XONA PARTNERS

RAN Virtualization:
Unleashing Opportunities for Market Disruption

Frank Rayal

June 2016

San Francisco • Singapore • Dubai • Paris

XONA PARTNERS

Table of Contents

List of Figures

List of Tables

Overview

Radio Access Network (RAN) virtualization is a highly disruptive technology that will radically impact how wireless services are delivered. It will change the current ecosystem and market structure; altering the way MNOs plan and roll out new services by providing a scalable, on demand alternative to the traditional architecture. Dedicated, on-site hardware to power the RAN is becoming expensive to build-out and maintain especially as more cell sites are required to keep up with capacity demand. Virtual Radio Access Networks (vRAN) moves the baseband modules away from the radio at the cell site to a data center. This enables intelligent scaling of computing resources as demand on capacity fluctuates, while reducing site lease costs, energy usage, and maintenance expenses. The evolution of LTE and advent of 5G networks increases bandwidth requirements further. This makes increased fronthaul requirements and the inflexibility of the legacy CPRI serial interface the primary challenges to vRAN deployments. Resolving the fronthaul challenge enables the Internet giants and fixed access service providers to enter the wireless market with lower cost basis, a move that is highly disruptive in a market dominated by telecom incumbents entrenched through massive equipment install-base.

The Genesis

Mobile network operators (MNOs) in Japan and Korea were first to centralize the radio access network by moving base stations baseband units to fiber centers, leaving only the remote radios and antennas at the cell site. This network architecture is possible provided fiber is available to link the baseband units to the remote radio – a link called fronthaul. Operational cost savings from this architecture range between 30 – 40% due to lower site lease, simplified support and maintenance, as well as lowered energy expenses. Operators without their own fiber assets would find it cost prohibitive to implement this architecture because of the high fronthaul performance requirements of legacy protocols used to connect the baseband to the radio (e.g. CPRI). Improvements to this link will make fronthaul feasible to service providers without their own fiber assets.

At the turn of the decade, LTE deployments were burgeoning and data traffic was doubling year over year. Unfortunately for MNOs, the average revenue per user (ARPU) did not increase, falling in many markets and leading to lower EBITDA margins. Some of MNOs, such as China Mobile, saw virtualized RAN as an opportunity to lower costs and improve financial performance. Together with other Asian operators, China Mobile promoted the concept of Cloud RAN, which virtualizes the centralized baseband processing to achieve further cost savings. The term Cloud RAN has since become a buzzword, and many vendors with different solutions began using the term liberally, a few with little relationship to actual Cloud RAN. We will use the term vRAN to denote a fully centralized and virtualized baseband implementation (Figure 1).

On top of cost savings, vRAN also brings performance benefits. This is owing to features such as coordinated multipoint and network MIMO, which become possible due to centralization, and are utilized to lower interference and improve throughput. The result is enhanced user experience, especially at the cell edge where performance is most lacking (up to 100% throughput gain at the cell edge has been demonstrated in field trials). In fact, centralization becomes more important in heterogeneous networks (HetNets) where low-power small cells are deployed in the service area of high-power macrocells. Centralization reverses the LTE distributed architecture which places the entire protocol stack at the base station leading to high overhead and timing requirements for coordination among base stations to mitigate interference. Future network architectures planned for 5G intend on implementing a flexible architecture, where part of the intelligence is centralized to reduce the coordination overhead.

