

WHITE PAPER

An Introduction to O-RAN

CONTENTS

Introduction.....	2
How to Split the RAN	3
Interoperability Testing (IOT)	6
Conclusion	8

Introduction

As global cellular data usage continues to increase, the way we build telecommunications systems must change to keep up. While the 5G standard meets demands for more cellular throughput and aims to address new use cases, many applications identified in the 5G charter may not happen without network evolution. Specifically, the ultrareliable low-latency (URLLC) use case demands that networks meet a latency spec—impossible without network changes. Future networks need to be more flexible and take advantage of new technologies such as artificial intelligence. Network operators want to move toward software-defined networking to customize and manage their deployed networks. Mobile network operators also need equipment interoperability, or the ability to choose different network equipment components regardless of vendor. On the whole, there is massive room for improvement to both radio access network (RAN) and telecommunication system hardware.

In Release 15, the 3GPP identifies three distinct gNodeB functions: Centralized Unit (CU), Distributed Unit (DU), and Radio Unit (RU). There are several ways to configure these three components, and deciding which configuration is best depends on each individual network. Single-vendor gNodeB options may choose proprietary interfaces between these components. The O-RAN alliance, or O-RAN, focuses on achieving an unprecedented level of openness in the way 5G RANs are built. The group's charter conveys how having open interfaces between the CU, DU, and RU translates to building networks from white boxes instead of being vendor-locked for the whole system, and that, by doing so, the RAN is more flexible, and network operators have more options. Also, this approach encourages more innovation from smaller companies that have not traditionally provided network hardware. More innovation and more options could mean potentially lower costs to deploy new networks. O-RAN also hopes to integrate deep learning techniques into every RAN architecture, making communications systems more intelligent. O-RAN's reference architecture, shown in Figure 1, illustrates how to build an O-RAN-compliant RAN.

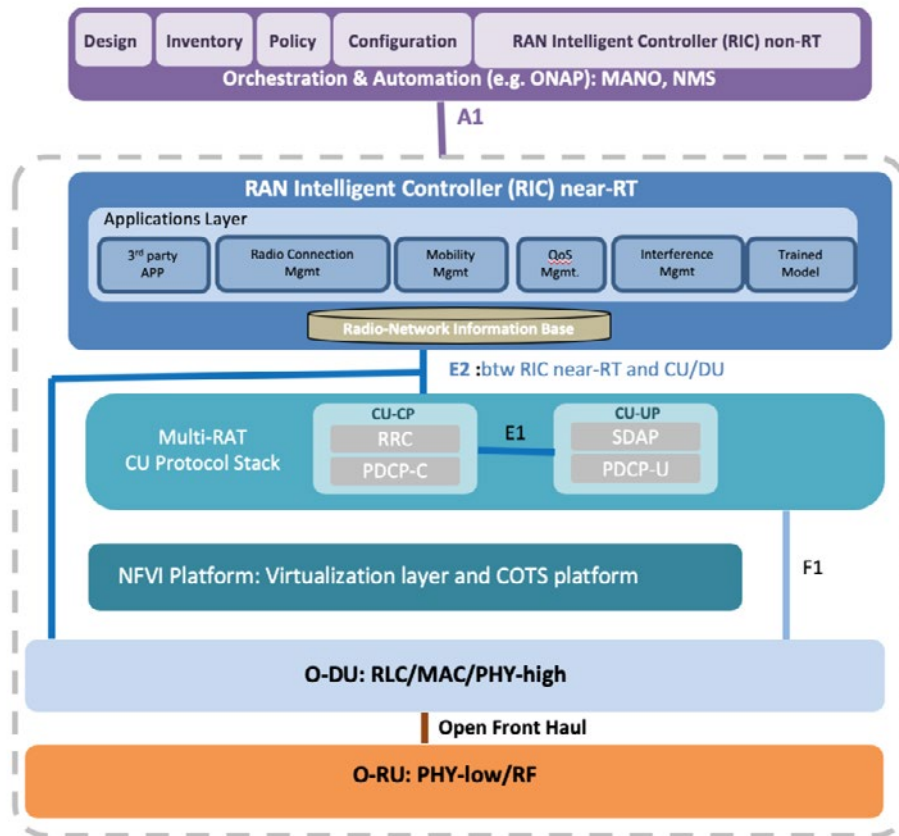


Figure 1. O-RAN Alliance Reference Architecture (Source: O-RAN White Paper)

How to Split the RAN

O-RAN’s proposed concepts and architectures use a split-RAN concept. There are eight known ways to functionally split the RAN, and each one proposes splitting the processing so that different parts of the protocol stack process on different hardware. Figure 2 summarizes the eight options.

