White Paper

Everything You need to Know About 5G xHaul

This white paper is based on the ***Transport Network Aspects*** book chapter, authored by RAD's CTO Dr. Yaakov Stein, Dr. Yuri Gittik and Ron Insler, Head of RAD's Innovation Lab. The book ***5G Radio Access Network Architecture: The Dark Side of 5G***, by Sasha Sirotkin, will be published in 2020.

Your Network's Edge®

The focus of 5G deployments so far has entirely been on the **Enhanced Mobile Broadband (eMBB)** flavor of 5G, which delivers increased data rate, mostly to support entertainment applications, such as ultra HD and 3D video, Work & Play in the cloud, as well as all sorts of alternative realities – AR, IR, VR and MR.

That said, 5G is the first mobile generation that was specifically designed with real-world applications in mind. Hence, many of the technologies being developed for 5G center around solving significant real-world problems, like vehicle accident prevention (V2V), reducing traffic congestion (V2I), automatic meter reading (AMR), crime reduction (Smart City), increasing manufacturing efficiency (IIoT), etc.

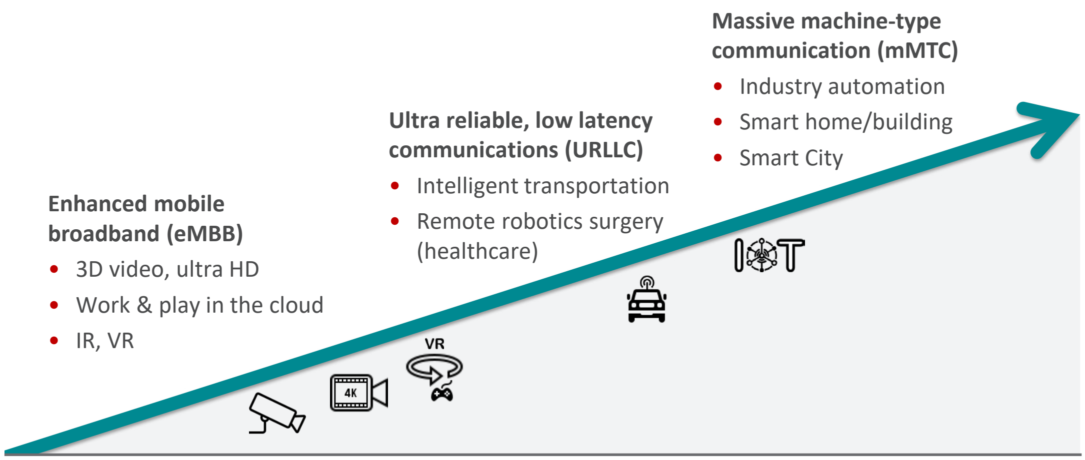

5G hype is everywhere. As of November 2020, we have seen more than 16,000 5G deployments world-wide, with varying degrees of readiness – from pre-releases to full commercial availability[1]. As the race to 5G is on, communications service providers (CSPs) are facing competitive pressures to be the first to launch 5G services. Given market expectations, spectrum auctions are gaining pace without necessarily coinciding with the CSPs' budgetary plans for new network buildouts. These developments need to be reconciled with technology uncertainties, as 5G represents a departure from previous cellular generations and the associated standards are still under development.

Before we move on to discuss what all of this means for mobile operators, network wholesalers and transport providers, let's first take a look at some basic 5G principals that are affecting both technological and financial aspects of the new cellular generation.

[1]Source: Speedtest's OOKLA 5G MAPTM

The 5G RAN – Decomposing the gNB

At the highest level of abstraction, a 5G network consists of three entities:

- 5G user equipment (UE)
- 5G radio access network (RAN including gNB base stations)
- 5G core network (5GC)

5G communications are carried out by interconnecting these entities. 3GPP specifications deal in minute detail with all aspects of the connection between the UE and gNB, but severely under-specify the connection between the gNB and the 5GC, viewing connections inside the RAN and to the core as ideal transport pipes.

The transport of data between the RAN/gNBs and the core is conventionally known as backhaul, in line with the terminology of previous generations of mobile communications.

