How VCMTS Paves The Way For 5G Over DOCSIS

Exploring Software-centric Solutions for 5G Xhaul and FMC

A Technical Paper prepared for SCTE by

Brendan Ryan
System Architect
Intel Corporation

Ed Dylag
Marketing Development Manager
Intel Corporation

Thushara Hewavithana,
Wireless Systems Architect
Intel Corporation

Eric Heaton,
Platform Solutions Architect
Intel Corporation

Table of Contents

List of Figures

1. Introduction

Network convergence has been gaining a lot of attention in the telecoms industry recently with the architecture shift in both wireline and wireless communication networks to be more centralized in terms of core and baseband processing, and more distributed in terms of the analogue and RF functions.

This has also led to the disaggregation of the hardware and software functions of the network. Furthermore, the standardization of the interfaces in this network, for example the O-RAN, 3GPP wireless wireline convergence (WWC), Broadband Forum fixed mobile convergence (FMC), cable distributed access architecture (DAA) and the more recent flexible mobile architecture (FMA) initiatives enables wider participation of ecosystem vendors leading to competitive solutions for operators.

Centralized functions of the network such as 5G-Core for wireless and cable modem termination system (CMTS) for data over cable system infrastructure specifications (DOCSIS) are well suited to being implemented in software and to run on common off the shelf (COTS) hardware, giving better scalability over time and adaptability for different deployment scenarios.

There are many aspects of network convergence as shown below (from reference [1]), at business, infrastructure, service, and architecture level.

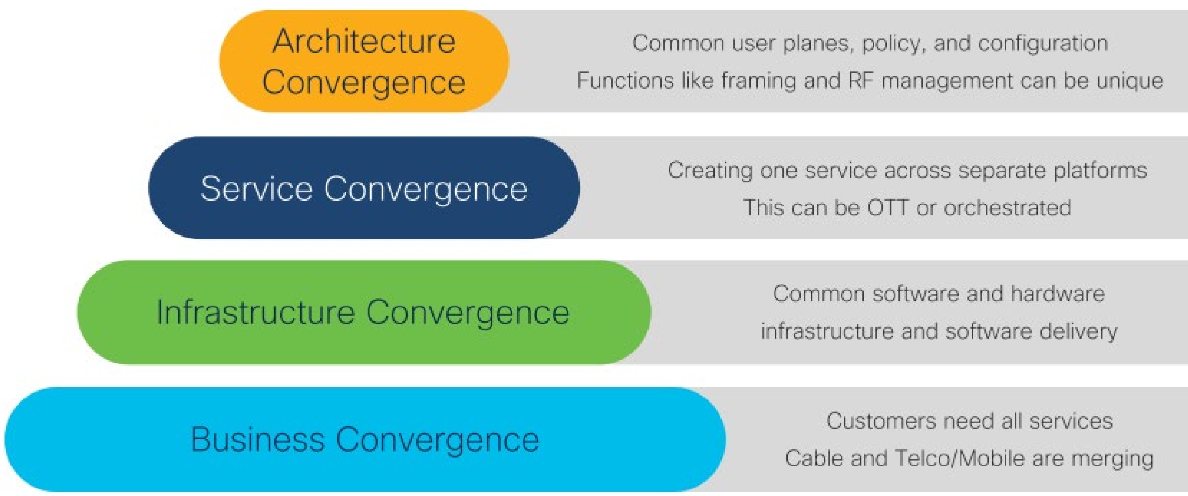

Source: Cisco & CableLabs, "Cable and Mobile Convergence, A Vision from the Cable Communities Around the World"

Figure 1 – Four Levels of Convergence

This paper will focus on infrastructure and architecture convergence and specifically, how the advent of 5G has created a great opportunity to leverage the existing DOCSIS network for two key purposes:-

- Xhaul of 5G small-cell traffic over DOCSIS
- Fixed mobile convergence (FMC) using the Cable access network

It will be shown how the flexibility of a software-based vCMTS makes it much easier to adapt the Cable access network for these new 5G use-cases and to evolve over time.

This paper will explore the rationale outlined above as to why Cable MSOs should prioritize vCMTS deployment in preparation for 5G xhaul and FMC, while identifying some gaps that need to be addressed by the Cable industry to prepare for 5G mobile convergence.

2. Cloud-native Infrastructure Convergence

Infrastructure convergence as shown in Figure 1 in the previous section is already happening in terms of common hardware platforms and software stacks being used to host wireless and wireline network functions as shown in Figure 2 below. Network function virtualization (NFV) and the subsequent progression to adoption of cloud technologies for communications network functions started in the Mobile core and is now gathering pace for central-office deployments of access network functions such as virtual broadband network gateway (vBNG), and the virtual cable modem termination system (vCMTS) in cable head-ends.

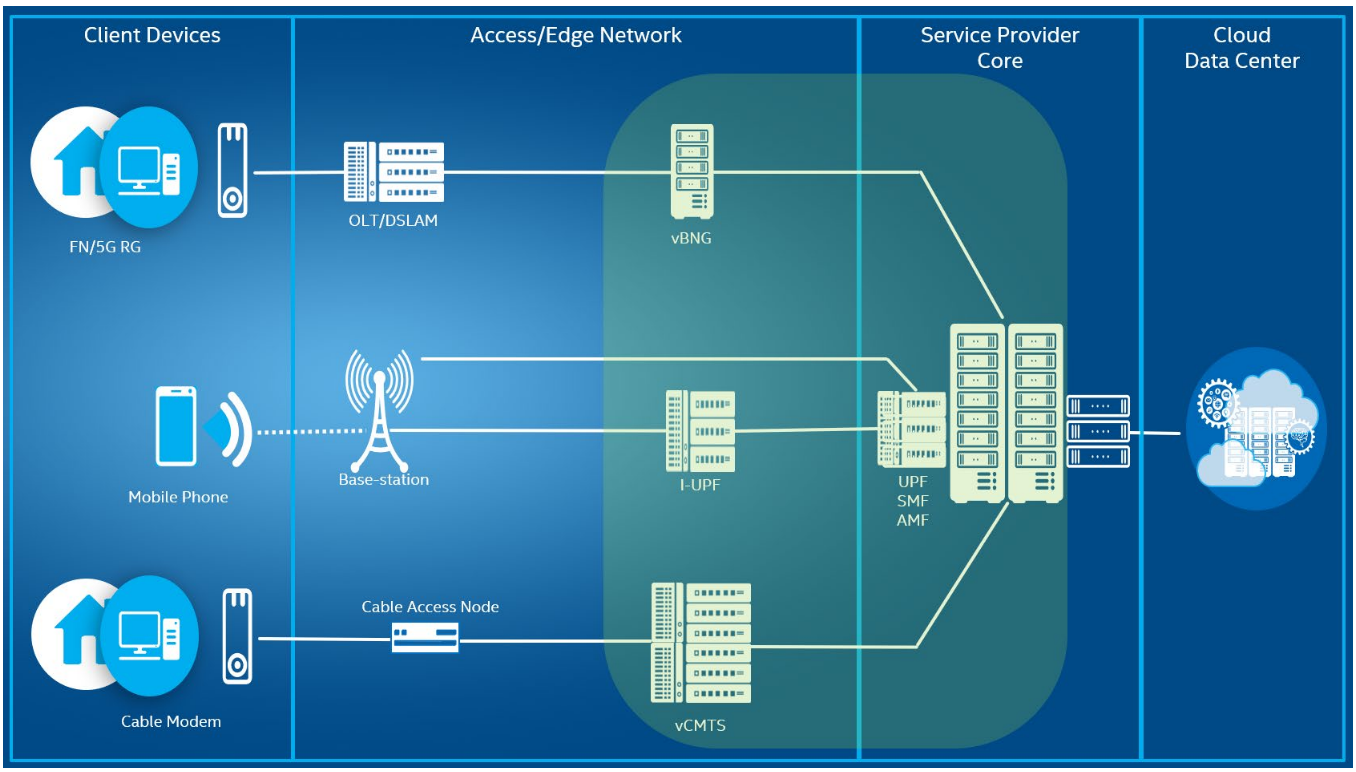

Figure 2 –Telecoms Network Infrastructure Convergence

Cloud-native is a term that has come to be used to describe the use of cloud technology for network function deployment on COTS hardware platforms. Significant work has been done in several mainstream open-source projects to enable high performance network function data-planes to be deployed on a cloud-native platform. Open-source packet-processing frameworks such as DPDK and FD.io/VPP (see references [2] and [3]) make the required data-plane performance possible in software while key features have also been added to open-source cloud computing platforms such as Kubernetes to enable network applications based on these packet-processing frameworks to be deployed with the scalability, flexibility, and observability of a cloud environment.

2.1. Cloud-native deployment of vCMTS

Significant effort has been invested in open-source projects to advance cloud-native network function virtualization by providing reference software which is highly optimized to run on general purpose COTS hardware platforms while also leveraging the benefits of cloud technology such as automation, scalability, flexibility and observability.

A Container Bare-metal Reference Architecture as shown in Figure 3 below (see reference [4] for details) is provided by Intel which combines general purpose hardware and open-source software components to provide a reference cloud-native platform which may be configured appropriately with Ansible playbooks to host network functions at various network locations - mobile core, central-office, access network edge and on-prem.

