Future Proofing Mobile Network Economics

Assessing the TCO for Cloud RAN and Centralized RAN

BY MONICA PAOLINI
SENZA FILI

SENZA FILI

SPONSORED BY

MAVENIR

1. How to virtualize the RAN

Mobile networks are becoming virtualized end to end – from the core up to and including the radio access network (RAN). As mobile operators embark on the transition, we see a variety of approaches, which vary in the depth and speed of the virtualization process as well as in the chosen topology. A critical decision is where to place functions physically within a virtualized network.

Which functions should operators keep in a centralized data center and which ones should they move toward the edge? What RAN functionality should operators keep at the cell site, and what at the data center? And with multi-access edge computing (MEC) and other edge computing initiatives, where should the edge equipment go? Depending on capacity, coverage, synchronization and latency requirements, operators can use the core, RAN and edge computing equipment at different network locations to optimize performance and cost efficiency.

We are still learning about the best locations in different environments and for different purposes, as well as for different operators' strategies, because the flexibility in the choice of location of network function is a new thing. In legacy networks, the functions are tied to hardware elements, and those, in turn, are tied to a location – it is a model that forces a rigid topology. The new freedom of choice that comes with virtualization gives operators the opportunity to extract more value from their networks, but it also challenges them as they select new topologies for their networks.

Operators are clearly on board to move away from a traditional Distributed RAN (DRAN), which has the baseband unit (BBU) located at the cell site, toward virtualized solutions with the BBU at a remote location. Where should that remote location be? Should it be close to the remote radio unit (RRU) or reside with the evolved packet core (EPC)? How will that choice impact the fronthaul (FH) and backhaul (BH) costs – and hence the overall total cost of ownership (TCO)? In this paper, we explore these questions by comparing the TCO for a Centralized RAN topology and a Cloud RAN topology.

RAN architectures

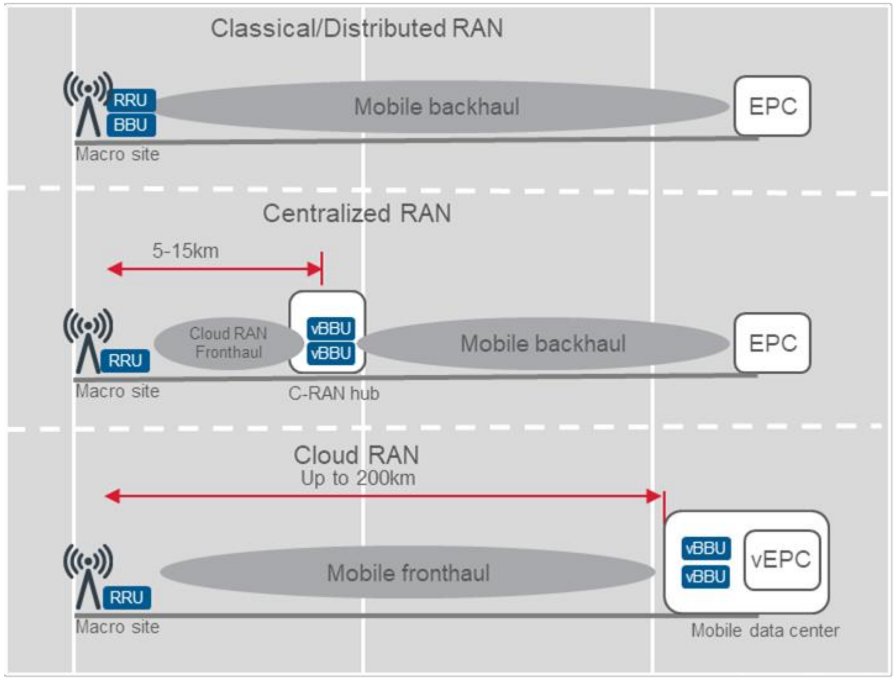

Source: Mavenir

This paper is a companion to "How much can operators save with a Cloud RAN? A TCO model for virtualized and distributed RAN," which compared the TCO for DRAN and Cloud RAN. In this paper, we look at two remote-BBU topologies – Centralized RAN and Cloud RAN – using that same TCO model, and we then compare them to the DRAN TCO.

