Mobile anyhaul

White paper

Mobile backhaul is evolving from a static, linear connection to a programmable mesh interconnecting all mobile and cloud elements dynamically. It must evolve because the radio access and packet core elements it interconnects are also changing to support new services and industry verticals. In the course of this evolution, traditional distinctions between fronthaul, backhaul, and backbone are fusing toward an end-to-end 'anyhaul' architecture. There is some urgency as transport capabilities must be deployed before mobile layer innovations can take place. With the acceleration of 5G and many of its capabilities being advanced into 4.5G now and 4.9G by the end of 2017, planning must begin now.

Contents

Introduction

This white paper discusses the evolution of mobile transport. New use cases for mobile and wireless services are driving innovations in the radio access network and the packet core. Since mobile backhaul interconnects these mobile elements, it too must evolve to support them and the new services they will offer. There is some urgency to the transport evolution because to realize the benefits of mobile layer innovations, refreshed transport capabilities must be deployed first. Like the introduction of 4G/LTE almost a decade ago, the path to 5G offers new opportunties to gain a competitive advantage, but on a much larger scale. With the acceleration of 5G and many of its capabilities being advanced into 4.5G now and 4.9G by the end of 2017, planning must begin now.

Mobile layer innovations that are emerging include mmWave spectrum, centralized RAN, and the cloudification of RANs and packet cores. Together these innovations drive new requirements into the underlying transport network including:

- delivering connectivity to more end points, in more diverse locations, and with more capacity

- carrying more traffic types with more challenging latency, synchronization, and security requirements

- enabling more automated network provisioning to match the agility and dynamism of the cloud.

As these requirements are addressed, the traditional distinctions between fronthaul, backhaul and backbone will blur sufficiently to warrant a new description. Nokia calls this new, end-to-end transport architecture 'mobile anyhaul'.

Expanded use cases on the path to 5G

The proven financial benefits of being first to market with a new service are driving a flurry of activity across the communications industry. And rightly so: 5G has the potential to change the way people live and transform the way businesses work in nearly every vertical industry.

Recognizing the potential for 5G performance to run innovative services is one thing; squaring these against the financial commitments needed to turn in a profit is more difficult, simply because no commercial deployments exist yet. At the same time, delaying investment risks losing a competitive advantage that will be hard to regain.

Nokia has conducted in-depth modelling to give communications service providers (CSPs) and verticals realistic insights into the technical and commercial factors that affect a business case's profitability and investment return[1]. Some of the earliest business cases[2] have focused on three specific use cases—immersive video delivery at an event, smart cities and the Internet of Things (IoT), and Industry 4.0 factory automation.

Immersive video experiences to a large number of subscribers

Any immersive video service aims to enhance the experience of spectators attending a major sporting event taking place in a large stadium. Subscribers can use their mobile devices to access live video coverage in the stadium by switching instantly between multiple camera angles, or experience the action through virtual reality. Peak data rates of 100 Mb/s and thousands of concurrent users can be expected.

Ultra-dense IoT in a mega-city

A smart city or smart mega-city uses a vast range of IoT applications and devices to improve quality of life and efficiency of urban operations and services, and to achieve high economic, social and environmental performance. The underlying connectivity technology must meet very diverse requirements in terms of throughput, latency, mobility and reliability.

Ultimately, there could be as many as one million IoT connections per square kilometer. Such an extreme density of devices and sensors is expected to create more signaling on networks than is generated by human mobile broadband use. Handling this signaling load more efficiently is a growing priority for network operators.

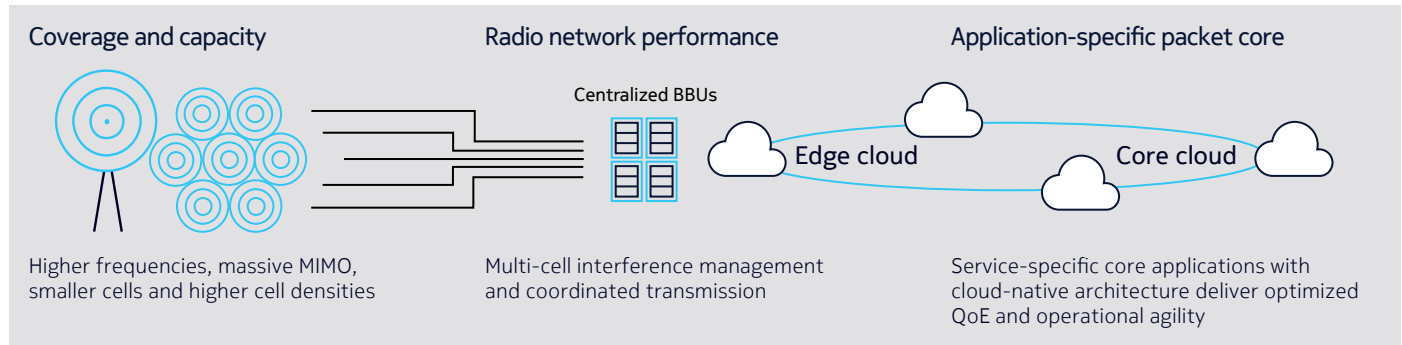## Wireless connectivity to support the Industry 4.0 factory

Industry 4.0 describes smart factories in which miniaturized processors, storage units, sensors and transmitters are built into machines, products, materials and smart tools that are networked to interact seamlessly and reliably to optimize production. The connectivity underlying the smart factory must be ultra-reliable and deliver extreme low latency, but also support a more dynamically configured factory floor than is possible today.

