provided in Section V. Finally, the paper's conclusions are presented in Section VI.

2. Related Works

2.1. O-RAN Architecture

As mentioned before, the focus of this work is the logical O-RAN architecture illustrated in Figure 1 below including its components and interfaces as defined by O-RAN Alliance [13]. The Service Management and Orchestration (SMO) deals with network orchestration, automation, and optimisation, supporting the Non-Real-Time RIC (non-RT RIC). The non-RT RIC operates with latency above 1 second, optimising the network through rApps, and interacts with the Near-Real-Time RIC (near-RT RIC) via the A1 interface for policy control and model updates. The near-RT RIC operates with latency between 10ms and 1s, and implements real-time optimisation through xApps, and controls the O-RAN Centralised Unit (O-CU) and O-RAN Distributed Unit (O-DU) via the E2 interface. The O-CU is composed of O-CU-CP (Control Plane) for signalling and O-CU-UP (User Plane) for user traffic. The O-DU is linked to the O-CU via the F1 interface, and to the Radio Unit (O-RU) via the Open Fronthaul (OpenFH) interface to support different functional split options. The O1 interface connects the SMO to all O-RAN elements for Fault, Configuration, Accounting, Performance, and Security (FCAPS) management. Finally, the O2 interface allows orchestration, management, and automation of cloud-based RAN (O-Cloud) components.

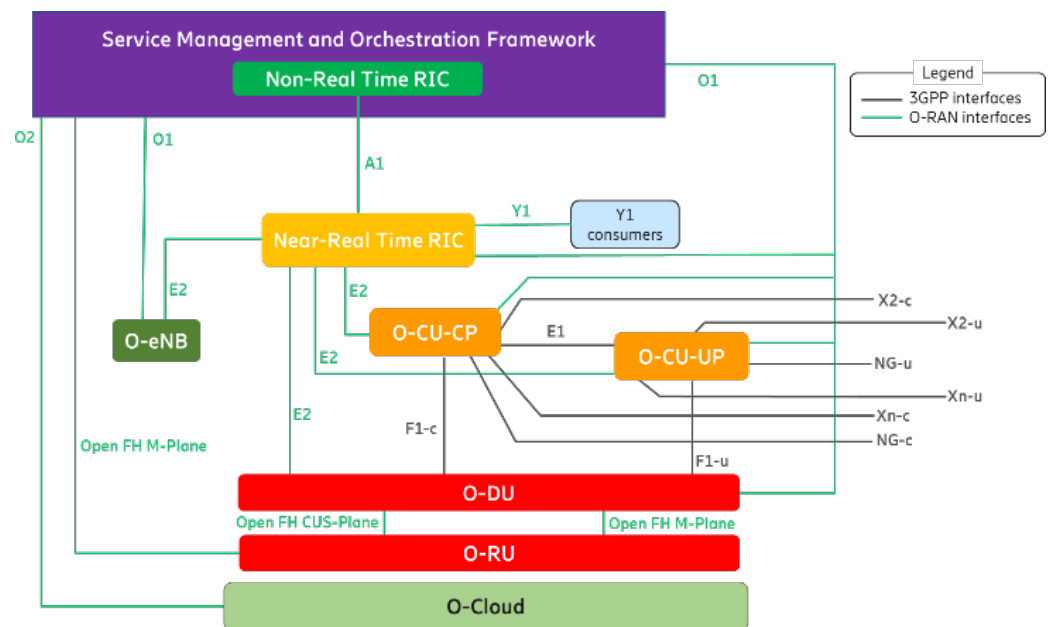


Figure 1. Logical Architecture of O-RAN [13].

2.2. O-RAN Security Challenges & Standards

In this section, we focus on the latest advancements in O-RAN security, particularly those aligned with ZTA principles. In [15], the authors analysed the security threats and challenges associated with O-RAN architecture and presented a comprehensive threat taxonomy of O-RAN components. They categorised risks into three types: process, technology, and global. The process risks are related to challenges to standardising critical processes in O-RAN where security controls can be applied to ensure privacy and integrity of assets. Technology risks mainly come from the use of open-source software and potential vulnerabilities in different components. Finally, global risks, in the broader sense, are related to securing telecom infrastructure against attacks and espionage from international/global attackers. The paper concludes by stressing the need for a comprehensive and robust security standards and governance to ensure the benefits of O-RAN are realised. In relation to ZTA, the authors touched briefly on how ZTA represents a simpler security model that can enhance the security of O-RAN. However, they did not explore the practical aspects of implementing ZTA in O-RAN nor provided practical security measures to advance O-RAN security towards ZTA.

In [49], Polese *et al.* have shown that AI-driven security is pivotal in O-RAN to address dynamic threats. AI could enable real-time anomaly detection, adaptive policy enforcement, and predictive threat mitigation, which are critical for ZT's continuous verification requirements in 6G RAN. This can be achieved via AI-driven anomaly detection engines that reside in O-RAN control loop to continuously analyse interactions across the A1 and E2 interfaces to identify deviations from baseline behaviour, enabling rapid isolation of suspect network elements and thwarting lateral movement. AI can also aid in further adapting security policies in real time (e.g., tuning mTLS cipher suites, IPSec parameters or sandbox permissions) based on evolving threat patterns and contextual telemetry. This AI-based proactive, data-driven policy enforcement will accelerate threat detection and aligns naturally with ZT tenets of continuous verification.

Challenges and solutions to securing O-RAN interfaces were presented in [16]. However, the authors focused their work on two open interfaces only, the E2 and the Open FH. Both interfaces carry sensitive data that should be protected against threats such as data tampering and unauthorised access. Therefore, the paper investigated the impact of encryption in terms of latency and overhead on the performance of these two interfaces. More specifically, the authors presented an experimental analysis using IPSec for the E2 interface and Media Access Control Security (MACsec) for the Open FH interface. The results indicated that securing the E2 interface via IPSec encryption introduces minimal latency and performance overhead, making it a feasible security solution for this interface. However, when using MACsec on the Open FH interface, there was a significant impact that could affect the whole system performance especially in high traffic scenarios. In terms of ZTA, the authors mentioned it very briefly without any details of how it can be used to secure O-RAN.

In the same context, through its WG11, the O-RAN Alliance continues to advance O-RAN security via their work in three security specifications and a technical report that form the pillars of O-RAN security. These are 1) O-RAN Security Threat Modelling and Risk Assessment [17]; 2) O-RAN Security Requirements and Controls Specifications [18] 3) O-RAN Security Protocols Specifications [19]; and 4) O-RAN Security Tests Specifications [20]. The specified security requirements and security controls aim to mitigate risk from external and internal threats in pursuit of a ZTA, which is one of thirteen active security work items in WG11. For instance, WG11 has suggested that Service Management and Orchestration (SMO) functions must support authentication of both internal and external systems, as well as authorisation of service requests [21]. Mutual authentication is also required for SMO internal communications. On the other hand, the non-RT RIC should support authorisation, and for the transport layer security, mutual TLS (mTLS) is specified to be used to provide API authentication and authorisation. JSON Web Token (JWT) is specified for application layer authentication, and strict authorisation protocols for API access and discovery.

Finally, the latest white paper published by O-RAN Alliance WG11 confirmed that all NIST's seven tenets of ZTA are applicable to O-RAN [22]. These tenets are [10]:

1. All data sources and computing services are considered resources.
2. All communication is secured regardless of network location.
3. Access to individual enterprise/operator resources is granted on a per-session basis.
4. Access to resources is determined by dynamic policy.
5. The enterprise/operator monitors and measures the integrity and security posture of all owned and associated assets.
6. All resource authentication and authorisation are dynamic and strictly enforced before access is allowed.
7. The enterprise/operator collects information about the current state of assets, network infrastructure and communications and uses it to improve its security posture.

The white paper, however, did not provide technical details of security controls and/or deployments that are required to realise these tenets. It mentioned that an upcoming report, which has not been published yet, will address key techniques expected to play an important role toward achieving security goals in 6G use cases without any further details.

While these O-RAN security specifications lay a strong foundation, they stop short of prescribing a concrete ZT deployment model. The next subsection reviews related works that apply ZT specifically to 6G networks or O-RAN.

2.3. Zero Trust in 6G & O-RAN

Abdalla et al. introduced zero trust RAN (ZTRAN) that includes service authentication, intrusion detection, and secure slicing subsystems that are encapsulated as xApps [23]. The authors highlighted the need for a ZT approach in O-RAN, emphasising continuous verification and authentication of users, devices, and applications. The authentication xApp implements multi-factor authentication (MFA) for securely identifying and verifying user equipment (UE) requesting access to O-RAN services. The intrusion detection xApp monitors the network activities continuously to detect suspicious behaviour. This is done by analysing Key Performance Metrics (KPM) and creating user behaviour profiles. Finally, the slicing xApp implements dynamic network slicing to isolate malicious users by assigning them to restricted instances to ensure legitimate users are not affected by any malicious activities. These three xApps were implemented and tested in the Open Artificial Intelligence Cellular (OAIC) platform to demonstrate its feasibility and effectiveness in terms of improving throughput for legitimate users, latency, and intrusion detection performance (i.e., isolating malicious users). However, some gaps still exist in the proposed ZTRAN framework such as scalability concerns (e.g., with many users and devices generating large-scale traffic) and resolution of xApps conflicts (i.e., where multiple microservices may interfere with each other especially when multiple xApps manage resources simultaneously).

In [24], the authors proposed the OZTrust framework, a ZT security system specifically designed for the O-RAN Near-RT RIC. OZTrust has two components: 1) an access control module that performs per-packet tagging and verification for each packet sent between xApps; and 2) a policy management module that automatically generates access control policies by tracing communication patterns between xApps through distributed tracing. Note that, in the O-RAN context, a unique tag represents the identity of an xApp derived from its authentic contexts. Therefore, OZTrust offers mitigations against compromised xApps and provides fine-grained xApp's access control policies without relying on static credentials or permissions. The OZTrust prototype is implemented based on eZTrust approach [25] and compared with Kubernetes Role-Based Access Control (RBAC) [26], and two Container Network Interfaces (CNIs) Cilium [27] and Calico [28]. OZTrust demonstrated that it can successfully block unauthorised API access and lateral movement attacks, where other systems failed. Besides its narrow focus on xApps security, OZTrust requires complex instrumentation processes to enable distributed tracing in an xApp. This involved inserting tracing APIs and their dependencies into the xApp's source code manually. This becomes challenging when the number of xApps increases and they come from different open-source platforms and different vendors, which might not necessarily collaborate among each other and with carriers.

