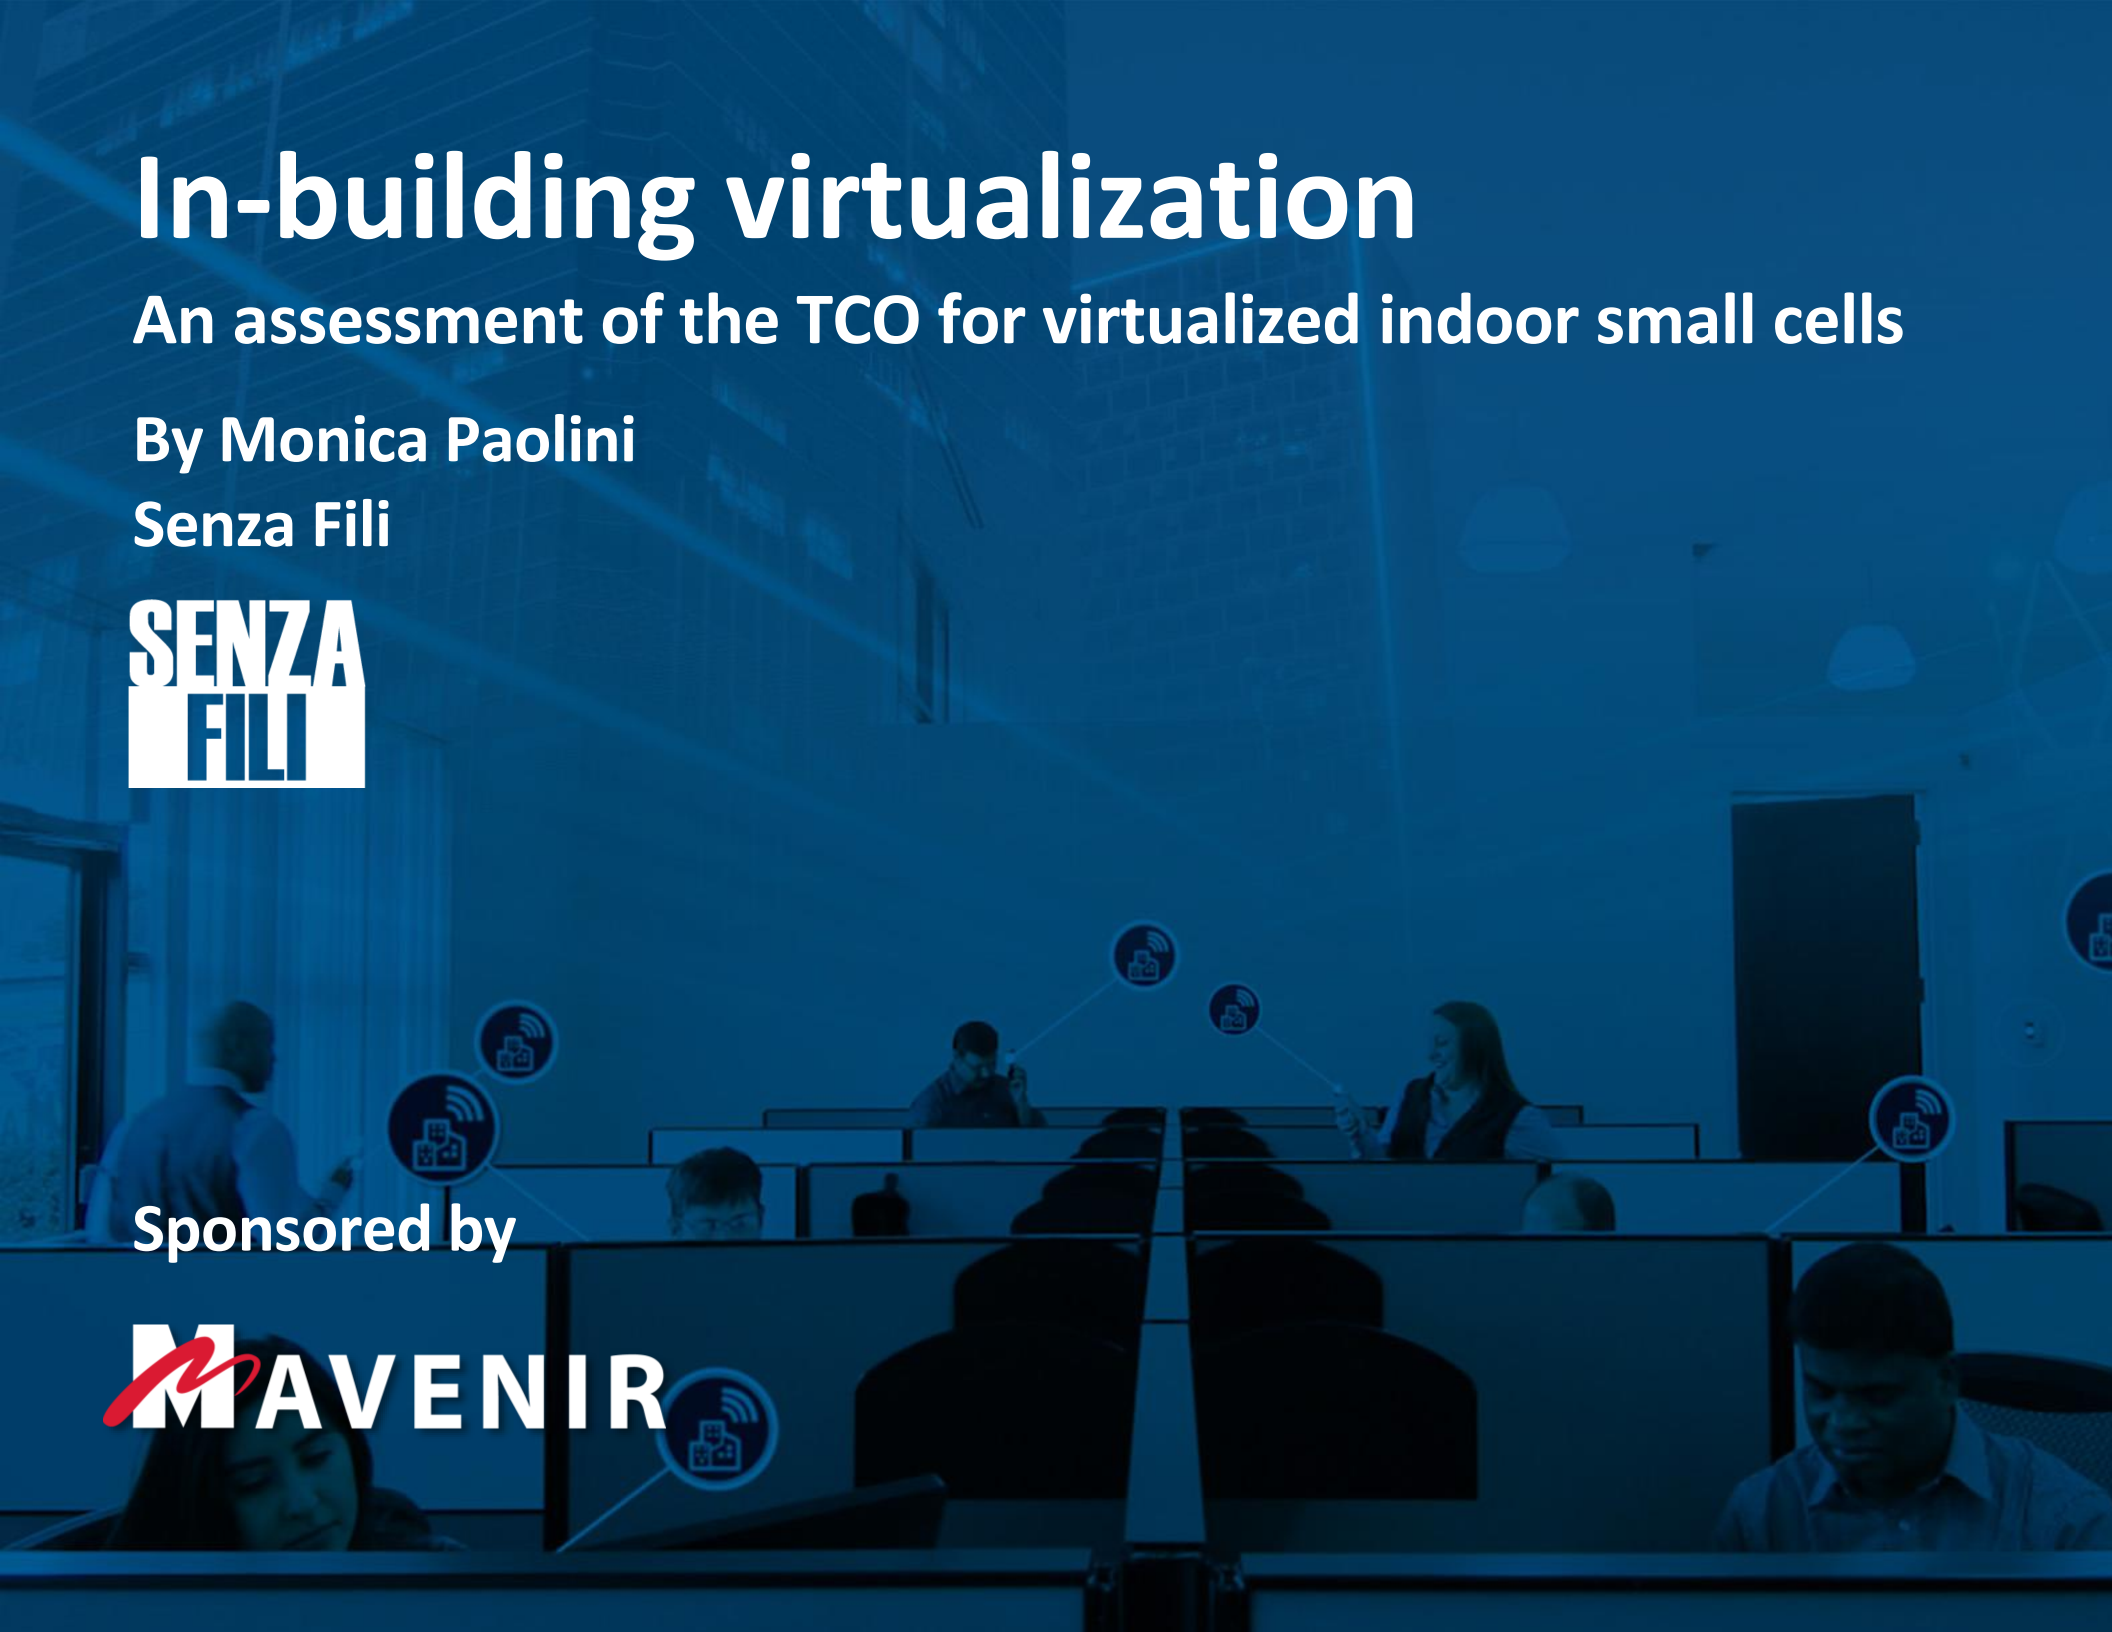

In-building virtualization

An assessment of the TCO for virtualized indoor small cells

By Monica Paolini
Senza Fili

SENZA
FILI

Sponsored by

MAVENIR

1. Indoor small cells to become deeply rooted in the venue and the enterprise

Small cells have entered the wireless ecosystem to enable the densification that is needed to improve coverage and capacity. Macro cells cannot provide the capillary indoor coverage we expect throughout buildings, because their signal does not penetrate buildings well. Macro cells also struggle to keep up with the concentration of traffic inside buildings, which is where most wireless usage takes place. Small cells are deployed closer to the users and provide coverage and capacity where needed.

Yet, despite small cells having demonstrated their ability to provide the coverage and capacity needed for in-building wireless (IBW) environments, their deployment has so far been below expectations. They are mostly concentrated in venues and enterprises with stringent requirements that are difficult to meet in a cost-effective way without densification.

