Demo: Attack Execution and Multi-Layer xApp-based Intrusion Detection in an O-RAN Environment

Farah Abed Zadeh, Bartlomiej Siniarski, Shen Wang, Madhusanka Liyanage
Network Softwarization and Security Labs (NetsLab), School of Computer Science, University College Dublin, Ireland
Email: {farah.abedzadeh, bartlomiej.siniarski, shen.wang, madhusanka}@ucd.ie

***Abstract*—As Open Radio Access Networks (O-RAN) accelerate 5G deployments through intelligent, software-defined architectures, they concurrently expand the cellular attack surface. Deploying and validating Intrusion Detection Systems (IDS) against attacks such as Radio Resource Control (RRC) signaling storms and volumetric network floods requires scalable, realistic testbeds. In this demonstration, we present a software-simulated O-RAN environment capable of multi-UE connectivity and executing coordinated multi-layer attacks. By leveraging srsRAN, Open5GS, FlexRIC, and a custom GNU Radio ZeroMQ multiplexer, we demonstrate a dual-layer defense framework featuring a real-time Monitoring xApp that bridges physical telemetry and network flows to AI-driven IDS xApps, enabling the rapid detection of malicious activities within the Near-Real-Time RAN Intelligent Controller (Near-RT RIC).**

***Index Terms*—O-RAN, Intrusion Detection, xApp, RRC Signaling Storms, Network Security, srsRAN**

I. Introduction

The disaggregation and open interfaces of O-RAN have significantly expanded the cellular attack surface, exposing networks to threats such as malicious xApps, adversarial AI/ML, and RRC signaling storms [1]. These vulnerabilities result in service outages, privacy breaches, and degraded quality of service, ultimately jeopardizing network availability for end users. While recent literature has made strides in O-RAN security, existing approaches exhibit critical limitations. For instance, security framework 5G-SPECTOR [2] relies on static, rule-based systems and upper-layer telemetry, lacking physical layer insights. Other works [3] depend on synthetic traffic generators or hardware-constrained testbeds that limit multi-UE scalability and reproducibility.

Our work bridges this gap by proposing the execution of realistic general network attacks and O-RAN-specific RRC signaling storms in a highly scalable, simulated O-RAN testbed. This setup supports the simultaneous connection of multiple User Equipments (UEs) and hosts a dual-layer IDS framework on a Near-Real-Time (Near-RT) RAN Intelligent Controller (RIC). By leveraging a foundational monitoring xApp, the two IDS xApps cross-validate one another using distinctly different metrics (PHY/MAC vs. Network layer) to ensure highly accurate, near-real-time detection of anomalous behavior.

ISBN 978-3-903176-82-9 ©2026 IFIP

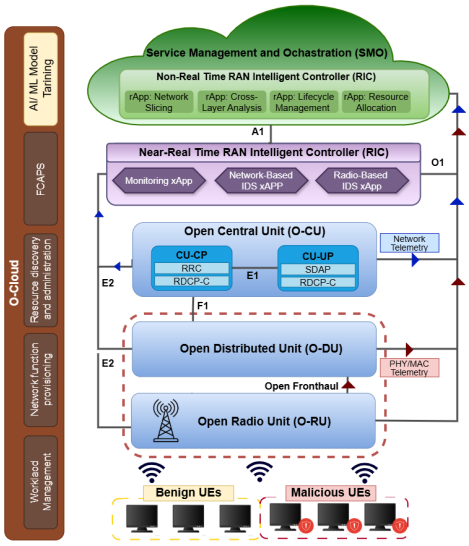

Fig. 1. Proposed O-RAN architecture with dual-layer IDS xApps deployed within the Near-RT RIC and telemetry extraction points across the stack.

Contributions

- **Reproducible Multi-UE O-RAN Testbed:** We provide a software-simulated O-RAN testbed built on srsRAN, Open5GS, FlexRIC, and GNU Radio, including attack scripts, UE configurations, custom ZMQ multiplexer, and xApp artifacts.
- **Multi-Modal Attack Execution:** Concurrent execution of data-plane volumetric floods and O-RAN-specific RRC signaling storms via UE reboot loops from independent network namespaces.
- **Dual-Layer Cross-Validating IDS:** A Monitoring xApp feeds multi-layer telemetry to two AI-driven IDS xApps, enabling cross-validation of physical-layer radio anomalies against network traffic behaviors.
- **Real-Time Dashboard:** A web-based interface correlating PHY/MAC degradation with network-layer anomalies for intuitive cross-layer visualization.

II. System Architecture

Fig. 1 illustrates the O-RAN architecture adopted in our system, spanning the full stack from the Open Radio Unit (O-RU) through the Open Distributed Unit (O-DU) and Open Central Unit (O-CU) to the Near-RT RIC, where the proposed xApps are deployed for dual-layer intrusion detection.

A. Virtual O-RAN Infrastructure

- **5G Core:** The core network is deployed using Dockerized Open5GS, providing AMF, UPF, and SMF functionalities.
- **RAN Setup:** We utilize the *srsRAN Project* gNB to simulate the O-CU and O-DU. The gNB interfaces with the core via the N2/NGAP interfaces.
- **Multi-UE Emulation:** User Equipment is emulated using *srsRAN 4G (srsUE)*. To allow multiple UEs to connect to a single gNB ZMQ port simultaneously, we introduce a **GNU Radio flowgraph**. This flowgraph acts as an RF channel emulator, multiplexing IQ samples between the gNB's single ZMQ port pair and multiple per-UE ZMQ port pairs. Each UE is isolated within its own Linux network namespace to ensure independent data-plane routing.
- **Near-RT RIC:** *FlexRIC* is deployed to handle E2 interface communications with the gNB, managing E2SM-KPM (Key Performance Metrics) and E2SM-RC (RAN Control) service models.

