How much can operators save with a Cloud RAN?

A TCO model for virtualized and distributed RAN

By Monica Paolini
Senza Fili

Sponsored by

1. Evaluating the cost savings of moving to Cloud RAN

Virtualization opens up new ways to architect, deploy and operate wireless networks, and its flexibility allows mobile operators to experiment with new network topologies. Virtualization has a profound impact on mobile networks end to end. As it shifts some of the traditionally centralized core functionality closer to the edge with initiatives like multi-access edge computing (MEC), it also pulls the radio access network (RAN) infrastructure in the opposite direction, away from Distributed RANs (DRAN) in the edge to a shared, centralized location – with the Centralized RAN and Cloud RAN architectures (see diagram). In a Cloud RAN, the Common Public Radio Interface (CPRI) fronthaul (FH) is replaced by a lower-bandwidth functional split, and baseband units (BBUs) can be moved farther from the edge and become virtual BBUs (vBBUs).

Through virtualization, the combination of a more distributed core and a more centralized edge in a Cloud RAN creates an intermediate area where processing can be shifted. From a performance perspective, we can expect improved traffic and interference management, more advanced quality of service (QoS), lower latency and more efficient use of network resources. From a financial perspective, operators can save money – and squeeze more value from their wireless networks – when they place processing where it is most effective and least expensive.

In this paper, we present a total cost of ownership (TCO) model to show the cost savings an operator can expect in a Cloud RAN deployment over 5 years. We drill down into the specific financial benefits that macro cells and indoor and outdoor small cells contribute to the overall network. The specific cost savings vary across markets, operators, and environments (e.g., rural vs metropolitan, dense indoor vs low-density suburban environments), but the drivers are the same. Although our model enables us to look at how variations in cost assumptions and deployment strategy affect the business model, in this paper we direct our attention to a base case model and explore the joint impact that these drivers have in motivating the shift to a Cloud RAN architecture.

RAN architectures

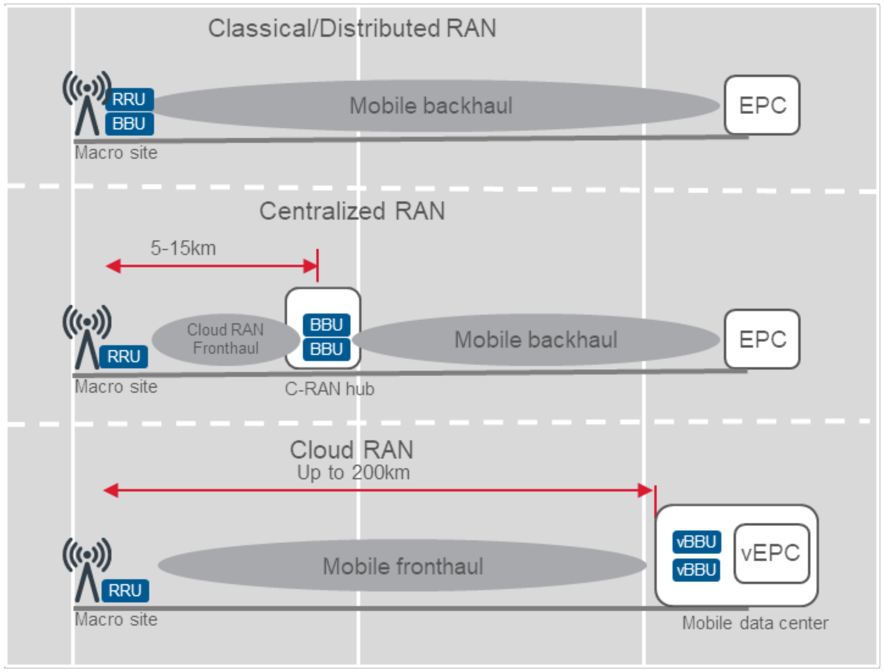

Source: Mavenir

Distributed RAN, Centralized RAN and Cloud RAN

Most networks today use a DRAN architecture in which the two base station components – the remote radio unit (RRU) and the BBU – are both located at the network-edge cell site. Virtualization makes it possible to physically separate them in a Centralized RAN or Cloud RAN: the RRU remains at the cell site, but the BBU/vBBU moves to a central location, where BBU/vBBU processing can be pooled for multiple RRUs. vBBU pooling contributes to operational efficiency and cost savings, and improves traffic and interference management. Having less equipment at the cell site speeds up deployments, and lowers the capex and opex. Centralized RAN and Cloud RAN require a high-reliability and low-latency FH link between RRU and BBU/vBBU. High costs for CPRI – the default FH interface today – have so far limited the adoption of Centralized RAN, but functional splits in the FH allow a sharp reduction in FH costs in the Cloud RAN.