In a **Cloud RAN** the vBBUs can be colocated with the vEPC, and operators need only mobile fronthaul to connect the RAN to the EPC.

In a **Centralized RAN** the physical BBUs are remote from the radio, but located closer to the RAN. To reach the EPC, a fronthaul link and a BH link are necessary.

In a **DRAN**, the BBU is at the cell site, so a backhaul link connects to the EPC from the RAN.

2. Choosing a functional split: latency and cost tradeoffs

Latency and cost are two main and linked factors in deciding how to virtualize the RAN. The closer the BBU is to the edge, the lower the latency, but also the lower the cost savings. So it is crucial to pinpoint the location that provides the desired cost/performance balance. The choice of FH plays a crucial role in this decision.

Most operators' efforts at RAN virtualization are still in the first generation, which we call Centralized RAN here. The dominant solution today is a functional split between the physical layer (PHY) and the radio frequency (RF), using Common Public Radio Interface (CPRI) for FH. This is defined as option 8 by 3GPP.

Several alternatives to option 8 are emerging, using other functional splits across the protocol stack. The higher the option number, the more processing is done at the BBU. Options 7 and 8 are the ones we consider in this paper.

In first-generation RAN virtualization, operators have to pay for FH from the RAN to the BBU location, and then also for BH from the BBU to the EPC. The combination of FH and BH, along with the use of CPRI in the FH, greatly reduces the scope for the cost savings expected from RAN virtualization.

With option 8, using CPRI for FH, performance is good, but the bandwidth requirements are so high that the technology does not scale to 5G or to an increasing number of 4G cell sites. CPRI is also limited in distance: when using CPRI, operators have to locate the BBUs within 15 km of the RRU in the cell site.

Option 7, the other scenario assessed in this paper, is an intra-PHY split. It reduces the latency and bandwidth requirements. At the same time, it allows for centralized and coordinated traffic management and efficient allocation of network resources, both for today's 4G networks and for future 5G networks. According to Mavenir, option 7 requires less than 1/10 of the bandwidth than option 8 does – e.g., an FH link that requires 2.4 Gbps with option 8 would need only 170 Mbps with option 7.

We did not consider lower splits because, although they reduce the FH performance requirements and thus the overall TCO, they limit the operator's ability to use new tools to manage traffic across the virtualized RAN.

3GPP functional splits

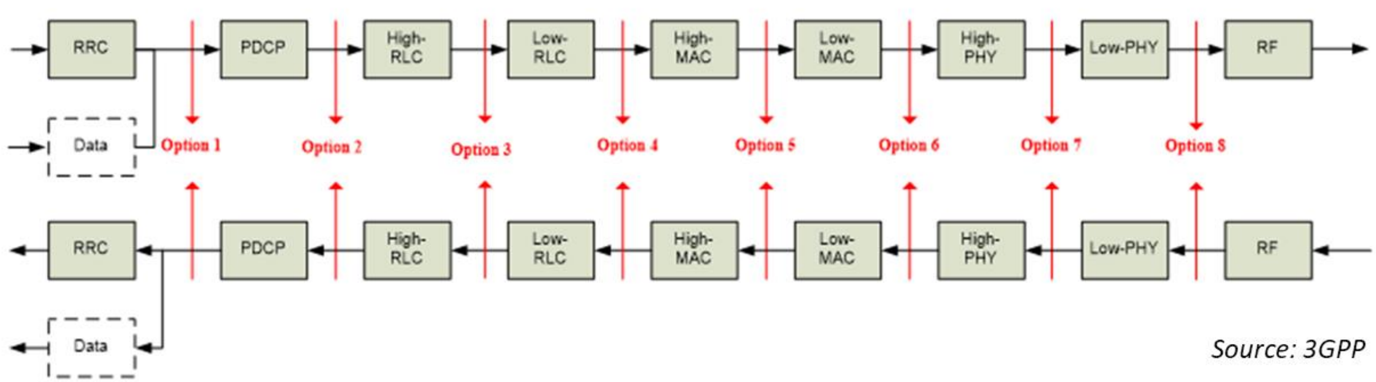

Source: 3GPP

3. TCO model assumptions

We built a TCO model to look at the financial differentiators of Cloud RAN adoption over a period of 5 years. In a previous paper, we compared the TCO for a Cloud RAN and DRAN network, both with the same type and number of RRUs. In this paper, we shift the focus to comparing Centralized RAN and Cloud RAN, to assess different RAN virtualization solutions.

