Managed IP Video Service: Making the Most of Adaptive Streaming

John Ulm & John Holobinko
Motorola Mobility

Abstract

The paper describes how an operator can leverage adaptive streaming protocols that are used today for unmanaged over-the-top (OTT) content for a complete managed IP video service. The paper describes how this solution is simpler and without some of the challenges imposed by implementing multicast delivery. Motorola's IP video modeling data shows compelling results regarding the relative benefits of adaptive versus multicast.

The conclusions and illustrations presented in this paper will help operators better understand how to: 1) initially deploy managed IP video services via DOCSIS, 2) plan their bandwidth and network resource requirements, 3) support existing video services in IP, and 4) optimize the network resources required as IP video viewership grows from small numbers to ultimately become the predominant means of video delivery in cable networks.

INTRODUCTION

Adaptive streaming is the primary technology for delivering over-the-top (i.e., unmanaged) IP video content to IP devices such as tablets, smartphones and gaming devices through the operator's Data Over Cable Service Interface Specification (DOCSIS) network. Adaptive streaming is the defacto delivery mechanism for OTT services. For managed services however, there is a popular assumption that multicast streaming video should be the principal delivery format to primary screens, not adaptive streaming. However, delivery of managed video in multicast format creates significant complexities for the operator, not the least of which are how to duplicate existing and planned services such as targeted advertising and network-based DVR, amongst others, and managing different segregated service group sizes compared to data services.

This paper presents a proposal to employ a comprehensive *managed* IP video services solution using adaptive streaming protocols with appropriate enhancements. An end-to-end multi-screen IP video architecture is presented, including the role of these adaptive bit rate (ABR) protocols.

The trade-offs of using adaptive streaming versus multicast for delivering managed video services are discussed. One of the other major concerns of operators is the bandwidth that will be required to deliver managed IP video services. Many factors come into play with the introduction of IP video, and our modeling results show that multicast gains may evaporate, so there is no penalty for using unicast-based adaptive protocols.

MANAGED IP VIDEO ARCHITECTURE

Multi-screen IP video delivery requires an end-to-end ecosystem that must encompass data, control and management planes. It must interact with legacy encoding, ad insertion, and content management systems while operating in parallel with traditional linear broadcasting. Operators will migrate towards multi-screen IP video to deliver content to a new generation of consumer devices such as tablets, smartphones and gaming devices; and to enable new cloud based services to attract and retain customers.

[Ulm_CS_2012] described an end-to-end conceptual architecture to support the evolution to IP video delivery. This architecture is segmented into Application, Services & Control and Media Infrastructure layers. Each of these layers is further decomposed into functional blocks.

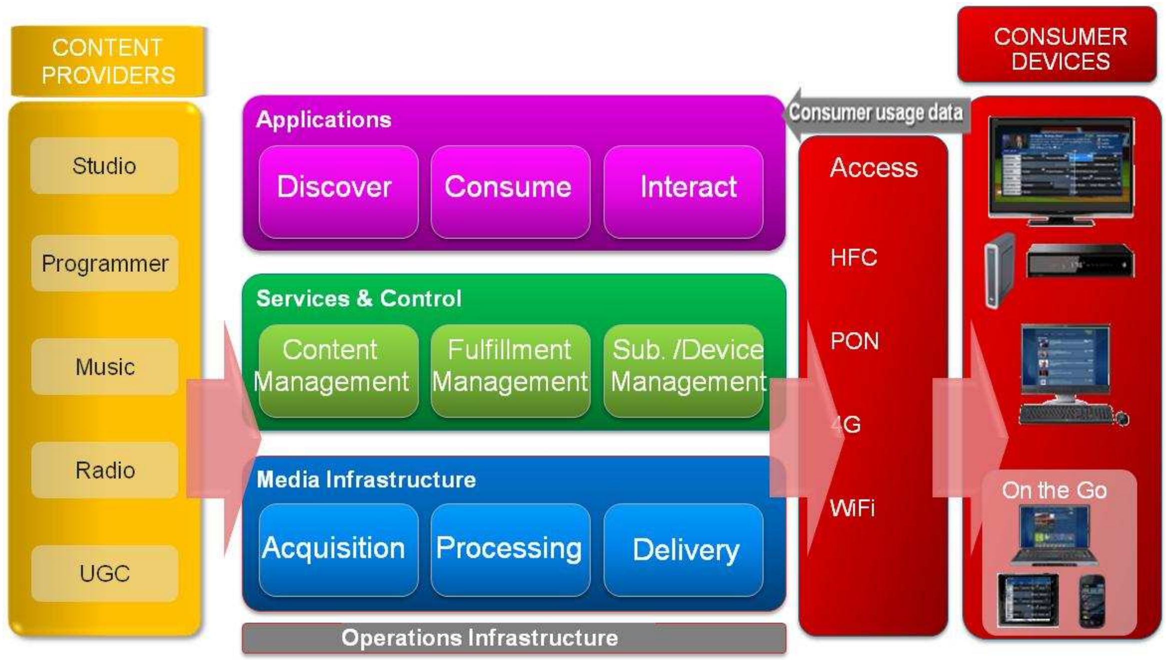

Figure 1: High Level Conceptual Architecture

Figure 1 shows a high-level abstraction of an end-to-end functional architecture for the delivery of 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.