---

1   Translating 5G use cases into viable business cases. Nokia White Paper, 2017.
    www.networks.nokia.com/innovation/5g-use-cases

2   The 5G advantage in real network scenarios. Nokia White Paper, 2017.
    www.networks.nokia.com/innovation/5g-use-cases

Mobile layer innovations

Figure 1. New services and industry segments drive mobile layer innovations

Coverage and capacity	Radio network performance	Application-specific packet core
Higher frequencies, massive MIMO, smaller cells and higher cell densities	Multi-cell interference management and coordinated transmission	Service-specific core applications with cloud-native architecture deliver optimized QoE and operational agility

These use case examples drive mobile layer innovations illustrated in Figure 1. At the center of every mobile and wireless service is spectrum, and there are only three ways to get more of it: add more, get more out of it, and re-use it more effectively. 5G helps on all fronts with its support for mmWave spectrum and its massive capacity and peak data rates. Unfortunately, the poor propagation characteristics of 30 GHz and higher frequencies limit cell sizes to about 200 meters. Massive multiple in, multiple out (MIMO) helps offset propagation issues and is more feasible given the smaller, millimeter-sized antennas. Cells, however, will ultimately be smaller and denser so the spectrum will be used and re-used in smaller footprints serving fewer devices. Smaller cells means fewer devices contend for spectrum resources than they would in a traditional 10km-wide macro.

An issue with denser, smaller cells is an increase in cell edge interference. This must be managed or it will detract from the spectrum's abilities. This is a big driver for C-RAN architectures where baseband processing is centralized. The concept of pooling processing functions in hubs is an effective way to improve multi-cell interference management, thereby ensuring the most spectrum is available for payable traffic.

Interference can also imply a benefit: when a device receives signals from multiple cells, the transmission to it can be coordinated (that is, multiple radios transmit to a device at the same time). Like interference management, it requires inter-cell processing, and a centralized architecture is ideal for this. Centralized processing is also a natural step toward cloud RAN and the use of cloud compute infrastructure to scale and reduce operational cost. The cloud also helps support the expansion of IoT devices, whose data transmission may be small, but whose control plane requirements are not, given the potential for millions and millions of devices.

The mobile packet core is also changing. It is being deconstructed, disaggregated and distributed to support the expanded range of new use cases, including Industry 4.0 and the specific case of critical machine

ccommunication. These applications demand ultra-low latency: the round trip delay for one device to transmit and receive signals to another will need to be less than 1 millisecond. To appreciate the challenge of hitting this number, consider that it is common in LTE to have a highly centralized packet core located as much as 1,000 km away from the physical cell sites it serves. It takes 5 microseconds for light to travel across 1 km of fiber. For 1,000 km, that means 5 milliseconds up to the packet core and another 5 milliseconds back down. If the up and down path transit delay must be half a millisecond, the specialized packet core must be located within 100km of the radio. It will likely be much closer in practice as operators address hyper-local opportunities in the first instance, thereby providing more momentum toward distributing cloud infrastructure, moving it outward to the edge of mobile networks.

Now, we will examine the impact of these mobile layer innovations on the underlying transport network.

Classical backhaul

Figure 2. Interconnecting cell sites and the mobile core with classical backhaul
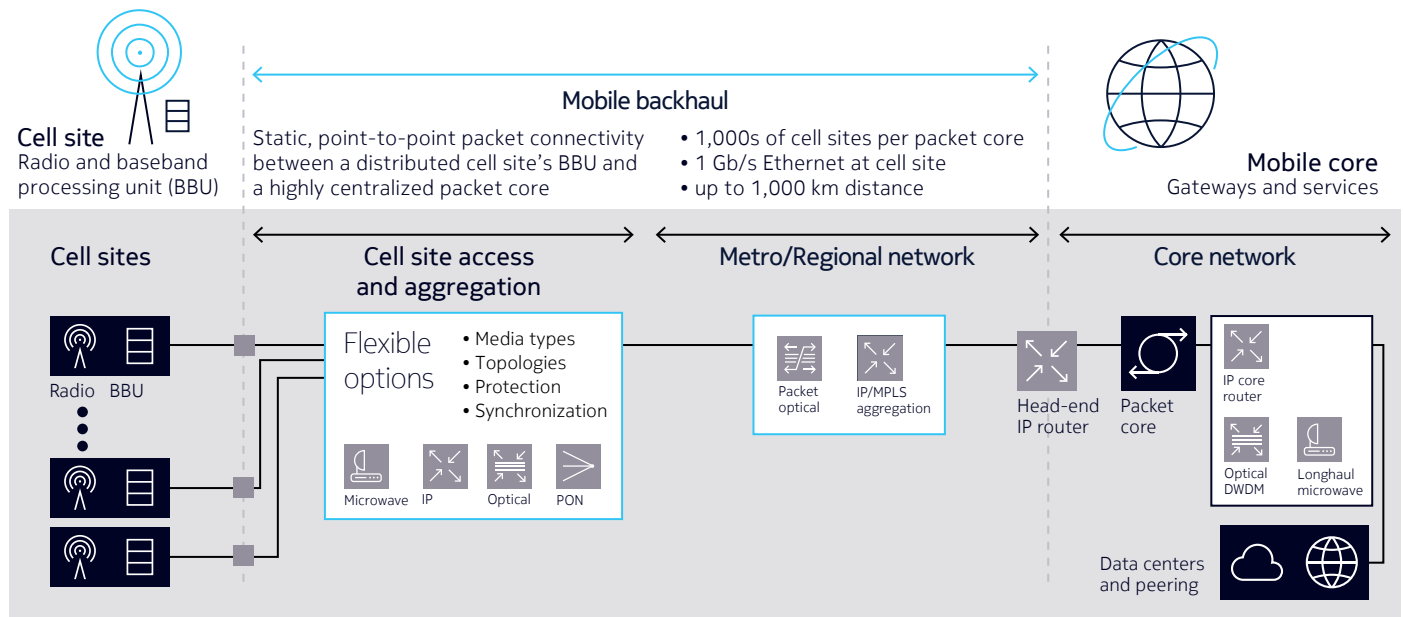
Co-located 2G/3G elements use same backhaul infrastructure.