Finally, Chen et al. [29] proposed a software-defined ZTA for 6G networks reinforcing the need to establish elastic and scalable security regime, and scalable access control. The authors proposed a distributed, community-based security framework in which each SDN-controlled community acts as its own ZT domain, enforcing access control at border switches and treating roaming devices like home/visited networks. Identity management

is decentralised via per-community certificate authorities and a hierarchical identifier scheme stored on a private blockchain, preserving privacy while enabling cross-domain verification. Trust is evaluated at both the device and community levels combining self-reported device posture, third-party feedback, patch history, and anomaly scores. Nonetheless, their work does not address O-RAN's unique multi-vendor, interface-centric architecture. It also lacks both enclave-based segmentation and application sandboxing to contain lateral movement at the macro and micro levels. The validation is done via abstract trust-metric simulations rather than a concrete Kubernetes/Istio/NIST Policy Machine proof-of-concept.

2.4. Comparative Analysis of O-RAN Security Methodologies

Modern O-RAN security solutions span a range of techniques from static rule-based enforcement (e.g., traditional firewalls and Access Control Lists (ACLs)) through sophisticated AI/ML schemes. Table 1 provides different categories of mainstream security approaches for O-RAN and 6G including their advantages and disadvantages. While cryptographic methods provide strong guarantees, they struggle with dynamic threats. AI/ML solutions offer adaptability but lack explainability and require trusted data pipelines. Hardware-assisted techniques (e.g., Trusted Execution Environments (TEEs)) excel in low-latency scenarios but do not address network-level threats. Our hybrid ZTA integrates the strongest aspects: cryptographic controls for data-in-transit (IPSec/mTLS), hardware acceleration where feasible, fully adaptive for dynamic policies while avoiding blockchain's latency penalties.

Table 1. O-RAN Security Methodologies – Comparison

Methodology	Key Techniques	Advantages	Disadvantages
Cryptographic	IPSec, MACsec, mTLS	Provable security, Standardised	High compute overhead, Key management complexity
Signature Based	IDS/IPS that match network traffic against a database of known threat signatures	High accuracy for known attacks, Low false-positive rate on familiar patterns	Ineffective against zero-day exploits, Signature database must be constantly updated
AI/ML Based	Supervised or unsupervised models trained on normal operations to flag anomalies	Adapts to new threats, Real-time response	Training data vulnerability, May suffer false positives under traffic shifts
Hardware Assisted	TEEs, Trust zones	Tamper-proof execution, Low latency	Limited to supported platforms, Additional cost, Does not address network-level threats
Blockchain Based	Distributed Ledger Technology (DLT), Decentralised identity	Immutable audit trails, No single Point of Failure, Transparent compliance	High latency, Energy intensive, Scalability concerns
Zero Trust Based (Proposed)	Dynamic PEP/PDP, Micro-segmentation	Least-privilege enforcement, full interface coverage, Fully adaptive	Policy management complexity, ~1-10ms latency overhead

Table 2 summarises related works' approaches, security methodology, limitations, and how our work differs.

Table 2. Comparison with Related Works.

Related Work	Approach	Security Methodology	Key Limitations	Our Work
Liyanage et al. [15]	<ul style="list-style-type: none"> Threat taxonomy and high-level security challenges in O-RAN. Brief mention of ZT as a conceptual model. 	<ul style="list-style-type: none"> Conceptual 	<ul style="list-style-type: none"> No concrete ZT deployment model tailored to O-RAN. Lacks practical enforcement or interface level controls. 	<ul style="list-style-type: none"> Proposes a detailed hybrid ZT model covering all O-RAN interfaces. Specifies PEP and PDP placement. Develops a PoC with Kubernetes/Istio/NIST PM.
Groen et al. [16]	<ul style="list-style-type: none"> Experimental analysis of IPSec (E2) and MACsec (OpenFH) performance overhead. 	<ul style="list-style-type: none"> Cryptographic 	<ul style="list-style-type: none"> Focuses only on two interfaces (E2, OpenFH). Mentions ZT only briefly without design details or policy mapping. 	<ul style="list-style-type: none"> Covers every major O-RAN interface. Integrates ZT tenets into one unified deployment strategy.
O-RAN WG11 (Threats, Security Requirements, and White Paper) [17-22]	<ul style="list-style-type: none"> O-RAN security threat modelling, requirements, controls, protocols, and tests. White paper confirms ZT tenets apply. 	<ul style="list-style-type: none"> AI/ML Based 	<ul style="list-style-type: none"> Standards/specs describe controls but do not prescribe deployment architecture. No implementation guidance or PoC. 	<ul style="list-style-type: none"> Provides a reference architecture with detailed mapping to NIST ZTA tenets and real-world PoC using cloud-native tools.
Abdalla et al. (ZTRAN) [23]	<ul style="list-style-type: none"> Zero Trust RAN (xApps for service authentication, intrusion detection, slicing). 	<ul style="list-style-type: none"> Zero Trust Based 	<ul style="list-style-type: none"> Narrow focus on xApps within Near-RT RIC. Does not address macro-level enclaves or multi-interface enforcement. 	<ul style="list-style-type: none"> Introduces both macro (enclave) and micro (sandbox) isolation. Covers all O-RAN logical groups (SMO, RICs, DU/CU, O-Cloud) and every interface.
Jiang et al. (OZTrust) [24]	<ul style="list-style-type: none"> Fine-grained access control for xApps via per-packet tagging and distributed tracing. 	<ul style="list-style-type: none"> Signature Based 	<ul style="list-style-type: none"> Limited to xApps, requires manual instrumentation. No enclave-level segmentation or policy engine discussion. 	<ul style="list-style-type: none"> Demonstrates a central policy engine (i.e., NIST PM) enforcing attribute-based rules across enclaves and sandboxing without manual code changes in xApps.
Chen et al. [25]	<ul style="list-style-type: none"> Software-defined ZTA for 6G. Distributed “communities” with border enforcement, Blockchain IDs, and trust metrics. 	<ul style="list-style-type: none"> Blockchain Based 	<ul style="list-style-type: none"> Abstract, non-O-RAN-specific. No mapping to O-RAN interfaces or components. 	<ul style="list-style-type: none"> Tailor ZT deployment specifically to O-RAN’s multi-vendor, interface-centric architecture. Validates via Kubernetes/Istio PoC rather than abstract models.

3. The Hybrid ZT Deployment Model For O-RAN

3.1. ZT Approaches & Potential Implementation in O-RAN

Within the ZT framework, there are several approaches that can be considered for implementation in O-RAN. Each approach implements all tenets of ZT, and a full ZT solution includes elements from all these approaches. In the following, we briefly explain each approach and how it can be applied to an O-RAN architecture.

3.1.1 Enhanced Identity Governance-Driven Approach

This approach focuses on managing and securing identities within an organisation. It ensures that only the right individuals have access to the right resources at the right times, and for the right reasons [30]. Authentication and authorisation are effective measures in disaggregated networks to ensure only authorised identities have access to resources. It mitigates the risk of unauthorised access and potential security breaches. Implementing identity-based controls capabilities could integrate additional capabilities such as compliance to regulatory requirements [31].

Implementation in O-RAN – the enhanced identity governance-driven approach in O-RAN can be utilised to enforce strict identity and access management across the disaggregated and virtualised components of the O-RAN architecture. This approach involves the use of advanced identity management systems to ensure that only authenticated and authorised entities (e.g., RIC administrators, vetted third-party xApp vendors, devices, and services) can access and interact with the O-RAN components. This can be achieved through mechanisms such as MFA, RBAC, access tokens issued by the centralised PDP, and continuous monitoring of identity usage patterns. This is important to mitigate risks of compromised Commercial-off-the-shelf (COTS) servers or rogue devices in virtualised RAN [22]. Furthermore, the automated identity management reduces administrative overhead and minimises the risk of human error associated with manual identity and access management processes.

3.1.2 Logical Micro-segmentation

This approach involves dividing the network into smaller, isolated segments at a granular level where access and movement between segments are controlled by a policy [32]. By isolating different parts of the network, micro-segmentation limits the lateral movement of attackers, thus containing potential breaches and reducing the attack surface. This containment is particularly advantageous in a disaggregated and highly dynamic environment like O-RAN. The micro-segmentation also enhances the ability to detect and respond to threats more effectively by providing enhanced visibility in the network [33]. Finally, micro-segmentation enables the application of specific security policies tailored to the needs of each segment, enhancing overall network security.

Implementation in O-RAN – logical micro-segmentation can be applied to O-RAN by dividing the network into smaller, isolated segments based on logical parameters such as function, user group, or data sensitivity. This involves creating virtual network segments within the O-RAN infrastructure, where each segment has its own set of security policies and controls. Each main O-RAN function (e.g., SMO, RIC) resides in its own segment, with strictly scoped east-west policies enforced by PEPs. For example, only the SMO orchestration service which possess an access token can send policy updates on A1 interface to the RIC while other SMO components are limited to read-only on O1 interface. Micro-segmentation can be implemented using software-defined networking (SDN) and network function virtualisation (NFV) technologies to dynamically enforce segmentation policies. This approach provides detailed visibility into network traffic and enables precise control over data flows. Therefore, enhancing the ability to detect and respond to

threats more effectively. Moreover, this aids in containing lateral movement in a distributed O-RAN cloud.

3.1.3 Network-based segmentation

It involves dividing the network into distinct segments based on physical or virtual boundaries. It typically relies on traditional network infrastructure, such as firewalls and Virtual LANs (VLANs). This approach provides strong isolation between different network functions and segments, preventing unauthorised access and data breaches [34]. It also allows for centralised management and enforcement of security policies across different network segments, simplifying the overall security posture. This offers further flexibility and control in managing security policies across the entire network infrastructure.

