CONFIGURATION RECOMMENDATIONS FOR DOCSIS TRANSPORT OF MANAGED IP VIDEO SERVICE

WILLIAM T. HANKS

With special thanks to the following contributors:

Jim Allen, Tom Cloonan, Michael Dehm, Amit Eshet, J.R. Flesch, Dwain Frieh, Will Jennings, Tal Laufer, Tushar Mathur, Dan Torbet, John Ulm, Bill Ward, and Ian Wheelock

TABLE OF CONTENTS

INTRODUCTION 4
TARGET AUDIENCE 4
VENDOR AND PRODUCT AGNOSTICISM 4
TECHNOLOGY BACKGROUND 5
Video Services 5
Linear Broadcast TV vs. On-demand Content 5
Television Channel Viewership 5
TV Channel Change Expectations 8
Internet Protocol (IP) 8
Connection-Oriented vs. Connectionless 8
IP Multicast – One Stream, Multiple Destinations 9
IP Unicast – One Stream, Single Destination 9
Adaptive Bitrate Video 9
IP Unicast ABR Video 9
Multicast-assisted ABR Video 10
Home Equipment Clarification 11
Cable Modem 12
Home Gateway 12
IP Video Player 12
Smart TV 13
Console 13
Tablet & PCs 13
IP Video Client w/multicast 13
Data over Cable System Interface Specification 14
Cable Modem Capabilities 14
Channel and Flow Assignment 15
Service Flow Types 17
MAC Management Dynamic Services Transactions 18
Quality of Service 19

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Primary-Capable Downstream Channels 21
Upstream Supervision 22
Energy Management 22
IP Multicast 23
QoS Control Plane Signaling 23
NACK-Oriented Reliable Multicast (NORM) 26
RECOMMENDATIONS - TYING IT ALL TOGETHER 26
CM Capabilities 27
Channel Assignment to CMs 27
Consolidate Multicast Program Streams 27
Load-Balance Remainder of Channels 30
CM Unicast IP Video Service Flow QoS Parameters 30
Downstream 31
Upstream 32
QoS Settings across Different Services 32
CONCLUSION 34
ABBREVIATIONS & DEFINITIONS 35
RELATED READINGS 38
REFERENCES 39
APPENDIX: SAMPLE IP VIDEO CONFIGURATIONS 40
Minimum 32 DS Channel CMs/HGs for IP Video 41
Minimum 24 DS Channel CMs/HGs for IP Video 42
Minimum 16 DS Channel CMs/HGs for IP Video 42
Minimum 8 DS Channel CMs/HGs for IP Video 43

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

INTRODUCTION

The internet protocol (IP) has been used to transport television (TV) programing in some form for over a decade. The telco companies have been using IP to commercially distribute television over Digital Subscriber Line (DSL) technology for at least the past five years. In order to do so, the network and subscriber technologies have been refined to provide a very acceptable video service over an IP transport network – generally without regard to the intricacies of the access network in the last mile to the subscriber.

Primarily due to their significant investment in Motion Pictures Expert Group (MPEG)-based video technology, cable television companies have been much slower to adopt an IP video architecture. However, it seems like nearly all of these Multiple System Operators (MSOs) have plans to "eventually" replace MPEG transport technology with IP transport technology to provide a managed video service for their subscriber base. The predominant way to transport IP over a cable television network is to use equipment that supports the Data Over Cable System Interface Specification (DOCSIS). Unlike DSL, DOCSIS uses a shared medium but provides a rich set of tools to ensure Quality of Service (QoS) to each subscriber. This paper will address some of the issues in using DOCSIS to transport IP video.

TARGET AUDIENCE

This document is intended to be read by people with a wide variety of backgrounds and experience. While some of the audience may have a broad understanding of the technologies, others may have deep experience with certain parts but have little experience with others. Still others may be new to most of these topics.

The first several sections will try to provide everyone with the requisite information to understand the recommendations and the reasoning for them. If one or more of the background sections are already understood by the reader, please feel free to skip to the next section.

VENDOR AND PRODUCT AGNOSTICISM

The ideas and recommendations expressed in this work are believed by the authors to be applicable on most products from most industry vendors. Some background assertions may prove to be untrue on some products and some recommendations may be unrealizable on some products. If issues arise when attempting to tune an IP video system, please refer to product information and contacts from the component equipment vendors.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

TECHNOLOGY BACKGROUND

Video Services

Linear Broadcast TV vs. On-demand Content

Since the very first television network broadcast over the air, television has followed the linear TV model – that is that the program lineup has been pre-scheduled so that an audience must tune into a particular television channel at a particular time to view a particular program. This is also known as appointment TV. From a transmission technology standpoint, a large number of viewers were able to view the same program at the same time because the program was (and continues to be) broadcast via radio frequency (RF) carriers to all users in the viewing area at once. Once these transmissions were distributed via a cable service, the cable TV (once known as community antenna television) companies followed this model by broadcasting the received signals to all subscribers.

