The Economics of IP video in a CCAP World

John Ulm & Gerry White
Motorola Mobility

Abstract

The paper outlines an IP video architecture and determines the relative cost contributions from the major components. For current equipment, the DOCSIS® downstream channel is shown to be the major contribution to infrastructure cost. As next generation Converged Cable Access Platform (CCAP) systems are deployed this will fall enabling a cost effective IP video platform to be realized. At this point other cost contributors become more significant. CDN and nDVR trade off options are discussed. Finally the paper looks at spectrum migration options to release the bandwidth needed to deliver IP video service.

INTRODUCTION

Delivery of IP video will be a major factor driving cable infrastructure during the next few years. Studies of Internet traffic patterns [SAND], [VNI] show that video has become the dominant traffic element in the Internet consuming 50 to 60% of downstream bandwidth. Cable's "Over The Top" (OTT) competitors account for much of this traffic, with Netflix alone constituting almost 33% of peak hour downstream traffic in North America.

To remain competitive cable operators need to deliver IP video to the rapidly expanding tablet, PC, smart-phone and gaming device market. They must leverage the same cost effective technologies as OTT competitors for this and for video delivery to the primary TV. Thus service providers must deliver two forms of IP video: unmanaged OTT off-net and managed video services on-net. This has caused the industry to become very focused on the implications of offering IP video over a DOCSIS® (Data Over Cable Service Interface Specification) channel.

Over the years, the relatively high cost of a DOCSIS channel has impacted potential solutions for IP video. In the past this spawned multiple alternate proposals. Bypass architectures such as DOCSIS IPTV Bypass Architecture [DIBA] were proposed as alternatives and bandwidth saving mechanisms such as multicast and variable bit rate (VBR) technologies investigated. These have become somewhat redundant with the recent surge of adaptive bit rate (ABR) protocols among consumer devices. This unicast delivery mechanism based on HTTP has become the defacto standard for IP video services to this class of devices. In fact, ABR may be used for all IP video traffic including primary screen [CS_2012].

Thus a critical question for operators is how to deliver unicast based IP video cost effectively. It is important to understand the cost implications for DOCSIS downstream channels in the future.

IP VIDEO ARCHITECTURE

To understand the economic impact of migrating to IP video, the system must be separated into key elements. Components of a Managed IP video Architecture are detailed in [CS_2012] and [Ulm_NCTA_2012].

Figure 1 shows a high level abstraction of an end to end functional architecture for delivering IP video from content providers to content consumers. The video service provider must ingest content from multiple content providers, process it appropriately and then transport it over multiple types of access networks to the destination consumer devices.

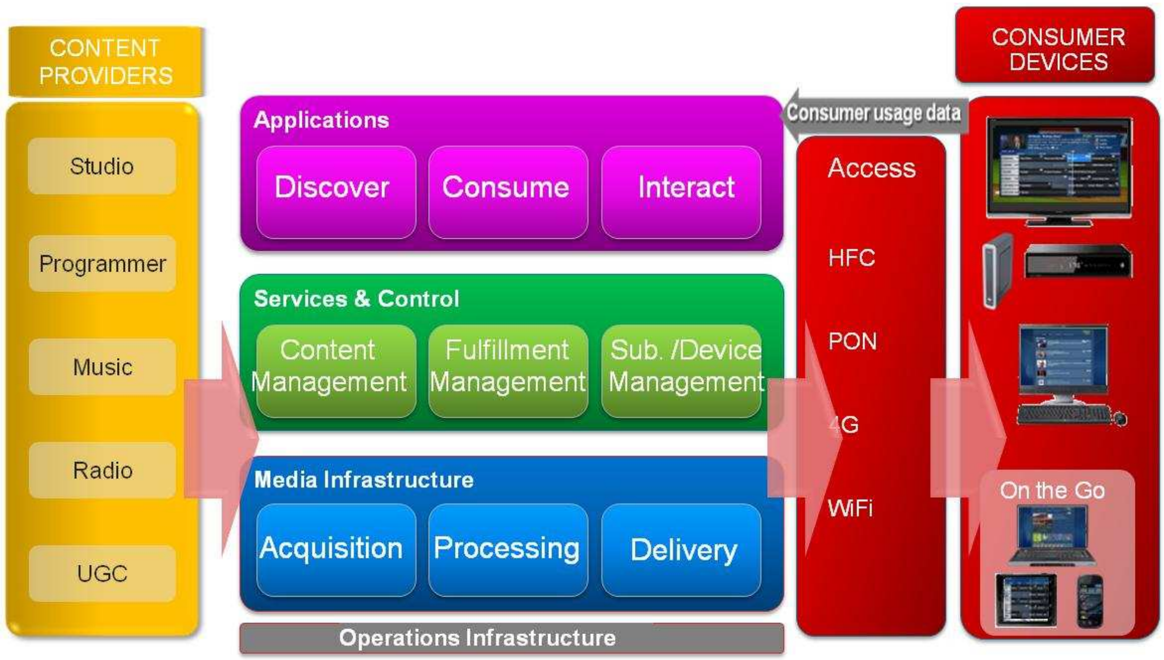

Figure 1 IP video Functional Model

This functional model is used to develop a high level breakdown of the costs for IP video delivery and to compare the relative contribution of each component. This will enable operators to understand the impact of the major cost drivers and make intelligent system trade-offs in their IP video architecture.