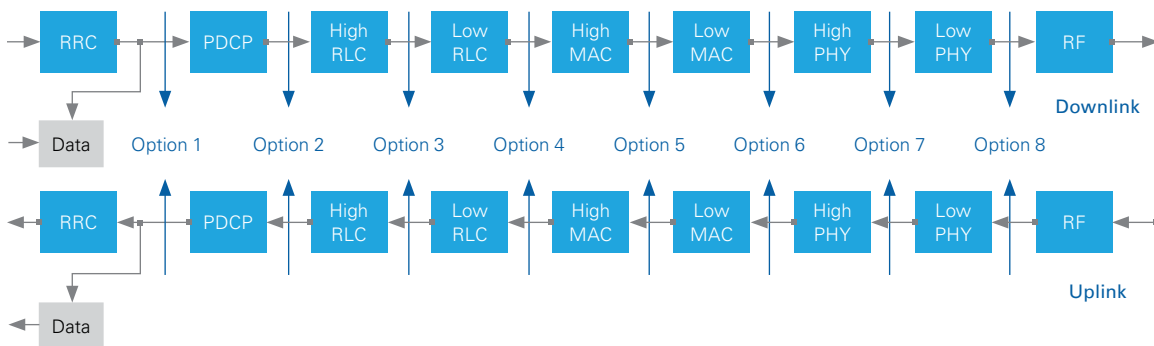


Figure 2. RAN Split Options (Source: NGMN 2018)

O-RAN proposes using option 7-2 which, as shown in Figure 2, splits the physical layer (PHY) into a high-PHY and a low-PHY. For option 7-2, the uplink (UL), CP removal, fast Fourier transform (FFT), digital beamforming (if applicable), and prefiltering (for PRACH (Physical Random Access Channel) only) functions all occur in the RU. The rest of the PHY is processed in the DU. For the downlink (DL), the inverse FFT (iFFT), CP addition, precoding functions, and digital beamforming (if applicable) occur in the RU, and the rest of the PHY processing happens in the DU.

2G, 3G, and 4G use Common Public Radio Interface (CPRI), which is passed on an option 8 split. Moving to the 7-2 split reduces traffic between the DU and RU. O-RAN has specified a version of the 7-2 split. Figure 3 below illustrates the 7-2 split, as well as how other parts of the protocol stack are split between the CU and the DU. The 7.2x split is the best balance between bringing this technology to market quickly and deployment cost. It reduces confusion about split specifics while making traffic reduction gains and improvements. Some 5G systems use evolved CPRI (eCPRI) as the DU-RU interface. eCPRI offers a vendor-specific split between the high and low PHY. Thus, you can optimize for traffic or flexibility by supporting multiple splits, accommodating different deployment environments due to unique antenna physical environments. With this comes the freedom to cost-optimize for connections specific to different operators.

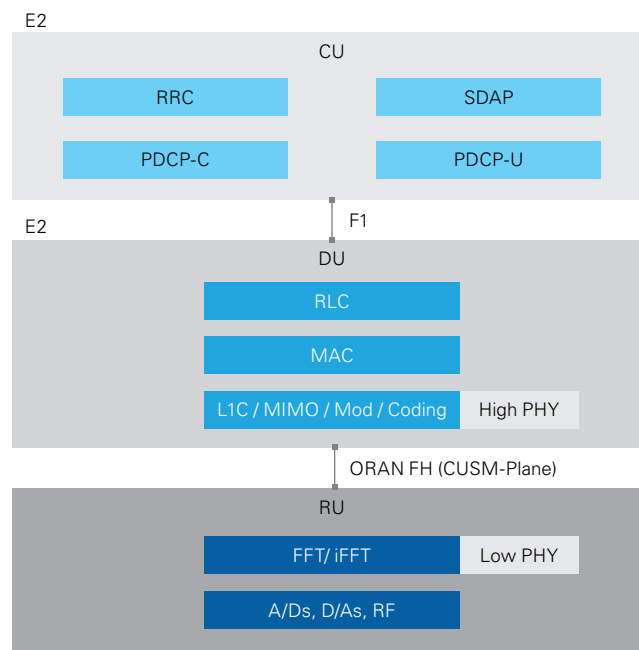


Figure 3. Protocol Layer Split Between CU, DU, and RU for Option 7-2

For new 5G RAN architectures (called NR-RAN), the 3GPP has defined and standardized on a new interface, the F1 interface, for communication between the CU and the DU. When the CU and the DU are physically split, it is called a higher-layer split (HLS). While not defined by the 3GPP, the lower-layer interface between the DU and the RU is called a lower-layer split (LLS). You can configure the CU and DU in relation to each other and in relation to RUs in several ways. Figure 4 diagrams various NR-RAN configurations. Note that the F1 interface is delay-tolerant, while the DU-to-RU interface needs to be low-latency. Given the challenges of creating low-latency interfaces, the rest of this paper details a central RAN with a lower-layer split use case.

