Received 5 June 2024; revised 8 July 2024; accepted 16 July 2024. Date of publication 22 July 2024; date of current version 5 August 2024.

Digital Object Identifier 10.1109/OJCOMS.2024.3431681

Attacking O-RAN Interfaces: Threat Modeling, Analysis and Practical Experimentation

PAU BAGUER[1], GIRMA M. YILMA[2], ESTEBAN MUNICIO[1], GINES GARCIA-AVILES[1], ANDRES GARCIA-SAAVEDRA[2], MARCO LIEBSCH[2] (Member, IEEE), AND XAVIER COSTA-PÉREZ[1,2,3] (Senior Member, IEEE)

[1]AI-Driven Systems, i2CAT Foundation, 08034 Barcelona, Spain

[2]NEC Laboratories Europe GmbH, 69115 Heidelberg, Germany

[3]Institucio Catalana de Recerca i Estudis Avancats, 08010 Barcelona, Spain

CORRESPONDING AUTHOR: E. MUNICIO (e-mail: esteban.municio@i2cat.net)

This work was supported in part by the European Commission through the BeGREEN Project under Grant 101097083; in part by the ORIGAMI Project under Grant 101139270; in part by the Spanish Ministry of Economic Affairs and Digital Transformation and the European Union NextGeneration EU through the Framework of the Recovery Plan, Transformation and Resilience (PRTR) (Call UNICO I+D 5G 2021) under Grant TSI-063000-2021-3-Open6G; in part by the CDTI and the European Union NextGenerationEU/PRTR within the Call "Ayudas Cervera para Centros Tecnológicos 2023" under Grant 6G-DIFERENTE (CER-20231018); in part by the CERCA Programme from the Generalitat de Catalunya; and in part by the SPIRS Project under Grant 952622.

ABSTRACT A new generation of open and disaggregated Radio Access Networks (RANs) enabling multi-vendor, flexible, and cost-effective deployments is being promoted by the Open Radio Access Network (O-RAN) Alliance. However, this new level of disaggregation in the RAN also entails new security risks that must be carefully addressed. The O-RAN Alliance has established Working Group 11 (WG11) to ensure that the new specifications are secure by design. Acknowledging the new security challenges arising from the expanded threat surface, O-RAN WG11 provides procedures to identify threats and assess and mitigate risks. Reportedly, as of 2024, 60% of found risks are related to Denial of Service (DoS) and performance degradation. Therefore, in this work, we analyse a vanilla O-RAN deployment and evaluate the endurance of different O-RAN interfaces under attacks in scenarios involving DoS and performance degradation. To do so, we use a reference O-RAN open source deployment to report, risks found, weak points, and counter-intuitive recommended design choices for both control plane (A1, E2, and F1-c) and user plane (F1-u) interfaces. Consequently, we map O-RAN WG11's threat model and risk assessment methodology to our considered DoS and performance degradation scenarios, and dissect existing threats and potential attacks over O-RAN interfaces that may compromise the security of O-RAN architectural deployments. Finally, we identify mechanisms to mitigate risks and discuss approaches aimed at improving the robustness of future O-RAN networks.

INDEX TERMS 5G, denial-of-service attacks, O-RAN, security.

I. INTRODUCTION

CURRENT mobile networks use novel technological concepts such as Software-Defined Networking (SDN), Network Function Virtualization (NFV), Multi-access Edge Computing (MEC), and public/private clouds to operate their services for billions of customers and trillions of devices [1]. However, making sure these technologies are secured is still a day-to-day challenge. Until recently, the approach for mobile network security has been based on risk analysis rather than incorporating security as a design element, leading to a number of potential vulnerabilities that could be exploited.

Open Radio Access Network (O-RAN) is the latest arena in the virtualization of network functions for 5G and beyond ecosystems, which is gaining significant momentum because of the support from both Mobile Network Operators (MNOs) and vendors [2]. O-RAN activities are led by the O-RAN Alliance, initially founded by AT&T, China Mobile, Deutsche Telekom, NTT DOCOMO, and Orange.

Currently, O-RAN is actively supported by more than 335 companies including academia, major cloud providers, and startups. O-RAN builds on top of 3GPP's specified Radio Access Network (RAN), by defining an open architecture and interfaces for the RAN space, decoupling hardware and software to foster innovation and competition, and running RAN network functions on a shared cloud infrastructure, which leverages virtualization to reduce CAPEX and OPEX.

Recently, at a hacker conference held in the Netherlands, a team of hackers breached live 5G networks in a series of "red teaming" exercises. The attacks were primarily directed to poorly configured "containers" and managed to demonstrate the likelihood of turning down a Kubernetes-based 5G core [3]. Such occurrences pose a significant challenge as operators are not yet equipped to properly secure and manage cloud workloads. Since O-RAN is also expected to run in similar environments, it is crucial to bridge the technical gaps required to securely operate virtualized workloads in a public/private cloud for the secure operation of telco services. This is especially relevant in O-RAN, since the disaggregated nature of its architecture adds new interfaces and functions, expanding the threat surface and bringing new security challenges [4].

Traditionally, 3GPP standards focused on service, protocol and operational security whereas deployment-specific security aspects were treated out-of-scope. Attempts to address such security aspects in subsequent versions of the standard have been laborious and often resulted in partially unsafe or complex solutions [5]. One notable example coming from the 2G era is the vulnerability in the A5/1 encryption algorithm. Initially designed to secure communication, A5/1 was later found to be vulnerable to various cryptography attacks. The response to this vulnerability was the introduction of A5/3, based on the more robust KASUMI algorithm. However, this fix was challenging to implement due to the wide already existing deployment and the need for backward compatibility (see 3GPP TS 55.216 [6]). In the 3G transition, 3GPP introduced mutual authentication and stronger encryption algorithms such as KASUMI for UEA1. Again, and despite these improvements, vulnerabilities were identified in the integrity protection mechanism using the f9 algorithm. The response involved revisiting the integrity algorithms and adding more secure options such as SNOW 3G in UEA2. Accordingly, such threats have been recognised and addressed throughout all releases. Finally, about a year ago, the first version of 3GPP TS 33.527 [7] came out. The document goes beyond the so far treated security aspects and addresses the Security Assurance Specification (SCAS) for 3GPP Virtualized Network Products to ensure the same level of security as non-virtualized physical network nodes.

These examples underscore the inherent challenges in addressing vulnerabilities reactively within the 3GPP framework. The transition from GSM to LTE and beyond has shown that while improvements are continuously made, the initial exclusion of broader security considerations needed complex and often delayed fixes. This historical context reinforces the importance of a proactive, security-by-design approach adopted in later 3GPP releases, particularly with the comprehensive security measures integrated into 5G-New Radio (NR) (3GPP TS 33.501) [8].

To prevent this from happening, O-RAN prioritizes and addresses IT security from the start of the development process. To do so, the O-RAN Alliance has established Working Group 11 (WG11) to ensure that the new specifications are secure by design. WG11 defines a security analysis methodology and a threat model that identifies vulnerabilities, risks, and threats. Remarkably, to date, 60% of those identified risks by WG11 are related to Denial-of-Service (DoS) and performance degradation [9]. In fact, the O-RAN Test and Integration Focus Group (TIFG) extensively includes DoS attacks as part of the O-RAN End-to-end Test Specification [10]. Despite this, most of current works have studied O-RAN security in a generic manner, without experimental deployments, and with little emphasis on DoS and performance degradation attacks [11], [12], [13].

In this article, we aim to comprehensively review the O-RAN WG11's threat model and further unveil the security challenges of existing O-RAN reference implementations. To do so, we study DoS and performance degradation-associated threats and experimentally evaluate in a reference open-source deployment their impact by doing performance degradation stress tests on the various O-RAN interfaces, concretely the A1, E2, and F1-c/F1-u interfaces, which would be equivalent to those impacts caused by real DoS or performance degradation attacks. The rationale behind selecting these interfaces while leaving out O1, O2, E1 and Open Fronthaul (OFH), is low maturity level for O1/O2/E1 implementations and lack of support on commercial Software-Defined Radio (SDR) hardware for the case of OFH. For example, 3GPP's vRAN interface E1 has so far not been aligned with the rest of the O-RAN specifications, and O1/O2 implementations are only partially implemented, which makes them not ready to be operationally used in a production O-RAN enabled environment. We follow a top-down approach. First, we study the O-RAN WG11's threat model and its security analysis methodology, identifying those risks associated with DoS or performance degradation. Subsequently, we deploy the different O-RAN components, and identify problems and issues (sometimes also surprisingly good performance) arising when pushing the O-RAN interfaces to the limit in terms of latency and drop rate, as it would happen in a DoS attack.

Our experimental results show robustness against latency and packet loss in the A1 interface (see Fig. 1 as architecture reference). In contrast, the E2 interface was found more vulnerable to network performance degradation conditions, and therefore more sensible to DoS attacks. Unsurprisingly, the time sensitivity in the procedures carried by the F1-c interface, ascertained its sensitivity to network latency and packet loss, which results in the F1-c being significantly vulnerable to subtle DoS attacks (where the attacker may

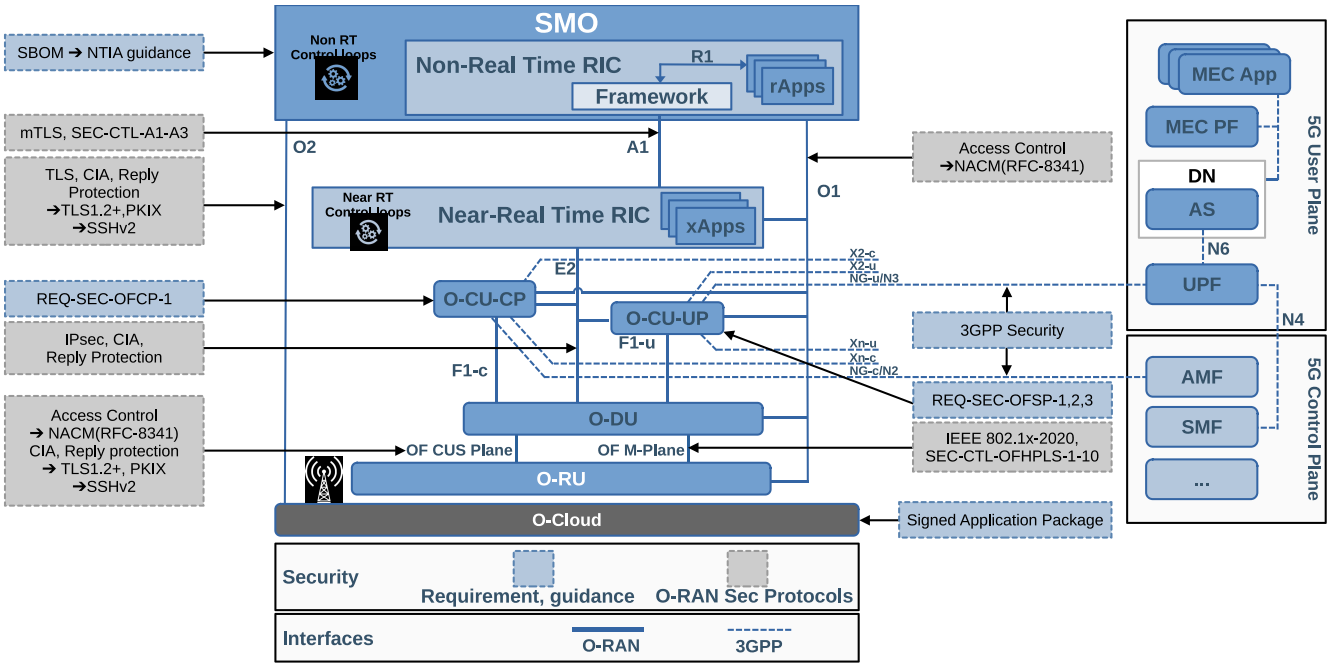

FIGURE 1. O-RAN Security requirements, guidance and selected protocols summary. Interfaces and components are discussed in detail in Section III.

remain easily unnoticed since the required disturbance required to deny the service is rather small). Finally, we also review the implementation of O-RAN interfaces and identify security risks that may arise in virtualization environments such as the O-RAN Cloud. Therefore, our main contributions are:

- We review O-RAN security analysis methodology and threat model, putting special emphasis on DoS and performance degradation threats impacting the A1, E2, and F1-c/F1-u interfaces. By doing so, we comprehensively map each identified threat to each affected interface.
- We experimentally assess the impact of the previously identified DoS and performance degradation threats to the A1, E2 and F1-c/F1-u interfaces. To do so, we perform performance degradation tests on such interfaces, identifying the resilience of the overall O-RAN deployment as well as of their associated O-RAN components.