Classical backhaul interconnected traffic between cell sites and the packet core as illustrated in Figure 2. The connection was static in the sense that once the network service was provisioned, it did not change. Change was not required because the baseband unit (BBU) and packet core elements ran perpetually on dedicated devices in fixed locations. A large, nationwide LTE network might have a dozen packet cores in total.

The point-to-point connection between a BBU and packet core uses Ethernet interfaces with GigE at the cell site. A highly centralized packet core serves thousands of cells as much as 1,000km away. Once traffic has passed through the packet core gateway, it is routed to application servers, peering points and Internet gateways.

The transport details beneath this simple concept are complex. The two main challenges are physical cell site access and several stages of aggregation as traffic is 'backhauled' between the BBU and packet core.

Cell site access is inherently problematic given the distributed nature of cell sites and the construction issues of bringing suitable network connectivity to them. Fiber is the media of choice but it is not always available or affordable. As a result, microwave transmission tends to be used more than fiber, but the ratio varies. The decision is ultimately an economic one, taking into account the total operational cost. Passive optical networking (PON) is also an option in neighbourhoods where fixed broadband infrastructure runs next to cell sites.

Some operators extend IP out to cell sites and use microwave, optical, and fixed networks or dark fiber for physical connectivity. Doing so provides a consistent, end-to-end IP OAM across a combination of layer 1 and 2 networks including spans owned by third parties. Another benefit favours converged services operators who have established IP operational models for their residential and enterprise business services and want to extend it to their mobile networks for consistency and simplicity.

Regardless of the technologies deploy to reach cell sites, topologies are as diverse as the cell site locations themselves, including chains, trees and rings— whatever provides the most economical connection. Topologies usually include one or more aggregation points as well.

Protection mechanisms are layered on top to provide failover and route diversity, an increasingly important aspect for mobile service differentiation in a world dependent on always-on connectivity. Security mechanisms add encrpytion to ensure transport integrity.

Synchronization is critical, of course, and ensuring frequency, time of day and phase alignment with cell sites over Ethernet transport require advanced and sophisticated algorithms.

Once connectivity is achieved to cell sites and traffic reaches the metro or regional network, another aggregation occurs as traffic is converged and packaged at the most efficient layer to utilize 100G IP/MPLS and packet optical facilities, and is ultimately transported to the packet core.

The core network then interconnects IP packets to and among data centers, peering points and the Internet. This involves the interconnection of massive, terabit core routers over equally massive optical DWDM switches. Longhaul microwave is also available, providing another economical option for connectivity over long distances.