MAJOR COMPONENTS AND COST IMPLICATIONS

Content Providers

The number of content sources is increasing. Traditional streamed linear television broadcasts from studios and programmers may be received over satellite or terrestrial links in MPEG-2 and MPEG-4 formats. User-generated content and other Web based multimedia sources must also be supported, but will more typically be delivered as file-based assets.

Costs associated with ingesting content from content providers scale based on the number of program sources ingested and the cost of the material. Once purchased and ingested, these programs are shared across all subscribers. The cost per subscriber is not materially impacted by changes within the delivery infrastructure and thus these costs are not considered further in this paper.

Consumer Devices

One of the principal drivers towards a service provider IP video infrastructure is to be able to support generic IP-based consumer devices such as smart-phones, tablets and gaming devices. The range of consumer devices appears to be almost limitless in terms of screen sizes and resolution, network data rates, processing power, mobility, media format support and DRM support.

Most of these consumer devices are owned directly by the consumer. The one exception to consumer-owned devices might be the IP set-top box. For this analysis, it is assumed the operator will have some leasing revenue associated with the IP STB so it is not considered as part of the infrastructure costs. There are some cost tradeoffs in the use of home gateways and hybrid video gateways which will be considered later in the paper.

Access Networks

A primary reason to move to an IP video infrastructure is that it can be access network independent in contrast to existing MPEG/RF video infrastructure. For the purpose of this investigation only the hybrid fiber coaxial (HFC) access network will be considered.

Core IP Network

The components of the IP video architecture interconnect via the same generic IP core network used for all video and high speed data service delivery. The costs of the core network are amortized over multiple services and all subscribers. Thus the cost contribution to IP video service on a per subscriber basis is relatively low.

Application Layer

The Application layer provides interaction with the end user and is largely responsible for the user experience. It includes functions that: 1) discover content through multiple navigation options such as user interfaces (UI), channel guides, interactive search, recommendation engines and social networking links; 2) consume content by providing applications for video streaming, video on demand (VOD) and network DVR (nDVR) consumption; and 3) provide companion applications which enable user interaction in conjunction with media programs such as interactive chat sessions.

Applications are typically implemented in software running on servers in the data center with a thin client application on the consumer device. The applications may be provided by the service provider, the device provider or a third party. Costs associated with the applications layer are thus shared between these entities. On a per subscriber basis these are relatively small as they are amortized over a large number of subscribers.

Services & Control Layer

The Services & Control Layer is responsible for assigning resources within the network and for enforcing rules on content consumption that ensure compliance from a legal or contractual perspective. It includes functions that manage: content work flow from ingest through to delivery; the resources needed to ensure content is delivered to users when requested; and subscribers and devices to ensure that content is delivered to authorized consumers in the required format.

The Services & Control Layer is implemented as a set of software applications running on servers in the service provider network. These applications are typically licensed on a per subscriber or per session basis. Thus costs are a combination of hardware platform and software licensing. The basic control components required include: workflow and session management, DRM control, and resource management.

Media Infrastructure Layer

The Media Infrastructure Layer is responsible for video content delivery from the content provider to the consuming devices over the access network. This includes acquiring content from satellite or terrestrial sources (as either program streams or files), encoding it for ingest into the system and processing the content to prepare it for delivery. This is where the heavy video processing occurs and functions such as trans-coding, multiplexing, advertising insertion, encryption and publishing to a content delivery network (CDN) are found. This layer must also deliver the content to the target device through mechanisms such as web servers, CDNs, and streaming servers.

Costs for content reception and encoding scale on a per content stream basis. The content is shared across many subscribers so that the per subscriber cost is low. Packaging costs may scale based on content streams for a pre processing model or on subscribers if a just-in-time model is used. The choice between these is based upon a trade-off between packaging, storage and transport costs [PACK]. CDN costs do scale on a per subscriber basis.

RELATIVE END TO END INFRASTRUCTURE COST

The relative end to end cost of delivering IP video to a subscriber includes contributions from all of the components mentioned above and each component can have a wide range of variability. The Application Layer and Services & Control Layer products tend to be software on standard server platforms in a data center where costs are shared over a very large number of subscribers. The Media Infrastructure Layer is the component that contains the specialized hardware products and is where most of the operator investment occurs. Rather than attempt a detailed investigation of all of these components, the focus of the paper is on how changes in the network access pieces of the Media Infrastructure Layer impact the cost model.