2. Drivers to the virtualized RAN: a TCO model

We built a TCO model to look at the financial drivers to Cloud RAN adoption over a period of 5 years, and compared it to a DRAN network with the same type and number of RRUs. The model covers a set of macro and small cells that share a vBBU pool in the Cloud RAN scenario, and have BBUs at the cell site in the DRAN scenario. It allows us to look at the cost savings that operators can achieve with different Cloud RAN topologies and in different environments. The drivers are the same across environments, but their relative impact on the business case differs. The model helps us to find out where and under what conditions a move to a Cloud RAN architecture makes financial sense.

In this paper, we focus on a base case that covers 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 used in the model from inputs from operator customers. We chose high-density areas because this is where operators initially plan to deploy Centralized RAN and Cloud RAN.

We assumed three-sector 2x2 MIMO macro cells, 4x4 MIMO outdoor small cells, and 2x2 MIMO indoor cells, but expect that the relative cost savings of Cloud RAN and DRAN are preserved as we move to new MIMO configurations or to 5G, as the relative cost differences between Cloud RAN and DRAN are comparable. For the FH, we used a functional split in which some of the vBBU functionality stays in the RAN and some resides in the vBBU pool. A functional split eliminates the need for CPRI-based FH, and hence it reduces the capacity and cost requirements of the FH, and it makes Cloud RAN cost-effective in a wider set of environments. We assumed an option-7 split (intra-PHY split), which leaves most of the baseband processing in the remote vBBU pool but allows the operator to use Ethernet-based FH or other FH solutions. See the companion white paper for the TCO analysis for different FH solutions and functional splits.

One of the advantages of Cloud RAN is the ability to virtualize the BBU pool and allocate resources as needed. As a result, the same vBBU server can support different RRUs, as demand shifts from one location to another through the day or because of specific events. For instance, during rush hour, activity in outdoor locations grows as people go to or leave their offices, but during the workday, more traffic is generated from indoor workplace locations.

TCO model assumptions

Framework. Our model compares the TCO of a DRAN vs a Cloud RAN greenfield network with vBBUs, over 5 years. All capex is in year 1, during deployment. It covers the RAN all the way to the Evolved Packet Core (EPC).

Network. 100 macro cells, 200 outdoor small cells, 250 indoor small cells.

Technology. Macro cells: 3-sector LTE 2x2 multiple-input, multiple-output (MIMO). Outdoor small cells: single-sector LTE 4x4 MIMO. Indoor small cells: single-sector LTE 2x2 MIMO.

Fronthaul/backhaul. DRAN uses backhaul (BH). Cloud RAN uses an option 7 intra-physical layer (PHY) functional split in the FH, which does not need a CPRI interface.

vBBU multiplexing. In the Cloud RAN, 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 Cloud RAN case, the RRU is at the cell site, and the vBBU pool is at the remote site.

3. Cloud RAN can save 37% in costs compared with DRAN

The base case in our TCO model demonstrates a 37% reduction in deployment and operational costs over 5 years, from a 49% savings in capex in year 1 and a cumulative 31% savings in opex over the 5 years.

Capex savings primarily come from a reduction in equipment costs in the vBBU. The RRU costs are largely the same in both the DRAN and Cloud RAN scenarios, but the vBBU costs are lower in the Cloud RAN scenario because the BBUs are virtualized. Virtualization makes it possible to use both less-expensive non-proprietary hardware, and BBU pooling. With pooling, the efficiency in the use of vBBU resources increases, and the vBBU pool needs less baseband processing capacity (and hence less hardware). The reduced need for equipment at the cell site not only lowers capex, it enables faster deployment and more flexibility of equipment location. Planning and installation are also cheaper for Cloud RAN, but the cost reduction for them is less pronounced, because mobile operators still have to deploy the RRU at the edge.