The model covers a set of macro and small cells that share a vBBU pool in a high-density area with a mix of macro cells, outdoor small cells, and indoor small cells, using cost assumptions that are within the typical range in a North American or European market. Mavenir provided the cost assumptions, based on inputs from its operator customers. We chose high-density areas because this is where operators initially plan to deploy virtualized RAN solutions.

To look at the opex impact of the functional split, in this analysis we assumed that RRU and BBU equipment is the same in the Centralized RAN and Cloud RAN cases. We assumed three-sector 2x2 MIMO macro cells, 4x4 MIMO outdoor small cells, and 2x2 MIMO indoor cells. We expect the relative costs of Centralized RAN and Cloud RAN to remain constant as we move to new MIMO configurations or 5G.

The difference between the two cases is in the costs and requirements for BH and FH.

In the Cloud RAN case, we used the option 7 intra-PHY functional split for the FH. This eliminates the need for CPRI-based FH, reducing the bandwidth and cost requirements of the FH and making Cloud RAN cost-effective in a wider set of environments. The option 7 split allows the operator to use Ethernet-based FH or other FH solutions that are cheaper than CPRI.

In the Centralized RAN case, we used option 8 for the FH functional split. This is more expensive than option 7 because it requires CPRI. Also, BH is required to transport traffic from the BBU pool to the EPC. (In the Cloud-RAN scenario, the BBU is co-located with the EPC, so there is no need for BH.)

TCO model assumptions
Framework. Our model compares the TCO, over five years, of a Centralized RAN versus a Cloud RAN greenfield network with vBBUs. All capex is in year 1, during deployment. The model covers the RAN all the way to the EPC.
Network. 100 macro cells, 200 outdoor small cells, 250 indoor small cells.
Technology. Macro cells: three-sector LTE 2x2 multiple-input, multiple-output (MIMO). Outdoor small cells: single-sector LTE 4x4 MIMO. Indoor small cells: single-sector LTE 2x2 MIMO.
FH/BH. Centralized RAN uses CPRI and the 3GPP option 8 split. It requires a CPRI FH connection to connect RRU to the BBU, and a BH connection from the BBU to the EPC. Cloud RAN uses the option 7 intra-PHY functional split in the FH, which does not need a CPRI interface. Without CPRI, the BBU can be located farther away and co-located with the EPC. As a result, there is no need for a BH link from the BBU to the EPC.
vBBU multiplexing. In both scenarios, vBBU resources can be dynamically allocated to RRUs with multiplexing. We estimate that, when used, multiplexing reduces the BBU capacity requirements by 50%.
Equipment. In the DRAN case, the RRU and BBU are at the cell site. In the Centralized RAN and Cloud RAN cases, the RRU is at the cell site, and the vBBU pool is at a remote site.

4. RAN virtualization beyond CPRI

Our TCO analysis shows that a Cloud RAN using the option 7 functional split delivers a 23% cumulative TCO (capex and opex) savings over a CPRI-based, option 8 Centralized RAN solution over five years. Nearly all the cost savings in the Cloud RAN case come from the opex, because the equipment and installation costs are the same as Centralized RAN, except for the BH equipment, which contributes to a 1% capex savings.

The opex reduction in the Cloud RAN case comes exclusively from the BH and FH recurring costs and creates a 29% cost savings over the Centralized RAN case. In the Cloud RAN case, not only is the BH not required, but FH costs are lower. The results show how crucial the selection of FH technology and the location of the BBUs are for cost-effective virtualized RAN solutions.

