

Virtualization in Cable Networks

Use cases and design considerations

Table of contents

Introduction	3
Background	3
Five Cases Involving Virtualization.....	3
1. Distributed Access Architecture	3
2. Video on Demand and Network PVR	4
3. Content Distribution Networks	5
4. 5G	6
5. Multi-Access Edge Computing	8
Workload Placement Options.....	8
1. Regional and National Datacenters	9
2. Headend sites	9
3. Remote PHY	9
4. New Definition of the Edge	9
Framework Requirements and Common Platform	10
Conclusion.....	11
References	15

Introduction

The virtualization of software workloads to provide network functions is a concept that has arisen in multiple industries, including the cable Multi-System Operator (MSO) / Multichannel Video Programming Distributor (MVPD) community. Most initial implementations of Network Function Virtualization (NFV) have followed a pattern of centralized sites that correspond to regional serving area or national data centers. But driven by the need to bring services, resources, and intelligence deeper into the Hybrid Fiber Coax (HFC) plant in closer proximity to subscribers the industry is promoting virtualization deeper in the access network. A leading example would be Distributed Access Architecture (DAA), and the associated Remote PHY Devices (RPDs). But there are other use cases and architectures driving demand for edge network intelligence including Cloud DVR, CDNs, as well as mobility driven multi-access edge computing (MEC) and 5G; all of which bear consideration for deployments as well. Going forward, a platform strategy and framework for virtualization, which anticipates multiple applications and software driven technologies spanning access and centralized datacenters will enable operators to enable new revenue streams and drive operational efficiencies across their service portfolio.

Background

The Cable industry has considered virtualization in various forms for approximately a decade. In addition to leveraging web-based architectures to provide services for video subscribers – which generally depend upon virtualization - the industry has for several years been planning to apply virtualization to the broadband network through the Distributed Access Architecture (DAA) initiative.

The high-profile DAA initiative has now advanced into commercial implementations and deployments. At the same time, we can point to other areas where virtualization intersects existing operations, such as Video on Demand (VOD), network Personal Video Recordings (PVRs) and Content Distribution Networks (CDNs). Two additional developments originating outside the industry – 5G and Multi-access Edge Computing (MEC) – also merit attention for their potential applicability. After first reviewing these five cases (DAA, cDVR, CDN, 5G and MEC) we will share some strategic considerations about the growth of virtualization within the industry's evolving HFC networks; and then conclude with thoughts about how a second-wave virtualization framework that extends beyond large data centers can continue to drive efficiency and reduce complexity.

Five Cases Involving Virtualization

1. Distributed Access Architecture

This framework has roots in the industry's earlier modular headend architecture (MHA), which separated the Physical (PHY) downstream and DOCSIS Media Access Control (MAC) components. ¹ (See Figure 1.) Using a new digital link extending from the CCAP core to an RPD, DAA enables the distribution and virtualization of network functions. In the new model, the CCAP core could reside at the headend or hub; and the RPD at a hub or node. Officially known by CableLabs® as the Distributed CCAP Architecture (DCA), DAA aligns with other industry initiatives, including Full Duplex DOCSIS (FDX), Extended Spectrum DOCSIS (ESD), and the extension of fiber to points deeper into the network. CableLabs also associates it with higher spectral efficiency, Gigabit services, and increased access network performance and a much smaller footprint. By transforming the CCAP (or CMTS) from purpose-built hardware into software that could potentially run in a data center on COTS equipment in a private cloud, DAA becomes a classic case of a virtualized and software-defined infrastructure.

1. Data-Over-Cable Service Interface Specifications, DCA – MHA v2, Remote PHY Specification, CM-SP-R-PHY-I12-1903307, March 7, 2019

