Part 1: Accelerating 5G virtual RAN deployment

With open, cost-optimized O-RAN Alliance whitebox implementations

May, 2020

This whitepaper provides an overview of the O-RAN Alliance and the value provided by O-RAN to the carrier industry along with a brief introduction to Comcores, our O-RAN solutions and our role in the O-RAN eco-system.

Executive summary

5G is currently being deployed across the globe with carriers expected to invest over $480 billion in infrastructure deployment not including spectrum investments. 5G discussions on business cases still continue, but the focus for carriers now is cost optimization.

The Radio Access Network (RAN) is the most critical part of new 5G networks. 5G ambitions to deliver more data, faster and with lower latency require a complete rethink on how the RAN is designed. It also requires a rethink on how the RAN is implemented. 5G RAN infrastructure Total Cost of Ownership (TCO) can grow by 65% compared to current 4G RAN costs. Energy costs can grow by as much as 140%. Managing these costs is critical to the carrier business case.

According to studies by the GSMA, virtualization of the RAN can reduce TCO by 25% while opening critical interfaces, such as the fronthaul interface can enable network sharing leading to a further 40% saving. Ensuring that solutions are compact and energy efficient will also be critical to the 5G business case.

Developing a 5G RAN that can meet future service needs cost-effectively is more than one carrier or vendor can address on their own. Collaboration on the development of open, standards-based, virtualized solutions has become a necessity as carriers seek opportunities to reduce costs, while also enabling a more flexible, dynamic and responsive RAN.

A good example of this kind of collaboration is the O-RAN Alliance, which is focused on accelerating the implementation of standards and reference designs for virtualized 5G RAN solutions. The O-RAN Alliance reference designs are provided as "whitebox" or un-branded implementations that can be deployed by carriers today thereby reducing costs and vendor lock-in.

The O-RAN Alliance whitebox group (WG7) has released a whitebox implementation of both a 5G RAN O-DU and O-RU. Comcores has provided the core functionality of the O-RU whitebox implementing layer 1 offload and the eCPRI interface to the Intel® FlexRAN-based O-DU, which is the first implementation of an open front-haul interface.

This is not just a reference design, but a significant step in enabling cost-effective 5G RAN deployments.

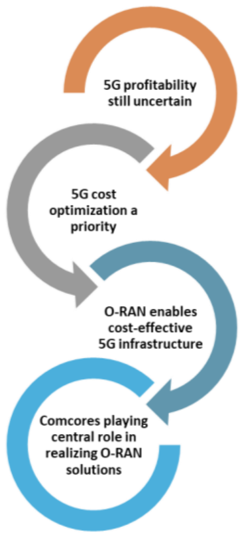

Figure 1: O-RAN solutions critical to 5G business case

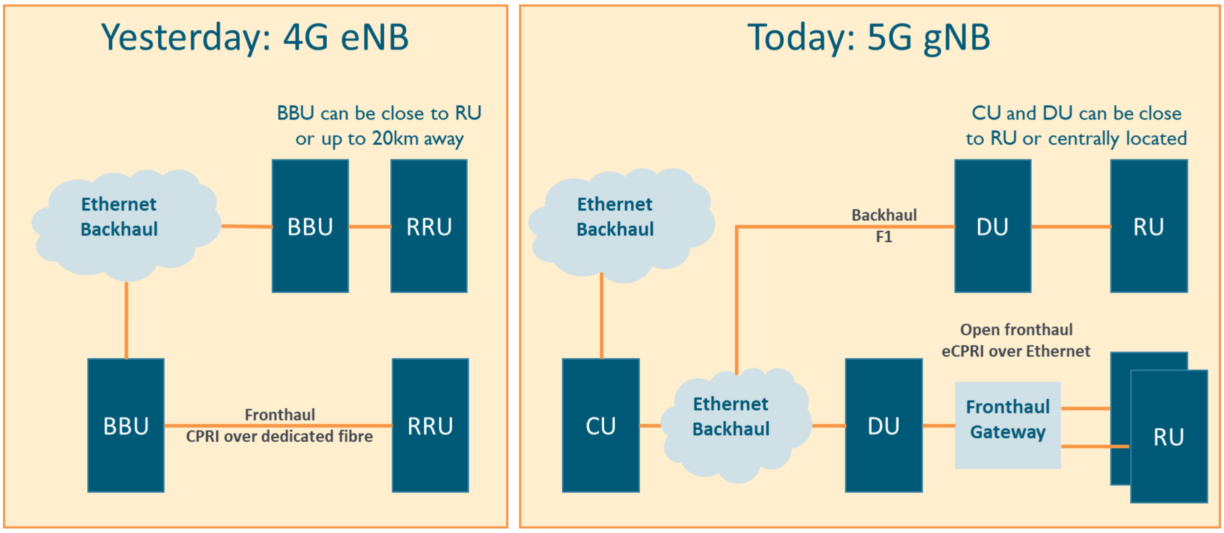

Figure 2: 5G RAN functions

5G is a revolution not an evolution

5G is designed to deliver high-quality voice and faster Internet connectivity, like previous generations. But it is also designed to support emerging service needs, such as connecting billions of Internet of Things (IoT) devices, supporting 4K/8K video streaming, autonomous vehicles, eHealth and industrial automation.

