



ONF's Software-Defined RAN Platform Consistent with the O-RAN Architecture

## Executive Summary

SD-RAN is ONF's new exemplar platform for 3GPP compliant software-defined RAN that is consistent with the O-RAN architecture. It is cloud-native and is built on ONF's well-established, operator-approved and deployed platforms, such as ONOS and Aether. Starting with an ONOS-based near real-time RAN Intelligent Controller (RIC), the exemplar platform aims to develop open source components for the control and user planes of the Central Unit (CU-C and CU-U, respectively) and the Distributed Unit (DU) of the disaggregated RAN in close coordination with the O-RAN Alliance and O-RAN Software Community.

Oğuz Sunay  
Shad Ansari  
Sean Condon  
Jordan Halterman  
Woojoong Kim  
Ray Milkey  
Guru Parulkar  
Larry Peterson  
Adib Rastegarnia  
Thomas Vachuska

## Introduction

Cellular communication has been shaping the society for the last 40 years. To support the increasing demand, as well as new services and applications that emerge continuously, approximately every 10 years, a new generation of cellular standards is being defined. Until recently, this evolution has focused on first enabling and then scaling two distinct services: voice and broadband data. As a consequence, we have been observing an increasing surge in the volume and variety of cellular traffic, with a corresponding decline in the growth rate of new mobile subscriptions and a flattening of the ARPU<sup>1</sup>. We observe two significant consequences to this development:

- For the first time, the new cellular standard, 5G is being developed with the understanding that the network needs to be adaptive and dynamic to optimally support a variety of new business verticals, such as low-latency, mission-critical and massive machine-type communications, and not just voice and broadband data as has been in the past,
- The operators are transforming their deployment strategies, based on disaggregation, NFV, SDN and cloud principles to allow for rapid innovation and onboarding of new services while lowering their CAPEX and OPEX spending.

While the transformation in the operator deployment strategies has already spurred significant changes in the cellular core towards a disaggregated, user plane-control plane separated, services-oriented architecture allowing for dynamic creation and lifecycle management of use-case optimized network slices, to a large extent, the RAN has remained untouched until recently. To change this, a large group of mobile network operators formed the O-RAN Alliance with the goal of providing an open and intelligent RAN architecture pursuing adaptive, dynamic, highly performant and cost-effective operation in 2018<sup>2</sup>.

In parallel, ONF has been pursuing the software-defined programmability of RAN since 2015. A novel, ONOS-controlled RAN slicing solution, called ProgRAN, was developed by the ONF community and was integrated into ONF's mobile edge cloud platform, M-CORD

---

<sup>1</sup> GSMA Association, "The 5G Era in the US," White Paper, 2018.

<sup>2</sup> O-RAN Alliance, "O-RAN: Towards an Open and Smart RAN," White Paper, 2018.

by 2016<sup>3</sup>. Subsequently, ONF has become an O-RAN Alliance member in 2019. Sharing many of the same mobile network operators on its board with O-RAN, and building on its prior experience and expertise, the ONF ecosystem has recently embarked on the task of developing a 3GPP compliant exemplar RAN platform, called SD-RAN that is consistent with the O-RAN architecture. ONF's SD-RAN is built around open platform and interfaces, disaggregation, virtualization, software-defined control and cloud principles. ONF will conduct this development in close coordination with the O-RAN Alliance and O-RAN Software Community. In the remainder of this white paper we overview RAN disaggregation, its software-defined control as well as ONF's SD-RAN vision and platform development efforts.

## RAN Disaggregation

In the cellular network, RAN provides wide-area wireless connectivity to mobile devices. Towards this end, it conducts two fundamental tasks:

1. It converts IP packets to Physical Layer packets suitable for transmission over the time-varying mobile channel using packet and signal processing techniques.
2. It conducts radio resource management (RRM) control to determine how best to use and manage the precious radio resources to provide connectivity to active end devices.

To conduct these tasks, 3GPP has architected the RAN using a protocol stack as illustrated in Figure 1.

Disaggregation of the RAN takes place in three-tiers across two dimensions. The first-tier is the horizontal disaggregation that effectively splits the RAN protocol stack so that the individual components can be realized independently. This aims to deal with the challenges of high total cost of ownership, high energy consumption, better system performance by intelligent and dynamic radio resource management, as well as rapid, open innovation in different components while ensuring multi-vendor operability.

---

<sup>3</sup> Oğuz Sunay, ProgRAN: SDN-Based Programmable Radio Access Network Architecture, Argela White Paper, 2017.