Centralized RAN and fronthaul

Figure 3. Interconnecting remote radios and pooled processing resources
with fronthaul
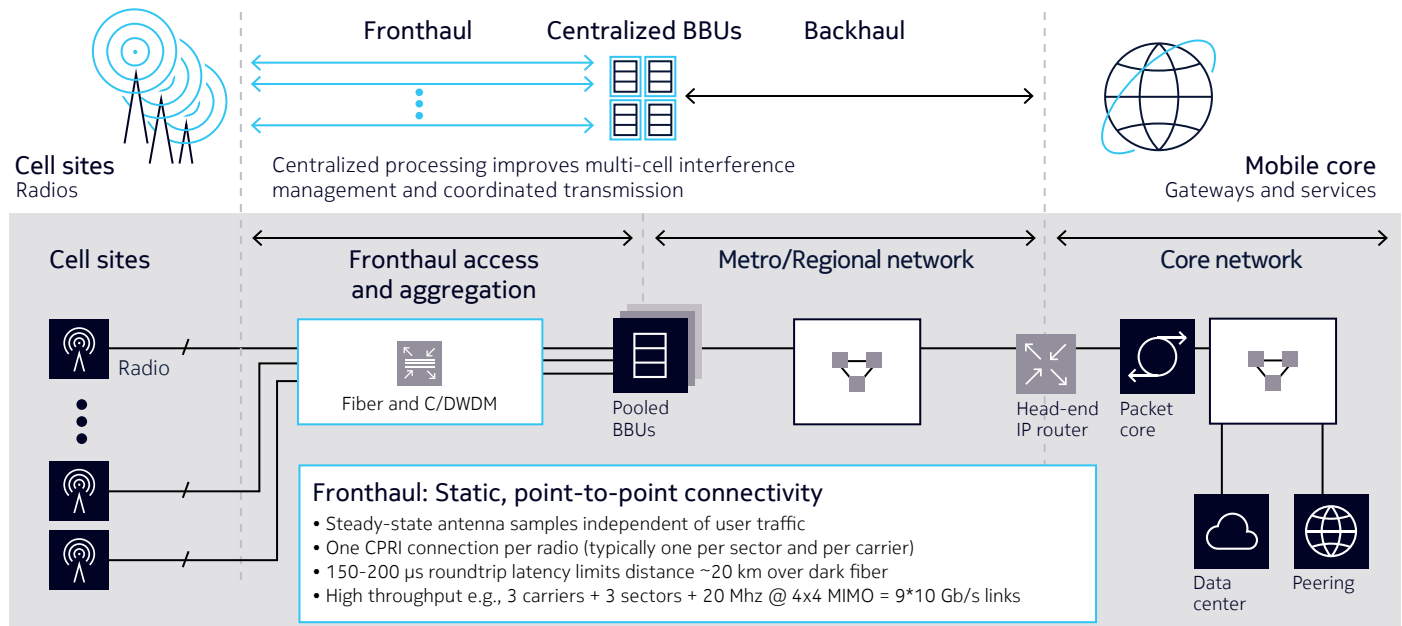
Centralized RAN pools processing functions and is an effective and efficient solution for interference management and coordinated transmission since many cells are processed in the same hub. Centralized RAN deployments have been underway for several years, especially in aggressive LTE-Advanced markets like Korea, Japan, China and the United States.

Relocating the BBU to a centralized hub requires fronthaul. Illustrated in Figure 3, the concept is simple enough and builds on an approach to cell site design more than a dozen years ago when radio heads were removed from the hut and placed at the top of the mast, next to the antenna. The interface that ran from the remote radio head at the antenna to the BBU in the hut below was called CPRI, and later another specification called OBSAI. In centralized RAN, the CPRI interface is simply extended from the radio site to the centralized hub.

The challenge with CPRI is it was never intended to do this. First, the roundtrip latency budget is only 150 to 200 microseconds. Given it takes 5 microseconds for light to transit 1km of fiber, the centralized hub location must be within 15km of the cells it supports.

Secondly, CPRI is a high and steady-state flow of antenna samples independent of actual cell traffic. One CPRI link is required for each radio, which can lead to significant transport capacities. For example, 9 CPRI links at 10 Gb/s each are required for a cell site with three carriers, three sectors and 20 MHz using 4x4

MIMO. It works within these constraints using dark fiber and WDM and is being aggressively deployed in dense urban environments. While it is LTE today, these operators are also preparing for mmWave spectrum in 5G and the benefit of centralized processing. Getting transport infrastructure in place today, where they can, is prudent planning and forward thinking.

CPRI, however, is not an appropriate solution.

The emergence of next gen fronthaul

Figure 4. Next gen fronthaul
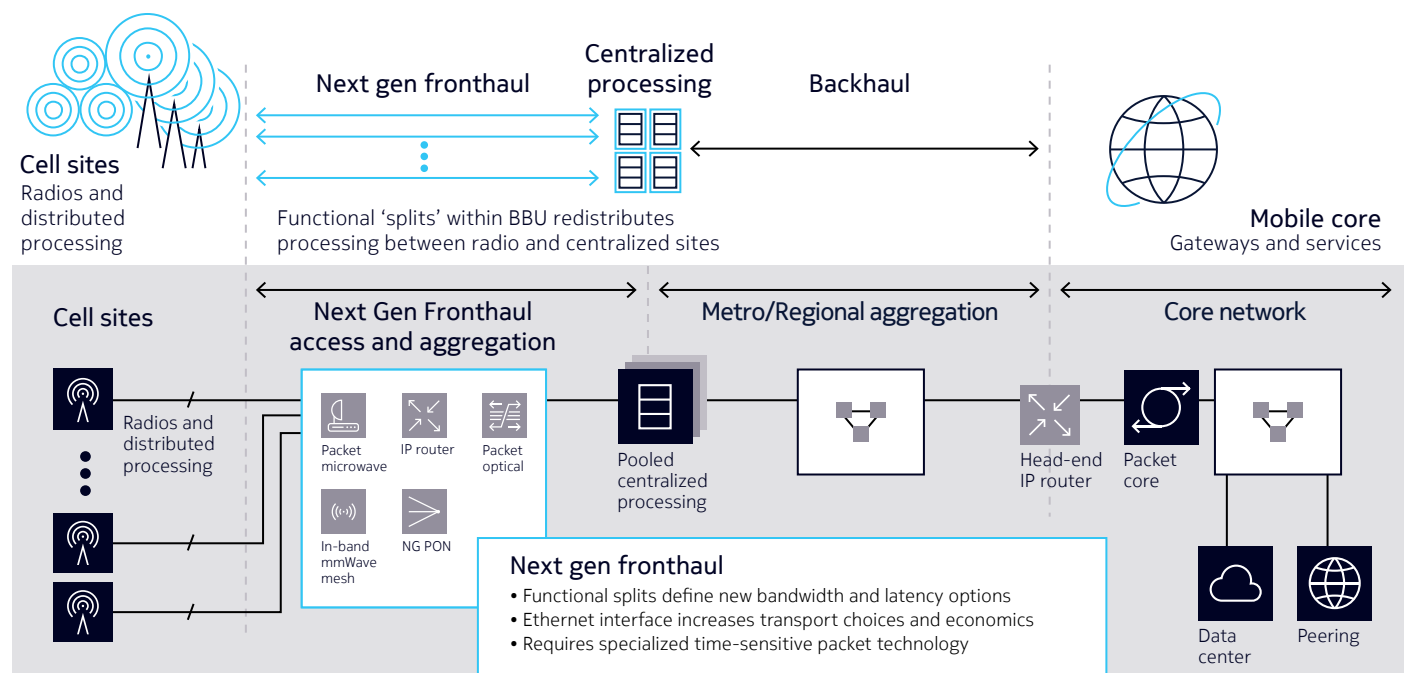
Many industry and standards organizations are pursuing a new RAN specification that will redistribute the BBU functions between cell site and centralized hub locations. The technical jargon is 'split processing'.The vertical, multi-layer BBU processing hierarchy is split horizontally at various points: the top half is centralized, the bottom half is distributed to the cell site.