Over time, some cable content companies came up with the idea of on-demand television. In this model, a single subscriber requests a particular program to be sent. Based on the idea that each viewer wants to be in control of the timing of their viewing experience (not to mention being able to control the “trick modes” of pause, rewind, skip, fast-forward, etc.), each on-demand session is presented as if it were unicast; regardless of the transmission technology used to deliver it. Thus, two viewers in close proximity who are watching the same on-demand program are each likely to be viewing a different transmission copy of the content.

Television Channel Viewership

The following graphs (Figures 1 and 2) show a snapshot of the actual observed TV viewership for a U.S. major metropolitan area during December 2015. Figure 1 shows that about 95% of the viewers were watching a total of only ten TV channels. When an additional ten TV channels are added to the first ten, only about 2.5% more (see Figure 2) of the viewers can be served. Ten more channels captures only an additional 1% of the viewers.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Figure 1: TV viewership vs. number of channels viewed

Figure 1 points to an opportunity to optimize a small number (ten) of TV channels to be carried in a "broadcast fashion" such that all viewers have access to these TV channels at all times. Furthermore, whenever a viewer tunes to a new channel, there is about a 95% chance that they will tune to one of these ten channels. In historic video circles, this small group of channels would be referred to as the "short tail" as they are a small number of channels with the bulk of the viewers. Far to the right of the graph, as more channels are included, a very small percentage of subscribers can be serviced by adding many additional TV channels – thus the notion of a "long tail". Somewhere between the short tail and the long tail would be considered the "medium tail". The graph shown here does not have much of a medium tail.

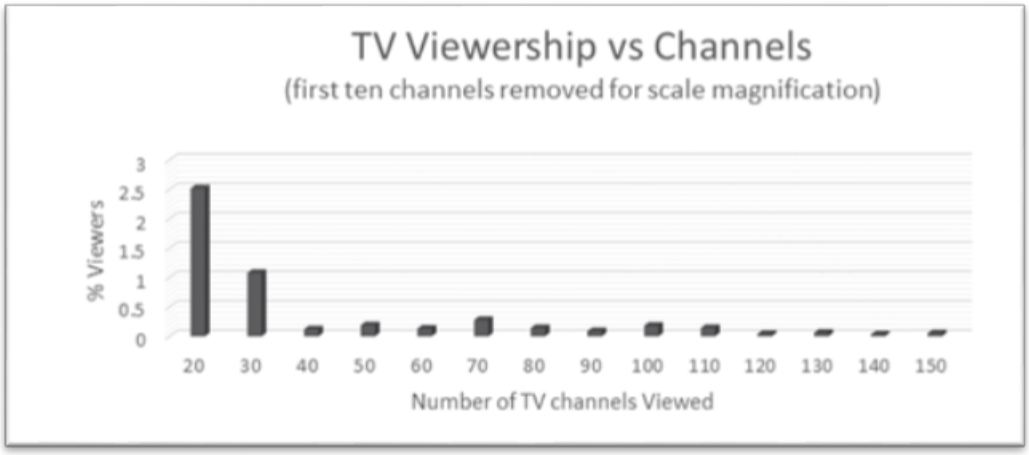

Figure 2: TV viewership vs. number of channels viewed (first ten removed for scale magnification)

The graphs here show the viewing patterns of a particular audience with a particular demographic. While the pattern tends to be the same across groups of subscribers, a more diverse audience may produce a more "flattened" viewership graph with a smaller short-tail and larger medium-tail and long-tail components.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Short Tail: Linear, Multiple Viewers (Broadcast)

Historically, all cable television video are served by a set of MPEG streams that were carried on a single quadrature amplitude modulation (QAM) channel per stream. This set of MPEG streams is broadcast to every television viewer at all times. The viewer's set-top box (STB) tunes to a different preexisting QAM channel frequency when the viewer changes the TV program on the set-top box. Every QAM channel is always available and there is no need for program change signaling to the network. This sort of broadcast approach (albeit modified for IP multicast) might still be useful to carry the short tail content – especially true when the short tail content might be carried on as few as ten TV program streams.