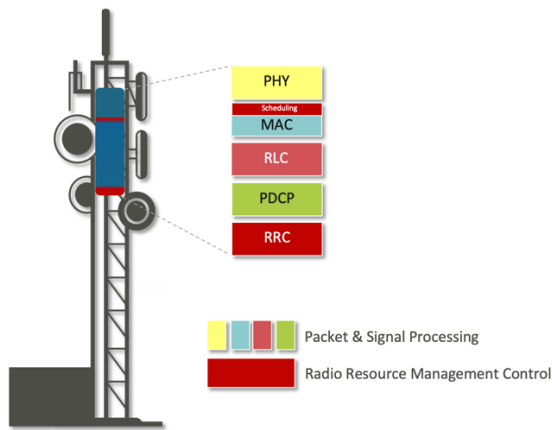


Figure 1: RAN Protocol Stack and the RRM control

3GPP has already defined a number of horizontal disaggregation options. These are summarized in Figure 2.

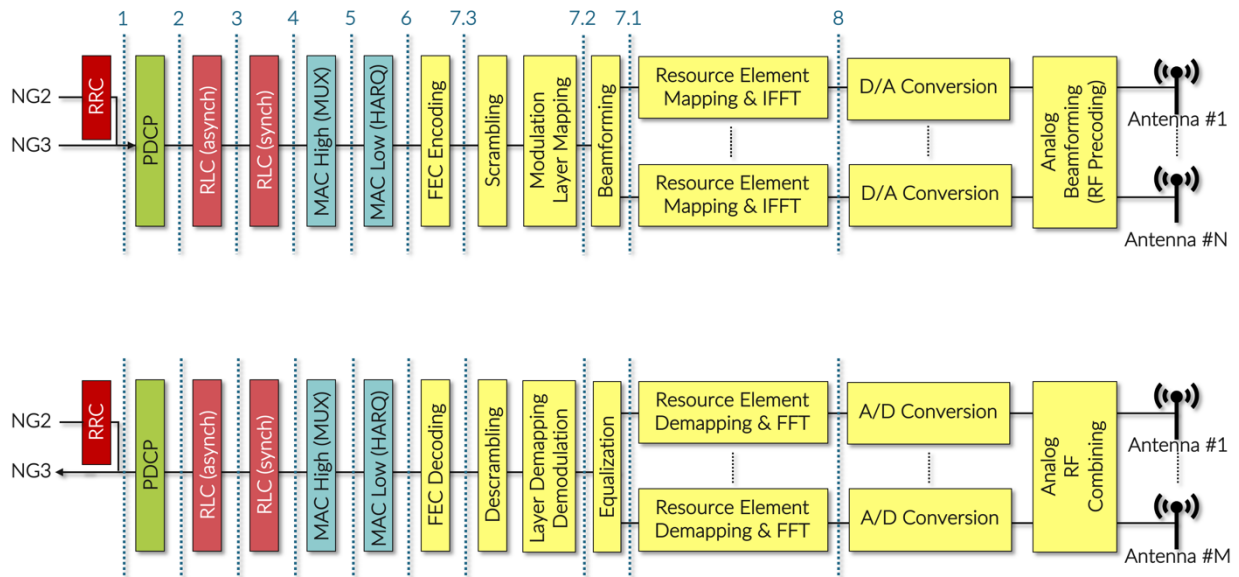
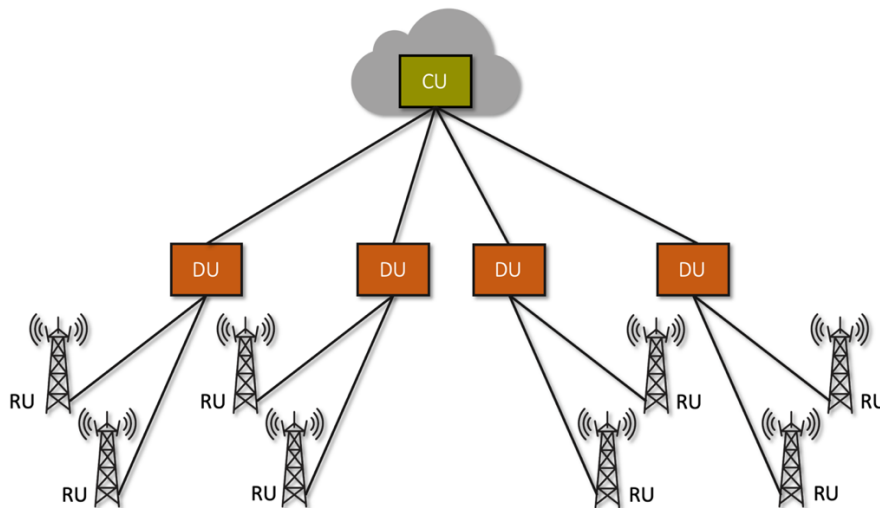