Implementation in O-RAN – network-based segmentation can be used to divide the O-RAN infrastructure into distinct zones or segments, each with its own security controls and policies. This can be done using firewalls and VLANs to create secure boundaries between different network segments. Network-based segmentation is particularly useful for separating critical network functions, such as control plane and user plane functions, to enhance security and manageability.

3.2. ZT Deployment Variations for O-RAN

According to NIST SP 800-207 [10], there are four deployment variations/models for ZTA. In the following, we discuss each deployment model including its advantages, disadvantages, and applicability in an O-RAN architecture.

3.2.1 Device Agent/Gateway-Based Deployment

This model splits the PEP into two components: a device agent and a gateway. The device agent, installed on each enterprise-issued asset, handles connections and forwards access requests to the policy administrator. The gateway component, located directly in front of a resource, acts as a proxy. This model is suitable for enterprises with a strong device management programme. Hence, it prevents Bring Your Own Device (BYOD) policies, allowing access only via enterprise-owned assets.

In O-RAN, this could translate to deploying device agents on network elements like RUs, DUs, and CUs, with gateways managing access and communication. This model ensures that only authorised devices can communicate. The agent installed on each network element manages the local network traffic, performs initial authentication, and forwards access requests to the policy administrator. It also ensures that each device adheres to the security policies before initiating any communication. The gateway acts as a mediator between the network elements and enterprise resources and enforces security policies based on the decisions made by the policy administrator. Moreover, it is responsible for maintaining secure communication channels and ensures that only authorised requests are processed. The advantages of this model are a) ensuring that unauthorised devices cannot access critical network resources; b) each communication request is thoroughly authenticated and authorised, reducing the risk of malicious activities; and c) this model allows for granular control over network access, ensuring that only necessary and authorised communications are permitted.

However, the disadvantages are a) this model requires a robust device management programme to handle the deployment and maintenance of device agents across all network elements. The need for continuous monitoring and updates to ensure that all agents and gateways are functioning correctly can add to the management overhead; b) implementing this model may incur higher costs due to the need for specialised hardware (gateways) and the development and deployment of device agents; c) the model does not support BYOD policies limiting its flexibility; d) enterprises must ensure that all network elements and devices are enterprise-issued, which might not be feasible in all scenarios; and

e) it requires significant planning and testing to ensure seamless integration without disrupting existing services.

3.2.2 Enclave-Based Deployment

Unlike the device agent/gateway-based deployment model, this one places the gateway at the boundary of a resource enclave rather than individual resources. Resources within an enclave, such as data centre or cloud-based microservices, are protected by a single gateway. This makes it ideal for environments where resources serve a single business function or for legacy systems that cannot directly communicate with a gateway. This model still needs an agent installed on all devices that need access to resources.

In O-RAN, the network can be segmented into enclaves where each represents a different network segment (e.g., cloud, edge). Gateways at these enclaves' boundaries can manage/control access and enforce security policies for devices/resources within the enclave ensuring that only authorised requests are processed. Resources within an enclave are grouped based on their function or service type. Each enclave operates independently, with the gateway providing a secure access point.

The advantages of this model are a) centralised security enforcement at the enclave boundary, simplifying the management of access policies; b) reduces the complexity of deploying and maintaining individual gateways for each resource; c) provides a protective barrier for legacy systems that cannot support direct communication with a gateway; d) easily scalable by adding new enclaves and gateways as the network grows; and e) cost effective as it reduces the need for deploying multiple gateways, lowering the overall cost of implementation. On the other hand, the disadvantages are a) the gateway at the enclave boundary can become a single point of failure. If compromised, it can impact all resources within the enclave; b) the gateway requires robust redundancy and failover mechanisms to ensure continuous availability; and c) offers less granular control compared to the device agent/gateway model, as the gateway enforces policies at the enclave level rather than on individual resources.

3.2.3 Resource Portal-Based Deployment

A resource portal (i.e., gateway) acts as the PEP, in front of the resource(s), integrating with the enterprise's authentication and authorisation systems. Hence, the PEP is a single component in this model. Users can access resources through the portal, which enforces access policies. In O-RAN, the SMO can act as a portal, providing an interface for managing and accessing various O-RAN resources, ensuring compliance with security policies. The SMO serves as a central access point for users, consolidating access to multiple resources. It can integrate with enterprise authentication systems (e.g., Single Sign-On, Identity Providers) to authenticate users. Thus, the SMO provides a unified access interface for users to access diverse resources, simplifying the user experience.

The advantages of this model are a) there is no need to install a software component on users' devices; b) centralised authentication and authorisation processes, simplifying management and ensuring uniform policy enforcement; and c) easily scalable by adding new resources behind the portal. However, the disadvantages are a) the portal represents a single point of failure which can be susceptible to DoS attacks; b) the portal requires robust provisioning to ensure high availability; c) as all access requests pass through the portal, it can become a performance bottleneck, especially under high load conditions; and d) devices can only be scanned for compliance when they request access making it difficult to continuously monitor them for malware or unpatched vulnerabilities.

3.2.4 Device Application Sandboxing

This deployment involves running applications within a sandboxed environment on devices. It ensures that applications adhere to enterprise security policies by restricting their access to system resources and data via the PEP. Thus, it is effective for managing

and securing applications on BYOD devices or other environments where device control is limited. In O-RAN, sandboxing can be used to isolate x/rApps to ensure they do not interfere with each other and adhere to security policies. Each sandbox operates independently, ensuring that applications do not access unauthorised resources. Enterprise security policies are enforced within the sandbox, controlling what system resources and data the application can access.

The advantages of this model are a) ensuring that applications are contained within a secure environment segmented from the rest of the asset; b) isolating applications from each other, preventing conflicts, and ensuring that one application's malfunction does not affect others. This is particularly important for O-RAN as it has multiple critical applications running simultaneously; and c) allowing the enterprise to enforce security policies on BYOD devices without needing full control over the device. However, the disadvantages are a) running applications in sandboxes can introduce performance overhead due to the additional layer of isolation and control; and b) managing and maintaining sandbox environments for multiple applications can be complex and resource intensive especially when full visibility into client assets is not available for the enterprise.

3.3. The Hybrid ZT Deployment Model for O-RAN

Considering the O-RAN architecture and the advantages/disadvantages of different ZT deployment models described above, we propose a hybrid ZT deployment model that integrates an enclave-based approach with device application sandboxing. This hybrid model gives us the ability to segment major functional groups via enclave borders, and isolate multi-vendor xApps/rApps via sandboxing. Thus, our hybrid model provides both macro-level (inter-enclave) and micro-level (intra-application) isolation. Figure 2 below shows how management, control, radio, and O-Cloud enclaves interconnect via agents and gateways, with sandboxed xApps/rApps operating inside each RIC as to be illustrated in Figures 3 and 4 below. Note that, in comparison to Figure 1, we omitted some 3GPP interfaces for simplicity and focused only on O-RAN interfaces.

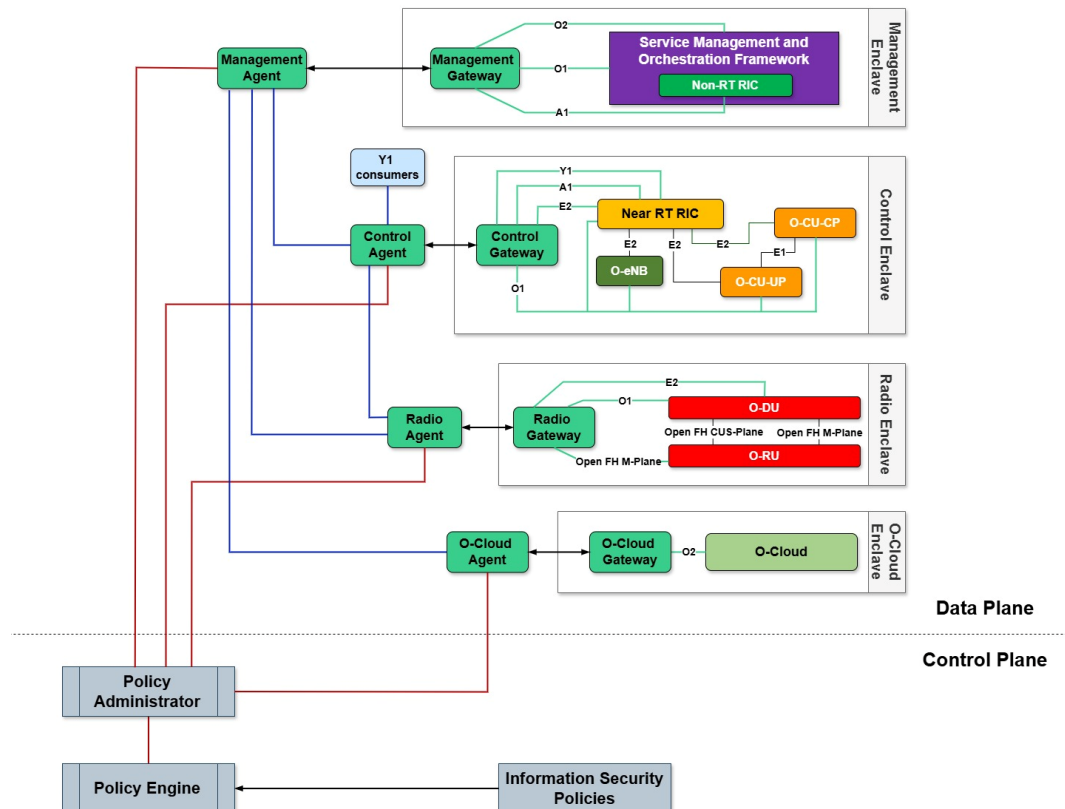


Figure 2. The Hybrid ZT Deployment Model for O-RAN, illustrating enclave boundaries (Management, Control, Radio, O-Cloud) and agent/gateway deployment.

