CHALLENGES IN MIGRATING CABLE NETWORKS TO IP VIDEO

First Presented at the SCTE Canadian Summit

Jay Chambers, Principal Engineer, Network Solutions Group, BigBand Networks

Scott Loveland, Senior Director, Product Management, BigBand Networks

Introduction

A major challenge facing the cable industry is the need to evolve existing networks into a centralized IP video delivery system where video is video, regardless of its destination. This paper discusses several challenges and considerations in this evolution:

- Choosing the correct platform
- Finding the bandwidth on the HFC network
- Taking advantage of universal video edge QAM (EQAM) technology to alleviate bandwidth bottlenecks
- Changing the in-home equipment strategy to lower capital expenditures (CapEx), and providing services to unmanaged consumer devices
- Managing issues relating to home networking and connectivity
- Customizing and personalizing the video content
- Implementing a platform that delivers all of these services, without compromising the reliability and quality of deployed services.

Ultimately, implementing the right platform, bandwidth management, and Customer Premises Equipment (CPE) strategy will play a critical role in yielding a more flexible cable IPTV delivery mechanism that can address the personalization of video for the consumer while providing a superior video viewing experience.

CHALLENGES IN MIGRATING CABLE NETWORKS TO IP VIDEO

The Emergence of IP Video

Today's cable operators are in the midst of evolving their networks, changing both their business models and the delivery of services to subscribers. The core foundation of this change is IPV. Video systems can deliver both traditional RFTV as well as IP video, and eventually evolve to provide a personalized viewing experience.

RFTV provides traditional video services, including linear broadcast and Switched Digital Video (SDV) as well as Video On Demand (VOD) services. The content delivered is usually in an MPEG-2 format, but is rapidly moving to include MPEG-4 format to traditional and advanced Set -Top Boxes (STBs).

IP Video, however, provides linear broadcast, multicast and unicast services as well as on demand services. IP video is generally assumed to be delivered in MPEG-4 format but could also be delivered in MPEG-2 format.

Video delivery has been evolving towards a personalized experience in which the video program is not only targeted for the viewer but also for the device the viewer is using to consume the programming. Personalized video services could include addressable advertising, specific overlays (e.g. stock ticker, weather ticker, sports ticker, etc.) as well as the ability to form personalized mosaics. Personalization for the device could include the delivery of the video stream with specific codec and encoding profile, bit rate, screen size, resolution, content protection, and other elements without compromising the quality viewing experience.

Although cable operators are relatively new to the CPE side of IPV, they are fully experienced with the transport and back-office aspects of this technology. It has become widely accepted that operators will make a complete shift to IP Video on the network and in the home. It is a little more contested as to the technology and timeframe for which this will happen. There are many similarities between the digital video offering of today's networks, and the next generation of video services that will be offered to any device, anywhere, and anytime. There may be different acronyms and terms for the various components in the network, but the functions are typically the same. While current CATV networks use terms like "Conditional Access Systems (CAS)", broadcast or switched, they have thick clients or "guides", and relatively higher cost set-top devices; IP video networks speak of "Digital Rights Management (DRM)", multicast and unicast, have thin clients and browsers, and offer relatively lower cost, feature rich IP-STBs.

In addition to the STB strategy, there are additional benefits stemming from the migration to IP video technologies including cost efficiency and diversity. One example is all the devices in the home that have multimedia capabilities. Consumer electronics (CE) "IP connected" devices include TVs, Blu-ray players, gaming consoles (i.e. PlayStation®3, Xbox, Wii™), iPods, PCs, and a slew of media extenders and front-ends. Both consumers and the technology powering these CE devices are pushing for new abilities and services when using these devices. For operators, a major opportunity exists to extend their core business devices to differentiate their brand and protect and expand revenue streams. There are always two sides, and with this great opportunity comes great risk. If operators do not move and adapt to these technologies, others will - whether YouTube, Hulu™, or Netflix

With the discussion of what devices consumers will watch their video services on, there is also a question of how they will use the service. Many believe that with the proliferation of IP video technology, there will also be a shift in subscriber viewing habits. One of the notable habits is how the consumption of linear, broadcast video will convert to non-linear forms. The industry has already seen developments of the matter with applications like Digital Video Recording (DVR) and Time-shifted Television (TSTV). Viewing habits will not be the only element that changes how services are delivered, but the devices delivering the video will also contribute to this. For example, most of the CE devices previously listed are not capable of supporting broadcast or multicast transmissions, but rather only unicast transmissions.