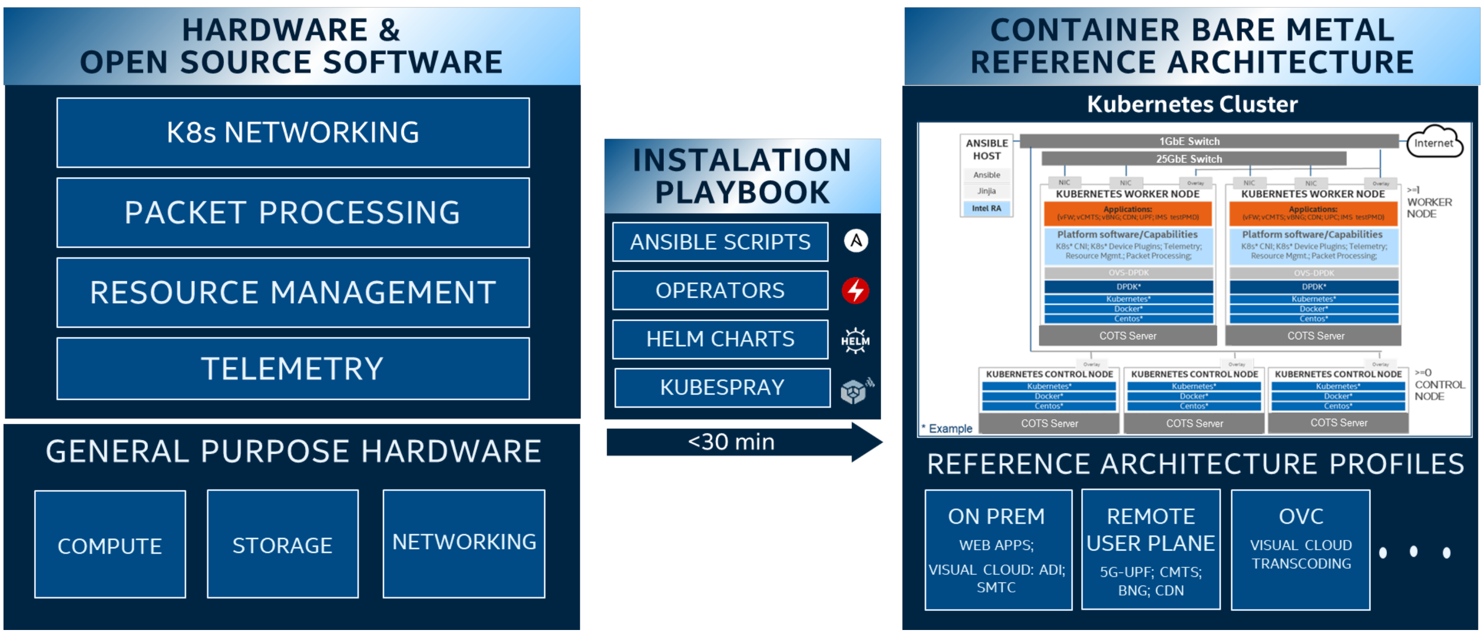

Figure 3 - Container Bare-metal Reference Architecture

A cloud-native software stack such as this is key to preparing the cable access network for mobile xhaul and FMC, as both require new functionality to be added. The network function adaption needed for these new DOCSIS use-cases are greatly simplified by having a modular extensible cloud-native vCMTS software architecture with rich telemetry and observability already in place.

A reference vCMTS data-plane implementation based on DPDK as shown in Figure 4 below has also been provided by Intel (see reference [6] for details). Such an implementation demonstrates the software performance capability of this key component of the cable access network on a cloud-native platform. DPDK forms the foundation for any high-speed packet processing software. And for this reference DOCSIS MAC data-plane, 80% of the code comes from DPDK, with each function in the DOCSIS MAC pipeline leveraging existing CPU-optimized DPDK library functions.

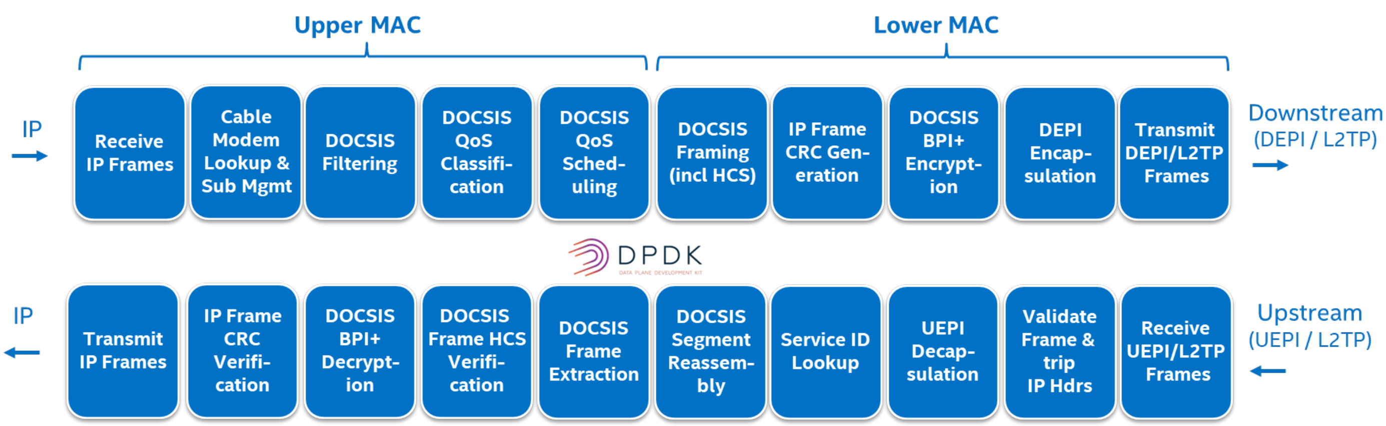

Figure 4 –vCMTS Data-plane based on DPDK

An example of a cloud-native software stack based on open-source components for deployment of vCMTS is shown below. Such a stack may be deployed in a lab environment for performance and TCO analysis using Intel's container BMRA (reference [4]) and vCMTS reference data-plane (reference [6]).

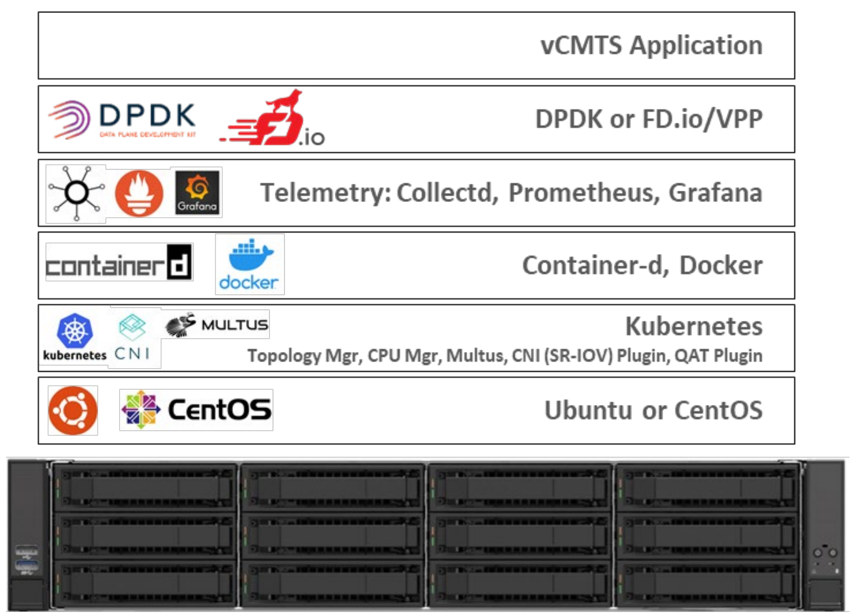

COTS Server with 100G NIC's

Figure 5 - Sample Software Stack for vCMTS Cloud-native Deployment

A cloud-native vCMTS runtime environment is shown in Figure 5 below. The vCMTS software architecture shown below is an example of how a monolithic CMTS may be decomposed into multiple Pod's (for example one per cable service-group) and each Pod decomposed into containers running distinct functional parts. This is based on the cloud concept of micro-services which enables optimum extensibility and resilience while simplifying the maintenance and upgradability of network function deployments.

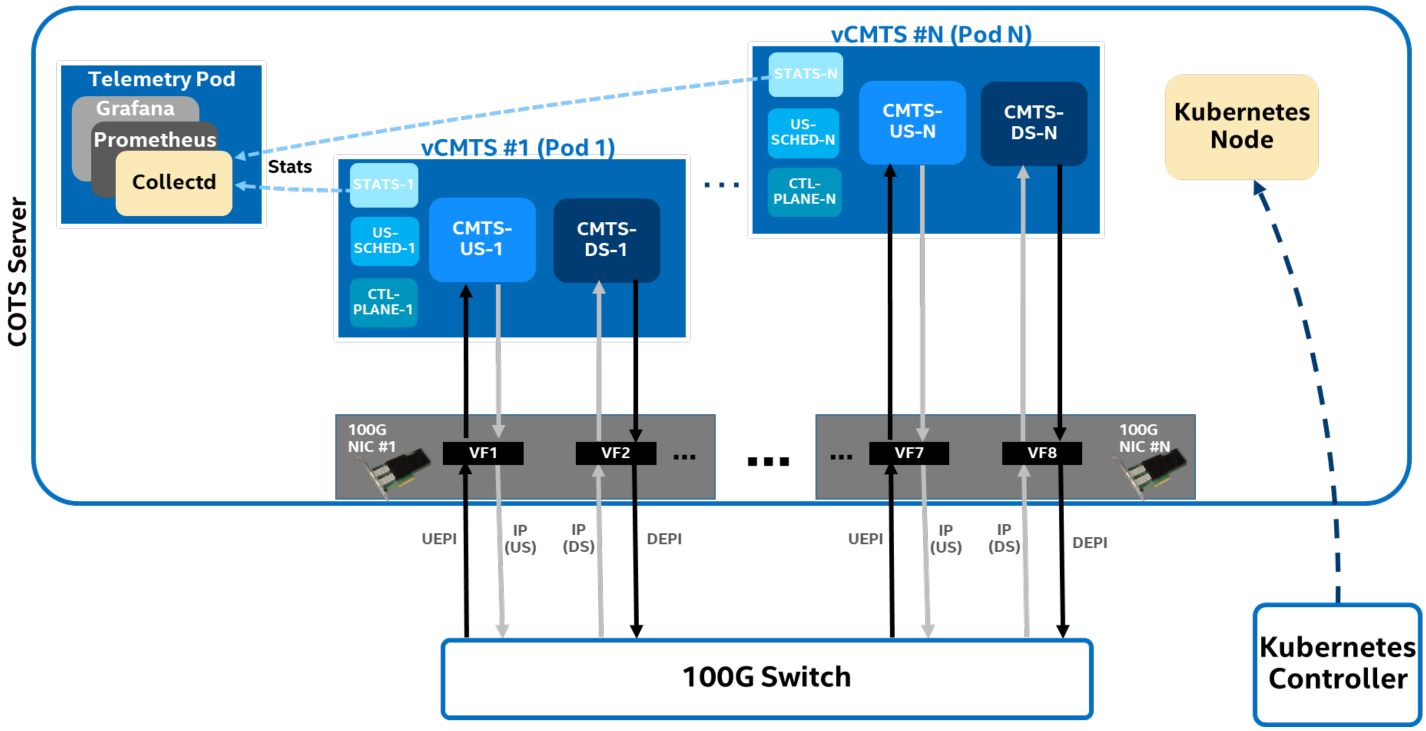

Figure 6 - Cloud-native vCMTS Simulation Environment

Reference [5] provides details of the performance capability of a cloud-native vCMTS data-plane. Empirical performance measurements such as shown in Figure 6, taken using the Intel vCMTS reference data-plane, prove that a cloud-native vCMTS is more than capable of achieving the throughput required for future DOCSIS 3.1 and 4.0 capabilities, and indeed to support transport of 5G small-cell traffic.

