PRACTICAL VIDEO OVER DOCSIS® IMPLEMENTATIONS WITHOUT FORKLIFTS

Doug Jones, Chief Architect
Roger Slyk, Director, Product Marketing
BigBand Networks

Abstract

This paper describes a scalable IPTV (video over IP over DOCSIS) solution for providing video content to a wide array of consumer devices that span TVs and PCs, and other device capable of receiving and processing video streams over IP. The solution is extensible to other last-mile edge networks including 3GPP, Wi-Fi, and WiMAX wireless networks, as well as going "over the top" on IP access networks managed by third-party providers. It defines and describes a distribution network that sources live broadcast and stored content from VOD (Video on Demand), PPV (Pay Per View), and Internet sources, with elements of centralized and distributed processing and management. The proposed distribution network utilizes existing network infrastructures and CMTS (cable modem termination system) platforms as the basis for the introduction of advanced services delivery, thus avoiding significant equipment change out, capital expense and operational disruption.

The authors describe key business and technology requirements that are intended to be a reference when making decisions about the IPTV system design in order to maintain the business viability of the solution. These include:

- *Ensuring the solution is agnostic of the last mile network;*
- *Enabling scalable, flexible and cost effective systems for wide scale deployment;*
- *Maximizing the seamlessness of the customer experience;*
- *Placing the intelligence in the network, and minimize the intelligence required in set-top boxes, to protect CPE (customer premise equipment) investment and maximize flexibility.*

The authors advocate that, while there may be many paths to an IPTV reality, the best evolution plan utilizes M-CMTS (Modular CMTS) designs that leverage existing DOCSIS 2.0 infrastructures while accelerating towards DOCSIS 3.0 broadband speeds and capabilities. The authors describe an evolutionary strategy based on an M-CMTS approach that can drive the availability of downstream DOCSIS 3.0 channel bonding, enable more flexible allocations of upstream to downstream traffic flows, and provide additional performance benefits. Additionally, other parts of the solution are already specified by the cable industry, including PacketCable™ Multimedia and Embedded DOCSIS.

INTRODUCTION

Consumers are rapidly demanding an experience that blurs the line between "lean-back" consumption of TV-based entertainment and "lean-forward" multimedia activities on their personal computers. It is not a single killer application, nor an intentional move towards new technology that is driving this rapid change in consumer behavior, but a merging of technology, content and consumer acceptance. The ability to watch similar content in a home theater and on a portable device is increasingly being demanded; the capability to access a broad array of content from many service providers is rapidly becoming expected as well. Cable operators must plan for the flexible delivery of any content from any source to any consumer device in order to remain competitive in the long term. An open architecture based on IP (Internet Protocol) is the preferable long-term solution.

Cable operators need the capability to deliver their services to more than just traditional STBs (set-top boxes) over more than just the HFC (Hybrid Fiber-Coax) network. Cable services are evolving to include to new services that will be delivered to a new generation of devices such as mobile phones, wireless devices and personal computers. The network infrastructure will need to be capable of obtaining content from broad range of sources including broadcast programming, VOD, PPV and emerging "new media" outlets such as YouTube, MetaCafe and Ziddio.

Figure 1: Delivery of video content over DOCSIS infrastructures is increasingly popular

But there is a great incentive to have as many of the new services as possible delivered to the existing base of MPEG-2 STBs as well as to a new generation of CPE with advanced encoding technology and high-speed interactive communications channels based on DOCSIS and DSG (DOCSIS Set-Top Gateway).

The proposed network architecture to deliver these services is based on a core network of video and communications services which provide IP-in, IP-out transport and processing and any number of edge networks which provide IP transport to the customer. The delivery of IPTV services over the existing HFC edge network will be based on high capacity, cost-efficient M-CMTS components, as well as low cost, scaleable MPEG-2 edge QAMs. Other edge networks could include 3GPP (3rd Generation Party Partnership) wireless networks, or any kind of high-speed IP access network which the operator does not even own.

To incent customers, this solution will offer broader ranges of content, time-shifting and place-shifting and greater opportunities for personalization

IPTV SERVICES

An IP service is a cable service delivered to an end user over an edge network capable of transporting IP. This makes no statement to the origination of the content; that is, it could be a traditional video service that an operator offers today, for example linear television service, or it could be a video clip from a so-called new media outlet such as a website or some other Internet-based source. These are both video services and they can be packaged into a service based on an operator business plan or objective, but either can be delivered to the consumer using IP.

IPTV services generally exhibit several characteristics including:

- content personalization through time-shifting, place-shifting and addressable advertising;
- expansion of programming choices by broadening the universe of content to include new media services available over an Internet connection;
- ubiquity of services due to the prevalence of networks which support IP transport.

Today, HSD and Voice are offered over the HFC network as IP services. The only cable service not offered over IP on the HFC is video although video services are distributed around cable networks (for example backbones and regional area networks) using IP, but delivered to the end user over the HFC edge network using MPEG-2 transport.

