



O-RAN Minimum Viable Plan and Acceleration towards Commercialization

White Paper, 29 June 2021

Acknowledgements

Special thanks must be given to the numerous people contributing to the White Paper, including

Co-editors: Arda Akman (Juniper Networks), Chih-Lin I (CMCC)

Major contributors: Arda Akman (Juniper Networks), Suzy Gu (CMCC), Jinri Huang (CMCC), Ian Wong (Viavi), Haseeb Akhtar (Ericsson), Scott Poretsky (Ericsson), Jan Balas (DT), Qi Sun (CMCC), Burcu Yargicoglu-Sahin (Juniper Networks), Yuxuan Xie (CMCC), Jens Voigt (Amdocs), Gabor Hannak (Nokia), Balaji Raghothaman (Keysight), Shinobu Nanba (KDDI), Kevin Tang (Sterlite), Joey Chou (Intel), Yaron Koral (AT&T), Bei Yang (China Telecom), Yang Liu (China Telecom), Jun Song (Samsung), Sukhdeep Singh (Samsung), Shankar Narayanan (AT&T), Lyndon Ong (Ciena), Rittwik Jana (VMware), Andrea Buldorini (TIM)

Major reviewers: Vikas Dixit (Reliance JIO), David McKechnie (Telstra), Dhruv Gupta (AT&T), Mike Garyantes (Qualcomm), Leifeng Ruan (Intel), Hank Kafka (AT&T)

Contents

EXECUTIVE SUMMARY	4
INTRODUCTION	5
O-RAN ARCHITECTURE	6
O-RAN MVP AND PRIORITIZED USE CASES FOR 2021	8
O-RAN MVP	8
“Open” Package	8
“Intelligent Package” - Prioritized Use Cases for 2021	9
Traffic Steering	9
QoS and QoE Optimization	10
RAN Slicing and SLA Assurance	11
SMO (Service Management and Orchestration framework)	12
Massive MIMO Optimization	13
NEW USE CASES AND UPDATED DEPLOYMENT SCENARIOS	14
New Use Cases	14
Multi-vendor Slices	14
O-RAN DSS	14
RAN Slice Resource Allocation Optimization	15
Local Indoor Positioning in RAN	15
Massive MIMO SU/MU-MIMO Grouping Optimization	16
O-RAN Signaling Storm Protection	16
Congestion Prediction and Management	17
Industrial IoT Optimization	17
O-DU Pooling to Achieve RAN Elasticity	18
Updated Deployment Scenarios	18
O-RAN OPEN-SOURCE COMMUNITY	19
O-RAN TESTING AND INTEGRATION – OTIC, GLOBAL PLUGFESTS AND MORE	21
Open Test and Integration Centers (OTICs)	21
Certification and Badging	21
End-to-End Test Specification	22
Global Plugfests	22
SECURE O-RAN	23
CONCLUSION	25
REFERENCES	26

Executive Summary

The O-RAN ALLIANCE is committed to evolving radio access networks (RAN), making them more open, smarter, interoperable, and scalable than contemporary deployments. The first whitepaper “O-RAN: Towards an Open and Smart RAN” introduced key O-RAN concepts: (1) Openness to build a more cost-effective network through open interfaces and open hardware, and (2) Intelligence to suit complex, dense, and future-ready networks through embedded intelligence in every layer.

The second whitepaper “O-RAN Use Cases and Deployment Scenarios” introduced an initial set of use cases that leverage the unique benefits of the O-RAN architecture. This includes utilization of machine learning (ML) and artificial intelligence (AI) modules operating through open and standardized interfaces in a multi-vendor network, along with white-box hardware, open hardware reference design, and cloud native deployment through the O-RAN cloudification and orchestration platform, O-Cloud.

This third whitepaper “O-RAN Minimum Viable Plan and Acceleration towards Commercialization” extends the previous whitepapers and outlines the Alliance’s focus on several key areas to accelerate and enable the introduction of the rich capabilities of O-RAN architecture in commercial networks and provides a minimum set of end-to-end specifications for selected use cases for the deployment of a secure, multi-vendor interoperable network. This whitepaper starts with the latest snapshot of O-RAN Architecture, followed by an introduction to the O-RAN Minimum Viable Plan (MVP) for prioritized delivery of a minimum viable set of end-to-end O-RAN solutions applicable in commercial networks using the input of Operator members. Additional use cases defined since the last use cases whitepaper are summarized in the subsequent section. The O-RAN Software Community (OSC) effort to help accelerate the realization of O-RAN vision through open-source projects is followed by the progress of Test and Integration Focus Group (TIFG) for testing and integration of multi-vendor O-RAN solutions through end-to-end test specifications, the guidelines for Open Test and Integration Centers (OTICs) and the organization of global plugfests. This whitepaper is concluded with the Security Focus Group’s (SFG’s) work in progress to ensure secure O-RAN through the development of a security architecture, enabling 5G service providers deploy and operate O-RAN with the same level of confidence as existing networks.

Introduction

O-RAN ALLIANCE is committed to evolving radio access networks (RAN), making them more open, smarter, interoperable, and scalable than contemporary deployments. The Alliance has grown significantly over the last two years and has surpassed 294 companies and over 3,500 people actively participating as of writing this whitepaper.

As described in the executive summary section, the first whitepaper “O-RAN: Towards an Open and Smart RAN” introduced the key O-RAN concepts: (1) Openness to build a more cost-effective network through open interfaces and open hardware, and (2) Intelligence to suit complex, dense, and future-ready networks through embedded intelligence in every layer.

The second whitepaper “O-RAN Use Cases and Deployment Scenarios” introduced an initial set of use cases that leverage the unique benefits of the O-RAN architecture. This includes utilization of machine learning (ML) and artificial intelligence (AI) modules operating through open and standardized interfaces in a multi-vendor network, along with white-box hardware, open hardware reference design, and cloud native deployment through the O-RAN cloudification and orchestration platform, O-Cloud.

Over the last year, along with the progress on a significant number of technical specifications, the Alliance has focused on ways to accelerate and enable the deployment of the O-RAN architecture in commercial networks, which is the focus of this whitepaper “O-RAN Minimum Viable Plan and Acceleration towards Commercialization”.

Two new Working Groups have been formed for better alignment with the industry needs: WG9 to focus on the requirements, solutions, and test specifications for Fronthaul, Midhaul or Backhaul (Xhaul) transport networks and WG10 to focus on the overall OAM architecture, Information Models and Data Models (IMs/DMs).

The Alliance has developed finer details of the architecture, as captured in the Architecture section. One key development has been the adoption of the O-RAN Minimum Viable Plan (MVP) approach where input from Operator members is used for use case prioritization and focus across all working groups to deliver the necessary minimum set of end-to-end specifications for these use cases. In addition to focusing on the delivery of prioritized use cases, member companies continued to innovate and define additional use cases, highlighting the intelligence aspects of O-RAN architecture. The O-RAN Software Community (OSC) continued development of open-source code to help accelerate realization of the O-RAN vision through open-source projects. The Test and Integration Focus Group (TIFG) worked towards the definition of end-to-end test specifications together with guidelines for Open Test and Integration Centers for verification of multi-vendor interoperable O-RAN networks. The TIFG also worked with plugfest organizations to support the ecosystem in joint testing and demonstration of O-RAN architecture and use cases. The Security Focus Group (SFG) was formed for ensuring secure deployment of O-RAN networks through coordination across all working groups toward the definition of a security architecture.