In the following sections, we first briefly present the main activities and players driving O-RAN security standardization in Section II, and then, we follow up with a system overview of O-RAN in Section III. In Section IV, we introduce the O-RAN's security analysis methodology, which includes the identification, assessment, and treatment of security risks. Section V presents the O-RAN threat model, identifying the DoS and performance degradation attacks and the scenarios that will allow us to test their effects in the network. In Section VI, we present experimental results of how the previously selected attacks would impact the network performance. Finally, Section VII provides an overview of known vulnerabilities and security recommendations for O-RAN deployments, and Section VIII concludes this work.

II. STANDARDIZING O-RAN SECURITY

The O-RAN Alliance Security Task Group is responsible for defining O-RAN security protocol specifications, requirements, test specifications, threat modelling, and remediation analysis. Figure 1 depicts selected protocols, inherited 3GPP security protocols, security requirements, guidance, and recommendations for the different components and interfaces of O-RAN. Recently, the O-RAN specification development process moved towards the paradigm of security/privacy by design (and by default). This task group comprises several task forces for specifying and validating the different O-RAN security aspects.

During the early stages of the O-RAN specifications, the security risk analysis conducted by the German government [5] played a pivotal role in addressing security concerns. Given the significant security issues identified in O-RAN, there was a risk that O-RAN might not fulfil the security requirements of regulatory bodies, potentially delaying its deployment. To address these concerns, the German government undertook a comprehensive risk analysis and provided recommendations for corrective measures. This analysis is now being taken into consideration by the O-RAN Alliance Security Task Group.

The risk analysis of O-RAN, as well as the likelihood of the security risks, is extensively discussed in the threat model analysis of O-RAN [9] and other recent works such as [13]. Experience has shown that in any case, a late addition of security measures leads either to very high costs or too

insecure solutions. According to the authors of the O-RAN risk analysis [5], there are a number of known security measures that can be implemented in the design phase relatively easily and at low cost, which can contribute to reducing the risk associated with individual security threats.

Despite this, although ensuring the security of O-RAN is crucial for operators as it impacts reputation, trust, and compliance with regulations [14], the O-RAN Alliance acknowledges that the use of open and cloud-based architectures increases the potential attack surface of RAN systems [14], [15], and that greater transparency in new open interfaces will lead to a more thorough examination of vulnerabilities.

The European Union (EU) 5G Cybersecurity Framework Security Focus Group (SFG) [16] lays out security requirements, architectures, and frameworks to support open interfaces, such as the ones defined by other O-RAN Working Groups (WGs) [17]. This includes guidelines, protocols, and specifications that encompass the entirety of the O-RAN architecture, with the ultimate objective of ensuring that O-RAN systems are secure prior to commercial deployment [18]. Implementing security measures can be time-consuming and resource-intensive, but it is necessary to ensure the protection of sensitive information and reduce the risk of security breaches. Regular assessments to security measures can help to identify and address potential vulnerabilities before they can be exploited by malicious actors.

III. THE O-RAN ECOSYSTEM

In this section, we briefly introduce the O-RAN system architecture. Figure 1 depicts the O-RAN architecture according to the O-RAN Alliance specification [19]. The O-RAN specification implements the disaggregation of RAN components and deployments, which are based on virtualized and software-based components to implement interoperability across different vendors.

O-RAN embraces the 3GPP NR 7.2 split for base stations, which splits RAN functionalities into Central Unit (O-CU), Distributed Unit (O-DU), and Radio Unit (O-RU) [11]. Moreover, it connects these functions through open interfaces to logical intelligent controllers namely; the non-Real-Time Radio Intelligent Controller (non-RT RIC), which performs management and control of RAN functions in a non-RT RIC closed-loop (longer than 1s), and the near-Real-Time Radio Intelligent Controller (near-RT RIC), which performs management and control of RAN functions in a near real-time closed-loop (10ms to 1s) such as radio resource optimization. It also comprises the Service Management and Orchestration Framework (SMO), which is responsible for hosting the non-RT RIC and overseeing the life cycle of RAN functions. Next, we briefly summarize the main components of O-RAN.

A. O-RAN COMPONENTS

- *O-Cloud.* The O-Cloud comprises a collection of physical infrastructure computing nodes that meet O-RAN requirements to host the relevant O-RAN components such as near-RT RIC, O-CU, O-DU, and other functions. The O-Cloud also contains supporting software components (such as the Operating System (OS), Virtual Machine (VM) Monitor, Container runtime, etc.) and other management and orchestration functions.
- *SMO.* The SMO framework is responsible for handling the automation, control, management, and orchestration of RAN components including the non-RT RIC.
- *rApps & xApps.* O-RAN introduces third-party applications called rApps for network automation services (which work in the non-real-time range, i.e., in a range longer than 1s) and xApps (which work in the near-real-time range, i.e., in the range of 10ms to 1s).
- *non-RT RIC.* As depicted in Figure 1, the non-RT RIC is part of the SMO, which manages resources and the life cycle of virtualized RAN functions. The set of SMO functions devoted to the A1 interface termination and the exposure of R1 services is identified as the non-RT RIC framework. The non-RT RIC orchestrates rApps to perform tasks such as interference management, optimization, and Machine Learning (ML) applications for closed-loop operations.
- *near-RT RIC.* As depicted in Figure 1, the near-RT RIC terminates three interfaces (O1, A1, and E2), and orchestrates xApps. The xApps are expected to work on the sub-one-second level and are responsible for enforcing policies received via A1.
- *O-CU.* The O-CU hosts the Radio Resource Control (RRC) and Packet Data Convergence Protocol (PDCP) layers of the protocol stack. The O-CU communicates with the near-RT RIC to report near-real-time information and to receive radio resource management policies. Its functions are divided into the control plane (O-CU-CP) and the user plane (O-CU-UP).
- *O-DU.* The O-DU hosts the Radio Link Control (RLC) layer, the Medium Access Control (MAC) layer, and the Physical (PHY) layer, where the latter only comprises part of the PHY layer functionality namely *high-PHY* depending on the selected functional split [20].
- *O-RU.* The O-RU provides the Low-PHY layer and Radio frequency (RF) processing. It performs tasks such as the Fast Fourier Transform (FFT), cyclic prefix insertion, or precoding, and offloads the remaining functions to the O-DU.

B. O-RAN-DEFINED INTERFACES

In this section, we present the standard interfaces maintained by O-RAN:

- The **O1 interface** connects all O-RAN managed elements (MEs) to the SMO framework and enables the SMO framework to access the O-RAN network functions and the O-RAN compliant eNB (O-eNB).
- The **O2 interface** is used for running open management and orchestration services and is responsible for communication between the SMO framework and the O-Cloud platform.

- The **A1 interface** connects the non-RT RIC and near-RT RIC enabling policy-driven guidance of near-RT RIC applications/functions (xApps) and supports Artificial Intelligence and Machine Learning (AI/ML) workflows.
- The **E2 interface** connects the near-RT RIC with the E2 nodes. E2 node is a collective term for all units that are controlled by the near-RT RIC, namely O-CUs, and O-DUs and O-eNBs. The E2 interface is also responsible for streaming telemetry from the RAN and providing feedback and control from the near-RT RIC.
- *OFH* communicates O-DUs and O-RUs through the Control, User, Synchronization, and Management planes, typically using a packet-based enhanced CPRI (eCPRI) interface. The User Plane (U-Plane) and Control Plane (C-Plane) are used for the transport of data and PHY-layer control commands respectively. The Synchronisation Plane (S-Plane) takes care of frequency and phase synchronization between O-DU and O-RU clocks using protocols such as Synchronous Ethernet (SyncE) or Precision Time Protocol (PTP). Finally, the Management Plane (M-Plane) is used to configure settings in O-RUs.

C. 3GPP-DEFINED INTERFACES

O-RAN utilizes the protocol stack defined by the 3GPP and adapts it to support open interfaces:

- The **F1-c interface** is used for C-Plane communication between the O-CU Control Plane (O-CU-CP) and O-DU functions.
- The **F1-u interface** is used for transferring application data between O-DU and O-CU.
- The **NG-u interface** connects the O-CU to the 5G User Plane Function (UPF). It is used to manage the U-Plane for 5G user services and to support the efficient transfer of user data, ensuring the reliability and Quality of Service (QoS) for user data transmission.
- The **NG-c interface** connects the O-CU to the 5G Core Service Management Function (SMF) and is responsible for context and mobility managing of User Equipment (UE)s.
- Additional interfaces, including **X2-c, X2-u, Xn-c, Xn-u, and Uu**, adhere to the 3GPP protocol stack and play vital roles to communicate 5G gNBs and 4G eNBs, and to facilitate functions like handovers, dual connectivity, and load balancing.

D. REQUIREMENTS AND PROTOCOLS OVERVIEW OF O-RAN SECURITY

This section is dedicated to explaining the main concepts in O-RAN's security strategy as well as their integration and requirements within the existing 3GPP ecosystem. The O-RAN architecture, as illustrated in Figure 1, integrates robust security measures to foster an open and interoperable ecosystem. A multi-layered security approach is employed, blending O-RAN-specific protocols with the established 3GPP security standards to strengthen the security posture and ensure compatibility with existing cellular networks following security requirements [21].

1) SECURITY IMPLEMENTATION IN O-RAN

In O-RAN, security is enforced through various mechanisms:

- *Confidentiality, Integrity, and Authentication (CIA):* These fundamental security principles ensure i) preserving authorized restrictions on information access and disclosure, ii) guarding against improper information modification or destruction, and iii) ensuring timely and reliable access to and use of information (see [22]).
- *Network Access Control Model (NACM):* NACM provides a sophisticated framework for access control within network configuration protocols, harmonized with 3GPP's access control mechanisms.
- *Software Bill of Materials (SBOM):* SBOM serves as a list of components in a software build, critical for tracking vulnerabilities and ensuring compliance with standards, as recommended by the NTIA [23].

2) SECURITY PROTOCOLS THAT EXTEND ACROSS DIFFERENT OPERATIONAL PLANES WITHIN THE O-RAN STRUCTURE

1) *C-Plane:* The C-Plane leverages IEEE 802.1X-2020 for Port-based Network Access Control to authenticate and authorize communication between O-DUs and O-RUs, as specified by REQ-SEC-OFCP-1.
2) *S-Plane:* Authentication and authorization of PTP nodes are secured as per REQ-SEC-OFSP-1, with additional measures like spoofing prevention and redundancy for master clocks outlined in REQ-SEC-OFSP-2 and SEC-CTL-OFSP-1 respectively.
3) *LAN Segment Security:* Between OFH network elements, mechanisms are in place for authenticating and authorizing Local Area Network (LAN) segments (REQ-SEC-OFHPLS-1), detecting and reporting segment status (REQ-SEC-OFHPLS-2), and blocking access to unused Ethernet ports (REQ-SEC-OFHPLS-3).
4) *IEEE 802.1X-2020 Compliance:* Support for 802.1X-2020 supplicant functionality across the OFH network elements is mandated to ensure robust authentication and authorization, with every Terminal Network Element (TNE) required to support authenticator functionality (SEC-CTL-OFHPLS-3) and perform Port-based Network Access Control (PNAC) as defined in IEEE 802.1X-2020 (SEC-CTL-OFHPLS-4).