The representation breaks the functions into three primary layers: Applications layer; Service & Control layer; and Media Infrastructure layer. A fourth functional block called Operations Infrastructure overlays the three primary layers.

<u>Application Layer</u>

The Applications layer provides interaction with the end user and is largely responsible for the user experience. It includes functions that discover content through multiple navigation options such as user interfaces (UI), channel guides, interactive search, recommendation engines and social networking links. It enables the user to consume content by providing applications for video streaming, video on demand (VOD) and network DVR (nDVR) consumption. These applications integrate with the Service & Control layer to authenticate the user, confirm access rights, establish content protection parameters and obtain resources for delivery as required.

The Application layer also provides companion applications which enable user interaction in conjunction with media programs. These may be as simple as allowing interactive chat sessions among viewers watching the same program or enable more complex integration with social media applications. It also enables enhanced monetization with new advanced advertising capabilities such as telescoping ads.

<u>Services & Control Layer</u>

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 through all phases of its lifetime including ingest, transcoding, digital rights management (DRM) and advertising insertion policy. Other functions manage the fulfillment of user requests for content delivery by providing resource and session management, nDVR and VOD management and Emergency Alert System (EAS) and blackout support. Finally, it must manage subscribers and devices to ensure content delivery to authorized consumers in a format compatible with the consuming device.

The Services & Control Layer provides a unified approach for managing entitlements, rights, policies and services for the multitude of devices and DRM domains expected in the emerging adaptive streaming IP video service model. This solution must provide a mapping function between the billing system and the DRM system interfaces, recognizing that leveraging existing billing interfaces provides for a more seamless transition from legacy solutions. Billing should focus on account level transactions – allowing the network and associated DRMs to determine if content viewing is allowed on a specific account or a specific device. A tight integration with compelling DRM solutions is a necessity. By abstracting the complexity of a multi-DRM system, the Service & Control layer efficiently manages entitlements, rights, policies and services for a multitude of devices across a number of DRM domains. These unified provisioning functions will provide an essential building block for end-to-end multi-screen video solutions. For a detailed discussion on this topic, see [Falvo_2011].

Media Infrastructure Layer

The Media Infrastructure layer is responsible for managing video content flow and delivering the media. It includes content ingest, preparations, and delivery to the devices. Functions in this layer acquire content from satellite or terrestrial sources as either program streams or files and encode it for ingest into the system. It processes the content to prepare it for delivery. This includes functions such as transcoding, multiplexing, advertising insertion, EAS, black outs and encryption. Finally, this layer delivers the content to the target device through mechanisms such as Web servers, content delivery networks (CDNs), and streaming servers.

It is in the Media Infrastructure layer where the decision is made on video delivery protocols. For ABR distribution models, this layer includes packaging into appropriate file formats, manifest creation and publishing to a CDN origin server.

The remainder of this paper takes a detailed look at managed content delivery using adaptive bit rate (ABR) protocols.

ABR BENEFITS FOR MANAGED IP VIDEO SERVICE

Using ABR for IP video delivery can be considered a “pull” delivery model in which the end client requests the video data. With ABR, the video content is broken up and stored in a CDN as a series of small files at multiple different bit rates. The end client uses standard HTTP “get” requests to download each file segment into a local buffer from which the content is played out. The client monitors the rate at which downloads are occurring and the available locally buffered content to determine which bit rate to request. If the network is fast, a high quality high bit rate will be selected. If the network is slow, a lower quality, lower bit rate option will be requested. This is an inherently unicast service as there is no coordination between clients (even if they are watching the same content at the same time, two clients would download it independently). A tutorial on ABR for cable may be found in [Ulm_2010]. Below is an in-depth look at many key considerations and benefits in using ABR for a managed service.

CPE: Right Choice for Second/Third Screens

A key driver for migrating to IP video delivery is the ability to deliver services to a wide range of IP devices, in particular personal computers, tablets, smartphones and gaming devices. Operators want to offer these services to remote subscribers who are "off-net" as well as managed IP video services to devices inside their own network. The protocols are applicable to both linear television and on-demand delivery.

ABR protocols are the best choice for these smaller screen devices and off-net operations. They have very simple customer premises equipment (CPE) clients that adapt dynamically to changing internet resource availability. With extremely high churn on CPE devices, it is very important from an operational perspective to support the embedded client on new devices. ABR protocols are becoming the de facto standard for IP video delivery to these devices. With ABR, the operator will not become the long pole in the tent while trying to provide device drivers for the newest gadget of the week.