In this hybrid model, we have 1) PDP: a central controller that manages dynamic policies for all O-RAN components; (2) Agents: lightweight PEPs on each network element that forward access requests to the PDP; (3) Gateways: boundary PEPs at enclave edges, enforcing approved traffic paths; and (4) Enclaves: logical groupings of related O-RAN functions (e.g., Management, Control, Radio, O-Cloud). The control plane, which includes the PDP, is composed of the policy engine (PE) and the policy administrator (PA), and the information security policies that should be enforced. This ensures a centralised management of these policies that are pushed to the agents associated with the various enclaves, ensuring consistent policy enforcement across the data plane. On the other hand, the PEP is distributed across three key components:

- Enclaves: These are collections of resources, comprising elements of the O-RAN architecture that share similar functionalities and support comparable control loops. Enclaves provide isolation and resource management for these services.
- Agents: These software agents facilitate connections and communication between different enclaves. An agent is responsible for directing a portion (or all) of the traffic to the policy administrator for evaluation.
- Gateways: These act as intermediaries for resources in their respecting enclaves, ensuring that ingress and egress traffic (i.e., communications) between resources across different enclaves occur exclusively through the gateways. The gateway communicates with the PA, via its dedicated agent, enforcing approved communication paths as configured by the PA.

Our hybrid model leverages a combination of traditional network-virtualisation (e.g., VLAN and NFV) for macro-segmentation and modern service-mesh controls for micro-segmentation. For instance, NFV orchestration can be used to instantiate O-RAN functions such as SMO, RIC, O-CU, and O-DU as a set of VNFs or cloud-native network functions (CNFs). Each enclave sits in its own VLAN where access controls ensure that for instance, the O-CU-VNF network segment cannot directly reach the O-DU-UP segment. On the other hand, in terms of micro-segmentation, within each enclave, every micro-service (e.g., xApp, telemetry collector) runs alongside a sidecar proxy. mTLS certificates, which are issued by the PDP's PKI authenticate both client and server before any call is allowed. The PDP guarantees a unified identity governance where all enclaves' PEPs fetch policies and updated access controls from the PE.

In Figure 2 above, the operational flow commences as follows. When a resource intends to communicate with another resource, the relevant agent intercepts the access request and forwards it to the PA. The PA then submits this request to the PE for evaluation. Upon authorisation, the PA configures a secure communication channel between the agent and the resource gateway via the control plane, including the necessary security artifacts, port information, and session keys. Encrypted data transfer begins once the agent and the destination gateway establish the connection, which is terminated either upon workflow completion or as directed by the PA.

In terms of device application sandboxing, we propose that ML and AI applications as well as RAN-specific applications (rApps/xApps), operate within secure sandbox environments. These applications can directly communicate with the policy administrator to request access to required resources, allowing them to remain segmented from other assets within the system. The decision whether to run each xApp, rApp in its own sandbox, or bundle all xApps, and all rApps in their respective single sandbox, or bundle similar xApps, rApps functionality-wise in their respective sandboxes, is left for the deployment stage. The decision should consider balancing the security requirements versus the

complexity and overhead from managing too many sandboxed environments. More details will be provided regarding rApps and xApps sandboxing in the following sections where we explain each enclave in Figure 2.

In terms of synergy with AI/ML security controls, our hybrid ZT model enables AI-enhanced security in three keyways. First, the PDP can act as an AI policy orchestrator that ingests AI-driven threat scores from sandboxed xApps to dynamically adjust access policies. Secondly, AI workloads are sandboxed in their respective xApps/rApps and isolated via micro-segmentation, preventing compromise of adjacent functions. Thirdly, Kiali metrics, as explained in Section 4.4, provide trusted data streams to train ML models without exfiltration risks.

3.3.1 Management Enclave

It combines SMO and the Non-RT RIC since they are interdependent. The management gateway/agent enforces policies over A1 (to Near-RT RIC), O1 (to O-DU and O-eNB), and O2 (to O-Cloud). All rApps run in a single sandboxed environment. This ensures their segmentation and isolation from other system components while reducing management complexity. As identified above, rApps can communicate directly with the policy administrator as shown in Figure 3.

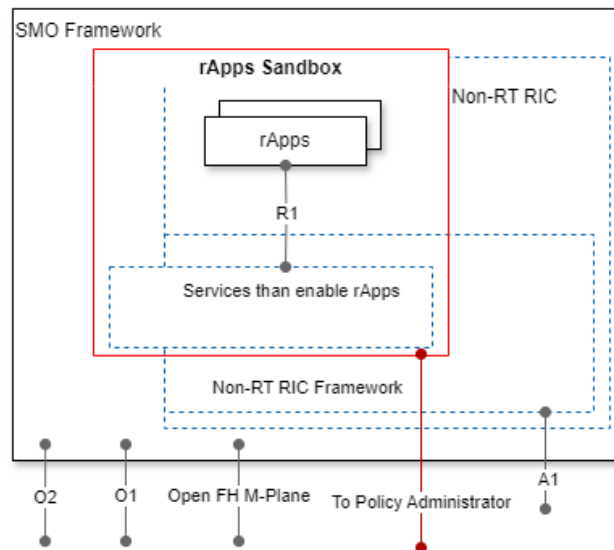


Figure 3. Sandboxed rApps.

3.3.2 Control Enclave

It contains the Near-RT RIC, O-CU-CP, O-CU-UP, and O-eNB. Within this enclave, all xApps run in a single sandboxed environment providing an additional protective layer and ensuring their segmentation from other system components as shown in Figure 4 below. The control gateway handles the E2, Y1, A1, and O1 interfaces. Inside the enclave, the Near-RT RIC establishes connectivity with the O-eNB, O-CU-CP, and O-CU-UP via the E2 interface directly (i.e., without having to go through the control gateway.) Nonetheless, the security policies are still enforced inside the enclave via the control gateway. One of the main differences compared with the O-RAN architecture in Figure 1 is that Y1 consumers do not have direct access to the Near-RT RIC, which improves access control. Note that the control gateway also handles communications via the 3GPP interfaces including F1-c, F1-u, NG-c, Xn-c, Xn-u, NG-u, X2-u and X2-c, which were omitted from Figure 2 for simplicity. Note that for implementation purposes, the control gateway can be split in terms of the functionalities and interfaces they support. We will discuss this in more details in Section 4.

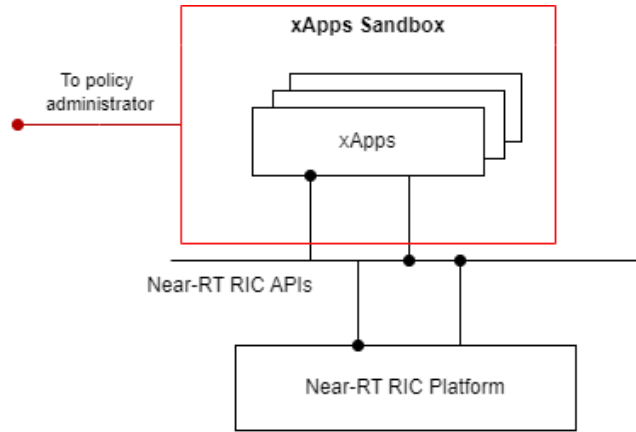


Figure 4. Sandboxed xApps.

3.3.3 Radio Enclave

It includes the O-DU and O-RU. The radio gateway supports E2, O1, and Open FH M-Plane for communications with other enclaves.

3.3.4 O-Cloud Enclave

It contains the O-Cloud environment where the O-Cloud gateway supports O2 interface for communication with other resources in different enclaves.

4. Implementation Strategy of The Hybrid ZT Deployment Model For O-RAN

In this section, we focus on the security controls and specifications for communications among different components in the proposed hybrid ZT deployment model. We also provide high-level systematic plans to implement a ZT deployment model followed by details of our PoC implementation using NIST PM [35].

4.1. Multi-Layered Security Across Control & Data Planes

To provide defence in depth, a multi-layered security approach should be implemented. This involves using both IPSec and mTLS to create a robust security framework that ensures secure communication both at the network and application layers. This combination leverages the strengths of both protocols, providing comprehensive protection against a wide range of threats. This multi-layered approach is consistent with ZT principles. IPSec operates at the network layer, securing all IP packets between the communicating parties. It provides encryption, integrity, and authentication for all data transmitted over the network. IPSec will handle the secure transportation of IP packets between agents and gateways at the network layer. IPSec encrypts the entire IP packet, ensuring that the data remains confidential while in transit. By using hashing algorithms, IPSec ensures that the data has not been altered during transmission, providing protection against tampering. Moreover, IPSec can authenticate the source of the IP packets, ensuring that they come from a trusted source.

On the other hand, mTLS operates at the application layer, providing end-to-end security between applications. It ensures mutual authentication and encrypts the data between the source and the destination. mTLS will secure the application data by providing encryption, mutual authentication, and integrity verification at the application layer. Both the communicating agents authenticate each other using certificates, ensuring that both parties are trusted. mTLS also establishes a secure session that protects the data throughout the entire communication process.

4.2. Agent/Gateway Security Controls

To comply with the specified O-RAN security controls in [18], the agents and gateways will implement these controls as follows. For any communication over the A1 interface, the management and control agent/gateway will employ 1) TLS for confidentiality, integrity, and replay prevention; 2) mTLS for authenticity and data origin authentication; and 3) OAuth for authorisation. Accordingly, management and control gateways/agents enforce these controls, with the PDP/PE pushing the necessary certificates and tokens at session setup. The O2 interface uses identical constraints.

On the other hand, the E2 interface must utilise IPsec for all security controls, including confidentiality, integrity, authenticity, authorisation, data origin authentication, and replay prevention. Note that the E2 interface between the Near-RT RIC and O-DU operates within a control loop between 10ms and 1s. Therefore, the radio and control gateways and agents must ensure they facilitate this connection within these time constraints. This requirement should be considered while implementing the radio and control gateways and agents to meet the E2 interface performance requirements. It is worth noting that enclave-based deployments have been demonstrated to introduce minimal communication overhead [4], performance overhead [5], and memory usage [6] during operation. These enclaves will therefore introduce minimal effects on the communication speed. The secure communication channel will make use of IPsec, which will be provided by PA/PE.