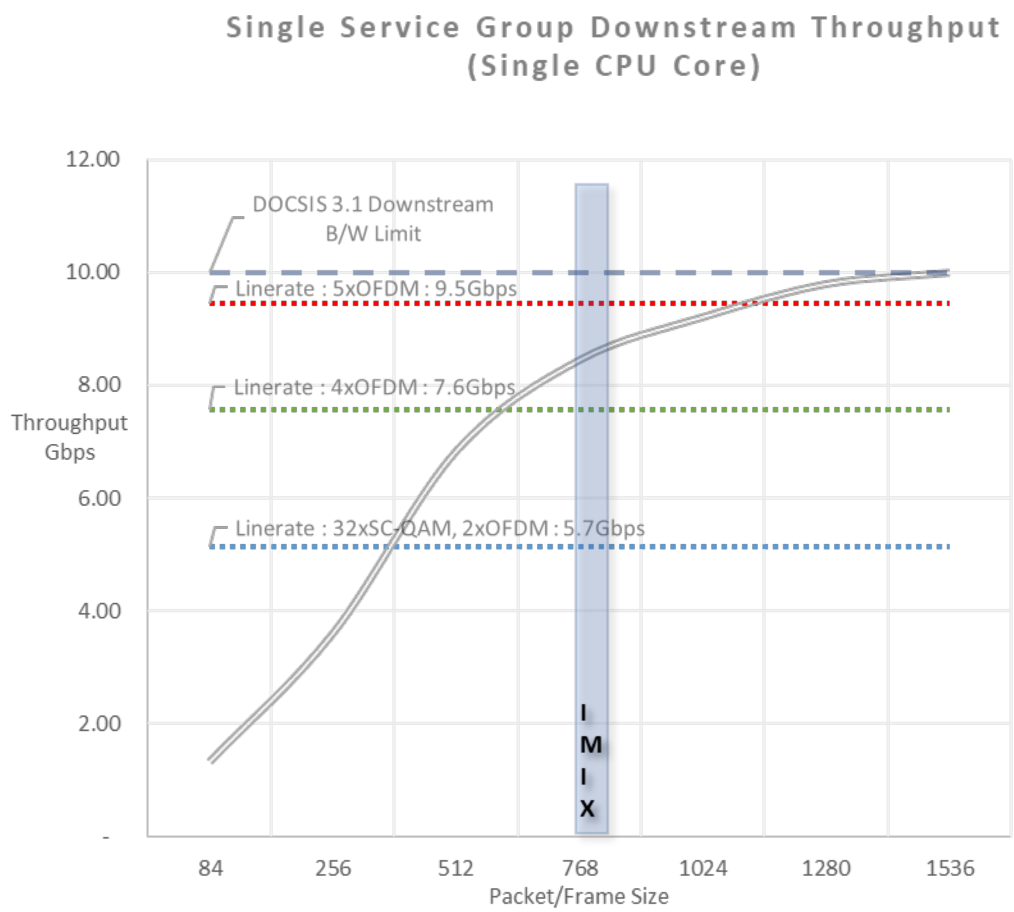

Figure 7 –vCMTS Data-plane throughput capability of a single CPU core

Scalable performance is a key aspect of vCMTS that directly paves the way for 5G xhaul and FMC use-cases. Each have their own particular performance requirements which will be covered in the respective sections that follow.

3. Extending vCMTS to support 5G Xhaul

Multiple systems operators (MSOs) are rolling out 5G deployments for a number of different business drivers and using a number of different convergence approaches. Operators like Shaw Communications and Videotron have acquired or built wireless assets to become mobile network operators (MNO). Charter, Comcast and Cox have entered into mobile virtual network operator (MVNO) agreements with MNOs such as Verizon and T-Mobile. In an MVNO agreement, the MVNO pays the MNO a fee to carry subscriber traffic over the MNO network. Still other operators like Vodafone began as an MNO and later acquired fixed network assets.

One common element across these operators is that there is a certain cost associated with carrying subscriber traffic and there is a goal to minimize that cost, thereby increasing operating margin. Minimizing cost to a large extent means leveraging existing assets to their full extent. A key existing asset in the cable industry of course is the hybrid fiber coax (HFC) network. Charter and Comcast plan to use the HFC network to build a network of 5G small cells to offload MVNO traffic, thereby reducing service cost. They and other operators will also market their HFC network to carry traffic for other operators planning to deploy their own small cell networks.

Three possible 5G functional split options are shown in Figure 7 below. Throughput and latency requirements placed on the HFC network will vary depending upon the split. These requirements are categorized as fronthaul, midhaul and backhaul with fronthaul having the most stringent bit-rate and latency constraints and backhaul the least.

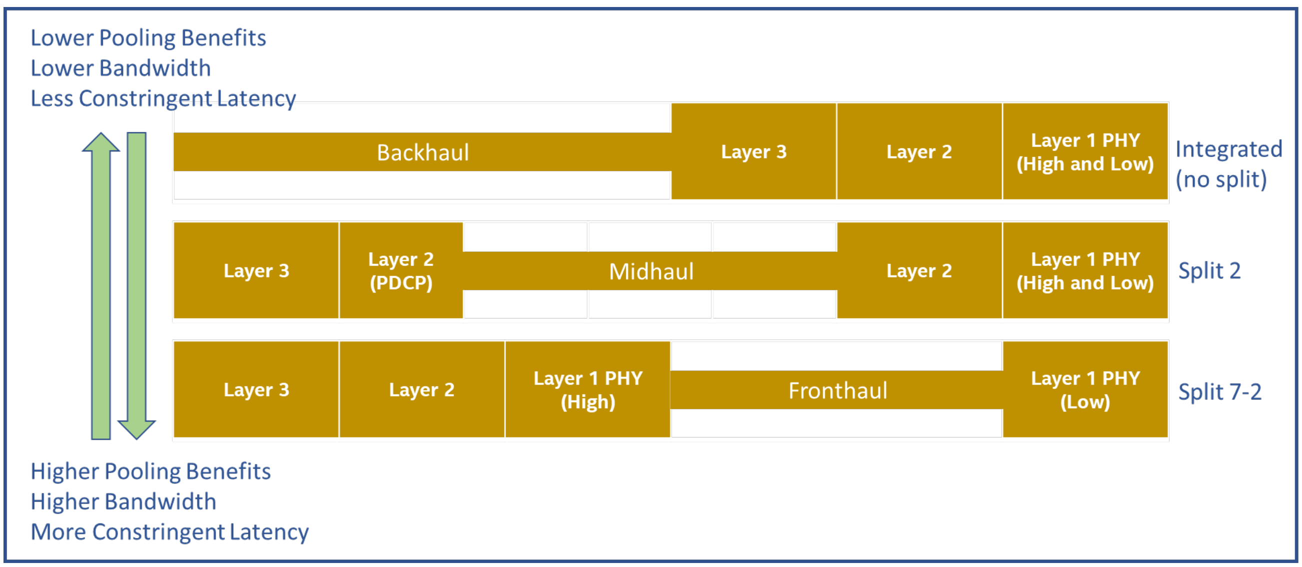

Figure 8 – Possible 5G Functional Split Options

Because HFC networks reach over 90% of North America and many parts of Europe, MSOs are uniquely positioned to quickly and cost-effectively roll out 5G if the HFC network can meet these latency and bit-rate requirements. Unfortunately, a 'plain vanilla' DOCSIS 3.1 network is hard pressed to meet all but basic mobile backhaul requirements.

But several recent technology advancements will greatly enhance the ability of the HFC network to meet 5G xhaul requirements - namely:

- The recent release of the DOCSIS 4.0 standard enabling multi-gigabit capacity in both uplink and downlink
- Distributed access architecture (DAA) which enables a software-based DOCSIS MAC to be deployed at the Cable head-end or remote HFC node
- The development of flexible MAC architecture (FMA) which enables flexible deployment of DAA options
- Low latency DOCSIS (LLD) which provides latency-sensitive traffic management
- Low latency xHaul (LLX) which streamlines upstream scheduling across 5G and DOCSIS schedulers, further reducing latency
- Finalization of the generic access platform (GAP) hybrid fiber coax (HFC) node specification which enables operators to deploy modern CPUs in nodes to implement evolving access protocols and optimization techniques in software much more cost effectively

In the remainder of this section, we will focus on LLX and GAP to showcase how implementing DOCSIS in software (i.e. implementing vCMTS) paves the way for 5G xhaul.

3.1. How to achieve low latency for 5G xhaul over DOCSIS

Figure 8 below shows how a DOCSIS access network may be used as a transport for 5G backhaul, midhaul and fronthaul cases. This is an example of infrastructure convergence which includes the transport network as well as common hardware and software.