O-RAN Architecture

The high-level view of the O-RAN architecture consists of the network functions, a Service Management and Orchestration framework (SMO) to manage the network functions and an O-Cloud (O-RAN Cloud) to host the cloudified network functions. Figure 1 below shows that the four key interfaces – namely, A1, O1, Open Fronthaul M-plane and O2 – connect the SMO to O-RAN network functions and the O-Cloud. It also illustrates that the O-RAN network functions can be VNFs (Virtualized Network Function), i.e., VMs or Containers, sitting above the O-Cloud and/or PNFs (Physical Network Function) utilizing customized hardware. All O-RAN network functions except the O-RU are expected to support the O1 interface when interfacing to the SMO.

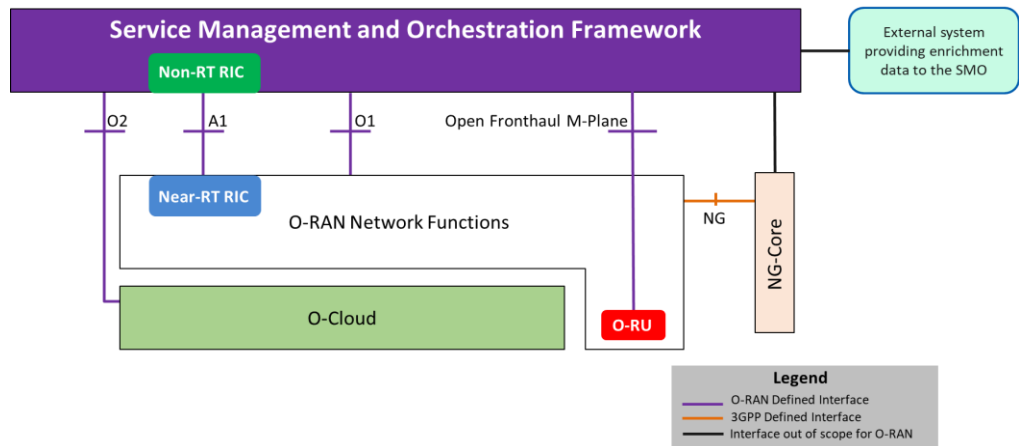


Figure 1: O-RAN High-Level Architecture

The Open Fronthaul M-plane interface, between SMO and O-RU, is to support the O-RU management in hybrid mode, as specified in [7]. O-RU termination of the O1 interface to the SMO is for future study.

Within the logical architecture of O-RAN, as shown in Figure 2 below, the radio side includes Near-RT RIC, O-CU-CP, O-CU-UP, O-DU, O-eNB, and O-RU functions. The E2 interface connects E2 Nodes (i.e., O-eNB, O-CU-CP, O-CU-UP and O-DU) to the Near-RT RIC. Please refer to [5] for details on Near-RT RIC architecture and [4] for details on E2 interface. Although not shown in this figure, the O-eNB does support O-DU and O-RU functions with an Open Fronthaul interface between them.

As stated earlier, the management side includes the SMO containing a Non-RT-RIC function [3] with rApps and R1 interface (not shown in the figure). The O-Cloud, on the other hand, is a cloud computing platform comprising a collection of physical infrastructure nodes that meet O-RAN requirements to host the relevant O-RAN functions (such as Near-RT RIC, O-CU-CP, O-CU-UP and O-DU etc.), the supporting software components (such as Operating System, Virtual Machine Monitor, Container Runtime, etc.) and the appropriate management and orchestration functions. Virtualization of the O-RU is for future study.

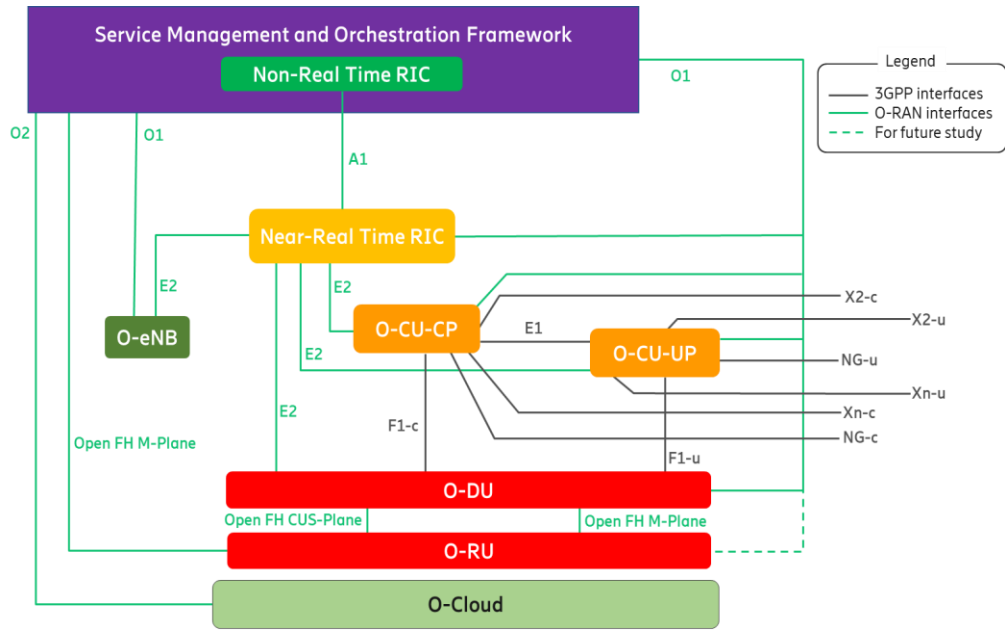


Figure 2: O-RAN Logical Architecture

As shown in Figure 2, the O-RU terminates the Open Fronthaul M-Plane interface from the O-DU and the SMO as specified in [7]. The Open Fronthaul interface (O-DU to O-RU interface) is the Lower Layer Split (LLS) Split Option 7-2x [6].

Please refer to the O-RAN Architecture Description [1][2] for more details.

O-RAN MVP and Prioritized Use Cases for 2021

O-RAN MVP

The O-RAN ALLIANCE set of specifications, open software development, testing, integration, and demonstrations are all driven by the priorities set by O-RAN's operator members. The input of the operator members is used to prioritize delivery of a minimum viable set of end-to-end O-RAN solutions applicable in commercial networks.

The O-RAN Minimum Viable Plan (MVP) provides a prioritized roadmap for these solutions. The process considers the backlog of O-RAN ALLIANCE use-cases, operator surveys and existing cross working-group proposals, in addition to any contribution-driven work items raised in the Alliance working groups. These are synthesized into the Release Plan and Feature Plans for the O-RAN ALLIANCE. Future O-RAN releases will extend the MVP with new features and functionalities as these inputs and priorities evolve.

In MWC Shanghai 2021, the O-RAN MVP published the "Open" package which included a series of open fronthaul, open transport, open hardware, open stack, open cloud features, testing and integration criteria and guidelines for OTIC (Open Testing and Integration Centers).

In the coming MWC LA 2021, the O-RAN MVP will deliver an "Intelligent" package, providing solutions for traffic steering, QoS/QoE optimization, and initial RAN slicing use cases. A baseline will also be setup for security and open cloud infrastructure, with initial security and open cloud API phases complete, as well as testing and integration guidelines and specifications which will continue to support more use cases.

By the end of 2021, the O-RAN MVP can be expected to offer an end-to-end solution for the SMO and a second phase of RAN slicing. Additionally, security and the open cloud API will be enhanced with further updates and new MVP features such as Shared O-RU will be considered for future releases.