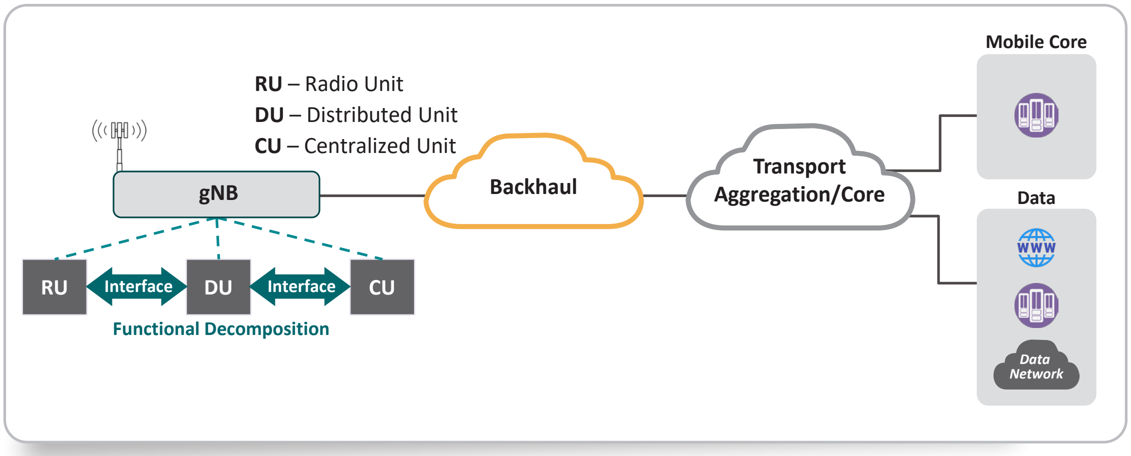

However, the gNB may be further subdivided. Typically, a 5G base station can be decomposed into a radio unit (RU), a distributed unit (DU) and a central unit (CU). There are several economic and operational benefits to splitting a 5G base station's functionality, including:

- Flexibility and modularity to better fit various use cases with different requirements (topology, latency, capacity, density, space availability, etc.)
- Simplified engineering of the radio site and reduced space and power requirements
- Better coordination for efficient use of radio resources and performance features
- More efficient use of the fiber plant
- Facilitating use of NFV and SDN technologies and migration to vRAN/cRAN

The transport segment between a remote RU and the rest of the gNB is called ***fronthaul***. If the CU is physically detached from the DU, the transport segment between them is termed ***midhaul***. These segments are collectively termed ***"xHaul"***.

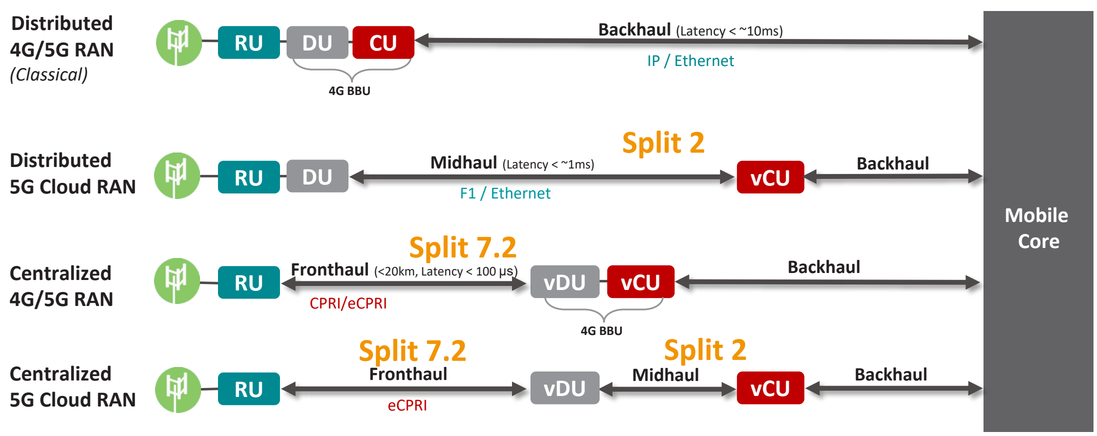

CU = Central Unit, DU = Distributed Unit, RU = Remote Unit

As we previously alluded, 3GPP standards tend to consider transport as a trivial function that effortlessly delivers data over the named interfaces with no availability failures, data-rate restrictions, burdensome latency, synchronization glitches, or other degradations. Unfortunately, this is not the case. Even well-engineered transport networks have limitations and occasionally fail to live up to design requirements.