Figure 2: 3GPP Specified RAN Disaggregation Options

Following the O-RAN architecture, the horizontal disaggregation solution needs to enable a distributed deployment of RAN functions over the coverage area, as illustrated in Figure 3:

- Central Unit (CU) will centralize the “packet processing functions,” realize them as virtualized network functions running on commodity hardware, and place them in geographically centralized telco edge cloud locations,
- Distributed Unit (DU): will realize “baseband processing functions” across cell sites, realize them as virtualized network functions running on commodity hardware, allowing for possible hardware acceleration using FPGAs etc.,
- Radio Unit (RU): will enable geographical coverage using “radio functions” across antenna sites, realized on specialized hardware.

The horizontal disaggregation solution needs to be flexible, in that, based on the use-case, geography, and operator choice, in addition to the CU–DU–RU split, it should also allow for the possibility of realizing the base stations as i) two disaggregated components: CU and DU+RU, and/or CU+DU and RU, or ii) all in-one: CU+DU+RU.



**Figure 3:** RAN Disaggregation and Distributed Deployment

The O-RAN Alliance is developing open specifications for the interfaces between the disaggregated components. In this process, they have selected a subset of the 3GPP RAN split options to focus on. Specifically, O-RAN describes the RU (called O-RU) as the logical node hosting Low-PHY layer and RF processing (split option 7.2 in Figure 1), the DU (called O-DU) as the logical node hosting RLC, MAC and High-PHY layers (split option 2), and the CU (called O-CU) as the logical node hosting RRC, PDCP and SDAP layers<sup>4</sup>. This is illustrated in Figure 4.

<sup>4</sup> O-RAN Alliance, O-RAN Architecture Description, O-RAN-WG1-O-RAN Architecture Description, v1.0, February 2020.

The second tier of disaggregation is vertical, focusing on control/user plane separation (CUPS) of the CU, further disaggregating it into CU-U, the logical node hosting the user-plane part of the PDCP protocol and SDAP protocol, and CU-C, the logical node hosting the control-plane part of the PDCP protocol and the RRC protocol.

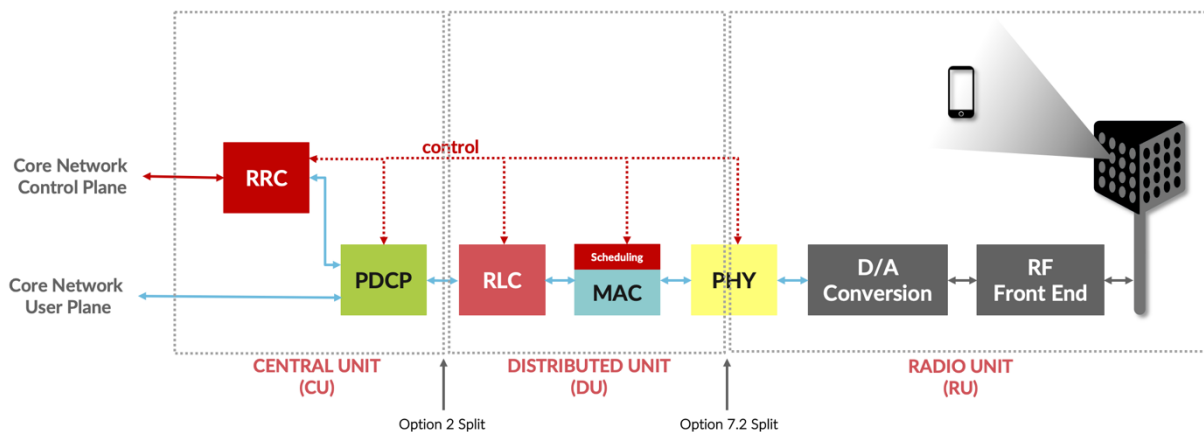


Figure 4: CU-DU-RU Disaggregation

The third-tier of disaggregation follows the SDN paradigm by carrying vertical disaggregation one step further. It does this by separating most of RAN control (RRM functions) from the disaggregated RAN components (mainly from CU-C), and logically centralizing them as applications running on an SDN Controller, which is labelled as the near real-time RAN Intelligent Controller (nRT-RIC) in the O-RAN Architecture. This SDN-based disaggregation will be further explored in the next section.

