

Which Open RAN is best for you?

TCO tradeoffs: transport versus location

Monica Paolini, Senza Fili



Which Open RAN is best for you?


When you explore how to deploy Open RAN, one of the first things you want to find out is whether and how you can reduce your RAN costs. Because there are many ways to deploy Open RAN, the answer depends on how and where you plan to do it, and what your specific costs are.

To find the most cost-efficient way to deploy Open RAN in your network, you need to assess multiple factors. A crucial factor is the tradeoff between transport and location.

We developed a financial model that allows you to compare the Open RAN TCO for three scenarios that use different transport cost assumptions and show how transport costs may drive network topology decisions.

With Open RAN, operators with high transport costs (HTC) can save 30% over 5 years, if they use a distributed topology with the distributed unit (DU) at the cell site, instead of a centralized topology with both the DU and centralized unit (CU) at remote locations.

Operators with low transport costs (LTC) are better off with a centralized topology, and can save 30% over a distributed topology.



Source: Senza Fili, Mavenir, HFR Networks

We demonstrated the TCO advantage of Open RAN architectures over traditional RAN architectures in three earlier papers, "[Future proofing mobile network economics](#)," "[How much can operators save with a Cloud RAN?](#)" and "[In-building virtualization](#)."

The new TCO model moves one step ahead and examines the financial impact of Open RAN architecture choices under variable costs and resource availability. In this paper we focus on transport costs, and upcoming papers will focus on other aspects of Open RAN deployments.

Open RAN gives operators flexibility in how they architect their RAN, allowing them to have distributed topologies with more hardware and processing toward the edge, and centralized topologies with DUs and CUs in remote locations in data centers.

The location-related costs vary across locations and operators. They depend on capex items such as site acquisition, deployment and data center set-up fees, and to a larger extent on opex items such as site leases, maintenance and power.

At the same time, transport costs may vary even more than location-related costs. As a result, the higher transport costs due to demanding fronthaul (FH) requirements increase the TCO in a centralized architecture.


We compared the location/transport tradeoffs in distributed and centralized architectures by keeping the location costs constant and varying the transport costs.

Our base case – high transport costs (HTC) – is more likely to apply to a brownfield mobile operator that does not own the transport infrastructure and has to pay market prices for transport. The low-transport-cost (LTC) case is more typical of an operator that owns a transport network and hence has a low transport cost basis.


Because the only difference between the HTC and LTC cases is transport costs (i.e., \$1,000 and \$100 per month for a one Gbps link, respectively, with per-Gbps costs declining as link capacity goes up), the overall TCO for the HTC case is higher than for the LTC case.

The next pages show the TCO results for both the HTC and LTC cases. The difference in transport costs determines which of three scenarios is more cost efficient. If transport costs are high, having the DU at the cell site (scenario 1) is the lowest-cost option. If transport costs are low, locating both the DU and CU in remote locations (scenario 2 and 3) reduces costs.

In some cases, locating the DU and CU at the same locations may not be a desirable topology. For instance, the DU location may be too far away from the cell site and there are limitations to the length of a fronthaul link. If DU and CU are not in the same remote location (scenario 2), there is a slight cost increase over scenario 3, due to the need to support more locations and for midhaul (MH) connections from the DU to the CU site. The choice between scenarios 2 and 3 will most likely be dependent on topology constraints (e.g., cell site locations and density, or distance to the CU), rather than cost considerations.



Model scenarios



TCO model: scenarios and assumptions

Our model compares the TCO for three scenarios:

- **Scenario 1 – Distributed topology:** DUs are located at the cell sites with RUs, and MH connects DUs to the CU.
- **Scenario 2 – Partially centralized topology:** DUs are at remote locations, separate from the CU's location. FH connects RUs to DUs, and MH connects DUs to the CU.
- **Scenario 3 – Centralized topology:** CU and DUs are in the same location, and FH connects RUs to the CU/DUs.

The results exclude the RU cost contribution because it is the same in all scenarios.

Our model covers Open RAN scenarios that include RU, CU, DU, MH and FH capex and opex costs over six years, with the capex incurred in the first year. Because the RU-related costs are constant across scenarios, we do not include them in the results shown in this paper as they do not affect the transport/location tradeoffs.

Cell sites: 3 sectors, 5G-NR 20 MHz channels, frequency division duplex (FDD) with 4x4 multiple input, multiple output (MIMO).

Network: 5,001 cell sectors, 1,667 cell sites, 10 DU locations (scenario 2), and one CU location.

Transport: The results shown here assume shared Ethernet transport, with star packet links, up to 15 km for FH and 100 km for MH, using radio over Ethernet (RoE) and supporting the eCPRI (Enhanced Common Public Radio Interface) 7.2x O-RAN Open Fronthaul Interface over colored wavelength-division multiplexing (WDM). The model also calculates the TCO for ring solutions.

Remote locations: DUs (scenarios 2 and 3) and CU are in data centers where hardware resources are shared across RUs, resulting in higher efficiency because of pooling gains due to DU and CU resource sharing across the Open RAN footprint.


Cost, requirements, and traffic inputs were from trials and customers of Mavenir and HFR Networks.

TCO for high transport costs (HTC) case

The HTC case favors scenario 1 (RU and DU at the cell site), with the cost benefits coming from a lower opex.

The higher costs of installing more equipment at the cell sites give scenario 1 the highest capex. However, the lower transport requirements in scenario 1 reduce the overall opex compared to scenarios 2 and 3.