"Open" Package

The Open Package was the first package the O-RAN ALLIANCE released in MWC Shanghai 2021. It included several published specifications and work items in the following areas:

- **Open fronthaul** specifications provide fully specified control, user and synchronization interfaces, a management plane interface between O-RU and O-DU and relevant conformance and interoperability tests to enable multi-vendor interoperability.
- **Open transport** defines requirements, solutions and test specifications for Fronthaul, Midhaul or Backhaul (Xhaul) transport networks. The Xhaul Transport Requirement document states requirements for an Open Xhaul transport infrastructure. The Wavelength Division Multiplexing (WDM) based Fronthaul transport document describes best practices for O-RAN fronthaul transport based on WDM technologies. The Xhaul Packet Switched Architectures and Solutions Work Item delivers best practices for deployment of end-to-end packet switching technologies. The Synchronization Architecture and Solutions Specification provides the architecture and solutions for Open Xhaul timing and synchronization and recommends best practices for relevant use cases in Xhaul. The Management Interfaces for Transport Network Elements (TNE) document covers best practices for TNE

management interfaces in O-RAN deployment cases. The Xhaul Transport testing document defines testing requirements for an Open Xhaul transport infrastructure.

- **Open hardware** offers Deployment Scenarios & BS Classes (DSC), Hardware Architecture and Requirement (HAR) and Hardware Reference Design (HDR) specifications for different modes of indoor picocell, outdoor microcell, outdoor Macro cell and open fronthaul gateway. The designs are used to accelerate the use of white-box designs and facilitate RAN ecosystem diversity.
- **Open stack** provides a software reference design for O-CU/O-DU architecture, interfaces, internal modules, and APIs to promote and guide the open-source software development.
- **Open cloud** presents the O-Cloud reference design with fundamental synchronization requirements and real-life deployment recommendations. While O-Cloud can support PTP (Precision Time Protocol) for applications that rely on very precise timing synchronization, these applications need to receive the status of the synchronization state from the O-Cloud. Open Cloud APIs let an application subscribe to an O-Cloud to receive status notifications related to the synchronization state, eliminating the need for proprietary ways to retrieve the synchronization state and keeping the O-Cloud as cloud native. The O-RAN ALLIANCE will continue to provide more open APIs to enable an open cloud for RAN solutions.
 - **The Open Cloud APIs** Phase I targets in MWC LA: Application functionality requires to select O-Cloud platform hardware resources which could provide notifications about the status, initial state and changing states even when they do not have the privilege mode to access it. A REST API was developed to allow Event Consumers (EC) such as a vDU or CNF to subscribe, then cloud infrastructure will provide Event Producers (EP) to enable cloud workloads to receive these events / statuses. The first phase of this API specification provides the status notifications related to PTP. Then Open Cloud API Phase II which targets the next MVP release, will be expanded to provide status notifications related to other hardware resources not just PTP.
 - Also, O-RAN ALLIANCE works on the **AAL (Acceleration Abstraction Layer)**. Application vendors' dependency on the acceleration hardware slows down deployments, makes it not easily portable and more expensive. The AAL interface provides the application vendors the flexibility to integrate their applications with any type of accelerator hardware. Phase I which targets MWC LA will support the Forward Error Correction (FEC) profile with the accelerators that support look aside architecture; Phase II which targets the next MVP release will support accelerators that support inline architecture with possibly additional profiles.
- **Open Testing and Integration Centers (OTICs)** ensures consistency and quality of the testing of O-RAN products and solutions.

“Intelligent Package” - Prioritized Use Cases for 2021

Traffic Steering

The rapid traffic growth and multiple frequency bands utilized in a commercial network make it challenging to steer traffic to achieve performance requirements for all users while considering changes in radio environment and application requirements. Typical controls include cell reselection and handover, modifying load calculations and setting cell priorities. These solutions are cell-centric rather than UE-centric and rarely

incorporate predictions of future network and UE performance. It is hard to adapt these techniques to diverse user demands and variable optimization objectives.

The Traffic steering (TS) use-case addresses this by allowing operators to flexibly configure optimization policies, utilize appropriate performance criteria and incorporate ML to enable intelligent and proactive traffic management and control.

A TS-aware resource optimization policy, such as TS targets and constraints, is generated by the Non-RT RIC and sent to the Near-RT RIC via the A1 interface. The Near-RT RIC determines the optimum changes it can make towards achieving those goals. It may also leverage A1 enrichment information (e.g., the radio fingerprint information for overhead reduction of UE measurement report), which is also sent from the Non-RT RIC. More specifically, the Near-RT RIC triggers E2 procedures to achieve improved network performance fulfilling the criteria identified in the A1 policies and feeds the result of the A1 policy fulfillment back to the Non-RT RIC to update A1 policies.

The current traffic steering use case, which is seen as a high priority of MVP, is defined by the WG1 Use Case Task Group. The feature of TS preference A1 policy has been defined by WG2. By the MWC LA timeframe, the features of mobility control over E2, UE context information report and cell and UE level measurement report from the E2 Node are foreseen to be finished by WG3 to support traffic steering use case phase I.

QoS and QoE Optimization

Highly demanding 5G native applications like Cloud VR and industrial automation are both bandwidth consuming and latency sensitive. Demanding contemporary 4G and 5G applications like online video streaming, multiplayer gaming and connected vehicles are today often handled in a best effort way with little or no application specific optimization. These traffic-intensive and highly interactive applications are not well served by the current semi-static QoS framework which does not efficiently satisfy diversified QoE requirements. These requirements can vary during an application lifetime, especially considering fluctuations of radio transmission capability and applications with dynamic performance requirements. It is expected that QoS or QoE estimation/prediction and service characteristics such as packet size and periodicity from application level can help deal with such uncertainty and improve the efficiency of radio resources, improve the user experience and yield a more efficient use of RAN resources.

Traditional approaches with per-site or per-cell parameter modification to create application or user specific handling policies are implemented without awareness of the application level QoE and the application characteristics, which will not scale with increasing QoE demands of mobile applications.

The O-RAN RIC architecture facilitates RAN and application cross layer QoS and QoE optimization in a real-time way with proactive closed-loop network optimization, both within the RAN and including input from external applications. This will improve the way congested cells are detected and automatically allocate resources based on the end user experience whose demands can change over time. The Non-RT RIC monitors long-term trends and provides QoS or QoE policy guidance to the Near-RT RIC. Empowered with ML techniques and RAN data reported from E2 nodes, the Near-RT RIC can dynamically predict the QoS or QoE performance of the application and then change the RAN behavior to assure the user experience.

This architecture also facilitates the deployment of QoS/QoE behavior as an automated policy control across the network, allowing operators to deploy and fine tune behavior rapidly. The policy generated from the Non-RT RIC and the control or policy from the Near-RT RIC can be instantiated as an rApp and xApp respectively, which will apply the operator's desired QoE configuration or modify behavior in response to external application requests. To further enable RAN aware application optimization, the application hosted by the external systems may leverage RAN information provided by the Non-RT RIC or Near-RT RIC to proactively adjust its behavior. For example, a cloud VR application may adjust the codec rate or resolution based on RAN reported Information.

The current QoS and QoE Optimization use case in O-RAN is defined by the WG1 Use Case Task Group. The feature of QoS and QoE related A1 policy has been defined by WG2. The E2 service model support is foreseen to be finished by WG3 in the MWC-LA timeframe; this will enable additional performance measurements in E2SM-KPM to monitor cell, UE group and UE level performance measurement and radio resource utilization and specify RIC services in E2SM-RC to configure and report event-triggered E2 node configuration, UE context information and execute control/policy-based actions over E2.