## Software-Defined RAN Control

On the path towards 5G, with network densification, and availability of different types of spectrum bands, it is increasingly a more difficult task to optimally allocate radio resources, implement handovers, manage interference, balance load between cells, etc. The current radio resource management (RRM) control that is distributed across the RAN nodes (base stations) is not optimal. Thus, it is necessary to bring software-defined controllability to RAN to increase system performance. This can be achieved by decoupling the associated intelligence from the underlying hardware and protocol stack.

In the disaggregated RAN architecture of Figure 4, the RRM functions (illustrated as red boxes) reside in the CU-C (in the RRC layer in the RAN protocol stack), and in the RU (in

the MAC layer in the RAN protocol stack). Then, the decoupling of the RRM intelligence from the underlying stack software effectively requires:

- Disaggregating the CU into CU-U and CU-C,
- Clearly defining open interfaces between CU-C and CU-U, DU, and RU
- Logically centralizing the RRM intelligence to run on the nRT-RIC,
- Explicitly defining a “Radio Network Information Base (R-NIB),” using
  - o RAN Nodes: CUs, DUs, RUs, Mobile Devices (UEs, IoTs, etc).
  - o RAN Links: all links between nodes that support data and control traffic,
  - o Node and Link Attributes: static, slow-varying and fast-varying parameters, that collectively define the nodes and the links.
- Maintaining and exposing the R-NIB using an open interface to own and 3<sup>rd</sup> party RIC applications: RRC-side RRM functions (e.g. handover control), SON applications (e.g. mobility load balancing, coordinated multi-point transmission) and ML-driven network optimization applications that are realized as SDN applications,
- Allowing for programmatic configuration of the MAC-side RRM functions using open interfaces,
- Allowing the RIC applications to exert control on the RAN based on the changes they observe on R-NIB by conveying their commands to the RAN nodes using the RIC southbound interface.

The corresponding architecture that allows for software-defined RAN control is illustrated in Figure 5.

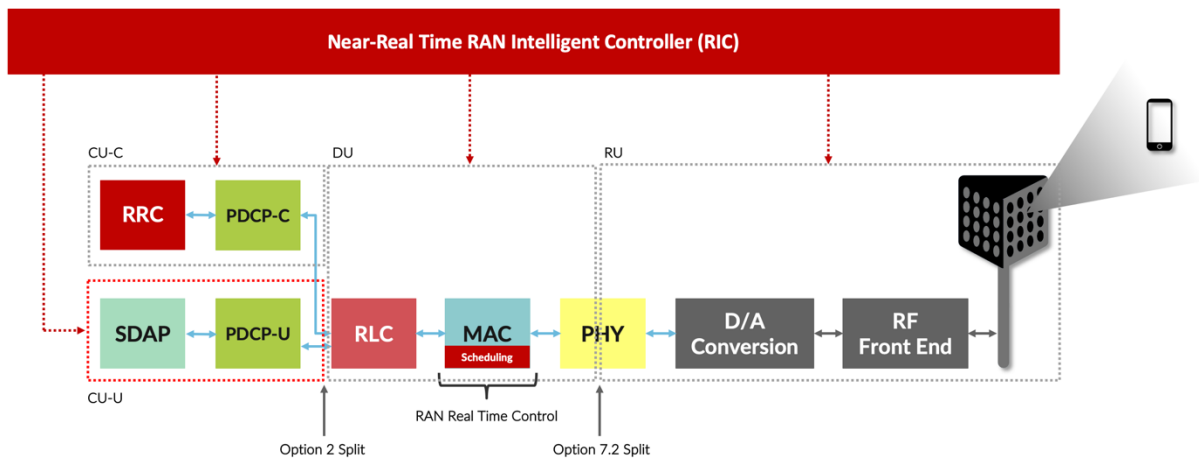


Figure 5: Software-Defined RAN Control

We note in Figure 5 that the MAC-side RRM intelligence (we call RAN Real Time Control) is not centralized in this architecture, but rather, it continues to run in a distributed manner

across the geography. However, using an open interface, we allow for these functions to be configured in real-time by the Near-Real-Time RIC.