3) LEVERAGING 3GPP SECURITY STRENGTHS

O-RAN's security strategy incorporates the field-tested security protocols developed by 3GPP, ensuring that enhanced O-RAN security measures align with the established security frameworks of existing cellular networks. This approach

caters to the unique requirements of an open architecture while preserving the proven security frameworks of 3GPP, equipping O-RAN with advanced defences against current security threats and fostering a secure evolution of RANs.

IV. O-RAN SECURITY

The O-RAN architecture introduces new challenges in maintaining the integrity, availability, and confidentiality of the whole RAN. They rise from the disaggregated nature of the proposed architecture where new interfaces and functionalities are introduced, expanding the attack surface compared with traditional RANs. Identifying new threats and their consequences, setting roles and responsibilities for involved stakeholders, and providing solutions, are key aspects to ensure the safety, trust, and success of O-RAN.

This section is dedicated to introducing O-RAN's security analysis methodology which will be described in the upcoming subsections. It is divided into three main stages: Risk identification, Risk assessment, and Risk treatment.

A. RISK IDENTIFICATION STAGE

The Risk identification stage revolves around the creation of a set of studies to gather the necessary information to properly execute the Risk assessment and eventually the Risk treatment. O-RAN WG11 [9] divides the identification stage into eight categories:

- Identification of stakeholders that are involved in the O-RAN system, defining roles and responsibilities of each. The main stakeholders include MNOs, vendors, administrators and infrastructure providers.
- Definition of assumptions and prerequisites required for the successful operation of an O-RAN system. This normally includes reliable timestamping, trustworthy and capable agents (e.g., administrators, operators, etc.), and protection of stored log files, secrets, and credentials in external systems.
- Identification of assets existing in an O-RAN system in terms of its type, security properties, and location.
- Identification of threats that apply to the newly introduced interfaces and components that are relevant. Along with the threats, the threat surface and participating agents are provided.
- Identification of vulnerabilities and weaknesses on O-RAN interfaces, components, and related technologies that would allow an attacker to pose a threat.
- Definition of security principles to be followed to reduce risk exposure.
- Elaborate new security principles so that each of them is detailed and refined into requirements, recommendations, and countermeasures.
- Identification of existing countermeasures which include all already O-RAN existing/ongoing controls that new countermeasures need to consider.

Among these, the identification of threats and vulnerabilities is a key step for a successful security analysis. The O-RAN definition for a *threat* is extracted from NIST SP 800-30 [24] and is referred to as any circumstance with the potential to adversely impact operations and assets, via unauthorized access, destruction, disclosure, or modification of information, and/or denial of service. Similarly, *threat surface* is defined as all the individual points of attack against a particular system. In this case, O-RAN's openness, new components and interfaces, and the decoupling of hardware and software, significantly increase the quantity of threat surfaces when compared to 3GPP-based RAN systems. O-RAN WG11 [9] identifies six groups of threat surfaces:

- *Functions:* SMO, non-RT RIC, and near-RT RIC are new O-RAN components bringing in new functions and hence, they may open new attack surfaces.
- *Open Interfaces:* A1, E2, O1, O2, and OFH are also new interfaces introduced by O-RAN, opening new attack surfaces within the architecture.
- *Architecture:* The O-RAN 7.2x [20], introduces threats inherited from defined functions and interfaces within the split 5G base stations (gNB).
- *Trust Chain:* Decoupling functionalities for disaggregated architectures usually forces the expansion of trust chains, hence increasing the available threat surface.
- *Containerization and Virtualization:* NFV brought flexibility and programmability to network functions at the expense of extending the threat surface.
- *Open-Source Code:* open-source code (usually coming from re-using parts of non-trusted repositories) may expose the components to public exploits. Special attention need to be paid to not enough mature O-RAN enabled implementations which may also lead to new threats.

Accordingly, the threat surfaces identified above lead to the definition of a number of threats originating from inside and outside of O-RAN and from the APIs existing between different planes. A description of all identified threats is fully given in [9]. Such threats can be grouped depending on the exposed surface and/or the involved elements:

- **Threats Against the O-RAN system**: they group vulnerabilities available within the O-RAN architecture combining the ones exposed by *Functions*, *Architecture* and, *Interfaces* threat surfaces. O-RAN provides preliminary guidelines to address different threats but security risks are not fully covered yet as it is still in work-in-progress [4].
- **Threats Against O-Cloud**: Actors involved in managing the O-Cloud must keep a relationship based on trust. However, many threats mainly affect the *Containerization and Virtualization* surface as studied in [25], [26].
- **Threats to Open-Source Code**: Within the open-source community, there are inherent, well-known security risks caused by the openness of the architecture that may affect some of the O-RAN software components. Examples of these are co-developed codes causing RAN

function collisions [27] or GenAI-assisted codes causing backdoor attacks [28].

- **Physical Threats**: Similarly to other architectures, physical access to O-RAN components exposes risks related to physical unauthorized access enabling attacks over components reachable from the physical device.
- **Threats Against 5G Radio Networks**: The wireless communication channel is usually affected by *Jamming*, *Sniffing* or *Spoofing* threats as studied in [29], [30], [31].
- **Threats Against ML System: O-RAN natively integrates ML** within its architecture, and hence exposing it to generic ML-based vulnerabilities as detailed in [32].
- **Protocol Stack Threats**: Different communication interfaces within the O-RAN architecture are based on the REST protocol stack and, equivalently to the previous category, generic attacks on the aforementioned stack will be exploitable [4].

Finally, vulnerabilities must be identified on the assets targeted by the above threats. A *vulnerability* is defined as any trust assumption that can be violated to attack a system due to a flaw or weakness in an asset's design, implementation, or operation and management. Ultimately, vulnerabilities will enable the attacker to infiltrate the system through one or more assets and pose a threat. WG11 identifies the following O-RAN-specific vulnerabilities:

- Unauthorized access to O-DU, O-CU and O-RU: this vulnerability involves the risk of unauthorized entities gaining access to critical RAN components. If exploited, attackers could manipulate network operations, intercept sensitive data, and disrupt/degrade the service.
- Unprotected S-Plane and C-Plane in OFH interface: an attacker could potentially disrupt the communication between the O-RU and the O-DU, leading to service degradation or DoS.
- Disabling over-the-air cyphers for eavesdropping: attackers could exploit this to eavesdrop on communications, leading to a breach of confidentiality and interception of sensitive data.
- Near-RT RIC conflicts with E2 nodes: having different vendors commissioning different O-RAN nodes may create conflicts on, e.g., because of having implemented different versions of a given E2 service model. This could potentially result in reduced network performance, disruptions or even additional vector attacks.
- xApp and rApp conflicts: similarly, these conflicts can lead to inconsistent or erroneous decision-making, adversely affecting RAN performance and reliability.
- xApp and rApp access to subscriber data: malicious or compromised applications could exploit this access to extract sensitive subscriber information, leading to privacy violations and data breaches.
- Unprotected management interfaces: attackers could exploit these interfaces to gain control over network elements, alter configurations, or launch DoS attacks.
- Injection of control messages to attack the U-Plane: these attacks could lead to unauthorized data access, service interruptions, or manipulation of U-Plane traffic, impacting both network integrity and user privacy.

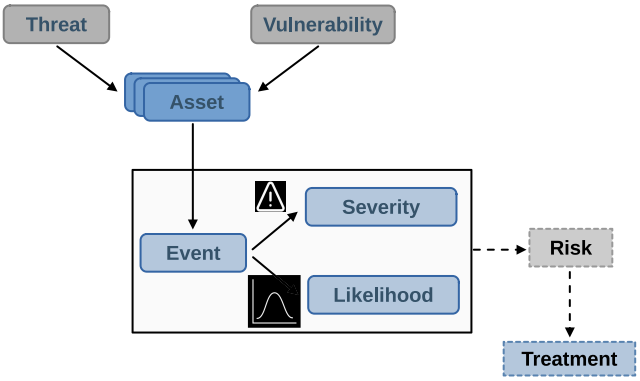

FIGURE 2. Risk assessment procedure assumed by the O-RAN Alliance.

TABLE 1. Risk assessment matrix [33].

		Likelihood		
		Low	Medium	High
Severity	Low	Low	Low	Medium
	Medium	Low	Medium	High
	High	Medium	High	High

In addition to these vulnerabilities, WG11 identifies system-wide, general vulnerabilities that include the decoupling of functions without required trust (i.e., hardware roots of trust and software trust chains), exposure to the public, open-source code exploits, and misconfiguration, poor isolation or insufficient access management in the O-Cloud. A comprehensive threat analysis is given in [13].

B. RISK ASSESSMENT STAGE

Once identified the threats exploiting the different vulnerabilities that may affect the assets of a given O-RAN system, the methodology defined by the WG11 quantifies the likelihood and potential impact of such events in the Risk assessment stage (see Figure 2). Two metrics contribute to the risk: the likelihood of occurrence of threats that successfully exploit vulnerabilities (likelihood measure), and the affectation severity (i.e., consequence) that such threat would pose (severity measure). Each of them is quantified with three levels: *low*, *medium* and *high*. Therefore, the final risk score is given according to the risk assessment matrix illustrated by Table 1. What follows is a description of how severity level and likelihood level ratings are calculated.

1) SEVERITY LEVEL RATING

In order to assess the severity level of a potential threat, the WG11 considers the level of impact on each of the security properties (i.e., Privacy, Confidentiality, Integrity

and Availability), the scale of the impact in terms of number of affected elements, the scale of the impact in terms of the affected Lower Layer Split (LLS) configurations (i.e., LLS-C1, LLS-C2, LLS-C3 and LLS-C4), and the possible adverse impacts. Following this approach, the *low* rating is constituted by a threat in which data leaked would unlikely reveal the subscriber's identity; would have minor effects on system running order; a single O-DU-O-RU pair is affected at most; DoS attacks on the S-Plane would only affect LLS-C2 or LLS-C4 configurations; and O-RAN's specifications already contain significant prevention actions towards the corresponding vulnerabilities. Additionally, a *medium* rating is given when personal data according to the European General Data Protection Regulation (GDPR) definition is disclosed to uniquely identify the end user; disclosure of privileged internal information such as access credentials or critical configuration data; the system operation can be interrupted in the order of hours or days; one O-DU with multiple O-RUs are affected; attacks are performed on LLS-C1 synchronization configurations; and O-RAN's specifications already contain significant prevention actions toward the corresponding vulnerabilities. Finally, a *high* classification corresponds to an event where sensitive personal data would be disclosed; the revelation of high-value internal information like trade secrets, IP, mission-critical data, or master keys; a complete loss of system integrity with an altered running order; long-term loss of availability measured in days or weeks; several O-DUs and O-RUs would be affected; DoS attacks affect S-Plane configurations based on LLS-C3; and O-RAN does not provide preventative actions against the particular vulnerabilities.

2) LIKELIHOOD LEVEL RATING

Parallelly, in order to estimate the likelihood level of a threat exploiting a vulnerability, the WG11 considers the potential adverse impacts, the type of initiation of the threat event, the exposure level, and the use or not of a Zero Trust Approach (ZTA). On the one hand, regarding the possibility of adverse impacts happening in the system, the likelihood level will be *high*, if no O-RAN security requirements and controls are put in place, while *medium* or *low* levels are achieved through the procurement of security countermeasures. Secondly, the threat event initiation factor takes into consideration the entry point for such vulnerability. For instance, the likelihood will be *high* if a vulnerability can be exploited from the public Internet, and less likely (e.g., *medium* or *low*), if access to a private network is previously required. Also, the exposure level will be directly related to the number of external interfaces and services exposed to the attacker. Finally, the use of ZTA involves assessing the likelihood of security incidents by considering a framework that assumes no inherent trust, even within the internal network. Unlike traditional security models relying on perimeter defence, ZTA acknowledges the potential threat from both external and internal actors. Consequently, likelihood scores are typically higher under a ZTA due to the expanded scope of potential threats.