RAN Slicing and SLA Assurance

Network slicing is a prominent feature of 5G that provides end-to-end connectivity tailored to specific service requirements such as the support for very high data rates, traffic densities or very low latency based on a Service Level Agreement (SLA). O-RAN slicing and RAN Slice SLA assurance features are prioritized in 2021 and planned to be delivered in phases. The current work items focus on achieving deployment of RAN slice subnets with O-RAN capabilities within the open RAN environment as well as bringing slice SLA assurance mechanisms to detect and prevent SLA violations.

O-RAN's open interfaces combined with its AI/ML based innovative hierarchical architecture enables challenging RAN slice SLA assurance mechanisms to be implemented, which are critical to allow operators to realize the opportunities of network slicing and offer these new products efficiently and at scale. Based on RAN specific slice SLA requirements, the Non-RT RIC and Near-RT RIC fine-tune RAN behavior to assure RAN slice SLAs dynamically. Utilizing slice specific performance metrics received from E2 Nodes, the Non-RT RIC monitors long-term trends and patterns regarding RAN slice subnets' performance and trains AI/ML models to be deployed at the Near-RT RIC. The Non-RT RIC also guides the Near-RT RIC using A1 policies, with possible inclusion of scope identifiers (e.g., S-NSSAI) and statements such as KPI targets. The Near-RT RIC enables optimized RAN actions through execution of deployed AI/ML models or other slice control/slice SLA assurance xApps in near-real-time by considering O1 configuration (e.g., static RRM policies), received A1 policies and received slice specific E2 measurements.

Support for RAN slice SLA assurance and O-RAN slice subnet management and provisioning are planned to be delivered in two phases. Phase 1 (MWC-LA timeframe) aims to deliver the detailed use cases, requirements, and initial specifications such as initial slice SLA A1 policies and basic TN slicing support. In addition, necessary interactions regarding O-Cloud, SMO and transport network will be addressed along with slice aware functioning of Non-RT RIC and Near-RT RIC. Phase 2 (next MVP release) builds on the first phase with features such as extended O-CU and O-DU slicing related information models, additional A1 policies, E2SM support, dynamic and orchestrated mapping of intra-DC slice links to inter-DC transport links and marking of slice flows between O-CU and O-DU for intra-DC and inter-DC transport mapping.

SMO (Service Management and Orchestration framework)

In the O-RAN architecture, the Service Management and Orchestration framework (SMO) is responsible for managing and orchestrating the multi-vendor disaggregated O-RAN Network Functions such as Near-RT RIC, O-CU-CP, O-CU-UP, O-DU, and O-RU. It also supports orchestration of O-Cloud infrastructure and management and the deployment of the O-RAN network functions on the O-Cloud as well as supporting SMO Life cycle management for RIC applications. Non-RT RIC enables “intelligence” in the SMO for controlling the RAN elements. The package includes but is not limited to the following:

The SMO provides an integration fabric that enables orchestration and management of multi-vendor RAN functions. It also supports data services to provide efficient data collection, storage and movement capabilities for multi-vendor RAN functions. Other domain orchestration functionality (e.g., core, transport) can be optionally installed on or integrated into the SMO. For example, the SMO can integrate third-party Network Management System (NMS) or orchestration platforms.

The SMO provides an **Operation and Maintenance (OAM) interface** towards O-RAN functions to achieve FCAPS management, software management, and file management capabilities. The common data model supports configuration parameters, performance counters, fault supervision and software management aspects of RAN functions, complementary to 3GPP SA5 parameters. The SMO utilizes the O1 interface for OAM towards the O-RAN Network Functions except for the O-RU. For hybrid management, the Open FH M-Plane is utilized between the SMO and the O-RU while for hierarchical management the SMO manages the O-RU indirectly via the O-DU.

The **Non-RT RIC** is a central component of the SMO, enabling non-real-time control and optimization of RAN elements and resources, and supporting AI/ML workflows. The Non-RT RIC provides fine grained policy-based guidance to the Near-RT RIC through the A1 interface and is comprised of the Non-RT RIC framework and Non-RT RIC applications (rApps):

- **A1** interface offers capabilities to inter-connect the Non-RT RIC and Near-RT RIC through an A1 Policy Management Service and an A1 Enrichment Information Service. This can provide different level guidance (cell/slice/UE) and enrichment information to enable various RAN intelligence use cases, enabling interoperability by different manufactures. By the MVP release after MWC LA, in addition to the current A1 specifications, an A1 ML Management Service is planned to be introduced.
- **R1** Interface allows the Non-RT RIC framework and the rApps to consume and provide R1 services which are a collection of services including, but not limited to, service registration and discovery services, authentication and authorization services, AI/ML workflow services and A1, O1 and O2 related services. Phase I focuses on R1 general aspects & principles, transport layer and application layer protocols to enable a collection of services such as service registration, discovery, AI/ML workflow and A1-related services.

The SMO also provides a unified approach for application and infrastructure **lifecycle management** which is applicable for all cloud-native applications and cloud-native network functions from different vendors. The O-RAN OAM Architecture defines a framework for common app lifecycle management, including common app package, onboarding, validation, model training (if needed) and deployment.

In addition, the SMO provides O2 interfaces towards the O-Cloud to support orchestration of O-Cloud infrastructure management and deployment of the O-RAN network functions on the O-Cloud.

- The O2 interface supports Infrastructure Management Services (IMS) services, i.e., services related to O-Cloud infrastructure resource management (e.g., inventory, monitoring, provisioning, software management and Lifecycle management).
- The O2 interface also supports Deployment Management Services (DMS) services, i.e., services for managing the life cycle of software deployments using O-Cloud infrastructure resources.

By the MVP release after MWC LA of 2021, the O-RAN will deliver an end-to-end SMO solution, which include framework, O1, O2 interface and related IM/DM, FM/PM dictionary. O-RAN will continue to enhance RAN automation, management and deployment of cloud native infrastructure and applications and continue to complete more data models to support more use cases.

Massive MIMO Optimization

Massive MIMO (mMIMO) is a key NR air interface technology. Thanks to its advanced beam forming techniques, mMIMO can improve coverage, capacity, latency and reliability. The objective of the mMIMO Optimization use case is to enable proactive and continuous improvement of cell- and/or user (group)-centric QoS in multi-cell and even multi-vendor mMIMO environments. mMIMO sub features include advanced network management technologies like beam-based load balancing and mobility handling, beam failure prevention and dynamic cell coverage adaptation, especially in highly dense urban environments. The high number of configuration parameters per antenna array, the massive volume of measurement data and the complexity, adaptivity as well as the partly near-real time requirements suggest the application of ML techniques in mMIMO analytics and inference, which are well supported by the O-RAN architecture and workflow.

Both the Non-RT RIC and Near-RT RIC ingest measurements descriptive of the current network state (e.g., traffic, coverage, interference, etc.) over a whole cluster of cells. Additionally, non-real time analytics and inference results are usually fed into E2 nodes which are real time rather than near-real time processes. This will be used to generate the appropriate mMIMO configuration and policy that optimizes the network over an entire coverage area and according to operator-defined objectives.

O-RAN MVP has identified mMIMO Optimization as one of the top priority use cases for the industry and has defined a mMIMO Optimization feature plan including the following sub features: a) the Non-RT Grid-of-Beams (GoB) Beam Forming optimization for coverage and capacity (e.g., number of beams, beam boresights, vertical beam widths, horizontal beam widths, transmission powers), b) the Near-RT adaptive beam shaping and Beam-based L3 Mobility Robustness Optimization (MRO, adapting mobility parameters such as beam-specific inter-cell handover threshold offsets on an inter-beam granularity), c) L1/L2 beam management optimization, and d) non-GoB Beam Forming optimization. This O-RAN MVP feature plan will progress towards specification by 2022.