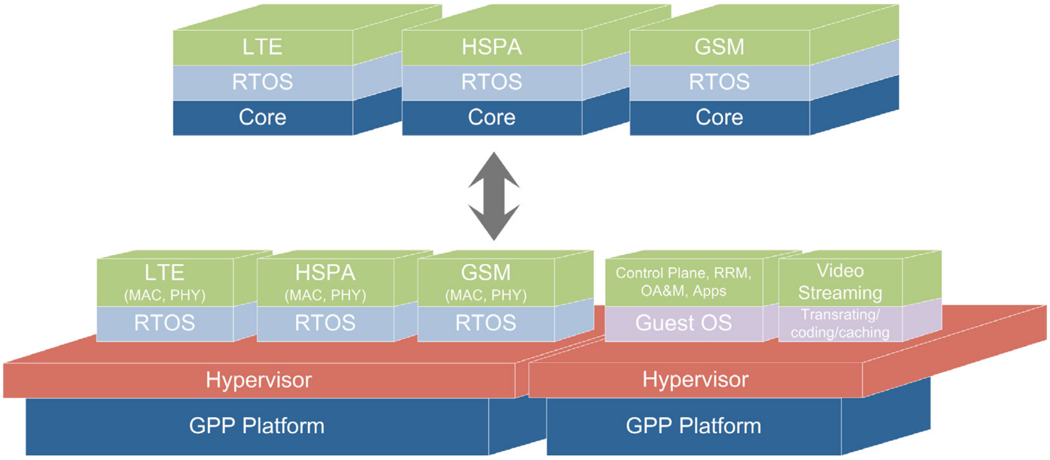

Figure 1 Virtual RAN: baseband virtualization.

The gain associated with virtualization is based on leveraging the cost structure and economies of scale of the IT/data center industry. Furthermore, the scalable and elastic properties of virtualization allow deploying processing power to provide capacity on demand when and where it is required in sharp contrast to distributed hardware architecture that is designed for peak capacity.

A Disruptive Idea

Virtualization decouples the software from hardware, enabling the use of commercial servers in the network. This profoundly alters the way MNOs plan, design, procure and roll out new services. They would no longer need to purchase hardware-optimized base stations from specific telecom equipment manufacturers (TEMs). Instead they would only need software and general purpose servers in data centers to run the wireless protocol stack as an application to power any remote radios on demand. Other applications can run on the same infrastructure to provide value added services, such as video optimization, caching and localization. TEMs could provide their applications in a software as a service (SaaS) setting, with an OPEX-based pricing model, instead of the CAPEX-dominant model of today. MNOs could control and manage large networks more efficiently to enable a HetNet architecture. Because wireless capacity is not in demand at peak level at all locations at the same time, MNOs could save substantial expenses by multiplexing wireless capacity to increase operational efficiency and reduce capital costs. The RAN market structure will be radically changed, altering the balance of power between vendors and operators; leading new entrants into a market that's becoming highly consolidated. Such is the disruptive nature of virtualization in the RAN.

The Challenges

The major challenge to implementing vRAN is the fronthaul interface between the baseband units and the remote radio. CPRI is the most common interface, which was designed in 2002 before the centralized architecture was advanced. It requires 10x the capacity of an LTE backhaul channel, which makes it prohibitively expensive for operators who don't own fiber assets. Unlike backhaul, CPRI fronthaul cannot be statistically multiplexed so its capacity requirements increase proportionally with the number of LTE carriers used. CPRI also has tight requirements for synchronization, latency and jitter that are difficult to meet when there is no direct connectivity between baseband and radio. As a result of these factors, fiber becomes the only media capable to implement fronthaul. While this is possible, especially as the cost and transmission capabilities of optical transceivers have been on a steep improvement curve, it remains a challenge to many operators who don't own fiber or where fiber penetration is thin.

Table 1 Backhaul and fronthaul requirements for a 20-MHz 2x2 MIMO LTE carrier.

# of Carriers	Backhaul (Mbps)	Fronthaul (Mbps)
1	236	2,547
3	248	7,641
6	496	15,282
9	744	22,923
12	992	30,564

A second challenge pertains to virtualization. The wireless protocol stack includes computationally intensive functions that are inefficient to run on general purpose processors (GPPs). Devices such as FPGAs, ASICs and SoCs are more efficient, and provide real-time response capability, which is required by some RAN functions. Such challenges are beginning to dissipate as new, more powerful, GPPs with vector acceleration functions are becoming available on the market. Additionally, there are different implementations of virtualization that can solve these challenges such as offloading complex functions to acceleration engines. It is now clear that challenges due to virtualization could be overcome as demonstrated in recent PoCs, where performance was near that of hardware-based implementations.