The challenge in 5G, which requires a revolution in network design and architecture, is to build an infrastructure that can meet diverse service needs, at the same time, cost-effectively.

Delivering higher data rates at massive scale with ultra-low latency requires more capacity and intelligence at the edge of the network. But how much capacity is required and for how long depends on the specific service. What is required is the ability to deploy exactly what is needed by a specific service, when it is needed without impacting the experience delivered for all other services. This requires a new Radio Access Network (RAN) architecture, new technologies, like virtualization and new techniques like network slicing.

5G RAN architecture

With 5G, a new architecture was introduced that allows the decision on whether to centralize or decentralize functionality to be made on a per service basis. Figure 2 shows the 4G enhanced NodeB (eNB) and the new 5G RAN architecture. In 5G, the 4G eNB equivalent is referred to as Next Generation Node B (gNB). This includes three distinct functions that can be separated geographically; a Central Unit (CU) that can support multiple Distributed Units (DUs) each connected to one or more Radio Units (RUs).

4G LTE introduced the option to split the eNB into a Baseband Unit (BBU) and Remote Radio Unit (RRU). The RRU that previously was installed together with the BBU at the base of the antenna could now be deployed on the mast close to the antenna and the BBU could be deployed centrally. This allows remote radio heads to be deployed in small cells where the macro cell is struggling to keep up with bandwidth demand or coverage dark spots.

As more intelligence is centralized, the bandwidth requirements for connecting functions increases as more “raw” data needs to be backhauled for processing. The flexibility to deploy the CU and DU centrally or remotely enables the best compromise between latency and backhaul bandwidth requirements to be achieved.

Fronthaul protocols and layer splits

In 4G networks, the connection between the BBU and RRU is referred to as the fronthaul network and is typically based on the Common Public Radio Interface (CPRI), as can be seen in Figure 2. This is a point-to-point connection running over a single fibre that can support distances up to 20km. But this means dedicated fibres for each RRU, which becomes a costly bottleneck. To address this, an ethernet based version of CPRI called eCPRI was introduced in 2017. With eCPRI, radio signals are transported over Ethernet, which allows multiple signals to be multiplexed and transported on a single fibre. An Ethernet switch or Fronthaul Gateway can be used to connect multiple RUs to a single DU thus reducing the number of fibres to be deployed.

As shown in Figure 3, CPRI provides an interface between the radio frequency and the physical layer of the communication stack. eCPRI defines five possible splits as well as options to split the PHY layer. 3GPP defines 8 split locations, which align with eCPRI splits.

The CPRI interface has to date been proprietary to the vendor implementing the radio equipment. With eCPRI, it is now possible to define an open interface that can allow other vendors to develop DU or RU equipment that can interoperate with other vendor equipment.

The choice of where to make the split is a trade-off between latency and backhaul requirements. Thus far, 3GPP has specified a “Higher Layer” Split 2 between the PDCP and RLC layers, which is effectively the interface between the CU and the DU while eCPRI proposes a “Lower Layer” split 7 or “intra-PHY” split between the higher and lower PHY layers, which is the interface between the DU and RU. 3GPP is considering this option for future revisions of their specifications.

Depending on where you place the RRU, DU and CU functions, it is possible to find the right trade-off between latency and backhaul requirements on a per service basis. But, to get the most benefit from the flexibility that the gNB offers, a mechanism is required to dynamically deploy functionality when and where it is required. This is where virtualization of RAN network functions and network slicing are essential.