Following the trends, video will also go over IP on the HFC and one reason to do this is to prepare to offer cable services on other edge networks which could include wireless networks or other wired networks, with the common denominator that they support IP transport of cable services. It is also expected that services will be delivered over networks which are not necessarily owned by the cable operators themselves (for example, these services would be provided "over the top" on a network owned by third-parties).

Getting back to HFC, with Modular CMTS, DOCSIS 3.0 and other specifications such as PacketCable Multimedia and Embedded DOCSIS, the toolset is available to begin laying the groundwork to deliver all cable services, including video, over the HFC edge network using IP.

NEXT GENERATION IPTV INFRASTRUCTURES: A VISION

The figure below provides a reference model for Internet TV that includes a core network and one or more last mile edge networks which has the goal of delivering cable services to a variety of end-user devices over a variety of edge networks. This network can deliver linear, switched and On Demand services. The diagram does not imply a one-to-one relationship between components.

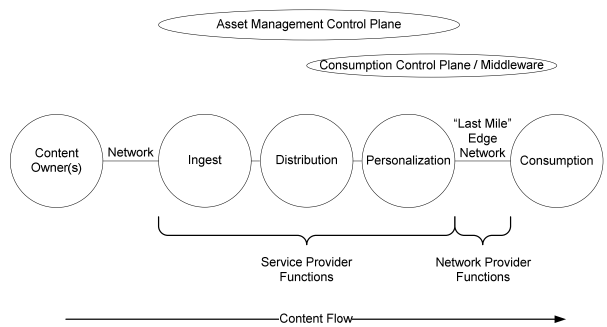

Figure 2: Internet TV reference network

The Ingest, Distribution and Personalization functions are considered part of a core network which can provide a variety of content and services to a variety of IP-capable edge networks. The system uses IP networking and transport as both the method for propagating content and as the method for staging content along the various steps of the system.

The paper first describes the HFC last mile edge network for supporting IPTV services, including M-CMTS and DOCSIS 3.0, to illustrate how these solutions work together to support high capacity, high bandwidth IP services to consumers. This section also describes how IPTV services could be delivered to existing MPEG-2 STBs.

The core network provides functions which are similar to existing cable networks. Later the paper posits how these core network functions are used to deliver cable services in an all-IP environment.

EDGE NETWORKS

As alluded to earlier, various edge networks are possible and several are shown in the diagrams that follow.

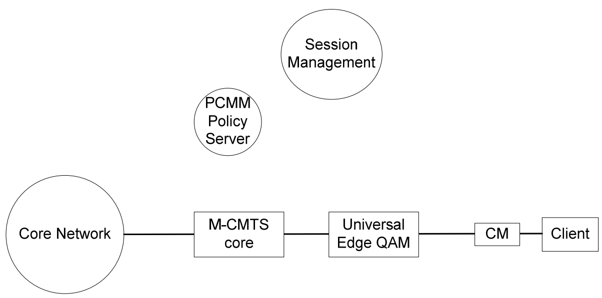

Figure 3: DOCSIS edge network

The figure above shows the DOCSIS edge network at a high level of detail, and includes the PCMM (PacketCable Multimedia) policy server and a session management function. The DOCSIS Edge network will be described in more detail in a later section.

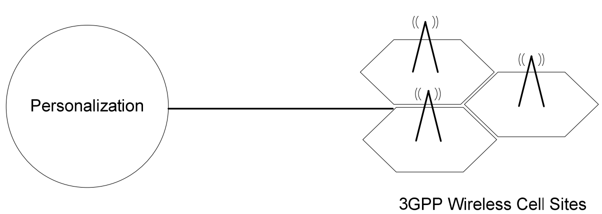

Figure 4: 3 GPP wireless edge network

The 3GPP wireless edge supports third generation mobile network services which include the ability to deliver higher data rates and advanced media features. The 3GPP wireless edge supports native IP services to customer premise equipment. 3GPP wireless technology delivers broadband-like data speeds to mobile devices while allowing consumers to send and receive services including data, voice, digital images, web pages, photographs and video three times faster than possible with a traditional GSM/GPRS network.

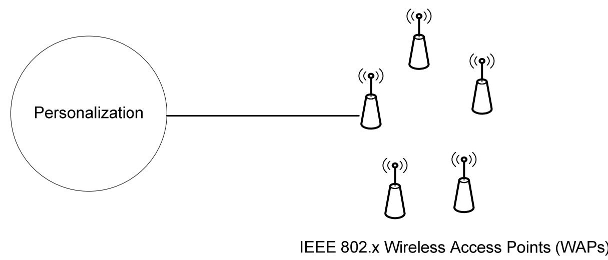

Figure 5: IEEE 802.x edge network

IEEE 802.x refers to a family of specifications developed by the Institute of Electrical and Electronics Engineers for wireless LAN technology. IEEE 802.11 (also known as Wi-Fi) specifies an over-the-air interface between a wireless client and a base station which supports data rates up to 54 Mbps per AP (Access Point) over a distance of up to 100 meters. IEEE 802.11 generally requires line of sight access and is suitable for hotspot access in areas of a city.