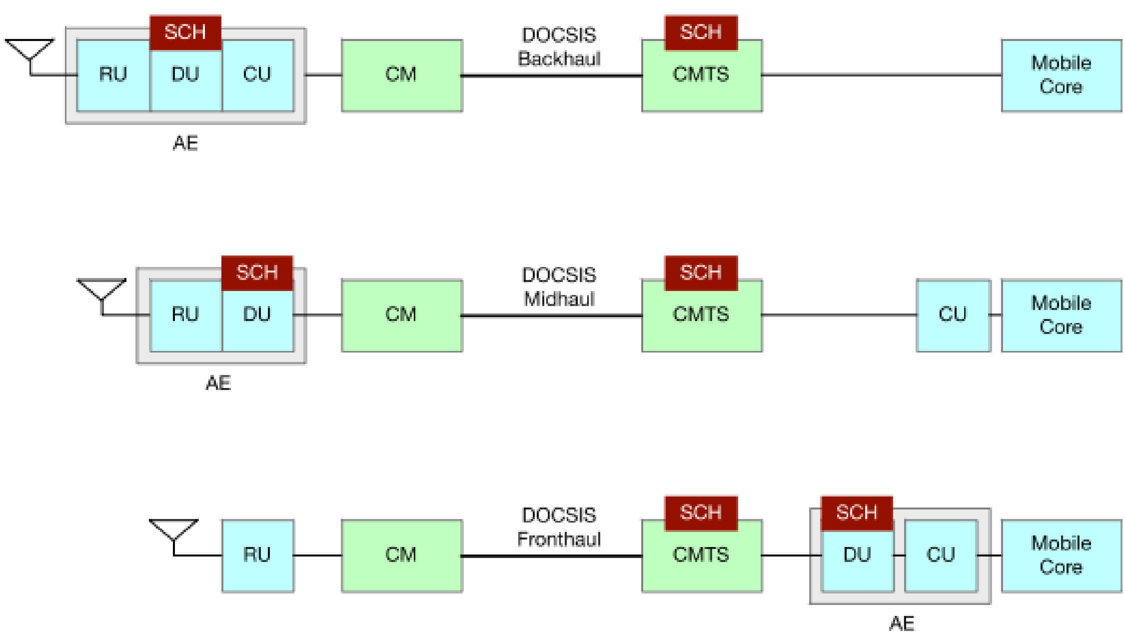

Source: CableLabs, "Low Latency Mobile xhaul over DOCSIS Technology"

Figure 9 – Access Entity Definition for Mobile Xhaul

As both DOCSIS and 5G are shared-medium networks, an upstream scheduler (SCH in the diagram above) must be present to ensure that traffic from subscribers does not collide. These respective schedulers introduce latency because a subscriber must ask and wait for permission to transmit through a request-grant message sequence. Without any optimization, the delay is compounded since each network has its own scheduler cycle.

Through use of LLX (as specified in reference [8]) the DOCSIS upstream bandwidth grant request-response delay is hidden behind the equivalent upstream traffic scheduling function for wireless access as shown below (noting that while the diagram shows LTE, the same principles apply 5G)

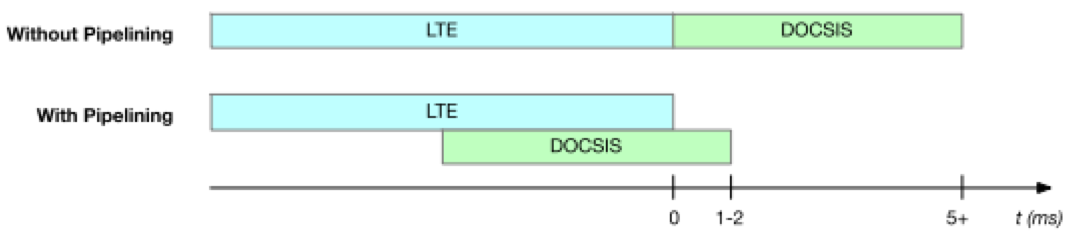

Source: CableLabs, "Low Latency Mobile Xhaul over DOCSIS Technology"

Figure 10 – Hiding DOCSIS Upstream Sheduling Latency behind Mobile

This is achieved by adding the bandwidth report (BWR) capability to the 5G DU and the DOCSIS Upstream scheduler in order to pre-determine the upstream grants needed on the DOCSIS network for Mobile xhaul.

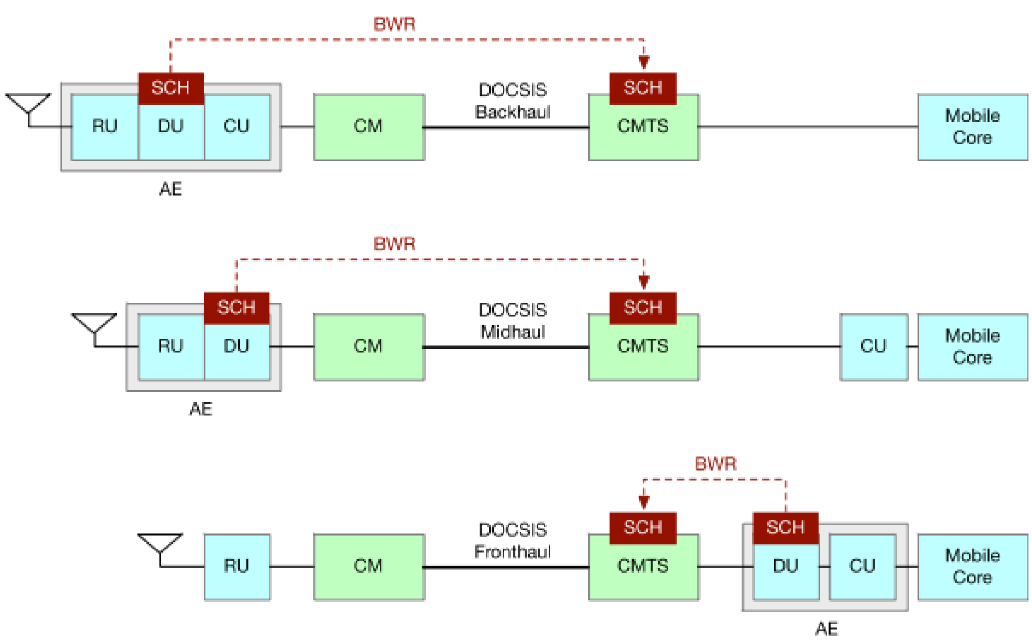

Source: CableLabs, "Low Latency Mobile Xhaul over DOCSIS Technology"

Figure 11 – BWR Transits Across Xhaul

As shown below, a cloud-native vCMTS software architecture can be relatively easily adapted to support the LLX BWR feature required for low-latency transport of 5G small-cell traffic. This is because the upstream scheduler is simply another algorithm implemented in software. As such, only software in the vCMTS system needs to be updated to implement LLX – and specifically only the software algorithm that implements the scheduler. Contrast that with a proprietary CMTS consisting of several sub-systems engineered to a specific set of requirements. Any change in the target requirements means that each sub-system needs to be carefully evaluated and re-tuned – which could include updates to proprietary hardware – a long and expensive undertaking.

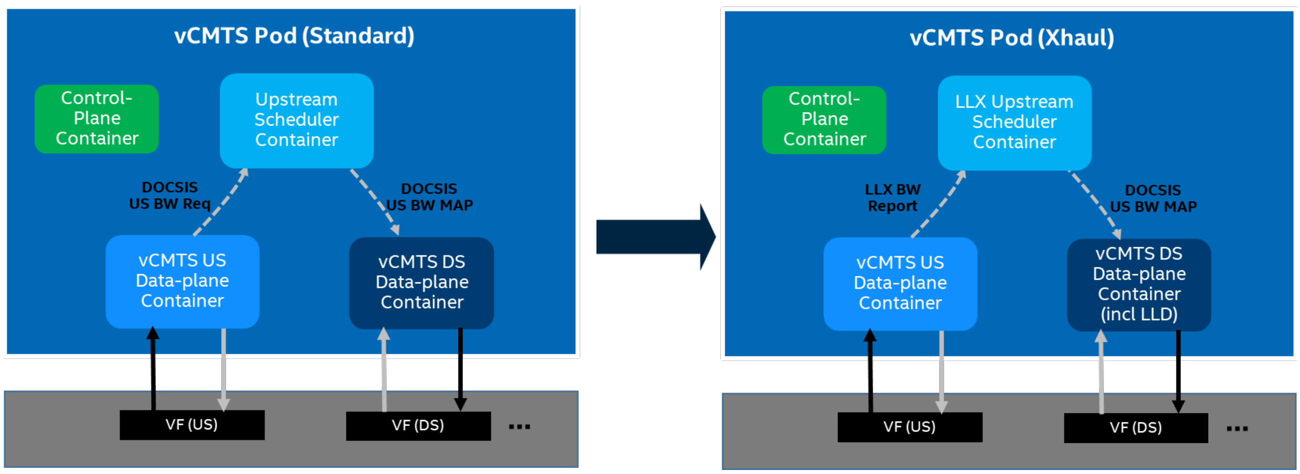

Figure 12 – Cloud-native Software Architectiture for LLX

But does implementing vCMTS (i.e. DOCSIS processing in software) actually introduce significant latency?

It can be seen in Figure 12 below from a CPU cycle count breakdown of the vCMTS upstream and downstream data-plane pipelines in software that it is possible to achieve ultra-low packet-processing latency in software with a single core run-to-completion pipeline based on a high-performance packet-processing frame-work such as DPDK or FD.io/VPP.

Considering average CPU cycles per packet for upstream and downstream vCMTS data-plane pipelines, based on measurement using the Intel vCMTS reference data-plane (reference [6]) with a 2.2 GHz CPU clock, software can perform DOCSIS MAC processing on upstream and downstream packets at an average of 615 nano-seconds and 967 nano-seconds respectively.