This doesn't necessarily change the fact that someone is watching a traditional broadcast program in real-time, but does change the delivery of that content.

Adapting the networks to enable managed video services delivery to these devices is challenging, but not impossible. Some would even argue that the end-game doesn't involve a managed STB at all, but rather it will use CE devices as the terminal to display content. One of the challenges that operators will face as a result of the migration to IP video delivery to any device, is how to achieve that without alienating existing digital services and subscribers or disrupting business models. Just as there are hundreds of devices that could be supported, there is also a similar number of protocols and interfaces that would need to be supported. While there are proposed specifications underway to standardize the abilities and interfaces of these devices, it would be wise to deploy solutions that can address the legacy digital STBs, the soon-to-be legacy IP devices, and the new generations of standardized CE devices.

Finally, an additional challenge is achieving a gradual, smooth manner of implementation that matches the network evolution and resource allocation as viewing patterns shift towards IP and non-linear. The ability to leverage and re-use network resources such as EQAM devices, optimize spectrum, and avoid having constant rewiring and reconfiguration of these network resources is an important element in any successful migration strategy to optimize CapEx and minimize operational expenditures.

Convergance

When we look at a traditional Cable network today, there are silos within silos for each service. IP services, such as High Speed Internet (HSI) and Voice over IP (VoIP) services, use Data Over Cable Service Interface Specification (DOCSIS®), which have their own control planes and hardware, and use their own spectrum. Because of its fragile nature and expectation of service, VoIP tends to be a silo within the DOCSIS tier. Video is not all that different. Video services (like broadcast video, SDV, and VOD also have their own control planes, hardware, and spectrum. Within video services, most still have silos for each type of service (broadcast, SDV, VoD). IP and video, rarely meet, but instead remain as two parallel networks from the headend to the home. To move forward with the IP video revolution, IP and video cannot continue to work in silos, as shown in Diagram 1. There simply is not