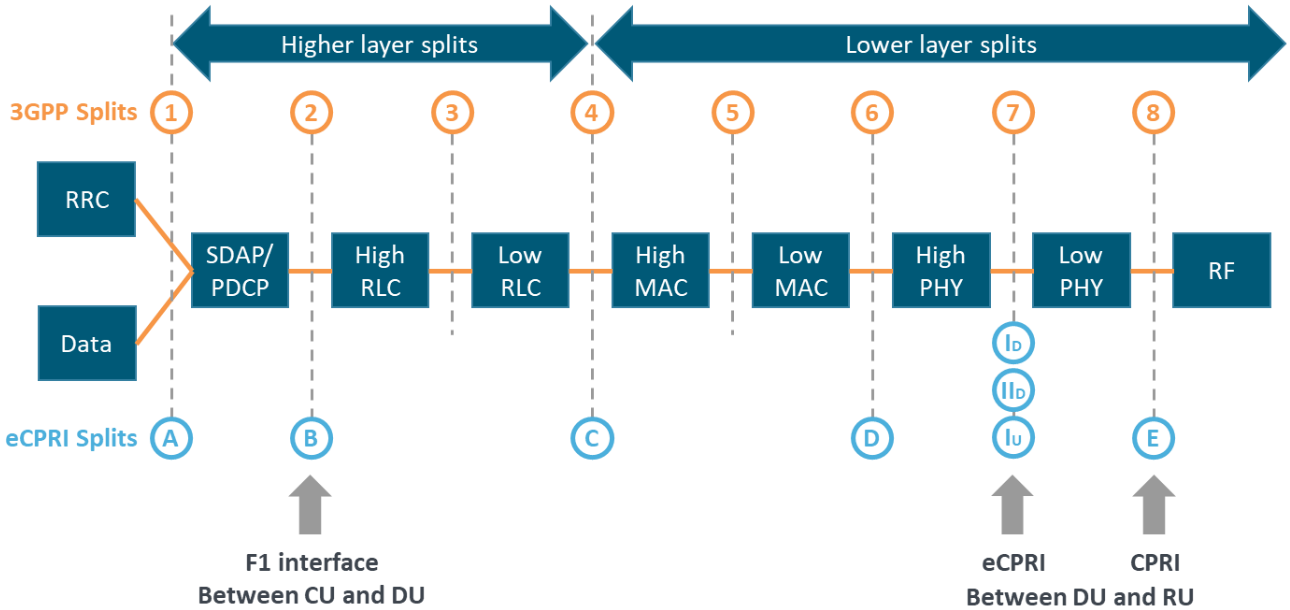

Figure 3: 5G RAN functional splits

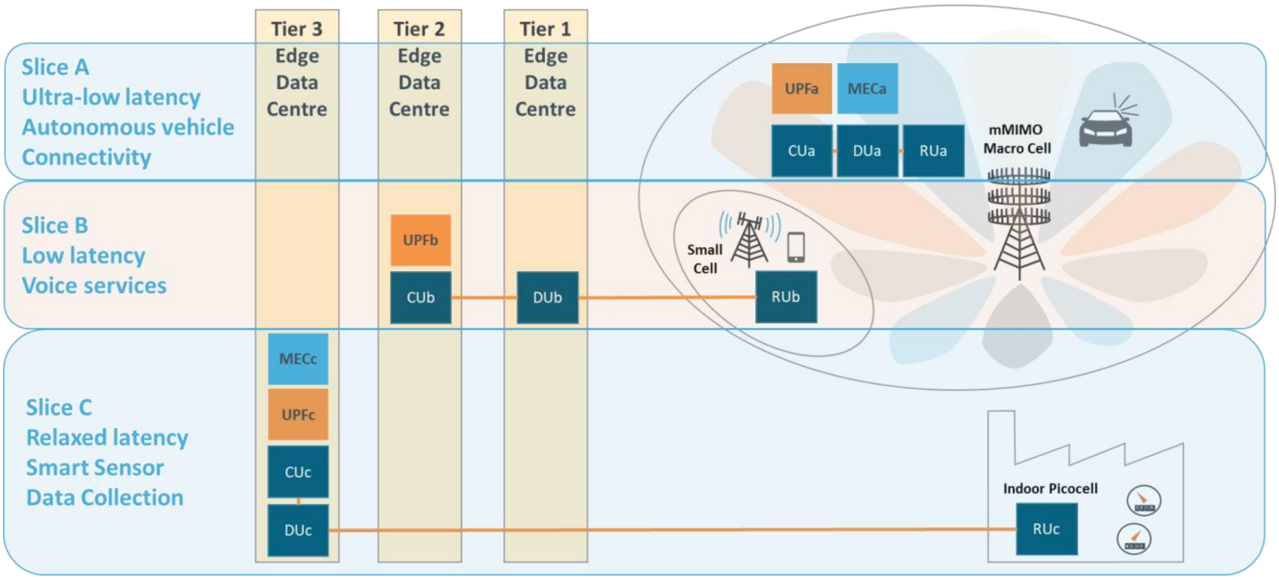

Figure 4: 5G Network slices

Virtualizing the RAN

To date, virtualization in mobile networks has been focused on core functions like Evolved Packet Core (EPC), where significant progress has been made. Virtualizing the RAN and the edge of the network is more demanding, but crucial to the success of 5G deployments. Since services are delivered with performance requirements end-to-end, the ability to manage how resources are deployed in the RAN in real-time becomes extremely important.