The actual 'splits' and their specific requirements remain a work in progress. Nokia believes the industry will likely define two splits, one high level and one low level. What is generally agreed by all is the bandwidth requirement will be significantly less than CPRI by a factor of 10 or more and the physical interface will be Ethernet. The ultra-low latency challenges of CPRI will remain but may be slightly relaxed in some models.

Illustrated in Figure 4, a new next gen fronthaul interface using Ethernet allows more packet technology choices updated to support low latency. The fundamental issue with latency and packet technologies is that packet networks

are designed for bursty traffic and cost-control through statistical gain. In practice that means deep buffers to 'smooth out' the peaks with minimal packet loss. Buffers, however, add delay and reducing it by enabling variable buffer depths is a major focus of next gen fronthaul support. To underscore the issue, the typical buffer size for an Internet router is 50 milliseconds! If a round trip latency target of 1 millisecond is needed, then clearly something has to change.

Nokia has done just this. Our entire transport portfolio of routers, optical and PON switches and microwave radios have been updated to support concurrent requirements for low latency, synchronization and security. The portfolio has also been optimized for 10 Gb/s cell site access, an increase from today's typical 1 Gb/s speeds.

All these packet technologies are available except for inband mmWave mesh, a proof of concept tied to 5G. So much bandwidth is available at 30 GHz and higher that some of it can be partitioned off for fronthaul. This makes a lot of sense when one considers the ultra-dense 5G hotspots whose street-level antennas will quite possibly be located on every city block. Using some of the spectrum to intelligently transport and aggregate traffic back to a fiber or microwave point adds another economic option to consider.

The cloudification of mobile networks

Figure 5. Cloudification of mobile networks and dynamic connectivity
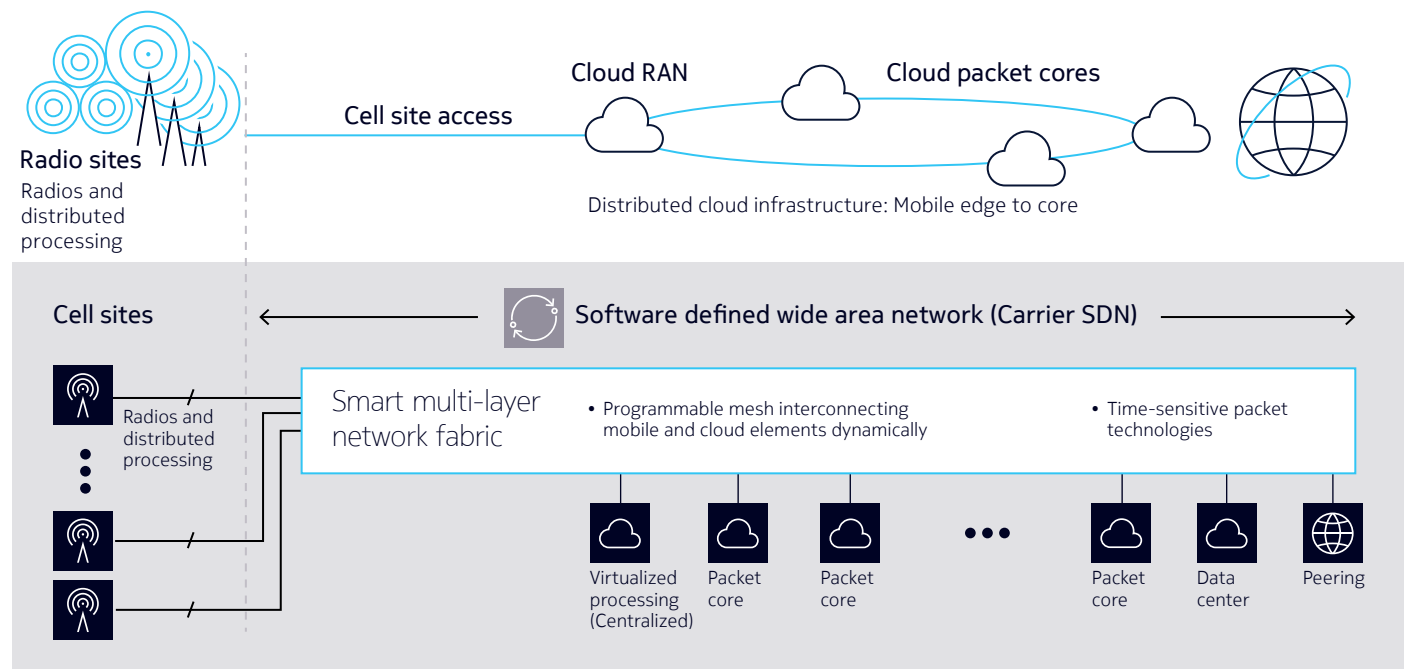
The cloudification of mobile networks has begun for RAN processing and the packet core. For the RAN, it is a natural evolution of centralized RAN, because whatever amount of RAN processing becomes centralized, hosting it on cloud compute infrastructure means operators can respond faster to the scaling, performance and cost challenges it entails.