New Use Cases and Updated Deployment Scenarios

While as described in the previous section, Traffic Steering, QoS/QoE Optimization, RAN Slicing and Slice SLA Assurance and Massive MIMO use cases are included in O-RAN MVP based on operator member's input, the O-RAN community continues to define new and innovative use cases, leveraging the benefits of the O-RAN architecture. This section summarizes new use cases introduced since the publication of "O-RAN Use Cases and Deployment Scenarios" in February 2020.

New Use Cases

Multi-vendor Slices

In this use case, multiple slices can be defined by different vendors, each with independent O-DU and O-CU functions. The O-RAN Architecture creates an opportunity for operators to select appropriate O-DU and O-CU functions to meet their service requirements. For example, a vendor may have a scheduler or other features optimized for eMBB service, while another may focus on URLLC services. This flexibility may also help to introduce a new vendor into an operator's network without a complete swap out. This use case also enables RAN sharing among different operators, even when their RAN vendors or placement of RAN functions may differ. In a multi-vendor situation, it relaxes restrictions among operators when it comes to sharing RAN equipment and resources. Finally, if an existing vendor providing a certain pair of O-DU and O-CU functions withdraws from the market, operators can deploy new components from other suppliers in a targeted way, reducing risk for operators' business continuity.

To realize multi-vendor slices, coordination between O-DUs & O-CUs is required to assign radio resources appropriately and without conflict. Resource allocation between slices or O-DU/O-CUs is provisioned through the O1/A1 interface and each pair of O-DU and O-CU will allocate radio resources to each customer within available radio resources allocated by Near-RT RIC and/or Non-RT RIC as a first phased solution.

O-RAN DSS

Dynamic Spectrum Sharing (DSS) technology dynamically allocates spectrum to 4G and 5G cells sharing the same spectrum block based on traffic demand. DSS helps mobile operators quickly rollout 5G services in shared 4G bands and over time seamlessly transition to full 5G networks adapting to the varying traffic loads of both networks in a cost-effective manner without the resource underutilization issue associated with spectrum re-farming. DSS can also improve 5G coverage as the networks evolve to 5G without the high cost and increased Inter Cell Interference suffered by dense deployment in C-band or mmWave bands.

In contrast to synchronizing 4G and 5G schedulers using proprietary resource management functions, O-RAN RIC based DSS uses the O-RAN open interfaces and third-party RIC applications to provide a vendor-agnostic DSS solution to co-existing 4G/5G deployments that use multi-vendor RAN equipment. RIC based DSS can also leverage near-real-time and non-real-time RAN data from a larger number of cells and AI/ML based intelligent processing capability, provided or supported by the O-RAN Architecture. This supports traffic load prediction and DSS resource management and control functions, which are the critical building blocks to help achieving better coordination gain and overall improved scheduling, throughput, and interference efficiency.

RAN Slice Resource Allocation Optimization

Network Slicing is essential to 5G, enabling new services across manufacturing, autonomous driving, gaming, smart city and many more via the provision of network functionality and features dedicated to specific users or even specific customers as a virtual overlay on the same physical hardware. Slices require different or even contrasting QoS requirements; it is a challenging task to allocate dynamically and efficiently among multiple network slice instances in overlapping geographies and logical networks.

Some 5G services have varying characteristics, with network traffic that can be sporadic and with varying usage patterns in terms of time, location, UE distribution and specific applications. For example, IoT sensor applications might be configured to run during off-peak hours or weekends, while special events such as sporting events or concerts can cause traffic demand to shoot up at very concentrated locations. Further ahead, connected vehicles will require URLLC-style services in the morning or afternoon rush hours to support automated driving and/or eMBB services for infotainment and other services, while mobile or fixed wireless users will tend to consume eMBB services to watch video streaming at night in residential areas.

The RAN Slice Resource Allocation Optimization function in O-RAN is designed to train an AI/ML model based on the huge volume of performance data collected from O-RAN nodes to predict the traffic demand patterns of different service types that are realized in network slice instances in different times and locations. This information is used to automatically optimize the resource allocation for network slice instances accordingly to improve user experiences and maintain optimal network usage efficiency.

Local Indoor Positioning in RAN

Cellular network-based positioning is an important technology for connected and automated industries, individuals and operators. It is especially useful in local indoor scenarios, for instance, in a shopping mall where value-added services such as local indoor navigation and shop recommendations can leverage this information, or as a mechanism to determine space utilization or equipment tracking in an office setting. In manufacturing or construction operations, real-time safety warnings based on user locations can be used to remind operators of impending danger or to provide automated lockout if users enter unsafe areas. NR positioning introduced in 3GPP Rel.16 has a long route for messages between eNB/gNB and a centralized Location Management Function (LMF). It may suffer network jitter and lead to inaccurate UE location results. For some real-time scenarios, it would be better to directly deploy the positioning function at the RAN location and expose UE locations to local or edge-hosted applications.

In the context of the O-RAN architecture, the positioning function can be deployed as a positioning xApp in the Near-RT RIC. The positioning xApp computes the UE location and velocity based on positioning measurements obtained via the E2 interface. One example is where a distributed indoor small base station in the operator's domain can provide a UE positioning service by adding a positioning service xApp in the Near-RT RIC. Positioning related information, based on uplink SRS data, may be provided by the small base station to the positioning xApp inside the Near-RT RIC. By leveraging multi-point positioning, field strength positioning and other positioning algorithms e.g., UT-DOA, the positioning xApp inside the Near-RT RIC can calculate a real-time position of UE. The positioning service can be leveraged by other RAN optimization xApps in the Near-RT RIC to improve the network and user experience. The xApp may

also pass the positioning results to the SMO for further exposure to edge applications nearby.

Massive MIMO SU/MU-MIMO Grouping Optimization

Given Massive MIMO's importance, in addition to the existing set of Massive MIMO use cases that are defined in "O-RAN Use Cases and Deployment Scenarios" and included in O-RAN MVP, a new set of use cases are being considered. One such use case, Massive MIMO SU/MU-MIMO Grouping Optimization, focuses on providing capacity enhancements through beamforming of transmitted signals and spatially multiplexing data streams for both single-user (SU) and multi-user (MU) MIMO. SU-MIMO can provide spatial multiplexing per user multi-layer or better coverage by beamforming gain. MU-MIMO provides higher spectrum efficiency, with up to 16 transmission layers transmitted across a population of users in the same area and using the same time/frequency resource. In a real deployment, SU-MIMO and MU-MIMO have variable benefits: for example, the performance of MU-MIMO is sensitive to UE velocity and gains are reduced in low traffic scenarios. Conversely, SU-MIMO is relatively tolerant to mobile UEs and can provide gains even with low traffic but might have lower overall gain in dense areas with static or nomadic users.

The objective of SU/MU-MIMO grouping optimization is to use information on the specific scenario to select an appropriate transmission method for users. The use case identifies and separates UEs into SU-MIMO and MU-MIMO types to save SRS resources, reduce implementation complexity, improve transmission efficiency and optimize RRC configuration. For example, based on a UE's GPS information such as velocity and location, and traffic-related information from the application server, combined with the RAN data from E2 nodes, the Non-RT RIC and Near-RT RIC can train the AI/ML model and infer which group (SU-MIMO or MU-MIMO group) is appropriate. The grouping list will be sent with O-RAN policy via the A1/E2 interface to help E2 nodes with real-time configuration, scheduling, and so on.