IEEE 802.16 (also know as WiMax) refers to a family of specifications developed by the IEEE for Broadband wireless access that provides secure, full-duplex, fixed wireless MAN (Metropolitan Area Network) service. Throughput can reach 75 Mbps and does not require line-of-sight to operate and reach can extend from one mile at full speed to thirty miles at reduced throughput.

IEEE 802.20 (also know as Mobile-Fi) is similar to WiMax but includes optimization for roaming and hand-off at speeds up to 150 miles per hour.

DOCSIS 3.0 / M-CMTS Edge Network

The advantages provided by DOCSIS 3.0 enable an operator to provide IP video services over their HFC networks. DOCSIS in general provides IP transport over HFC networks and DOCSIS 3.0 in conjunction with PCMM provides the functions needed to deliver high-speed video services over IP over an HFC network.

DOCSIS 1.1 provided the necessary QoS (Quality of Service) for IP services and DOCSIS 2.0 provided higher capacity return path channels. DOCSIS 3.0 provides channel bonding services which dramatically increase both upstream and downstream speeds and capacities as well as the IP Multicast services needed to deliver video. The DOCSIS 3.0 specifications also describe how existing DOCSIS 2.0 and DOCSIS 1.x cable modems operate on DOCSIS 3.0 upstream and downstream channels. The combination of DOCSIS 3.0 and M-CMTS provide both the technical and economic solutions to providing video services over DOCSIS.

The M-CMTS specification describes an economical method for implementing downstream channel bonding which is a key enabler for offering entertainment quality video over DOCSIS.

In the M-CMTS architecture shown below, a device referred to as the M-CMTS Core contains the DOCSIS MAC (Media Access Control) functions which include all DOCSIS signaling functions, downstream bandwidth scheduling, and DOCSIS framing. A UEQ (Universal Edge QAM) contains mainly physical layer related circuitry, such as QAM modulators, and packet tunneling logic to connect to the M-CMTS Core.

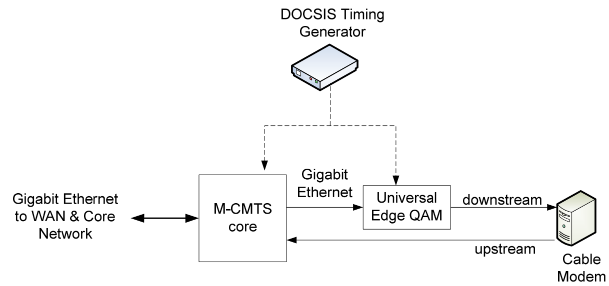

Figure 6: Conceptual view of M-CMTS architecture

Original CMTS implementations placed the CMTS core and QAM functions in an integrated assembly, fed with a common internal clock. The M-CMTS architecture creates interfaces to separate the downstream edge QAM from the M-CMTS core. It is this separation which allows massive, scaleable downstream capacity through the use of universal edge QAMs instead of downstream QAMs integrated into the CMTS as is the case today.

Since M-CMTS enables the separation of these components, the DOCSIS Timing Generator was created to supply the accurate and robust transport of a master clock signal between the M-CMTS components. The DOCSIS Timing protocol is structured to support all SCDMA and TDMA timing requirements.

An added benefit of an M-CMTS architecture is its ability to support flexible allocations of upstream and downstream channels to meet changing traffic pattern demands. That is, M-CMTS can offer more capacity in the downstream direction for video services during the evening hours, but symmetric upstream and downstream capacity during business hours when commercial data services are at peak usage.

ACHIEVING SCALEABLE DOWNSTREAM DOCSIS CAPACITY

Delivering traditional cable video services over DOCSIS requires a lot of downstream capacity. M-CMTS allows downstream capacity to be added to the system by adding scaleable modular UEQs to the M-CMTS core.

The original integrated CMTS implementations included cards that had both upstream and downstream capabilities. These implementations were suitable for the relatively low bit rate services of high speed Internet and voice. But video services changed the requirements of the system by requiring significantly more downstream capacity than before. Integrated CMTS implementations cannot provide massive economic downstream capacity because each time a card is added, the upstream capacity of those cards is left stranded. Even a downstream-only card is not optimal because of the opportunity cost associated with sparing the additional downstream-only card in the CMTS chassis to provide redundancy for services.

The M-CMTS specification was designed to address these two issues and bring a third solution to bear by supporting the use of dense, scaleable UEQs to provide the downstream capacity while offering the opportunity to share the UEQs with other services such as traditional digital video, switched video and video On Demand.

DOCSIS 3.0 describes channel bonding services which supports the concurrent scheduling of DOCSIS data over multiple channels.

As shown in the following figure for downstream channel bonding, the CMTS distributes individual IP packets over multiple QAM channels in such a way that permits the CM (cable modem) to resequence out-of-order packets. By bonding together multiple downstream channels, the aggregate downstream speeds to a single CM can be 100 Mbps, or more, and have potential to increase broadband subscribers' downstream access speeds by factors ranging from 2 to 20, to support services ranging up to HDTV (High Definition Television).