The Solutions

The solution to the fronthaul challenge takes different paths depending on the objective. If the goal is to ensure compatibility with installed base of remote radios, CPRI compression techniques may be used. These typically achieve between 50% – 66% savings in bandwidth. Alternatively, the protocol stack can be divided, with some functions virtualized at the center and others performed at the cell site. The functional split of the protocol stack trades off potential performance enhancement against fronthaul latency and capacity requirements (Figure 2, Table 2).

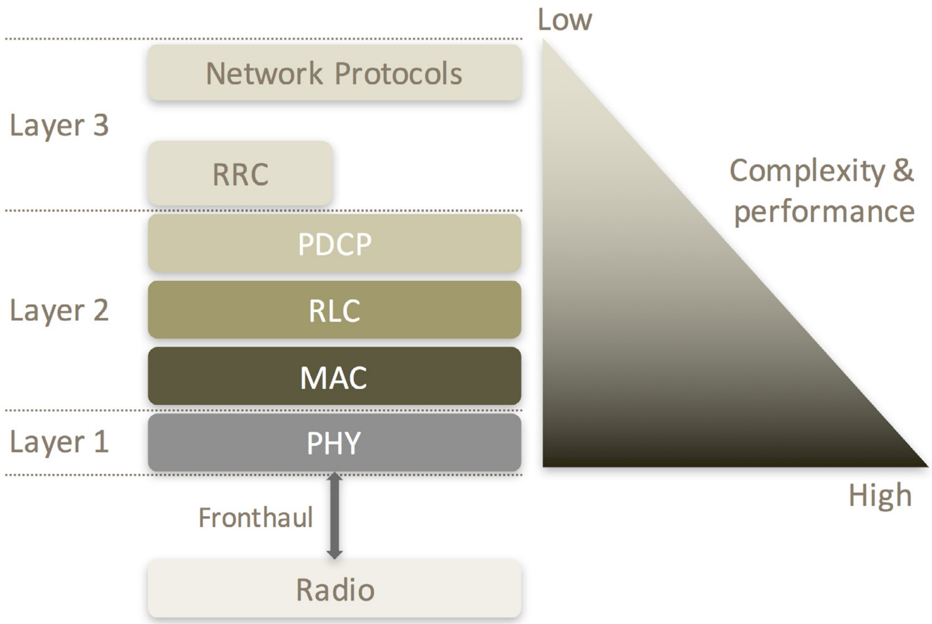

Figure 2 Functional partitioning of the LTE protocol stack.

While such approaches cater to accommodating legacy networks, it is possible to design new interfaces optimized to meet the requirements of future networks (high scalability, low cost). Such interfaces bring about the full benefits of RAN virtualization and revolutionize the wireless infrastructure market. While the technology has been demonstrated, achieving consensus in the industry is more challenging as incumbents work to protect their market share and position. Several industry forums have initiated studies to engineer a new interface – these efforts are still at a relatively early stage.

Table 2 Overview of functional split characteristics.

	High Functional Split	Low Functional Split
Fronthaul requirements	1 – 2x the capacity requirements of backhaul	Same as CPRI requirements, if CPRI is used
Performance enhancements	Limited in comparison to low functional split but better than a fully distributed architecture	Maximum performance enhancements though CoMP and network MIMO techniques
Cost of implementation	Low cost in comparison to distributed architecture	High cost if CPRI fronthaul is used
Compatibility with installed-base	High compatibility with current install-base of equipment: could be implemented with additional network elements	Limited compatibility with current install-base of equipment
Disruptive potential	None – similar fundamental building blocks to the current distributed architecture	Disruptive potential requires an efficient packet-based interface. Low disruptive potential with CPRI

Categorization of Architectures

In an effort to improve performance of the distributed LTE architecture in HetNets to meet future capacity demand, equipment vendors are beginning to centralize parts of the protocol stack. Virtualization is implemented in some centralized designs, but not all. This has led to a bifurcation of architectures that diluted the term Cloud RAN. From its original definition of fully centralized and virtualized air interface protocol stack, Cloud RAN is now even used to refer to solutions that include neither centralization nor virtualization. We introduce the following definitions while recognizing that different implementations exist within each category (Figure 3, Table 3):