The O1 interface utilises TLS for confidentiality, integrity, and replay prevention; mTLS for authenticity and data origin authentication; and the Network Configuration Access Control Model (NACM) for authorisation. Connecting the SMO with the O-DU, the O1 interface operates within a control loop of 10ms to 1s. Hence, all the relevant gateways and agents that support the O1 interface should be designed to function within these time constraints. As pointed out earlier, given that enclave-based models have been shown to have minimal impact, communication speed is expected to remain unaffected. Finally, although not included in the proposed hybrid model in Figure 2 for simplicity, 3GPP interfaces are expected to follow the security capabilities recommended by 3GPP and O-RAN specifications including TLS and mTLS.

4.3. Systematic High-level Implementation Plan for ZT Deployment in O-RAN

In this systematic plan, we explain the steps required before implementing a ZT deployment model for O-RAN architecture. This plan is utilised while implementing our proposed hybrid model as to be explained later in Section 4.4. First, we need to identify the various network elements, interfaces, software components, and data flows involved, along with the roles of human subjects and non-person entities in managing and interacting with these assets.

- **Assets.** These include O-RUs, O-DUs, O-CUs, Near-RT RIC, Non-RT RIC, SMO framework, and O-Cloud infrastructure.
- **Interfaces.** These include Open Fronthaul between O-RUs and O-DUs, E2 interface between Near-RT RIC and DUs/CUs, A1 interface between Non-RT RIC and Near-RT RIC, O1 interface between SMO and O-RAN nodes, and F1 interface between DUs and CUs.
- **Software Components.** These include xApps and rApps running on RICs for advanced control and optimisation, Virtual Network Functions (VNFs), and AI/ML algorithms for network management.
- **Data Flows.** This includes:
 - a. Control data flows between non-RT RIC and near-RT RIC via A1 interface for policy and control information, and between near-RT RIC and DUs/CUs via E2 interface for real-time control and optimisation.
 - b. User data flows from end devices to O-RUs, then to O-DUs, and finally to O-CUs before reaching the core network.

- c. Telemetry data flows from O-RUs/O-DUs/O-CUs to SMO for monitoring and orchestration via O1 interface. 648 649
- d. AI/ML data flows where data collected from network operations are fed into AI/ML models for training and inference. 650 651
- **Human Subjects.** This includes: 652
 - a. Network Operators. Manage, monitor and maintain the O-RAN infrastructure by ensuring efficient operation, performing troubleshooting, applying configurations, and monitoring performance metrics. 653 654 655
 - b. Service Providers. Offer telecommunication services and applications that utilise the O-RAN network. Their responsibilities include ensuring service quality, managing service delivery and addressing customer requirements. 656 657 658
 - c. Vendors. Supply hardware and software components for the O-RAN ecosystem. They develop, deliver, and support network equipment and software solutions. 659 660
 - d. Developers. Develop, maintain, and update xApps, rApps, and other O-RAN software components. This includes application development, testing, and ensuring compatibility with O-RAN specifications. 661 662 663
 - e. Regulatory Bodies. Ensure compliance with telecommunication standards and regulations. Regulatory bodies also have responsibility to oversee network operations, ensure adherence to regulatory requirements, and protect consumer interests. 664 665 666 667
 - f. End Users. Consumers and businesses use services provided over O-RAN networks. Use network services, report issues, and provide feedback on service quality. 668 669 670
- **Non-person Entities.** This includes: 671
 - a. Service Accounts. These are automated accounts for managing and interacting with O-RAN resources. They perform routine tasks such as data collection, software updates and resource provisioning. 672 673 674
 - b. AI/ML models deployed within xApps/rApps for network optimisation. They analyse data, predict network conditions and suggest or apply configurations for optimisation. 675 676 677
 - c. VNFs and PNFs. These perform specific tasks in the network. Tasks related to data processing signal transmission and network management. 678 679
 - d. Management and orchestration systems. Systems within the SMO that manage network resources and orchestrate services. They coordinate the deployment and operation of network components, ensure resource allocation, and maintain service quality. 680 681 682 683
 - e. Monitoring and Analytics tools. These collect, process, and analyse network data. Their responsibilities include providing insights into network performance, detect anomalies and support decision making process. 684 685 686
 - f. Security systems that enforce security policies and detect threats. They monitor network activities, identify security breaches and apply countermeasures when needed to protect the O-RAN network. 687 688 689

After identifying all the elements in O-RAN architecture, we can proceed with the design, deployment, and testing phases, culminating in continuous monitoring, maintenance, and periodic review to ensure the sustained effectiveness of the deployed hybrid ZT model. These steps are carried out during the preparation phase which we outlined earlier and can be summarised as follows: 690 691 692 693 694

- **Preparation.** This phase has been outlined earlier where assets are identified. It also includes defining the scope and objectives of ZT deployment in O-RAN, assessing the current security posture, and identifying gaps. 695 696 697

- **Design.** Develop a detailed ZT implementation plan, including policies, procedures, and tools required. It also includes designing the architecture for integrating IAM, PEPs, segmentation, monitoring, and encryption within O-RAN architecture.
- **Deployment.** This phase includes deploying IAM solutions, PEPs, monitoring tools, and encryption mechanisms across O-RAN components and interfaces. Besides that, the configuration of access policies, segmentation rules, and encryption protocols.
- **Testing.** Conduct thorough testing to ensure all ZT principles are effectively implemented and integrated. In addition, penetration testing and vulnerability assessments to identify and address any security weaknesses.
- **Monitoring & Maintenance.** Continuously monitor the network for security threats and policy violations. Moreover, regularly update security policies, tools, and configurations based on evolving threats and network changes.
- **Review & Improvement.** Periodically review the effectiveness of the ZT hybrid model implementation and make necessary improvements.

4.4. Policy Machine-based Implementation – Proof of Concept

In this section, we present a PoC implementation of the proposed ZT hybrid model for O-RAN using the following technologies 1) NIST PM; 2) Kubernetes [36]; 3) Istio service mesh [37]; and 4) Kiali for observability [38]. To simplify the implementation and mapping to the design in Figure 2, we define the following two elements.

- **Pod.** Following the terms used in Kubernetes, a pod acts as a logical boundary (i.e., enclave) encapsulating one or more related services. Pods provide isolation and resource management for these services. Each pod can contain multiple containers (i.e., services) that work together. These services are protected from external interference through policies and proxies in front or within the pod.
- **Proxy.** In this case, a proxy could serve as an agent, gateway and policy enforcer. The benefits of combining these components into a proxy include easier management, dynamic policy implementation, and reducing the attack surface. Each proxy serves as an intermediary between services within the enclaves and the outside world, or between internal services, enforcing policies. Proxies receive policy details from the control plane (i.e., the PA) and are responsible for applying them at the pod level, ensuring security and access control. The proxy functionality can be further split into ingress and egress proxies. An ingress proxy handles all incoming traffic to the pod, applying policies to ensure only authorised requests are allowed into the enclave. On the other hand, egress proxy handles traffic exiting the pod. It ensures that any data leaving the pod (i.e., enclave) complies with outbound security policies, preventing unauthorised data transmission or data leaks.

In terms of placing the proxies, two different options can be considered during the deployment phase. First, a central proxy gateway at each physical site, which is easier to harden and monitor, but concentrates all traffic through a single node, introducing a potential single point of failure. The second option is where proxies are embedded with every service. This delivers fine grained enforcement at the expense of higher operational overhead consuming higher resources in terms of memory and computational power. Our PoC implementation adopts a balanced approach whereby we deploy one proxy per enclave. This strikes a pragmatic compromise between visibility and manageability; while still allowing horizontal scaling should future traffic volumes or resilience requirements require finer granularity. Furthermore, during the implementation phase, different decisions can be made to split the functionalities of proxies, arrange services in their enclaves as pods or clusters, and split the gateway functionality to support different interfaces rather than one gateway for multiple interfaces as depicted in Figure 2.

To setup our PoC implementation, virtual machines (VMs) will be utilised. A controller VM will host the Kubernetes control plane, responsible for managing the cluster's state. It will act as the PA and the PE storing and deploying policies for access control. A worker node VM will host a service (e.g., non-RT RIC). It serves as the execution environment for the applications and services within the cluster (i.e., enclave). Test nodes VMs will be used to test access to the service in the worker node VM to ensure authorisation policies and access controls are enforced. We use Kubernetes to create and manage the services, and handle container orchestration, and resource allocation across the VMs. Istio, which is deployed as the service mesh within Kubernetes, is used to secure, connect, and control the services, enabling traffic routing, observability, and applying security policies. Kiali is used to provide GUI to visualise traffic flow, monitor the service mesh, and view connection graphs. It provides real-time insight into the services, traffic distribution, and health of the applications running.

The controller VM runs with 8 vCPUs, 32GB RAM, 60GB SSD, hosting the kube-api-server, istiod and both PA and PE. Two worker VMs of identical specification host two workloads. The first VM emulates the radio and RAN services as separate containers (i.e., RU, vDU, vCU-UP, vCU-CP), and the other VM emulates the router and core services (i.e., ethernet-router, core-network, near-RT RIC). mTLS is enabled within the network for all communications among services. Traffic generation follows the chain RU → vDU → {vCU-UP, vCU-CP} → ethernet-router → {core-network, near-RT RIC} with services issuing requests periodically to simulate real-life traffic.

To demonstrate the feasibility of our hybrid ZT model, in this PoC, we show two examples of authorisation security policy enforcement. Let this security policy be `AuthorisationPolicy` where we can layer and enforce different controls within different levels of the configuration. To clarify this concept, in the following, we explain two examples Policy 1. `AuthorisationPolicy: allow-ru-to-vdu` and Policy 2. `AuthorisationPolicy: deny-bad-cidr`, which are written in YAML. Policy 1 is scoped to the `ran` namespace (i.e., enclave) and vDU workload. This policy permits traffic only from the RU's SPIFFE¹ identity `spiffe://cluster.local/ns/radio/sa/ru-service-account` to talk to vDU service. By default, everything else is blocked.