The software-defined control of the RAN will allow for:

- Democratization of the innovation within the RAN: The control plane – user plane separation and open, clearly-defined interfaces between the disaggregated RAN components as well as between the RAN control and associated control applications, allow for innovative, third party control solutions to be rapidly deployed regardless of which vendors have provided the underlying hardware and software solution,
- Holistic control of Radio Resource Management: Software-defined control of RAN will allow for logically centralized (within limited-geography) control of radio resource management. Then, for a given active user, using innovative control applications, operators are empowered to conduct dynamic selection of any radio beam within reach across all network technologies, antenna points and sites using a global view that minimizes interference, and thus maximizes observed user quality of experience. This can be achieved by applying carrier aggregation, dual connectivity, coordinated multi-point transmission as well as selection of MIMO and beamforming schemes using a global view of the wireless network.
- Use-Case Based Management of the RAN: Software-defined RAN control will allow for the integration of performance-based decisions with policy-based constraints, with such constraints to be dynamically set, based on use cases, geographies, or operator decisions.

## ONF's SD-RAN Platform

O-RAN Alliance is in the process of developing a Reference Design specification for a disaggregated, open, virtualized, and intelligent RAN architecture. O-RAN builds on a subset of 3GPP specification options and further describes specifications for the corresponding components and interfaces for the disaggregated RAN architecture. In this Reference Design, as illustrated in Figure 6, all interfaces will be clearly specified to allow for interoperability.



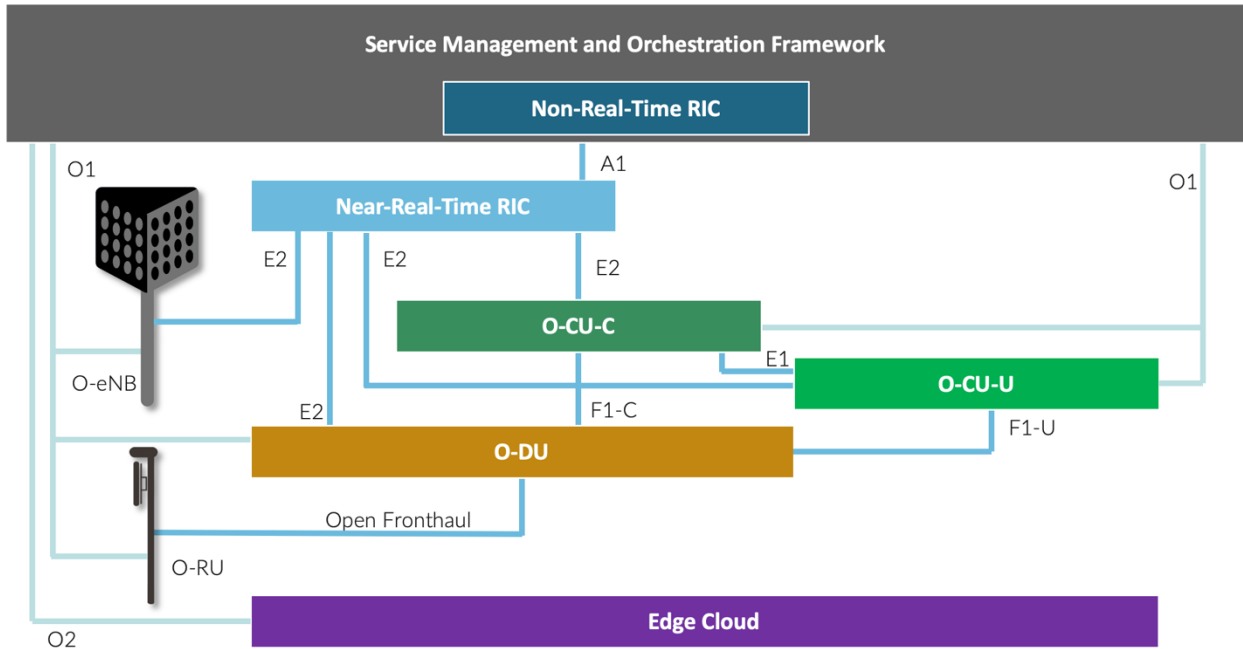


Figure 6: O-RAN Architecture

The O-RU, O-DU and O-CU specifications follow 3GPP option 7.2x and 2 splits, respectively. The Near-Real Time RIC is responsible for providing a logically centralized intelligent control of the RAN. As such, the Near-Real Time RIC should be able to control any compliant RAN platform, from the disaggregated O-RU, O-DU, O-CU-U and O-CU-C, to an all-in-one O-RAN-compliant small cell (O-eNB).

In the O-RAN architecture, the O-CU-U, O-CU-C and Near-Real-Time RIC are realized as virtual nodes that are all hosted in an Edge Cloud. O-DU can be realized using a line card, or as a virtualized node, hosted in the Edge Cloud, or outside of it, distributed across the cell sites, but connected to the edge. O-RUs are specialized components with RF modules and are distributed across the geography.