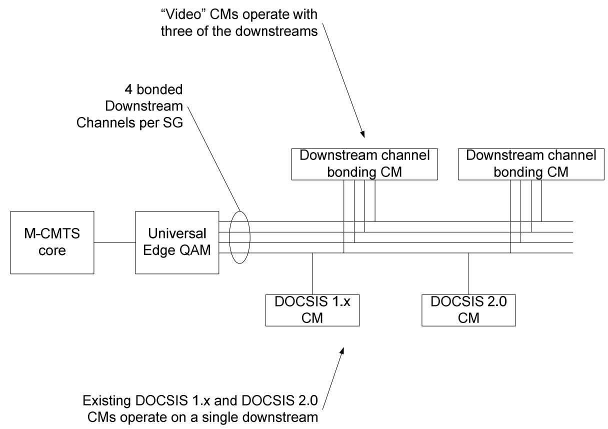

Figure 7: Downstream M-CMTS traffic flows

A key element of downstream channel bonding as described in DOCSIS 3.0 is the ability of existing DOCSIS 1.x and DOCSIS 2.0 CMs to operate transparently on channels which are bonded together for DOCSIS 3.0 operation.

Additionally, DOCSIS 3.0 allows the M-CMTS to assign individual downstream services to particular downstream channels. For example, an M-CMTS may assign a video-over-IP service flow to a downstream channel with deeper interleaving for higher reliability, while also assigning a voice service destined for the same modem to a different downstream channel with shallower interleaving for low latency.

DELIVERING IP VIDEO OVER HFC NETWORKS

An M-CMTS architecture can be combined with DOCSIS 3.0 to economically deliver cable video services over a DOCSIS network, including linear, switched digital and On Demand services.

Delivering Linear and Switched Digital IPTV Services

The DOCSIS 3.0 specifications describe a method for delivering linear IPTV services; however, it is the addition of PCMM which allows service guarantees to be offered through policy-based QoS management on the HFC network.

Policy-based QoS management ensures the bandwidth needed for the video service is allocated and managed on the HFC network. Without this management, it would be possible to put more services on the HFC than can be carried by the available bit rate, that is, video services could interfere with data or voice services, and visa versa. It is the PCMM function which allows the operator to manage the HFC resources and make intelligent business decisions in times of resource contention as to which services should be allowed on the HFC network.

There are a couple of models being investigated which use PCMM for delivering linear services and the intent is to provide the best customer experience by minimizing channel change times while providing the operator with the necessary control points to manage the service.

Both methods exhibit a key feature of DOCSIS 3.0 which is the ability to natively provide switched digital video as a by-product of the IPTV implementation using IP Multicast.

The diagram below shows the first method which follows a traditional IPTV model adapted to use PCMM where the IPTV client sends an IGMP Join to the network to “change channels” to a broadcast service.

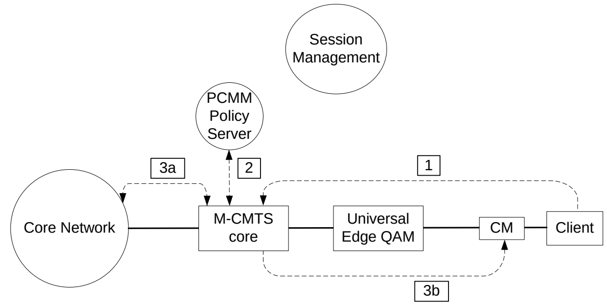

Figure 8: Synergy between PCMM and IPTV

To change to a linear channel, the IPTV client sends an IGMP Join (1) which is received by the M-CMTS core. If that channel is not already on a downstream accessible by that cable modem, but is already flowing into the M-CMTS core, upon receipt of the IGMP Join the M-CMTS will initiate a query (2) to the PCMM policy server to allocate capacity on the appropriate downstream for that service, then configure the CM for the IP multicast flow carrying that service (3b) and then forward the service to that cable modem.

If that channel is already on a downstream accessible by that cable modem, the M-CMTS core will configure the CM for the IP multicast flow which is already carrying that service (3b) on the appropriate downstream and the client will be able to view the content. This is an example of the inherent video switching capability of DOCSIS 3.0. Once a linear service is on a downstream channel, any client which can access that downstream channel can access that service; hence, only one copy of the service is needed on that downstream even if there are multiple viewers of that service. To complete the transaction, the M-CMTS core will signal the PCMM policy server of the channel change request.

If that channel is not already on a downstream accessible by that cable modem, and is also not yet flowing into the M-CMTS core, upon receipt of the IGMP Join the M-CMTS will first complete the query (2) to the

PCMM policy server to allocate capacity on the downstream for that service and then issue an IGMP Join to the core network (3a) to get the service to flow to the M-CMTS core, then configure the CM for the IP multicast flow carrying that service (3b) and then forward the service to that cable modem.

Notice that there is no explicit interaction with the session management function. Since all channel change requests are signaled to the PCMM policy server, session information can be queried from there. This is in contrast to the second way of implementing linear IPTV services with PCMM which explicitly uses the session management function.

The diagram below shows a linear IPTV model with a session-based controller adapted to use PCMM. Interaction with session management is very similar to switched broadcast. In this model, when a channel change occurs the IPTV client issues a session request to the session management function.