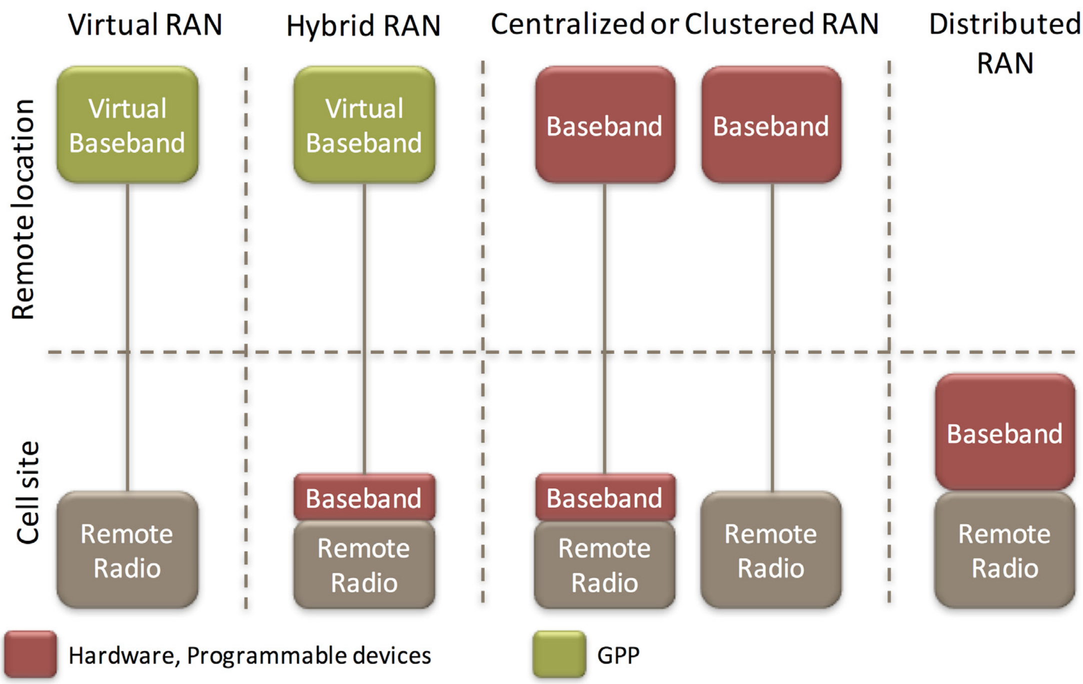

Figure 3 RAN architecture definitions.

Virtual RAN: An architecture where general purpose processors and servers are used to run air interface protocol stack in a central location (Figure 4). Various architectures and implementations of vRAN exist:

a. Architecture where all layers of the air interface protocol stack run on GPPs located in a central location.

b. Architecture where non-real-time functions in Layer 2 and Layer 1 run on GPPs while real-time functions run on hardware accelerators.

Some implementations run the protocol stack on a processor without capabilities for pooling and load-sharing of resources (i.e. bare metal).