O-RAN Signaling Storm Protection

Mobile device design and regulation mandate restrictions on the number of signaling events generated by a device, keeping the average volume relatively low. Still, occasionally, there are devices that violate these restrictions. Attackers take advantage of these devices by recruiting them into a botnet to launch multiple signaling storms that cause DDoS over the mobility network. There are no controls to mitigate such coordinated attacks in the RAN and signaling flow throttling in the core does not distinguish between benign and malicious devices. Furthermore, it is preferable to block these attacks at the network's edge as close to the source as possible.

The O-RAN architecture opens an opportunity to add the necessary logic at the RAN to detect and protect against signaling storms. Adding a Signaling Storm Protection (SSP) xApp to the Near-RT RIC provides detection and mitigation against such rogue devices. The SSP xApp applies ML techniques to identify anomalous signaling patterns and distinguish benign devices from harmful ones. Then, using the E2 interface, it throttles signaling from the rogue devices while allowing the messages from regular devices to go uninterrupted as much as possible. In some cases, the signaling storm may be widespread across different network regions that more than one Near-RT RIC covers. The Non-RT RIC can coordinate the event based on enrichment information by instructing and configuring parameters of the Near-RT RICs located in signaling storm regions. The Signaling Storm Protection capability is crucial, especially in times of increases in the number and variety of cellular IoT devices prone to vulnerabilities. The SSP xApp introduces programmable security capability that is

flexible to evolve alongside coming security threats. Having SSP at the RAN provides significant savings by avoiding unnecessary redundant core infrastructure to handle signaling storms.

Congestion Prediction and Management

Large-scale commercial cellular networks are in a constant state of flux, with traffic movement and cell congestion leading to radio link failures, handover failures, poor data rates and other undesirable effects across the coverage area. Network congestion is a crucial problem for telecom operators as it affects the Quality of Service (QoS) of the users directly, and many solutions are available to address this such as traffic steering or offload (across carriers, Wi-Fi, etc.) or antenna techniques (Cell Split, Higher order MIMO, etc.). Congestion mitigation is critical for operators to provide a good experience to their subscribers.

However, the congestion patterns in the network are often not fully characterized and mitigation is performed reactively over timescales dictated by operational and capital investment opportunities at the expense of a prolonged user experience degradation. The main objective of this use case is to use the embedded intelligence of O-RAN to predict congestion ahead of time so operators can invoke cell congestion mitigation solutions before the congestion is predicted to happen. Congestion Prediction & Management (CPM) provides a proactive and intelligent approach to congestion handling in the base station by analyzing radio resource utilization and taking timely corrective action to mitigate any potential congestion in the system. In the CPM architecture, E2 node statistics are collected by the data collector in the SMO via the O1 interface and then pre-processed with appropriate tagging and conversion into KPIs. After pre-processing, the data is shared with the Non-RT RIC which will invoke the corresponding training model/application in an AI server which will learn and predict future traffic for the next hour/day/month, with prediction windows configured by operators depending on available data and periodicity. The output inference from this process will contain the cell ids, information about whether the cell is congested or not, time stamp of cell congestion and future predicted KPI values.

Industrial IoT Optimization

Industrial Internet of Things (IIoT) applications which require high reliability, such factory automation or the transport and automotive industry, require new and diverse radio features and optimization techniques. Key features for IIoT in 5G are being designed and deployed such as data duplication and multi-connectivity enhancements, Time Sensitive Networking, QoS and scheduling enhancements.

The objective of the IIoT optimization use case is to guarantee reliability and resource utilization efficiency in the above scenarios by optimizing the enhancement mechanisms of these related features to meet the overall application requirements. The benefits of the O-RAN Architecture for IIoT use cases include the ability to use external enrichment information provided by MEC server, APPs server or industrial control platform as input to the SMO. Based on this external service information and RAN information, the Non-RT RIC uses an AI/ML model to infer IIoT key techniques configurations such as whether the QoS flow shall map to DRB supporting duplication and Ethernet Header Compression (EHC) or whether high priority traffic can preempt transmission resource by cancelling the low priority traffic transmission. The SMO deploys or updates AI/ML models to the Near-RT RIC via the O2 interface to infer and make decision of PDCP duplication on/off, RLC entity selection and high priority traffic preemption.

O-DU Pooling to Achieve RAN Elasticity

Pooling of gNB resources provides benefits like having to maintain only a single large, controlled environment for many cell sites, which can lead to better resiliency and agility to deal with unexpected failures and changes in demands. Capacity can be added incrementally and assigned to cell sites dynamically, resulting in better overall utilization. More importantly, gNB pooling provides significant statistical multiplexing benefits due to better resource consolidation. The total resources required from the shared pool will be less than total resources required across distributed locations, because the peak of the sum of the traffic will generally be lower than the sum of the individual cell site traffic peaks. gNB pooling enables more flexible capacity dimensioning based on different percentiles of resource consumption rather than for the peak consumption.

O-DU Pooling is enabled by the switched ethernet nature of the O-RAN Open Fronthaul with LLS (option 7-2x split) and decoupling of baseband hardware and software enabled by the O-Cloud infrastructure. Coarse-grained Class 1 pooling allows flexible re-mapping of O-RUs to the O-DU hardware periodically. The O-RAN Architecture standardizes the SMO functions needed to enable pooled gNBs including management of O-DU & O-RU (registration/de-registration of O-RU to O-DU, registration of RU locations to SMO), orchestration features via O1/O2 interfaces and procedures for remapping O-RUs from one O-DU to another O-DU and orchestration of the control loops provided by the Non-RT RIC. O-RAN also helps with the procedures for hitless draining of cell traffic from one O-DU to another O-DU, besides monitoring the performance/fault metrics to determine DU load and KPIs. Fine-grained Class 2 pooling dynamically assigns subsets of traffic to O-DU resources in the shared O-DU pool. The O-RAN Architecture helps with maintaining the affinity of control and downlink data through a cluster aware scheduler, maintaining affinity between control and uplink data through a Cloud Access Switch or in the O-RU, configuring the O-DU to support a wider range of sectors through the cluster aware scheduler, monitoring the performance/fault metrics to determine DU load and KPIs and triggering the scale-in/scale-out workflows in the SMO that manages the O-DU Pool.

Updated Deployment Scenarios

The upcoming Cloud Architecture and Deployment Scenarios document introduces an extension of the deployment scenario E for macrocell deployments (see Figure 3), targeting distributed cloudified RAN deployments where the cloudified O-DU is located at the cell site. It is an extension of the open hardware approach that is used in WG7, where the O-DU 7-2 (of O-RAN WG7 OMAC HAR 0-v01.00) can be a cloudified network function. In this small-scale scenario, HW acceleration is optional, and the cloud platform could be physically located near or at the bottom of the tower and be associated with several non-virtualized O-RUs implemented with the open HW design, possibly but not necessarily in the same chassis.




Figure 3: Deployment Scenario E.1

O-RAN Open-Source Community

Specification development is not the only work done by the O-RAN ALLIANCE. From the first day when the Alliance is set up, open source has been viewed as one of the powerful forces to accelerate the realization of the O-RAN vision. In April 2019, joining force with Linux Foundation, the O-RAN ALLIANCE set up an O-RAN Software Community (OSC) to develop software for O-RAN solutions in an open-source way. The OSC is focused on aligning with the O-RAN ALLIANCE's open architecture and specifications to achieve a solution that can be utilized for industry deployment and will enable the development of open-source software enabling modular, open, intelligent, efficient and agile disaggregated radio access networks.

To this end, the OSC has set up a series of projects to cover every critical O-RAN entity, including Near-RT RIC, Non-RT RIC, Cloudification and virtualization platforms, O-RAN Central Unit, O-RAN Distributed Unit, Operation and Maintenance, Simulations, Service Management and Orchestration, RIC Application and a test and integration effort to provide a working reference implementation. These projects are guided by a Requirements and Software Architecture Committee (RSAC) and a Technical Operating Committee (TOC). These projects also work with other adjacent open-source networking communities, so the O-RAN Software Community can enable collaborative development across the full operator network stack with strong community alignment.