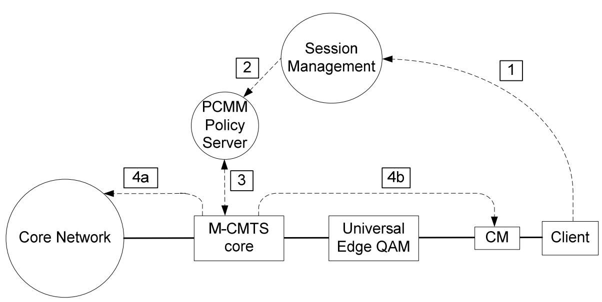

Figure 9: Linear TV model with session-based control

Upon receipt of the channel change request, the session management function interacts with the PCMM policy server to request bandwidth on the M-CMTS core for the service. If that channel is already on a downstream accessible by that cable modem, the M-CMTS core will configure the CM for the IP multicast flow which is already carrying that service (4b) on the appropriate downstream and the client will be able to view the content. This also is an example of the inherent video switching capability of DOCSIS 3.0. Once a linear service is on a downstream channel, any client which can access that downstream channel can access that service; hence, only one copy of the service is needed on that downstream.

If that channel is not already on a downstream accessible by that cable modem, but is already flowing into the M-CMTS core, the M-CMTS will configure the CM for the IP multicast flow carrying that service (4b) and then forward the service to the downstreams which services that cable modem.

If that channel is not already flowing into the M-CMTS core, the M-CMTS will first issue an IGMP Join to the core network (4a) to get the service to flow to the M-CMTS core, then configure the CM for the IP multicast flow carrying that service (4b) and then forward the service to that cable modem.

Note that in this case there is always an interaction with a session server, as well as interaction with the PCMM policy server.

Delivering On Demand IPTV Services

The DOCSIS 3.0 specifications support the delivery of On Demand IPTV services as unicast streams. As with linear services, the addition of PCMM allows service guarantees to be offered through policy QoS management on the HFC network.

The model for On Demand IPTV services is very similar to existing On Demand services, except rather than send the stream through a video QAM it is sent through a DOCSIS QAM.

The diagram below shows an On Demand IPTV model with a session-based controller adapted to use PCMM. In this model, the IPTV client issues a session request to the session manager when requesting On Demand content.

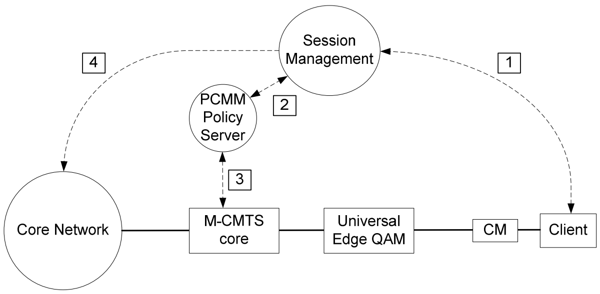

Figure 10: On Demand IPTV with session-based control

The session management then instantiates QoS with the M-CMTS core through the PCMM policy server (2, 3). If the bandwidth needed to deliver that service is available, the session manager will instruct the core network to begin streaming the service to the client.

Considerations for a Cable IPTV STB

An IPTV STB needs a connection to a DOCSIS cable modem. It also needs to be able to decode video from the CM to the decoder.

The following figure shows an example of an eDOCSIS (Embedded DOCSIS) STB solution where the DOCSIS CM follows the eDOCSIS specifications and the Client is considered an Embedded DOCSIS Embedded eSAFE (Service Application Functional Entity).

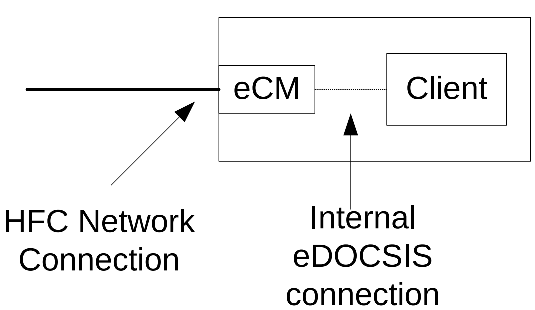

Figure 11: Embedded DOCSIS STB solution

The following figure shows a solution utilizing a stand-alone DOCSIS CM connected to the Client over a home network.

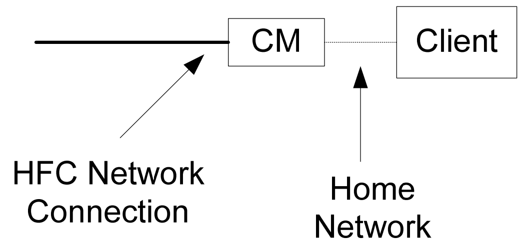

Figure 12: Stand-alone cable modem and client

In this case, the home network supports video streaming in order not to have a negative effect on the service.

With all DOCSIS CM solutions, an M-CMTS solution with downstream channel bonding would provide the most scaleable and economic solution. Downstream channel bonding provides additional capacity which will benefit video services. Additionally, DOCSIS/DSG is used for the out-of-band communication path thereby obviating need for legacy return path components in the STB.