Cable modem termination system (CMTS) ports to date have been deployed to provide high-speed data (HSD) and voice services. CMTS costs on a per subscriber basis have been relatively low. This cost point has been possible because HSD services could be heavily over-subscribed. IP video has a very different service model and cannot be over-subscribed to the same extent. A single CMTS channel can support anywhere from a half dozen high definition (HD) to a couple dozen smaller active video streams depending on the encoding rate used. The ratio of high definition to standard definition content now becomes very important. At historical CMTS pricing points, this translates to an order of magnitude of $100's per IP video stream.

Each of the components above is highly configurable which can result in wide variations in the end to end cost analysis. To understand the relative costs of these components required a nominal use case based on data from actual products and bid responses. This is shown in Figure 2 on a cost per active video user basis. For this example, the CMTS cost is roughly ten times the costs ascribed to the other major components. It is clearly the most significant cost driver for IP video and will be the primary focus of the rest of the paper.

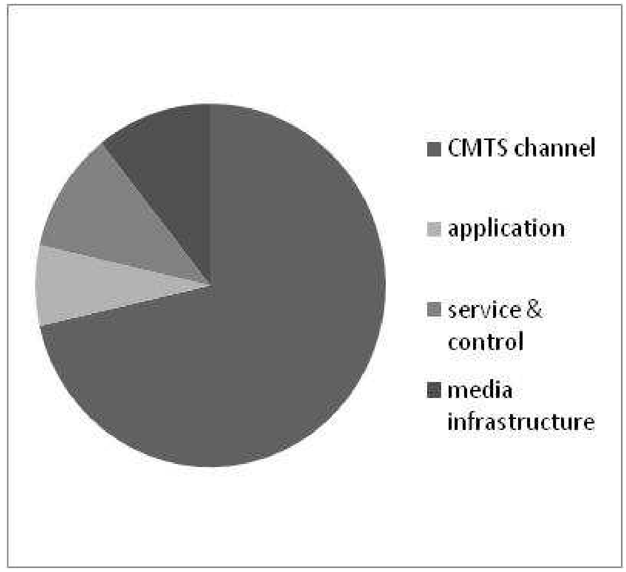

Figure 2 Per video user cost contributors

IP VIDEO & DOCSIS CHANNEL COSTS

CMTS Costs – Historical Perspective

DOCSIS is now 15 years old, having first been established in March 1997. Over that time, it has continued to evolve. In the early days, the cost per downstream channel was above $10,000. Early implementations had fixed downstream to upstream ratios (e.g. 2x8), so if more downstream bandwidth was needed, the system was burdened with the cost of more upstreams whether or not they were needed.

In addition to the fixed ratio, these early CMTS's were focused on offering a robust voice service for the operators. This introduced significant costs as these CMTS became carrier grade incurring the associated redundancy overheads.

Thanks to Moore's Law, these costs were reduced over time. Two architectural changes accelerated this trend. First, the DOCSIS 3.0 specification (D3.0) was developed and released. This laid the groundwork to enable multiple bonded channels per downstream port. At the same time, CMTS architectures shifted to decoupled architectures where upstream and downstreams could scale independently of each other. Some vendors chose a modular CMTS (M-CMTS) path for this while others implemented decoupled architectures within their Integrated CMTS (I-CMTS). As D3.0 was deployed, this helped to accelerate the reduction in cost per downstream channel as multiple channels were now implemented per port and the upstream burden per downstream channel was reduced.

So where are we today? Based upon recent research from Infonetics (Q4 CY2011), the revenue per downstream (channel) will decline in calendar year (CY) 2012 to approximately $1,600. After several years of significant reductions following the introduction of D3.0, the industry is starting to see price declines level out. Infonetics has forecasted that CY12 will see a 10% drop over CY11 which is substantially less than the previous two years.

As we move forward with unicast based IP video, it is very important to understand the cost implications for DOCSIS downstream channels going forward.

CCAP Disrupts DOCSIS Density & Pricing

Recent industry and CableLabs® efforts have defined a new specification called CCAP that is a high density combination of CMTS and edge QAM (EQAM) in a single unit. Current CMTS products may only support 4 or 8 channels per downstream port. The initial version of CCAP is defined to support 64 narrowcast channels per port, with a flexible channel mix between DOCSIS and EQAM. Future CCAP products may support 128 or even 160 channels per port, enough to fill the entire 1GHz downstream spectrum. Clearly, CCAP causes a disruptive shift in downstream densities, increasing by a factor of sixteen! With these densities, there will be a corresponding decrease in the cost per downstream channel. For IP video deployment, it is very important to understand how CCAP will affect access network costs.