For the packet core, the new use cases for mobile and wireless services require an expanded range of packet core applications that fit naturally with the shift to cloud infrastructure. For example, to support latency targets for critical machine communications, the packet core must be located within 100km of the radio site. With the distribution of cloud infrastructure out to the mobile edge, this becomes a logical spot to host it. Being cloud-based and without any physical workflows to delay service launches, operators have the agility to respond faster to new business opportunities.

Because services can be launched faster and QoE is driving the decision on where to deploy the packet core and application, operators can host specialized, ultra-localized services for a long weekend, for example, giving them a competitive advantage.

All this application agility is meaningless, however, if the provisioning of network services cannot keep up. Illustrated in Figure 5, this is where software defined networks come in.

To be clear, data centers and cloud compute infrastructures are full interconnected by fiber and microwave links. Physical connectivity is not the issue. The issue is that they comprise a packet network upon which a network service runs. This end-to-end network service needs to be instantiated on the cloud resources providing the RAN processing and packet core applications.

This is possible and is performed today. The orchestration system that establishes applications tells the SDN controller to create the network service at the same time. 'Carrier SDN' describes the functionality needed to provision end-to-end network services over multi-layer and multi-technology wide area networks.

With network provisioning (and optimization and assurance) being managed by the carrier SDN controller, the physical transport infrastructure becomes more of a programmable fabric which interconnects all mobile and cloud elements dynamically. This 'smart fabric' extends end-to-end to support connectivity among virtualized processing and applications wherever they may be located, using the transport technology and network service type that make most economic sense. Since this includes connectivity to radio sites, the transport's programmability must include time-sensitivity in the network service SLA.

Fronthaul, backhaul and the coexistence of LTE and 5G

Figure 6. Interconnecting cell sites with LTE and 5G
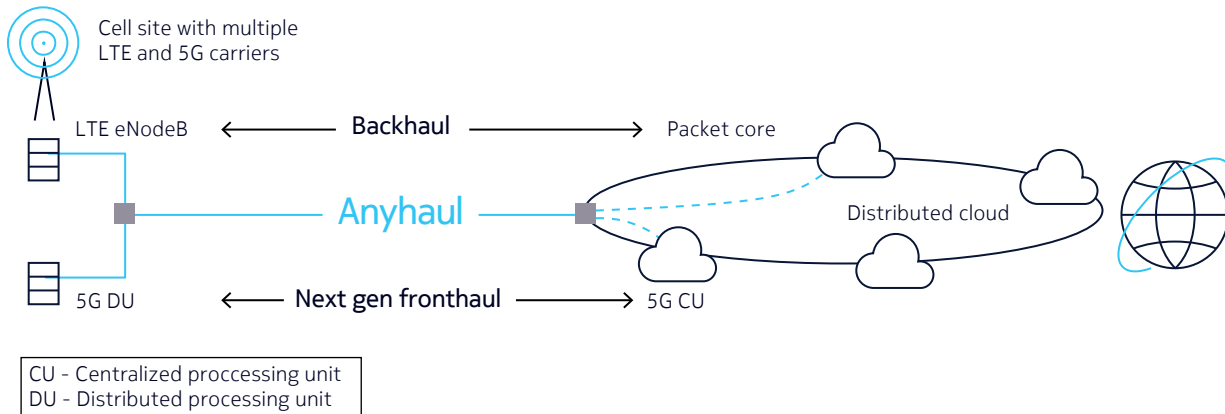
With all the industry talk on 5G and the history of 'new Gs' replacing 'old Gs', it is easy to think that 5G replaces LTE. It does not: 5G's new capabilities complement and are intended to coexist with LTE.

This coexistence and the continuing discussions around split processing have created an expansion of terms such as mid-haul. Even the traditional fronthaul, backhaul, and the core backbone networks are blurring as we embrace the dynamism of distributed cloud computing and the coexistence of multiple radio access technologies and architectures per site.