Diagram 1: Video is Video - The Converged IP Network

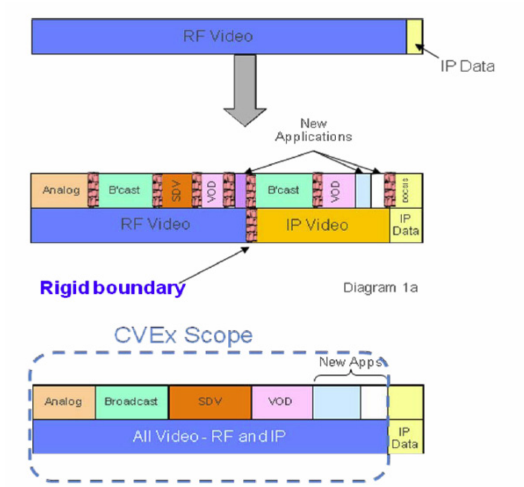

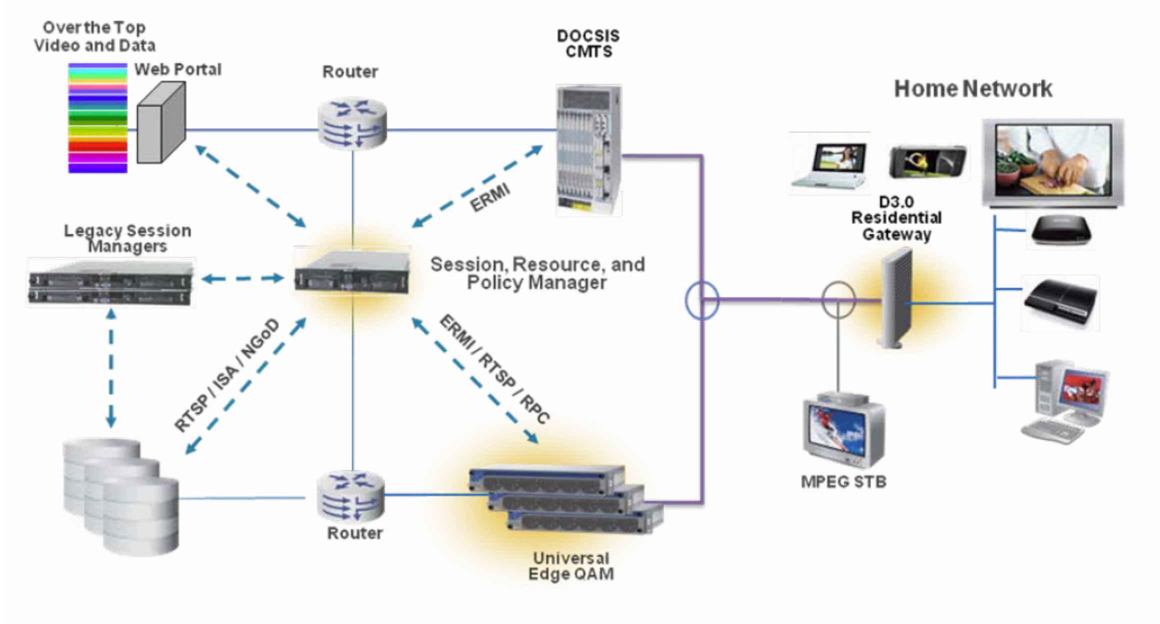

Diagram 2: Resource Allocation on a Common Control Plane

enough spectrum, operational systems and staff, or capital to continue on this path. There are opportunities for huge efficiency gains through pooling resources that have been traditionally built in separate parallel systems. Diagram 1 demonstrates how removing these silos will allow operators to centrally manage services, content and the associated bandwidth and resources across, what was up until recently, siloed parts of the network.

True convergence should address the video delivery to the millions of legacy subscribers nationally, as well as new IP video subscribers. It should support both linear and non-linear applications, from VOD, to Over the Top [OTT], to other services. In addition, it should identify devices, and their capabilities, and serve personalized streams acceptable for each device. All of this can be accomplished using common spectrum, common hardware, and a common control plane. Diagram 2 above demonstrates such an implementation, using the existing SDV control plane and EQAMs to deliver an operator's managed content, to any device in the home. We will next address the above in three distinct categories: Access and/or transmission systems, signaling and data control planes, and CPE strategies.

There is not a single technology that solves all aspects of video convergence. IP transport in the network currently relies on DOCSIS, with three approaches emerging in the IP video transition. The first approach is DOCSIS 3.0 Cable Modem Termination Systems (CMTS), which is widely deployed and the technology is well understood. But there is a significant cost associated with the increase in downstream ports. Second is the approach of Transport Gateways, where the traffic is not IP encapsulated in the headend but rather transported as an MPEG-2 transport stream, like standard video, and encapsulated in the modem. This requires non-standard capabilities in the CPE, whether modem or home gateway, which all comes at a cost.

Finally, a cost effective and less complex solution, known as DOCSIS Lite, offers the use of lower cost video QAMs while, at the same time, using a cost-effective DOCSIS 3.0 cable modem.