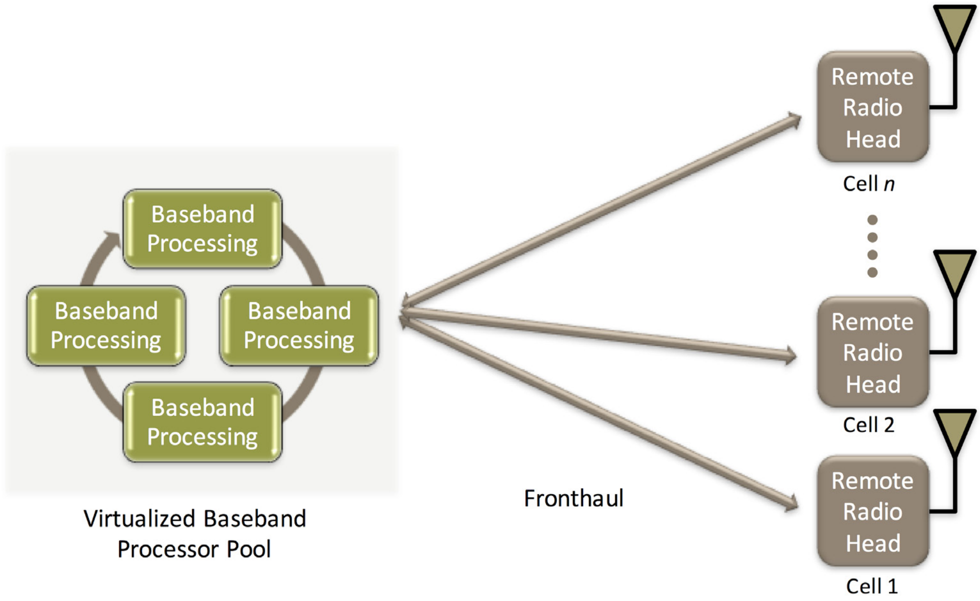

Figure 4 Simplified vRAN architecture.

Hybrid RAN: A split baseband architecture where some modem functions run on GPPs in the center while other baseband functions, such as Layer 1 or parts of Layer 2, run on programmable and hardware devices, such as FPGAs, DSPs, NPUs ASICs and SoCs, at the remote radio. The split can occur at different locations and is a vendor specific design. Hybrid RAN is an architecture that optimizes cost and performance but does not have the same disruptive potential as vRAN.

Clustered RAN: An architecture where baseband modules are located in a central location as is done in today's base station hotels. The air interface protocol stack runs on programmable and hardware devices. This is the most basic form of centralization, and is targeted for OPEX reduction in certain Asian markets. It is also used for practical considerations in other parts of the world where it is not possible to collocate the baseband with the remote radio due to different considerations such as space and access. Clustered RAN is the name given by SK Telecom to Phase 1 of their roadmap to implements vRAN.

Centralized RAN: An architecture where the baseband modules are located in a central location, similar to Clustered RAN, but with two variations:

a. All the baseband functions of the air interface protocol stack are centralized (full centralization). In this case, the difference from Clustered RAN lies in the integration of baseband processing to save cost among different modems and to improve performance through coordination of resources.

b. Part of the upper layers of the protocol stack are centralized while the lower layers are distributed at the remote radio (partial centralization) – essentially a split architecture without virtualized baseband.

In either case, the baseband processing is based on programmable devices running all air interface modem functions. The architecture supports a 1:1 relationship between a radio and its baseband modem. GPPs may be used to run Layer 3 functions in addition to different applications.

Table 3 RAN architecture definitions.

<table>
<tr><th></th><th></th><th colspan="3">Architecture: Baseband Centralization</th></tr>
<tr><th></th><th></th><th>Centralized</th><th>Split</th><th>Distributed</th></tr>
<tr><td rowspan="2">Technology: Baseband Virtualization</td><td>Virtualized</td><td>Virtual RAN
• Pioneered by startups
• High potential for market disruption
• Likely lead deployments in local-area coverage use cases (venues)</td><td>Hybrid RAN
• Supported by major vendors in wide-area deployments with a functional split high in the protocol stack</td><td rowspan="2">Distributed RAN
Architecture used in 490+ commercial LTE networks.</td></tr>
<tr><td>Not Virtualized</td><td colspan="2">Clustered RAN or Centralized RAN
• Deployed on wide-scale by leading carriers in Korea and Japan for network OPEX savings
• Deployed in select installation by operators worldwide for different reason: site acquisition challenges, zoning, security, power availability, theft prevention, etc.</td></tr>
</table>

Other terms are used in the industry to denote a level of coordination among base stations for interference management such as Cooperative, Collaborative and Elastic RAN (Ericsson) where the baseband processing is not necessarily virtualized. They can be classified according to one of the above categories.

Market Trends

Vendors' Strategies

Major equipment vendors are focusing on Hybrid RAN architectures that centralize and virtualize the upper layers of the protocol stack, typically the PDCP layer as it is a straight forward migration that utilizes existing infrastructure (Figure 5). This functional split allows the implementation of dual connectivity small cells, which improves mobility management in HetNet deployments.