5G's Paradigm Shift – From Mobile Backhaul to xHaul

When considering the key changes from the traditional 4G mobile backhaul to the evolving 5G xHaul it's clear that the former can be characterized by the following essential qualities:

- It connects radio sites to centralized locations where content and applications are hosted. The exceptional case of MEC, that brings content and processing closer to users (for example, with a local CDN) was not spectacularly successful, partially due to architectural mismatch and the requirement for complementary vendor-specific components.
- It supports a single class of service (CoS) for all applications.
- It is fashioned from static, pre-provisioned connections ("pipes") with star or ring topologies. Even when eNBs are inter-connected in the transport segment, this is accomplished using an aggregation node instead of direct X2 links.
- It is managed semi-statically, primarily using initial configuration with relatively rare updates throughout the network lifecycle.

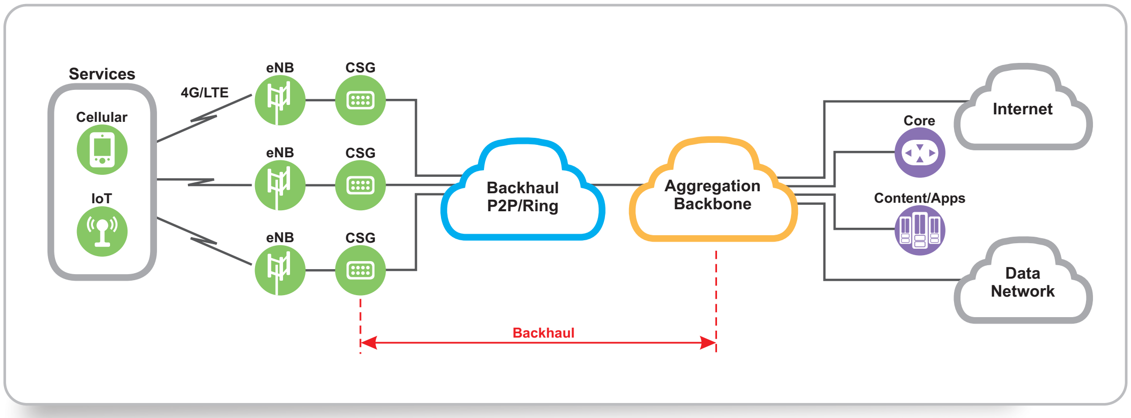

5G xHaul dramatically changes these basic features. This is driven by new use cases (with highly distributed content and applications) and a novel RAN architecture that mandates new transport attributes:

- Higher data-rates, lower latency and packet loss, higher reliability per application
- Hard isolation with end-to-end network slicing
- Dynamic connectivity for on-demand services encompassing multiple physical and virtual components