The O-RAN Reference Design specification allows for various implementation choices:

- While it is possible to develop a solution based on proprietary software components, it is also possible to consider an open source implementation.
- While it is possible to leverage CPUs to implement all virtualized components, it is also possible to develop solutions that leverage P4-based user plane components.
- While it is possible to leverage an edge cloud platform that will host all relevant network functions from open source components, it is also possible to leverage an existing vendor solution that is proprietary.

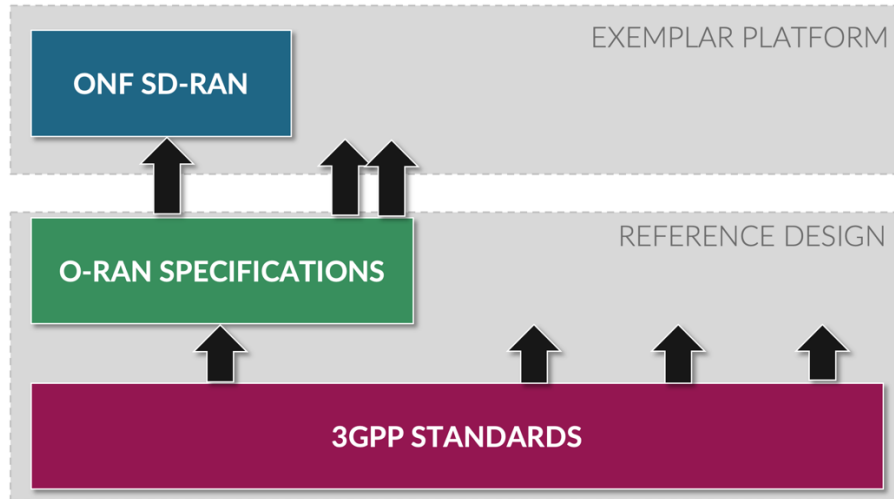


Figure 7: ONF’s SD-RAN Platform is an Exemplar Platform for O-RAN Based on Specific Design Choices

As illustrated in Figure 7, ONF has started to develop an Exemplar Platform consistent with the O-RAN architecture using a specific set of implementation choices:

- The solution will include open source implementations of O-DU, O-CU-UP and O-CU-CP,
- The solution will implement O-CU-UP using P4,
- The solution will include an open source Near-Real Time RIC Controller implementation that is based on ONF’s ONOS,
- The solution will likely expand on the E2 interface to allow for scheduler control and network slicing and contribute this expansion back to O-RAN for inclusion in the specifications,
- The solution will be inter-operable with third party RUs,
- The solution will leverage COTS and white box P4-programmable switches,
- The solution will use Aether<sup>5</sup> as the Virtualization Layer, VIM and Infrastructure Management Framework.

The corresponding ONF SD-RAN Exemplar Platform architecture is illustrated in Figure 8.

<sup>5</sup> Oğuz Sunay, et.al., Aether: Enterprise-5G/LTE-Edge-Cloud-as-a-Service, ONF White Paper, 2020.