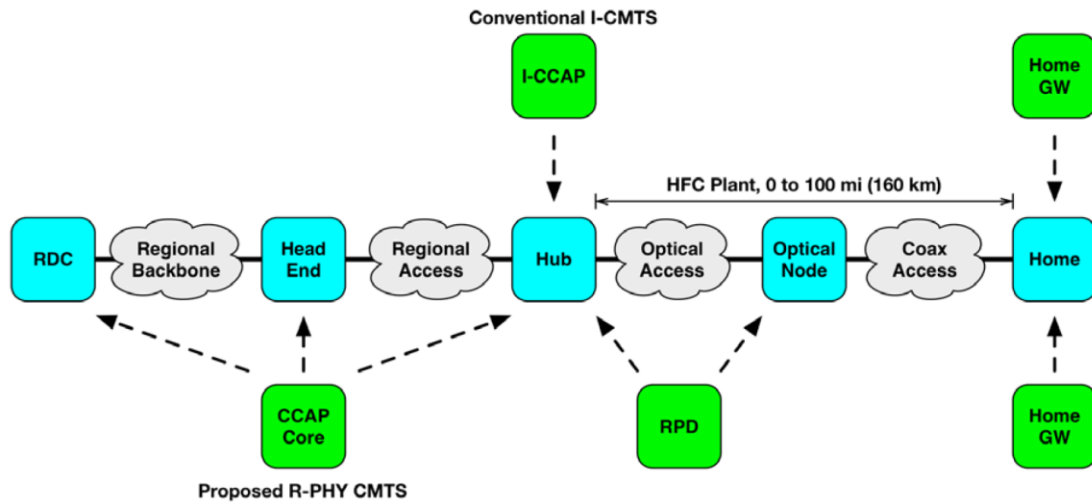


Figure 1 – Remote PHY System Source: CableLabs

2. Video on Demand and Network PVR

The architecture to deliver localized or personalized content through Digital Program Insertion (DPI) and Video on Demand (VOD) systems is generally implemented through a multitude of software-based elements, a number of which are naturally distributed to the network edge for functional reasons. See Figure 2.

Likewise, the migration of video content storage from purpose-built customer premises equipment (CPE) in the home to cloud infrastructure is another example of a service innovation enabled by virtualization. The network-based personal video recorder (or “cloud DVR”) approach facilitates simultaneous, efficient availability of private and catalog content in a multi-device and multi-network consumption model. Although control plane functions like schedulers, license servers, program guides and the like use web technology and interfaces that are readily centralized, the data plane and network traffic requirements of origin servers, packagers, transcoders scale quickly according to subscriber demand.

Higher resolution, rate, and quality media formats now in use for streaming (UHD, HDR) and the associated codecs (H.265) call for an increased proportion of resources per active subscriber. These content types, coupled with growth of non-streaming and less predictable services also give rise to characteristically more variable compute, storage, and network demands. Thus, the ability to dynamically allocate – and relocate – these resources for software workloads in response to network demand is highly desirable.

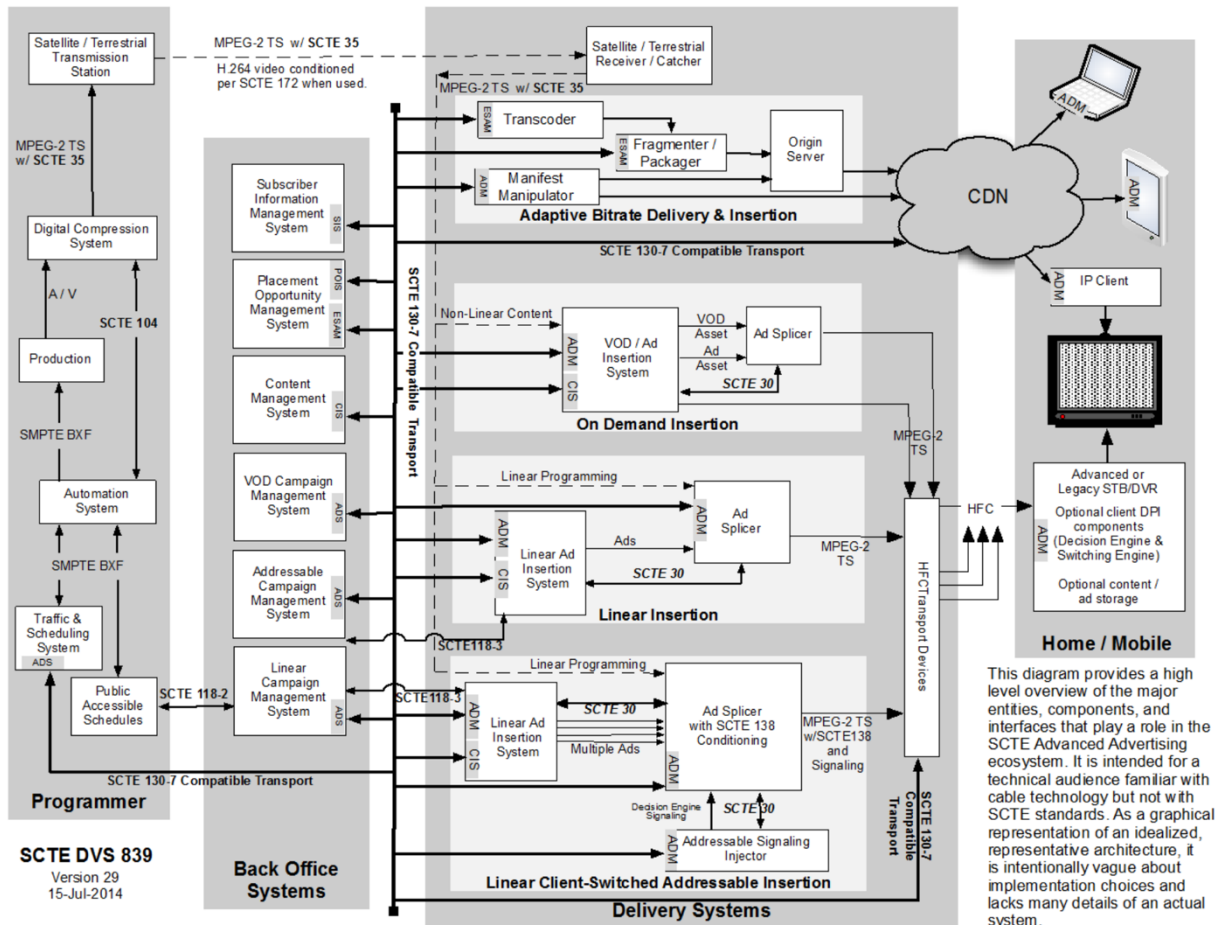


Figure 2 – Architecture for Advertising Supported VOD and Linear Video

3. Content Distribution Networks

Content distribution networks (CDNs) are another application category that can leverage virtualization, and a runtime environment at the network edge. MSOs have long-established patterns of working with or deploying CDNs, which provide web, media, OTT streaming content delivery. The CDN reference model created by the Apache Software Foundation highlights a cluster of traffic control functions connected to each other and related servers on the data plane. (See Figure 3.) In a web centric application paradigm, RESTful APIs are the prevalent way for applications to interact and access resources; they also facilitate network portability and remote interconnection between these systems. However, shield cache tiers (or, “traffic servers” according to this model) are typically positioned in proximity to demand sources to benefit network efficiency and scaling. These data plane candidates are natural candidates for distribution to the network edge.

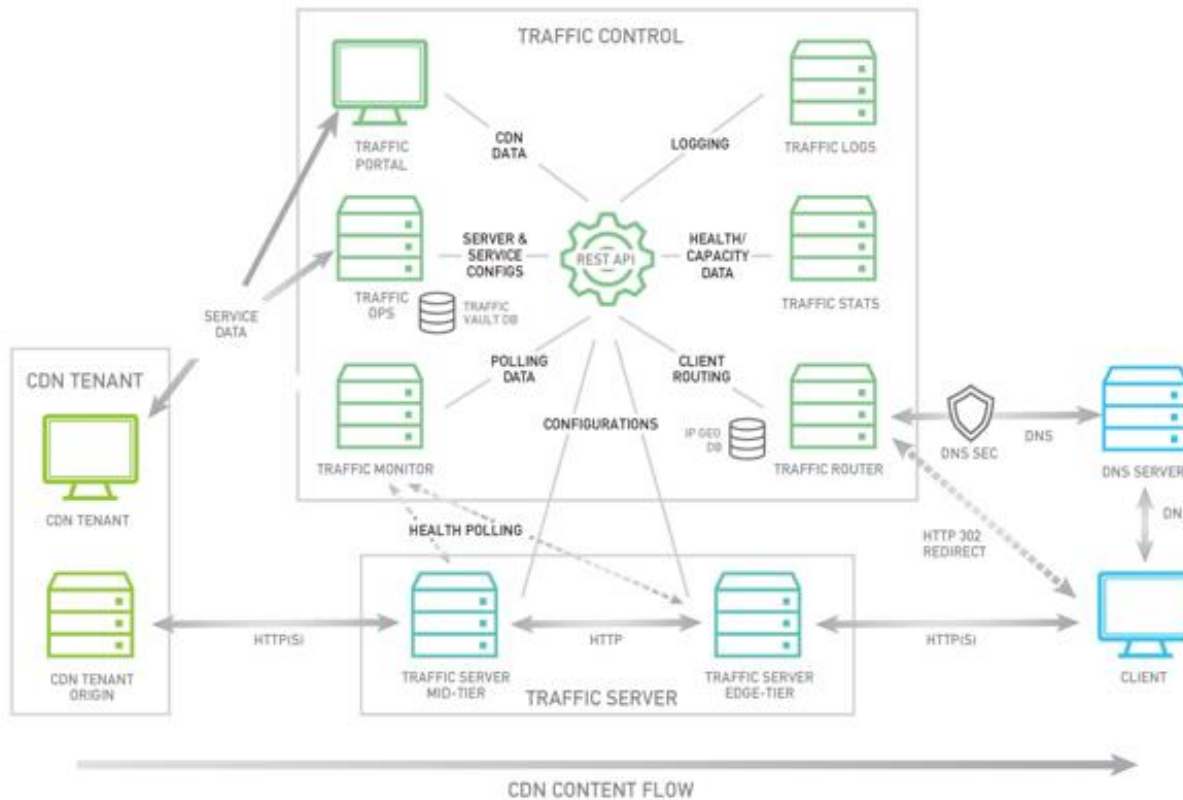


Figure 3 – CDN Reference Model Source: Apache Software Foundation

4. 5G

Using mobile networks and technology for service delivery within and outside the home has been a strategic focus for a number of MSOs^{2,3}, supporting new applications and modes of consumption. The 5G initiative is the next frontier of focus for the mobile industry building upon the mobile broadband framework provided by 4G and LTE. This technology and standards regime will provide new capabilities for ultra-reliable low latency communication, enhanced mobile broadband, and enhanced machine type communication for IoT, supporting both consumers and enterprise use cases. The question raised here is how, where and when to transition mobile virtual network operator (MVNO) or mobile network operator (MNO) operations to on-net solutions leveraging 5G. The new architecture associated with 5G differs from 3G and 4G, all the way down to the radio level and the new base station, or gNodeB (gNB). Given new frequency bands, air interface technology and propagation properties, as well as Radio Access Network (RAN) and network core architectures associated with 5G - architects and planners must determine where in the network footprint new user plane and radio units should be located. (See Figure 4-a.) The core and edge components of the 5G network are all expected to be IP-connected, software-defined and virtualized throughout... down to the radio baseband level. (See Figure 4-b.) So called “densification” of the gNB radio units is anticipated, due to higher frequency spectrum, higher levels of demand and concurrent use – that in turn will drive radio placement choices in the home gateway, NID, street or neighborhood level for small cell gNBs versus the macrocell-heavy network footprint that typifies the 3G and 4G footprint.

2. <https://newsroom.charter.com/press-releases/charter-launches-spectrum-mobile-a-smarter-network-designed-for-the-future>

3. <https://corporate.comcast.com/news-information/news-feed/comcast-xfinity-mobile>

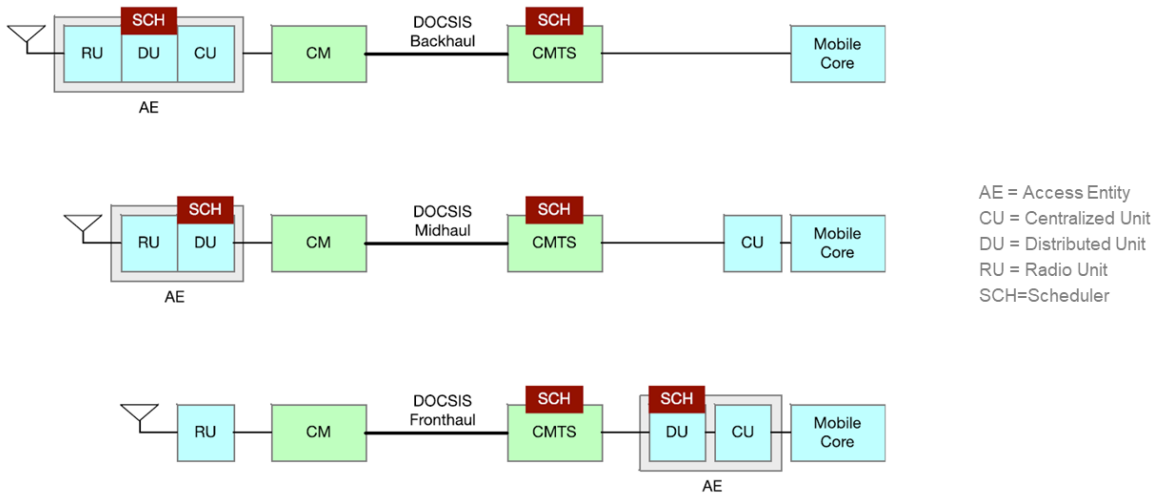


Figure 4-a – Virtualized Radio Access Network Topology Options Source: CableLabs

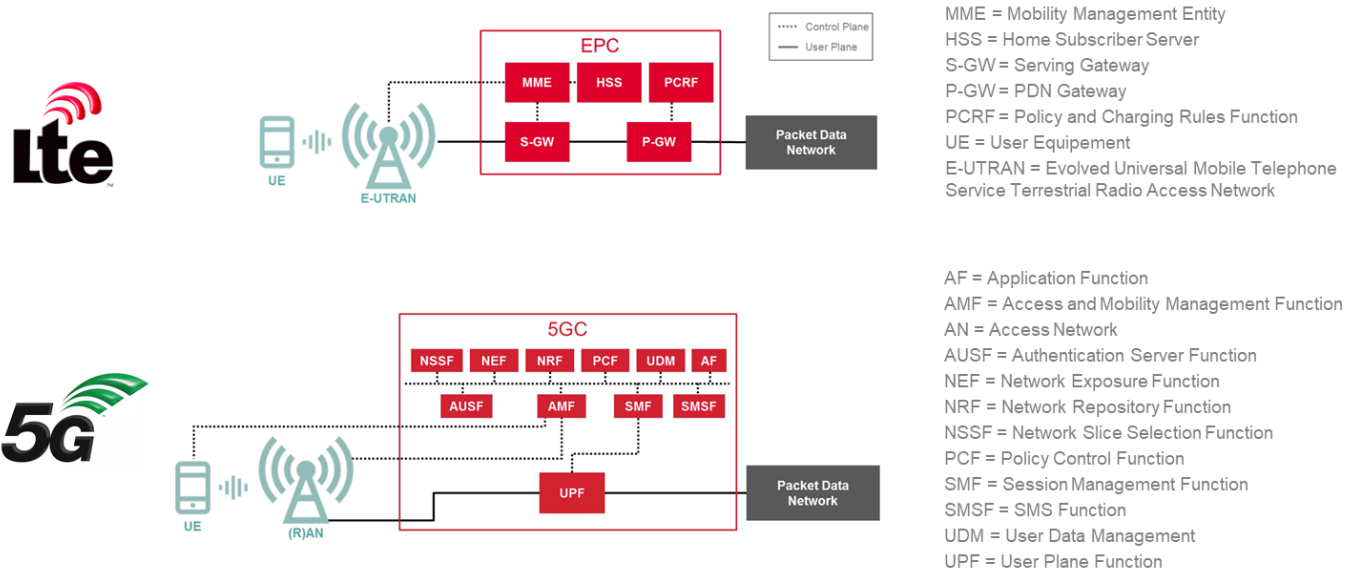


Figure 4-b – 5G Functional Elements Source: GSMA

5. Multi-Access Edge Computing

The Multi-access Edge Computing (MEC) model, established by ETSI, provides a framework for the distribution and coordination of edge computing intelligence (in the “MEC host”) to support software based network functions (“MEC applications”).⁴ Although originally conceived in anticipation of mobile applications and use cases, MEC acknowledges the notion of a “multi-access” paradigm, using varying access and transport technologies.

A MEC platform provides bandwidth management and prioritization, location, addresses the question of how to do edge computing and run workloads supporting mobile, but not necessarily over a mobile network.

While being developed by ETSI, this distributed computing model has implications for MSOs because it acknowledges the need for a platform strategy that pervades access technologies, and provides capabilities and services at the network edge. (See Figure 5.) This network edge could correspond with or complement the placement of Remote PHY Nodes (RPNs), Remote CCAP (MAC-PHY) elements, or Headend equipment, according to network and application requirements.