In-Home delivery of managed IP Video

Since ABR protocols use HTTP, they are extremely well suited for traversing home firewalls. This is in stark contrast to multicast delivery through consumer owned routers. This means that ABR is much better from an operations and support perspective.

The other issue with in-home delivery is that it may span a consumer's home wireless network with unpredictable latency and throughput. The ABR protocols are also well suited to adapt to this environment.

CDN Considerations

There are some CDN considerations that the operator must review when architecting an IP video delivery system. Traditional VOD systems today use a "push" model where streaming content 'pushes' through the system in real time. This approach supports multicast delivery, but requires session management and admission control to secure resources, guaranteed bandwidth from the server to the client, CBR-based video, and dedicated servers.

The server has the added constraints of maintaining correct timing for transmitting content. Any network-induced jitter must be removed by the edge device (edge QAM or set-top box). This approach uses a non-robust transport (e.g. UDP or RTP) which requires added complexity to detect and recover from errors. Because of all of this, a push CDN model cannot exploit general internet CDN technologies for access network delivery.

In an adaptive streaming world, clients "pull" content from the CDN as files or file segments using a reliable HTTP over TCP transport. The client pull approach is CDN friendly and allows operators to re-use HTTP-based Web caching technology that uses standard servers. The CDN caching reduces backbone capacity requirements for both linear and on-demand content. Multicast only reduces backbone traffic for linear content. All of this gives the operator significant cost benefits by leveraging internet technologies. Its state-less architecture also readily scales as needed.

To summarize, a pull CDN model provides the operator with a simpler, more cost-effective system that uses a single IP infrastructure. It leverages internet technologies for performance and resiliency. It supports ABR and enhanced quality of experience (QoE) from a common infrastructure. The operator is able to incorporate public and third party CDN services with its private CDN. Finally, this scales to a global delivery model.

Quality of Experience Considerations

In offering a managed IP video service, QoE is an important consideration for operators. One of the key factors is how the system reacts to congestion. With the high levels of compression in today's video streams, any lost packets can have severe impact on the user's experience. Implementing a multicast- based streaming service puts significant additional burdens on the operator's system. As mentioned earlier, multicast streaming is based on non-robust protocols, so in a heavily congested environment they might lose packets. The operator could choose to over provision the amount of bandwidth needed to prevent these conditions, in which case they are throwing away potential capacity gains from using multicast. The alternatives are to implement some combination of admission control and/or error recovery. An admission control algorithm will be further complicated if variable bit rate (VBR) video delivery is used to maximize bandwidth savings rather than constant bit rate (CBR). An error recovery system introduces new servers into the network and requires custom clients in the consumer devices. Overall, the design, deployment and operation of a multicast-based system are inherently complex.

ABR protocols were developed for Internet delivery with its constantly changing throughput. ABR seamlessly adapts to this varying environment. In a managed network with infrequent periods of congestion, ABR reduces its bit rates during these periods to compensate. The impact on QoE might be comparable to that of running legacy MPEG video through a statistical multiplexer (statmux), which is familiar to operators. ABR also is based on a reliable TCP protocol that has error recovery already built into it, so any packets lost during congestion are automatically retransmitted. Thus, it prevents blocking and other video artifacts that significantly impact QoE. In this case, no network resources need be reserved in advance for the service and ABR reduces or eliminates the potential for blocking. Using adaptive protocols for all IP video delivery helps the operator's overall system become much simpler. More on this topic can be found in [White_2012].

Another QoE consideration is the impact of channel change time. ABR protocols are well suited to fast change times as they can quickly load lower bit rate streams and then switch to higher bit rates as bandwidth is available. Using multicast delivery requires separate additional bandwidth and a proprietary protocol to quick start the video delivery in parallel with the multicast video.

Advanced Services

Another key reason for migrating to IP video services is the ability to offer new advanced services. In particular, this might include highly targeted advertising such as personalized advertisements and telescoping. The system must also support EAS and blackout identical to legacy video services. Using its playlist manipulation, ABR provides the service provider with tremendous capability to re-direct a client on-the-fly with minimal effort and equipment. Supporting these advanced services using multicast delivery becomes problematic.

Miscellaneous Considerations

IP video penetration will occur over a long period of time. This means that the operator's home gateway will continually change during that time as well. Today cable operators have DOCSIS D3.0 devices in the field with 3, 4 or 8 downstream channels. Over the next several years we will see this expand to include 16, 24 and perhaps 32 downstream channels. The operator needs to manage this DOCSIS modem transition. Using ABR and its unicast delivery allows every modem to be in a bonding group suited to its capabilities; multiple bonding groups can then overlap, allowing the cable modem termination system (CMTS) to fully utilize the bandwidth. Multicast delivery runs into multiple problems in a mixed bonding group environment as discussed in [Ulm_2009].