C. RISK TREATMENT

Once the perceived Severity level and Likelihood level ratings have been calculated and a final risk score has been given (see again Table 1), a decision must be taken to possibly mitigate such risk. WG11 organizes risk treatment actions that can be undertaken into four different groups:

- **Modify the Risk:** This involves taking proactive measures to reduce the likelihood or impact of a threat. By modifying the risk, the operator implements controls to lessen the probability of a risk occurring or diminish its effects, for instance, by implementing stronger authentication and encryption protocols.
- **Avoid the Risk:** This strategy involves stopping the activities that lead to the risk. For example, if a particular action or process is deemed too risky, it is entirely discontinued.
- **Share the Risk:** In this approach, the operator outsources the risk management to a third party. This may include partnering with other entities that can manage the risk more effectively.
- **Retain the Risk:** This involves accepting the risk when the cost of mitigating it is higher than the potential impact. In such cases, the risk is acknowledged but it is chosen not to take significant action against it. A common example is to not upgrade legacy systems.

While these strategies are vital for managing the security environment of an O-RAN-enabled network, as of early 2024, no specific actions for each threat and risk have been defined by the O-RAN Alliance WG11 yet.

V. THREAT MODELLING

This section identifies the different attacks emerging from the threats identified by the O-RAN Alliance as introduced in Section IV. With the aim of characterizing the consequences of suffering such events and experimentally assessing the impact in the network of such anomalies (i.e., delay and packet loss), we first provide an overview of the threat model defined by O-RAN. Later on, since performance degradation and DoS attacks are significantly frequent threats, we define three scenarios that will allow us to test some of the most relevant interfaces, such as A1, F1-c/u, and E2, particularizing the O-RAN threat models to our considered scenarios.

A. ANALYSING O-RAN'S THREAT MODELS

This section presents the most common attacks on each interface as identified by O-RAN's WG11. Figure 3 depicts the following threat models and agents for the O-RAN system illustrating the expected location of such attacks (from top to bottom).

Data/model poisoning: This might include deliberate manipulation of datasets, or data models used in ML

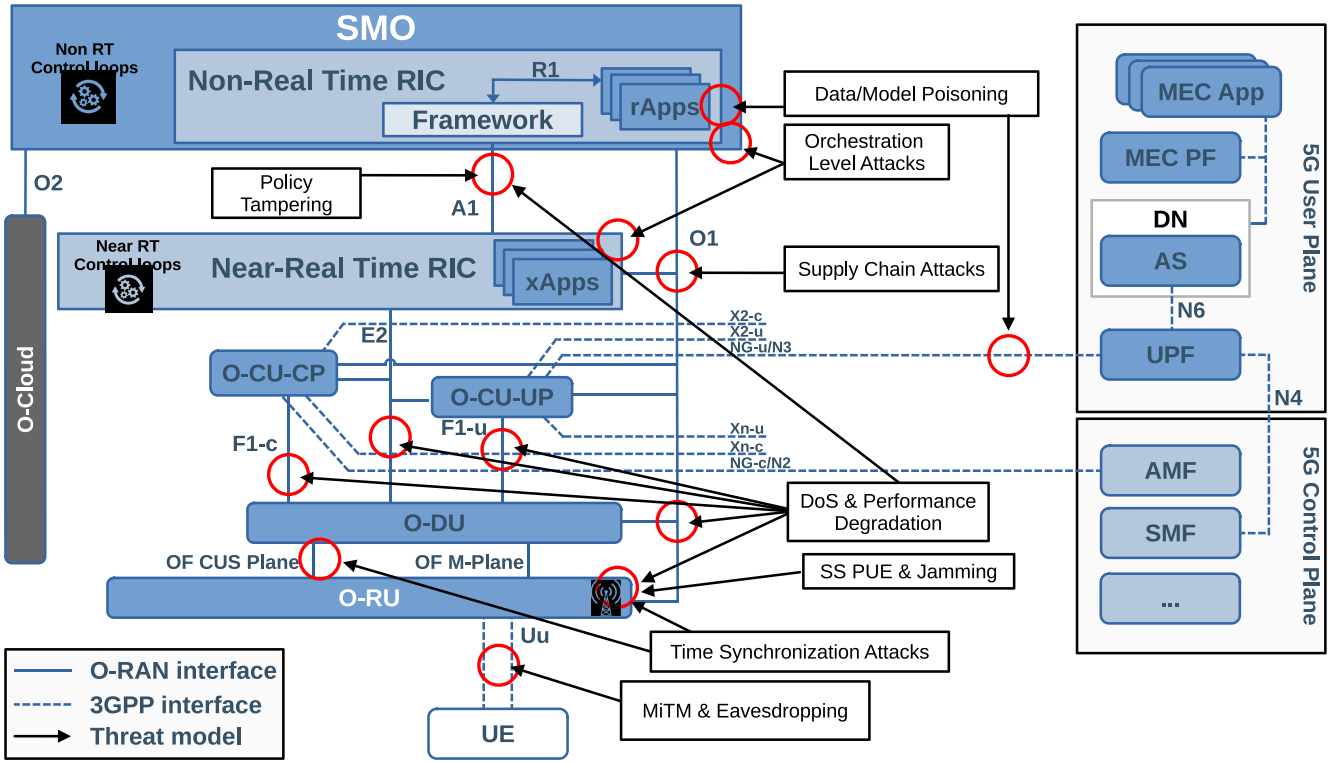

FIGURE 3. O-RAN Attacks in different interfaces as identified by WG11.

algorithms running as rApps, or affecting ML frameworks running both in the non-RT RIC and near-RT RIC, in order to cause the algorithm to make incorrect or malicious decisions. This could affect O-RAN services' leveraging ML algorithms for various tasks, such as self-organizing networks, traffic optimization, and interference management. As depicted in Figure 3, the data poisoning could happen also at the N3/NG-u interface or between MEC applications and rApps.

Orchestration level: In an Orchestration level attack, an attacker may try to take control of the orchestration layer, by exploiting vulnerabilities in the network management software, or by intercepting and manipulating network management messages. Once the attacker has control of the orchestration layer, they can cause a variety of negative consequences, such as:

- Disrupting the service: The attacker can cause the O-RAN to allocate resources incorrectly, leading to service disruption for users.
- Manipulating performance: The attacker can manipulate the O-RAN to allocate resources in a way that degrades overall network performance.
- Stealing information: The attacker can access sensitive information, such as user location data or network configuration details.
- Launching further attacks: The attacker can use their control of the orchestration layer to launch further attacks, such as launching a Denial-of-Service attack on other parts of the O-RAN system.

Policy Tampering: This is an attack in which an attacker alters the policies that govern the communication between different network elements, causing the network to make incorrect or harmful decisions. Policy Tampering attacks can be launched in several ways; an attacker can exploit vulnerabilities in the network management software to gain unauthorized access to the network and alter the policies, or to interfere, intercept, and manipulate network management messages. This can lead to a variety of negative consequences, such as service disruption, reduced performance, or security vulnerabilities. Policy Tampering can also be done by an attacker who can gain access to the network element and then physically manipulate the device's settings or firmware.

Supply Chain Attack: This is an attack that targets the supply chain of the O-RAN system. This can include introducing malicious hardware or software into the network or compromising legitimate hardware or software during the manufacturing or distribution process. For example, an attacker may introduce a malicious component into O-RAN, such as a rogue base station, that can be used to intercept and alter communications, disrupt the service, or steal sensitive information. Additionally, an attacker may also target the software used in O-RAN, e.g., by introducing malware or backdoors into the network management software, which can be used to gain unauthorized access to the network.

Supply Chain Attacks can be particularly difficult to detect and prevent, as the malicious components or software may not be easily identifiable.

DoS, performance degradation, Spoofed Station Physical User Equipment (SS PUE), and jamming: These are types of cyber attacks that aim to disrupt or jam access to a network or service by overwhelming it with traffic, using fake devices to gain unauthorized access, or intentionally disrupting radio communications by transmitting signals that interfere with the normal operation of the system. These attacks can cause significant damage in O-RAN by disrupting communication between network elements, leading to service disruption or causing the O-RAN controllers to make incorrect decisions.

Time Synchronization Attacks: Time Synchronization Attacks in O-RAN networks involve the manipulation of timing signals used by network elements to synchronize their operations, which can result in a DoS. In O-RAN, precise time synchronization is crucial for the coordinated operation of network elements, especially for the OFH. Synchronization is achieved through PTP and SyncE protocols in scenarios where the S-Plane goes over wire/fibre such as LLS-C3, and through Global Navigation Satellite System (GNSS) receivers in LLS-C4 scenarios [34]. Since the OFH interface is based on the eCPRI interface, it is the most time-and-packet-loss sensitive of the O-RAN interfaces. More precisely, packet delays of more than 100 μs will generally affect the performance and may trigger a connection reset [35]. Indeed, the loss of one C-Plane packet may trigger the loss of a slot's worth of data or can force the O-RU to move to a holdover state, (i.e., an O-RU relying only on a local clock source has the potential risk of drifting across the spectrum and increasing UE ranging errors). Thus, synchronization loss is a serious issue and can become a vector for DoS and performance degradation attacks.

Machine-in-the-Middle (MiTM): Finally, MiTM and eavesdropping attacks may also occur in O-RAN:

- MiTM Attack: A MiTM attack occurs when an attacker intercepts and alters communications between two parties. In the context of O-RAN, an attacker may intercept and alter control or U-Plane messages, causing the network to make incorrect decisions or disrupt service.
- Eavesdropping Attack: Eavesdropping attack is when an attacker intercepts and listens to communications without the parties involved being aware of it. In O-RAN, an attacker can intercept and listen to U-Plane data, allowing them to access sensitive information, such as user location data or private communications.

From this taxonomy, we can infer that DoS, outage, service disruption, and performance degradation threats are commonly depicted as the cause of the majority of threats identified by the O-RAN Alliance. In fact, as of the beginning of 2024, 62 of the 108 threats (around 60%) defined in WG11 are related to some kind of DoS or performance degradation. Therefore, in Table 2 we collect the threats identified by the WG11's Threat Modeling [33] with a given risk of "*high*" or "*medium*" that have DoS or performance degradation as direct consequences for the A1, E2 F1-c/u, O1, O2 or OFH interfaces. The table shows that generally, vulnerabilities that enable the attacker to gain access or control over an O-RAN component or interface are categorized with risk score "*high*", since that may allow for a severe disruption or degradation of the communications present in the corresponding compromised asset with a significant likelihood. Those marked with "*medium*" risk have either a lower severity level or a lower likelihood level. Furthermore, note that as the attacker gains access to the lower-level infrastructure elements (i.e., physical components, virtual machines, containers and virtualization platforms, etc.) anomalies in all interfaces must be expected, as can be observed in the last 13 threats of Table 2 (i.e., threats from T-OPENSRC-02 to T-OCAPI-01).

B. THREAT MODEL OF THE CONSIDERED SCENARIOS

To characterize the resulting effects associated with the exploitation of the security threats mentioned in Table 2, we experimentally assess the impact in the system performance of network anomalies (i.e., delay and packet loss) possibly caused by attacks affecting O-RAN components and interfaces. To do so, we deploy an O-RAN system fully based on O-RAN-compliant open-source initiatives. Then, we measure the criticality of such impact in terms of system reliability, resiliency, and Quality of Experience (QoE) impact on end-users.