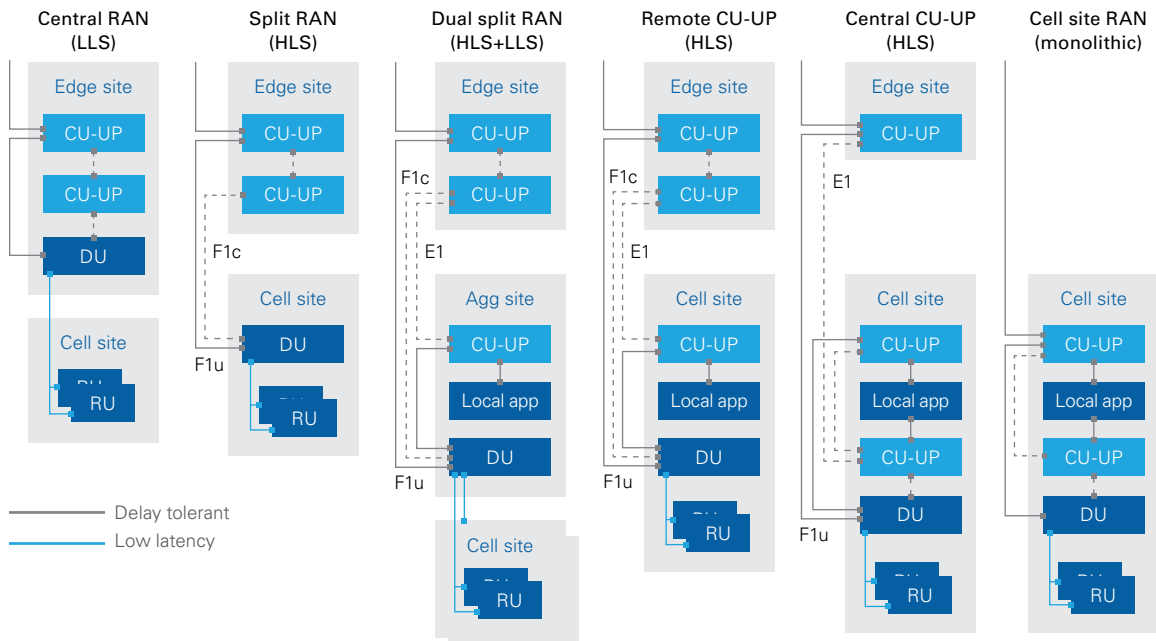
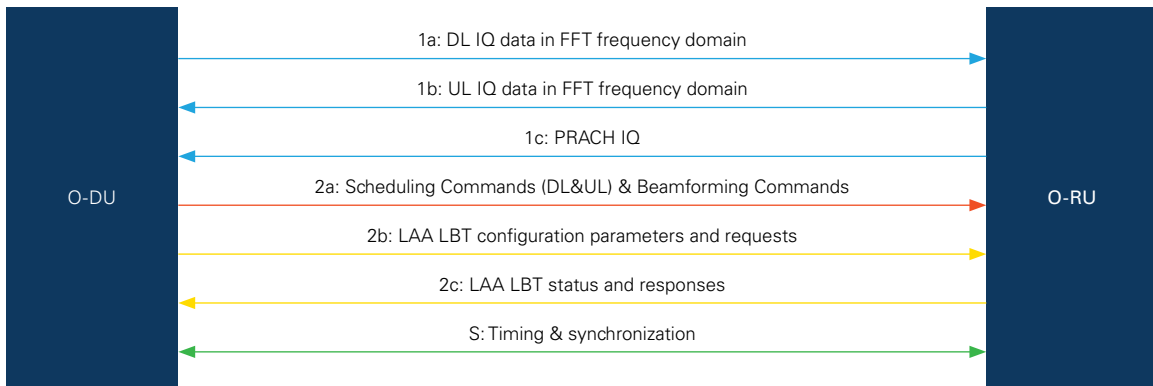