Universal Edge QAMs

In the access network, universal edge QAMs are used to deliver broadcast, SDV, and VOD services. Cable operators have already started to move toward the unification of the access network by selecting a few EQAM solutions for all services. In contrast, to the past method where operators deployed one EQAM model for VOD, one EQAM model for SDV services, and another for digital broadcast. Some operators are even moving their more expensive integrated CMTS QAMs to the universal EQAMs. Although Modular CMTS (M-CMTS) is a lower cost method of delivering CMTS downstream ports as compared to Integrated CMTS (I-CMTS), it is still not as cost-effective as ports used for video services. Using Session Resource Management (SRM) and/or Edge Resource Management (ERM) enable the support of dynamic QAM sharing using low-cost universal EQAMs. There has been a debate about whether M-CMTS traffic should be included in this mix. From an EQAM perspective, this is practical, given the number of QAMs being supported in a Service Group (SG) that uses DOCSIS 3.0 channel bonding. Likewise, there have been discussions about supporting both video and data within the same QAM channel. While this is technically feasible, it is not practical given the difference in traffic types. While data can be characterized as bursts of traffic that is unpredictable and tolerant to delay and loss, video is a high bit-rate, with continuous streams, which are susceptible to impairments from delay and loss. In addition, M-CMTS architecture still relies on allocating downstream ports within the CMTS, which adds to the cost. Reallocating the QAM portion of the spectrum does not relieve the need to pre-allocate the maximal resources needed for data in any given time on the CMTS shelf.

Not only are there differences in the type of content (data and video), but there are also differences in the systems currently used to transmit them. In the case of CMTS, the architecture is packet-based with packet processors, queuing and scheduling engines, switch fabrics, network transceivers, and QAM modulators and demodulators. In video QAMs, there are similar components, including network transceivers, QAM modulators, a packet bridge, and Digital Signal Processors (DSPs) used for video processing and multiplexing. Much of the video processing is done upstream in the network, using high-density MPEG multiplexers, and rate-shapers. Given the movement of service to a personalized stream offering more formats than before, it is logical to begin the migration of some of this processing to the edge of the network, closer to the subscribe

Video is Video

The video is video mantra continues with the concept of DOCSIS Lite. The premise being that the data plane for managed video services, linear or non-linear, are downstream only services. They only require two-way or upstream communication to setup the data plane. However, there is no need for most of the cost and complexity of today's DOCSIS downstream designs to support downstream video. This is accomplished by using a universal EQAM to ingest standard video streams from the network, and apply the DOCSIS header using the EQAM. This forgoes most of the processing done by today's DOCSIS MAC designs, and simplifies DOCSIS encapsulation for certain types of traffic that fall within this category.

To support legacy broadcast and narrowcast services, IP video multicast and unicast services, the number of QAMs required will push the limits of current technology for true convergence. Existing access and transport systems must evolve with the service offering. Next generation or high density EQAM platforms are designed to resolve pricing and operational issues while growing the access networks. The next challenge is addressing a more intelligent use of spectrum with technologies like SDV, DOCSIS Lite, Variable Bit Rate (VBR) and channel bonding. Intelligent spectrum technologies, combined with video processing and IP statistical multiplexing, calls for a new architecture. Such architectures are known by many names like Next Generation Access Architecture (NGAA), Next Generation Termination System (NGTS), and Narrowcast Termination System (NTS), but the general principal is the same. They combine upcoming DOCSIS MAC functionality, packet processors, and modulators or Physical Layers (PHYs) in the same chassis and support video and data services using high-density QAMs. With the change in EQAM modulation technology moving to Direct Digital Synthesis (DDS), there are up-conversion densities of 32:1 to 128:1. The next step is producing the entire spectrum using a single RF interface, or even optical outputs. This will help simplify hub wiring, service group splits, and the addition of QAM capacity for new services and allow for legacy video services to use the PHYs as needed for modulation. IP Video would get EQAM DOCSIS encapsulation for transmission, and DOCSIS data services would be based on a flexible MAC design using dense, lower cost PHYs.

Edge Resource Management in a Common Control Plane

Next generation QAM technologies only resolve the access hurdles; not the use of these QAMs. All narrowcast video services can be considered session based video, which requires an SRM to manage and set up these connections. Different services typically use control plane protocols and separate session managers. In addition, they usually have "dedicated" spectrum they operate within, where broadcast video consumes 20-40 QAMs, SDV usually has its own 8-24 QAMs, VOD operates within its own silo of 6-10 QAMs, and DOCSIS HSI service uses its own 4-8 QAMs. With IP Video, the natural tendency is to follow suit and silo IP services to their own QAM space.