The following figure shows an Off-Net CPE solution where the last mile access IP network is provided by a third party. In this case, a third party manages the network and the relationship with the NID (Network Interface Device). The cable operator is only a provider of content services to the client.

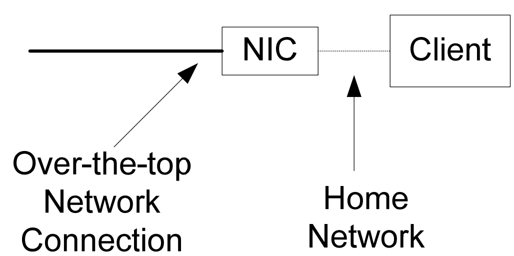

Figure 13: Off-Net CPE in over-the-top application

The client will need the ability to decode video in the format that is sourced from the cable operator. IPTV content can be encoded and delivered in several formats. The traditional format for cable digital video is MPEG-2 transport streams over UDP/IP, specifically packing seven 188-Byte MPEG-2 cells in one IP packet. This is the same format used to deliver On Demand or digital simulcast content over Gigabit Ethernet to a QAM. This same format can be used by an M-CMTS to deliver digital content to a DOCSIS QAM for consumption by the client.

An alternative method to deliver IPTV may be supported by "new media" sourced from Internet content providers. In this case, the content is coded as elementary streams over RTP/UDP/IP with no MPEG-2 framing. In this case, the timing for the content is recovered from the RTP layer instead of from the MPEG-2 transport stream layer. This method is how content is delivered to clients such as Windows Media and Real Player and there are standards available for delivering content using MPEG-2, H.264 and AC-3 codecs as well; hence, this may be an emerging trend for IPTV content.

Delivering IPTV Services to Existing STBs

The IPTV architecture could be considered incomplete if it did not address delivering IPTV services to legacy STBs.

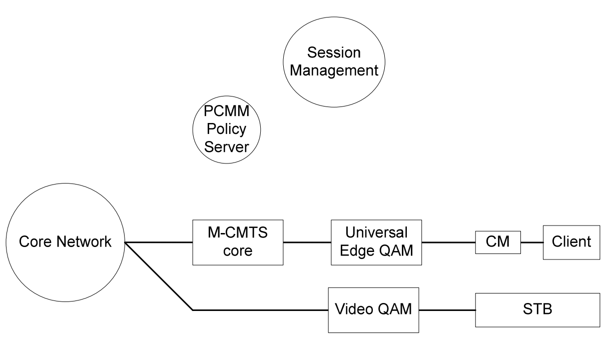

Figure 14: Video delivery to Edge QAMs

The above figure indicates how content available in the core network can be distributed to an existing video QAM just as easily as a DOCSIS QAM through an M-CMTS.

The session management function will know the capabilities of the entity requesting the service, whether it is an IP client or an existing STB application. When the entity places the request, the core network will stream the content over IP to the HFC network. If the content is destined to an IP client then the system will route the content over the M-CMTS QAM. If the content is destined for an existing STB, the system will route the content over IP to a video QAM just as happens today for digital simulcast and On Demand. The video QAM will take the content over IP as input and then strip off the IP and deliver the stream as native MPEG just as it does today to digital STBs.

Existing digital STBs can only render content which is encoded as MPEG-2 transport streams. Some IPTV content can be encoded as elementary streams over IP, with no MPEG-2 framing. For delivery of codec elementary streams to an existing STB, there will need to be a transcoding step to change the codec elementary streams over IP to an MPEG-2 transport stream over IP. This function is currently defined to happen as part of the personalization function.

With this transcoding function, it will be possible to deliver IPTV content, regardless of the encoding and transport, to both IPTV devices and to existing STBs.

CORE NETWORK

The functions of the core network show many similarities to existing cable networks, including linear and On Demand services. The differences are in the areas of making the existing and emerging services available to a wider array of end devices. In all cases it is assumed the transport of content to these end devices uses IP, but these devices will have characteristics which vary from those of the traditional STB.

The core network can be characterized by three functions including Ingest, Distribution and Personalization.

Ingest Functions

The Ingest functions of the core network are shown in the following figure.

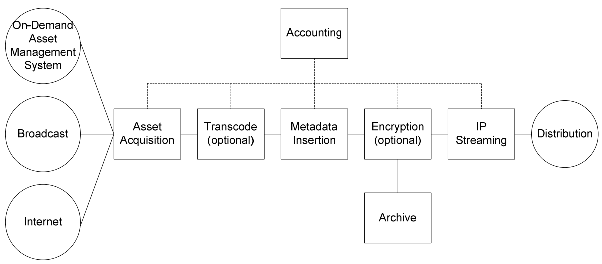

Figure 15: Ingest functionality in network model

The Ingest functions are designed to acquire content from a variety of sources including linear broadcast networks, On Demand storage and content from sources that might otherwise be associated with the Internet.