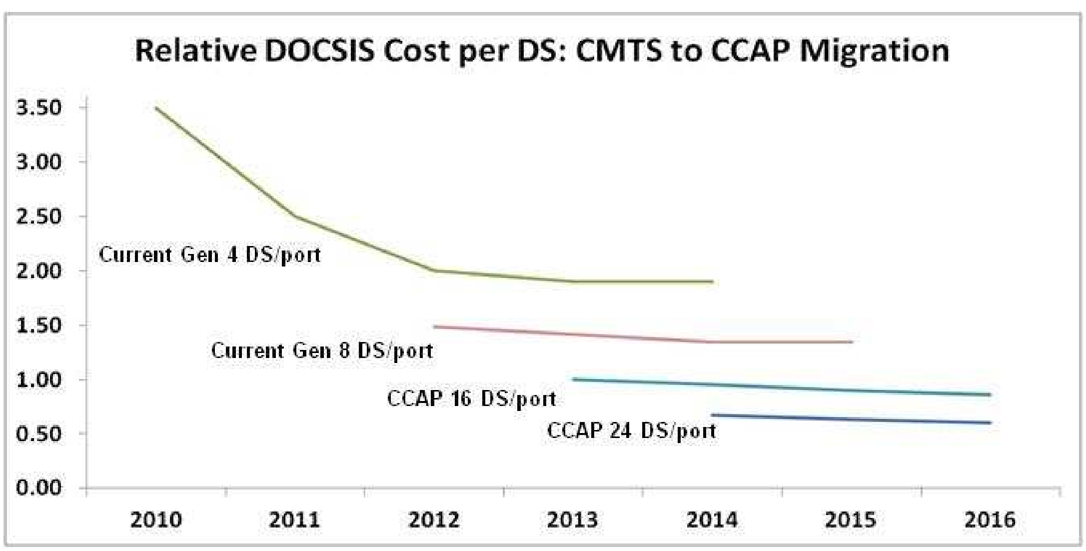

Figure 3 DOCSIS Downstream Cost

Initially, operators will only need a fraction of the CCAP capacity. Even if they wanted to deploy more channels, the spectrum required is a very scarce resource. For an operator to buy the full CCAP capabilities but only use a fraction of its capacity (e.g. 16 downstream channels) would cause a significant spike in the cost of downstream channels. CCAP would not be cost effective compared to current CMTS platforms. Therefore, vendors will need to license channels, similar to what is done today for high-density EQAM products. This allows CCAP products to be deployed while offering competitive downstream channel costs; vendors then defer revenue to a later time once additional channels are licensed and operators gain the benefit of deploying systems with longer lifetimes. Figure 3 above depicts the downstream channel cost trends over time for current CMTS with 4 and 8 downstream channels per port; then speculates where CCAP with 16 and 24 downstream channels per port might be positioned relative to current CMTS pricing.

To further explore this, Motorola developed an economic model for CCAP deployments around licensing algorithms. As discussed previously, a model where the full CCAP costs are paid up front will be difficult to justify on a cost per channel basis. On the other extreme, selling CCAP channels at the average price per channel based on a fully deployed product is also problematic. The system must be designed to support the full working load. If only a small number of channels are licensed to start, then vendors will lose money on initial deployments with no guarantees of future revenue. This would inhibit product development.

The ideal model required a licensing algorithm that would reflect the expected channel deployment. As referenced in [Howald], downstream capacity can be expected to continue at the 40-60% annual rate. Based on this along with an assumed starting point of 16 downstream channels per port, Table 1 shows how the downstream channel deployment is modeled.

Year	Total Downstreams	Incremental Downstreams
2013	16	-
2014	24	+8
2015	32	+8
2016	48	+16
2017	64	+16
2018	96	+32
2019	128	+32

Table 1 Downstream Growth

Note that this is reasonably close to the 50% growth per year that is often quoted.

Another factor that must be taken into consideration is that operators have a limited budget to spend in a given year. Infonetics forecasts show that CMTS revenue is only expected to grow 5% annually over the coming years while overall capacity above is growing at 50%. This implies that the CCAP downstream cost per channel must drop year over year (YOY) as larger number of channels are introduced in later years.

The results from our economic model are shown in Figure 4 and Figure 5 below. The baseline was 16 downstreams (DS) per port for the initial year and the average cost per downstream channel is shown for the sequence described in Table 1. Figure 4 shows the ratio with 16 DS being the 1.0 baseline. Figure 5 is interesting in that it plots the same data with a log scale. Even though Figure 4 shows each sequence getting progressively closer together, Figure 5 highlights that there is a roughly fixed percentage decrease YOY.