Scenario 3 is slightly better than scenario 2, with a 2% lower capex, a 3% lower opex, resulting in a 3% reduction in the cumulative TCO. This is due to scenario 2's additional costs incurred by having additional locations to host the DUs and the addition of MH costs from the DUs to the CU.




Source: Senza Fili, Mavenir, HFR Networks

TCO for low transport costs (LTC) case

Scenarios 2 and 3 are the best ones for LTC operators, combining the benefits of both a lower capex and a lower opex.

The percentage differences in capex among the scenarios is the same as in the HTC case. As in the HTC, a lower cell-site equipment cost favors the centralized scenarios 2 and 3. The lower transport costs drive most of the opex reduction. Further savings in opex come from the lower cost of concentrating the DU and CU capex in remote locations.

As in the HTC case, the difference between scenarios 2 and 3 is small, with a 2% lower capex and 3% lower opex, resulting in a 3% reduction in cumulative TCO in scenario 3.



Source: Senza Fili, Mavenir, HFR Networks


Comparing opex and capex in the HTC and LTC cases

The capex in both the HTC and LTC cases is the same. Across scenarios, the major difference is in the DU costs are 60% and 61% lower in scenarios 2 and 3, respectively, than in scenario 1, which has higher installation costs and hardware costs. CU costs are the same, because the CU is located remotely in one location in all scenarios. For the combined FH and MH transport, scenario 1 has the lowest capex because it requires only MH; scenario 2 is 22% higher because of its MH/FH combination, and scenario 3, using FH only, is 19% higher than scenario 1.

Not only does the opex total change across scenarios and across the HTC and LTC cases, its composition changes as well.

In the HTC case, transport costs account for a larger share of the opex: 76%, 97% and 98% in the three scenarios, respectively. The higher share of transport costs in scenarios 2 and 3 is due to the higher transport requirements from FH.

In the LTC case, DU costs also play a larger role across all scenarios, with DU costs accounting for 45% of opex in scenario 1, 7% in scenario 2, and 4% in scenario 3. Transport costs account for 52% of opex in scenario 1, 90% in scenario 2, and 92% in scenario 3.



Takeaways

Our model shows that the TCO crucially depends on the Open RAN topology the operator selects and, more specifically, on the resource and transport costs dictated by the chosen topology.

Transport costs can steer an operator toward different Open RAN topologies, either for the entire network or for specific locations within the network.

For an operator with relatively high transport costs (HTC), a distributed topology is more cost effective. Placing DUs at the cell site (scenario 1) reduces the transport requirements but does increase the equipment and operating costs for the DUs. DU-driven costs are higher because equipment at the cell site is typically more expensive to install and operate and because there are no pooling benefits from sharing resources at a remote location. However, if the increase in DU-related costs is lower than the increase in transport costs across scenarios, as it is in our model, then the DU should be at the cell site.

Operators with lower transport costs (e.g., they own the transport network) benefit from a more centralized topology (scenarios 2 and 3). In addition, the lower transport costs enable them to take advantage of the lower DU-driven hardware and operating costs.

Operators with lower transport costs and better transport resources stand to benefit more from Open RAN. This is not just because the transport costs are lower, but also because a centralized topology unlocks pooling gains that are not available in a distributed topology.

HTC operators may also benefit from centralized topologies. The confluence of evolution trends that are outside the scope of our model – virtualization, cloud-native and containerized architectures, edge infrastructure, network slicing – may change the tradeoffs between location and transport. For instance, more efficient pooling of network resources and lower costs for remote locations may make the move to a centralized Open RAN financially more attractive.

Finally, the crucial impact of transport and remote location costs creates an opportunity for cloud and transport providers to offer new services or expand the current ones using new network-as-a-service business models. Mobile operators and other wireless service providers can take advantage of new cost dynamics as they transition to Open RAN and want to explore new ways to manage their end-to-end wireless networks.

About Mavenir



Mavenir is the industry's only end-to-end, cloud-native Network Software Provider focused on accelerating software network transformation and redefining network economics for Communications Service Providers (CSPs) by offering a comprehensive end-to-end product portfolio across every layer of the network infrastructure stack. From 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. Leveraging industry-leading firsts in VoLTE, VoWiFi, Advanced Messaging (RCS), vEPC, and OpenRAN vRAN, Mavenir accelerates network transformation for more than 250+ CSP customers in over 120 countries, which serve 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 help CSPs with cost reduction, revenue generation, and revenue protection. Learn more at www.mavenir.com

About Senza Fili



Senza Fili provides advisory support on wireless technologies and services. At Senza Fili we have in-depth expertise in financial modeling, market forecasts and research, strategy, business plan support, and due diligence. 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, use 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.senzafili.com.

About Monica Paolini



Monica Paolini, PhD, founded Senza Fili in 2003. 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 frequently gives presentations at conferences, and she has written many reports and articles on wireless technologies and services. 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).

© 2020 Senza Fili. All rights reserved. We thank Mavenir for supporting the preparation of this paper. The views and statements expressed in this report are those of Senza Fili, and they should not be inferred to reflect the position of Mavenir or other parties mentioned in this document. The document can be distributed only in its integral form and acknowledging the source. No selection of this material may be copied, photocopied, or duplicated in any form or by any means, or redistributed without express written permission from Senza Fili. While the document is based on information that we consider accurate and reliable, Senza Fili makes no warranty, express or implied, as to the accuracy of the information in this document. Senza Fili assumes no liability for any damage or loss arising from reliance on this information. Mavenir and the M logo are trademarks owned by Mavenir Systems, Inc. Trademarks mentioned in this document are the property of their respective owners. Photos by Shutterstock.