To take full advantage of the large number of narrowcast QAM carriers offered by new EQAM technologies and the constant shift in viewing habits over time, operators need to be able to dynamically allocate edge resources between linear and non-linear services and between RF and IP networks. Traditional ERM technologies and offerings focus only on resource sharing within the RF spectrum. But as subscribers migrate to IPTV, resources must be reallocated to address this shift. Only a wider scope of edge resource management tools that can manage the entire narrowcast spectrum as a single pool can provide a true migration path to IPTV. The result is that from a narrowcast perspective, there are no QAMs allocated for each service, but rather a QAM SuperGroup™ for all services. Diagram 2 demonstrates what a resource context would look like, and how it might manage bandwidth.

In order to avoid disruption to existing technology, this "intelligent" control plane convergence needs to first consider the deployed SMs already in place. By initially acting as a front-end to the various SMs that already exist, the result is the native ability to act in a single context. Its platform architecture provides ERM and SM functions and allows IP Video, SDV, VOD, TSTV, and NPVR to all share common spectrum. For instance, within a single QAM channel, there could be an MPEG-4 IP video session, an MPEG-2 SDV session, and a TSTV session. This common control plane approach provides true capacity based migration and makes use of the new QAM platforms.

The Evolution of Customer Premise Equipment

CPE is the final category where convergence needs to take place. There are several approaches to convergence at the CPE level. A high-end digital set-top box acting as a gateway device is one approach, which then relies on tethering lower cost IP STBs with integrated Multimedia over Coax Alliance (MoCA®) interfaces. An IP pure play, where all digital STBs are phased out and are replaced with IP-based STBs is another approach. A third option is the cap and grow architecture of capping the deployment of digital devices into the field, and providing IP STBs to upper tiers and new subscribers. All of these methods are appropriate for today's cable operators. The CPE decision will likely be based on demographics, installed base, and what technology is available, thereby making them all viable options. To further complicate the situation, new CE devices have integrated IP connectivity and the ability to access OTT video content. The co-mingling of these devices in the home network presents additional challenges.

IP-based STBs for linear programming use Internet Group Management Protocol (IGMP) or multicast for channel selection. For unicast based services, like VOD, this is typically done with the use of Real Time Streaming Protocol (RTSP/RTP). In contrast, legacy based system use SDV protocols like Internet Security and Acceleration (ISA) or Next Generation on Demand (NGoD), and Digital Storage Media Commend and Control (DSM-CC) or other proprietary protocols for VOD. CE devices might use Digital Living Network Alliance (DLNA) standards, HTTP protocols, and RTSP protocols. To support the broad array of protocols originally supported in subscribers' homes, operators must adapt their networks. This can be done within the network or by using home networking and unified standard based protocols in the network, like SDV and VOD. This would be achieved by having a managed device or gateway in the home on CE devices using their native protocols and translating them into unified protocols that simplify the operation within the networks. Beyond just this simplification of protocols used in the network, there are other reasons that home-based protocols should not be extended into the operator's network. Some have never been designed as a routable protocol but rather a link-local protocol that speak to the devices within the home network. Security is another key concern as some protocols are less secure than others and might compromise the network.

There is some discussion around putting certain features, like transcription, transcoding, and other video processing features into a transport gateway. An area of concern, and for some, the motivating factor for moving from legacy systems to IP systems, involves the cost of a digital STBs with comparable features, and the Close Air Support (CAS) and back-office systems used to operate them. In other words, the aggregate costs of the legacy CPE devices are prohibitive. There are functions that absolutely should be put into the CPE device or in the home. Other features are better fulfilled within the network, allowing for cost efficiency, and maximum flexibility for growth. By attempting to incorporate too many features into a CPE device, the redundancy element is repeated. In addition, too many features often become obsolete or fall out of favor in light of new technologies. As the industry has replaced legacy STBs because of codec support, speed, memory, and storage reasons, it would make sense to replace a device in the home because it is faulty, not because of any missing features.