The virtualization of 5G networks has introduced the key concept of network slices. A network slice is a virtual network dedicated to a specific customer or service where the resources supporting the service are configured and deployed to meet the performance profile required. By logically separating the functions in the gNB and then implementing them as virtual software, the 5G RAN can support specific network slice requirements and the diverse demands of each service at the same time.

Optimizing the cost of 5G RAN

Virtualizing the RAN for 5G requires a new architecture and new technologies, which will increase costs. However, the new architectures and technologies also provide opportunities to manage and optimize these costs.

In a recent study by GSMA, it was reported that carriers, as a result of uncertainty about service revenue, have recently shifted their focus from building the business case for 5G services to optimizing the cost of 5G. In the study, it was found that the Total Cost of Ownership (TCO) associated with 5G RAN infrastructure, including capital expenditures for buying equipment and operational costs, can grow as much as 65% compared to current mobile networks. Optimizing these costs is therefore a priority for carriers.

In the study, it was concluded that by virtualizing the RAN, using more automation and intelligence and by adopting open implementations, such as those provided by the O-RAN Alliance, it would be possible to reduce these costs by 25%. This is made possible by the flexibility to deploy functions and capabilities when and where they are needed.

According to the GSMA, using open implementations also enables sharing of network ownership and costs, which can further reduce RAN TCO by up to 40%.

But the largest cost accelerator identified in the study is energy. The study found that the introduction of 5G can increase energy costs by as much as 140% of current needs. This is driven by massive MIMO, the need for more sites and mobile data traffic growth requiring more processing power.

This means that any possibility to reduce power consumption in mobile networks is valuable, which, in turn requires the right functionality to be deployed in the right locations at the right time to optimize energy usage while also meeting service requirements.

To read the study in more detail, see:
https://www.gsma.com/futurenetworks/wiki/5g-era-mobile-network-cost-evolution/

The O-RAN Alliance

Formed through the merger of the xRAN Forum and the C-RAN Alliance in 2018, the O-RAN Alliance is seeking an open, standards-based consensus for virtualizing the RAN. They are focused on building a more cost-effective and agile RAN through open interfaces, while using intelligence and deep learning techniques to automate operations.

The O-RAN Alliance proposes a new architecture based on continuous data collection and deep learning techniques. The RAN Intelligent Controller (RIC) is a key component in the new architecture and is responsible for managing and optimising the deployment and use of RAN resources. See Figure 5 for an overview of the O-RAN architecture.

The O-RAN Alliance is also focused on opening the front-haul interface between the O-DU and O-RU, which is currently not standardized. This is critical to enabling multi-vendor O-DU and O-RU deployments and network sharing.

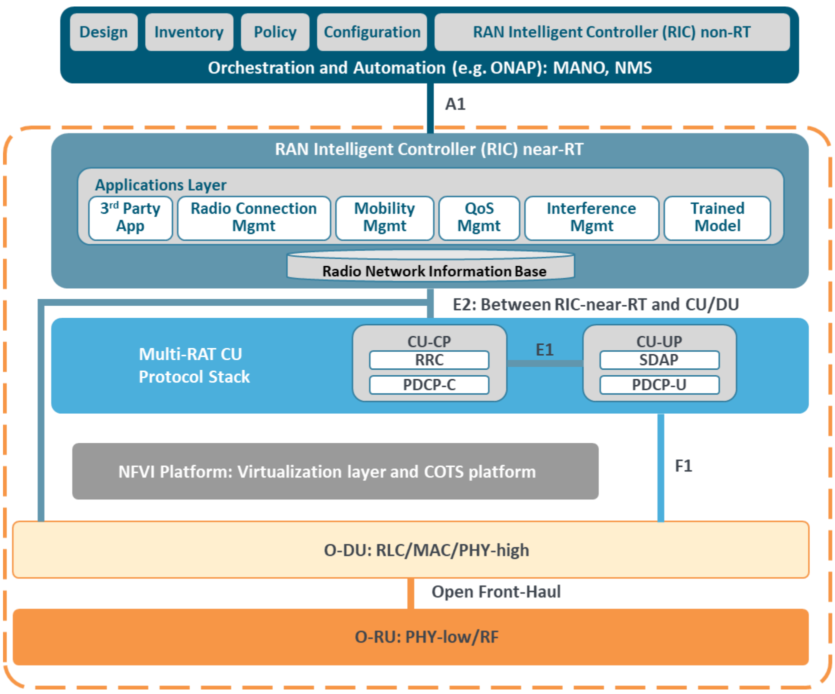

RAN Intelligent Controller (RIC): enables control and optimization of RAN elements and resources based on data collection and analysis.