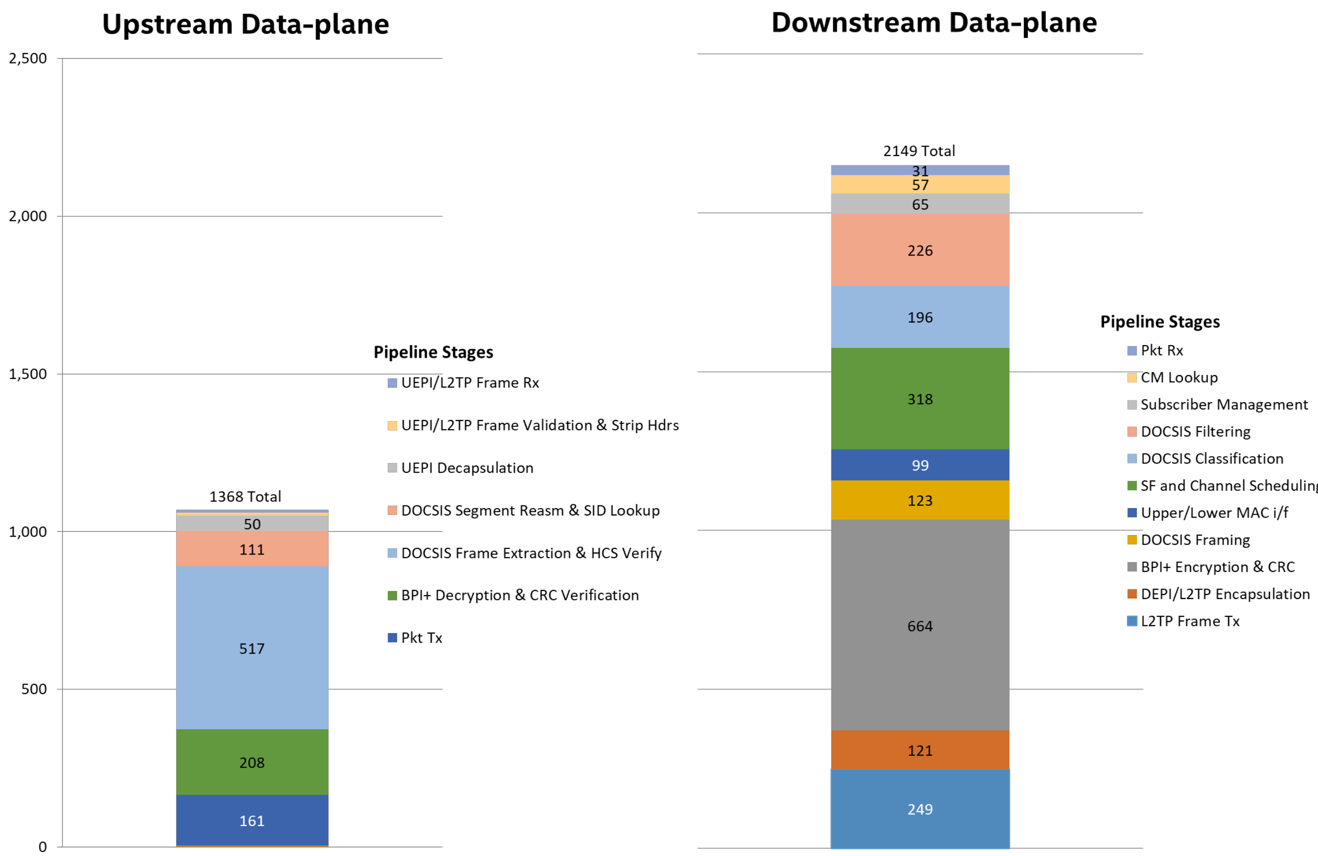

Figure 13 – vCMTS data-plane CPU cycle-count breakdown

However, latency is also added by packet buffering which occurs at three points in a DOCSIS MAC data-plane software pipeline - NIC Receive, Transmit and Downstream QoS scheduling. However, based on latency analysis using the Intel vCMTS reference data-plane, an optimally tuned system, can achieve sub 100 micro-second average latency per frame with a DOCSIS MAC data-plane implementation in software.

Low latency DOCSIS (LLD) functionality as described in a white-paper from CableLabs (reference [9]) may be used to reduce the impact of this by extending the standard Downstream Service-Flow QoS scheduling component to prioritize latency sensitive flows such as 5G signaling when vCMTS is being used for 5G xhaul over DOCSIS. LLD will most likely be needed to achieve the ultra low latency required for 5G fronthaul over DOCSIS.

The combination of an ultra-low-CPU-cycle software pipeline and the low-latency extensions to the DOCSIS standard described earlier should enable DOCSIS to meet the latency requirements for 5G backhaul and it should even be possible to meet the more stringent latency requirements for midhaul by combining LLX with proactive upstream scheduling.

Further work is needed to prove that DOCSIS is sufficient for 5G fronthaul, which has latency requirements in the order of hundreds of microseconds. At this low level of latency, even the PHY layer latencies become critical and need addressing. As for the DOCSIS upstream schedular related latency, LLX could potentially be augmented with AI models for better traffic pattern prediction resulting in even smaller request-grant cycles and less grant waste. Perhaps the solution lies in a custom grant-request cycle just for small cells, or a combination. Whatever the future solution brings, vCMTS is just a software upgrade away from implementing that solution which makes vCMTS a logical first step in 5G over DOCSIS and network convergence.

Not only is vCMTS more easily adapted to new innovations, but it can be readily scaled down and distributed as needed.

3.2. How FMA and GAP can help with 5G xhaul scenarios

As shown below, the portability of software enables the same DOCSIS MAC software to be used for all three flexible MAC architecture (FMA) scenarios being specified by CableLabs (see reference [10]).

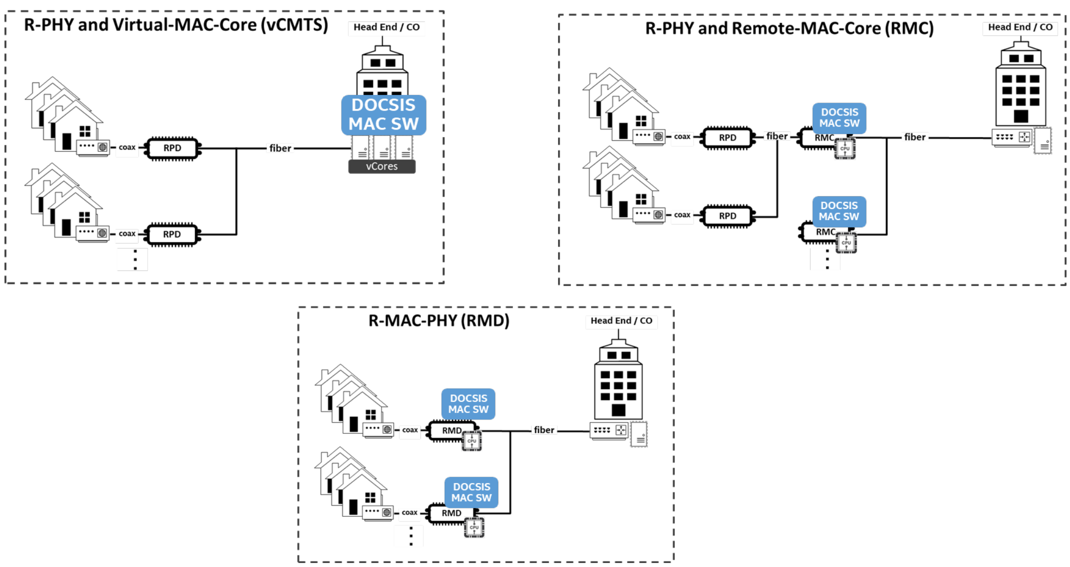

Figure 14 – Portability of Software for Flexible MAC Architecture Scenarios

Overlaying these three FMA options with the RU, DU, CU split options shown in Figure 7 produces a rather large matrix of deployment options for 5G over DOCSIS.

The industry is still in the process of evaluating these deployment options. While not all options ultimately will get deployed, it is most likely that at least several will get deployed because operators all have different assets and aspirations.

Analogous to the Intel vCMTS reference data-plane (see reference [6]), Intel has developed vRAN reference software (see reference [12]) for the upper PHY portion of a 5G base band unit (BBU). The industry eco-system adds layer 2 and layer 3 components to implement a full software based vRAN stack. This means both DOCSIS and RAN access implementations that are almost entirely software-based exist in the market today. In the case of vCMTS, the PHY is implemented in hardware and in the case of vRAN, the lower half of the PHY is implemented in hardware.

Because vCMTS and vRAN are both software based and can be scaled and moved as needed, implementing the matrix of FMA and RAN splits in combination becomes a much more tenable exercise. And as technology evolves both in the RAN and CMTS stacks to make them work better together, implementation of that new technology is a software upgrade away.

Of course, a suitable CPU host is required to run these software stacks. In the headend, the host will be a standard off the shelf rack mountable server with one or more x86 processors. But until recently a suitable CPU host in the HFC node did not exist.

The SCTE has recently released the generic access platform (GAP) specification. This specification defines the mechanical, electrical and thermal requirements of a modular, standardized strand mount HFC node enclosure. GAP includes the specification of a high speed PCIe/KR backplane which means that GAP is a suitable host for modern day x86 and other processors.

AOI, ATX, Charter, Cisco, Intel and Silicom demonstrated an early working prototype of a general purpose compute module at CableTec Expo in 2019 (see reference [13]). Coupled with an RPD module, the prototype was an early implementation of a software based RMD device. By adding the appropriate vRAN BBU components, a GAP RMD node becomes a key enabler for 5G over DOCSIS. While out of scope for this paper, by adding a pluggable PON SFP, a GAP node can similarly be used to drive 5G small cells over PON.