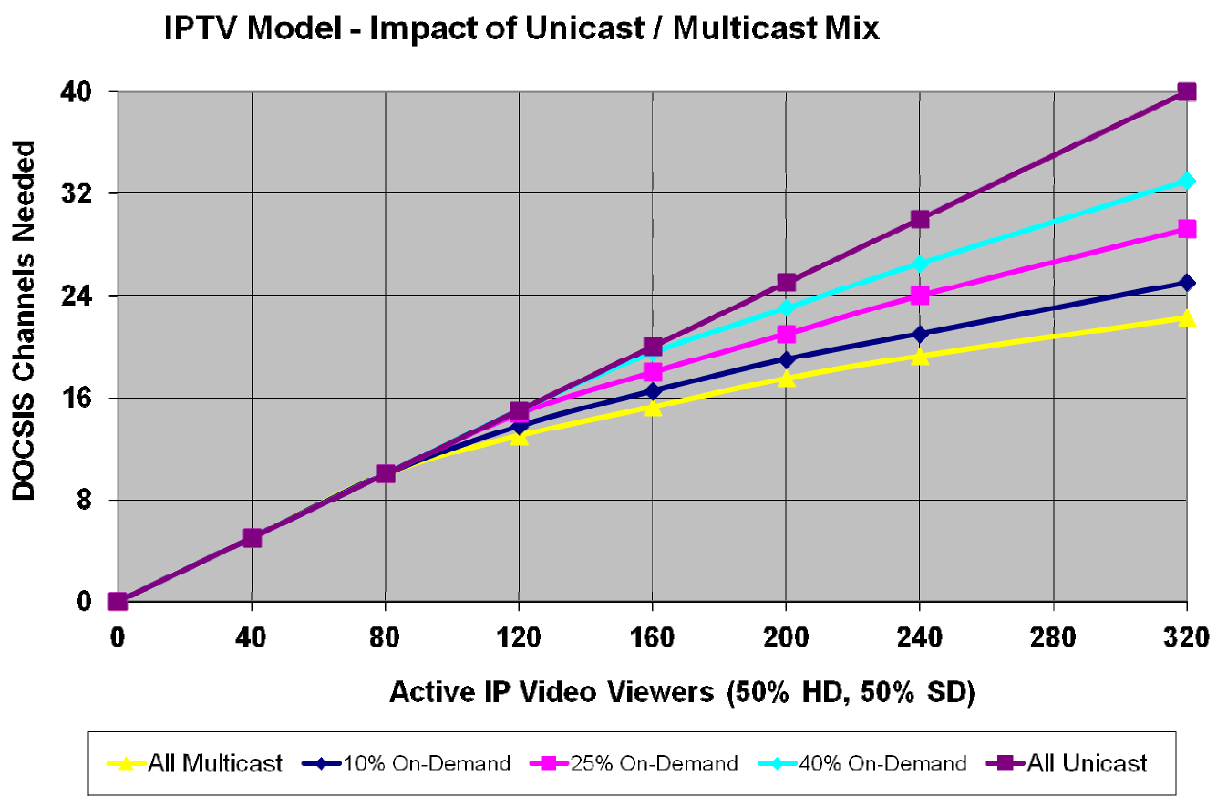

Figure 2: Impact of Unicast / Multicast Mix

ABR BANDWIDTH CONSIDERATIONS

A detailed analysis of bandwidth requirements for ABR compared to multicast was given in [Ulm_CS_2012]. The findings were that, under most conditions, multicast delivery will have little or no bandwidth capacity advantages over ABR unicast delivery. Figure 2 shows some results from that paper.

For early deployments of IP video, the penetration rate will be low. As indicated in this figure, there is no multicast benefit below 120 active viewers. With many operators considering phasing in IP video gradually, the operator also needs to factor in their plans for service group sizes. If the phasing takes 5-7 years, will the operator initiate node splits and cut service group sizes in half during that time? At the same time, increased VOD usage and the introduction of nDVR services might cause a shift from 10% to 25% or even 40% unicast usage. Figure 2 clearly shows what happens when the number of active viewers drops from 320 to 160 or 240 to 120 viewers.

Impact of Multi-Screen Delivery

This analysis was done for a two screen system: 50% of viewers watching high definition (HD) TV content and 50% of viewers watching standard definition (SD) TV content. With multi-screen delivery being a key impetus for IP video services, Motorola extended the IP video capacity modeling to see the effect of multi-screen viewing on capacity requirements.

Below are some sample outputs from the enhanced IP video capacity modeling. This looks at the bandwidth requirements for IP video for two different sized service groups as penetration grows.