On the other hand, Policy 2 below shows how network isolation can be achieved at the Container Network Interface (CNI) layer, then refined and applied at the proxy service layer via using the `Gateway` and `VirtualService` resources. First, the `Gateway` resource acts as an ingress proxy that allows HTTP traffic on port 80 only. Then, `VirtualService` steers incoming requests to the core-network service. Finally, the `deny-bad-cidr` policy blocks any IP in the specified CIDR.

Policy 1. AuthorisationPolicy: allow-ru-to-vdu

```
apiVersion: security.istio.io/v1beta1
kind: AuthorisationPolicy
metadata:
  name: allow-ru-to-vdu
  namespace: ran
spec:
  selector:
    matchLabels:
      app: vdu          # apply to the vDU workload only
  action: ALLOW
```

¹ SPIFFE – Secure Production Identity Framework For Everyone. Available at: <https://spiffe.io/>

```

rules:
- from:
  - source:
    principals:
    - spiffe://cluster.local/ns/radio/sa/ru-service-account

```

Policy 2. AuthorisationPolicy: deny-bad-cidr

```

# 1) Expose the ingress gateway
apiVersion: networking.istio.io/v1beta1
kind: Gateway
metadata:
  name: public-gw
  namespace: istio-system
spec:
  selector:
    istio: ingressgateway
  servers:
  - port:
    number: 80
    name: http
    protocol: HTTP
    hosts:
    - "*"
---
# 2) Route all "/" traffic to core-network
apiVersion: networking.istio.io/v1beta1
kind: VirtualService
metadata:
  name: route-to-core
  namespace: istio-system
spec:
  hosts:
  - "*"
  gateways:
  - public-gw
  http:
  - match:
    - uri:
      prefix: /
    route:
    - destination:
      host: core-network.core.svc.cluster.local
      port:
        number: 80
  ---
# 3) Block traffic from a bad CIDR
apiVersion: security.istio.io/v1beta1
kind: AuthorisationPolicy
metadata:
  name: deny-bad-cidr
  namespace: istio-system

```

```

spec:
  selector:
    matchLabels:
      istio: ingressgateway
  action: DENY
  rules:
    - from:
      - source:
          ipBlocks:
            - 192.0.2.0/24 # replace with the CIDR to block

```

Table 2 shows the attribute mapping that is used in our PoC. Since it is based on NIST’s PM, it defines security rules as logical expressions over user attributes (UAs), object attributes (OAs), and the requested operation. The PE evaluates these attributes at run-time and issues a single allow or deny verdict. Hence, in our PoC, every service account is treated as a UA, and each workload label such as `app=vdu` or `namespace=ran` is treated as an OA. At runtime, the proxy calls the PM via the standard `ext_authz` hook, passing the caller’s SPIFFE identity and the callee’s fully qualified domain name. The PM looks up the caller’s UA set, matches the callee to its OA set, evaluates the rule logic, and returns allow or deny. Proxies enforce decisions immediately, so no packet enters the target enclave unless the UA/OA rule permits the requested action. This delivers dynamic attribute-based access control which is fully aligned with the NIST PM framework.

Table 2. Attribute Mapping Table

Caller UA (SPIFEE)	Target (namespace/app)	Operation	Decision
<code>spiffe://.../ns/radio/sa/ru-service-account</code>	<code>ran, app: vdu</code>	any	Allow
any other principal (<i>implicit</i>)	<code>ran, app: vdu</code>	any	Deny
any source IP \in 192.0.2.0/24	<code>istio-system ingressgateway</code>	any	Deny
any source IP \notin 192.0.2.0/24 (<i>implicit</i>)	<code>istio-system ingressgateway</code>	any	Allow

Note that our PoC considered a single, centralised PA. While this consideration is advantageous for easier management, simplicity and auditability, there is a risk of a throughput bottleneck once requests or policies increase to hundreds or thousands. To mitigate this issue, we considered two options. The first one involves proxy level policy caching where the PA periodically pushes rule bundles to each proxy. This enables the proxy to decide locally on requests it has already seen, hence reducing the stress on the PA. However, this would need a robust policy update and revocation strategy. The other mitigation strategy is to have a central PA, and multiple PA instances, each responsible for a subset of enclaves. This would remove the single bottleneck point at the expense of extra infrastructure and more complex management.

5. Performance & Trade-Off Analysis of The Hybrid ZT Deployment Model for O-RAN

When applying ZT, it is essential to consider the potential trade-offs in performance and user experience alongside security [39]. It is worth noting that, since no substantial data is yet available from our work and PoC implementation, we discuss similar studies for baseline comparison in terms of latency and CPU/Memory usage. In future work, we will measure end-to-end loop latencies on an O-RAN Software Community (OSC) testbed to validate that our hybrid ZT model added per-hop latency ($\approx 1\text{--}10$ ms) does not exceed O-RAN control-loop constraints.

5.1. Baseline Performance from Published ZT Approaches

5.1.1 Latency

In [40], the authors applied an enclave model to a network and tested its performance. Their findings show that an enclave-based ZT model has trivial communication overhead during connection setup time between the client and agent, approximately 1% overhead. The enclave/IPsec adds ~1–10 ms overhead. This minimal overhead is attributed to the required policy rules in the firewall and the decryption of packets before connection acceptance.

The performance overhead was found to be small and tolerable in the enclave-based deployment applied in the following network [41]. When compared to traditional security frameworks, the enclave-based model in ZT architecture offers enhanced security with minimal performance trade-offs. Traditional models often involve complex key management processes and longer recovery times, which can significantly impact system performance. In contrast, the enclave-based model maintains a balance by providing robust security measures with only a slight increase in overhead, which is deemed tolerable given the substantial security benefits achieved.

Existing studies show that sidecar proxies can increase request latency by up to 166% in default service mesh configurations in default mode [42]. However, Istio's new 'ambient' mode nearly eliminated that overhead. This suggests a potential performance issue hence, in our proposed model, agents/gateways should avoid unnecessary hops or use kernel-bypass techniques to minimise delay.

5.1.2 CPU/Memory Usage

The researchers in [43] tested an enclave model in a multi-cloud environment. Their results showed a minimal impact on cluster CPU and memory usage. Segmenting the network using enclaves as microparameters within a ZTA allows for granular security controls and robust data protection while maintaining business productivity [44]. This approach ensures that security measures are precise, flexible, and adaptable to changing business needs, providing a secure foundation for digital transformation and operational efficiency. In terms of device sandboxing, the authors in [45] showed that applications running in isolated compartments can be deployed in such a way that it will have low impact on overall system performance [45]. Trusted execution environments were used to achieve ZT with minimal performance overhead in 5G core networks in [46]. A comparison between having no security, other security models and ZT application sandboxing showed that application sandboxing had less impact on network and disk access performance in comparison to other security models [47].

While introducing a ZT deployment model may incur additional overheads and costs, it is possible to optimise the deployment to minimise effects. In [48], the authors enhanced the verification flexibility in the PEP, saving verification cost and optimising the process. ZT deployment models enhance network security but can introduce performance overheads, such as increased CPU and memory usage, and potential latency issues. However, strategic configurations/techniques like policy caching, hardware acceleration (e.g., TLS offload), and selective enforcement on critical flows only can mitigate these effects, balancing security with performance improvements.

5.2. Trade-offs of Our Hybrid ZT Model

Although our proposed hybrid ZT deployment model provides comprehensive security coverage for all O-RAN interfaces, it also introduces specific performance and operational trade-offs that operators should consider.

5.2.1 Latency Overhead vs. Granular Isolation

By combining enclave-based segmentation (i.e., macro isolation) with application sandboxing (i.e., micro isolation), our hybrid ZT model can tightly contain lateral

movement at both inter-enclave and intra-enclave levels. This multi-layered isolation makes it much harder for an attacker to move laterally from one O-RAN function to another (e.g., from xApp to near-RT RIC, or from O-DU to O-CU). However, each additional boundary crossing (e.g., enclave gateway hops or sandbox entry/exit) adds small non-zero latency. In practice, enclave/IPsec handshakes incur $\approx 1\text{--}10$ ms extra per new connection [40], and sidecar proxies (e.g., Istio's default model) can increase request latency up to 166 % before tuning [42]. In a 6G network where sub-10 ms control loops (e.g., E2, O1) are critical, careful placement and optimisation of proxies (e.g., using ambient mode or kernel bypass) are needed to keep total round-trip time within the 10 ms limit.

5.2.2 CPU/Memory Overhead vs. Scalable Policy Enforcement

Istio sidecar proxies combined with the NIST PM enforce attribute-based policies at runtime. This allows a single PDP to push fine-grained, per-session rules into each enclave or sandbox without modifying xApp/rApp code. On the other hand, running a proxy in front of each pod (i.e., enclave or sandbox) consumes additional CPU cycles (e.g., for mTLS/TLS handshakes) and memory. Although the authors in [43] showed enclave models add only a few percentage points to CPU/memory usage under moderate load, at large scale (e.g., hundreds of pods), the cumulative overhead becomes significant. Mitigation strategies such as policy caching in proxies or splitting the PDP into regional instances can reduce control-plane bottlenecks but increase management complexity.

5.2.3 Operational Complexity vs. End-to-End Visibility

Via Kiali/Prometheus, our hybrid ZT model provides end-to-end telemetry of all traffic flows across enclaves and sandboxes, enabling real-time assessment of every O-RAN component. Operators can trace precisely which UA/OA attributes triggered a deny or allow decision, which is vital for dynamic policy adjustments. Nonetheless, designing, deploying, and maintaining multiple logical enclaves plus sandboxed xApps/rApps requires upfront configuration effort such as defining namespaces, service accounts, SPIFFE identities, policy rules, and inter-pod network policies. This demands a steeper operational learning curve and higher DevOps/CNI skillset.