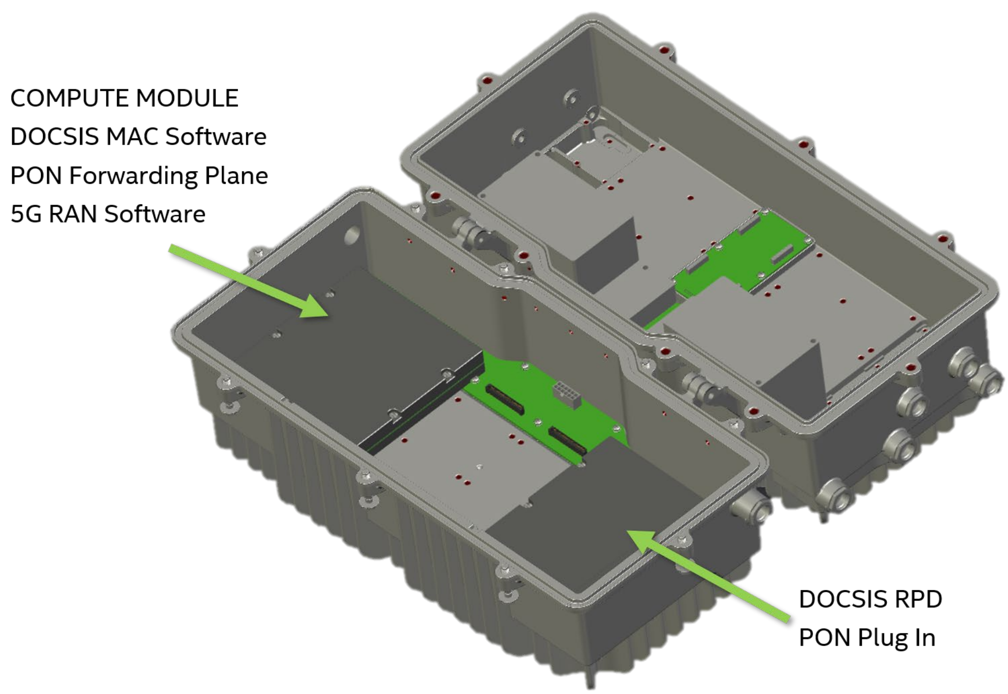

Figure 15 – GAP Configuration for RMD

Other Benefits of a software-based access network include enabling the highest degree of flexibility as it can be written once and re-targeted onto different platforms such as central-office servers or edge nodes.

But there are other more compelling reasons to move to software-based access that revolve around the availability of a tremendous amount of readily available software tools and packages.

Examples include:

- Telemetry frameworks and tools for faster fault detection and correction built around modern interfaces such as Yang models.
- Management and orchestration tools for easier and faster provisioning as well as 'self-service' customer provisioning.
- Network slicing applications to automate the creation of customer defined networks with customer specific characteristics around throughput, latency and quality of service.
- AI and Deep Learning frameworks to build sophisticated network optimization algorithms.

While it could be argued that all of the above are POSSIBLE to implement on networks with proprietary access nodes, software-based access nodes make the above much more PRACTICAL to implement. The reason is simple economics. Standard tools and utilities are developed for a broad range of markets and industries. This means that the cost of development is amortized over a much larger footprint and in the case of open source, that development cost is shared across several companies. Contrast that to a proprietary CMTS as an example where a single vendor historically had to develop 100% of these tools – including custom command line interface (CLI) protocols tasked with managing several internal sub systems such as ASIC based line cards. The latter is costly, time consuming to implement, costly to maintain and largely not re-usable.

4. Adapting vCMTS for Cable FMC

The goal of fixed mobile convergence (FMC) is to have a common core across wireless and wireline access networks. The term wireless and wireline convergence (WWC) is used by 3GPP to describe this convergence which may also be applicable to Cable access; for the purpose of this paper FMC will be used to describe this, whether for Cable or BBF Wireline.

As described in reference[16], the drivers for this include:

- Seamless multi-access connectivity and simplified service experience for end-users
- More effective use fixed and mobile assets for service providers
- Opportunity for service providers to properly integrate emerging technologies such as Connected Home, Virtualization/Cloud and High-speed Broadband Access

Figure 15 shows the desired end state for WWC/FMC – both wireline and wireless access served by a single 5G core.

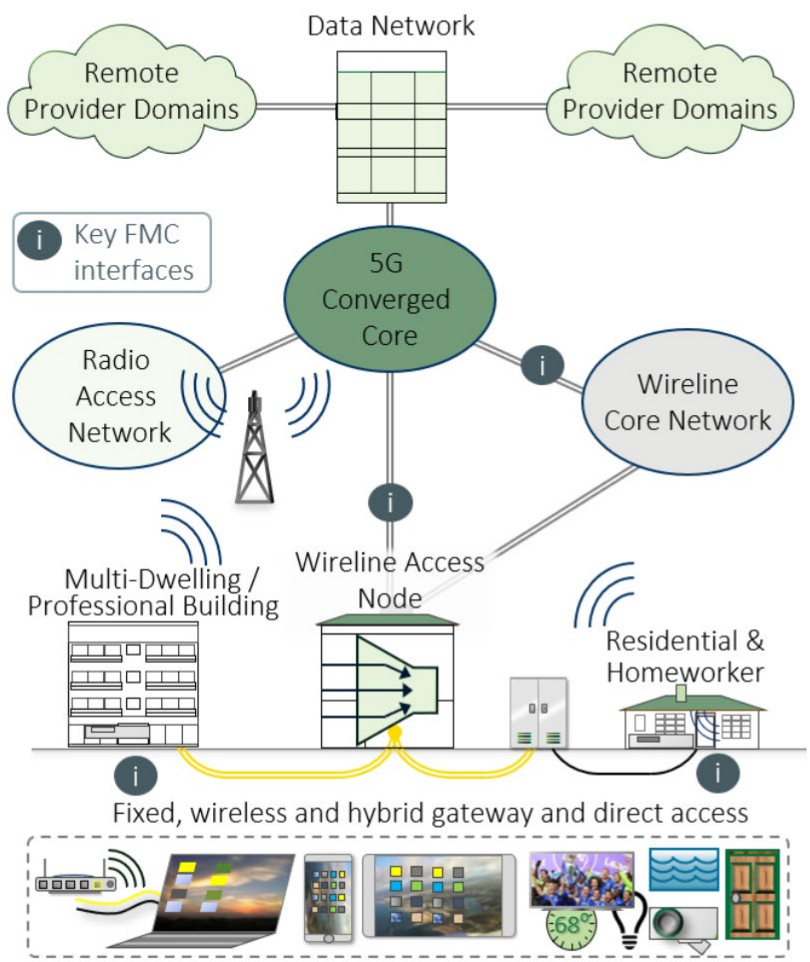

Source: Broadband Forum, "5G Fixed-Mobile Convergence – Marketing Report"

Figure 16 – 5G Fixed-Mobile Convergence

As stated in the 3GPP TS 23.316 specification for wireless and wireline access support (reference [17]) the key network function required for convergence is the Wireline AGF (W-AGF). This is a network function that mediates between the wireline access network and the 5G Core Network. In addition to supporting existing residential gateways (RGs), it is assumed that a new residential gateway (5G RG) will be introduced which must also be supported in the future (through the AGF). Both scenarios are shown below.