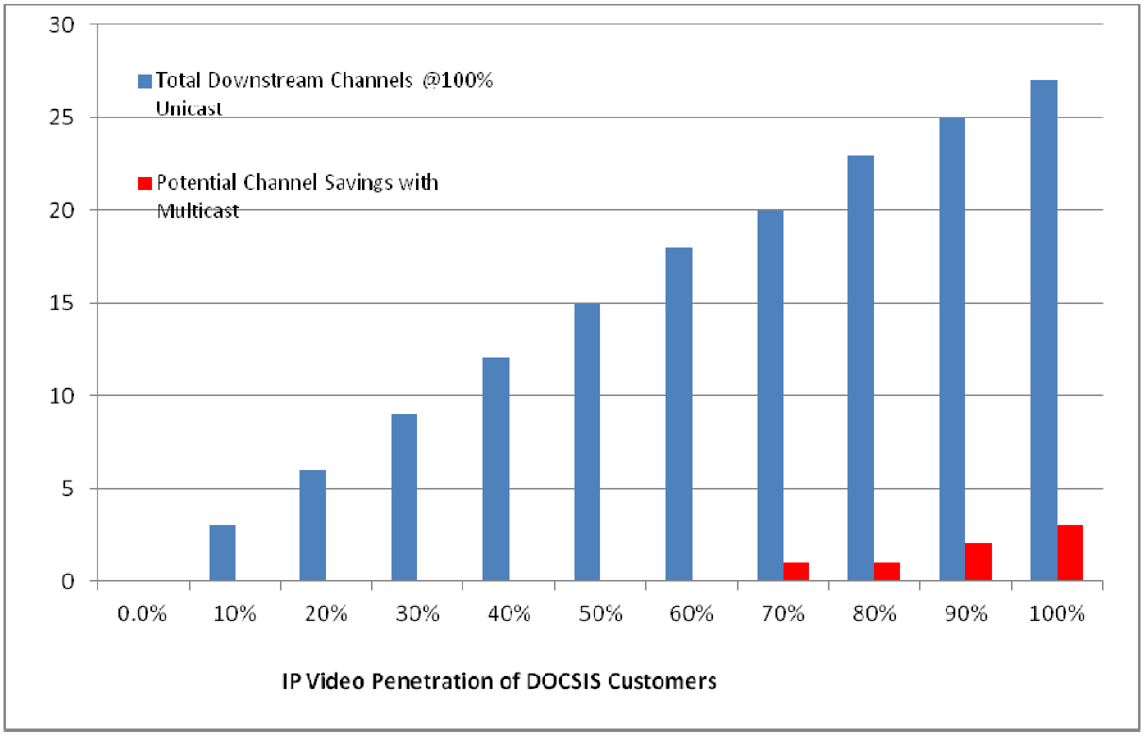

Figure 3: IP Video Bandwidth & Multicast Savings: 320 Active Viewers

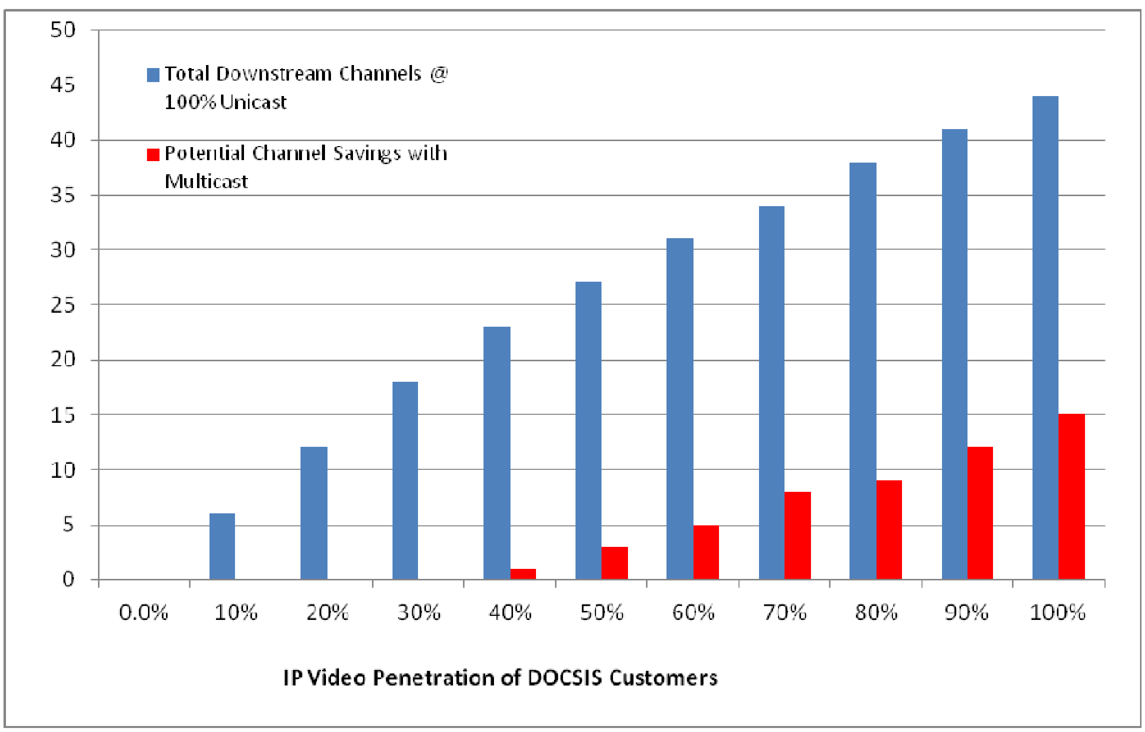

Figure 4: IP Video Bandwidth & Multicast Savings: 640 Active Viewers

In Figure 3, 100% IP video penetration corresponds to 320 active viewers which might represent a 500 homes passed (HP) service group, identical to the analysis above. Figure 4 doubles the service group size to 640 active viewers. In both these examples, viewership is spread across five different screen sizes: 30% HDTV, 30% SDTV; 20% tablets; and 10% each for two smaller screen sizes. It also assumes 25% on-demand usage which is reasonable if nDVR is deployed for the IP devices.

As indicated in Figure 3, the potential multicast gain is non-existent until the operator has reached 70% IP subscriber penetration. Even at 100% penetration, the multicast gain is only 3 channels or ~10% of capacity. This amount is almost negligible in a converged cable access platform (CCAP) environment capable of 64 channels per port.

In Figure 4, the serving group size is doubled. Perhaps the operator combined two fiber nodes to the same CCAP port to get additional multicast gains. Even with this extremely large service group of ~1000 HP, the multicast savings is still less than 20% at 70% IP penetration, yet it requires 34 DOCSIS channels of capacity for the large serving group. The small savings for multicast comes at a significant cost in spectrum used. It also comes in the late stages of the IP video deployment.

QoS in a Multicast Implementation

The purpose of implementing Multicast for delivering managed video content is to save bandwidth. By its very nature, a multicast system only makes sense if fewer channels of spectrum are required than a unicast implementation. Multicast designs are wholly dependent on the assumptions of multicast viewership during peak. At peak viewership, if more programs are being requested than the multicast service group was designed for, blocking occurs resulting in a denial of service. Therefore a prudent design calls for a safety factor in the number of QAMs reserved for the multicast service group. However this flies in the face of the rationale for implementing multicast, which is bandwidth savings.

In contrast, in a unicast implementation, if the bandwidth peak is achieved, the adaptive bit rates are lowered for the viewers in the service group. While video quality may lessen slightly in these cases, there is no denial of service. Therefore, unicast is a better choice for insuring a non-blocking service at peak usage times.