Given the importance of DoS attacks among the ones identified by O-RAN Security WG11, we focus on the case where an attacker directly (e.g., via interception) or indirectly (e.g., via a malicious component) deteriorates the performance of an interface to downgrade or deny the service of a user. To model these attacks, we consider the threat model represented in Figure 4, where an attacker may gain access to the RAN and is able to disturb different O-RAN interfaces and components by exploiting one or more of the previously presented threats in Table 2. For example, the attacker could cause an xApp to misbehave within the near-RT RIC, applying malicious configurations to deteriorate the access to specific users or harming the E2 interface (e.g., by flooding it with bulk traffic), or blocking or degrading control flows going through (although may leave the A1 or F1-c interfaces untouched). The attacker may achieve this by exploiting, e.g., the threats *T-NEAR-RT-02*, *T-xAPP-01*, *T-xAPP-02* or *T-xAPP-04* presented in Table 2. Another example may be an attacker compromising the VM infrastructure by gaining access to the SMO through a backdoor in its code (e.g., *T-SMO-03*) [28], which could lead to a number of DoS related attacks.

Note that unlike in TIFG [10], we do not perform *black box* tests to certify O-RAN compliance upon a specific attack (e.g., a component abiding a network flooding attack of up to 1 Gbps). Instead, we focus on the effects on performance

TABLE 2. DoS and performance degradation threats identified by the WG11 [33].

THREAT ID	RISK DESCRIPTION	RISK SCORE	AFFECTED O-RAN INTERFACES					
			A1	E2	F1	O1	O2	OFH
T-O-RAN-01(nearRT)	Lack of security design adoption	High	✓	✓		✓		
T-O-RAN-01 (nonRT + SMO)	Lack of security design adoption	High	✓			✓	✓	
T-O-RAN-01 (O-CU)	Lack of security design adoption	High		✓	✓	✓		
T-O-RAN-01 (O-DU)	Lack of security design adoption	High		✓	✓	✓		✓
T-O-RAN-01 (O-RU)	Lack of security design adoption	High				✓		✓
T-O-RAN-02	Exploit of misconfigured components	High	✓	✓	✓	✓	✓	✓
T-O-RAN-03	Weak authentication and access control	High	✓	✓	✓	✓	✓	✓
T-O-RAN-05	Attacker compromises O-RAN open interfaces	High	✓	✓		✓	✓	✓
T-O-RAN-06	Insufficient authentication/authorization	High	✓	✓	✓	✓	✓	✓
T-O-RAN-09	Compromise of component integrity/availability	High	✓	✓	✓	✓	✓	✓
T-FRHAUL-01	An attacker penetrates O-DU and beyond through O-RU	High		✓	✓	✓		✓
T-FRHAUL-02	Unauthorized access to the OFH Ethernet interface	Medium		✓	✓	✓		✓
T-SPLANE-01-C1	DoS attack on an O-DU to distribute timing toward O-RU	Medium						✓
T-SPLANE-01-C3	DoS attack against a Master clock in LLS-C3	High						✓
T-SPLANE-02	Impersonation of a Master clock within a PTP network	Medium						✓
T-SPLANE-03	Rogue PTP Instance sending malicious messages	Medium						✓
T-SPLANE-04	Selective interception of PTP timing packets	Medium						✓
T-SPLANE-05	Packet delay manipulation attack	High						✓
T-ORU-01-b	An attacker stands up a rogue O-RU attacking O-DU	Medium			✓			✓
T-NEAR-RT-02	Risk of deployment of a malicious xApp on near-RT RIC	High	✓	✓	✓	✓		
T-NEAR-RT-03	near-RT RIC APIs compromised due to weak authentication	High	✓	✓		✓		
T-NEAR-RT-04	Resources provided by near-RT RIC via APIs can be abused	High	✓	✓		✓		
T-NONRTRIC-01/03	Gain non-RT RIC access through SMO	High	✓	✓		✓	✓	
T-xAPP-01	Exploit xApp vulnerabilities and misconfiguration	High	✓	✓		✓		
T-xAPP-03	Attacker compromises xApp isolation	High	✓	✓		✓		
T-xAPP-04	Malicious A1 policies modify xApp behaviour	High	✓	✓				
T-rAPP-01	Conflicting rApps impact performance or trigger DoS	High	✓			✓		
T-rAPP-02	Attacker exploits rApp vulnerability	High	✓			✓	✓	
T-rAPP-03	Attacker exploits rApps misconfiguration	High	✓			✓	✓	
T-rAPP-05	Malicious rApp deployment	High	✓			✓	✓	
T-PNF-01	Attacker compromises a PNF	High	✓	✓	✓	✓	✓	✓
T-SMO-03	Overload DoS attacks at SMO	Medium	✓			✓	✓	
T-OPENSRC-02	Developer inserts backdoor in code	High	✓	✓	✓	✓	✓	✓
T-PHYS-01/02	Intruder gains physical access to components	High	✓	✓	✓	✓	✓	✓
T-GEN-04	Availability of O-Cloud services compromised	High	✓	✓	✓	✓	✓	✓
T-VM-C-01	Gain VM or container high privileges	High	✓	✓	✓	✓	✓	✓
T-VM-C-02	Deploy malicious VM or container to lunch DDoS	High	✓	✓	✓	✓	✓	✓
T-VM-C-04-a	VM or container migration flooding	High	✓	✓	✓	✓	✓	✓
T-VM-C-04-b	Sniff, monitor and modify packets	High	✓	✓	✓	✓	✓	✓
T-VM-C-05	Misguide virtualization layer to reduce resources, cause DoS	High	✓	✓	✓	✓	✓	✓
T-IMG-04	Build a custom image that includes malware	High	✓	✓	✓	✓	✓	✓
T-VL-01	Gain control over host server and control all VMs/containers	High	✓	✓	✓	✓	✓	✓
T-VL-03	DoS against the service discovery infrastructure	Medium	✓	✓	✓	✓	✓	✓
T-O2-01	Tamper/alter/disclose services sent over O2	Medium	✓	✓	✓	✓	✓	✓
T-OCAPI-01	MiTM attacks on O-Cloud interfaces	Medium	✓	✓	✓	✓	✓	✓

degradation that such attack would have on different O-RAN components and interfaces (e.g., the effects in a component enduring a network flooding attack that causes 500 ms of additional latency in the interface, regardless of its magnitude or scale).

Thus, to evaluate the impact of such attacks on the system performance, we define the following three scenarios:

i) *End-to-end Video Scenario*: A UE from a network operator is requesting video-on-demand. Then, an attacker is able to harm operators' communications.

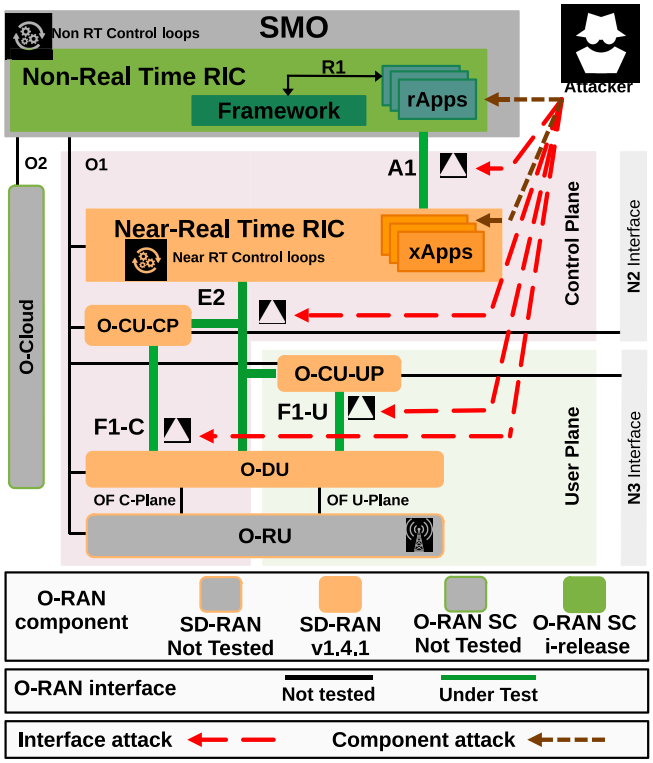

FIGURE 4. Threat model for three different attack scenarios involving A1, E2 and F1-c/u interfaces. We specify the open-source components comprising the O-RAN deployment used to test such scenarios.

- Threat Model: the attacker generates two different DoS attacks by injecting delay or packet loss to data flows at specific interfaces.
- Exploited Surfaces: A1, E2, F1-c and F1-u communication interfaces.
- KPI (U-Plane): Standardized QoE through the Peak Signal-to-Noise Ratio (PSNR) and Video Multimethod Assessment Fusion (VMAF) of the video received at the client [36].

ii) *Policy-Based Slice Configuration Scenario*: The flexibility introduced by O-RAN infrastructures enables RAN reconfiguration at different time scales. In this scenario, we trigger a RAN slice reconfiguration from the near-RT RIC, while a malicious attacker downgrades the control channel performance to delay the enforcement of this policy in the RAN.

- Threat Model: the attacker generates a delay-based DoS attack while a RAN reconfiguration is triggered.
- Exploited Surfaces: E2 communication interface.
- KPI (C-Plane): Policy reconfiguration within Operators' Service Level Agreements (SLAs).

iii) *Subscriber Attachment Scenario*: In this scenario, a UE is performing an attach procedure against the 5G core in order to get access to an external data network (e.g., the Internet). Simultaneously, an attacker selectively degrades the performance of the control channels involving O-CUs and O-DUs, aiming to prevent users from attaching.

- Threat Model: the attacker generates performance degradation attacks by selectively injecting delay and packet loss to control data flows.
- Exploited Surfaces: F1-c communication interface.
- KPI (C-Plane): Successful Attach Rate of a UE performing the registration process (%).

Policy-Based RAN Configuration and Subscriber Attachment scenarios (Scenario 2 and Scenario 3 respectively) are designed to study specific functionalities and behavioural situations in detail which were not fully estimable in the end-to-end Video On-Demand scenario (e.g., Scenario 1), but have a remarkable effect in the network performance (i.e., effective application of near-RT policies which only applies to the E2 interface, and reliable UE attachment, which only applies to F1-c). Table 3 maps which of the previously presented WG11 identified threats could be exploited to trigger network performance degradation issues as the ones experimentally evaluated in each of the proposed scenarios. As it can be observed, the covered threats are a subset of the threats presented in Table 2, excluding those only affecting the OFH interface. Also, while Scenarios 2 and 3 generally cover threats affecting E2 and F1-c interfaces respectively, Scenario 1 encompasses all the threat subsets that may affect A1, E2 or F1-c/F1-u.

VI. ATTACKING O-RAN COMMUNICATION INTERFACES

This section presents the outcomes derived from the experimental assessment. First, we describe the experimental deployment and present the methodology we follow. Then, in Section VI-B, we start reporting evaluation results obtained from *Scenario 1* for all the interfaces examined, using PSNR and VMAF metrics to describe the impact of both delay and packet loss on the interfaces. Subsequently, Section VI-C, showcases the results obtained from scenarios involving the E2 interface (i.e., *Scenario 1* and *Scenario 2*), focusing on measuring the delay experienced during a policy enforcement action when the E2 interface faces a DoS attack. Then, Section VI-D, exhibits the results acquired from scenarios involving the F1-c interface (i.e., *Scenario 1* and *Scenario 3*), where the primary objective is to characterize the failure probability of a user's attachment when F1-c is subjected to a DoS attack. Lastly, Section VI-E discusses the results from *Scenario 1* related to the F1-u interface. In order to comprehensively aggregate the obtained results from the three scenarios, Table 4 consolidates a summary of the results encompassed in the preceding subsections and includes outcomes related to the availability of the services involved in the communications.

The experimental scenario used for our evaluation combines two of the most relevant open-source O-RAN compliant initiatives, namely the O-RAN Software Community (O-RAN SC) and the Open Networking Foundation's (ONF) SD-RAN. The former is a collaboration between the O-RAN Alliance and the Linux Foundation comprising a community

TABLE 3. Mapping WG11 threat modelling to our evaluated scenarios.