Cross-collaboration among related open-source projects is key to the OSC's success. In this respect a collaboration with the Open Networking Foundation (ONF) on some related RIC topics and with ONAP on the SMO aspects are ongoing, which aim at harmonization of the ecosystem.

The OSC publishes two releases every year. The first release, "Amber", came out in December 2019 with the second release, "Bronze", published six months later. The latest "Cherry" release has been published in December of 2020 and the current focus is on the "Dawn" release targeting in June 2021.

The main features delivered in C release ("Cherry") include:

- Policy based Traffic Steering using the A1 and E2 interfaces
- Automated Monitoring and Health check of selected components using the O1 interface
- Configuration, Fault and Performance Management Services aligned with OAM models and specifications newly approved by O-RAN
- Life Cycle Management Framework of rApps and xApps using SMO
- O-DU Low and High pairwise testing in TIMER mode in O-RAN Software Community lab
- Simulators for testing and integration, such as E2, A1, Open Test Framework, and more

The upcoming D release features will include:

- Support for a new closed loop processing use case, i.e., front-haul failure detection and recovery
- Security and CII badging progress for all OSC projects
- New performance measurement xApp (Bouncer), Load Prediction xApp and other enhanced versions of Anomaly Detection, Traffic Steering and KPI monitoring xApps

- O1 command line utility to simulate DMS functionality, inclusion of Time-series database (InfluxDB), load and scalability testing using the new Bouncer xApp
- Support for A1 enrichment information (A1-EI), multi-version support of A1-AP and pre-spec A1, runtime configuration of A1 Policy Management Service
- Support for closed loop use case for Fault alarm notification (FM) and translation from YANG to VES, PNF registration using VES and support for CallHome functionality via TLS
- O-DU-HIGH support of DL and UL data transmission for FDD (20 MHz) and TDD (100 MHz)
- Integration of O-DU-HIGH and O-DU-LOW with an RU emulator and real UE attach, supporting cell stop and restart
- O1 simulator for O-DU and O-RU for supporting closed loop use case
- Preliminary support for O-Cloud DMS and IMS implementation
- Support for O1/VES interface (version 7.2.1), Implementation of NETCONF client in SMO, YANG model verification stub for DU to facilitate closed loop use case

For more information, readers are encouraged to visit OSC's website at www.o-ran-sc.org.

O-RAN Testing and Integration – OTIC, Global Plugfests and More

The O-RAN Test and Integration Focus Group (TIFG) defines O-RAN's overall approach for testing and integration, which includes coordinating test specifications across all the different Working Groups. The current focus areas in the TIFG are:

1. Setting guidelines for the **Open Test & Integration Centers (OTICs)**
2. Developing processes for the **Certification and Badging** of vendor equipment
3. Creating an **End-to-End Test Specification** for the full O-RAN system as a black box
4. Organizing the **O-RAN ALLIANCE Global Plugfests**

Open Test and Integration Centers (OTICs)

Open Test and Integration Centers (OTICs) are vendor-independent, open and qualified physical labs that provide a collaborative, open and impartial working environment to support the wide adoption of O-RAN specifications. OTICs promote the openness of O-RAN ecosystem via testing services, lab and field trials and community events (e.g., speaker sessions, workshops, tutorials). Specific services and activities that an OTIC can host include:

1. Test and verify the conformity of RAN equipment to O-RAN conformance test specifications.
2. Test and verify the interoperability of RAN equipment from different vendors (or the same vendor) based on O-RAN interoperability test specifications.
3. Conduct functional and performance (load, capacity) tests of end-to-end systems or sub-systems.
4. Demonstrate implementations and solutions based on O-RAN specifications via plugfests and proofs of concept (PoC).
5. Foster and develop the technical capabilities of integrators via workshops, tutorials, etc.
6. Provide feedback to the O-RAN community about the experiences with O-RAN specifications acquired during the testing, enabling implementation-driven specification.

O-RAN members and contributors that exhibit vendor neutrality can apply to host an OTIC. The TIFG has published criteria and guidelines for interested OTIC applicants in the specification "O-RAN Criteria and Guidelines of Open Testing and Integration Center 2.0 - November 2020 (O-RAN.TIFG.CGofOTIC.0-v02.00) [10].

Currently, there are two OTICs already approved (European OTIC in Berlin and European OTIC in Torino) and more OTICs are being considered in various regions across the globe.

Certification and Badging

Procedures for certification and badging of vendor equipment is a critical function to ensure the conformance, interoperability and performance of equipment based on O-RAN ALLIANCE specifications. The TIFG has recently published the "Certification and Badging Processes and Procedures 1.0 (O-RAN.TIFG.Cert-Badge.0-v01.00)" [11] document that specifies the processes and associated detail technical procedures that OTICs and various other testing entities should adopt:

1. Verify compliance of Devices Under Test (DUTs) with O-RAN interface or reference design specifications based on conformance test specifications (conformance certification).

2. Assess interoperability of DUT pairs with O-RAN interfaces based on IOT specifications (interoperability badging).
3. Assess End-to-End system integration of System Under Test (SUT) with O-RAN interfaces based on E2E system testing specifications (E2E system integration badging).

End-to-End Test Specification

In addition to coordinating the test specifications across all the O-RAN working groups, the TIFG is also responsible for the overall End-to-End test specification and has recently published the “End-to-End Test Specification v1.0 (O-RAN.TIFG.E2E-Test.0-v01.00) [12]. This test specification is focused on validating the end-to-end system functionality, performance, and key features of the O-RAN system as a black box. In this initial version, TIFG has focused on the most essential functional, performance, services and security tests, and continues to work on expanding the test cases to truly reflect real-world end-user experience to ensure that equipment provide uncompromising performance and security to any RAN networks based on O-RAN ALLIANCE specifications.

Global Plugfests

O-RAN plugfests support the ecosystem players in testing and integration of their solutions, allowing for a coordinated global effort that allows for timely advancement of the solutions from the various vendors based on the latest O-RAN ALLIANCE normative core and test specifications. TIFG is responsible for organizing the plugfests with the support of the O-RAN office and has already successfully conducted two global plugfests, the first in December 2019, and the second in September 2020. The TIFG is actively organizing the third global plugfest, which is currently planned to conclude in November 2021.

Secure O-RAN

The O-RAN ALLIANCE's goal is to achieve a secure, open and interoperable RAN. Following examples set by standards development organizations such as 3GPP and IETF, the O-RAN ALLIANCE has assembled a committed group of security experts in its Security Focus Group (SFG) to develop an O-RAN security architecture that enables 5G service providers to deploy and operate O-RAN with the same level of confidence as a 3GPP defined RAN. The SFG is leading and collaborating with the O-RAN ALLIANCE's other working groups to ensure that O-RAN is secure by design.

O-RAN's openness and disaggregated architecture provide inherent security benefits:

- Open-source software enables transparency and common control
- Open interfaces ensure use of and interoperability of secure protocols and security features
- Disaggregation enables supply chain security through diversity

The O-RAN Architecture includes new interfaces and functions, expanding the threat surface to introduce new security risks. O-RAN also shares common security risks with virtual and cloud-based deployments due to use of open-source software, white-box hardware and the multi-party relationship between the operator, cloud provider and system integrator. Providers must take a risk-based approach with O-RAN deployments. Recognizing these security challenges and the criticality of a secure RAN, the SFG is following 3GPP security design practices and industry best practices to identify security requirements and solutions that enable O-RAN to deliver the level of security expected by 5G network operators and users.