Figure 6 illustrates the coexistence of a new 5G radio and a traditional LTE macro with a distributed BBU. Both 'Gs' would use the same transport facilities. So what do we call it? Backhaul? Fronthaul?

At Nokia we use 'anyhaul' to describe the new scope of mobile transport requirements on the path to 5G. Anyhaul networks are comprised of the transport technologies needed to reach radio sites and provide interconnectivity between them and the cloud in the most economical and dynamic way.

Conclusion

Figure 7. Mobile anyhaul architecture interconnects mobile and cloud elements dynamically
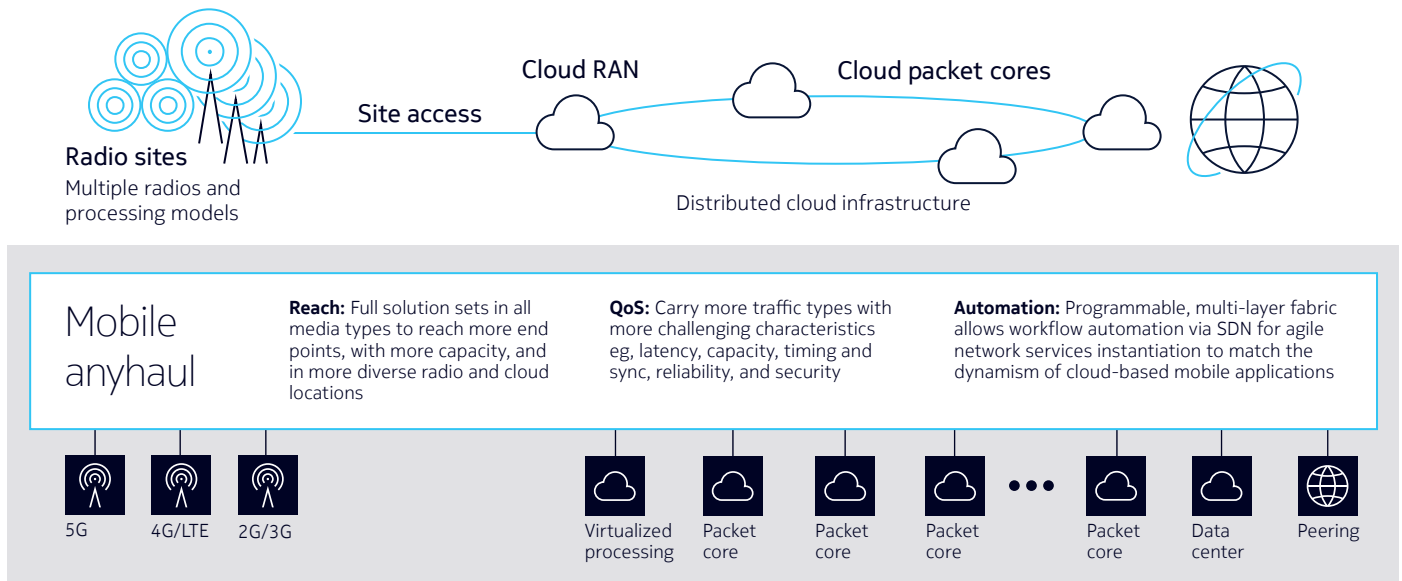
This white paper discussed the connections between new mobile and wireless use cases, the corresponding changes in the mobile layer to support them, and the ultimate consequences for the transport network. These consequences are significant, creating arguably the most exciting wide area network topic today on the path to 5G.

Mobile backhaul is evolving to support innovations at the mobile layer. Innovations such as cloud RAN and cloud packet core are happening today in 4G and will be foundational technologies in 5G.

The new requirements for mobile transport are changing sufficiently to warrant a new description, something broader than backhaul that also takes into account fronthaul and backbone as well as new requirements for extended reach, enhanced QoS for latency and strict phase synchronization etc., and SDN automation. Nokia calls this 'mobile anyhaul'. Illustrated in Figure 7, it is a programmable mesh interconnecting all mobile elements and the cloud. It supports multiple RAN technologies, processing models and sites, and distributed cloud infrastructure for virtualized RAN and packet core functions.

The anyhaul architecture reaches radio sites with full solution sets in all transport technologies to ensure the appropriate, most economical decision can be made. It carries more traffic volumes and more traffic types including low latency as well as perennial needs for synchronization, reliability and security. It consists of a programmable, multi-layer and multi-technology fabric that supports SDN workflow automation and agile network services instantiation.

These characteristics require changes in the way transport networks are designed and operated, and these changes affect all transport technologies. As a vendor with global reach and leadership in microwave, IP, optical, PON and mobile, Nokia is in a good position to provide a technology-neutral perspective on the broader, end-to-end architectural challenges, allowing each mobile network operator's requirements to dictate the appropriate resolution of them for their business.