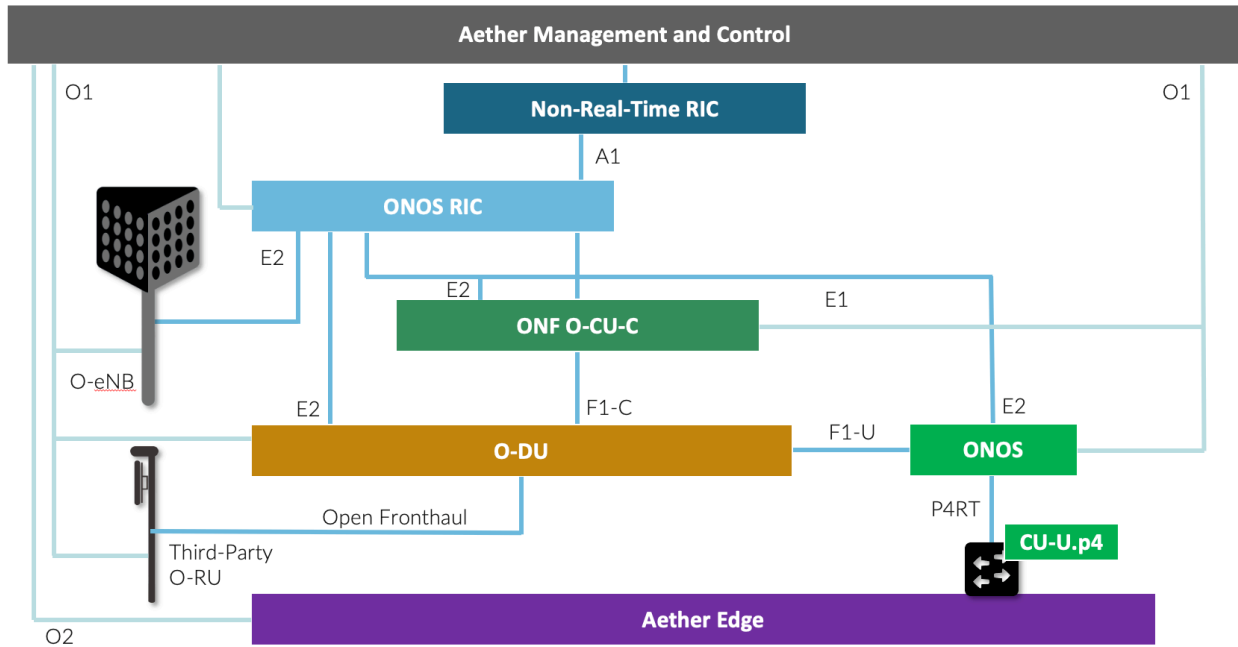


Figure 8: ONF's SD-RAN Exemplar Platform

ONF has started developing the individual components of the exemplar platform. The work will progress in stages and will be conducted in close coordination with the O-RAN Software Community. The timeline for the development of the different components is as follows:

- Aether: A pilot Aether network is operational in multiple sites today. Further development is in progress and will continue throughout 2020.
- ONOS RIC: Development work on optimizing the microservices-based ONOS ( $\mu$ ONOS) for RAN control has already started. Two sample RRM applications have also been developed for handover control and mobility load balancing. Development will continue throughout 2020.
- CU: Development work on ONOS-controlled and P4-based CU-U as well as containerized CU-C will start in 2020. CU-C development will leverage O-RAN Software Community's ongoing work.
- DU: Development work on DU will start on 2021. DU development will leverage O-RAN Software Community's ongoing work.

## $\mu$ ONOS RIC

A first-generation implementation of ONOS has been in use for 5+ years and is the leader in the leading open-source SDN control plane in terms of high availability, performance and scalability. Currently, ONOS controls Comcast's Trellis Open Source Networking Fabric

rollout in a number of geographies, serving a large number of live customers<sup>6</sup>. Development of the next generation of ONOS, based on a microservices architecture is underway and ONOS RIC is the main driver of this effort. This architecture will include the following innovations:

- Support a new generation of control and configuration interfaces and standards: P4/P4Runtime, gNMI/OpenConfig, gNOI.
- Cloud native, adopting best practices in micro-services, the use of polyglot interface mechanism (gRPC), and container management (Kubernetes).
- Highly available, dynamically scalable and highly performant, in terms of both throughput (control/config ops/sec) and latency for implementing control-loops.

μONOS RIC will run as a logically centralized SD-RAN controller, and adopt a microservices architecture that includes the assembly of the following components:

- **Certificate Manager:** Provides CA and certificate management required to securely interact with the network environment.
- **Topology Manager:** Tracks inventory of network infrastructure devices and their interconnects to provide a shared view of the network environment for the rest of the platform and applications.
- **Configuration Manager:** Facilitates issuing, tracking, rollback, and validation of atomic configuration operations (via gNMI and gNOI) on multiple network infrastructure devices to maintain consistent network operation.
- **RAN Control Manager:** Allows shaping the mobile network infrastructure devices (O-RU, O-DU, O-CU, O-eNB) and nodes, and subsequent control of RRM and network optimization using O-RAN interfaces within specified latency limits.
- **Distributed Store:** Provides services for other core systems and applications to store and access data using several strategies (e.g., Raft consensus, eventual consistency). The system will include client-side libraries to manage caches and to provide data structures for efficiency and ease of use.

The high level μONOS RIC architecture is illustrated in Figure 9.

---

<sup>6</sup> ONF, Comcast has Achieved Production Roll-out of Trellis Open Source Networking Fabric, ONF News Release, 2019.