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 hybrid fiber coaxial (HFC) to 1GHz or turned to Switched Digital Video (SDV). This available spectrum is being gobbled up today as more HD content is deployed, VOD requirements continue to increase and high speed data (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 – Hybrid Gateways

One way to significantly reduce spectrum requirements is to convert legacy MPEG-2 linear TV to IP video in a video gateway device that includes a transcoder. This approach requires no new spectrum for linear TV as this video gateway device appears as a set-top box (STB) to the system and uses legacy broadcast content.

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 hybrid devices can also operate as IP devices and are pivotal in the transition to an all IP system. Longer term, the transcoding capability and adaptive protocols supported by the gateway may limit the quantity and type of IP devices supported in the home. Eventually the operator will want to support IP devices directly from the "cloud" using their network infrastructure.

A detailed discussion of the home gateway migration is given in [Ulm_CS_2012].

Complete Recovery of Legacy Bandwidth

The previous section on video gateway migration plans helps the operator as they 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 even after DTA and 1GHz upgrades. Regardless of which path the operator initially took to free up spectrum, eventually they will need to install switched digital video (SDV) to reclaim all of the legacy digital TV bandwidth.

Adding SDV to the mix also increases the need for narrowcast QAM channels. This plays well into a CCAP migration. As the mix between legacy and IP subscribers changes, an operator will need to re-assign SDV bandwidth to IP video bandwidth. This is well suited for CCAP. For a detailed discussion on IP video economics in a CCAP world see [Ulm_NCTA_2012].

Some SDV capacity reclamation modeling results are shown in Figures 5 and 6. Figure 5 shows the total spectrum required for legacy video services as the number of legacy viewers is reduced to zero. It assumes a video service with 180 HD programs (3 per QAM) and 200 SD programs (10 per QAM), so full broadcast requires 80 QAM channels. Figure 6 shows the corresponding SDV narrowcast requirements.