Startup pioneers are leading in vRAN implementation, where different designs have emerged that promise to reshape the market landscape. vRAN lends itself to new ways of deploying small cells and distributed antenna systems (DAS).

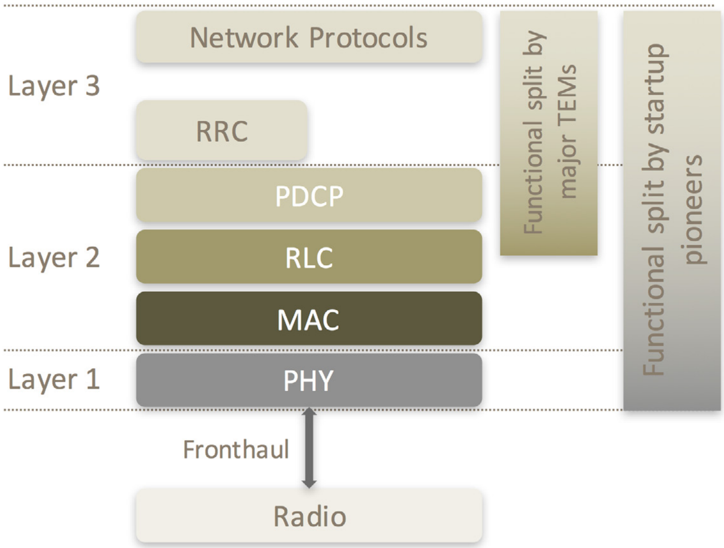

Figure 5 Functional split trends for LTE.

Telecom and Internet Ecosystem Convergence

Behind the vRAN pioneers stand major Internet players such as Facebook, who initiated the Telecom Infrastructure Project (TIP) to explore the benefits of vRANs and its potential to reduce the cost of connectivity. TIP participants joined the Open Compute Platform (OCP) which is a 5-year old initiative on data center technologies for telecom companies. This points to the confluence of the Internet/compute world with the telecom world which has significant ramifications.

Impact on DAS and Small Cell Ecosystems

Deployment of vRAN is likely to be driven by venues and indoor applications, where demand for capacity is highest. This would precede deployments in macrocells, where there is already a large install base of LTE equipment in over 490 networks worldwide and a change in architecture is unlikely to occur before a major technology upgrade to 5G. The vRAN market will take off, provided the fronthaul connectivity requirements are similar to those of backhaul. vRAN would be a substitute for small cells and DAS, which is not optimized to support MIMO technologies, a leading feature in LTE (4x4 MIMO is a key feature of LTE-Advanced Pro; 3GPP Release 12 & 13). This development means greater overlap and interdependency between DAS vendors and TEMs.

Fixed Access Service Providers and Neutral Hosts

As LTE expands to unlicensed bands (e.g. 5 GHz) and shared spectrum bands (e.g. 3.5 GHz CBRS and 2.3 GHz), third parties will have the option to roll out LTE services there, concentrating on the indoor and venue markets. This allows companies with fixed assets such as fiber or cable, as well as neutral hosts, to enter the access service market with wireless solutions complementing those of the MNOs who own the wide-area coverage market.

Evolution Towards 5G

RAN Virtualization is a major topic as the definition of 5G networks emerges with varying use cases including extreme broadband, massive machine-type connectivity and ultra-reliable communications. The ability to run services at the network edge to optimize bandwidth utilization and user experience requires a configurable architecture. The scale which 5G networks are required to support can only be implemented cost effectively with a scalable and elastic network architecture. RAN virtualization provides this capability. However, as 5G incorporates millimeter wave bands for access services, different architectures will be in play as millimeter wave systems rely on large antenna arrays to achieve the desired coverage range.

The Financial Business Case

Analysis of different RAN architectures shows that the centralization of baseband leads to high operational cost savings in Asian markets (26%). This is due to the structure of cell site leases, limited availability of space at the cell site, and high energy costs. In North America, the structure of site leases is beginning to change. Energy costs are relatively low, such that the business case for vRAN would not be positive in all cases, especially as dark fiber will be required to meet the requirements of CPRI fronthaul. This results in high financial uncertainty and risk that deployment requirements can be met.