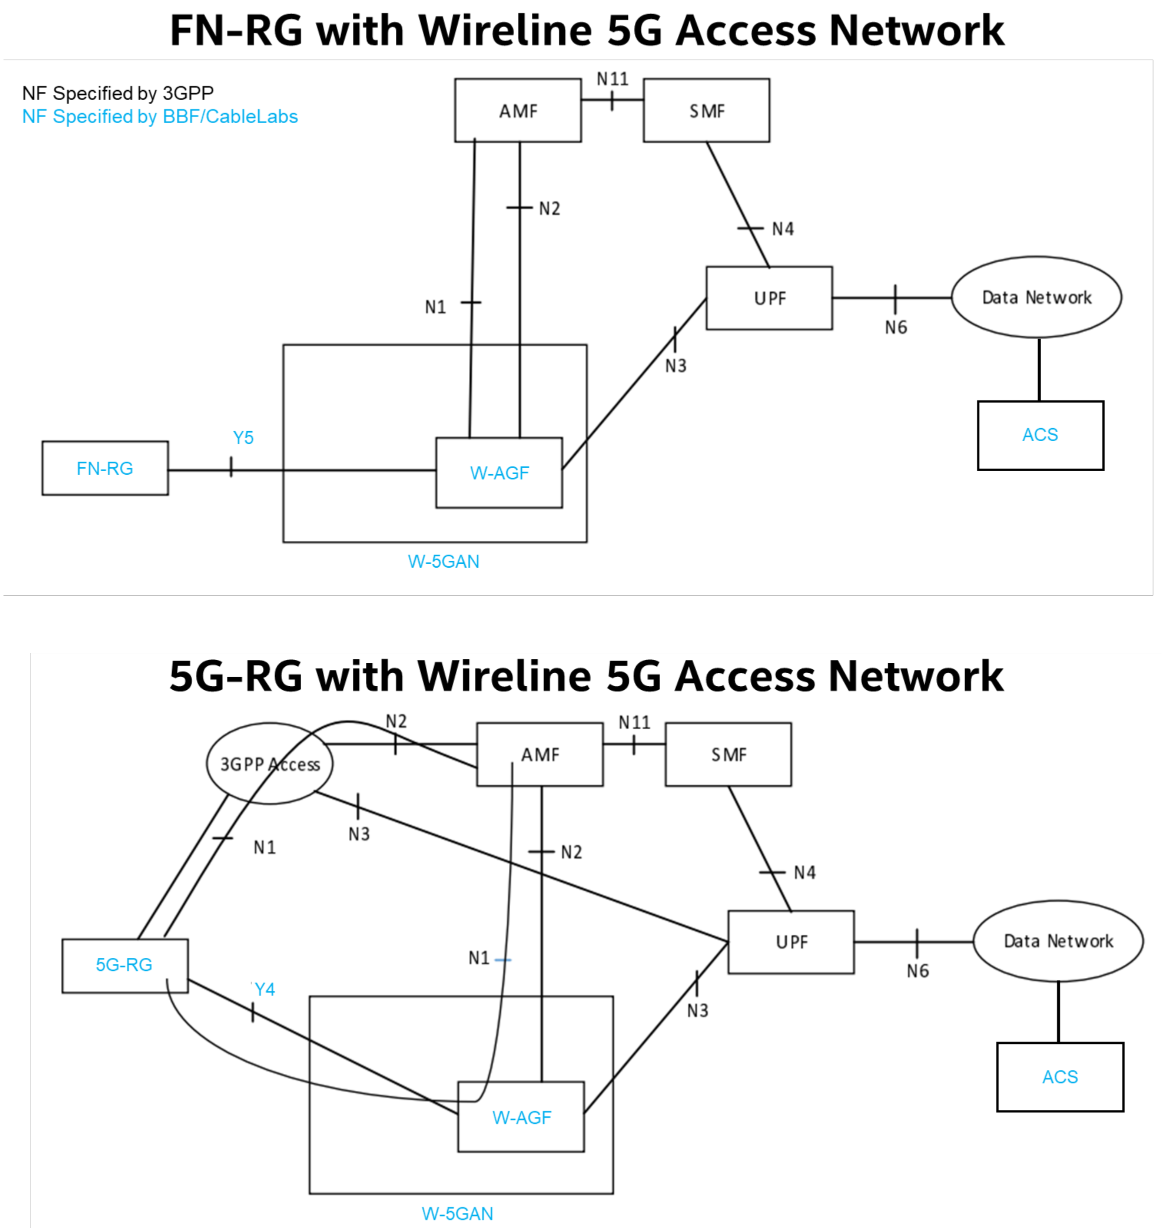

Source: 3GPP, "23.501 Rel 17: System architecture for the 5G System (5GS)"

Figure 17 – 5G Converged Core Architecture for FMC

While CableLabs was involved in the 3GPP Wireless Wireline Converged Core working group and published a technical report in relation to this (see reference [15]), the Broadband Forum (BBF) has made significantly more progress by providing FMC and AGF specifications for PON and DSL wireline access.

The requirements for initial convergence with existing RG's/Cable-modems should be similar to BBF: the integration of vCMTS functionality into a 5G-Core access gateway function (AGF) will be required to enable the use of the cable access network for FMC. At this point 5G and cable access networks become architecturally converged, the top level of convergence as shown in Figure 1.

Figure 17 below shows the five different FMC deployment scenarios/modes as defined by the BBF in TR-470.

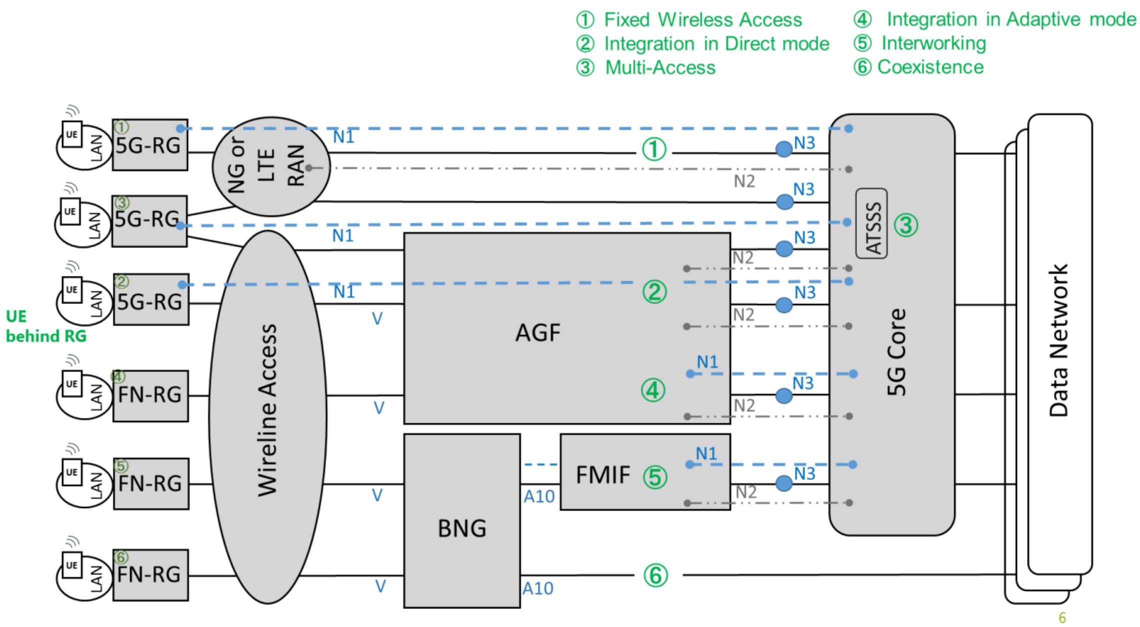

Source: Broadband Forum, "TR-470: 5G Wireless Wireline Convergence Architecture"

Figure 18 – FMC Deployment Scenarios

The scenario most likely to be used for initial BBF convergence is **Integration in Adaptive Mode** with FN-RG (i.e. existing residential-gateways/modems). It is proposed that a similar approach also makes most sense for Cable FMC.

In this case the RG (aka modem) is connected over the wireline access network and the AGF mediates layer 2 traffic with the 5G core network based on N2 and N3 interfaces. However, FN-RG does not support N1, so the AGF acts as end point of N1 on behalf of the FN-RG. The AGF is said to integrate access sessions in "adaptive mode".

In the case of "adaptive mode" for cable access the AGF will perform the same functions as the vCMTS, on the Cable access network (aka CIN) side, while providing the following additional functionality to integrate with a 5G-Core on the other side:

- Control Plane function:
  - map wireline information into 5GC information
  - map 5GC information into wireline information

- proxy N1 and N2 on behalf of the FN-RG and generate all the relevant N1 / N2 signalling for an FN-RG that has been identified / connected via the wireline access network
- map data-plane L3 connections from the Cable access network to PDU sessions in the 5GC and provide relevant signalling on the N1, N2, N3 interfaces
- manage access specific resources based on 5G QoS profiles

- User Plane functions:
  - for Uplink: encapsulate incoming FN-RG traffic with a GTP-U Header (mapping to the appropriate GTP-U Session), apply an appropriate DSCP value to the outer packet and forward it to the UPF
  - for Downlink: decapsulate incoming UPF traffic by removing the GTP-U Header, mapping the incoming traffic to the appropriate RG service-flow (using the GTP-U header, QoS flow indicator (QFI) and DSCP value if provided) and forward the extracted packet to the target FN-RG

Below is a high-level view of an AGF for DOCSIS, with control and user plane split, connecting cable residential gateways to the 5G core through the cable access network.

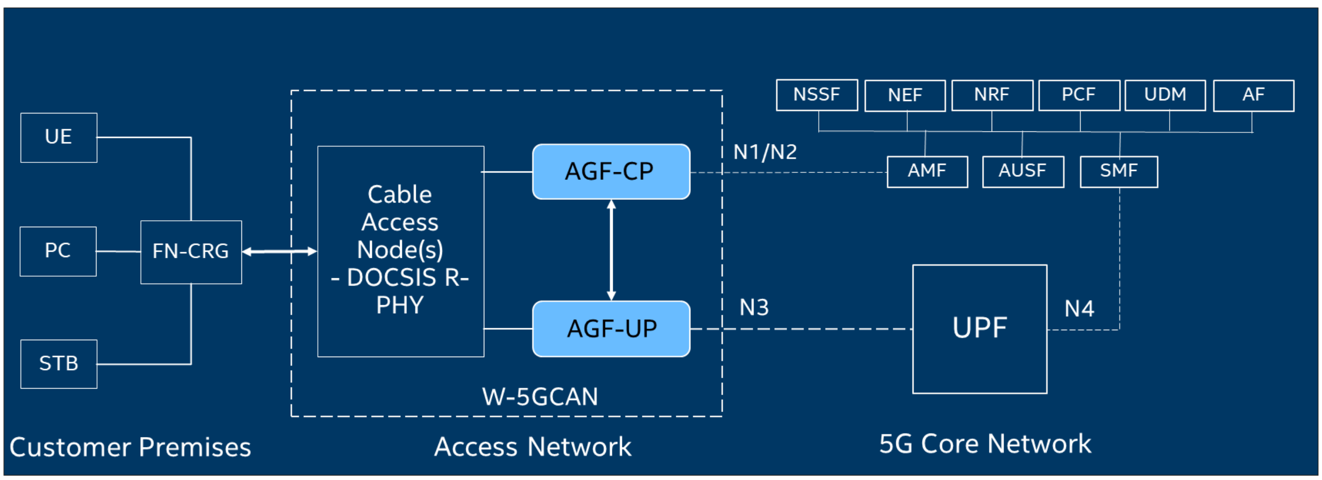

Figure 19 – Cable Residential Gateway Integration with 5G Core (in Adaptive Mode)

The existing software-based vCMTS network function will essentially need to re-purposed to become the vAGF network function as shown above. This will involve re-using vCMTS components such as DOCSIC MAC upstream and downstream data-planes, DOCSIS upstream scheduler and control-plane and integrating these with some new components for the 5G Core interface such as N2 control-plane, GTP-U encap/decap for 5G user-plane conversion and mapping of 5G QoS flows to DOCSIS service flows.

It is recommended that the Cable AGF adopt a control-plane/user-plane separation (CUPS) architecture which may require some re-architecting of an existing non-CUPS vCMTS architecture to deploy control-plane and data-plane functions in separate Kubernetes Pods on the same server or on separate servers in the same Kubernetes Cluster. Indeed, such an architecture has the added advantage of allowing control-plane and data-plane functions to be deployed at separate locations in the network. For example, it may be desirable to place the control-plane function in the cloud.

In terms of performance requirements for a Cable vAGF, data-plane throughput and latency are the critical KPI's that need to be considered. Upstream scheduling latency may be addressed in the same way as for 5G xhaul as shown in the previous section. The impact of adding additional processing stages for 5G user-plane adaption to data-pane also needs to be considered for FMC.