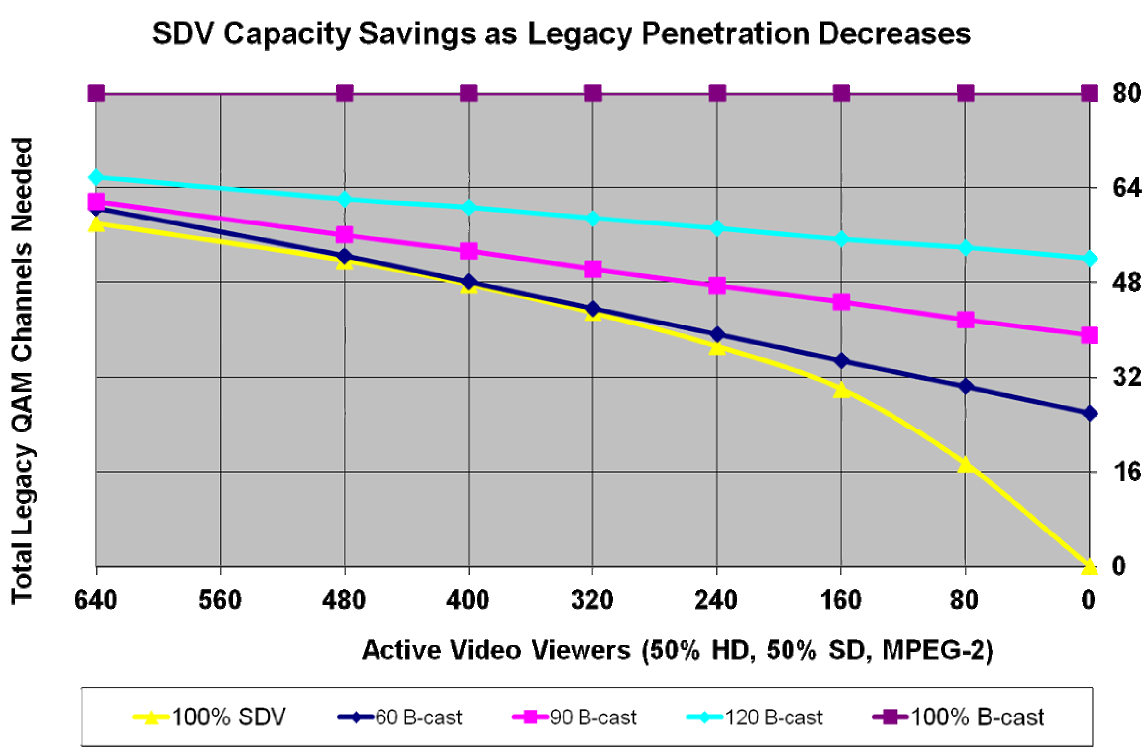

Figure 5: SDV – Total Capacity Savings with Decreasing Penetration

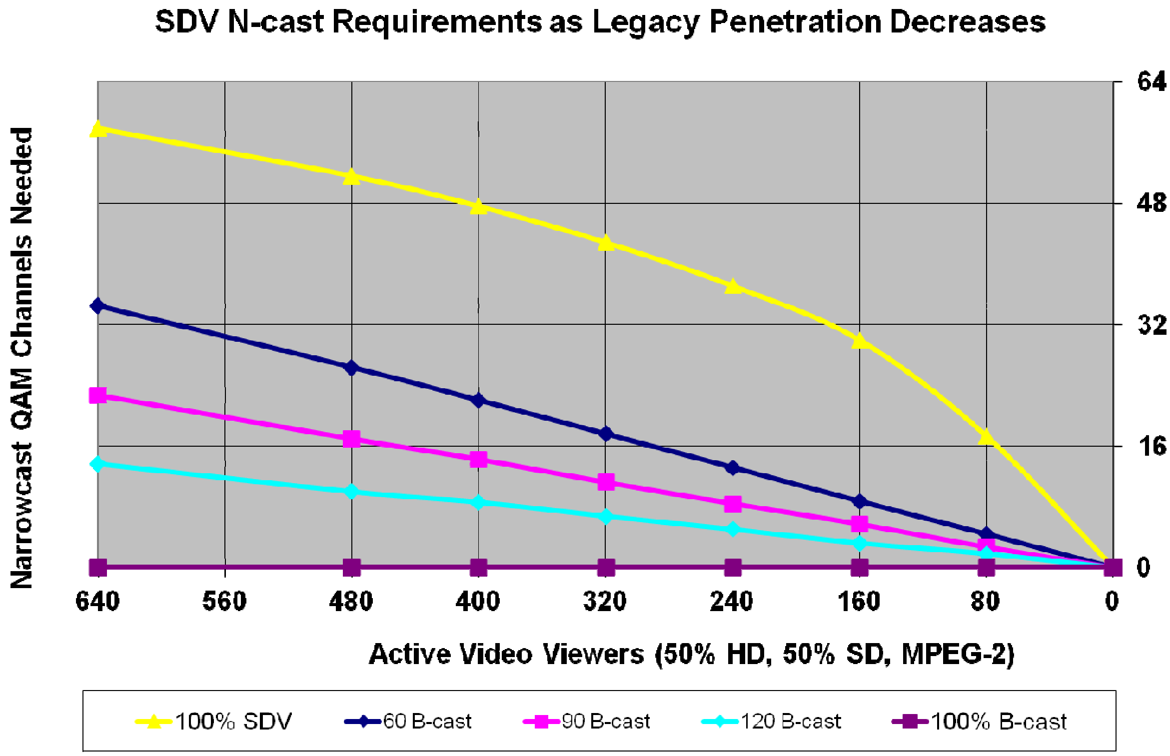

Figure 6: SDV – Total Narrowcast Requirements with Decreasing Penetration

Four scenarios are given varying the amount of switched content up to 100% switched. As shown, 100% switched provides the most bandwidth savings, but requires significantly more narrowcast. The operator has complete flexibility in trading off between spectrum saved and narrowcast QAM requirements. As can be seen in Figure 6, as the number of legacy viewers decreases, there is a corresponding decrease in narrowcast QAM requirements. This allows the operator to repurpose SDV QAM channels as they become freed for DOCSIS channels (HSD or IP video) or additional SDV savings.

It is informative to look at an example where the operator allocates twelve QAM channels for SDV and watch the impact as their legacy viewers are reduced. From Figure 6, the curve representing 120 broadcast programs and 60HD/80SD switched programs crosses 12 QAMs at 560 active viewers. Now looking at Figure 5, this scenario (i.e. 560 viewers, 120 B-cast) requires 64 channels of spectrum, freeing 16 channels (compared to 80 channels for 100% broadcast) for other usage such as IP video growth. As IP video penetration grows, legacy penetration shrinks. The next curve (90 broadcast with 90HD/110SD switched) on Figure 6 crosses 12 QAMs at 320 viewers. Mapping to Figure 5, this scenario (i.e. 320 viewers, 90 B-cast) only requires 50 channels of spectrum, so 30 channels are now available. The next scenario (60 broadcast with 120HD/140SD switched) crosses 12 QAMs around 200 viewers and requires ~36 channels for more savings.

As a result, the SDV spectrum savings are significantly more than multicast gains seen in the previous section. The SDV benefits are also available for small and large service groups. Every operator needs to consider SDV as a crucial part of its IP video migration.

CONCLUSION

Cable service providers will migrate from existing legacy video networks to a full end-to-end IP video system in a number of stages as new services are rolled out. They need to leverage the technology used for these intermediate stages into the final end-to-end system. Therefore, it is critical to have a layered architecture approach as presented in this paper that can isolate the changes between the various components.