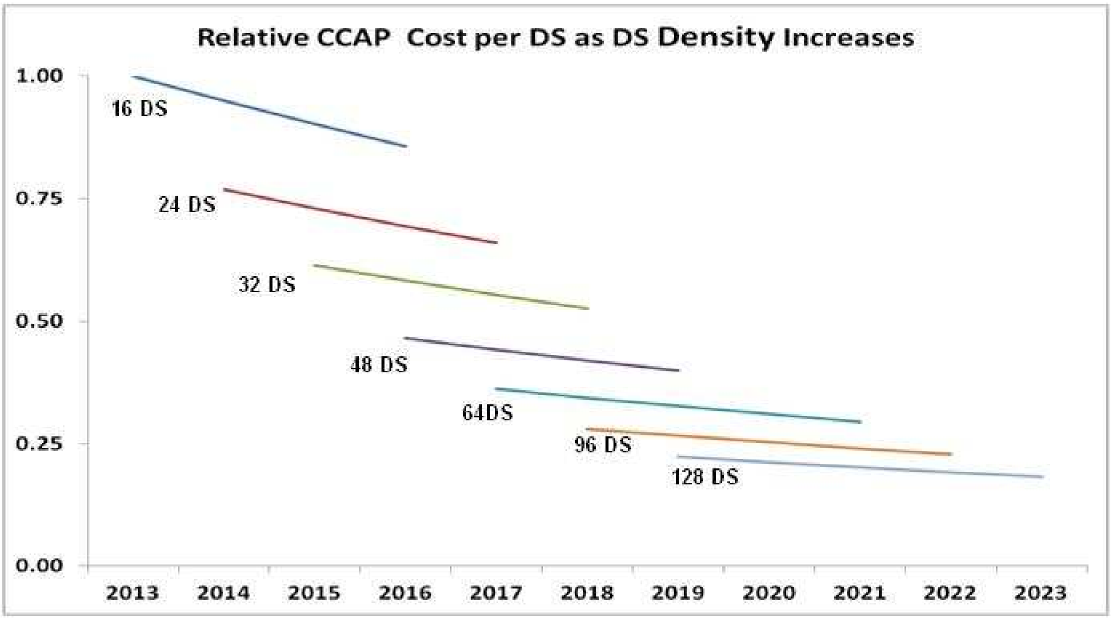

Figure 4 Cost Per DS at Higher Density

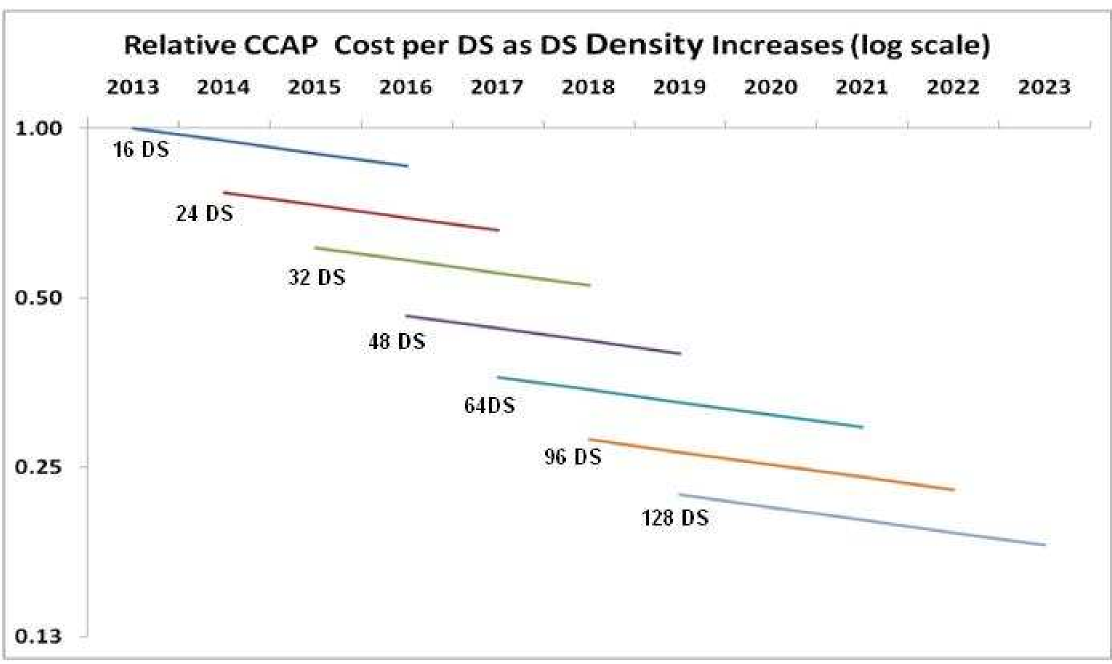

Figure 5 Cost Per DS at Higher density (Log Scale)

A licensing model like this is beneficial to both customers and vendors, assuming the initial starting point of 16 downstreams is sufficient to the vendor for initial installation and the YOY decrease in costs per downstream channel is sufficient to enable the operator to incrementally add channels in ever larger amounts within their budget.