The main sources of cost savings with Cloud RAN are BH and FH. In our model, we used median BH and FH costs, but there is a large amount of variability, depending on the country, location within the operator's footprint, and technology selected. In all cases, however, the cost savings with Cloud RAN are substantial because, unlike Centralized RAN, it does not require BH.

The availability and cost of fiber connectivity needed for the FH play a crucial role in the TCO. Here we assume the operator leases fiber. However, if the operator owns a fiber network, FH costs are typically much lower than for a lease from a third party. At the same time, if fiber is not available at the cell site and the operator decides to install a fiber network, then the capex will increase substantially, and the opex will decrease as the leasing fees go away. Typically, the capex increase will be much higher than the opex decrease over five years, because it usually takes more than five years to recoup the investment in deploying fiber. Finally, if fiber is not available and the operator is not willing to deploy it, Centralized RAN may simply not be a feasible option at that location.

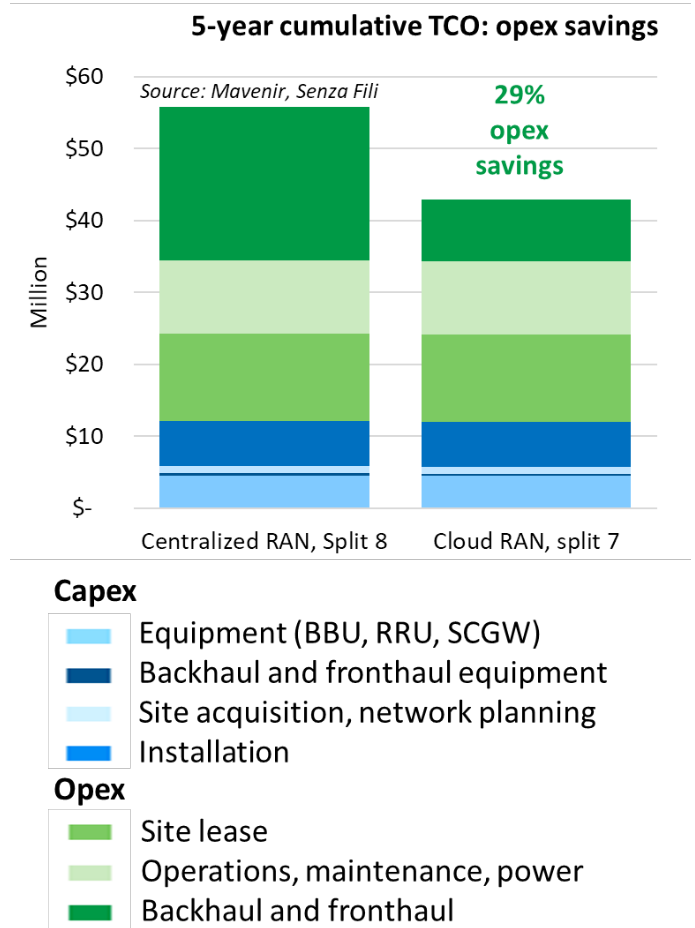

Figure 1

To put these results in a broader context, we have gone back to the DRAN analysis we did previously (see figure on the right) and compared DRAN to the two virtualized RAN cases, to see what gains there are from virtualizing the RAN. The overall TCO decreases by 18% as we move from DRAN to Centralized RAN, but the greater savings are for Cloud RAN – 37% over five years.

For Centralized RAN, the cost savings over DRAN come almost exclusively from the capex (49% versus 2% in the opex). Opex from site lease, operations, maintenance, and power are substantially lower in the Centralized RAN case, but the higher combined BH and FH is still high enough, however, to erase virtually all the other opex savings.

For the Cloud RAN, the savings were 49% for capex and 31% for opex, as previously reported. The move to a virtualized RAN allows operators to use less-expensive off-the-shelf hardware and to benefit from efficient use of baseband resources with vBBU pooling. Opex savings come from lower maintenance, power, and operational costs, because operators deploy less equipment at the edge, and the remote equipment is typically cheaper to manage.