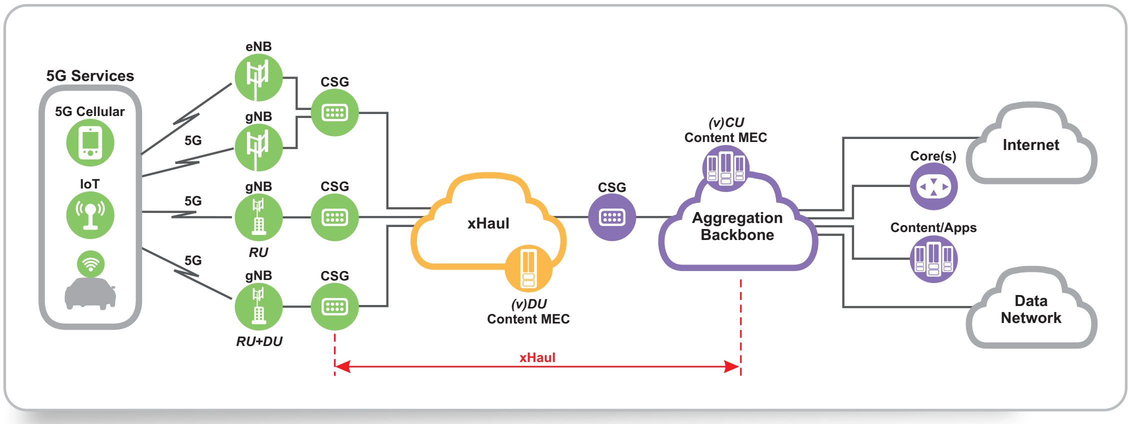

Virtualization and Edge Computing

Virtualization of functionality in the RAN applies to three different function types:

1. As mentioned above, gNB components themselves. These would most commonly be the vDU and vCU, although user plane function (UPF) and other mobile core functions may also be virtualized;
2. Networking functionalities required by the transport network itself, for example, security, fault management (FM) and performance monitoring (PM) probes and reflectors;
3. Service/user-centric **V**alue **A**dded functionalities and applications, such as firewalls, support for location-based services, CDN, and IoT aggregation.

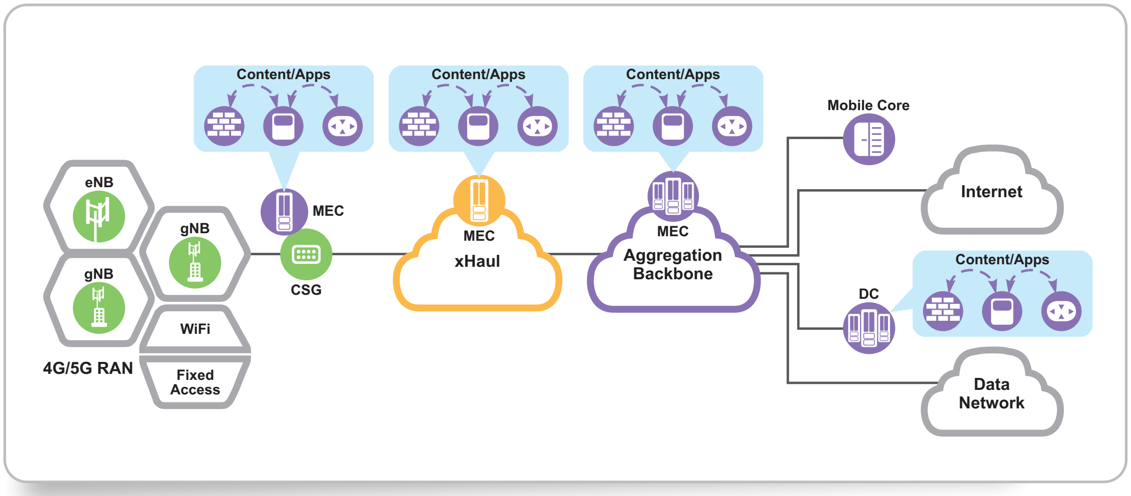

The first of these involves functions belonging to the mobile operator, although the transport provider may provide the computational platform and hosting services. The second function type focuses unambiguously on the needs of the transport provider (reducing costs, increasing automation), and is intended to be transparent to both mobile operator and end-user. Value-added functionalities benefit the end-user, and are frequently marketed by mobile operators, although they may be provided by third parties (a model facilitated by the 5G Service-based Architecture-SBA).

The first type clearly depends on the functional split option being used, while the second type is only indirectly influenced by the split (being directly susceptible only to traffic characteristics such as data rate or delay constraints). In some cases type 2 VNFs may be slice-dependent, in which case they must be able to map traffic to a slice.

User VAS are generally restricted to backhauling (split option 1) where user IP packets are discernible.

5G Slicing

Slicing is a key mechanism of 5G networks that allows using common networks for multiple services with different attributes such as bandwidth, delay requirements, required functionality, etc. In this context, a **network slice** is a logical network that provides specific capabilities and characteristics. The expected network slice behavior, in terms of features and services, is indicated by a slice/service type (**SST**) and each SST can be further divided into multiple slice differentiators (**SD**s).