Disclaimer: the above analysis is hypothetical and not based on any real products. It shows some possibilities for licensing algorithms that may be beneficial to vendors and customers. Every vendor may implement their own licensing algorithm and market conditions may cause these licensing algorithms to change over time.

As seen in Figure 4 and Figure 5, the economics around IP video deployment will vary over time. Costs will be higher initially but volumes will be lower. As IP video penetration ramps up, DOCSIS channel costs start to drop substantially.

Another important aspect is that IP video deployment is an incremental addition onto an existing DOCSIS HSD infrastructure. Therefore, it is critical to understand the incremental costs for downstream channels, not just the average costs which were previously discussed. This can be best explained by an example. Let's start with 16 downstreams as a baseline cost. Now suppose once there are 32 downstreams, the average cost per channel is 75% of the baseline cost per channel. In reality, the first 16 channels cost 100% and the incrementally added 16 channels were just 50% of baseline, giving a weighted average of 75%. So the incremental cost of 50% is the number that should be used for IP video economic analysis.

Taking the analysis further, CCAP leverages high-density EQAM technology. In the extreme, the incremental addition of a downstream channel could approach that of a high density EQAM product. Infonetics research shows that the average QAM cost was $163 in CY11 and forecasts that it will drop to $86 by CY16. Note that these are the average cost per QAM.

From our previous analysis, the incremental cost per channel could be substantially less. So it would not be a reach to suggest that the incremental cost per QAM several years from now may reach $40 per channel. This is an interesting number as the industry will approach $1 per Mbps for downstream bandwidth.

Working with this number for IP video economics, an IP video HD stream @ 5Mbps would therefore cost $5 to transport. Note that a few years ago this may have been $200-$400 using older CMTS downstreams. This radically changes the IP video economics. An updated chart with relative infrastructure costs is shown in Figure 6 below and shows the DOCSIS component has fallen from being the major cost contributor to become comparable to the other elements in the total cost. At this point other components become just as significant to the overall cost model.

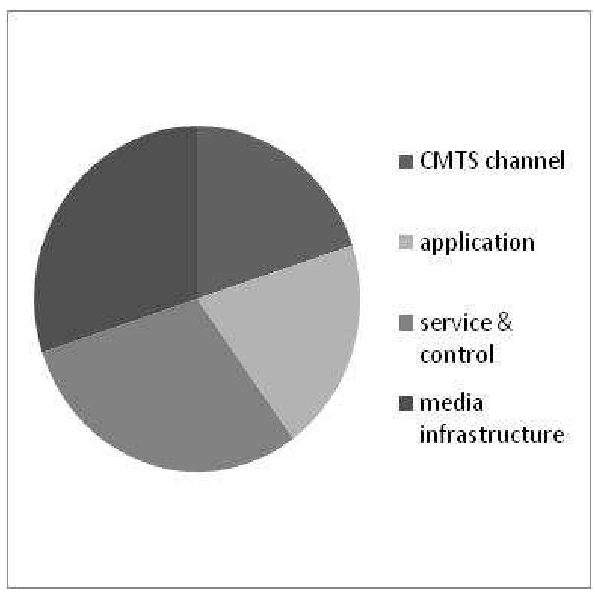

Figure 6 Post CCAP Cost Contributors

OTHER COST CONSIDERATIONS

CDN Options

As operators migrate to IP video services using ABR, they will be able to leverage internet CDN technology for video delivery. There are a wide range of options to achieve this with a corresponding range of costs.

Initially many operators may purchase CDN services from one of the worldwide CDN providers. Eventually an operator may enter into a wholesale relationship with that CDN provider in order to resell CDN capacity directly to content providers and web site servers. This may allow the operator to extend their brand to the CDN services as well.

A possible next step in the CDN progression would be to install a managed CDN. In this step CDN nodes are added inside the service provider network but are still managed by the CDN provider. This allows the service provider to deliver content internally on their own nodes and network while still leveraging global access through the CDN partner company. The service provider minimizes operational expenses (OPEX) since the CDN partner still manages the internal CDN.

Finally, the service provider can install a licensed CDN. Equipment and software are deployed on the service provider's network and the provider assumes responsibility for operations and support. At this stage, the service provider can participate in a federated CDN exchange with other CDNs to deliver content outside their own CDN.

Table 2 shows the various functions associated with each of the three approaches. From a cost perspective, the wholesale approach requires the least amount of up-front investment but it is also the most expensive on a per-bit-delivered basis. Each step then requires more investment from both a capital expenditure (CAPEX) and OPEX perspective, but continues to result in lower costs for delivering each bit of content.

Service Provider Investment in CDN offering			
	Wholesale CDN	Managed CDN	Licensed CDN
Sales	x	x	x
Billing	x	x	x
Hardware		x	x
Datacenter		x	x
Network		x	x
Support			x
Operations			x
Technology			x
NOC			x
Log Processing			x
Monitoring			x
Software			x

Table 2 Service Provider CDN Options

Transcoder and Storage Trade-offs

For linear television service, there is traditionally no storage costs associated with it. The content is encoded/transcoded, prepared and delivered to the consumer. With the new world of IP devices arriving, operators will want to go beyond simple linear television service to these devices and offer the ability to time shift. Consumers have become accustomed to their DVR for the television screen and will demand the same service for their IP devices. This will create a need for network based or "cloud" DVR services (nDVR).

Some current legal rulings based on existing content contracts require that nDVR content have a unique copy for each subscriber that records it. Other services offered today with re-negotiated content agreements allow single copy storage provided the fast-forward feature is disabled. The relative cost impact of nDVR is affected dramatically by the ratio between these.

Multi-rate ABR also exasperates the problem since a unique copy of a piece of content must now be stored in multiple bit rate formats. An example of this cost impact is shown in Figure 7.

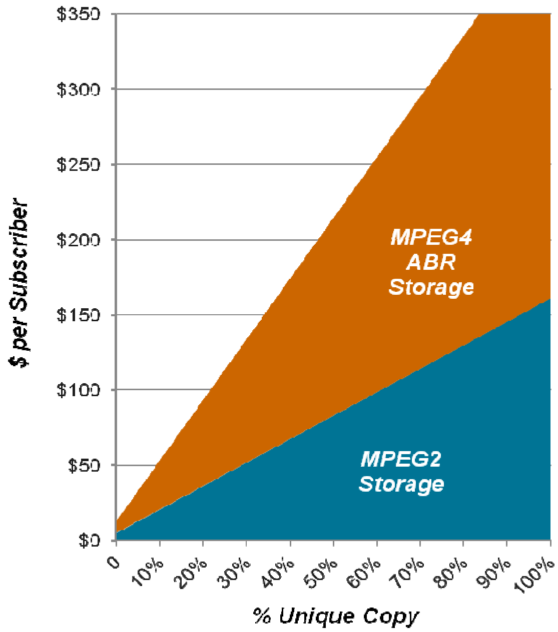

Figure 7 Storage Costs – nDVR

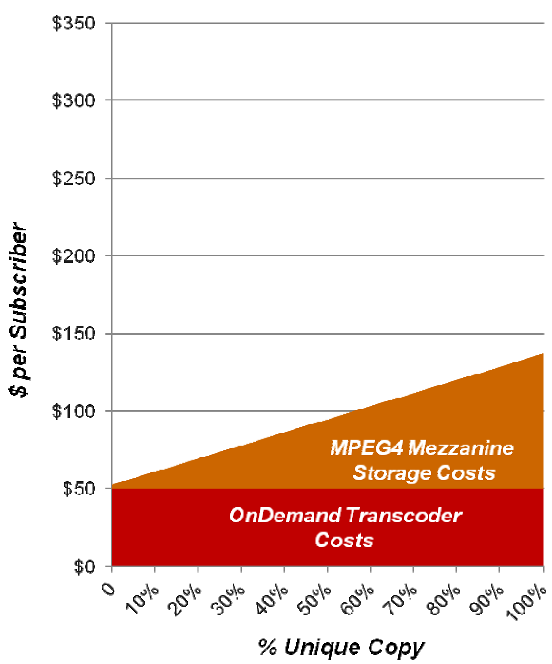

Figure 8 Transcoder vs. Storage

An alternative approach is to store a limited number of mezzanine formats in the nDVR storage and then transcode the content to the appropriate ABR bit rate on the fly when it is being viewed. Figure 8 shows an example of how costs may be impacted.

This creates a tradeoff between storage costs and transcoders costs that is constantly shifting. Many factors go into this analysis and the on-demand transcoder costs can vary significantly. This is an area where a disruptive change in transcoder costs could significantly change the landscape.

SPECTRUM MIGRATION STRATEGIES

Another very important aspect to IP video migration is finding sufficient spectrum. Some operators have already made more spectrum available by recovering analog TV channels using digital TV terminal adapters (DTA) while other operators have upgraded their HFC to 1GHz or used switched digital video (SDV). This available spectrum is being gobbled up today as more HD content is deployed, VOD requirements continue to increase and HSD services continue to grow at 50% annual rates. So there may still be a need for additional spectrum to ramp up IP video services with a corresponding economic impact.

Early Transition Plans

One way to significantly reduce spectrum requirements is to convert legacy MPEG-2 linear TV to IP video in a home gateway device. To support ABR devices in the home requires a transcoder in the home gateway device. Simple stand-alone devices are available today that accomplish this. This is an excellent approach for early deployments as it has almost no impact on infrastructure costs for rolling out linear IP video services. It also requires no new spectrum as this home gateway device appears as an STB to the system.