Figure 4. 5G RAN Functional Unit Location Flexibility (Source: NGMN, 2018)

The interface between the DU and RU also is known as the fronthaul (FH) interface. The FH interface, one of the most demanding system interfaces, is very latency-sensitive. Where the DU and RU come from the same manufacturer, most systems use CPRI or eCPRI (5G only) as the FH interface. While CPRI originally was intended as an open interface, in practice, each vendor implemented it in a slightly differently way to work with their own hardware, rendering multivendor interoperability difficult or impossible. While not encouraging an open, white-box hardware architecture, you can more easily achieve tight synchronization between the DU and RU. When the DU and RU are from one vendor, they are matched in terms of when to send and when to receive (the only variation being the distance between DU and RU).

One of O-RAN’s two goals is to create more open ecosystems, which requires defining a new FH interface. One of the seven O-RAN working groups, working group 4 (WG4), is dedicated to defining this interface. Called the Open Fronthaul Interfaces Workgroup, its objective “is to deliver truly open fronthaul interfaces, in which multivendor DU-RRU interoperability can be realized.” Figure 5 shows how the proposed DU-to-RU interface exchanges information in different planes. While the seven different flows—plus additional management (M)-plane flows—may seem daunting, at a higher level, there are only three data types (IQ data; timing and sync data; and command and control information) across four total planes (control, user, sync, and management).



Note: M-Plane flows not represented here

Figure 5. Lower Layer Fronthaul Data Flows (Source: O-RAN)

Compared to CPRI, there are significant differences in how the IQ data is transferred, packed, and unpacked in the FH interface, since CPRI is based on an option 8 split. The option 8 split splits the network at the RF, so the IQ samples have not undergone any PHY processing (FFT/iFFT). As networks evolved in the later stages of 4G and early 5G, eCPRI sought to reduce traffic caused by the increase in antennas and sampling rate (multiple samples per antenna) used in massive multiple-input and multiple-output (MIMO). System traffic overwhelmed the physical connection, and connections that can accommodate the traffic are expensive to implement. To reduce traffic across this interface, eCPRI moves certain parts of the PHY into the RU and adds compression algorithms. However, which pieces of the PHY are moved into the RU does not follow any specific split and differs from vendor to vendor. This could be a competitive advantage to some vendors, and a way to potentially reduce the link cost for operators. Since part of the lower-level PHY functions are in the RU, the DU needs to inform the RU how to perform these functions. This creates a very different command-and-control interface between eCPRI and O-RAN's FH interface as well. But having vendor-specific splits perpetuates a service-provider vendor lock. O-RAN's open FH interface aims to standardize on which pieces of the PHY are moved into the RU by using the 7-2x split so that you can integrate hardware from different vendors.

Interoperability Testing (IOT)

As WG4 works to finalize an FH interface, they must consider how to test it. For systems containing a DU and RU from different hardware vendors, system integrators and vendors need the ability to validate that the DU and RU interface properly. This type of testing is commonly referred to as interoperability testing. O-RAN is investigating how to test an O-RAN-compliant system. Figure 6 is an O-RAN diagram showing what a test setup might look like for testing an O-RAN-DU (O-DU) and O-RAN-RU (O-RU) with an O-RAN-CU (O-CU) and a UE, which could be emulated or commercial. There is a test point to look at the interface between the CU and DU, and one to look at the RU RF input/output, but the DU and RU are combined as the device under test (DUT). This leaves the FH interface between the DU and RU untested when using active stimulus and only considered for passive monitoring. O-RAN currently is looking at how to test the FH interface.