The Move to Personalization

As services like NDVR, TSTV, Mosaics, and other personalized services are offered, the appropriate solutions should be selected to promise these services are not only targeted to low-penetration of next-generation devices, but also to widely deployed legacy devices as well. Personalized services will become key to the converged network, and so will the new subscribers that will come with it. Personalization comes in a couple of different forms. Stream personalization will become important with the advent IP video and third screen applications. Just as there were hundreds of CE devices, and a subsequent number of protocols, there will also be requirements for content to be delivered and presented in different formats. These format changes might be the actual codec used, like MPEG-2, MPEG-4, DTS, or AC3, or they might simply target to keep the elementary stream intact and change the transport container. While MPEG2-TS is universal within cable systems, it is not within CE devices. The supported containers might have MOV, MP4, or MPEG-7 formats. Along with these changes, it is likely that they will need to be transcoded to support different resolutions and bitrates.

Along with the personalization of the streams, content personalization will also be offered. This includes addressable advertising, to any device, not just legacy STBs. With personalized ads, another widely used feature is customized guide environments, and personalized channels. These mosaic features are much less standardized as to how they are supported from device to device, driving the ability to the network.

As more video consumption devices are introduced, there is a growing need to customize and personalize the video content. This includes matching the resolution, encoding, DRM, transport and control protocols to the capabilities of the end device. In addition, by personalizing streams to include addressable advertising, overlays, personalized tickers and messages and overall look-and-feel, operators are able to generate more revenues both directly from subscribers as well as from advertisers.

The migration towards a more non-linear, session-based video delivery mechanism introduces the opportunity to deliver linear content as session-based personalized stream. The combination of a converged video control plane that fields any video session request and a unified data-plane device that can personalize video streams and turn multicast streams into unicast personalized streams, provides the migration path to video personalization.

Such architecture can alleviate the challenges of having a consolidated ad insertion system. Since all sessions are services from a single context using the emerging SCTE-130 advertising protocols, a single advertising back-office server can provide advanced advertising capabilities to all video applications on all devices, whether legacy or IP.

Summary

A graceful migration to an all IP network poses many new challenges, most notably delivering video to any device, anywhere, at any time. Video has no tolerance for even the smallest error, and is the ultimate "always-on" service with subscribers demanding the utmost in quality. IP dramatically changes both elements - it changes the experience of video (by making it personalized, on demand, with value added revenue-generating capabilities), and at the same time it transforms the network itself because it now has to deal with significant amounts of bandwidth and become media aware in order to meet the subscriber required experience.

Additionally, multi-screen delivery means new requirements for more dynamic monitoring and response to bandwidth and spectrum needs and new tools to monetize these more personalized services. Much can happen on the networking side, especially with respect to bandwidth management and media processing. There are opportunities for huge efficiency gains through pooling resources that have been traditionally built in separate parallel systems. This allows operators to centrally manage services, content and the associated bandwidth and resources across.

The points in this paper serve to address the challenges of migrating to an all IP network, including:

- IP video's ability to address lower cost CPE and the need to bring video to CE devices in the home with the growing demand for personalized video
- HSI and video convergence onto the same platform and moving away from the existing technology silos
- The opportunity for universal edge QAMs to provide a low cost solution for delivering video and data services
- Video is video and it is not necessary to use high cost CMTS downstream capacity as a delivery mechanism
- Converging ERMs and SMs into a single control plane, resulting in a capacity-based system, making use of new QAM platforms and providing a complete architecture
- The combination of a policy engine into the ERM and SM, providing a mechanism to identify a device, and its abilities
- The role of consumer electronics in driving the architectural considerations of the IP video network
- Video personalization, which will demand more non-linear, session-based video delivery mechanisms

United States Corporate Headquarters

475 Broadway Street
Redwood City, CA 94063
United States
phone: 1.650.995.5000
fax: 1.650.995.0060
bigbandnet.com

1413-0610