Selecting the correct technology is particularly important for the delivery component of the Media Infrastructure layer as it is hardware centric, widely deployed and capital intensive. In particular, this paper focuses on the selection of adaptive protocols as the primary video delivery mechanism and discusses its benefits. ABR enables:

- A wealth of new and constantly changing IP devices
- Easily handles the home environment
- Provides excellent QoE to consumers
- Adapts to congestion without requiring complex admission control or re-try mechanisms
- Leverages internet CDN technology
- Readily supports advanced services including personalized advertising.

The updated IP video capacity modeling results shows the impact of migrating to a multi-screen environment. A 500HP service group may only get 10% multicast gain even once its switched to all IP video delivery.

Understanding the migration plan is a critical piece of the IP video architecture, especially with respect to managing available spectrum. Hybrid video gateways enable the introduction of IP video delivery with minimal impact on an operator's infrastructure. As the system scales, these devices transition to full IP video delivery.

Finally, the operator needs to plan the reclamation of legacy spectrum as they migrate to an all-IP world. This migration will eventually require the use of SDV. The modeling results show that the benefits of SDV are actually greater than the savings from multicast delivery.

In conclusion, the operator needs ABR for its first IP video steps when delivering content to second and third screens; i.e., tablets, smartphones, PCs and gaming devices. Adaptive streaming is the final solution the operator needs once there is an all-IP world with any content, anywhere, anytime, anyplace. We have shown that ABR also handles the transition years and is the only delivery mechanism needed for a managed IP video service.

REFERENCES

[Ulm 2009]	J. Ulm, P. Maurer, "IP Video Guide – Avoiding Pot Holes on the Cable IPTV Highway", SCTE Cable-Tec Expo, 2009.
[Ulm 2010]	J. Ulm, T. du Breuil, G. Hughes, S. McCarthy, "Adaptive Streaming – New Approaches For Cable IP Video Delivery", The Cable Show NCTA/SCTE Technical Sessions, spring 2010.
[Falvo 2011]	B. Falvo, D. Clarke, C. Poli, *"Supporting Multi-CAS and DRM Entitlements"*, SCTE Cable-Tec Expo, Nov 2011.
[Ulm 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, G. White, *"The Economics of IP Video in a CCAP World"*, NCTA Technical Sessions, May 2012.
[White 2012]	G. White, J. Ulm, *"Reclaimng Control of the Network from Adaptive Bit Rate Video Clients"*, NCTA Technical Sessions, May 2012.