THREAT ID	STUDIED INTERFACES			SCENARIOS		
	A1	E2	F1	1	2	3
T-O-RAN-01 near-RT RIC	✓	✓		✓	✓	
T-O-RAN-01 NonRT RIC + SMO	✓			✓		
T-O-RAN-01 O-CU		✓	✓	✓	✓	✓
T-O-RAN-01 O-DU		✓	✓	✓	✓	✓
T-O-RAN-02	✓	✓	✓	✓	✓	✓
T-O-RAN-03	✓	✓	✓	✓	✓	✓
T-O-RAN-05	✓	✓		✓	✓	
T-O-RAN-06	✓	✓	✓	✓	✓	✓
T-O-RAN-09	✓	✓	✓	✓	✓	✓
T-FRHAUL-01		✓	✓	✓	✓	✓
T-FRHAUL-02		✓	✓	✓	✓	✓
T-ORU-01-b			✓	✓		✓
T-NEAR-RT-02	✓	✓	✓	✓	✓	✓
T-NEAR-RT-03	✓	✓		✓	✓	
T-NEAR-RT-04	✓	✓		✓	✓	
T-NONRTRIC-01/03	✓			✓		
T-xAPP-01	✓	✓		✓	✓	
T-xAPP-03	✓	✓		✓	✓	
T-xApp-04	✓	✓		✓	✓	
T-rAPP-01	✓			✓		
T-rAPP-02	✓			✓		
T-rAPP-03	✓			✓		
T-rAPP-05	✓			✓		
T-PNF-01	✓	✓	✓	✓	✓	✓
T-SMO-03	✓			✓		
T-OPENSRC-02	✓	✓	✓	✓	✓	✓
T-PHYS-01/02	✓	✓	✓	✓	✓	✓
T-GEN-04	✓	✓	✓	✓	✓	✓
T-VM-C-01	✓	✓	✓	✓	✓	✓
T-VM-C-02	✓	✓	✓	✓	✓	✓
T-VM-C-04-a	✓	✓	✓	✓	✓	✓
T-VM-C-04-b	✓	✓	✓	✓	✓	✓
T-VM-C-05	✓	✓	✓	✓	✓	✓
T-IMG-04	✓	✓	✓	✓	✓	✓
T-VL-01	✓	✓	✓	✓	✓	✓
T-VL-03	✓	✓	✓	✓	✓	✓
T-O2-01	✓	✓	✓	✓	✓	✓
T-OCAPI-01	✓	✓	✓	✓	✓	✓

of key actors within the telecommunications field, while the latter is an initiative addressing the architectural challenges associated with modularity, openness, and multi-deployment interoperability. Both reference initiatives provide open-source components following O-RAN standards to stimulate innovation within the ecosystem.

Figure 4 depicts the deployed architecture at which components from different initiatives have been used. The non-RT RIC is part of O-RAN SC (i-release), and the near-RT RIC, O-CU, and O-DU are part of SD-RAN (v1.4.1). This setup has been selected due to the absence of O-CU/O-DU split support in O-RAN's near-RT RIC (i-release), which would have impeded the evaluation of the F1-c and F1-u interfaces. On the other hand, SD-RAN does not include a non-RT RIC implementation, similarly impeding the evaluation of the A1 interface. Therefore, we combine both projects to build a complete O-RAN deployment. It is important to highlight that SD-RAN's O-CU/O-DU is based on Open Air Interface (OAI), one of the most important open-source implementations of a 4G/5G protocol stack capable of providing cellular connectivity to off-the-shelf end-devices.

To perform end-to-end tests, we use *Live555MediaServer* as Real-Time Transport Protocol (RTP) server for 1080p video located at the edge of the cellular network by the UPF, and *VideoLAN* as a video client located at the UE. In order to generate QoE metrics (i.e., PSNR and VMAF) we use *FFmpeg*. With the aim of performing such end-to-end evaluation, the aforementioned deployment also includes software UEs from OAI and a core network (*OMEC*), which effectively communicates with the O-CU and O-DU. To avoid undesired and non-O-RAN related RF behaviours from Layer 1 (e.g., external interference or multi-path fading), and to improve the repeatability of the obtained results, the wireless link is simulated by using the *L2 nFAPI Simulator* from OAI, a simulator interconnecting the base station and the UE which short-cuts the PHY layer. The deployment configuration for SD-RAN's has been done following the *SDRAN-in-a-Box (RiaB) v1.4.0 "OMEC/CU-CP/OAI nFAPI Emulator for DU/UE"* documentation, where we subsequently included the non-RT RIC from O-RAN SC. For other open-source platforms that may be used to build an equivalent O-RAN setup, we refer the reader to FlexRIC [37] for the near-RT and srsRAN [38] for the O-CU and O-DU.

Overall, this setup allows us to evaluate the A1 and E2 interfaces defined by O-RAN and the F1-c and F1-u interfaces defined by 3GPP. All the components have been deployed in a single-node Kubernetes cluster (v1.20) on top of Ubuntu 20.04 LTS powered by a Haswell Intel Xeon Processor with 8 cores and 16GB of RAM. Performance degradation or DoS attacks have been enforced using the Trafic Control (TC) tool widely available in Linux kernels to create the adversarial effects through the *netem* network emulator. The effects are directly applied to the interfaces under study by means of the *delay* and *loss* options.

A. EXPERIMENTAL DEPLOYMENT

B. A1 INTERFACE

The A1 interface is based on HTTP REST and interconnects the non-RT RIC and the near-RT RIC, enabling the former to send policies and enriched information to the latter. It relies on TCP and hence, it is relatively tolerant to delay/loss attacks due to the retransmission and congestion control procedures.

End-to-end Video Scenario: The results of attacking A1 interface in this scenario are depicted in Fig. 5 and Table 4. The figure shows that regardless of the attack performed over the A1 interface (delay-based or packet-loss-based), the end-users are not affected as long as the policies that are being

TABLE 4. Interface evaluation summary. An *Interface* service refers to the availability of the interface, the rest of the services are sustained by the corresponding interface but may have failures of their own.

Interface	Service	Reaction to delays		Reaction to losses		Recovery
		Low (d ≥ 100ms)	High (d ≥ 2s)	Low (e ≥ 5%)	High (e ≥ 50%)	
A1	Availability	✓	✓	✓	✲	🕒
	A1-P	✓	✓	✓	✲	🕒
E2	Availability	✓	✓	✓	✓	✲
	onos-kpimon xApp	✓	✓	✓	✓	✗
	onos-rsm xApp	✓	✗	✓	✗	✗
F1-u	Availability	✓	✓	✓	✓	✲
F1-c	Availability	✗	✗	%	✗	✗
	UE attach.	✗	✗	%	✗	✗
	UE reconfig.	✗	✗	✗	✗	✗

✓ Unaffected | ✲ Temporarily unavailable | 🕒 Slow recovery (∼5 min) | % High failure chance (∼20%) | ✗ Failure

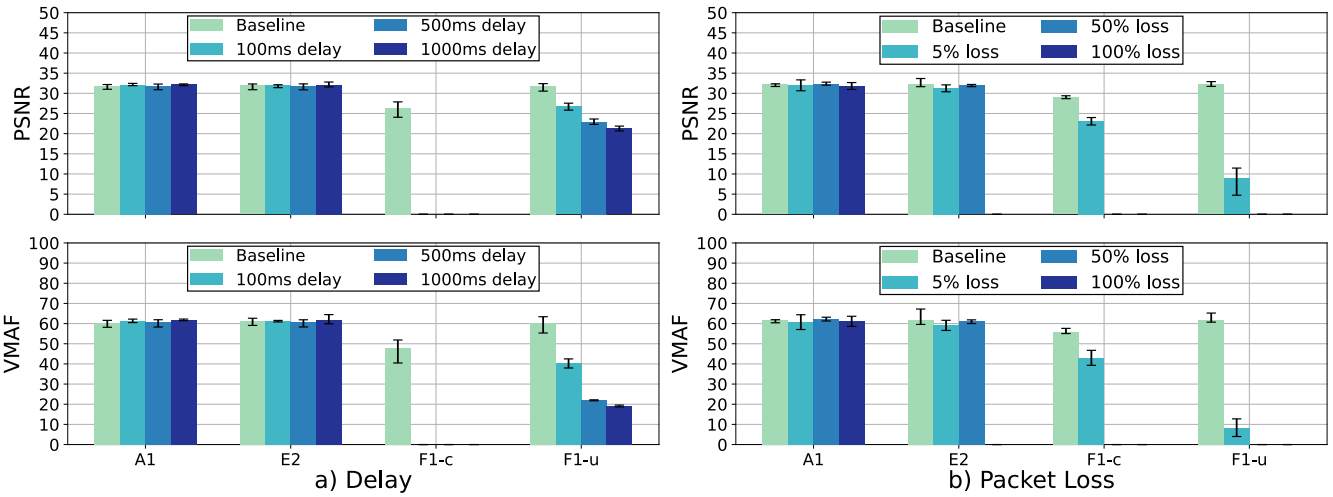

FIGURE 5. PSNR and VMAF video metrics per interface for different levels of a) packet delay (left), and b) packet loss (right).

transmitted through, do not actively manage their U-Plane. However, when the attack reaches 50% of packet loss, the policy management service is marked as *UNAVAILABLE* by the non-RT RIC and hence, disabled for any policy management (*"Reaction to losses column"*). In this case, our experiments show that the A1 interface remains unavailable for ∼ 5 minutes on average, a situation in which the overall architecture remains unmanaged. However, this does not affect the users' QoE as long as A1 policy updates are not required. Hence, A1 proves to be resilient and stable to non-ideal network conditions, and only heavy and continued DoS attacks may be effective.

C. E2 INTERFACE

As introduced in Section III-B, the E2 interface interconnects the near-RT RIC with the E2 nodes, enabling flexible and near real-time configuration of the nodes through the so-called *xApps*. The xApps are software modules included within the near-RT RIC implementing specific functionalities such as data collection or configuration enforcement. In our evaluation, we select two xApps: i) *onos-kpimon*, which periodically collects information from the E2 nodes; and ii) *onos-rsm*, which actively enforces network slicing configurations at the E2 nodes affecting the available bandwidth of specific end-users. E2 runs over the Stream Control Transmission Protocol (SCTP) transport protocol, using similar techniques to those of TCP for flow and congestion control.

End-to-end Video Scenario: Fig. 5 shows the impact of different QoS degradation attacks on the E2 interface in terms of PSNR and VMAF video metrics for the end-user. We can see that even if the E2 interface suffers from small delays (< 1s) or low packet loss (< 50% loss), traffic and services

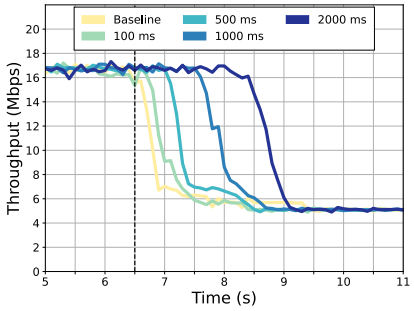

FIGURE 6. E2 interface delay evaluation for a network slicing use case.

will successfully recover and the RAN will not suffer any disruption. However, if the E2 interface has more than 50% of packet loss (or is not available for more than 225 seconds), an SCTP ABORT message will be triggered to terminate the association. Despite E2 is reestablished when the connection is restored, the O-CU carries the failure to the F1-c interface, where both the O-CU and O-DU will restart. This results in a temporal disruption of the UE traffic (see $PSNR = 0$ and $VMAF = 0$ for 50% of packet loss in Fig 5 b)). Still, the correlation between E2 and U-Plane performance is not strong, as long as the near-RT RIC is not directly optimizing QoS-related metrics or updating slices, a situation evaluated in the next scenario.