RIC near-RT: RIC that enables near real-time control

RIC non—RT: RIC that enables non real-time control

Multi-RAT CU: Central Unit (CU) supporting multiple Radio Access Technologies (RAT)

CU-CP: CU Control Plane (CP)

CU-UP: CU User Plane (UP)

RRC: Radio Resource Control protocol layer

SDAP: Service Data Adaptation Protocol

PDCP: Packet Data Convergence Protocol

PDCP-C: PDCP control plane

PDCP-U: PDCP user plane

O-DU: O-RAN Distributed Unit

O-RU: O-RAN Radio Unit

Figure 5: O-RAN Architecture

O-RAN O-RU whitebox solution

One of the nine work groups formed by the O-RAN Alliance is the whitebox hardware workgroup (WG7). The goal of the O-RAN whitebox hardware working group is to “specify and release a complete reference design of a high-performance, spectral and energy efficient whitebox base station”.

The first test platform for the O-RAN O-RU whitebox solution is now available. Full details of the solution are available including reference design schematics and open software so that others can build their own O-RU solutions. The O-RAN O-DU is expected to be available in late 2020.

The WG7 implementations are the result of a collaboration between Analog Devices, Comcores, Intel, Radisys, and Whizz Systems. The O-RU is based on a combination of Intel® Arria 10® System-on-Chip (SoC) and Analog Devices ADV902x Digital Front End (DFE).

Comcores provides the FPGA software or Intellectual Property (IP) that implements the functions of the O-RU on the FPGA. The IP includes implementation of the eCPRI front-end protocol as well as layer 1 offload functionality.

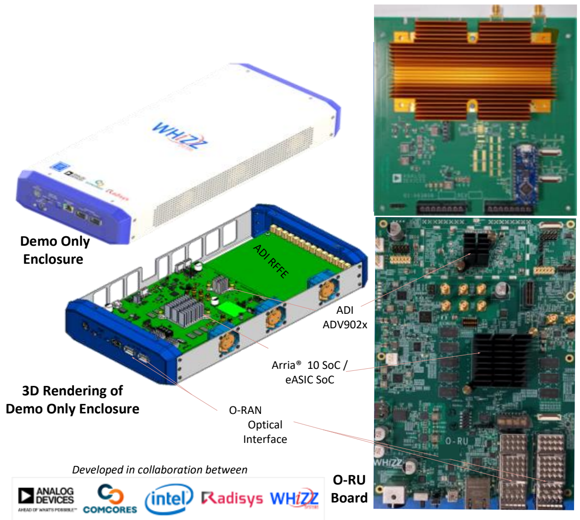

Figure 6: O-RAN O-RU Reference Platform

The Comcores eCPRI solution is the first implementation of an open fronthaul interface. Now, for the first time, it is possible for carriers to deploy a vendor-neutral RAN architecture.

However, it is possible to independently upgrade the Comcores eCPRI solution to accommodate new eCPRI split options or alternative fronthaul technologies. This provides carriers with a flexible, but future-proof solution that can adapt to new requirements and needs.

For more information on specs and availability contact Comcores at sales@comcores.com

The role of Comcores in the acceleration of 5G RAN

Carriers are already making early-stage 5G investments and trials, with initial commercial service launches being deployed alongside existing 4G network assets. However, carriers also need to plan and prepare for the next step; evolution to fully stand-alone 5G networks.

Managing and optimizing costs are paramount. This is the one area of the business case that is fully within carrier control. This requires cost-effective, yet flexible implementations of 5G functions that can support carrier cost goals without limiting revenue opportunities.

The Comcores solution portfolio is designed to support the new functions, technologies and protocols required for 5G deployments along with the options and enhancements required to support flexibility. Comcores reference designs and platforms enable cost-effective implementation and testing of O-RAN solutions providing carriers with open and reliable 5G RAN solutions. This provides carriers with the confidence to accelerate 5G deployments.

Through our participation in collaborations like the O-RAN Alliance, Comcores ensures that we continue to remain at the forefront of developments so that the right capabilities are available when you need them.

Read Part 2 of the Whitepaper series at: