TRANSFORMING CABLE INFRASTRUCTURE INTO A CLOUD ENVIRONMENT

Gerry White
Motorola Mobility Network Infrastructure Solutions

Abstract

The paper outlines a methodical evolution strategy from today's RF centric, headend-based infrastructure to a digital-centric one, taking advantage of mature Internet, cloud computing and data center technologies. It presents a phased approach, identifying incremental steps leading to the ultimate goal of an efficient delivery infrastructure, and most importantly one that is aligned with Internet and data center technologies, and henceforth is able to leverage their continued development. Each transitional step is evaluated in the context of current and expected changes in technology, products and services. For each step the advantages provided are highlighted and potential risks are noted.

A number of practical options to deploy subsets of the phases are provided, depending on the individual circumstances of an operator, such as service needs and timing, network characteristics, and risk tolerance.

INTRODUCTION

One of the most significant developments in service delivery in the last few years has been the evolution of cloud computing and the massive data centers used to deliver it. The combination of high bandwidth network connections together with low cost computing and storage platforms has enabled companies such as Amazon and Google to deliver sophisticated services at very low cost points. To date, cable operators have used some of this technology for services such as TV Everywhere but in general it has been competitors such as over the top (OTT) video providers who have best leveraged the new technology. This paper proposes a way for the cable industry to take better advantage of data center technology and outlines a number of steps to achieve the transition.

EVOLUTION

In order to speculate on the evolution of the cable network infrastructure we need to consider the evolution of the services which it must support. Figure 1 shows these parallel evolutionary paths.

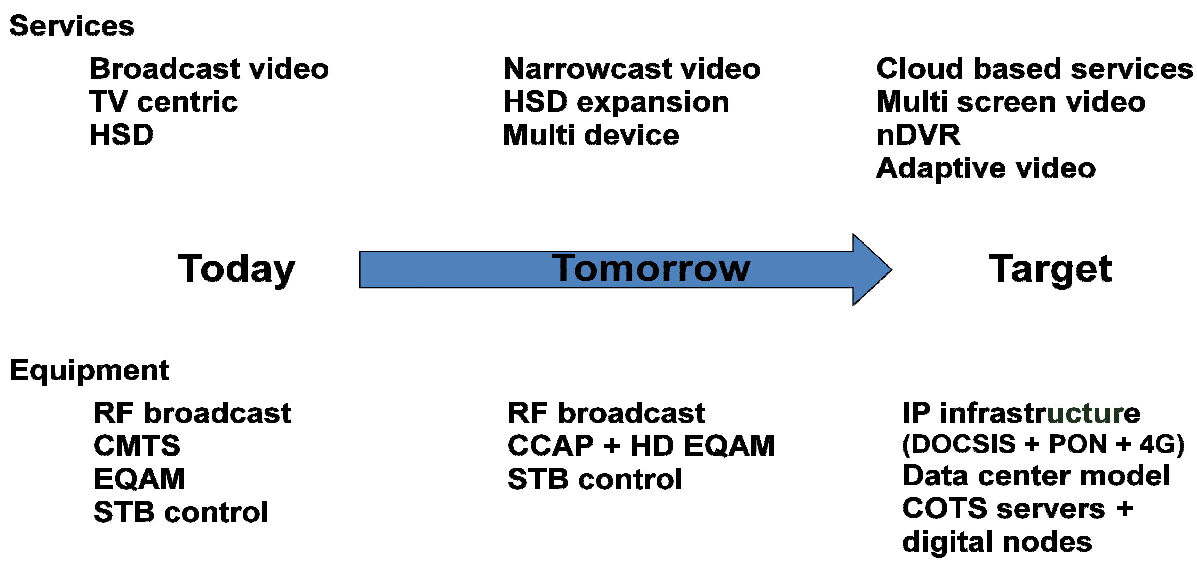

Figure 1: Service and Equipment Evolution

Service Evolution

Current services are heavily focused on delivering linear video programming to a broadcast audience via an STB/TV combination.

In the immediate future we expect to see significant expansion of this service set as narrowcast video services delivering video on demand (VOD) and network based DVR to STB/TV platforms are deployed in parallel with the broadcast service. At the same time high speed data service expansion with compound annual growth rates in the order of 50% is being driven by over the top video services delivered to both TVs and other screens [SAND].

Looking further into the future operators will continue to expand on demand and nDVR services increasingly moving from broadcast to narrowcast services. At the same time competition with OTT video providers will require MSO's to deliver equivalent or better services to multiple types of CPE devices, in the home and on the road. Thus cloud based adaptive video services to multiple devices will become a key component of the service mix.

Infrastructure Evolution

The cable infrastructure must evolve in parallel with the services. Today video services are delivered over an RF broadcast infrastructure, primarily to set top boxes using a proprietary control system. High speed data services are delivered via a parallel CMTS based infrastructure sharing the same physical HFC network as video services but little else.

In the immediate future, as high speed data and narrowcast video expands, existing CMTS and EQAM equipment will be augmented or replaced by higher density platforms to add more narrowcast channels. Systems based on the CCAP specifications will combine CMTS and EQAM functions into a single edge platform but retain the same frequency division multiplexing to share the HFC network and data and video will continue to use independent control systems.

As multi screen video delivery expands the rapid deployment and short lifetimes of the CPE devices will require that service delivery to these devices be based on standard internet protocols with minimal or no changes for the cable infrastructure. The infrastructure must evolve to support IP video delivery to these devices. Thus over time more video will move to an IP delivery mechanism sharing the same resources and technology as high speed data services. This will use existing IP backbone technology for distribution to access networks. The access networks will initially be based on DOCSIS but will include PON and wireless alternatives. This move to a standard IP solution enables the use of standard data center and cloud based services to reduce costs and increase service velocity.

CLOUD SOLUTION

Currently the cloud based environment provides a platform for applications such as OTT video which run over the cable operator's IP broadband service. OTT vendors have taken advantage of cloud services to improve efficiency. Operators have followed suit to some extent with their own OTT like offerings but in terms of leveraging cloud technology are at best on a par with their competition. The remainder of the paper examines how the operator can gain additional advantages from cloud technology by migrating some components of the broadband service itself into this same cloud environment. The technology steps required and the benefits and impacts it may have will be reviewed.

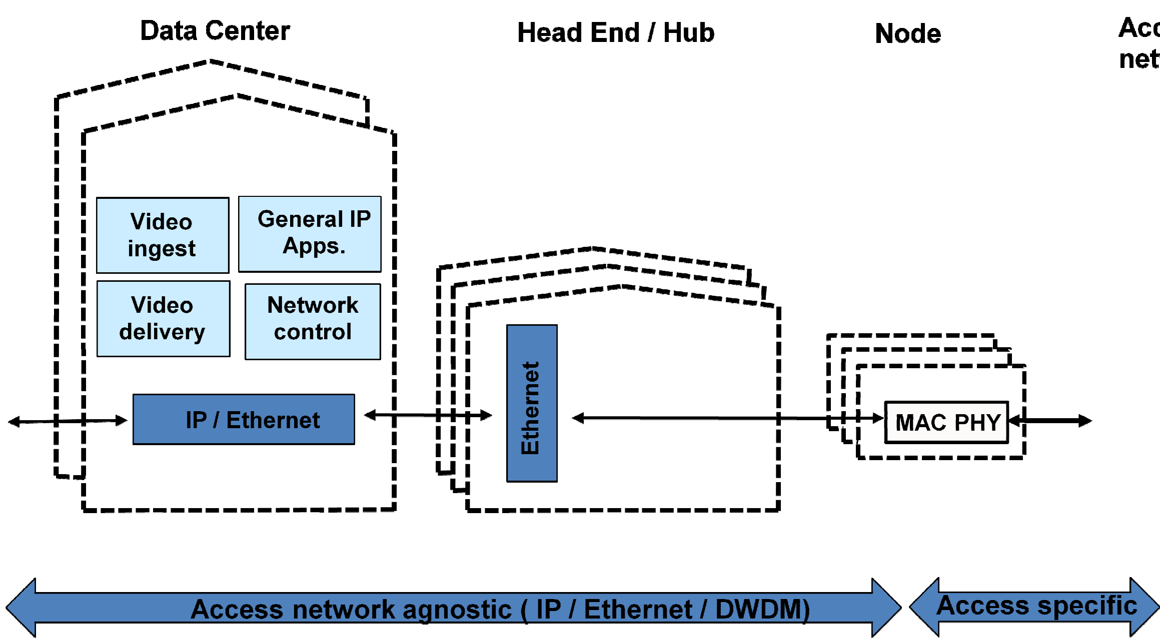

Figure 2: Cloud Centric Solution

Figure 2 shows a possible implementation of this type of architecture illustrating the major components and their location in the network. The philosophy behind this approach is simple; to put as much functionality in the data center as practical, to leverage Ethernet and digital optics where possible and to keep the node simple. The reasons to migrate functions to the data center are to leverage off the shelf hardware and software for reduced cost, to leverage virtualization for redundancy and scaling and to centralize complexity for simpler operation. Ethernet and digital optics are used to provide low cost and long distance options. The node is kept simple for low cost and ease of operation. As a consequence of the migration of functions to the data center and node the head end and hubs become much simpler.

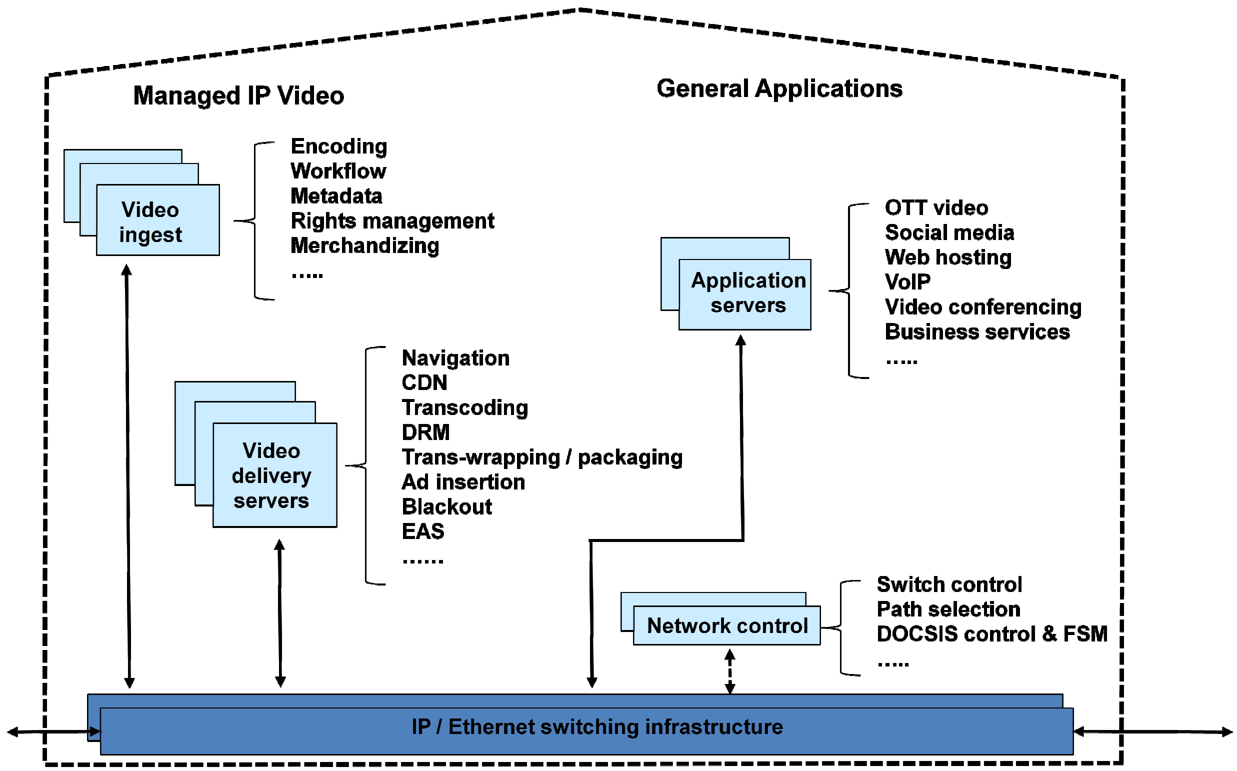

Figure 3: Data Center Functions

<u>Data Center</u>

As with existing cloud based services the data center is based on off the shelf servers running general applications software in a virtual environment as shown in Figure 3. These include general applications software such as OTT video, social networking and Internet access.

Additional servers provide the functions needed to create a multi-screen video service. Video ingest servers receive content from multiple sources and provide encoding, metadata and content management functions.

Video delivery servers provide transcoding and packaging functions to create multiple bit rate program streams along with additional functions such as advertising insertion, and content protection that are required by a full service video provider. A description of a layered architecture for an end to end IP video system can be found in [MSIPD].

This is a simplified description in that these functions may be implemented in a central data center or in a more distributed model based on multiple data centers interconnected with CDN distribution networks. For the purposes of this paper the simple model will suffice as the evolutionary steps proposed for the network are independent of the model selected for data center deployment. The video service architecture described above is in fact deployed currently in both centralized and distributed modes and can be used with both existing HFC delivery networks and the evolved network proposed.

The next step in evolving the network is to use additional servers in the data center to run the access network (e.g., DOCSIS) control plane. In a traditional router or CMTS the data plane typically runs in specialized hardware while the control plane runs in a general purpose CPU embedded in the platform. In this case the general purpose CPU has migrated to a server and standard IP/Ethernet switches provide the data plane forwarding within the data center and from the data center to the head end. A control protocol between the server and the switch such as OpenFlow [OPENF] is used to control the forwarding path.

Head End

The head end in this architecture continues to support traditional analog and MPEG based broadcast video. High speed data and narrowcast video processing has moved to the data center as described above. HSD and IP video traffic passes through the head end via an IP/Ethernet network and is forwarded to the node using the existing fiber links which it shares with the broadcast video using WDM.

Node

In the node the broadcast video is converted from optical to coax media as today. The Ethernet traffic is converted from baseband Ethernet to the protocol to be used for the node to home portion of the network. This may be DOCSIS, other Ethernet over coax (EoC) access technologies, point to point Ethernet or even wireless technologies such as WiFi or LTE. Operation of the node is controlled from the data center (via the head end).

ADVANTAGES

The architecture described above has multiple advantages over a traditional HFC video delivery infrastructure.

Leverage Data Centers

It leverages the work done to provide massively scalable Internet based applications over the last decade to provide cost savings and efficiency:

- The use of COTS technology reduces costs by using general purpose servers as processing engines and standard Ethernet networking equipment for connectivity.
- It leverages standard virtual environments to provide horizontal scaling and redundancy
- It provides a simpler and more robust environment for networking software development.
- It provides a friendly platform for application level software development where familiar toolsets and environments enable software to be created by operators, equipment vendors and third parties to accelerate service velocity.

Head End Simplification

The head end becomes a much simpler environment leading to lower operational costs:

- Complex software functions are centralized in the data center where expertise can be concentrated.
- Power, cooling and rack space needs at the head end are reduced as devices such as CMTS are removed.

- IP services are delivered via a standardized Ethernet network as low level DOCSIS and QAM functions migrate to the node and the need for RF combining in the head end is reduced or ultimately removed.
- Multiple head end / hub locations may be collapsed or run remotely as a "lights out" operation saving operating and real estate costs.

Network Simplification

The use of IP transport to the node simplifies the network in the following ways:

- Standard Ethernet and digital optics can be used between the head end and the node eliminating distance limitations and enabling distribution hubs and small head ends to be consolidated.
- Ethernet switching is used in the data center, head end and hub to reduce costs in the transport network.
- The network from the data center to the hub is independent of the last mile technology from the hub to the home allowing these parts of the network to evolve independently and more cost effectively.

TRANSITION STAGES

From the above it appears that there are significant advantages to moving to the proposed architecture but as always the problem is in how to facilitate the transition at a reasonable cost while continuing to provide service. The following sections of the paper address this transition and break it down into a number of potential stages. The discussion identifies six possible transition stages between today's HFC network and that shown in Figure 2. These are:

1. Introduction of CCAP
2. Split packet processing from physical media dependent and physical layer (PMD/PHY) processing
3. Move PMD/PHY to the node
4. Move MAC processing to the node
5. Move narrowcast processing from head end / hub to data center
6. Retire broadcast processing and remove from head end /hub

Not all stages are required and each operator can select the most appropriate path based on their existing network, operational needs and competitive demands.

Stage 1

This first phase of the transition, shown in

Figure **4**, addresses the change in services from broadcast to narrowcast. Narrowcast channels for MPEG and DOCSIS are added to the head end using high density CCAP based equipment. This may augment or replace existing EQAM and CMTS equipment. The CCAP platform connects to the core network through the Ethernet distribution network in the head end as for existing equipment but will use 10Gbps rather than 1Gbps Ethernet links. Connectivity to the HFC remains RF based and connects to an optical shelf through the existing RF distribution /combining networks in the head end. The optical shelf converts the signals to analog optics and transmits them to the fiber nodes in the outside plant.

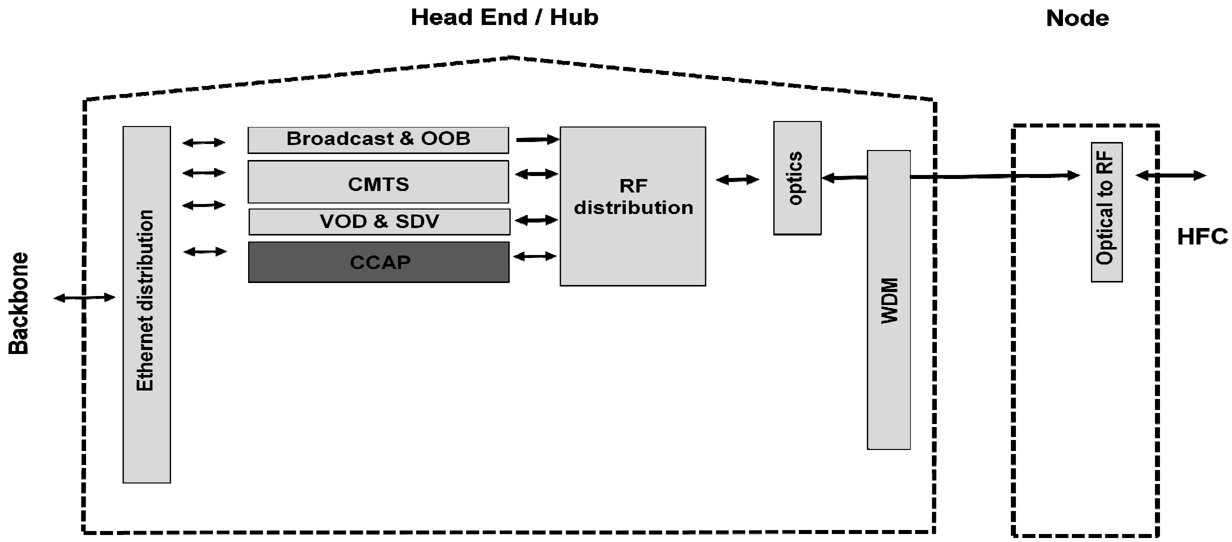

Figure 4: Phase 1, CCAP

The advantages of this transition are the savings in cost, real estate and power consumption provided by the increased density of next generation platforms [CCAP].

The technology changes required and risks associated with this change are those associated with the CCAP program and are well understood at this time.

<u>Stage 2</u>

The next stage of the proposed transition requires a slightly more radical change and is shown in Figure 5. It leverages some of the work done for modular CMTS [M-CMTS] and distributed CCAP architectures [D-CCAP] to move to an Ethernet based distribution system in the head end. The core principal is to decouple the packet processing from the physical media dependent (PMD) portion of the MAC and the PHY. The PHY and PMD dependent functions move out of the core processing engines to be collocated with the optical shelf which may be within the head end or in a distribution hub. The CMTS and EQAM core engines output MPEG streams over Ethernet using standard framing [M-CMTS]. The downstream modulation function is included with the optical transmitters and the upstream demodulation function is included with the optical receivers.

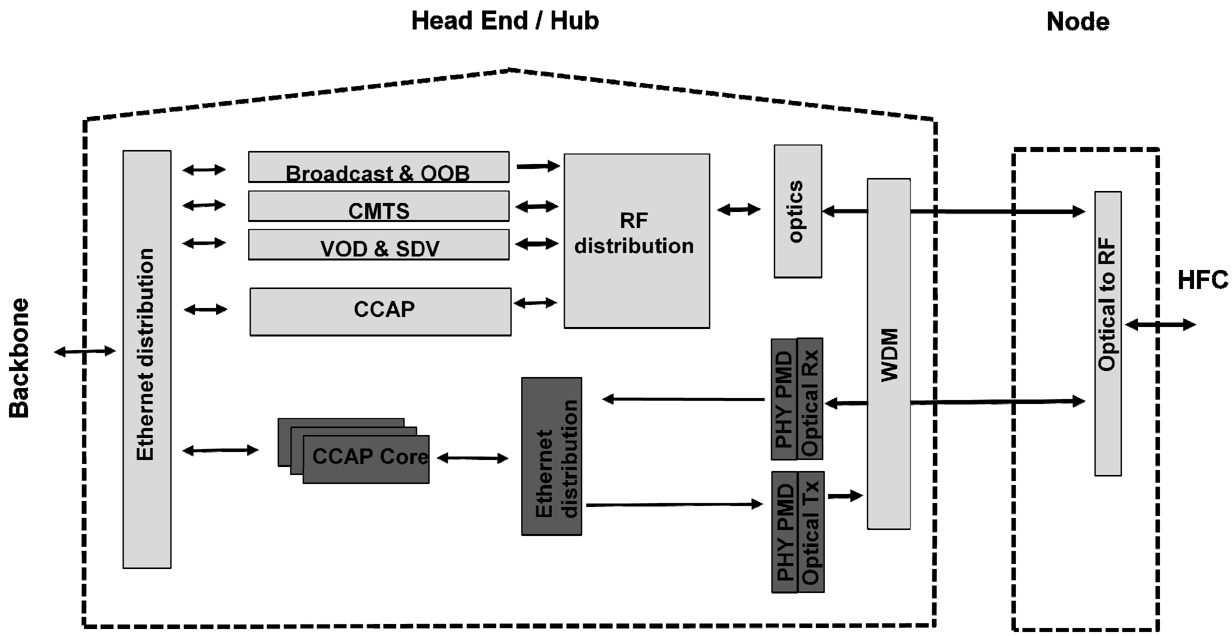

Figure 5: Phase 2 with Local CCAP core

This separation of digital processing from modulation has several advantages. Scaling the system becomes more efficient as CMTS and EQAM core engines scale based on processing needs while the PMD and PHY layer functions scale based on port counts. Thus the core engines scale up as the total number of DOCSIS and EQAM channels to be delivered increases while the PMD/PHY functions scale as the number of serving groups increases. While these factors (channel count vs. serving group count) are certainly related it is typically not a strictly linear relation so there is benefit to the separation.

Redundancy may also be simplified in this case. Ethernet based redundancy can be used for the core processing engines so that there is no need for complex RF switching logic in the core chassis in the data center. RF redundancy can be provided in the optical shelf for large serving groups but for smaller serving group sizes the failure group in the optical shelf may be small enough that RF redundancy is not needed providing further cost reductions.

With the replacement of the RF combining network by Ethernet switching changes such as node splits are made simpler; ports are added to the optical shelf for the new node and processing engines added to the core shelf if required. The operations and management network for the head end can share the Ethernet infrastructure to provide central configuration and monitoring using standard IP tools. Thus path traces and testing across the head end become trivial using tools such as ping and traceroute.

The most interesting feature of the PHY-PMD separation is the increased flexibility it enables for equipment deployment. The standard Ethernet links between the core processing engines and the optical shelves effectively remove any distance restrictions between them. Thus the optical shelves could be deployed in a remote head end or hub while the processing engines reside in a data center. This enables “lights out” operation of the remote facility and significantly reduces power and cooling needs at these locations. Centralizing the complex equipment also reduces the operations skill sets needed at the remote locations.

Figure 5 above shows a potential deployment scenario in which additional services are added via the new model as needed. Existing broadcast, CMTS and EQAM equipment can remain unchanged or be replaced over time.

Figure 6 shows an alternative deployment model with the CCAP core platforms located in a data center remote to the head end.

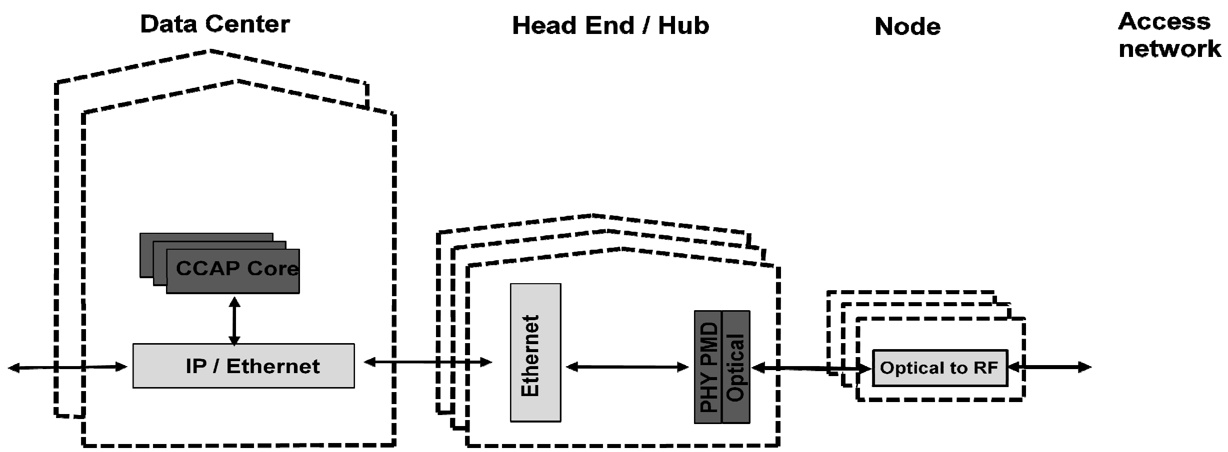

Figure 6: Phase 2 with Remote CCAP core

The technology risks associated with this phase are primarily the timing issues associated with the separation of the DOCSIS MAC and PMD functions.

Stage 3

As mentioned above the use of an Ethernet distribution network between the core processing engines and the optical shelf eliminates distance restrictions between them. The next phase takes advantage of this fact to move the PHY-PMD function out of the head end to the node and extend the Ethernet distribution to the node using standard Ethernet optics as shown in Figure 7 . New services are added using the Ethernet to the node transport while legacy services continue to be supported using an RF overlay. Ethernet and analog wavelengths are multiplexed onto the same fibers using existing WDM equipment.

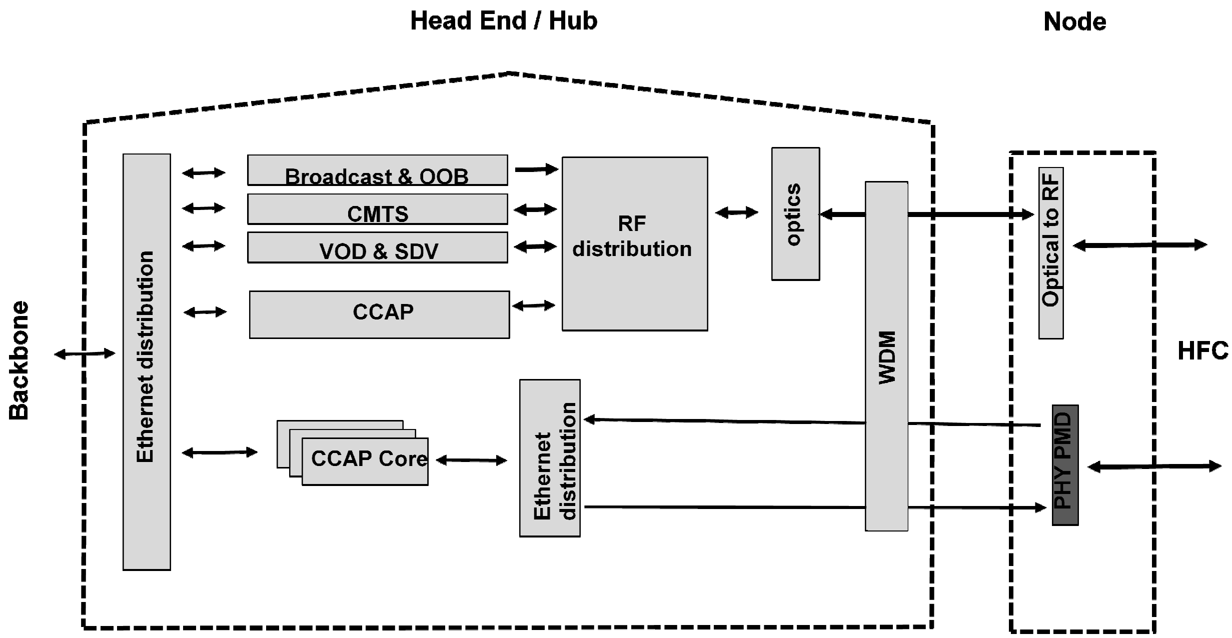

Figure 7: Stage 3, Remote PHY-PMD

Deployment of this phase of the transition plan brings several advantages to the operator. Service expansion is no longer limited by the head end to node transport as the use of digital optics and DWDM provide essentially unlimited bandwidth on this link. The links leverage the costs, speed and density trajectory set by Ethernet systems and standard DWDM platforms.

If the RF overlay can be removed then distance limits between the head end and the node are eliminated. It may then be possible to centralize head end functions and retire some distribution hubs and small head ends providing savings in real estate and operational costs.

The technology risks associated with this phase are the density, powering and cooling of the components in the fiber node and the provision of timing services to the node.

Stage 4

This phase, shown in Figure 8 is a conceptually small change in which the upper layer MAC functions are moved to the node so that they are co-resident with the PMD and PHY functions. As in the previous phase new services are added using the Ethernet to the node transport while legacy services may continue to be supported using an RF overlay

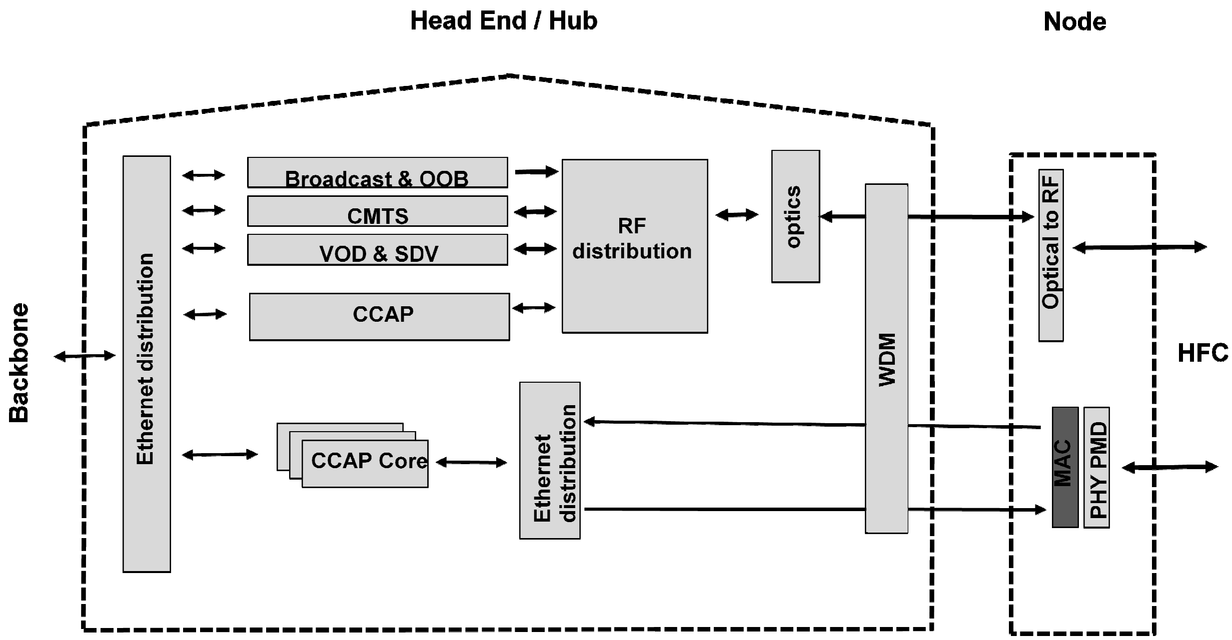

Figure 8: Stage 4 Remote MAC

The advantage of moving the MAC to the node is that it simplifies the timing issues relative to the split implementations described previously. More importantly it decouples the technology for the head end to node transport from that used between the node and the home. This allows the technologies to evolve at their natural and different paces. The head end to node link uses general enterprise networking technology while the node to home link remains specific to the HFC plant. Node splitting and the progression towards an n+0 architecture becomes a process of replacing the DOCSIS MAC/PHY module in the node with an Ethernet switch and moving it further downstream. The impact of transitions to next generation technologies such as PON or EoC is restricted to the node to home portion of the network.

The technology risks are similar to the prior phase as more functions are moved to the node increasing powering and cooling needs.

Stage 5

In this phase, shown in
Figure 9 the legacy narrowcast equipment in the head end is retired and all processing elements migrate to a central data center. The HFC specific MAC elements have migrated to the node and IP/Ethernet transport is used throughout the network from the core to the node. Application and control plane logic are no longer in the head end which is only used for legacy broadcast services. It still acts as a pass through for narrowcast services but this is reduced to an IP/Ethernet switching function.

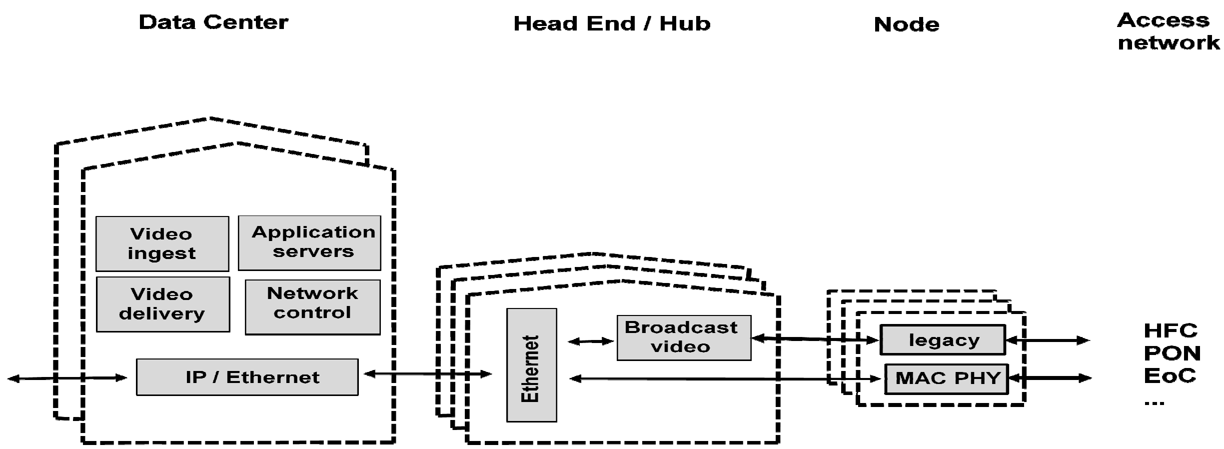

Figure 9: Stage 5 Narrowcast Equipment Removal

Stage 6

With the removal of traditional broadcast the head end in its current form can go away completely and be replaced by a data center, an Ethernet distribution hub and a simple node as shown in Figure 10. At this point the advantage of centralization, standard IP/Ethernet transport and isolation of HFC specific functions described in the earlier phases are now fully realized.

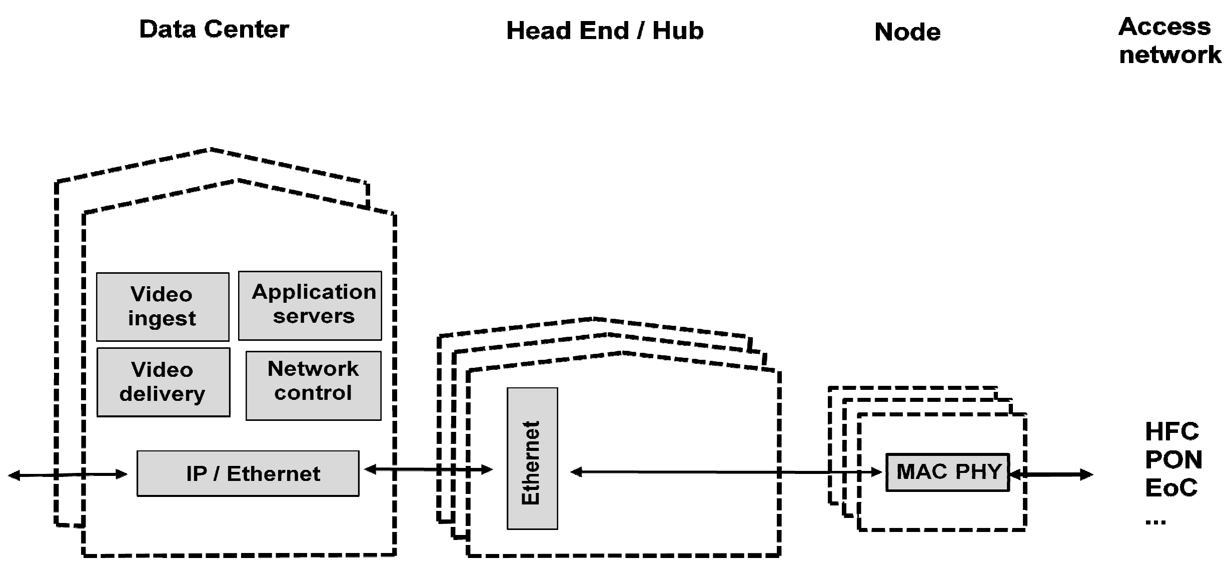

Figure 10: Stage 6 Broadcast Equipment Removal

Moving the MAC and PHY functions from the head end to the node allows the use of standard Ethernet optics and enables distributed processing but results in a more intelligent outside plant architecture. Operators who do not wish to take this step and prefer to keep a simpler outside plant can elect to deploy the MAC-PHY components in the remote hub rather than the node as shown in Figure 11. They still retain the advantages of the move to the data center and a significant reduction in hub complexity. Readers interested in an in depth comparison of traditional and intelligent HFC architectures are referred to [HFCDFC].

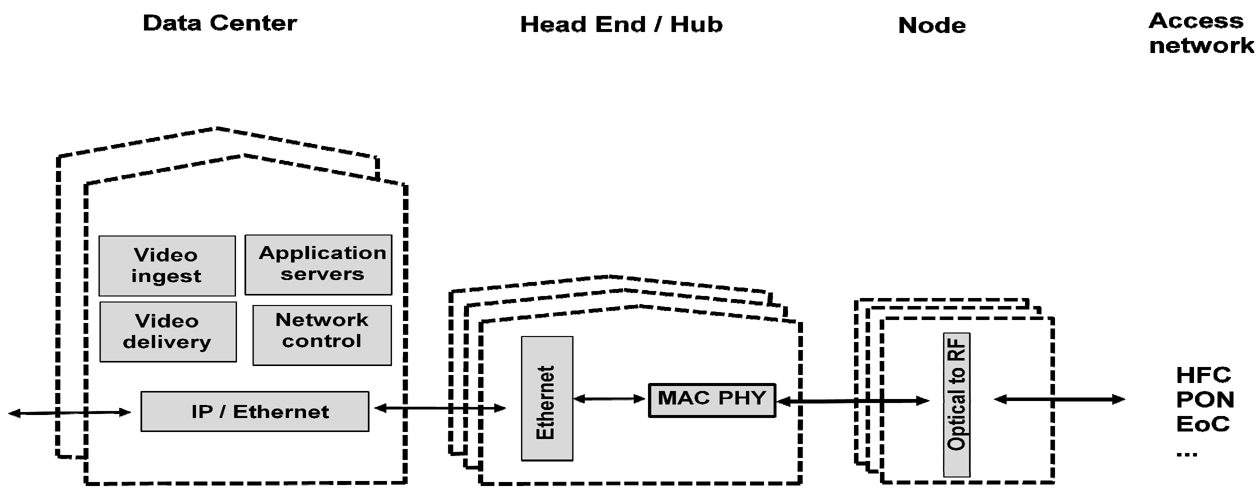

Figure 11: Passive HFC Architecture

DEPLOYMENT

The transition stages previously described illustrate a logical roadmap to the data center architecture. They are based on technology evolution but are essentially independent. Which stages are deployed by an operator will depend on their customer needs, timing, risk profile, budget and network architecture.

Table 1 shows a summary of the risks and benefits associated with each phase together with an indication of when it would be appropriate to be used.

Stage	Change	Benefits	Risks	When Appropriate
1	Move to CCAP platform	Increased density, lower cost, simplified combining	Minimal, well understood problem	Need to add high density narrowcast services in HE/hub
2	Decoupled MAC / PHY in head end	Processing and port scaling separated, simpler redundancy, simpler combining	DOCSIS MAC/PHY timing split.	Gain HE/hub benefits without touching node
3	PMD+ PHY move to node	Ethernet Digital optics to node	Power & cooling in node	Consolidate small HE/hub; retain existing core platforms
4	MAC to node	Data center to node links all Ethernet for narrowcast traffic	Power & cooling in node	Standardize on Ethernet transport to the node to set up for stages 5 & 6
5	Narrowcast removed from HE	Head end space savings	minimal	Consolidation from HE/hub to data center for lower OPEX
6	RF broadcast removed from HE	HE becomes simple switching center or is removed	minimal	Consolidate or retire HE/hub

Table 1: Risks and Benefits of Each Stage

CONCLUSION

Moving functions from the head end or distribution hub into a data center has many advantages and has the capability to provide significant capital and operational savings. To transition to a network architecture which can take full advantage of this move is not trivial but can be achieved through a series of stages as technology evolves and service needs demand. The transition stages described are largely independent and any given operator can select the transition path best suited to their specific needs.

REFERENCES

[CCAP1]	J. Salinger, "Proposed Next Generation Cable Access Network Architecture", SCTE Conference on Emerging Technology, 2009.
[CCAP2]	J. Salinger, "Understanding and Planning CMAP Network Design and Operations", SCTE Cable-Tec Expo, 2010.
[CCAP3]	J. Finkelstein, J. Salinger, "IP Video Delivery using Converged Multi-Service Access Platform (CMAP)", SCTE Canadian Summit, 2011
[D-CCAP]	John Ulm, Gerry White New Converged Access Architectures for Cable Services, NCTA 2011 Spring Technical Forum
[HFCDFC]	M. Emmendorfer, S. Shape, T. Cloonan & Z. Maricevic Examining HFC and DFC (Digital Fiber Coax) Access Architectures, SCTE Cable -TEC Expo 2011
[M-CMTS]	Cablelabs DOCSIS® Specifications — Modular Headend Architecture (MHA)
[MSIPD]	John Ulm, Gerry White Arch & Migration Strategies for Multi-screen IP Video Delivery, SCTE Canadian Summit 2012
[OPENF]	www.openflow.org
[SAND]	Global Internet Phenomena Report Fall 2011; Sandvine

ABBREVIATIONS AND ACRONYMS

CCAP	Converged Cable Access Platform
CDN	Content Delivery Network
CMTS	DOCSIS Cable Modem Termination System
COTS	Commercial Off The Shelf
CPE	Customer Premise Equipment
DOCSIS	Data over Cable Service Interface Specification
DRM	Digital Rights Management
DVR	Digital Video Recorder
DWDM	Dense Wave Division Multiplexing
EAS	Emergency Alert System
EoC	Ethernet over Coax
EPoC	EPON over Coax
EPON	Ethernet Passive Optical Network
EQAM	Edge QAM device
FSM	Finite State Machine
Gbps	Gigabit per second
HFC	Hybrid Fiber Coaxial system
HSD	High Speed Data; broadband data service
HTTP	Hyper Text Transfer Protocol
IP	Internet Protocol
MAC	Media Access Control (layer)
Mbps	Megabit per second
MPEG	Moving Picture Experts Group
MPEG-TS	MPEG Transport Stream
nDVR	network (based) Digital Video Recorder
OTT	Over The Top (video)
PHY	Physical (layer)
PMD	Physical Medium Dependent (layer)
PON	Passive Optical Network
RF	Radio Frequency
STB	Set Top Box
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VOD	Video On-Demand
VoIP	Voice over IP
WDM	Wave Division Multiplexing