Below are the additional packet processing stages that would be added to a standard vCMTS data-plane to re-purpose it to a Cable vAGF data-plane processing pipeline. The additional stages are highlighted.

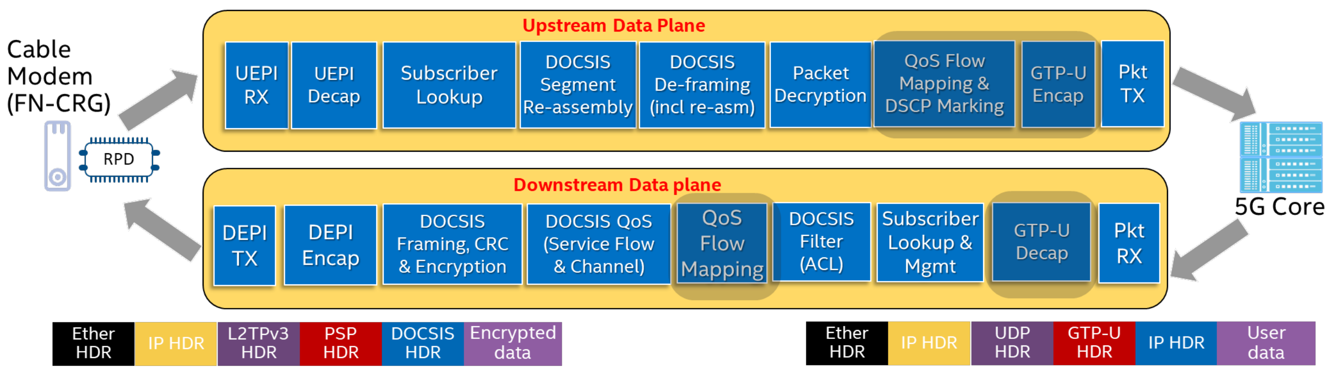

Figure 20 – Comparison of vCMTS and Cable vAGF data-pane processing pipeline

It is estimated that the additional processing shown above for 5G user-plane conversion adds 15 to 20% more CPU cycles per frame when compared to standard vCMTS data-plane processing. Given the performance capability of modern CPU architecture shown in Figure 6 and Figure 12 the required packet-processing performance for Cable access FMC is considered feasible in software.

It remains to be seen how broadly applicable AGF will be to cable operators; it may be a more desirable approach for mobile operators rather than cable operators. Furthermore, there are other viable approaches to fixed mobile convergence. For example, it could be said that MSOs such as Comcast have achieved this by deploying Wi-Fi Access Gateways in their cable network. This coupled with a high density of Wi-Fi hotspots allows subscribers to seamlessly switch between mobile and cable access networks.

However, it is clear that when starting with a cloud native vCMTS software architecture which is modular, extensible, and scalable, implementing FMC is greatly simplified versus starting with a proprietary hardware based solution.

As was shown for the xhaul case, software offers the flexibility to enable deployment at the pace of innovation and standardization of FMC while advancements in CPU architecture and the accompanying high-performance packet-processing frameworks ensure that throughput and latency requirements can also be satisfied. Such flexibility will be key to MSO's leveraging their Cable access networks to take advantage of advantage of such a major new use-case as 5G FMC.

5. Conclusion

As shown in this paper, the flexibility of a software based vCMTS makes it much easier to adapt for new 5G use-cases such as xhaul and FMC and to evolve over time.

It has also been shown that a software-based vCMTS is more than capable of meeting the throughput and latency requirements to transport 5G small-cell traffic by leveraging new low-latency DOCSIS methods. These should be used in conjunction with high-performance packet-processing frameworks such as DPDK and FD.io/VPP which are highly optimized for modern CPU architectures, which in turn are continuously improving gen-on-gen to meet ever-increasing network processing performance demands.

A software-centric architecture also has the benefit of leveraging open-source and the thousands of engineering person-hours of optimization and feature-development in mainstream projects such as Kubernetes, DPDK and FD.io/VPP. Additionally, there exist off the shelf tools for telemetry, management, orchestration, and AI which can be leveraged to improve network performance and resiliency.

Ultimately it is clear that starting with a modular, extensible, and scalable cloud-native software based vCMTS greatly reduces development effort and time-to-market for the two major new 5G use-cases of Xhaul and FMC as well as future-proofing for further standardization and innovation.

Abbreviations

3GPP	Third generation partnership project
5GC	5G core
BBU	Base band unit
BNG	Broadband network gateway
Bps	Bits per second
BWR	Bandwidth report
CM	Cable modem
CMTS	Cable modem termination system
CNF	Cloud native function
CO	Central office
COTS	Common-off-the-shelf
CP	Control plane
CPE	Customer premise equipment
CPRI	Common public radio interface
CU	Central unit
CUPS	Control and user plane separation
DAA	Distributed access architecture
DC	Data center
DL	Downlink
DOCSIS	Data over cable system interface specification
DP	Data plane
DPDK	Data plane development kit (open-source project)
DS	Downstream
DTP	DOCSIS Time Protocol
DU	Distributed unit
eCPRI	Enhanced CPRI
eMBB	Enhanced mobile broadband
eNB	eNodeB
EPC	Evolved packet core
FD.io/VPP	Fast Data I/O (open-source project)
FMA	Flexible MAC architecture
FMC	Fixed mobile convergence
FTTH	Fiber to the home
FWA	Fixed wireless access
gNB	gNodeB
HFC	Hybrid fiber-coaxial
LAN	Local area network
LLX	Low latency xhaul
LTE	Long term evolution
MIMO	Multiple in multiple out
MNO	Mobile network operator
ms	Millisecond
MSO	Multiple system operator
MVNO	Mobile virtual network operator
NFV	Network function virtualization
NGA	Next generation access

OLT	Optical line termination
Opex	Operating expense
O-RAN	Open RAN
O-RU	O-RAN RU
OS	Operating system
PON	Passive optical network
PGS	Proactive grant service
PTP	Precision time protocol
QoE	Quality of experience
QoS	Quality of service
RAN	Radio access network
RF	Radio frequency
RG	Residential gateway
RGW	Residential gateway
RMACPHY	Remote MAC and PHY
RMD	RMACPHY device
RPD	Remote PHY device
RPHY	Remote PHY
RTT	Round trip time
RU	Radio unit
SCTE	Society of Cable Telecommunications Engineers
SDN	Software defined network
SG	Service group
SP	Service provider
TCO	Total cost of ownership
UE	User equipment
UL	Uplink
UP	User plane
UPF	User plane function
US	Upstream
URLLC	Ultra-reliable and low-latency communications
vCCAP	Virtualized CCAP
vCMTS	Virtualized CMTS
vEPC	Virtualized EPC
VM	Virtual machine
VNF	Virtual network function
VPP	Vector packet processing (project)
vRAN	Virtualized RAN
WWC	Wireless and wireline convergence

Bibliography & References

[1] *Cable and Mobile Convergence: A Vision from the Cable Communities Around the World*, Cisco & CableLabs, https://www.nctatechnicalpapers.com/Paper/2020/2020-cable-and-mobile-convergence

[2] Data Plane Development Kit project, https://www.dpdk.org

[3] FD.io VPP project, https://fd.io/vppproject/vpptech

[4] *Container Bare-metal Reference Architecture,* Intel Corporation, https://networkbuilders.intel.com/solutionslibrary/container-bare-metal-for-2nd-3rd-generation-intel-xeon-scalable-processor

[5] *Maximizing vCMTS Data Plane Performance with 3rd Gen Intel® Xeon® Scalable Processor Architecture,* Intel Corporation, https://builders.intel.com/docs/networkbuilders/maximizing-vcmts-data-plane-performance-with-3rd-gen-intel-xeon-scalable-processor-architecture.pdf

[6] *Intel vCMTS Reference Data-plane project,* Intel Corporation, https://01.org/access-network-dataplanes

[7] *View on 5G Architecture*, 5G PPP Architecture Working Group, https://5g-ppp.eu/wp-content/uploads/2019/07/5G-PPP-5G-Architecture-White-Paper_v3.0_PublicConsultation.pdf

[8] *Low Latency Mobile Xhaul over DOCSIS Technology*, CM-SP-LLX, CableLabs, https://www.cablelabs.com/specifications/CM-SP-LLX

[9] *Low Latency DOCSIS: Technology Overview,* CableLabs, https://www.cablelabs.com/technologies/low-latency-docsis

[10] *Flexible MAC Architecture Specification,* CableLabs, https://www.cablelabs.com/specifications/CM-SP-FMA-SYS

[11] *Generic Access Platform(GAP),* SCTE, https://www.scte.org/generic-access-platform

[12] *Intel FlexRAN Project,* Intel Corporation, https://github.com/intel/FlexRAN

[13] *Can cable bridge the GAP?,* LightReading Article, https://www.lightreading.com/cable/can-cable-bridge-the-gap-/d/d-id/754620

[14] *What is network convergence and why do we need it?*, BT (IEEE BMSB 2018 Key-note), https://www.mcg.upv.es/wp-content/uploads/2018/06/IEEE_BMSB_2018_Keynote_Day2_MariaCuevas_Public.pdf

[15] *5G Wireless Wireline Converged Core Architecture Technical Report,* CableLabs https://www.cablelabs.com/specifications/WR-TR-5WWC-ARCH

[16] *MR-427: 5G Fixed-Mobile Convergence Marketing Report*, Broadband Forum (BBF) https://www.broadband-forum.org/download/MR-427.pdf

[17] 23.316 *Rel 17*: Wireless and wireline convergence access support for the 5G System, 3GPP https://www.3gpp.org/ftp/Specs/archive/23_series/23.316/23316-h00.zip

[18] *23.501 Rel 17: System architecture for the 5G System (5GS),* 3GPP, https://www.3gpp.org/ftp/Specs/archive/23_series/23.501/23501-h00.zip

[19] *TR-470: 5G Wireless Wireline Convergence Architecture*, Broadband Forum (BBF), https://www.broadband-forum.org/technical/download/TR-470.pdf