5.2.4 Single PDP vs. Policy Scalability

A central PDP simplifies auditability and ensures a single source of truth for all ZT policies across O-RAN domains. However, as the number of enclaves and sandboxes grows, the central PDP can become a throughput bottleneck. Caching strategies or sharing the PDP into multiple instances (e.g., each responsible for a subset of enclaves) can mitigate this but at the cost of added network traffic for policy propagation and more complex policy synchronisation. Considering these trade-offs, it can be noted that while our proposed hybrid ZT model maximises security coverage and policy granularity, it requires careful tuning and capacity planning. This is especially important for ultra-low-latency O-RAN loops and large-scale, multi-tenant deployments. Table 3 highlights how each benefit (e.g., fine-grained isolation) pairs with its corresponding cost (e.g., CPU overhead), guiding network architects in making informed implementation and deployment decisions.

5.2.5 Power & Bandwidth Consideration

In practical deployments of O-RAN, especially in edge and rural environments, power efficiency and operational bandwidth constraints remain critical factors. As highlighted by Nadeem *et al.* [51], narrowband communication protocols such as NB-IoT require lightweight and energy-efficient security mechanisms to support long device lifespans and preserve channel availability. Similarly, in O-RAN, where DUs, RUs, and edge components may operate under tight bandwidth budgets, different options are available. For instance, introduce lightweight PEP agents that implements minimal-

overhead policy enforcement (e.g., DTLS instead of mTLS). Moreover, the PDP can issue long-lived access tokens for low-power devices to reduce frequent re-authentication, thereby saving power. However, this must be balanced against security requirements, as longer token lifetimes may increase vulnerability if tokens are compromised. In future work, we will explore PDP-driven energy-security trade-off algorithms for massive IoT deployments, extending the pairing concepts in [51] to ZT policy orchestration.

Table 3. Hybrid ZT Deployment Model – Advantages vs. Limitations.

Feature	Advantages	Limitations
Isolation Granularity	<ul style="list-style-type: none"> Macro-level enclave segmentation for inter-function isolation (SMO, RIC, DU/CU, O-Cloud). Micro-level sandboxing for xApp/rApp isolation. 	<ul style="list-style-type: none"> Each new boundary hop (gateway or sandbox) adds latency ($\approx 1\text{--}10$ ms per connection). Potential for small but cumulative delay in sub-10 ms control loops (e.g., E2).
Policy Enforcement	<ul style="list-style-type: none"> Attribute-based decisions via NIST Policy Machine allow dynamic, per-session rules. Uniform enforcement across all O-RAN interfaces. 	<ul style="list-style-type: none"> Sidecar proxies (e.g., Istio) introduce CPU cycles for mTLS/TLS handshakes (benchmarks show up to +166 % latency in default mode). PDP can become bottleneck at scale. Extreme network conditions robustness causing potential policy decision latency
Security Coverage	<ul style="list-style-type: none"> Covers all O-RAN interfaces (A1, E2, O1, O2, Y1, OpenFH). Multi-layered defence (IPSec at network layer, mTLS at application layer). 	<ul style="list-style-type: none"> Requires careful key/certificate management for IPSec + mTLS across 100s of pods. Increased attack surface on the PDP and control channel if not properly hardened.
Operational Visibility	<ul style="list-style-type: none"> Fine-grained telemetry via Kiali/Prometheus for every UA\rightarrowOA decision. Real-time posture assessment of enclaves and sandboxes. 	<ul style="list-style-type: none"> Higher operational overhead to configure and interpret monitoring data. DevOps skillset required for Kubernetes/Istio/CNI management.
Scalability/Manageability	<ul style="list-style-type: none"> Enclaves can be scaled independently (horizontal pod autoscaling). Policy caching or PDP sharding can improve performance. 	<ul style="list-style-type: none"> Increased complexity: defining namespaces, service accounts, and SPIFFE identities for each enclave. PDP sharding adds synchronization and orchestration overhead.
Performance Overhead	<ul style="list-style-type: none"> Minimal additional latency in well-tuned ambient-mode service mesh or kernel-bypass IPSec. CPU/memory overhead limited ($\sim 5\text{--}10$ %). 	<ul style="list-style-type: none"> Out-of-the-box service mesh modes may incur sizable latency spikes unless optimised. Memory footprint of Istio sidecars ($\sim 100\text{--}200$ MB per pod).

5.3. Adapting Security Policies under Non-Ideal Channel Conditions

In key 6G use cases with mission-critical applications such as Industrial IoT (IIoT), the accuracy of Channel State Information (CSI) directly impacts the efficiency of dynamic security policy enforcement. Non-ideal channel conditions (e.g., due to interference, noise, or sparse multipath propagation common in industrial settings) can degrade the performance of security protocols that rely on timely communication, such as policy updates or authentication exchanges. To address this challenge, we propose integrating the efficient channel estimation techniques proposed by Wang *et al.* [52] into our model. The authors developed a low-complexity sparse Multiple-input-multiple-output filter bank

multicarrier (MIMO-FBMC) channel estimation scheme tailored for sparse industrial channels, achieving higher estimation accuracy and reduced computational complexity compared to traditional methods. Reliable CSI enables our hybrid ZT model to dynamically adjust security policies in real time as follows. The PDP can consider real-time channel quality metrics (e.g., signal-to-noise ratio, delay spread) as an additional attribute in policy decisions. For example, if the channel quality drops below a threshold, the PDP may temporarily grant extended session durations to avoid frequent re-authentication that would otherwise exacerbate delays. On the other hand, the PEPs can switch between security protocols based on channel conditions. For instance, during periods of high interference, they might use lighter but still secure cryptographic algorithms (e.g., AES-128 instead of AES-256) to reduce processing and transmission time, ensuring that the control loop deadlines are met. These adaptations ensure that our hybrid ZT deployment model maintains robust security without compromising the operational requirements of IIoT applications, even under non-ideal channel conditions.

5.4. Robustness in Extreme Network Conditions

While our PoC demonstrates the feasibility and effectiveness of the proposed ZT deployment model in standard operational settings, extreme radio conditions such as high user mobility (e.g., vehicular scenarios), and strong interference (e.g., industrial settings) may introduce channel estimation errors that may impact the enforcement of security policies. As part of future work, we plan to extend our validation with performance simulations of the following scenarios to validate the robustness of our ZT mode model.

- High Mobility: simulate UE moving at speeds up to 120 km/h in an urban environment, causing rapid channel variations and beam misalignment. We will measure policy decision latency and impact on PEP session establishment success rate.
- Strong Interference: introduce controlled interference sources (e.g., at 2.4/5 GHz bands) to SINR levels from 0 to 10 dB. We will measure channel estimation error impact on policy update delays and security policy violation rates (e.g., expired tokens due to delayed PDP responses). Channel estimation errors will be mitigated by integrating robust algorithms such as MIMO-FBMC sparse channel estimation [52] and the training channel-based method for mmWave systems [53].

These simulations will be conducted using the OSC testbed with hardware-in-the-loop emulation of adverse RF conditions.

6. Conclusion & Future Work

Securing O-RAN in 6G networks requires moving beyond perimeter-centric defences toward a Zero Trust paradigm. This paper presented the first dedicated ZT deployment model for O-RAN that integrates enclave-based segmentation with application sandboxing, mapped comprehensively to all O-RAN interfaces and aligned with NIST ZTA tenets. By centralising policy decision via a PDP and distributing enforcement across logical enclaves, agents, and gateways, our hybrid ZT deployment model secures all O-RAN interfaces using industry-standard controls such as mTLS, IPsec, and NACM. The Kubernetes/Istio PoC with the NIST PM validates that pods can act as logical enclaves, sidecar proxies enforce per-session, attribute-based policies, and centralised PDP/PE architecture remains practical. Through trade-off analysis, which is summarised in Table 3, operators can understand latency overheads, resource costs, and operational complexity versus security benefits. We also supply a high-level implementation plan covering preparation, design, deployment, testing, and continuous monitoring stages.

For future work, we will extend our ZT deployment model to address the challenges of millimetre-wave (mmWave) based O-RAN deployments. The adoption of mmWave technology in 6G introduces unique channel characteristics (e.g., high path loss, narrow

beams, and rapid spatial variations) that impact security enforcement. The rapid channel variations and potential beam misalignment could affect security deployment, particularly in terms of maintaining continuous authentication and policy enforcement. A promising research direction will be incorporating real-time channel estimation to optimise security policies and maintain reliability (e.g., the method proposed in [53]). Furthermore, future work will include quantitative evaluation on an O-RAN testbed (e.g., OSC) to measure the robustness of our ZT model under extreme network conditions (e.g., high mobility, interference), end-to-end latency and throughput under ZT controls, experimentation with distributed PDP shards, as in [47], to alleviate control-plane bottlenecks, and automated conflict resolution for concurrent xApp/rApp policies. Moreover, we will explore AI-trustworthiness techniques (e.g., explainable AI for policy decisions) and federated learning to preserve privacy across multi-vendor xApps. Ultimately, by integrating our comparison tables and architecture figures, this paper delivers a holistic blueprint for deploying Zero Trust in O-RAN toward full Zero Trust maturity.

Author Contributions: Conceptualization, M.H.E. and B.K.; methodology, M.H.E. and B.K.; software, B.K.; validation, M.H.E. and B.K.; formal analysis, M.H.E. B.K., A.R., M.M. and M.K.M; writing—original draft preparation, M.H.E. and B.K.; writing—review and editing, A.R., M.M. and M.K.M; All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Department for Science, Innovation and Technology (DSIT), UK Government, Liverpool City Region HDD (High Demand Density) project.

Data Availability Statement: Data is unavailable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. D. Chen Larsson, A. Grövlén, S. Parkvall and O. Liberg, “6G standardization – an overview of timeline and high-level technology principles,” 22 March 2024. [Online]. Available: <https://www.ericsson.com/en/blog/2024/3/6g-standardization-timeline-and-technology-principles>.
2. O-RAN ALLIANCE e.V., “O-RAN Specifications,” [Online]. Available: <https://www.o-ran.org/specifications>. [Accessed 10 September 2024].
3. 6G Smart Networks and Services Industry Association (6G-IA), “Open RAN and 6G Future Networks Development,” 6G SNS IA, 2024.
4. O-RAN next Generation Research Group (nGRG), “O-RAN Towards 6G Report ID: RR-2023-01,” O-RAN Alliance, 2023.
5. O-RAN next Generation Research Group (nGRG), “Architecture principles for a cloud-friendly future 6G RAN architecture,” ORAN Alliance, 2024.
6. US Government, “National Cybersecurity Strategy,” The White House, Washington, 2023.
7. National Cyber Security Centre (NCSC), “Zero trust architecture design principles,” NCSC, 2021.
8. O-RAN ALLIANCE, “Zero Trust Architecture for Secure O-RAN v1.0,” O-RAN Alliance, May 2024.
9. National Institute of Standards and Technology (NIST), “Advanced Security Architectures for Next Generation Wireless,” 15 April 2024. [Online]. Available: <https://www.nist.gov/programs-projects/advanced-security-architectures-next-generation-wireless>. [Accessed 13 September 2024].
10. S. Rose, O. Borchert, S. Mitchell and S. Connelly, “Zero Trust Architecture NIST Special Publication 800-207,” NIST, 2020.
11. NSA & CISA, “Security Guidance for 5G Cloud Infrastructures, Volumes 1 thru 4, Enduring Security Framework (ESF),” US DHS CISA, October-November 2021.
12. US DHS CISA, “Zero Trust Maturity Model (ZTMM), version 2.0,” CISA, April 2023.
13. O-RAN ALLIANCE, “O-RAN Architecture Description v12.0,” O-RAN Alliance, 2024.

14. O-RAN Alliance, "The O-RAN ALLIANCE Security Working Group Continues to Advance O-RAN Security," 09 Feb 2024. [Online]. Available: <https://www.o-ran.org/blog/the-o-ran-alliance-security-working-group-continues-to-advance-o-ran-security>. [Accessed 27 Sept 2024].
15. M. Liyanage, A. Braeken, S. Shahabuddin and P. Ranaweera, "Open RAN security: Challenges and opportunities," *Journal of Network and Computer Applications*, vol. 214, no. 103621, 2023.
16. E.A. J. Groen, "Securing O-RAN Open Interfaces," *IEEE Transactions on Mobile Computing*, vol. Early Access, pp. 1-13, 2024.
17. O-RAN Alliance e.V., "O-RAN Security Threat Modelling and Risk Assessment 3.0," O-RAN Alliance, 2024.
18. O-RAN Alliance, "O-RAN Security Requirements and Controls Specification 9.0," O-RAN Alliance, 2024.
19. O-RAN Alliance e.V., "O-RAN Security Protocols Specifications 9.0," O-RAN Alliance, 2024.
20. O-RAN Alliance e.V., "O-RAN Security Test Specifications 7.0," O-RAN Alliance, 2024.
21. O-RAN Alliance e.V., "O-RAN Study on Security for Service Management and Orchestration (SMO) 4.0," O-RAN Alliance, 2024.
22. O-RAN ALLIANCE WG11 for Security V1.0, "Zero Trust Architecture for Secure O-RAN," O-RAN Alliance, 2024.
23. A. Abdalla, J. Moore, N. Adhikari and V. Marojevic, "ZTRAN: Prototyping Zero Trust Security xApps for Open Radio Access Network Deployments," *IEEE Wireless Communications*, vol. 31, no. 2, pp. 66-73, 2024.
24. H. Jiang, H. Chang, S. Mukherjee and J. Van der Merwe, "OZTrust: An O-RAN Zero-Trust Security System," in *IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, Dresden, Germany, 2023.
25. Z. Zaheer, H. Chang, S. Mukherjee and J. Van der Merwe, "eZTrust: Network-Independent Zero-Trust Perimeterization for Microservices," in *Proceedings of the 2019 ACM Symposium on SDN Research (SOSR '19)*, San Jose, CA USA, 2019.
26. The Kubernetes Authors, "Using RBAC Authorization," 28 June 2024. [Online]. Available: <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>. [Accessed 11 October 2024].
27. The Cilium Authors, "Cilium - Cloud Native eBPF-based Networking, Observability, Security," [Online]. Available: <https://cilium.io/>. [Accessed 11 October 2024].
28. Tigera, Inc., "Calico," [Online]. Available: <https://docs.tigera.io/calico/latest/about/>. [Accessed 11 October 2024].
29. X. Chen, W. Feng, N. Ge and Y. and Zhang, "Zero Trust Architecture for 6G Security," *IEEE Network*, vol. 38, no. 4, pp. 224-232, 2024.
30. I.-A. Dumitru, "Zero Trust Security," in *International Conference on Cybersecurity and Cybercrime*, Bucharest, Romania, 2022.
31. Y. He, D. Huang, L. Chen, Y. Ni and X. Ma, "A Survey on Zero Trust Architecture: Challenges and Future Trends," *Wireless Communications and Mobile Computing*, vol. 2022, p. 13 pages, 2022.
32. N. Basta, M. Ikram, M. Kaafar and A. Walker, "Towards a Zero-Trust Micro-segmentation Network Security Strategy: An Evaluation Framework," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*, Budapest, Hungary, 2022.
33. N. Sheikh, M. Pawar and V. and Lawrence, "Zero trust using Network Micro Segmentation," in *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, Vancouver, BC, Canada, 2021.
34. N. Mhaskar, M. Alabbad and R. Khedri, "A Formal Approach to Network Segmentation," *Computers & Security*, vol. 103, 2021.
35. NIST, Computer Security Resource Centre (CSRC), "Policy Machine | CSRC," 06 June 2016. [Online]. Available: <https://csrc.nist.gov/Projects/Policy-Machine>. [Accessed 10 Feb 2025].
36. The Kubernetes Authors, "Kubernetes," 2025. [Online]. Available: <https://kubernetes.io/>. [Accessed 2024].
37. Istio Authors, "The Istio service mesh," 2025. [Online]. Available: <https://istio.io/latest/>. [Accessed 2024].
38. Kiali, "Kiali - The Console for Istio Service Mesh," 2025. [Online]. Available: <https://kiali.io/>. [Accessed 2024].
39. B. Paul and M. and Rao, "Zero-Trust Model for Smart Manufacturing Industry," *Applied Sciences*, vol. 13, no. 1, 2023.
40. F. Ruambo, D. Zou, I. lopes and B. and Yuan, "Securing SDN/NFV-Enabled Campus Networks with Software-Defined Perimeter-Based Zero-Trust Architecture," *SSRN*, <http://dx.doi.org/10.2139/>, 2023.
41. Y. Bello, A. Hussein, M. Ulema and J. and Koilpillai, "On Sustained Zero Trust Conceptualization Security for Mobile Core Networks in 5G and Beyond," *IEEE Transactions on Network and Service Management*, vol. 19, no. 2, pp. 1876-1889, 2022.
42. A. Barr, O. Lavi, Y. Naor, S. Rampal and J. and Tavori, "Technical Report: Performance Comparison of Service Mesh Frameworks: the MTLs Test Case," *arXiv*, <https://arxiv.org/html/2411.02267v1>, 2024.
43. S. Rodigari, D. O'Shea, P. McCarthy, M. McCarry and S. McSweeney, "Performance Analysis of Zero-Trust multi-cloud," in *IEEE 14th International Conference on Cloud Computing (CLOUD)*, Chicago, IL, USA, 2021.

44. C. Cunningham, D. Holmes and J. Pollard, "The Eight Business and Security Benefits of Zero Trust Business Case: The Zero Trust Security Playbook," Forrester Research, <https://www.kennisportal.com/wp-content/uploads/2022/06/Akamai-the-eight-business-and-security-benefits-of-zero-trust-report.pdf>, 2019. 1143 1144 1145
45. F. Brasser, D. Gens, P. Jauernig, A. Sadeghi and E. and Stapf, "SANCTUARY: ARMing TrustZone with User-space Enclaves," in Network and Distributed Systems Security (NDSS) Symposium, San Diego, CA, USA, 2019. 1146 1147
46. M. Vomvas, N. Ludant and G. and Noubir, "Establishing Trust in the Beyond-5G Core Network using Trusted Execution Environments," arXiv, <https://arxiv.org/abs/2405.12177>, 2024. 1148 1149
47. J. Zhang, J. Zheng, Z. Zhang, T. Chen, K. Qiu, Q. Zhang and Y. and Li, "Hybrid isolation model for device application sandboxing deployment in Zero Trust architecture," International Journal of Intelligent Systems, no. <https://doi.org/10.1002/int.23037>, 2022. 1150 1151 1152
48. D. Zheng, H. Xing, X. Cao and J. and Xu, "Efficient Zero-Trust-enabled Service Function Chain Deployment in Multi-Vendor Networks," TechRxiv, <https://doi.org/10.36227/techrxiv.171778732.29298095/v1>, 2024. 1153 1154
49. M. Polese, L. Bonati, S. D'Oro, S. Basagni and T. Melodia, "Understanding O-RAN: Architecture, Interfaces, Algorithms, Security, and Research Challenges," IEEE Communications Surveys & Tutorials, vol. 25, no. 2, pp. 1376-1411, 2023. 1155 1156
50. L. Alevizos, V. Ta, and M. Hashem Eiza, "Augmenting Zero Trust Architecture to Endpoints Using Blockchain: A State-Of-The-Art Review," Security and Privacy, vol. 5, no. 1, 2022. 1157 1158
51. A. Nadeem, M. Hussain, A. Iftikhar and S. Aslam, "Narrowband IoT Device to Device Pairing Scheme to Save Power," in Proc of IEEE 23rd International Multitopic Conference (INMIC), Bahawalpur, Pakistan, pp. 1-5, 2020. 1159 1160
52. H. Wang, L. Xu, Z. Yan, and T.A. Gulliver, "Low-Complexity MIMO-FBMC Sparse Channel Parameter Estimation for Industrial Big Data Communications," IEEE Transactions on Industrial Informatics, vol. 17, no. 5, pp. 3422–3430, May 2020. 1161 1162
53. H. Wang, P. Xiao, and X. Li, "Channel Parameter Estimation of mmWave MIMO System in Urban Traffic Scene: A Training Channel-Based Method." IEEE Transactions on Intelligent Transportation Systems, vol. 25, no. 1, pp. 754 – 762, Jan 2024. 1163 1164