In HetNet deployments, fronthaul can overcome the advantage of wireless backhaul cost effectiveness ($/Mbps) only if we consider high utilization of the remote cell. While Virtual and Hybrid RAN boost capacity, the average utilization of small cells over time is generally low, which erodes the return on investment. This issue is endemic to the HetNet architecture irrespective whether it is based on small cells or low-power remote radio.

In HetNets, fiber fronthaul is attractive in connecting remote small cells that are close to the macrocell. This is where interference between the HetNet layers is highest due to proximity. The breakeven point is about 75m: any remote cell at greater distance than 75m is better connected through wireless, if possible.

The major financial implications with vRAN is with regards to capital expenses. CAPEX reduction is driven by the baseband pooling gain of vRAN, however, that will depend on a number of factors. Primarily CAPEX savings depend on the deployment scenario and size of vRAN cluster, which is an MNO design option. Among other factors is the pricing model from vendors.

The Ecosystem

The Cloud RAN ecosystem comprises a wide cross section of vendors from the entire wireless ecosystem (Figure 6). However, we consider that a critical element of the ecosystem includes the Internet giants who are looking to reduce the cost of access to reach more subscribers and provide better quality of OTT services. Another important element of the ecosystem are the cable and fiber operators, whose fiber and other fixed access assets will have a major role in providing fronthaul services. These service providers already operate Wi-Fi as an extension to their fixed access services and some have looked

Figure 6 Cloud RAN ecosystem.

Conclusions

vRAN is a forward facing disruptive technology that is rapidly becoming more feasible as it garners support from Internet giants and startup pioneers. Current architectures being pursued by the TEMs, such as Hybrid RAN, will allow MNOs to improve the performance of HetNets specifically related to interference and mobility management, but will fall short of having a disruptive impact on the industry. Disruption will come from vRAN technologies when the fronthaul challenge is solved. This will alter the MNO-TEM relationship and market structure, and will allow new entrants into the market such as the fixed access service providers who can leverage their infrastructure for fronthaul services. The advent of RAN virtualization becomes especially potent when coupled with shared spectrum regulations, which increases the service possibilities and market opportunity.

Acronyms

3GPP	Third generation partnership project
5G	Fifth generation
ARPU	Average revenue per user
ASIC	Application-specific integrated circuit
CAPEX	Capital expenditure
CBRS	Citizen Band Radio Service
CoMP	Coordinated multipoint
CBRI	Common Public Radio Interface
DAS	Distributed antenna system
DSP	Digital signal processor
EBITDA	Earnings before interest tax depreciation and amortization
FPGA	Field programmable gate array
GPP	General purpose processor
HetNet	Heterogeneous network
LTE	Long Term Evolution
MIMO	Multiple input multiple output
MNO	Mobile network operator
MVNO	Mobile virtual network operator
NPU	Network processing unit
OCP	Open Compute Platform
OPEX	Operational expenditure
OTT	Over-the-Top
PDCP	Packet data convergence protocol
PoC	Proof of concept
RAN	Radio access network
RLC	Radio link control
RRC	Radio resource management
SoC	System on chip
TEM	Telecom equipment manufacturer
TIP	Telecom Infrastructure Project
vRAN	Virtual radio access network

Xona Partners (Xona) is a boutique advisory services firm specialized in technology, media and telecommunications. Xona was founded in 2012 by a team of seasoned technologists and startup founders, managing directors in global ventures, and investment advisors. Drawing on its founders' cross-functional expertise, Xona offers a unique multi-disciplinary integrative technology and investment advisory service to private equity and venture funds, technology corporations, as well as regulators and public sector organizations. We help our clients in pre-investment due diligence, post investment life-cycle management, and strategic technology management to develop new sources of revenue. The firm operates out of four regional hubs which include San Francisco, Paris, Dubai, and Singapore.

www.xonapartners.com advisors@xonapartners.com @xonapartners