What’s the best way to get to Open RAN?

Optimizing pooling gains can reduce TCO by up to 42%

Monica Paolini, Senza Fili

What’s the best way to get to Open RAN?

There are multiple migration paths to Open RAN. Each operator has the flexibility to pick the path best suited to its needs and what is available and cost effective in its network.

In this paper, we compare the TCO for two operator cases. In the first case, the operator has access to low-cost transport and may own a fiber network that reaches most of its macro cell sites. In the second case, the operator is in the opposite situation: it faces high transport costs and limited fiber availability at the macro cell sites. This is likely to change over the next few years, but the operator does not want to wait until then to deploy Open RAN.

How should the adoption of Open RAN differ in the two cases? And what does this tell us about all the operators that find themselves in a midway situation?

In an Open RAN deployment, should the DUs be at the cell site or co-located in a remote data center with the DUs? It depends on transport costs, which unlock different sources of pooling gains.

Operators with low transport costs can reduce the TCO by 36% to 42% by locating the DUs in a centralized location with the CU.

Operators with high transport costs can reduce the TCO by 25% to 12% by moving the DUs to the cell site, and leaving the CU in a centralized location.

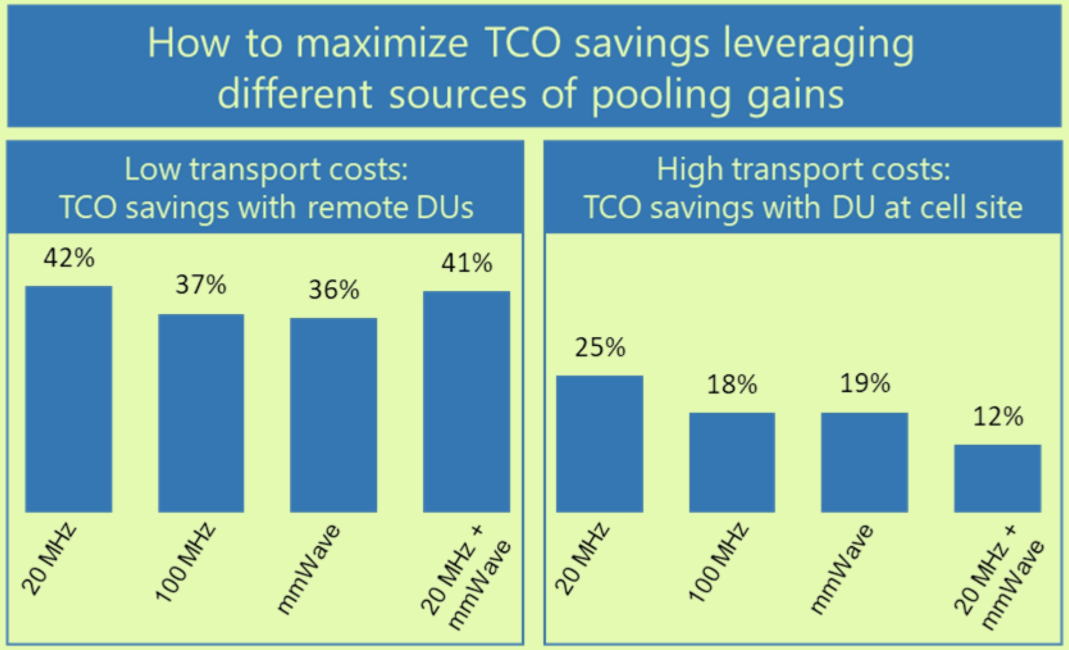

Sources: Mavenir, Intel, HFR Networks, Senza Fili

Which Open RAN is best for you?
Get the companion paper

This is a follow-up to our earlier “Which Open RAN is best for you?” paper based on the same TCO model. The model examines the financial impact of Open RAN architecture choices under variable costs and resource availability. The first paper was on how transport costs affect the topology of Open RAN deployments, and it provided a more detailed description of the underlying model. The current paper looks at how operators can maximize pooling gains as they transition to Open RAN.

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.”

Open RAN is a fundamentally new approach to planning for, deploying and operating the RAN. In Open RAN, virtualization, disaggregation, open interfaces and multi-vendor coexistence create a more flexible and dynamic wireless network, coupled with a core that is undergoing a similar transformation toward openness. These benefits of Open RAN are truly valuable only to the extent to which they also increase efficiency in the use of network resources – e.g., spectrum and equipment – and reduce the cost of providing wireless connectivity.

Our analysis shows how the network topology leverage different sources of Open RAN pooling gains and how these, in combination with the transport costs, affect the TCO. In the first paper, we showed that an operator with high transport costs is better off keeping the DU at the cell site, and one with low transport costs benefits from locating the DU at a remote data center with the CUs. Here we go one step further and look at how the pooling gains differ for operators with high versus low transport costs.

Pooling gains accrue because the hardware for the DU and CU can be shared by multiple RRUs. In a traditional, fully distributed RAN, RU, DU and CU processing is all allocated to a single cell – or sector – thus eliminating the potential for pooling gains. In an Open RAN environment, pooling gains come from different sources:

- DU at the cell site – Hardware serving multiple RRUs is shared, and this leads to a reduction in equipment-related capex and opex. One server at the cell site may act as the DU for multiple RRUs. As the cell site's capacity increases, the hardware costs may stay the same or increase at a slower rate than the capacity. This results in a lower per-bit cost as the cell site's capacity grows.
- DU at a centralized data center – Hardware serving multiple cell sites is shared, and the reduction in equipment-related capex and opex is higher than when the DU is located at the cell site, because sharing is more efficient and larger servers can be used – processing-related costs are lower for larger servers.
- CU at a centralized data center – Hardware is shared, and the CU can serve multiple cell sites. In our TCO model, CUs are always remote, so the pooling gains are the same across scenarios and cell types.

To realize a net TCO benefit, however, pooling gains have to be balanced with the transport costs. The pooling gains may not be sufficiently high to justify centralized DUs, which typically increase transport costs. If, on the other hand, transport costs are low, moving the DUs to a centralized location can increase the pooling gains enough to lower the overall TCO.

Migration paths to Open RAN: What are the options?

Mobile operators have increasingly committed to moving to Open RAN, but the timing and deployment options vary among them. Crucial to the success of implementing Open RAN is the choice of a migration path toward a fully virtualized and disaggregated architecture.

Both greenfield and brownfield operators may initially choose to have a more distributed Open RAN topology and later move gradually toward a more centralized one. In this case the DU is initially at the cell site, and eventually it may get relocated to a remote data center with the CUs.

A main reason to have DUs at the cell site is that this reduces the transport costs. Remote DUs require a high-capacity fronthaul (FH) link from each cell site. DUs located at the cell site need a lower-capacity midhaul (MH) link to connect to the CUs at a central location.

An operator with high FH costs or limited fiber availability may initially prefer to install the DUs at the cell site. When fiber becomes available or cheaper, the operator can move DUs to a centralized location and enjoy the pooling gains from virtualized and shared DU servers. Many European operators – Telefonica is one example – are deeply committed to Open RAN, and yet they are proceeding with caution in their migration to centralized DUs because a fiber connection may be unavailable or overly expensive.

At the opposite side, an operator that owns a fiber network that reaches the cell sites is likely to co-locate DUs and CUs in a remote data center from the beginning, because this is the most cost-effective path. This is the path chosen by Rakuten – the greenfield operator in Japan that leverages a $30/month fiber connection cost and has access to its own fiber network.

Many operators will find themselves in an in-between situation, and hence they may at different paces to a virtualized and centralized topology. Furthermore, the transition path and speed may change across different areas within a network, depending on the transport assets the operator owns or can lease in those areas.

Model scenarios

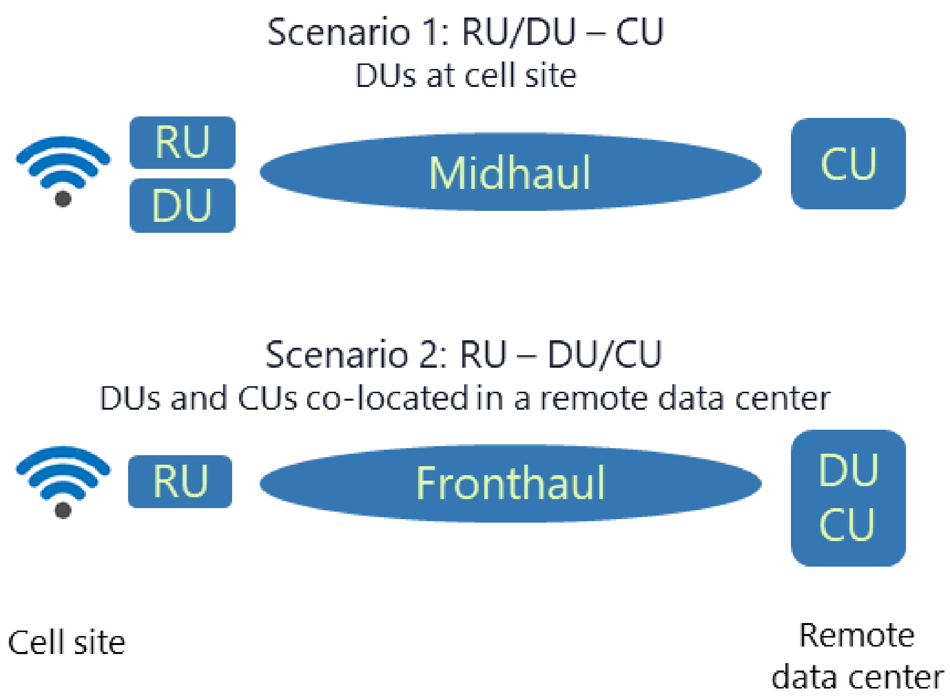

Operator cases

Operator 1
Low transport costs

- Low-cost fiber transport
- Fiber available in most cell sites
- $30/month per fiber link (up to 100 Gbps)

Operator 2
High transport costs

- Expensive fiber leases
- Limited fiber availability
- $1,000/month 1/Gbps lease (lower per-Gbps cost with higher capacity)

TCO model: Scenarios and assumptions

Scenarios Our model compares the TCO for two scenarios:

- Scenario 1 – Distributed topology: DUs are located at the cell sites with RUs, and MH connects DUs to the CU.
- Scenario 2 – Centralized topology: CU and DUs are in the same location, and FH connects RUs to the CU/DUs.

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

Cell sites We compared four cell profiles:

1- 5G-NR, 20 MHz channels, frequency-division duplex (FDD) with 4T4R multiple input, multiple output (MIMO).
2- 5G-NR, 100 MHz channels, time-division duplex (TDD) with 32T32R MIMO, 8 layers.
3- 5G-NR, 400 MHz channels, mmWave, 4 layers.
4- Combination of profile 1 and profile 3.

Each site has 3 cells (sectors) in the first three profiles, and 6 cells in the fourth.

Network 1,667 cell sites.

Transport Ethernet transport, with star packet links, using radio over Ethernet (RoE) and supporting the Enhanced Common Public Radio Interface (eCPRI) 7.2x O-RAN Open Fronthaul Interface over colored wavelength-division multiplexing (WDM).

Remote data centers DUs (scenario 2) and CU (scenarios 1 and 2) are in data centers where hardware resources are shared across RUs, resulting in higher efficiency because of pooling gains.

Inputs Cost, requirements, and traffic inputs are from trials and customers of Mavenir, Intel and HFR Networks.

Low transport costs: Move the DUs with the CUs to remote data centers

For an operator that has easy and low-cost access to fiber for FH over the 7.2 split, co-locating the DUs with the CUs (scenario 1) is the most cost-effective solution, delivering a cumulative TCO cost reduction ranging from 36% to 42% compared with the second scenario with DUs the cell site. The cost savings are highest for the cell with the lowest capacity (cell profile 1, 20 MHz), because in this case the pooling-gains difference between the two scenarios is the largest: the DU server at the cell site is underused and cannot be shared across cell sites. If the DU at the cell site supported more than the 3 cells that we assume here, the TCO for scenario 1 would look better, because DU resources would be shared more efficiently. However, cell profile 4 combines 6 cells, yet the TCO savings are similar to cell profile 1 with 3 cells. This is because this fourth profile includes 20 MHz and mmWave cells, and the TCO cost savings for mmWave are lower. As a result, the overall cost saving for profile 4 is slightly lower than that for cell profile 1. As capacity increases, the costs savings for moving the DU to a centralized location decrease (i.e., from 42%, to 37% and 36%), because with the higher throughput, the pooling gains in scenario 2 remain largely unchanged, but the pooling gains in scenario 1 increase, reducing the overall cost savings.

TCO savings for an operator with low transport costs by keeping the DUs at a remote data center

Low transport costs, cell profile 1 (20 MHz)

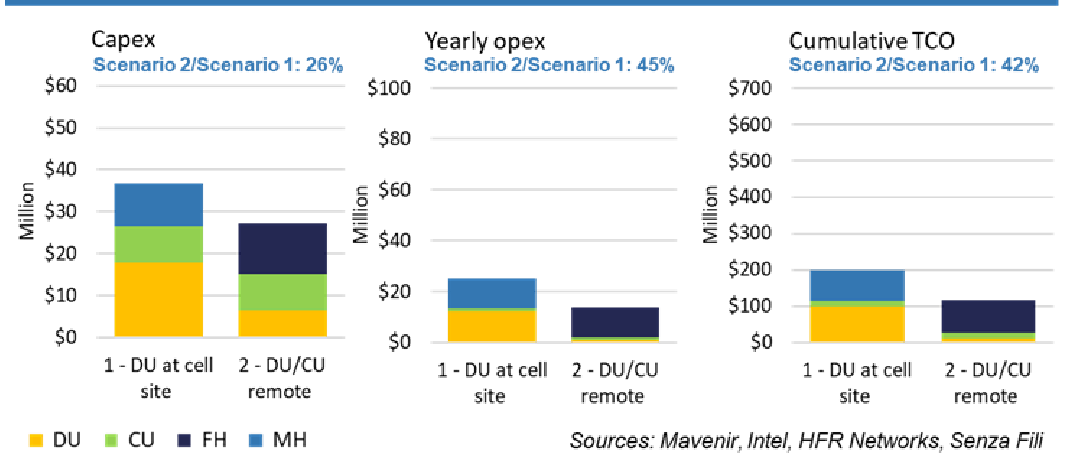

High transport costs, cell profile 2 (100 MHz)

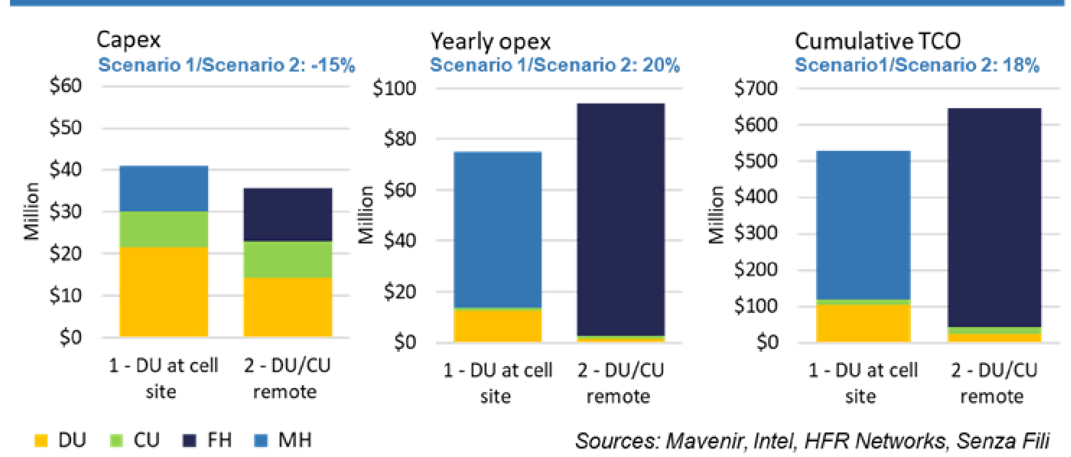

High transport costs, cell profile 3 (mmWave)

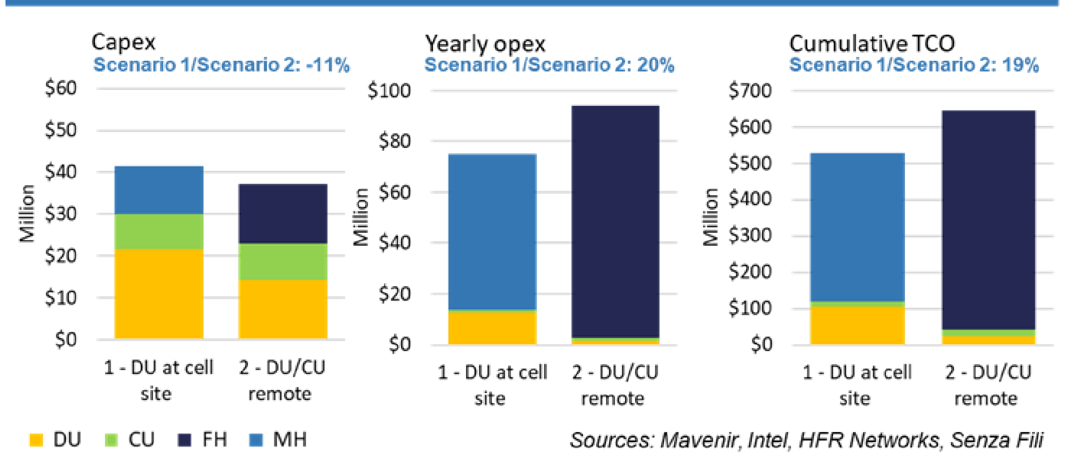

High transport costs, cell profile 4 (combination 1 +3)

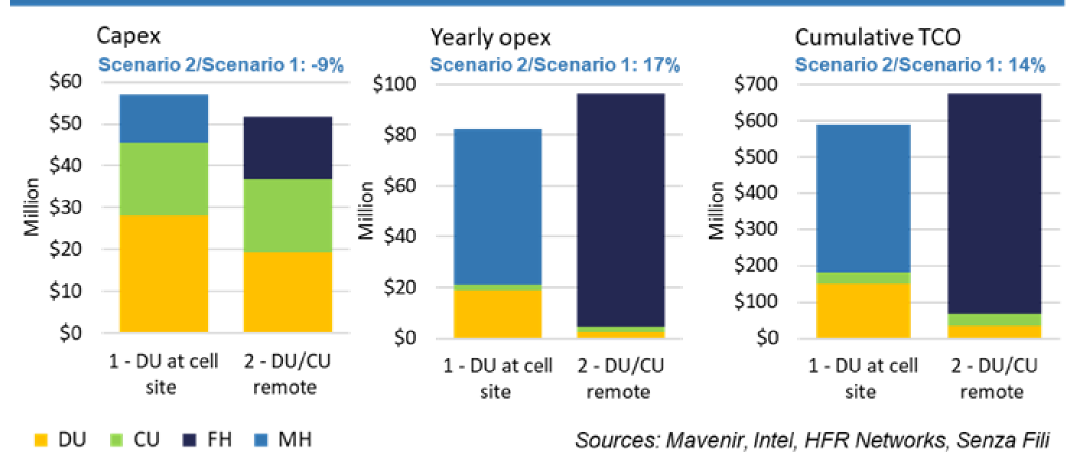

High transport costs: Keep the DUs at the cell site

Operators with high transport costs can get TCO savings ranging from 12% to 15% by deploying the DUs at the cell site. This lowers the transport costs and operators benefit from pooling gains at the cell site. In the high-transport case, the pooling gains have less impact, because monthly transport lease costs overshadow the CU and DU costs. The FH requirements in scenario 2 (DU and CU remote) add more in costs than the pooling gains save. As the cell capacity goes up, the costs gap between scenario 1 and scenario 2 decreases. This is because transport costs as a percentage of the overall TCO decrease, as a result of the lower per-Gbps cost of transport links of higher capacity.

Many of the initial Open RAN adopters are deploying DUs at the cell sites initially to benefit from these initial cost savings and other benefits, such as disaggregation, avoidance of vendor lock-in and wider equipment choice. In a later phase of the transition path to Open RAN, we expect them to move the DUs to a centralized location when transport prices go down or they have deployed its own fiber network.

TCO savings for an operator with high transport costs by keeping the DUs at the cell site

High transport costs, cell profile 1 (20 MHz)

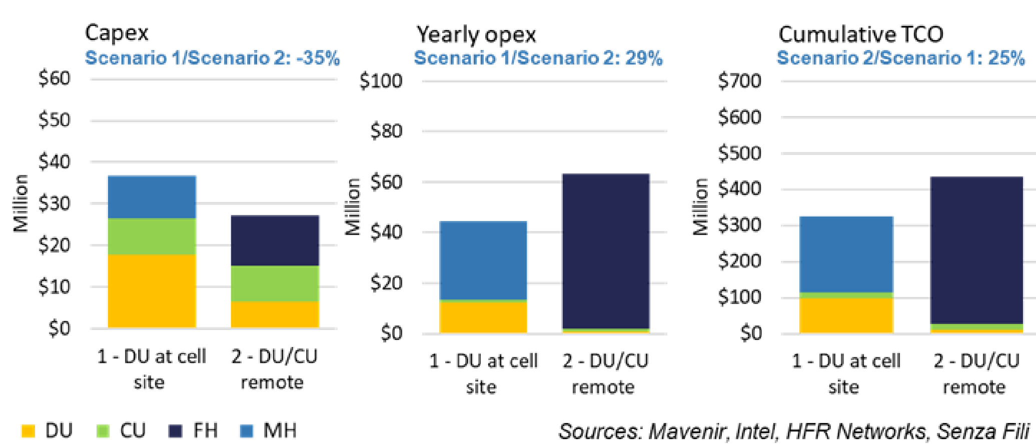

Sources: Mavenir, Intel, HFR Networks, Senza Fili

High transport costs, cell profile 2 (100 MHz)

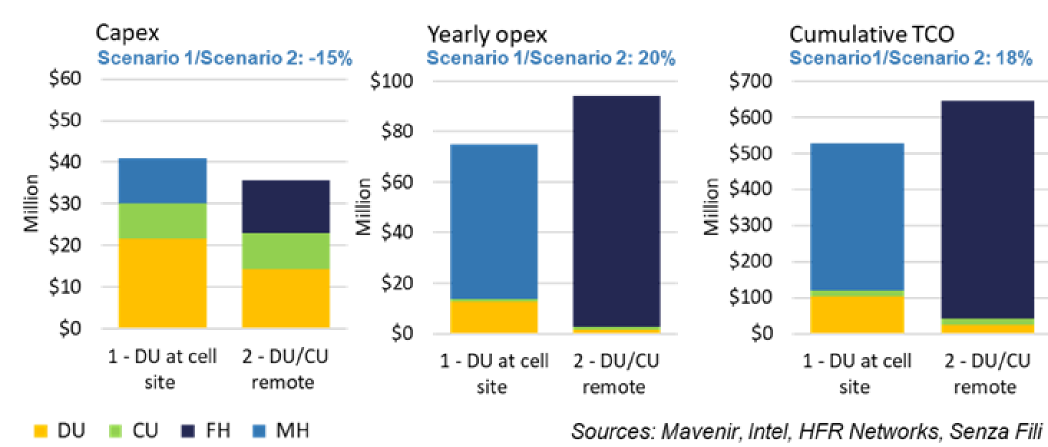

Sources: Mavenir, Intel, HFR Networks, Senza Fili

High transport costs, cell profile 3 (mmWave)

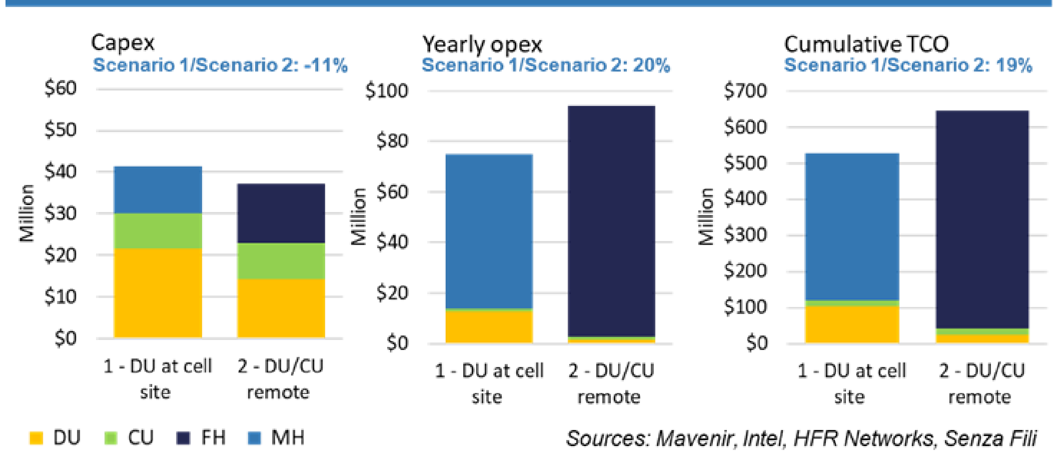

Sources: Mavenir, Intel, HFR Networks, Senza Fili

High transport costs, cell profile 4 (combination 1 +3)

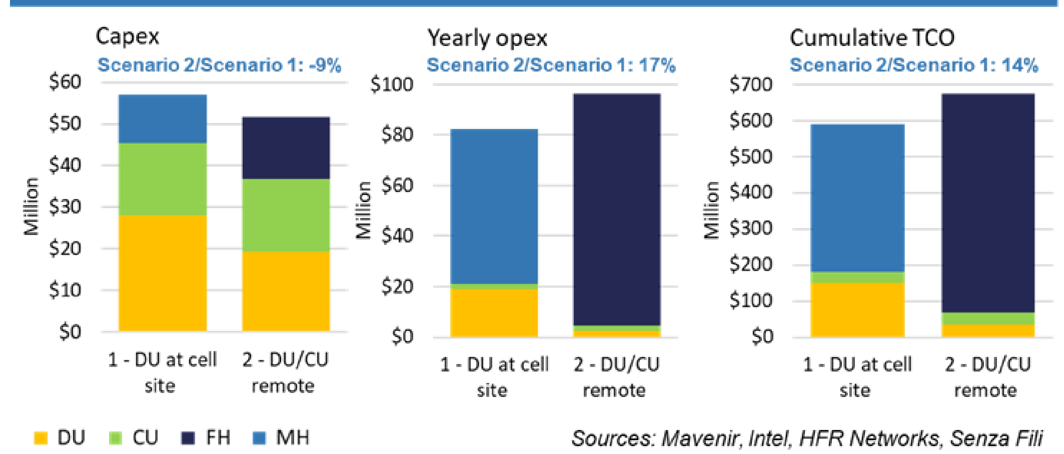

Sources: Mavenir, Intel, HFR Networks, Senza Fili

Per-bit comparison: Better TCO with more capacity

We have looked at the TCO for different scenarios and cases. But, as we noted, the total TCO changes significantly for different cell types and, more importantly, transport costs.

To compare the impact of these changes more directly, we looked at the per-bit TCO, which shows the TCO as a function of throughput. To simplify the comparison, we present the per-bit TCO for each case as a percentage of the highest per-bit TCO case, which is for scenario 2 with high transport costs, for cell profile 1 (20 MHz).

Across the four groups (two scenarios with high and low transport costs), the per-bit TCO is highest for cell profile 1, because the throughput is lowest. In these cases, the deployment costs are lower, but the capacity difference is higher than the cost difference. This high-level result would hold in traditional RAN deployments as well and simply shows that deploying cells with higher capacity improves the cost-effectiveness of the network.

The per-bit cost is consistently higher for operators with high transport costs. This is to be expected, because the opex contribution is higher due to the FH and MH lease costs.

The most interesting comparison is between scenario 1 and scenario 2, which shows the impact of pooling gains that depend on DU locations – at the cell site or in a remote data center. For the low-transport-cost operator, the per-bit TCO decreases for scenario 2, as pooling gains from the remote DU come into play. For the high-transport-costs operator, the per-bit TCO is lower for scenario 1, which is the most cost effective.

Per-bit TCO comparison

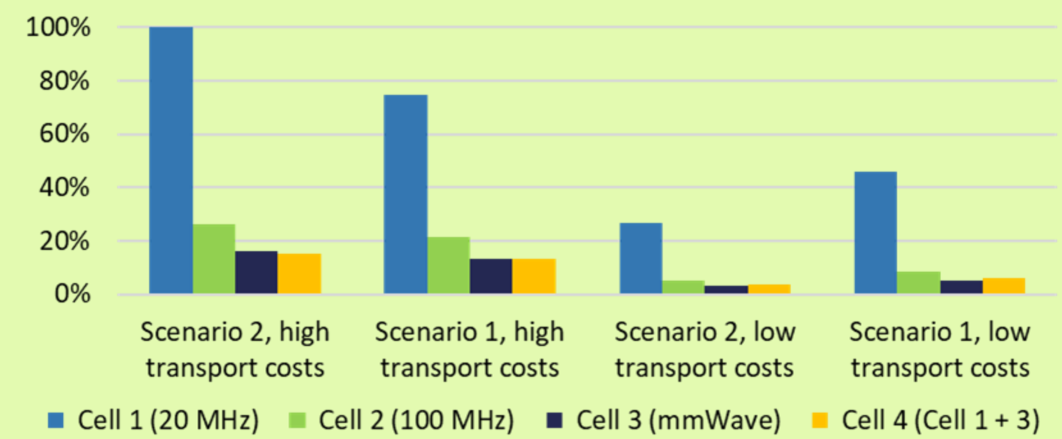

Scenario 1: DUs at cell site Scenario 2: Remote DUs, co-located with CUs

Sources: Mavenir, Intel, HFR Networks, Senza Fili

Takeaways

Operators with low transport costs can maximize Open RAN's pooling gains by locating the DUs in remote data centers with the CUs.

Operators with high transport costs are better off with the DU located at the cell site to lower the transport costs. In their case, minimizing transport costs is more effective than maximizing pooling gains.

The cost efficiency of Open RAN deployments is higher as the capacity of the cells increases.

About Mavenir

Mavenir is building the future of networks and pioneering advanced technology, focusing on the vision of a single, software-based automated network that runs on any cloud. As the industry's only end-to-end, cloud-native network software provider, Mavenir is transforming the way the world connects, accelerating software network transformation for 250.
Learn more at http://www.mavenir.com

About Intel

Intel is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.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).