B. Dual-Layer xApp Defense Framework

To detect both control-plane and data-plane anomalies, the Near-RT RIC hosts a modular security framework:

1) **Monitoring xApp:** This module subscribes to the gNB via E2SM-KPM to extract real-time PHY/MAC telemetry (e.g., RSRP, SINR, CQI, BLER). Simultaneously, it correlates this data with higher-layer network flow metrics (byte counts, protocols) captured via host network namespaces. The parsed and aligned data is streamed to the analytics layer.
2) **Radio IDS xApp:** An ML-driven application that consumes PHY/MAC telemetry to detect physical footprints of attacks, such as the scheduling instability and random-access collisions.
3) **Network IDS xApp:** An ML-driven application evaluating higher-layer network traffic metrics to identify the executed flooding attacks.

III. Attack Execution

To validate the IDS xApp models, we execute two primary categories of attacks within the simulated environment.

A. Data-Plane Volumetric Attacks

Data-plane floods aim to overwhelm the 5G core network or the O-DU/O-CU packet processing pipelines. From compromised UE namespaces, we launch simulated DoS/DDoS attacks. Utilizing penetration testing tools such as *hping3*, *nping*, and *slowloris.py*, we generate TCP SYN/ACK floods, UDP/ICMP floods, and HTTP floods to exhaust connection tables and saturate resources.

TABLE I
Dual-layer IDS performance comparison using XGBoost for RRC signaling storm detection.

	Radio-Level	**Network-Level**
Detection Performance		
Accuracy (%)	99.43	97.88
Precision (%)	99.17	97.67
Recall (%)	99.52	97.17
F1-Score (%)	99.35	97.42
Computational Efficiency		
Training Time (s)	1.31	1.43
Detection Time (ms)	0.13	0.17
CPU Effort (%·s)	43.56	44.94
Inference Energy (mWh/1000)	25.10	8.04

B. RAN-Related Control-Plane Attacks

Unlike volumetric floods, RRC signaling storms use protocol-compliant messages to exhaust state-machine resources at the gNB. Using a synchronized bash-scripted automation, we introduce rapid UE reboot cycles of multiple srsUEs. This behavior floods the gNB with repeated PRACH preambles and RRC Setup Requests.

IV. Demonstration

The demo proceeds in three phases. We first demonstrate the automated orchestration of the environment through initialization of the Open5GS core, FlexRIC, and gNB. Subsequently, a batch script launches 15 instances of srsUE. The audience will observe the GNU Radio flowgraph routing IQ samples, the successful assignment of IP addresses to the UE namespaces via the UPF, and the seamless communication between the network and the UEs. Once stable, we trigger the attacks detailed in Section III. We will show the terminal execution of both data-plane and control-plane attacks, and how specific UEs can cause latency delays and congestion within the network. Next, the Monitoring xApp interface will be displayed, visualizing the extraction of E2 telemetry and network flows. The Network IDS xApp will flag anomalies based on network telemetry, while the Radio IDS xApp will simultaneously detect these attacks, identifying sharp degradation in radio-level signals and frequencies. To visualize the detection process, Fig. 2 shows the interface that provides a unified view of multi-layer network behavior across the O-RAN stack.

V. Experimental Evaluation

Table I compares XGBoost-based binary classification of RRC signaling storm attacks across both IDS layers. The Radio IDS outperforms the Network IDS in detection performance, due to richer PHY/MAC telemetry capturing the immediate physical footprint of attacks. From a computational perspective, both models demonstrate efficient execution suitable for near-real-time deployment, with the Radio IDS achieving faster detection. The Network IDS shows lower inference energy per 1000 instances, indicating upper-layer detection can

Fig. 2. O-RAN multi-layer IDS dashboard showing radio-layer degradation, detected threat alerts, network-layer anomalies, and a correlated attack timeline.

be more energy-efficient. These results validate the complementary nature of the dual-layer framework: early physical-layer detection augmented by network-layer validation. All results were derived from prior experimental deployments in a realistic O-RAN environment. This demonstration further facilitates reproducibility and scalability to larger multi-UE scenarios.

VI. Conclusion

This demo presents a scalable, fully software-simulated 5G O-RAN testbed capable of emulating complex, multi-UE attack scenarios without requiring physical RF hardware. By integrating GNU Radio ZMQ multiplexing for multi-UE connectivity, FlexRIC for real-time E2 telemetry extraction, and a novel dual-layer xApp IDS framework, we demonstrate that cross-validating PHY/MAC radio anomalies with network-layer traffic analysis enables highly accurate, near-real-time intrusion detection. To the best of our knowledge, this is the first zero-hardware O-RAN security testbed to support concurrent multi-layer attack execution and detection, lowering the barrier for researchers and operators to prototype and validate ML-driven security solutions for next-generation open cellular networks.

Acknowledgment

This research was funded by the ROBUST-6G Project through the SNS JU under the EU's Horizon Europe Program (Grant: 101139068), and CONNECT phase 2 (SFI grant 13/RC/2077_P2).

References

[1] A. Tabiban, H. A. Alameddine, M. A. Salahuddin, and R. Boutaba, "Signaling storm in o-ran: Challenges and research opportunities," *IEEE Communications Magazine*, vol. 62, no. 6, pp. 58–64, 2024.

[2] H. Wen, P. A. Porras, V. Yegneswaran, A. Gehani, and Z. Lin, "5g-spector: An o-ran compliant layer-3 cellular attack detection service." in *NDSS*, 2024.

[3] M. Hoffmann and P. Kryszkiewicz, "Signaling storm detection in iiot network based on the open ran architecture," in *IEEE INFOCOM 2023-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2023, pp. 1–2.