The SFG is developing a series of documents to secure O-RAN architecture. The Security Threat Modeling and Remediation Analysis document [9] provides an in-depth analysis of threats to O-RAN's assets including data, interfaces and components. The Risk Assessment document provides an impact assessment and severity level for each risk, based upon ISO 27005 [1], with consideration of internal and external attacks while applying the U.S. NIST SP 800-207 Zero Trust Architecture [14] which assumes there is no implicit trust granted to assets or user accounts based solely on their physical location, network location or ownership. Security Protocol Specifications [8] and Security Requirements Specifications documents provide high-level requirements for use and configuration of SSHv2, TLS 1.2 and 1.3, DTLS 1.2, and IPsec on O-RAN's interfaces. The SFG will be producing a security guidelines document for vendors to securely develop, contribute and use open-source software in the O-RAN Software Community (OSC). The SFG is also currently evaluating the application of leading-edge technologies such as blockchain for mutual authentication and distributed identity management, zero-touch provisioning and AI to secure O-RAN.

The SFG actively collaborates with other working groups to address O-RAN's security risks for potential attacks on confidentiality, integrity, and availability:

- Near-RT-RIC and its Third-Party xApps, as shown below
- Open Fronthaul Interface [O-RAN LLS(7-2x)] on the M-Plane and C/U/S-Plane, as shown below
- O1 Network Access Control Model (NACM)
- UE identification for privacy, as shown below

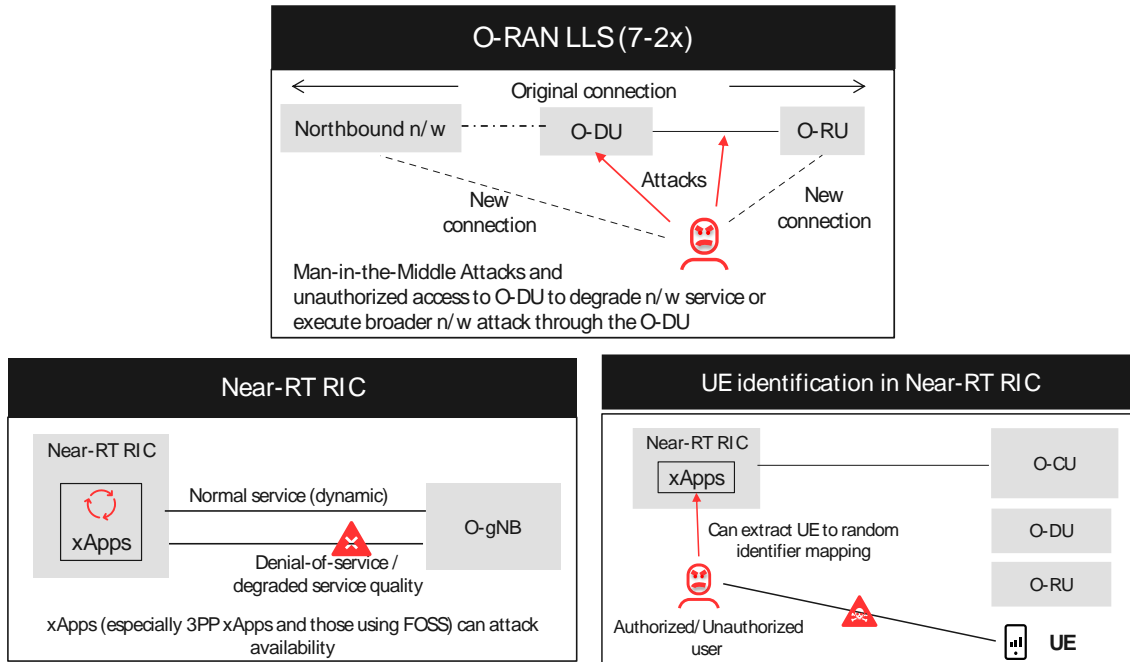


Figure 4: O-RAN Security Risks [8] [13]

Network security is a continuous process. Changes to risk due to evolving threats, attack vectors and security control technologies should be periodically re-assessed by all stakeholders, including vendors, operators, cloud providers and system integrators. O-RAN will be a complementary part of 5G's broader progression to greater security and the O-RAN ALLIANCE's SFG is working to ensure it is secure.

Conclusion

The O-RAN ALLIANCE continues to evolve the RAN through intelligence and openness in the RAN with standardized open interfaces and white-box hardware powered by an innovative architecture with AI/ML based hierarchical controllers.

This third whitepaper extends the previous whitepapers and outlines the Alliance's focus on several key areas to accelerate and enable the introduction of the rich capabilities of O-RAN architecture in commercial networks and to provide a minimum set of end-to-end specifications for selected use cases for the deployment of a secure, multi-vendor interoperable network.

This whitepaper has covered the latest snapshot of O-RAN Architecture, O-RAN MVP for prioritized delivery of minimum viable set of end-to-end O-RAN solutions, additional use cases defined since the last use cases whitepaper, the efforts of O-RAN Open Software Community (OSC), Test and Integration Focus Group (TIFG) and Security Focus Group (SFG).

The O-RAN ALLIANCE will continue to work towards the vision of a fully open and intelligent RAN through the definition of innovative use cases and a secure network architecture that can be deployed commercially with interoperable verified multi-vendor solutions.

References

- [1] ISO/IEC 27005:2018, Information technology — Security techniques — Information security risk management, <https://www.iso.org/standard/75281.html>, July 2018
- [2] O-RAN.WG1.O-RAN-Architecture-Description-v04.00: “O-RAN Working Group 1; O-RAN Architecture Description v04.00”
- [3] O-RAN.WG2.Non-RT-RIC-ARCH-v01.01: “O-RAN Working Group 2; Non-RT RIC Functional Architecture v01.01”.
- [4] O-RAN.WG3.E2GAP-v01.01, “O-RAN Working Group 3, Near-Real-time RAN Intelligent Controller, Architecture & E2 General Aspects and Principles”
- [5] O-RAN.WG3.RICARCH-v02.00, “O-RAN Working Group 3, Near-Real-time RAN Intelligent Controller, Near-RT RIC Architecture”
- [6] O-RAN-WG4.CUS.0-v06.00: “O-RAN Fronthaul Working Group; Control, User and Synchronization Plane Specification v06.00”
- [7] O-RAN-WG4.MP.0-v06.00: “O-RAN ALLIANCE Working Group 4; Management Plane Specification v06.00”
- [8] O-RAN Security Protocols Specifications-V1.0, O-RAN ALLIANCE, March 2021
- [9] O-RAN Threat modeling and remediation analysis, v1.0, O-RAN ALLIANCE, March 2021
- [10] O-RAN.TIFG.CGofOTIC.0-v02.00: “O-RAN Criteria and Guidelines of Open Testing and Integration Centre 2.0,” Nov. 2020
- [11] O-RAN.TIFG.Cert-Badge.0-v01.00: “O-RAN Certification and Badging Processes and Procedures 1.0,” March 2021
- [12] O-RAN.TIFG.E2E-Test.0-v01.00: “End-to-end Test Specification v1.0,” March 2021
- [13] ‘Security Considerations of Open RAN’, Ericsson, <https://www.ericsson.com/4a4b77/assets/local/security/security-considerations-open-ran.pdf>, August 2020
- [14] U.S. NIST SP 800-207, Zero Trust Architecture, <https://csrc.nist.gov/publications/detail/sp/800-207/final>, August 2020