In other words, when implementing network slicing an operator needs to decide on granularity. 3GPP Release 15 defines three SSTs mapping to eMBB, URLLC, and MMTC. In some networks, there will only be three slice instances (one per each SST), and traffic will be mapped into the slice instance of the required SST. In other networks, there may be many slice instances further differentiated by SD and traffic will be mapped accordingly, for example:

SST	SD	Slice Instance
eMBB	Corporate	eMBB/Corporate
	Residential	eMBB/Residential
URLLC	Factory	URLLC/Factory
	Utility	URLLC/Utility

In still other networks, there may be a huge number of slice instances in order to isolate individual user flows.

A key requirement to achieve this is to ensure an end-to-end network slicing support, not only in the RAN but also in the xHaul segment.

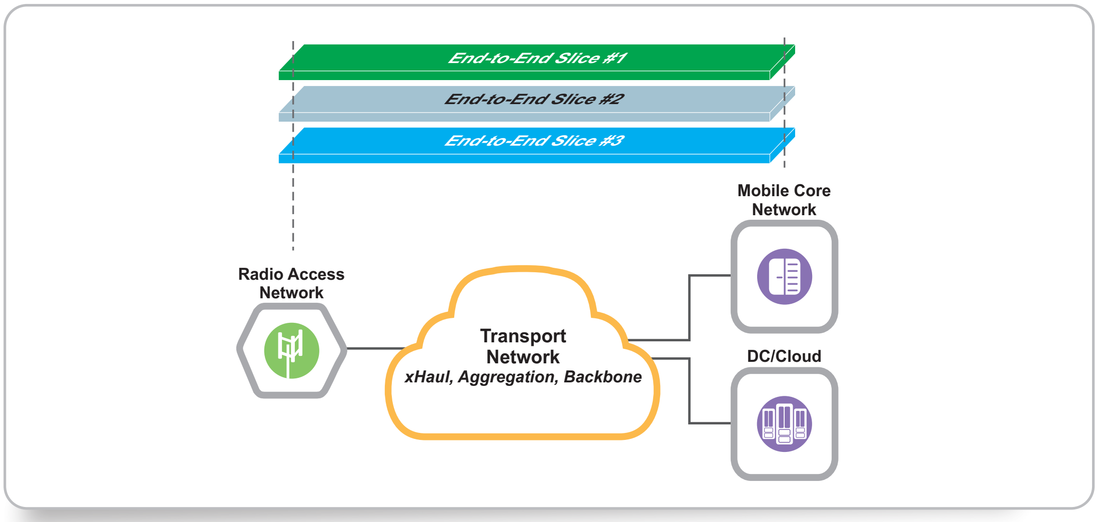

In transport networks, there are two main slicing implementation options – **hard** isolation slicing and **soft** isolation slicing, distinguished by the possibility of resource contention impacting service attributes. Hard isolation often involves provisioning **dedicated** resources to a slice instance, such as a dedicated transport fiber, lambda or TDM channel, etc., whereas in the latter slices may share multiplexed resources such as when employing IP/MPLS tunnels or pseudowires.

It's important to note that while soft isolation network slices might be negatively impacted by resource contention, they can never exchange information. For this reason VPN technologies have been suggested for soft isolation network slicing. After all, VPNs provide forwarding table isolation between different connections and an overlay topology for the connectivity between different customer sites. VPN methods could hence furnish a valid alternative for eMBB and some mMTC slices, if they augmented by dynamic set-up procedures (most VPNs today require lengthy configuration times). However, they fall short when considering transport requirements of URLLC, such as service-tailored network functions and data forwarding mechanisms, as well as compute resource handling.