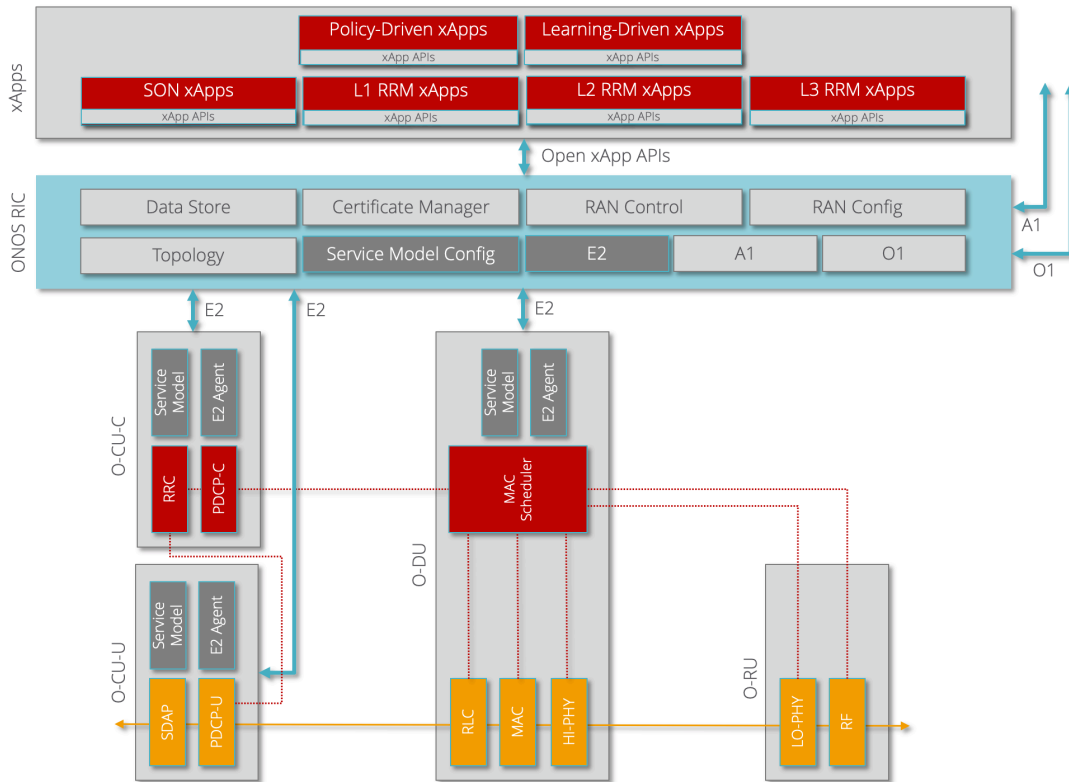


Figure 9: High-Level μONOS RIC Architecture

μONOS RIC enables a multi-cluster operation for high-availability and scalability. The southbound interface of ONOS RIC is the O-RAN specified E2 interface. ONOS provides a distributed data store to maintain the Radio Network Information Base (R-NIB), messaging infrastructure as well as topology, control, and configuration managers in a microservices environment. It intends to support conflict resolution to resolve conflicts emanating from multiple RAN applications.

μONOS RIC provides open APIs to host 3<sup>rd</sup> party RAN applications, called xApps in the O-RAN architecture. These APIs enable the xApps to interface with the E2 Nodes (O-DU, O-CU-C, O-CU-U, all-in-one gNB) in the south, the operators' OSS/BSS platforms in the north, the non real-time RIC in the north, as well as the distributed data store within the μONOS RIC. These applications will vary from basic RRM functions, to Self-Organizing Network (SON) applications, to ML-driven network optimization applications and policy-based control and configuration applications as illustrated in Figure 9.

## Conclusion

With the explosive ongoing demand in mobile connectivity via an ever increasing diversity of end devices for a wide variety of use cases each with vastly different demands, the mobile network, including the RAN needs to be open, disaggregated, virtualized, software-defined, controllable and should operate using cloud principles. 3GPP and O-RAN collectively provide the inter-operability specifications for this vision. ONF's SD-RAN is an open exemplar platform that is cloud-native and is built on ONF's well-established, operator-approved and deployed platforms, such as ONOS and Aether. A vibrant ecosystem composed of operators, xApp developers, RAN vendors and stack vendors has already formed for SD-RAN, including AT&T, China Mobile, Deutsche Telekom, China Unicom, Facebook, Intel, Sercomm, Baicells, and Radisys. We invite everyone from the ONF ecosystem to join us in this group to actively participate, collaborate and adopt this exciting new endeavor.

## About ONF

The Open Networking Foundation (ONF) is an operator led consortium spearheading disruptive network transformation. Now the recognized leader for open source solutions for operators, the ONF first launched in 2011 as the standard bearer for Software Defined Networking (SDN). Led by its operator partners AT&T, China Unicom, Comcast, Deutsche Telekom, Google, NTT Group and Turk Telekom, the ONF is driving vast transformation across the operator space. For further information visit <http://www.opennetworking.org>