Further information

Figure 8. End-to-end mobile anyhaul leverages full solution sets in microwave, IP, optical, NG PON and mobile.
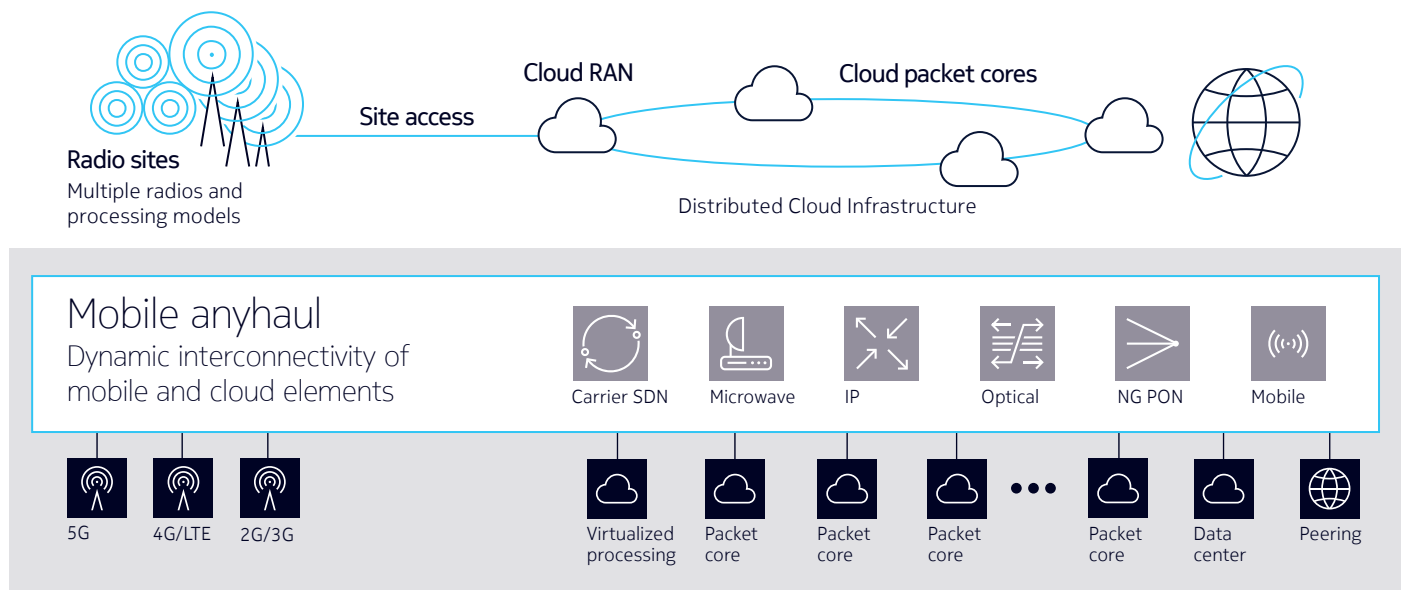
Nokia has expanded its industry-leading mobile anyhaul transport portfolio with new products and innovations specifically designed to address the critical requirements on the path to 5G. Illustrated in Figure 8, this expansion aims to satisfy the increasingly complex and diverse challenges of today's evolving mobile infrastructure with the industry's most comprehensive range of programmable, end-to-end networking solutions.

Built upon Nokia's deep pedigree in mobile, microwave, IP, optical, and fixed access technologies, the solutions leverage Nokia's early investments in SDN and virtualization to deliver the programmable IP interconnectivity needed to assure the heightened service requirements, and with an evolutionary path to support past, present and future mobile technologies on a common infrastructure.

NOKIA

More information on mobile anyhaul can be found at https://networks.nokia.com/solutions/mobile-transport. A summary of specific product updates for mobile anyhaul include the following:

The **Nokia Wavence Microwave Portfolio** is a new family of ultra-broadband transceivers. The ultra-broadband transceiver twin, a 'dual carrier in a box' radio, and the ultra-broadband transceiver 80, a compact E-band radio, support carrier aggregation and carrier SDN to deliver multi-gigabit, low latency and programmable microwave transport. Nokia Wavence is the evolution and the new brand of the Nokia 9500 Microwave Packet Radio family.

https://networks.nokia.com/products/Wavence

The **Nokia 7250 Interconnect Router R6 (IXR-R6)** is a new IP/MPLS router with terabit-scale, low latency, improved port densities and support for next-generation interfaces, such as Ethernet fronthaul, that expands and complements Nokia's IP/MPLS mobile transport solutions including the 7705 SAR, 7210 SAS and 7750 SR product families.

https://networks.nokia.com/products/7250-interconnect-router-r6

The **Nokia 1830 family** introduces new functionality to further enhance its optical anyhaul mobile transport solutions. New Integrated Packet Transport cards for the Nokia 1830 Photonic Service Switch address Ethernet anyhaul applications by combining the scalability of 100G packet aggregation and coherent DWDM with ultra-low latency and industry leading time synchronization. Additionally, the Nokia 1830 Versatile WDM Module (VWM) adds optical protection switching, assuring high availability of WDM fronthaul traffic.

https://networks.nokia.com/solutions/optical-networking

The **Nokia 7360 ISAM FX** access node and **7368 ISAM ONT** optical network termination devices extend the performance of Passive Optical Networks (PON) and are designed to cost effectively introduce more bandwidth with 10Gbps for PON and point-to-point technologies. Including network synchronization support, the new capabilities enable operators to easily leverage existing fiber-to-the-home deployments for 'anyhaul' applications.

https://networks.nokia.com/products/7360-isam-fx-shelf

https://networks.nokia.com/products/7368-isam-ont