A key functions of the Ingest network is asset acquisition which includes the ability to connect and retrieve real-time and non-real-time content feeds including the reception and distribution of content from multiple independent sources to multiple distribution servers and applications.

The asset acquisition system streamlines and automates the process of receiving, staging, storing and propagating content including creating index files (for example trick mode files, for On Demand content.)

Another important function of the Ingest network is transcoding. The transcoding system processes MPEG-2 content and generates content in the proper format/profile and with the proper metadata for the media devices supported in the network.

Minimally the transcoding system must support the transcoding of MPEG-2 video and audio to H.264 and MP3 or AAC standards, respectively. In addition, the transcoding system may support conversion from H.264 to MPEG-2 for providing internet-sourced content to existing STBs. Both SDTV and HDTV formats should be supported.

For use with mobile networks, the transcoding system should support the media formats that are described in the PacketCable Release 2.0 Codec and Media specification, available through CableLabs. This specification requires the H.263 codec as well as recommends various profiles of H.264 as optional.

Depending on operator plans, additional codecs could be considered including the transcoding of MPEG-2 video into the following formats:

- Flash
- Windows Media
- Real
- Quicktime

Though these formats are available with CDN systems, there are many system implication to supporting these formats which need to be investigated.

Video content carried over IP can take one of two basic forms. The first encapsulates an MPEG-2 transport stream over IP as is used today to transport video streams over backbone and regional area networks. The second encapsulates codec elementary streams directly into RTP over IP with no MPEG-2 framing. These two methods are fundamentally different and are not necessarily interchangeable at end devices such as STBs and mobile video players. The core network should include the capability to transcode from one to the other while either maintaining the same codecs or when switching between codecs.

The next function includes the insertion of metadata which should comply with the CableLabs metadata specification, which may need an update to support the new types of content which can be ingested including obtaining content and metadata from various internal and external (for example the Internet) sources. This function should support both automatic and manual insertion and include the capability to translate from proprietary metadata formats to the format specified in the CableLabs metadata specification.

Finally, an encryption function should be present to protect content on the core network and an archive function could be present to maintain copies of content.

Distribution Functions

The distribution functions are shown in the following figure.

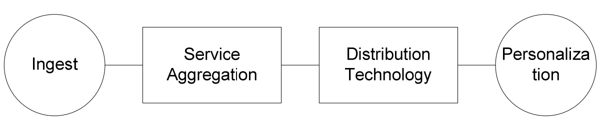

Figure 16: Service aggregation and distribution

Service Aggregation defines the set of service packages and bundles that are ultimately sold to the subscriber. Packages can be composed of real-time transcoded programs, multicast time shifted programs, unicast On Demand programs or a combination of all of the above.

The Service Aggregation system understands the current promotional offers and provides intelligent up-sell opportunities by interfacing with the billing system. The Service Aggregation system makes it easy to create new services, push personalized offers, and recommend add-on services and up-sell.

The Service Aggregation function interacts with the Ingest system and obtains information necessary for the On Demand application from other components, such as the asset metadata from the Asset Management System. This function of Service Aggregation interacts with the Host to present navigation menus and related application features to the On Demand Client and exchanges messages with the On Demand Client to enable the navigation functions, including parental controls.

The distribution technology will be chosen to provide a dynamic system for managing content in the network.

The most common method by which files are transferred on an IP network is the client-server model where a central server sends the entire file to each client that requests it. In this model the clients only speak to the server, and not to each other. The main advantages of this method are that it's simple to set up, and the files are usually always available since the servers tend to be dedicated to the task of serving, and are always on and connected to the Internet. However, this model has a significant problem with files that are large or very popular, or both. Namely, it takes a great deal of bandwidth and server resources to distribute such a file since the server must transmit the entire file to each client. The concept of mirrors partially addresses this shortcoming by distributing the load across multiple servers but it requires a lot of coordination and effort to set up an efficient network of mirrors and it's usually only feasible for the busiest of sites.

On the newer end of the spectrum are Torrent-based technologies which work by taking a large file and breaking it into pieces which are spread out over a network of servers. When the file is downloaded, it travels over the network in pieces from multiple sources which takes less time than downloading the file from a single source. The issue of scale with popular downloads is somewhat mitigated because there's a greater chance that a popular file will be offered by a number of servers. Overall, distributing the download in this fashion reduces the aggregate load on the servers and makes for a more efficient network.

Personalization Functions

The personalization functions are shown in the following figure.

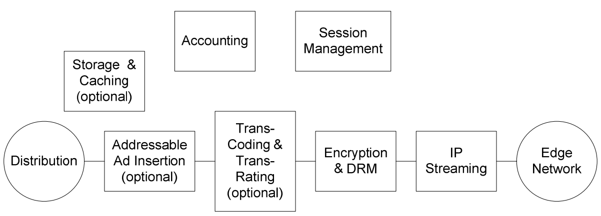

Figure 17: Personalization functions

Session Management

The Session Manager should support the ability for components to scale to large numbers of active sessions and many session setups / teardowns per second. The Session Manager implements the service and business logic, including policy decisions, associated with Internet TV. The number of session may grow very large and the architecture needs to support optimized resource allocation when many potential resources for a particular session exist.

The architecture should be designed to result in minimal session setup latency by minimizing the number of messages used to set up sessions and allowing as much parallelism as possible between the components involved in setting up the session.

The architecture allows service and business logic to be incorporated into the resource selection process. For example, it may be more important to fulfill requests for sessions associated with a pay per view movie service than for sessions associated with a free On Demand Service. The architecture should be able to translate service and business logic into a set of rules to be incorporated in the resource selection process.

The architecture should make it possible to incorporate new classes of resources along with new resource managers without having to make major modifications to other components of the system. The architecture should allow the ability to modify or add resources to an existing session.

The architecture and resulting protocols should enable the components to be implemented using simple state machines. The architecture should also minimize the amount of state that needs to be maintained across components.

Addressable Ad Insertion

Addressable advertising provides the capability of addressing specific ads to specific customers. An AA (addressable advertising) architecture is designed to make near real-time decisions on addressability that can differentiate between individual STBs and different users watching that STB. For example, if a home has just one STB the system should be able to discern if it is a parent or a child using that STB and address the appropriate ads. For example, the system could use near real-time analysis of remote control key presses to determine if the user is watching sports, dramas or cartoons and make the appropriate near real-time decisions on which ads to insert. The system can also make use of outside information such as billing profile, demographic and economic data.

Addressable advertising is available for all services provided over the network.

Instant Channel Change

Instant Channel Change enables channel changing times that are less than traditional cable network tuning of digital channels. ICC (Instant Channel Change) effectively eliminates the delay associated with tuning channels in a digital cable system by exploiting the large video buffers in the STB available with an IPTV system.

Trans-Coding & Trans-Rating

The trans-coding function is a final opportunity to modify the codecs used to deliver the service.

The trans-rating function is available to “down rate” the service from the bit rate used for transport on the network to a bit rate that is suitable for delivery to the client. For example, a 1.5 Mbps H.264 encoded service could be trans-rated to 300 kbps for delivery over an edge network which cannot support higher data rates.

Encryption

The encryption function is used to make the service unreadable by users without knowledge of the keys to decrypt the service.

The encryption may be specific to the client and possibilities include Windows Media, Real Helix or any other scheme used by the cable operator.

Digital Rights Management

Digital rights management is a comprehensive and flexible platform that enables rights holders to create business models for delivered content through the secure distribution of rights to trusted players.

The DRM (Digital Right Management) technology consists of two components including metadata associated with a piece of content which describes the business rules associated with that content. These rules generally indicate how the content can be consumed, by whom, and for how long. Services are built using the metadata. Additionally DRM includes cryptographic algorithms used to protect the metadata and the content from unauthorized use.

The DRM secures both the content and the business rules associated with that content. The devices enforcing the DRM must be trusted, e.g., STBs, PCs, portable devices, etc. and must be manufactured in such a way that the DRM is not easily defeated.

The DRM metadata is used to describe business models and should be flexible enough to allow an operator to implement a wide array of services. For instance, DRM metadata could require a piece of content to be deleted after 5 plays could disallow trick-mode playback during a certain window of time (for example a theatrical release).

DRM could be designed as an extension of the operators existing conditional access systems, and this should include multiroom capabilities, i.e., sharing content over a home network and possibly even over large geographic distances.

IP Streaming

The IP Streaming function will support outputting content in either IP multicast or IP unicast streams. The IP Streaming function must support IPv4 and must provide an upgrade path to IPv6.

Quality of Service

The QoS function must include support for DSCP (Differentiated Services Code Point) as well as possibly other IP-based QoS mechanisms as specified by the operator.

Bandwidth Management

The Bandwidth Management function will monitor and police services and bandwidth usage.

Storage and Caching

The Storage and Caching functions are used to keep select assets close to the edge. Examples can be assets which are stored as an extension to the Distribution System including movies or advertising content.

Storage at the network edge is also used for remote storage DVR (Digital Video Recording) in conjunction with the Consumption control plane.

CONCLUSIONS

The paper presents a Video over DOCSIS solution using several available industry specifications including:

- Modular CMTS
- DOCSIS 3.0
- PacketCable Multimedia
- Embedded DOCSIS

Modular CMTS in conjunction with DOCSIS 3.0 provide the most economical means to deliver linear, switched and On Demand digital programming over an HFC IP network. The proposed solution maintains backward compatibility with both existing DOCSIS 1.1 and DOCSIS 2.0 infrastructure and describes a method for delivering IPTV content to existing STBs.

The solution is applicable to not only DOCSIS networks, but also other edge networks including 3GPP, Wi-Fi, and WiMAX wireless networks as well as going over the top on networks owned by third party service providers.

Functions for a core network are proposed that are not that different from cable networks today. Content ingest remains essentially the same with the added ability to transcode content to different codecs. Content distribution capabilities become more sophisticated to replicate traffic around the network where it is needed. Lastly, a group of personalization functions are described which allow cable services to be delivered to customers in a way which matches their method used to attach to the network, whether that be with a STB, mobile device, or some other customer premise equipment.