The two comparisons show that RAN virtualization can deliver cost savings over five years in the scenarios we used in our TCO model for both the Centralized RAN and Cloud RAN. However, the choice of RAN virtualization model has a significant impact on the overall cost savings, with Cloud RAN delivering the largest cost savings – 37% over DRAN (Figure 2) and 23% over Centralized RAN (Figure 1).

5-year cumulative TCO:
DRAN, Centralized RAN and Cloud RAN

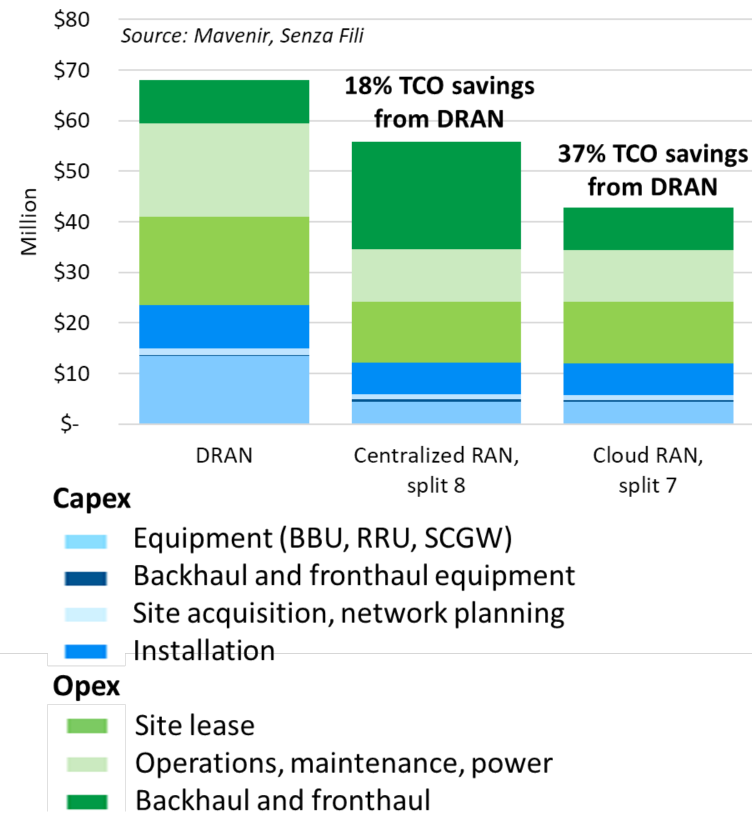

Figure 2

References

3GPP (2016) 3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Study on New Radio Access Technology: Radio Access Architecture and Interfaces, 3GPP TR 38.801.

Monica Paolini (2017) How much can operators save with a Cloud RAN? A TCO model for virtualized and distributed RAN, white paper.

xRAN Forum (2017) Fronthaul Working Group, white paper.

About Mavenir

Mavenir is purpose-built to redefine mobile network economics for Communication Service Providers (CSPs). Our innovative solutions pave the way to 5G with 100% software-based, end-to-end, Cloud Native network solutions. Leveraging industry-leading firsts in VoLTE, VoWiFi, Advanced Messaging (RCS), Multi-ID, vEPC and Cloud RAN, Mavenir accelerates network transformation for more than 250+ CSP customers in over 130 countries, serving over 50% of the world's subscribers. Mavenir embraces disruptive, innovative technology architectures and business models that drive service agility, flexibility, and velocity. With solutions that propel NFV evolution to achieve web-scale economics, Mavenir offers solutions to CSPs for cost reduction, revenue generation, and revenue protection.

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modeling, market forecasts and research, white paper preparation, business plan support, RFP preparation and management, due diligence, and training. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, leverage these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com, or contact us at info@senzafiliconsulting.com or +1 425 657 4991.

About the Monica Paolini

Monica Paolini, Ph.D., is the founder and president of Senza Fili. She is an expert in wireless technologies and has helped clients worldwide to understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She has frequently been invited to give presentations at conferences and has written several reports and articles on wireless broadband technologies. She has a Ph.D. in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy). You can contact Monica at monica.paolini@senzafiliconsulting.com.