Policy-Based Slice Configuration Scenario: Results of performing delay-based DoS attacks to the E2 interface within Scenario 2 are depicted in Fig. 6. It highlights the delay sensitivity of the E2 interface when applying a RAN policy generated at the near-RT RIC. In this scenario, the *onos-rsm* executes a policy for reducing the resources available for a specific user (i.e., reducing its throughput from 17 Mbps to 5.5 Mbps). Since E2 RAN policies are near-real-time (<1 s), an attacker can effectively delay the enforcement of a specific policy. For instance, this may have undesired effects in the RAN, destabilising the scheduling algorithm or potentially generating congestion on other users. Fig. 6 shows how different delays in the E2 messages affect the throughput of the user. Since the policy mandates to reduce the slice size, we observe that the drop of throughput occurs only after the corresponding delayed E2 messages arrive to the gNB. However, for packet delays larger than 2400 ms, E2 message timeouts appear (see Table 4). Consequently, *onos-rsm* xApp enters a "blocking" state that forces its redeployment, keeping the last used configuration for the slice policies and opening a new surface to be exploited. Similarly, we have also tested the behaviour of *onos-kpimon* upon performance degradation in the E2 interface. While *onos-kpimon* proved to be resilient to performance degradation (low and high) both in terms of delay and loss, the xApp fails when there are major disruptions on the E2 interface, and therefore, monitoring traffic is never re-established (noted as *Failure* in Table 4). This may be problematic in use cases with a tight closed-loop control relying on *onos-kpimon*.

Overall, these experiments evidence the need for a strict trust chain for xApp onboarding, not only applying to malicious ones but also to malfunctioning ones.

D. F1-C INTERFACE

The F1-c interface interconnects the O-CU-CP component with the O-DU (see again Section III), using SCTP as the transport protocol. During experimentation, the SD-RAN deployment unveiled the usage of the nFAPI protocol on its lower split. Considering the available 3GPP 5G functional splits, SD-RAN uses double split Option 2 & 6 which departs from O-RAN's specified 2 & 7x split [20]. In either option, the F1-c interface interconnects the Midhaul split option 2, transporting PDCP and RLC traffic with a latency tolerance of 1.5∼10 ms during both attachment and GPRS Tunnelling Protocol (GTP) session phases. This means that if F1-c packets suffer delay (or are dropped too frequently), the UE may not be able to even attach, and therefore a possible vector for a DoS attack.

End-to-end Video Scenario: Fig. 5 shows the results for this scenario, clearly revealing the criticality of the F1-c interface when affected by DoS attacks, scoring zero PSNR and VMAF when the delay is increased from the baseline, and zero PSNR and VMAF score when packet loss is higher than 5%. The reasoning behind the zero score resides in the attach failures caused by these attacks which evidently results in zero scores for both metrics. This behaviour is further studied next in *Scenario 3*. Finally, such sensitivity to network issues in F1-c makes it even more vulnerable to stealth, subtle DoS attacks.

Subscriber Attachment Scenario: This scenario further studies the effects observed in *Scenario 1*, where we zoom in to evaluate the correlation of the UE attachment failure rate as a function of F1-c packet delay and packet loss. Results for this scenario are depicted in Fig.7. We experience a baseline value of 4% for RRC Setup Fail rate, very close to 3.3% stated by [39]. Fig. 7 a) shows a maximum tolerated latency of about ∼3 ms. For higher delays, the SCTP session fails to issue an SCTP shutdown, leaving the UEs detached from the network and requiring the restart of F1-c. This is a 3GPP related issue directly related to LTE's maximum UE association latency of 3 ms, as previously identified in [35] and [40]. Regarding packet loss, SCTP takes care of the re-transmission of the missing packets. However, if these re-transmissions occur on certain critical attachment packets (i.e., RRC Setup Request) the UE will not successfully attach because the 3 ms time limit will be exceeded.

Consequently, Fig. 7 b) illustrates a clear linear relation between packet loss probability and UE attachment failure rate. Remarkably, Fig. 7 b) also shows that with 30% of packet loss all attachment attempts fail, and that few drops (e.g., 5%) will cause 1 of every 5 attempts to fail.

E. F1-U INTERFACE

The F1-u interface carries U-Plane traffic and interconnects *O-CU User Plane (O-CU-UP)* and *O-DU* as introduced in

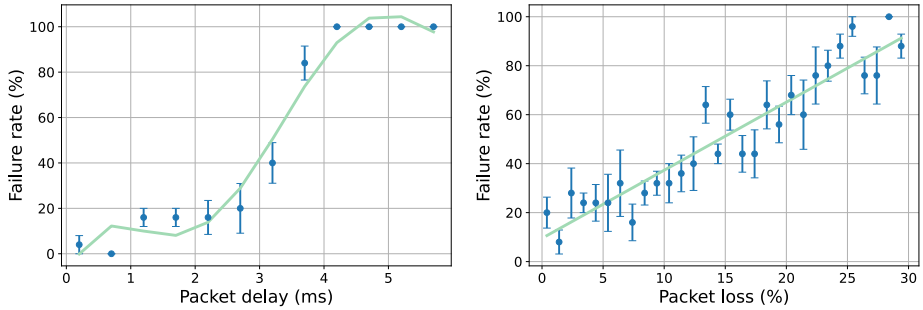

FIGURE 7. Subscriber attachment failure as a function of a) F1-c packet delay (left), and b) F1-c packet loss (right).

Section III. The F1-u interface specification defines GTP-U and UDP as the reference transport stack for their communication [41]. Surprisingly, the implementation present in SD-RAN deployment relies on *Protobuf* and UDP which, in turn, does not provide any delivery guarantees either. However, given the similarities between the standard option and the implementation included within the deployment under study, the results are applicable to both configurations.

End-to-end Video Scenario: Fig. 5 depicts the results of attacking the F1-u. As a result, the delays and packet loss introduced in this interface directly affect the U-Plane, certainly downgrading the user's QoE. Indeed, delay and loss tests from Fig. 5 confirm this behaviour. On the one hand, packet delay causes congestion in the video buffer, thus dropping the PSNR and VMAF. On the other hand, the video's PSNR and VMAF are dramatically reduced for 5% packet loss, and for 50% of packet loss, the video client is not even able to establish the streams.

VII. VULNERABILITIES AND SECURITY RECOMMENDATIONS ON O-RAN DEPLOYMENTS

A. VULNERABILITIES

In this section, we present a summary of how different O-RAN components and interface implementations behave upon performance degradation or DoS attacks, and the different security threats they are vulnerable to. Table 4 summarizes the results obtained in Section VI, highlighting how the effects of such attacks impact the different O-RAN services such as xApps, or RAN procedures and the availability of the underlying interfaces.

First, we can see that the A1 and E2 interfaces, and the service in charge of managing A1 traffic within the non-RT RIC (i.e., A1-P) are relatively robust against delay and packet loss, mainly due to TCP. However, the xApps (i.e., the *onos-rms* xApp in our tests) may misbehave when latency or packet loss is high, effectively impacting key RIC components or functionalities, and hence, opening a clear attack surface to be exploited by malicious actors operating within the infrastructure. Please note that xApps issues may vary depending on the specific xApp. Such events demonstrate the need for the near-RT RIC Security Component [9] to effectively validate and test xApp's compliance within the trust chain before they are on-boarded. Otherwise, the near-RT RIC is susceptible to easy-to-perform, repeatable, silent DoS attacks.

In addition to this, Table 4 shows that for the case of the F1-u interface, the resulting QoS degradation coming from a harmful DoS attack will only affect the QoE perceived by the end-user, without compromising any other architectural components (F1-u carries only user data). In contrast, F1-c is the most sensitive interface in the evaluation, since even small variations in latency or packet loss affect the service (please note that this is a feature inherited from 3GPP, and not only specific from O-RAN). This is critical since such variations cause the SCTP session to abort, and impede the normal operation of the interface for a period of time, including UE initial attachment and reconfiguration. This makes F1-c also a vulnerable target for easy-to-perform, subtle, and silent DoS attacks difficult to detect even with AI/ML-based anomalous behaviour detection algorithms.

Finally, while we focus on A1, E2 and F1-c/u interfaces, other O-RAN interfaces are not devoid of vulnerabilities. On the one hand, the OFH is highly vulnerable to DoS and performance degradation attacks which can interfere with the S-Plane by jamming and spoofing PTP, SyncE, or GNSS signals (we refer to [12] for a OFH preliminary security study). On the other hand, the O1 and O2 interfaces are still prone to generic attacks targeting the existing RESTful APIs [4] and only meet industry best practices when Transport Layer Security (TLS) is enabled (currently O-RAN sets it as optional).

B. SECURITY RECOMMENDATIONS

Our tests show that O-RAN deployments are subject to attacks that can not only affect and downgrade the performance of the interfaces but also disable crucial components such as the near-RT RIC or the O-DU. In order to prevent this, we identify several mechanisms that can help to mitigate attack situations and increase system security and availability:

1) RESULTS-BASED RECOMMENDATIONS

- **Strict Traffic Engineering.** Time Sensitive Networking (TSN) mechanisms such as strict priority with frame preemption and time-aware shaping [42] may not only help to strongly isolate malicious flows in very

high latency-sensitive O-RAN interfaces such as the OFH [43], but also in E2 and F1-c, (see Table 4). Also, zero trust approaches such as the one proposed by [45] may help to isolate intruders.

- **AI-based Anomaly Detection Systems.** Such systems may automatically detect malicious flows that, by causing packet losses, impact the services shown in Table 4, and proactively act against them [44].

2) GENERAL RECOMMENDATIONS

- **Secured Provisioning and Certificate Enrollment.** Since the O-RAN ecosystem opens up the number of vendors in the RAN, security gaps coming from integration issues are more prone to happen. A bottom-up, certificate-based trust chain for the provisioning, compliance, and conflict resolution of different O-RAN components (e.g., O-RUs, O-DUs, xApps, rApps, etc.) [4] can be used to ensure stability on the interfaces and services depicted in Table 4.
- **Secure Failure-proof Virtualization of O-RAN.** Besides cloud-native environments with chipsets enabling trusted computing and TLS enabled by default, secure carrier-grade virtualization in the O-Cloud may involve redundant instances and hard-swapping of O-RAN components in parallel infrastructure upon failures or anomalies, as exposed in [45].
- **Migration to Standalone (SA) 5G.** Non-Standalone (NSA) 5G inherits a number of legacy 4G security threats. O-RAN deployments based on SA 5G address most of the issues as reported in [31], and among other things, hinders UE impersonation attacks (that may lead to a high signalling load in F1-c de-registration requests) by making integrity protection mandatory.
- **S-Plane Attacks Mitigation.** Mitigation of S-Plane attacks includes strict traffic engineering to protect PTP and SyncE traffic, anti-jamming and anti-spoofing systems to secure GNSS receivers, as well as deploying PTP/SyncE/GNSS backup sync sources.

VIII. CONCLUSION

In this work, we revise and analyse the O-RAN threat model defined by the O-RAN Alliance in terms of risk identification, assessment and mitigation. Bearing in mind that most vulnerabilities reported by the O-RAN Alliance lead to DoS and performance degradation attacks, we study an O-RAN vanilla deployment and experimentally the effects of some of these threats in several O-RAN interfaces. Our hands-on evaluation depicts the actual consequences of such attacks happening on different interfaces and components. We found that A1 is robust against network issues (despite having large recovery times). However, E2, and especially F1-c, are very sensitive to link conditions, heavily impacting the end-users' QoE and even causing services to malfunction and drop. Accordingly, we endorse DoS exposure as a major threat toward a carrier-grade O-RAN and discuss preventive measures to improve the security and robustness of O-RAN architectures. Future work should extend this work to remaining O-RAN interfaces, i.e., O1/O2, E1 and the OFH.

REFERENCES

[1] I. Ahmad, T. Kumar, M. Liyanage, J. Okwuibe, M. Ylianttila, and A. Gurtov, "Overview of 5G security challenges and solutions," *IEEE Wireless Commun. Stand. Mag.*, vol. 2, no. 1, pp. 36–43, Mar. 2018.

[2] A. Garcia-Saavedra and X. Costa-Pérez, "O-RAN: Disrupting the virtualized RAN ecosystem," *IEEE Commun. Stand. Mag.*, vol. 5, no. 4, pp. 96–103, Dec. 2021.

[3] "5G Networks are worryingly hackable a shift to the cloud is opening the industry up to new attacks." IEEE Spectrum. Nov. 2021, Accessed: Nov. 11, 2023. [Online]. Available: https://spectrum.ieee.org/5g-virtualization-increased-hackability

[4] D. Mimran et al., "Security of open radio access networks," *Comput. Secur.*, vol. 122, Nov. 2022, Art. no. 102890.

[5] "Open RAN risk analysis 5GRANR v1.0: Federal office for information security." Accessed: Nov. 11, 2022. [Online]. Available: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Studies/5G/5GRAN-Risk-Analysis.pdf

[6] *Specification of the A5/3 Encryption Algorithms for GSM and ECSD; (Release 18), Version 18.0*, ETSI Standard TS 55.216, 2024.

[7] *Security Assurance Specification (SCAS) for 3GPP Virtualized Network Products; (Release 18), Version 18.2*, 3GPP Standard TS 33.527, 2024.

[8] *Security Architecture and Procedures for 5G System; (Release 15), Version 15.2*, 3GPP Standard TS 33.501, 2018.

[9] "WG11 O-RAN security threat modeling and remediation analysis 4.0." 2022. [Online]. Available: https://orandownloadsweb.azurewebsites.net/download?id=356

[10] "O-RAN test and integration focus group end-to-end test specification." 2022. [Online]. Available: https://orandownloadsweb.azurewebsites.net/download?id=361

[11] M. Polese, L. Bonati, S. D'Oro, S. Basagni, and T. Melodia, "Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges," *IEEE Commun. Surveys Tuts.*, vol. 25, no. 2, pp. 1376–1411, 2nd Quart., 2023.

[12] A. S. Abdalla and V. Marojevic, "End-to-end O-RAN security architecture, threat surface, coverage, and the case of the open fronthaul," *IEEE Wireless Commun. Stand. Mag.*, vol. 8, no. 1, pp. 36–43, Mar. 2024.

[13] F. Klement, W. Liu, and S. Katzenbeisser, "Towards securing the 6G transition: A comprehensive empirical method to analyze threats in O-RAN environments," *IEEE J. Sel. Areas Commun.*, vol. 42, no. 2, pp. 420–431, Feb. 2024.

[14] (Telefónica Telecommun. Co., Madrid, Spain, TID, Turlock, CA, USA, Deutsche Telekom Telecommun. Co., Bonn, Germany, Orange Telecommun. Co., Paris, France, and Vodafone Group PLC Telecommun. Co., Berkshire, U.K.). *Under The Open Ran MOU: Notes on the Technical Priorities (Release 4)*. (2021). Accessed: Jul. 29, 2024. [Online]. Available: https://telecominfraproject.mediavalet.com/galleries/dd6dc2e3-1a6e-4392-adc4-cad179cc539c_75e00c33-1f1b-4be5-b8b1-32663e0db872-ExternalUser

[15] "Sustainable trust v1.0." NGMN. Accessed: Nov. 11, 2022. [Online]. Available: https://www.ngmn.org/wp-content/uploads/210726-NGMN-Sustainable-Trust-V1.0.pdf

[16] N. C. Group. "Report on the cybersecurity of open ran." 2022, Accessed: Nov. 11, 2023. [Online]. Available: https://d110erj175o600.cloudfront.net/wp-content/uploads/2022/05/11160610/OPEN.pdf

[17] *O-RAN Security Requirements And Controls Specification 7.0*, O-Ran Alliance, Alfter, Germany, Accessed: Nov. 11, 2023.

[18] *O-RAN Security Protocols Specifications 7.0*, O-Ran Alliance, Alfter, Germany, Nov. 2023, Accessed: Nov. 11, 2022.

[19] *O-RAN Architecture Description 10.0*, O-RAN Alliance, Alfter, Germany, Accessed Nov. 11, 2023.

[20] L. M. P. Larsen, A. Checko, and H. L. Christiansen, "A survey of the functional splits proposed for 5G mobile crosshaul networks," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 146–172, 1st Quart., 2019.

[21] (Telefonica Telecommun. Co., Madrid, Spain). *Rec-Sec Open-Ran Technical Priority Document Release*, 2023, Accessed: May 15, 2024. [Online]. Available: https://www.telefonica.com/en/communication-room/reports/open-ran-technical-priorities-release-3

[22] W. Stallings, *Cryptography and Network Security: Principles and Practice*. 6th ed. Hoboken, NJ, USA: Prentice-Hall, 1995.
[23] "National telecommunications and information administration (NTIA)." 2023, Accessed: May 15, 2024. [Online]. Available: https://www.ntia.gov/
[24] *NIST SP 800-30 Revision 1: Guide for Conducting Risk Assessments*, Nat. Inst. Stand. Technol. Gov. Agency, Gaithersburg, MD, USA, 2012.
[25] "Security Assurance Methodology (SECAM) and Security Assurance Specification (SCAS) for 3GPP virtualized network product," 3GPP, Sophia Antipolis, France, Rep. TR 33.818, 2021. [Online]. Available: https://www.3gpp.org/ftp/Specs/archive/33_series/33.818/33818-h10.zip
[26] M. Souppaya, J. Morello, and K. Scarfone, "Application container security guide," U.S. Dept. Commer., Nat. Inst. Stand. Technol., Gaithersburg, MD, USA, Rep. NIST SP 800-190, 2017.
[27] D. Je, J. Jung, and S. Choi, "Toward 6G security: Technology trends, threats, and solutions," *IEEE Commun. Stand. Mag.*, vol. 5, no. 3, pp. 64–71, Sep. 2021.
[28] Z. Yang et al., "Stealthy backdoor attack for code models," *IEEE Trans. Softw. Eng.*, vol. 50, no. 4, pp. 721–741, Apr. 2024.
[29] M. Lichtman, R. Rao, V. Marojevic, J. Reed, and R. P. Jover, "5G NR jamming, spoofing, and sniffing: Threat assessment and mitigation," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, 2018, pp. 1–6.
[30] N. Ludant, P. Robyns, and G. Noubir, "From 5G sniffing to harvesting leakages of privacy-preserving messengers," in *Proc. IEEE Symp. Security Privacy (SP)*, 2023, pp. 3146–3161.
[31] O. Lasierra, G. Garcia-Aviles, E. Municio, A. Skarmeta, and X. Costa-Pérez, "European 5G security in the wild: Reality versus expectations," in *Proc. 16th ACM Conf. Security Privacy Wireless Mobile Netw.*, 2023, pp. 13–18. [Online]. Available: https://doi.org/10.1145/3558482.3581776
[32] M. Xue, C. Yuan, H. Wu, Y. Zhang, and W. Liu, "Machine learning security: Threats, countermeasures, and evaluations," *IEEE Access*, vol. 8, pp. 74720–74742, 2020.
[33] *WG11 O-RAN Security Threat Modeling and Risk Assessment 1.0*, O-RAN Alliance, Alfter, Germany, 2023. [Online]. Available: https://orandownloadsweb.azurewebsites.net/download?id=554
[34] *Synchronization Architecture and Solution Specification 1.00 (O-RAN.WG9.XTRP-SYN.0-v01.00)*, O-RAN Alliance, Alfter, Germany, 2021.
[35] E. Municio, G. Garcia-Aviles, A. Garcia-Saavedra, and X. Costa-Pérez, "O-RAN: Analysis of latency-critical interfaces and overview of time sensitive networking solutions," *IEEE Commun. Stand. Mag.*, vol. 7, no. 3, pp. 82–89, Sep. 2023.
[36] R. Rassool, "VMAF reproducibility: Validating a perceptual practical video quality metric," in *Proc. IEEE Int. Symp. Broadband Multimedia Syst. Broadcast. (BMSB)*, 2017, pp. 1–2.
[37] R. Schmidt, M. Irazabal, and N. Nikaein, "FlexRic: An SDK for next-generation SD-RANs," in *Proc. 17th Int. Conf. Emerg. Netw. Exp. Technol.*, 2021, pp. 411–425.
[38] "Open-source 4G and 5G software radio suites developed by software radio systems." SRSRAN, Accessed: May, 21, 2024. [Online]. Available: https://www.srslte.com/
[39] H.-S. Park and Y.-S. Choi, "Taking advantage of multiple handover preparations to improve handover performance in LTE networks," in *Proc. 8th Int. Conf. Future Generation Commun. Netw. (FGCN)*, 2014, pp. 9–12.
[40] G. Garcia-Aviles, A. Garcia-Saavedra, M. Gramaglia, X. Costa-Perez, P. Serrano, and A. Banchs, "Nuberu: Reliable RAN Virtualization in Shared Platforms," in *Proc. 27th Annu. Int. Conf. Mobile Comput. Netw.*, 2021, pp. 749–761.
[41] *5G NG-RAN F1 General Aspects and Principles; (Release 16), Version 16.2*, ETSI Standard TS 38.470, 2020.
[42] S. Haas et al., "Trustworthy computing for O-RAN: Security in a latency-sensitive environment," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, 2022, pp. 826–831.
[43] S.-H. Liao, C.-W. Lin, F. A. Bimo, and R.-G. Cheng, "Development of C-plane DoS attacker for O-RAN FHI," in *Proc. 28th Annu. Int. Conf. Mobile Comput. Network.*, 2022, pp. 850–852.
[44] K. N. Qureshi, G. Jeon, and F. Piccialli, "Anomaly Detection and trust authority in artificial intelligence and cloud computing," *Comput. Netw.*, vol. 184, Jan. 2021, Art. no. 107647. [Online]. Available: https://doi.org/10.1016/j.comnet.2020.107647
[45] I. Tamim, A. Saci, M. Jammal, and A. Shami, "Downtime-aware O-RAN VNF deployment strategy for optimized self-healing in the O-cloud," in *Proc. IEEE Glob. Commun. Conf. (GLOBECOM)*, 2021, pp. 1–6.

PAU BAGUER received the double engineering degree in aerospace systems and networking from the Polytechnic University of Catalonia, Barcelona, in 2023. He joined the AI-Driven Systems Group, i2CAT, June 2022, and became a Junior Researcher in 2023.

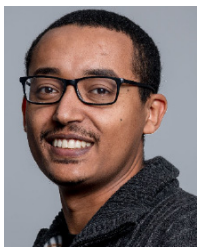

GIRMA M. YILMA received the M.Sc. degree in telecommunications engineering from the University of Trento, Italy, in 2016. He is currently a Senior Research Engineer with NEC Laboratories Europe.

ESTEBAN MUNICIO received the Ph.D. degree from the University of Antwerp, imec, Belgium, in 2020. Then, he joined imec as a Postdoctoral Researcher for two years. Since January 2022, he has been with i2CAT, where currently he is a Research Scientist with the AI-Driven Systems Group.

GINES GARCIA-AVILES received the Ph.D. degree in telematics engineering from the IMDEA Networks Institute, University Carlos III of Madrid. Since January 2021, he has been with i2CAT, where he is currently a Research Scientist with the AI-Driven Systems Group.

ANDRES GARCIA-SAAVEDRA received the Ph.D. degree from the University Carlos III of Madrid in 2013. He is currently a Principal Researcher with NEC Laboratories Europe. His research interests lie in the application of fundamental mathematics to real-life wireless communication systems.

MARCO LIEBSCH (Member, IEEE) received the Ph.D. degree from the University of Karlsruhe, Germany, in 2007. He is a Chief Researcher with NEC Laboratories Europe GmbH, and is working in the area of 5G evolution, mobile edge computing, and cloud networking.

XAVIER COSTA-PÉREZ (Senior Member, IEEE) received the Ph.D. degree in telecommunications from the Polytechnic University of Catalonia, Barcelona. He is an ICREA Research Professor, a Scientific Director of the i2CAT Research Center, and the Head of the 6G R&D at NEC Laboratories Europe. He has served on the organizing committees of several conferences, published papers of high impact, and holds more than 80 granted patents. He was the recipient of a National Award for his Ph.D. thesis.