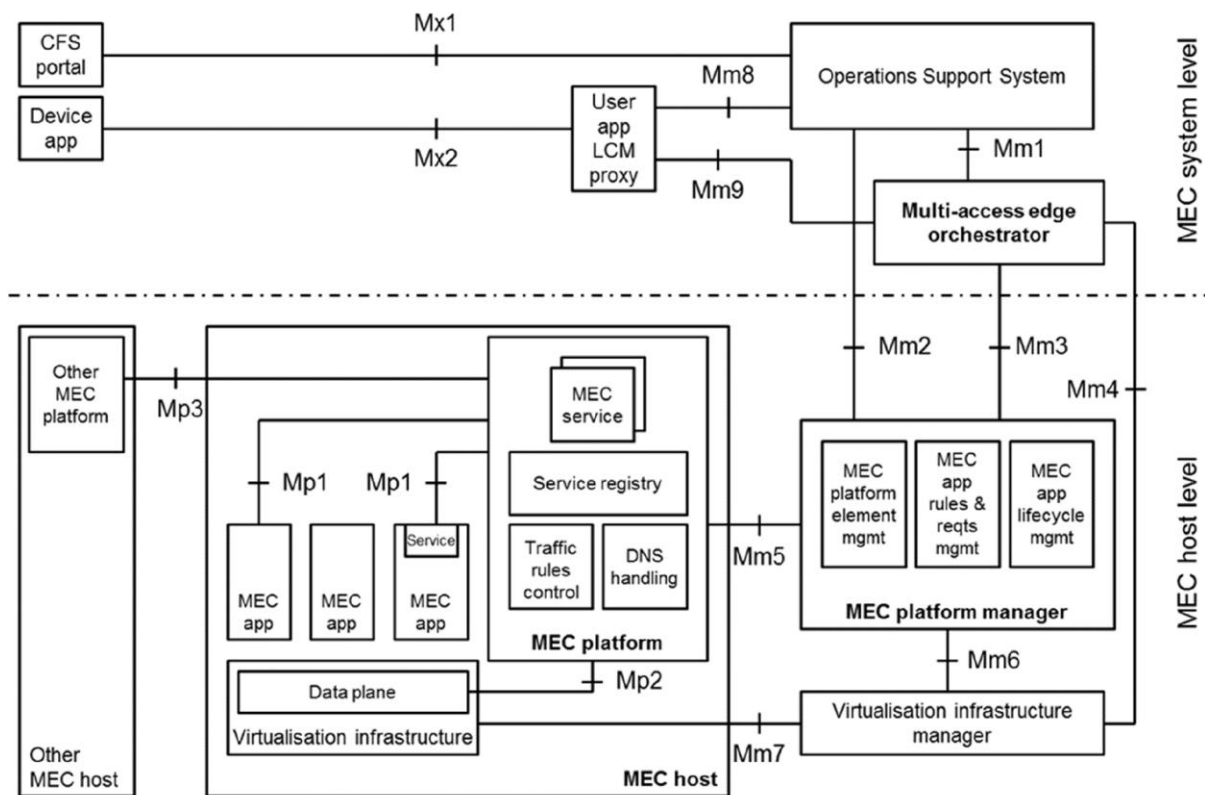


Figure 5 – ETSI Multi-access Edge

Workload Placement Options

4. Multi-access Edge Computing (MEC); Framework and Reference Architecture, ETSI GS MEC 003 V2.1.1 (2019-01)

In light of the use cases we have considered, engineering and planning questions naturally arise about what level within (or outside of) the HFC network is the correct one for a given functional element. In turn, if the majority of these elements are software-defined, and virtualized, a decision is required as to what framework shall be used to facilitate the orchestration, management and lifecycles of this collection of elements.

1. Regional and National Datacenters

The primary environment in the operator service network where virtualization has been implemented to date is the large, centralized regional (or super-regional) datacenter. Various architectures and frameworks have been conceived to support this paradigm, which include both open source community maintained and commercially supported implementations.

As the notion of software-defined network workloads is applied at scale in this context, specialized network capabilities, performance demands, and operational considerations have become issues for many deployments.

Resolutions have been brought forward for these issues, either from the community in reference implementations, or by commercial vendors. However, many of these implementations presume a deployment model that only favors a small number of discrete sites with a large, dense population of resources, and unconstrained “East - West” network bandwidth for the software-defined workloads.

Considering only this deployment type drives a model that may not be extensible in the opposite case – “sparse” network edge sites that have a low density of nodes or hosts across a large quantity of sites.

2. Headend sites

Today’s headend sites are where edge-QAM devices, CCAPs and other legacy systems reside. Virtualized and software-defined functions are starting to be deployed at this level as well.

There are industry initiatives that seek to extend this trend to optimize and re-architect headends as datacenters for IP-based software workloads (as well as the edge sites of other operators, such as mobile network base station sites, telco central offices, and the like).

Regardless of the network type, these sites differ from the regional or national datacenter in various engineering and design parameters, including their geographical distribution, access to transport, available space, power and compute capacity. Although technical innovations continue to change the scope of these limitations, these differences ultimately determine the type and quantity of workloads that are suitable for headend site.

3. Remote PHY

The Remote PHY node is the new edge site for intelligent software devices defined in a DAA network. A site that might simply have been associated with an amplifier or regeneration in the past can now become a site for one or more bona fide compute elements running software.

Whether within the Remote PHY Device itself, or collocated in the enclosure or site equipment package, this can be a natural location for a small cell radio site or gNB, virtual RAN components or MEC hosts.

4. New Definition of the Edge

Going beyond RPN sites, there are other candidates for software defined network functions, including the last active, the NID, and even home gateways, set-top boxes, or other CPE devices – these are locations where equipment can be installed or upgraded with elements that provide compute capacity for additional software and network capabilities, and even this class of devices is now capable of supporting virtualization.

Because of the cost, environment, and various technical considerations, these devices often have less resources and capacity than network elements at the RPN, or any point upstream. However, there is still a requirement to provide management, security, software deployment, and lifecycle capabilities for these elements.

For operators who are implementing mobile, these “new edge” locations are obvious candidates for a small cell or combined device providing licensed and unlicensed radio access alongside the DOCSIS network.

With the advent of FDX, the asymmetric balance of the HFC network starts to become more upstream-oriented. This will enable new consumption patterns, as AR/VR hardware and applications gain traction, for example. This will lead to new equipment and software-defined capabilities in the “customer edge”.

Framework Requirements and Common Platform

To enable a common platform for virtual functions – what we might call a “reference environment” – any such framework must provide a consistent set of capabilities:

- Disaggregated functions. Services are highly disaggregated so that control, data, and management planes can be deployed across the distributed topology. Edge clouds offer the performance advantages of low latency and data plane intensive workloads. While control and management plane components can be centralized with a regional and global scope.
- Functional isolation. Network slicing provides network and service isolation across different tenancy models in the reference environment. However, resource management considerations need to be made for shared network functions such as DNS, policy, authentication, and so on.
- Data intensive workload acceleration. The demand for throughput has increased exponentially with smart devices and immersive media services. Networking and compute expenditures
 - continue to grow to meet traffic throughput demands. Support for acceleration technologies like DPDK, VPP, and hardware offload are required to make virtualization of data intensive applications feasible
- Cloud-native and hybrid form factor execution environments. Cloud-native approaches are dictating a new CI/CD paradigm and micro services application architectures. Container technology is a new lightweight execution environment option for delivery of these applications. While the fine-grained abstraction of applications might be a good fit for control plane functions in the reference environment, user plane functions may be required to execute as native VM functions. This requires a cloud infrastructure environment to be heterogeneous enabling such hybrid execution environments for native VM and containerized applications.
- Federation options. The reference environment must provide a diverse set of federation options for end-points, private and public clouds, each with distinct ownership and management domains. Virtualized end-points provide better control and manageability, however they are not suitable for all types of use cases. Likewise, service functions need to be distributed and managed across private and public clouds.
- Service placement. The highly distributed topology allows for flexibility in the workload placement. Making decisions based on proximity, locality, latency, analytical intelligence, and other criteria are critical to enable an intent-based placement model.
- Workload life cycle management. Each cloud is elastic with workload mobility and how applications are deployed, executed, and scaled. An integrated operations management solution can enable an efficient life cycle management to ensure service delivery and QoS.
- Platform lifecycle management. The platform must be patched and upgraded by using optimized change management approaches for zero to minimal downtime.

- Carrier grade characteristics. Because Communications Service Providers (CSPs) deliver services that are often regulated, carrier grade aspects of these services, such as high availability and deterministic performance are also important.

The solution then (as shown in Figure 6) must be a multi-tiered hierarchical platform capable of addressing the requirements and workload types at each level within the service provider cloud – at the regional or national datacenter, the headend or “near edge”, and emerging “far edge” as well.

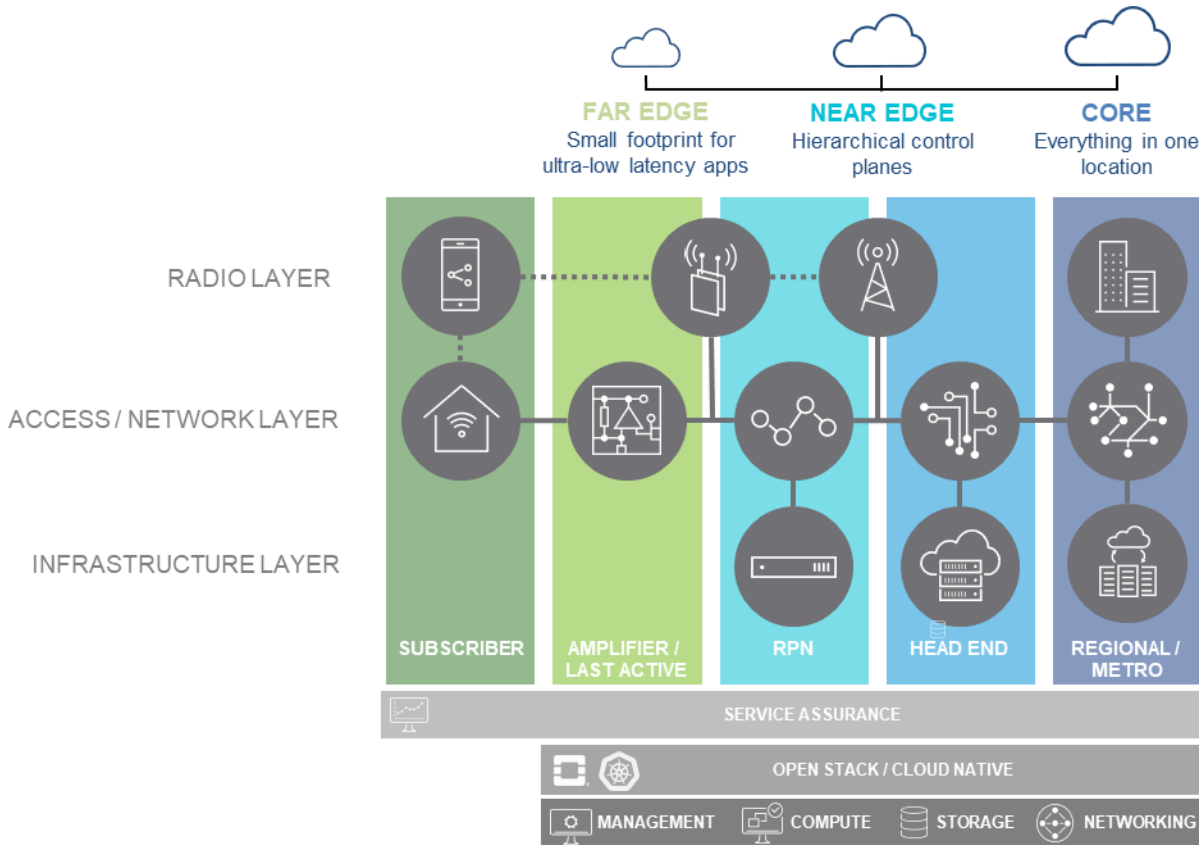


Figure 6 – A Distributed Architecture for the CSP Core / Edge / Access Network Source: VMware

Conclusion

Specific applications are what drove many initial virtualization deployments, sometimes for the use cases discussed here. That led, in turn, to these clouds being customized, tuned, or optimized in unique ways for specific workloads. When there are multiple instances of these clouds, and each is a bespoke environment with a diverging architecture the opportunity to realize a common platform across these applications is lost.

Whenever the architecture of a network changes, or a new cloud is implemented, a key consideration also becomes the visibility, operational tools, troubleshooting and service assurance framework that enables the environment to be managed. Each instance or cloud then requires a solution set for these capabilities which then makes the associated operational practices and support systems potentially different as well.

A fundamental reason for the drive toward virtualization and a common platform for network functions is the principle that it is no longer necessary to solve for the platform and runtime layer below network applications in a different and particular way for each additional application – with all of the attracted cost, complexity, and operational management overhead that differentiation implies.

Learn More

To learn more about VMware's cable cloud solutions, please visit [vmware.com/cablecloud](https://www.vmware.com/cablecloud) or contact your VMware representative.

Abbreviations

5GC	5G core network
AE	access entity
AF	application function
AMF	access and mobility management function
API	application programming interface
AR	augmented reality
AUSF	authentication server function
CBRS	citizens broadband radio service
CCAP	converged cable access platform
CD	continuous development
CDN	content distribution network
CFS	customer facing service
CI	continuous integration
CMTS	cable modem termination system
CU	centralized unit
CPE	customer premises equipment
DAA	distributed access architecture
DCA	distributed CCAP architecture
DOCSIS	data-over-cable service interface specification
DNS	domain name system
DNSSEC	domain name system security extensions
DVR	digital video recorder
DU	distributed unit
EIR	equipment identity register
ESD	extended spectrum DOCSIS
ETSI	European Telecommunications Standards Institute
E-UTRAN	evolved UMTS terrestrial radio access network

FDX	full-duplex DOCSIS
FMA	flexible MAC architecture
HFC	hybrid fiber-coax
HSS	home subscriber server
IoT	internet of things
ISBE	International Society of Broadband Experts
LCM	lifecycle management
LTE	long-term evolution
MAC	media access control layer
MBR	maximum bit rate
MEC	multi-access edge computing
MHA	modular headend architecture
MME	mobility management entity
MNO	mobile network operator
MSO	multiple systems operator
MVNO	mobile virtual network operator
NEF	network exposure function
NID	network interface device
NFV	network functions virtualization
NG-RAN	next-generation radio access network
NSSF	network slice selection function
OTT	over the top
OVP	online video platform
PCF	policy control function
PCRF	policy charging and rules function
PHY	physical layer
P-GW	packet data network gateway
QAM	quadrature amplitude modulation

RCA	root-cause analysis
RDK	reference design kit
REST	representational state transfer
RF	radio frequency
RPD	remote PHY device
RPN	remote PNY node
RS-DVR	remote storage DVR
RU	radio unit
SCH	scheduler
SCTE	Society of Cable Telecommunications Engineers
SMF	session management function
UDR	unified data repository
UDSF	unstructured data storage function
UDM	user data management
UE	user equipment
UMTS	universal mobile telephone service
UPF	user plane function
VIM	virtual infrastructure manager
VR	virtual reality

References

1. Data-Over-Cable Service Interface Specifications, DCA – MHA v2, Remote PHY Specification, CM-SP-R-PHY-I12-1903307, March 7, 2019
2. <https://newsroom.charter.com/press-releases/charter-launches-spectrum-mobile-a-smarter-network-designed-for-the-future>
3. <https://corporate.comcast.com/news-information/news-feed/comcast-xfinity-mobile>
4. ETSI. Multi-access Edge Computing (MEC); Framework and Reference Architecture, ETSI GS MEC 003 V2.1.1 (2019-01)



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 vmware.com Copyright © 2020 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at [vmware.com/go/patents](https://www.vmware.com/go/patents). VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: vmw-wp-tech-temp-uslet-word-101-proof 6/20