Opex savings are mostly due to the reduction in maintenance, power and operations costs, in more-centralized vBBU locations that are typically easier to access and cheaper to operate. Leases at the cell site also cost less, because of the reduction in equipment located there. Because the model assumes a functional split, we assumed the cost for the FH in the Cloud RAN scenario to be the same as the cost for BH in the DRAN scenario. Had a CPRI-based FH been used instead, the Cloud RAN FH costs would have been substantially higher, and the opex savings reduced – to 11%, from the 31% we demonstrated in the base case.

5-year cumulative TCO: capex and opex

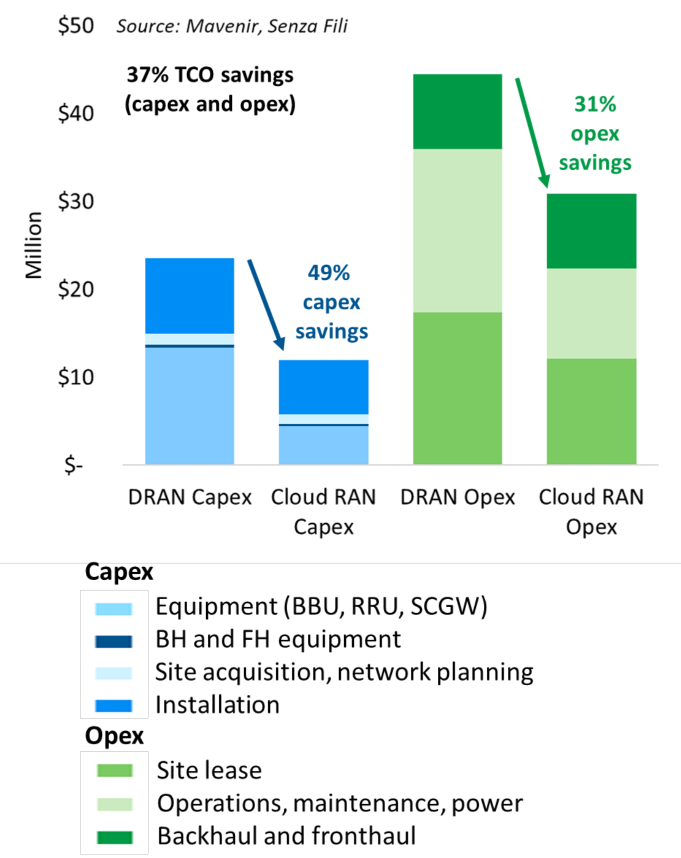

Beyond the TCO base case

Our TCO model focuses on a base case that reflects cost savings levels that mobile operators can achieve in many markets. However, in addition to the cost savings, operators stand to benefit from improved performance. Improved performance does not lower the TCO, but lowers per-bit costs (and can improve revenues and QoE). To keep the assessment conservative, we have excluded the performance gains from the TCO base case. Also, operators that own, or otherwise have low-cost access to, an FH/BH network can significantly lower their operational costs. In a scenario in which the operator has free access to FH/BH, the cost savings can reach 42% (an increase from 31% in the base case). A neutral-host model can also lower costs to an individual operator, because operators can share the costs of the network deployment and operations with other operators. In a Cloud RAN neutral-host scenario, we project 48% TCO savings (54% from capex, 45% from opex).

4. Cloud RAN for macro and small cells

One of the appeals of Cloud RAN is that this architecture integrates different network layers, because the remote, virtualized BBUs support both macro and small cells. Eventually, we may not even think of wireless networks as multi-layer, but rather as a collection of RRUs with different settings (e.g., power, capacity, location, range, radio frequency) whose transmission has to be coordinated at the BBU location. As we look at the TCO today, however, the financial proposition for macro and small cells is different, because the cost drivers have different weights.

Perhaps not surprisingly, the biggest cost savings of Cloud RAN come from indoor small cells. TCO savings for macro cells are 35% over 5 years, while they are 36% for outdoor small cells and 44% for indoor small cells. The breakdown for macro cells is 43% from capex and 31% from opex. For outdoor small cells, it is 46% from capex and 31% from opex. And for indoor small cells, it is 66% from capex and 27% from opex.