New RAN Requirements and the Essential CSG Checklist

So what does 5G mean in terms of the xHaul networks and transport services?
The following table lists some of the key requirements:

xHaul Needs	Transport Network Requirements
High capacity	High speed connections (n x 10G, n x 25G, 100G)
Densification	High speed connectivity for an increased number of nodes
Prompt on-demand connectivity	Dynamic traffic steering for network and inter-site connectivity
Service separation for multiple use cases and customers	Support for CoS and end-to-end slicing
URLLC compatibility	Support special transport technologies (e.g., TSN, DetNet) and architectures/functional splits for low delay and high reliability assurance
New Cloud RAN topology	High speed and stringent latency requirements to enable various C-RAN/D-RAN deployments
Virtual RAN and MEC support	Edge Computing infrastructure (distributed NFVI)
Service and network automation	Integration with SDN controllers and service and network orchestrator(s), zero-touch provisioning

As 5G technology is still being defined, the above list is likely to be supplemented with additional required functions not yet defined today.

A key enabler for CSPs to meet such requirements is the cell site gateway (CSG)

5G deployments require re-defining of the functionality of the CSG. The 5G-ready CSG performs aggregation of multiple traffic types originating in the cell, including 4G, 5G sub 6 GHz (FR1), 5G mmWaves (FR2), small cell transit traffic, WiFi hot spot backhauling, etc. The CSG is responsible for homogenizing traffic across generations and technologies, minimizing transport expenses (including energy efficiency and multiplexing/duplexing in order to minimize OOF fiber expenses), constructing a single frequency and time-of-day reference clock, providing a cell-wide heart-beat to ensure connectivity to the cell and all its components, and initiating fail-over self-healing procedures as needed. It may additionally incorporate a micro-data center platform for edge computation.

Ideally, CSGs should meet this checklist to fully support efficient 5G deployments:

- Full LOW-PHY CPRI to eCPRI conversion
- Assured minimal delay fronthaul
- Network slicing support
- Field-programmable hardware for 5G evolution
- Upgradable protocol evolution: TSN, DetNet, Segment Routing, SRv6, etc.
- Carrier-grade SLA assurance: Programmable network slicing and segmentation, MEF 3.0 traffic prioritization (H-QoS) and hardware-based Ethernet OAM/SAT
- SDN automation: NETCONF for lifecycle management, Zero Touch provisioning
- Timing synchronization: 1588 Grandmaster with Sync-E, GNSS holdover (APTS)

This set of capabilities allows operators to support highly distributed content and applications with multiple classes of service for various end-to-end 5G slices.

Choosing an Effective Migration Path

5G migration poses significant challenges for mobile operators from both engineering and financial perspectives. Early 5G adopters are embracing an evolutionary approach and implement 5G xHaul in phases. The first stage inserts an NR-capable in an existing 4G cell site, exploiting existing (4G) backhaul and aggregation networks. The next stage adds new functionalities to these networks to meet new requirements, incrementally upgrading them towards the 5G xHaul end-game by matching infrastructure investments to revenue flows. Later the 5G core network is activated alongside the 4G EPC, the latter eventually being deactivated. Here is an example for initial stages of evolution of the RAN transport network:

Phase 1: Focus on capacity increase by enabling higher data rates to support eMBB services and prepare the foundation for other services. This phase would typically include only coarse slice granularity.

Phase 2: Add edge computing (NFVI) and integrate it with the transport network. One example could be the virtualization of gNB components, such as vDU and vCU, which allows operators to introduce functional network splits as they increase cell site densification to deploy more advanced services.

Phase 3: Introduce fine slicing granularity and new technologies, such as TSN and DetNet, in specific network segments to significantly lower latency and increase reliability. These are key enablers for new URLLC services.

Takeaways

5G, and xHaul/transport in particular, are essentially a moving target and operators are still formalizing their deployment plans. Yet they face immense pressure to make early market impact. Early adopters are focusing on increasing capacity of their transport networks as well as on cell site densification.

A phased approach to 5G deployment allows CSPs to make early market impact while delaying some planning decisions such as fronthaul implementation until revenues ramp up. In this manner, the pace in which network slices are being activated can be matched to TCO considerations, to ensure a minimal risk when venturing on a new technology that is not yet full defined.

At the end of the day, much of such a high level of flexibility hinges on the CSG. Choosing wisely and asking the proper questions is the first step in the right direction.