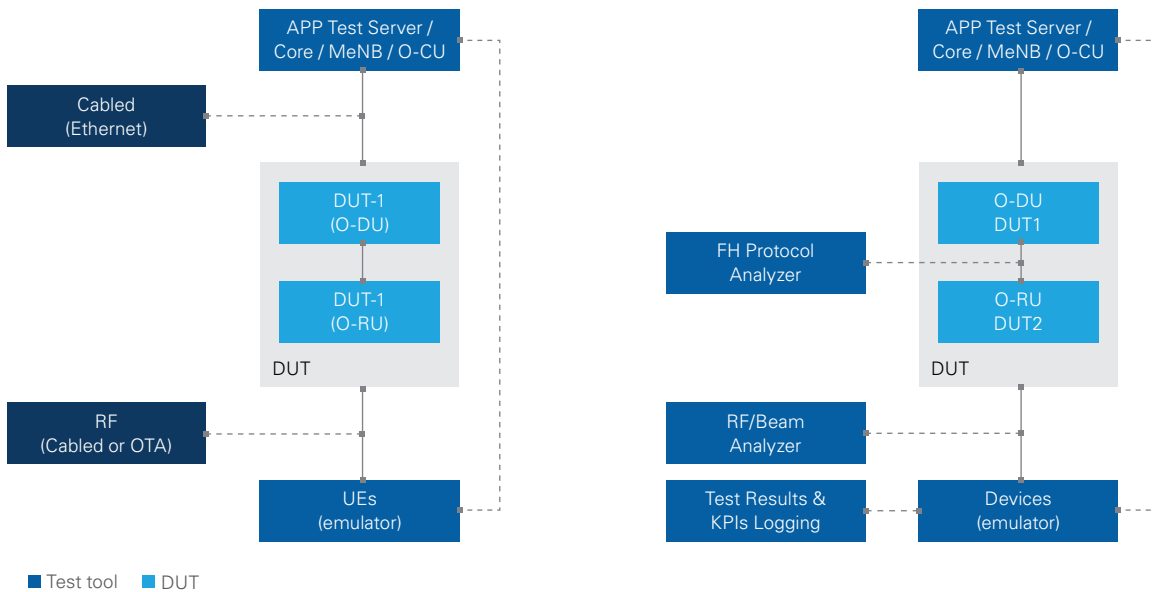


Figure 6. O-RAN Test Setup, Active (Left) and Passive (Right) (Source: O-RAN)

There are two kinds of active testing to consider for the FH interface: Protocol testing and parametric testing. O-RAN has shown that protocol testing is necessary for test-case validation and troubleshooting. Having test tools to validate designs during development is key to interfacing successfully with other O-RAN-compliant devices. After the design is finalized and DUs and RUs enter the validation and production stages, parametric tests make sure that each unit performs as expected.

For an RU tester, the E1 interface between the CU and DU does not need to be tested. An RU tester must emulate both CU and DU functionality and be able to monitor the FH interface. Depending on required testing level, adding a test UE or UE emulator to the system can create a full end-to-end (E2E) tester. NI provides 5G NR IP based on its software defined radio (SDR) and FPGA hardware. An example E2E tester is shown below in Figure 7.

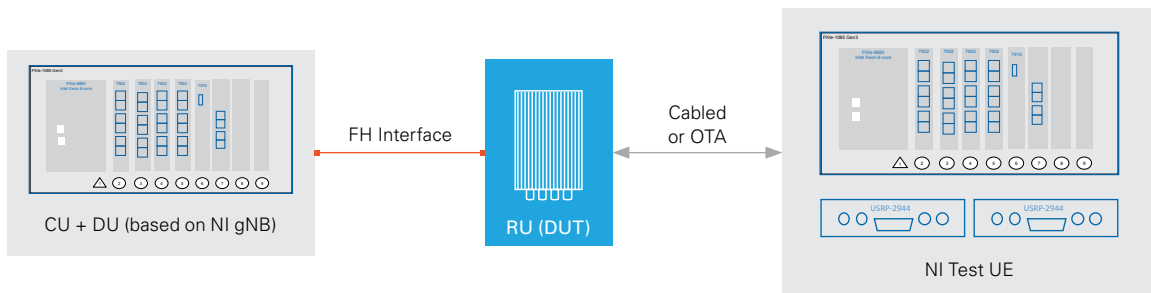


Figure 7. E2E DU-RU Tester Example

Conclusion

O-RAN maintains three goals:

- Creating a smarter RAN that takes advantage of being able to virtualize pieces of the network for maximum efficiency gains
- Supporting white-box hardware for multivendor network solutions
- Creating standardized interfaces between network components

By delivering on these key initiatives, O-RAN aims to evolve networks to be more future-proof and incorporate new 5G-promised features and use cases, such as URLLC. The FH interface, in particular, is challenging because of the low latency required for DU-to-RU communications. O-RAN's WG4 is making progress to define this interface, and companies are beginning to build O-RAN-compliant RUs to connect to other O-RAN-compliant hardware. Being able to validate and test the DU-RU interface is important in both the design and validation stages, as well as in production test, as this new technology comes to market. By offering IOT testing hardware and software, NI can help accelerate bringing new O-RAN-compliant RUs to market. Whether O-RAN is widely adopted and used in 5G networks remains unclear; however, as of now, the alliance actively is working towards defining new interfaces and seeking ways to use multivendor hardware solutions to build new networks for the sake of improving and advancing the way that networks are built.