Capex savings are comparable for macro cells and outdoor small cells. The bigger capex savings for indoor cells come from equipment costs, due to the smaller, less intrusive hardware and the ability to pool the baseband processing remotely. Cost savings from installation are lower for indoor small cells because we expect indoor small cells to be deployed mostly in environments where installation costs are low. In environments where it costs a lot to install small cells, the business case for them is not robust, regardless of whether it is a DRAN or a Cloud RAN, and we expect limited deployments of indoor small cells in those locations. However, cost savings from installation in outdoor locations are larger than for indoor cells, because the reduction in the amount of equipment has a deeper impact outdoors.

Opex savings are larger for the outdoor infrastructure than the indoor, and they come primarily from operations, power and maintenance. Moving the baseband processing to a remote, indoor location, such as a central office, reduces the expense of operating and maintaining the wireless infrastructure. With co-location of equipment and centralization, there is less need for truck rolls. Also, power at a data center or a central office may be cheaper than at a cell tower.

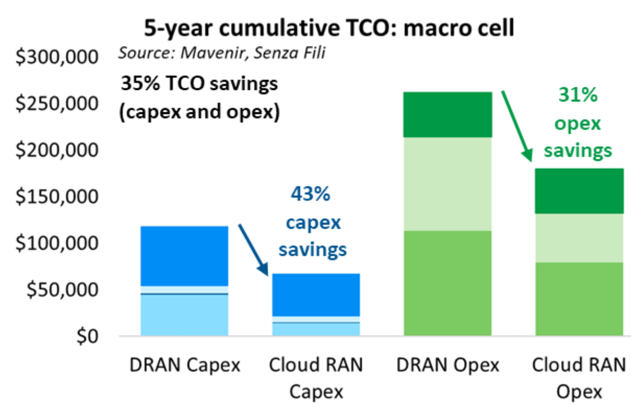

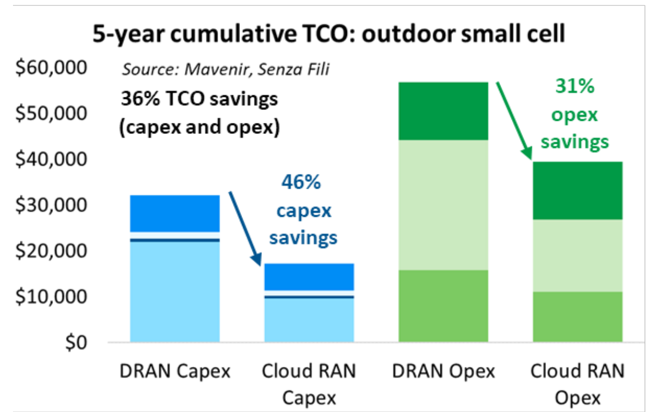

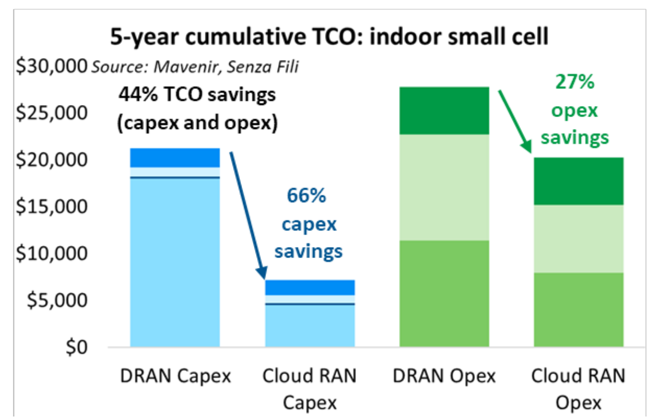

About Mavenir

Mavenir is a leader in accelerating network transformation and redefining network economics for Service Providers, by offering a comprehensive product portfolio across every layer of the network infrastructure stack. From 4G to 5G application/service layers to packet core and RAN – Mavenir leads the way in evolved, cloud-native networking solutions enabling innovative and secure experiences for end users. Through its industry first VoLTE, VoWiFi, Advanced Messaging (RCS), Multi-ID, vEPC, Visual Voicemail and Cloud RAN solutions, and serving 350+ customers across 3.5 billion subscribers, Mavenir's platform enables Service Providers to successfully deliver next-generation vision today and realize new revenue streams and operational efficiencies.

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modelling, 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 author

Monica Paolini, PhD, 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 PhD 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.