The next step in this migration is to introduce hybrid video gateways that also incorporate transcoding technology. These perform the same IP video conversion for linear TV described above for delivery to multi-screen IP devices. The video gateway also has the advantage that it is the single point of entry for video services and allows IP STBs to be deployed elsewhere in the home behind it. These devices can also operate as IP devices and are pivotal in the transition to an all IP system. As above, it can have a minimal impact on infrastructure costs to start and allows the operator to grow its IP video infrastructure at their own rate.

A detailed discussion of the home gateway migration is given in [CS_2012].

Complete Recovery of Legacy Bandwidth

The previous discussion on home gateway migration plans helps the operator begin the IP video transition. However, the end game is to eventually get to an all-IP system. Legacy MPEG digital TV services may continue to consume 50% to 80% of the available spectrum. Regardless of which path the operator took to free up spectrum, eventually they will need to install switched digital video (SDV) to reclaim all of the legacy bandwidth.

Adding SDV to the mix also increases the need for narrowcast QAM channels. This plays well into the previous CCAP analysis in this paper. Also, as the mix between legacy and IP subscribers change, an operator will need to re-assign SDV bandwidth to IP video bandwidth. This is also well suited for CCAP. A more detailed analysis of the SDV migration is in [Ulm_NCTA_2012].

CONCLUSION

Operators must deploy unicast ABR video to remain competitive. The infrastructure costs of providing this service are currently dominated by the cost of the downstream DOCSIS channels needed. With the development and deployment of high density CCAP platforms the cost per downstream is expected to fall dramatically, enabling the operator to deploy sufficient channels to meet demand while remaining within budget.

In the early days of CCAP deployment, not all channels will be used creating a potential disconnect between the capacity of the platform and the cost per channel deployed. The paper offers a framework for licensing which should be mutually acceptable to vendors and operators to circumvent this hurdle.

With the DOCSIS channel cost reduced significantly other cost components become more significant. ABR video is conveniently and cost effectively delivered via a standard internet CDN. A range of options to implement this are available from complete outsourcing to in house each offering different trade-offs in OPEX and CAPEX.

As nDVR is deployed into the ABR infrastructure another set of trade-offs will be required. For each recorded asset the multiple bit rate versions required can either be created at record time or created at play out from a recorded mezzanine format. In this case the trade off is between storage capacity and real time transcoding costs.

Operators will need to find the downstream bandwidth required for IP video delivery. Several options are available to do this. Home gateways may be used for early deployments in parallel with legacy MPEG-2 video. As the move to all IP video progresses the amount of MPEG-2 channels will decrease so that they can be economically delivered using SDV. CCAP is well suited for this.

The operator has multiple choices to make but will be able to deploy the technology required to remain competitive in an IP video environment.

REFERENCES

[SAND]	Global Internet Phenomena Report Fall 2011; Sandvine
[VNI]	Cisco® Visual Networking Index (VNI) 2011
[DIBA]	M. Patrick, J. Joyce, "*DIBA – DOCSIS IPTV Bypass Architecture*", SCTE Conference on Emerging Technology, 2007
[CS 2012]	J. Ulm, G. White, *"Architectures & Migration Strategies for Multi-Screen IP Video Delivery"*, SCTE Canadian Summit, March 2012.
[Ulm NCTA 2012]	J. Ulm, J. Holobinko, *"Managed IP Video Service: Making the Most of Adpative Streaming"*, NCTA Technical Sessions, May 2012.
[PACK]	Unified Content Packaging Architectures for Managed Video Content Delivery, Santosh Krishnan, Weidong Mao, SCTE Cable-Tec Expo 2011
Howald 2011]	Dr. Robert Howald, "*Looking to the Future: Service Growth, HFC Capacity, and Network Migration"*, 2011 SCTE Cable-Tec Expo Capacity Management Seminar,, Atlanta, Ga, November 14, 2011
Howald 2010]	Dr. Robert Howald, "*Boundaries of Consumption for the Infinite Content World"*, 2010 SCTE Cable-Tec Expo, sponsored by the , New Orleans, LA, October 20-22, 2010
[Howald 2012]	Howald, Ulm, *"Delivering Media Mania: HFC Evolution Planning"*, SCTE Canadian Summit, March 27-28, 2012, Toronto,

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
EAS	Emergency Alert System
EQAM	Edge QAM device
Gbps	Gigabit per second
HFC	Hybrid Fiber Coaxial system
HSD	High Speed Data; broadband data service
HTTP	Hyper Text Transfer Protocol
IP	Internet Protocol
nDVR	network (based) Digital Video Recorder
OTT	Over The Top (video)
STB	Set Top Box
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VOD	Video On-Demand