The main cause of the slow adoption of small cells is that operators have found it difficult, expensive and effort intensive to roll out small-cell deployments in venues and enterprises. They have preferred to enhance the macro infrastructure instead of investing in small cells. To date, indoor small cell deployments have followed a model in which a mobile operator pays for, operates and controls its own access infrastructure. This is a model that works for macro cells, but not for small cells.

The business model is rapidly evolving in a direction that makes small cells more attractive to both operators and their physical hosts – venue owners and enterprises. The change is driven by a more active involvement of venue owners and enterprises, which are increasingly willing to pay for the small-cell infrastructure and work with neutral hosts to open access to their premises to multiple operators. In exchange, they expect to exert more control over the infrastructure they want, have visibility into the operations, and be able to use the network for their internal services, including Internet of Things (IoT).

Indoor small cells and the edge:
The synergy of radio access network (RAN) virtualization and edge computing

Improved in-building coverage and capacity: spectrum management and reuse can be more efficient with virtualized small cells.

Lower latency: moving processing and content to the edge improves latency for time-sensitive traffic.

Local breakout: enterprises and venues can run their applications and store their content locally and can use their small-cell infrastructure to deliver their own services in addition to operator and over-the-top (OTT) services.

Enterprise/venue-funded: the enterprise or venue can time, size and plan the IBW network to meet its own requirements and funding availability. It can also control network operations and usage, either directly or, more likely, through a neutral host or other third party.

More control for the enterprise, less effort for carriers: the enterprise/venue pays for the infrastructure it needs, while the carrier provides the connectivity and integration to the wide-area network.

Stronger security: the local edge-computing infrastructure decreases the vulnerability to network-wide attacks.

Integration with enterprise and venue Wi-Fi: Wi-Fi will continue to provide most connectivity through a strong and cost-effective ecosystem, but Wi-Fi and small cells can both improve performance when integrated.

Support for IoT and Industrial IoT (IIoT): indoor small cells can support enterprise or venue-based IoT or IIoT applications in private networks that are not part of a mobile network.

IBW has become crucial for many entities that control real estate: some need the functionality it provides to run location-based services and basic connectivity, others see it as an amenity that increases the value of their real estate – and many of them want both.

In parallel, new small-cell solutions and edge-computing capabilities enable and strengthen the business model evolution, by facilitating the involvement of the enterprise and venue owners.

Small-cell equipment is increasingly designed to support sharing between the venue owner/enterprise and multiple operators – often with the mediation of a neutral host, which may plan, deploy and operate a network and establish relations among the involved parties. In-building deployments of this new generation of small cells make it possible to support multiple frequencies (including unlicensed access), combine multiple cells or sectors in the same enclosure, share backhaul, have centralized architectures, and use Ethernet for backhaul. Crucially, in this deployment model, operators can participate in a multi-operator IBW network, but retain exclusive use of their spectrum assets, as they currently do with distributed antenna systems (DAS).

Private LTE networks that venue owners and enterprises can deploy independently of mobile operators have started to gain great traction in the enterprise today. With 5G, private networks adoption will grow, and they will support additional functionality. Private networks can support a wide range of services that are local to the enterprise and that may have specific requirements – for instance, for reliability, latency, or security – but they can be operated with stripped-down, less complex core functionality.

In the US, the introduction of spectrum sharing in the 3.5 GHz Citizens Broadband Radio Service (CBRS) band using OnGo creates a major incentive for enterprises to deploy private networks, with up to 150 MHz of clean spectrum without having to purchase a license.

At an end-to-end network level, virtualization in the RAN and in the core provides the flexibility that operators and the enterprise/venue owners need to establish a deeper collaboration framework that is mutually advantageous and that delivers improved network and financial efficiency.

With virtualization in the RAN with Cloud RAN, small cells can use a split architecture, in which the remote radio unit (RRU) sits at locations of high traffic inside buildings, and a virtual BBU (vBBU) is located within a remote location, either physically within the venue or in the cloud. The first option may appeal to venues that are larger or have demanding requirements. The second option may be more attractive to mid-size and small venues that would find it too complex and expensive to host dedicated BBUs.

Core virtualization is a great enabler for moving functionality to the edge. For the enterprise and venue owners, this greatly expands the ability to deploy and control services and functionality that is local to them, and over which they have full access and control, if they need it.

Edge computing and RAN virtualization are complementary and mutually reinforcing, because the vBBUs can be co-located with the edge infrastructure. Not only does this help reduce the costs, but it also improves the performance benefits of both as traffic management and resource optimization can be jointly performed at the edge/vBBU location.

This paper is a companion to two previous papers. The first white paper, “How much can operators save with a Cloud RAN? A total cost of ownership (TCO) model for virtualized and distributed RAN,” compares the TCO for distributed RAN (DRAN) and Cloud RAN. The second white paper, “Future proofing mobile networks economics. Accessing the TCO for Cloud RAN and Centralized RAN,” looks at two remote-BBU topologies – Centralized RAN and Cloud RAN – using the same TCO model we use here, and we then compare them to the DRAN TCO.

While we built the TCO model within the context of these new trends in technology, business models and new solutions, our analysis has a narrow focus to look at the impact of RAN virtualization in isolation. It does not include the financial benefits that concurrent changes, such as the introduction of new hardware solutions, CBRS or edge computing, may bring – which expand the cost savings and financial benefits of Cloud RAN.

RAN architectures

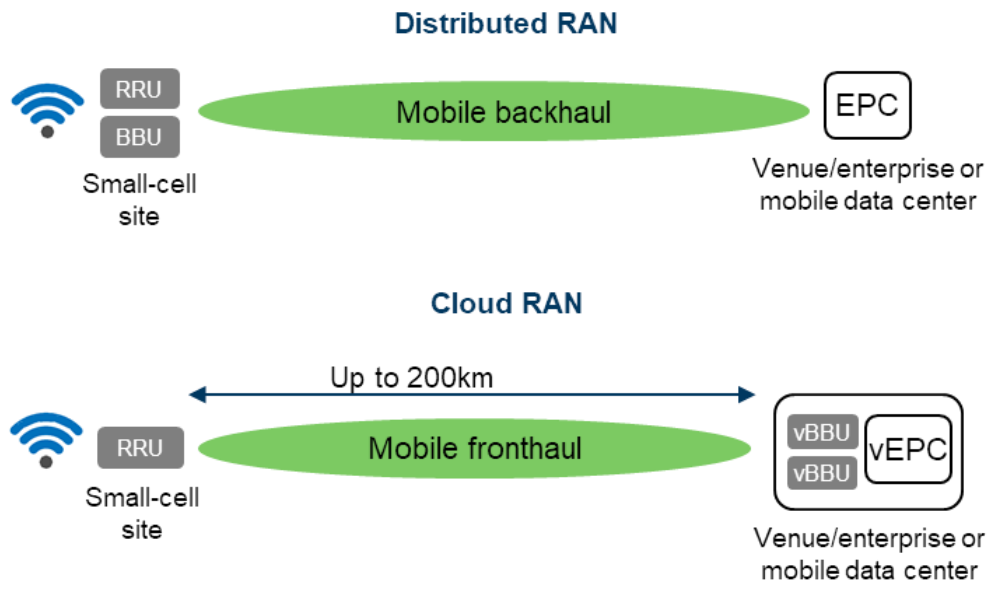

Source: Mavenir

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 Cloud RAN. The RRU remains at the cell site, but the vBBU moves to a central location, where vBBU processing can be pooled for multiple RRUs.

In a Cloud RAN, the vBBUs can be co-located with the virtual Evolved Packet Core (vEPC), and operators need only mobile fronthaul to connect the RAN to the Evolved Packet Core (EPC). In a DRAN, the BBU is at the cell site, so a backhaul link connects to the EPC from the RAN.

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.

Cloud RAN requires a high-reliability and low-latency fronthaul (FH) link between RRU and vBBU. High costs for Common Public Radio Interface (CPRI) – the default FH interface today – have so far limited the adoption of Centralized RAN, but functional splits in the FH introduce a sharp reduction in FH costs in the Cloud RAN case, making FH costs comparable to backhaul (BH) costs in the DRAN.

In this paper we use the term “small cell” to indicate a RAN element that has less power and coverage than a macro cell, and typically has a single sector. As with macro cells, small cells can be deployed in distributed or centralized architectures (i.e., DRAN and Cloud RAN in our model). More specifically, here we only consider indoor small cells.

2. TCO model assumptions

We built a TCO model to look at the financial benefits from Cloud RAN compared to DRAN over a period of 5 years. In the first white paper based on this TCO model, we compared the TCO for a Cloud RAN network and a DRAN network, both with the same type and number of RRUs. In the second paper, we shifted the focus to comparing Centralized RAN and Cloud RAN, to assess different RAN virtualization solutions. Here, we narrow the scope of the analysis to specifically compare the TCO for Cloud RAN to that for DRAN, in IBW small-cell deployments. The model looks at a network that covers a single venue or enterprise. It makes no assumption about the business model used (e.g., whether the deployment is driven by the operator or by the enterprise/venue owner), so that we can concentrate on the changes introduced by the Cloud RAN architecture. We discuss later the financial impact of the business model and other deployment options.

The model covers a set of small cells that share a vBBU pool in a high-density area, which could be an enterprise campus, an industrial location, a warehouse, an educational institution, a hospital, a multi-unit residential venue, or a retail center. We used 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 and enterprise customers.

The RRU and BBU equipment is the same (2x2 MIMO indoor small cells) in the Centralized RAN and Cloud RAN cases. 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. The option 7 split allows the operator to use Ethernet-based FH or other FH solutions that are cheaper than CPRI.

TCO model assumptions

Framework. Our model compares the TCO, over five years, of a Centralized RAN versus a Cloud RAN greenfield IBW small-cell network with vBBUs. All capex is in year 1, during deployment. The model covers the RAN all the way to the EPC.

Network. 250 indoor single-sector LTE 2x2 multiple input, multiple output (MIMO) small cells.

FH/BH. DRAN uses backhaul (BH). 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. vBBU resources can be dynamically allocated to RRUs with multiplexing. We estimate that, when used, multiplexing reduces the BBU capacity requirements by 50%. Multiplexing is not an option for DRAN.

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 a remote site.

Leasing. We did not assume any leasing costs for the infrastructure, as we assumed that the enterprise/venue owner actively works with the mobile operator in the deployment.

3. Financial benefits from small-cell Cloud RAN

Our TCO model shows that in a greenfield IBW small-cell deployment, an operator can save 48% over five years when choosing a Cloud RAN architecture instead of a DRAN architecture. The cost savings reflect a 66% capex reduction and a 25% opex reduction.

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.

In the Cloud RAN scenario, the operator uses both less-expensive non-proprietary hardware, and vBBU 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 small-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 RRUs at the edge.

The model indicates that with a Cloud RAN deployment, mobile operators can save on opex through a reduction in maintenance, power and operations costs by aggregating vBBUs in more-centralized locations, which are typically easier to access and cheaper to operate. 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.

The Cloud RAN cost savings for indoor small cells are higher than those for macro cells, which we presented in the initial white paper, entitled "How much can operators save with a Cloud RAN? A TCO model for virtualized and distributed RAN." In the macro scenario, cumulative TCO savings are 37% (versus 48% here, for indoor small cells), derived from a 49% decrease in capex (versus 66%) and a 31% decrease in opex (versus 25%). Capex cost savings come primarily from the lower cost of indoor equipment and installation that Cloud RAN brings. Opex savings are higher in the macro case, because the macro scenario includes leasing costs that we exclude here.

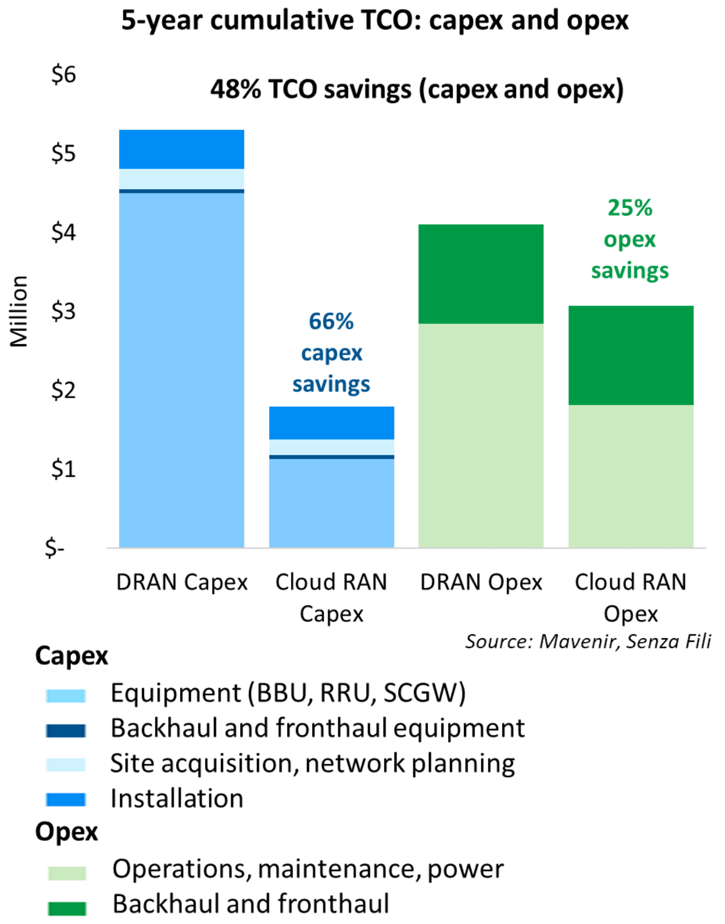

Source: Mavenir, Senza Fili

4. Implications

When deploying indoor small cells, our TCO model shows that operators, venue owners and enterprises could all benefit from a 48% decrease in combined capex and opex over five years. These cost savings come from the cost efficiencies that a virtualized and centralized RAN architecture brings: lower equipment and installation costs, and lower costs to run the network. With a functional split, the cost of fronthaul needed to connect the RRU to the BBU is comparable to the cost of backhaul in the DRAN case. This is a key improvement over centralized RAN models that use CPRI, which increases FH costs and thus erases some of the cost savings over DRAN.

Other factors that we discussed at the beginning of this paper and that were beyond the scope of the TCO can bring additional savings, and they strengthen the case for Cloud RAN. The overall cost of deploying small cells decreases even further when a combination of enterprise/venue owners fund an IBW small-cell network incorporating equipment that facilitates network sharing and they open it to multiple operators. All parties benefit from the cost savings from Cloud RAN and network sharing, and the two sources of savings complement each other, because infrastructure sharing is easier to implement in a Cloud RAN environment.

Similarly, edge computing provides benefits for wireless networks and, specifically, for indoor small cells that do not depend on the adoption of Cloud RAN. But as we noted at the beginning, there are strong synergies between edge computing and Cloud RAN that result in improved performance and network utilization, and better support for services and applications. This translates into further cost savings that depend on the integration of the two. For instance, co-locating edge-computing functionality and BBUs will reduce the opex for both, as it enables network operators to consolidate the physical location of their network elements.

Virtualization provides the foundation and trigger for a chain of changes in the network architecture that enable operators to increase the use of network resources and their cost efficiency. Cloud RAN is a critical part of this evolution to increase network flexibility and optimization capabilities that deliver the short-term cost savings we have presented in this paper, as well as longer-term ones that have a wider impact on wireless networks.

Download the companion white papers

[How much can operators save with a Cloud RAN? A total cost of ownership (TCO) model for virtualized and distributed RAN]

[Future Proofing Mobile Network Economics]

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.

© 2018 Senza Fili Consulting LLC. All rights reserved. This white paper was prepared on behalf of Mavenir. The views and statements expressed in the white paper are those of Senza Fili, and they should not be inferred to reflect the position of Mavenir. 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 relies 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. Trademarks mentioned in this document are the property of their respective owners. Cover page photo by chuyuss/Shutterstock.