Medium Tail: Linear, Multiple Viewers (SDV)

Linear MPEG video service took a step forward in conserving bandwidth with the innovative approach known as Switched Digital Video (SDV). This technology might be described as switched broadcast linear television in that a program stream is assigned to a QAM channel within a video service group (SG) only when a STB requests the program stream from an SDV controller. Once the program stream is assigned to a QAM channel, the STB retunes a receiver to the QAM channel and delivers the video. In this way the linear video stream is delivered on-demand but subsequent requesters for the same program stream from the same video SG can listen to the same exact broadcast MPEG stream by tuning to the same QAM channel. This sharing of the linear but on-demand stream greatly reduces the amount of bandwidth capacity used to deliver the program stream to multiple viewers over the bandwidth that would be needed to transmit to each requester individually. This economy comes at the price of subscriber control (no start, stop, or trick modes) of the on-demand stream. Of equal importance is the fact that this approach to program delivery also "turns off" the program stream to the video SG whenever the system has detected that the last viewer of the program stream has tuned to another program- as a result, bandwidth on the cable plant is not wasted by delivering un-viewed programs.

Long Tail: On-demand, Single Viewer (VOD)

On-demand video has long been associated with home viewing of single program content (i.e. movies). However, the same technology can be used to serve individual subscribers who wish to view special interest or significantly time-shifted (longer than normal jitter buffers) linear content. In the MPEG video world, a requested program stream might be placed onto an amount of previously unallocated QAM channel bandwidth for use by a single subscriber. Such a technology might allow the use of trick modes on the stream to provide the interface as one might find on a home video player.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

TV Channel Change Expectations

Video subscribers expect their video services to be responsive to their inputs. While a number of articles indicate that acceptable channel change times are on the order of a couple of seconds, one study of viewer behavior concluded that in order to guarantee an acceptable quality of service for channel zapping, the zapping time (time from stimulus to change a channel until video and audio for new channel appear) needs to be less than half a second (430 milliseconds [1]).

Internet Protocol (IP)

The internet protocol (IP) was created to get packets from one place to another, possibly through multiple devices or "hops". IP does not guarantee that packets will not be lost. IP does not guarantee that packets will arrive in order. IP does not guarantee that an application will only receive the packets that are useful to that application. All of these extra features are left to the next higher layers – either transmission control protocol (TCP) or user datagram protocol (UDP).

Connection-Oriented vs. Connectionless

Both TCP and UDP will properly route packets to the proper destination application (and filter the packets from other applications). However, other features depend on whether there is a controlled virtual connection between the sender and receiver or whether there is an uncontrolled mass of packets being transferred.

Transmission Control Protocol (TCP)

TCP is a connection-oriented protocol. It has the ability to ensure reliable, in-order transport from a sending application to a destination application. These features are accommodated by the use of sequence numbers, resequencing buffers, acknowledgements, and retransmissions. TCP connections typically present a stream of transferred payload bytes to the receiver without any indication of where the actual transport packet boundaries may have occurred.

User Datagram Protocol (UDP)

UDP is a connectionless protocol. It can properly route transport packets to the proper application but it does not share the reliability or packet ordering features of TCP. It also tends to present transport packets as packets – header and all - to the receiving application.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

IP Multicast – One Stream, Multiple Destinations

While internet protocol version 4 (IPv4) does support the concept of a broadcast destination address, and this broadcast address is used by BOOTP and DHCP, broadcast destination address use by other protocols is typically not found. Instead, multicast group addressing tends to be used in place of broadcast addressing. In fact, when internet protocol version 6 (IPv6) was defined, the concept of a broadcast destination IP address was replaced with multicast destination addressing to the all-hosts multicast group. While the distinction is subtle, there is the advantage that bandwidth does not need to be consumed for group traffic unless a group member is present (similar functionality to SDV). If a TV program stream is being transmitted to multiple viewers using multicast addressing, then there is an opportunity to have all viewers receive the same copy of the TV program stream (also similar functionality to SDV). While this mode can lead to some complicated retransmission request scenarios, it is overall the most efficient IP-based transmission mechanism in terms of bandwidth used. Multicast is best used for short-tail programming where many viewers are receiving the same program stream. It may also be used for medium-tail programming in an effort to share media streams across multiple viewers. Multicast transmission usually uses UDP transport.

IP Unicast – One Stream, Single Destination

The destination addressing mode used most often in IP (both IPv4 and IPv6) is the unicast or single-host address. A single host would be assigned this address and any IP packets with this host's destination IP address should be ignored by all hosts except the one that has been assigned the unicast address (except for a few special cases in which packets are being snooped). Thus, unicast traffic cannot be shared across destinations. If a TV program stream is being transmitted to multiple viewers using unicast addressing, then each viewer gets its own copy of the TV program stream. Due to the one-to-one service, this form of IP communications is easy to control and the use of a TCP (see section: Transmission Control Protocol (TCP)) can guarantee transmission of the media stream, albeit at the expense of replicas of single-user streams. Unicast destination IP addressing (a.k.a. IP unicast) is best used for long-tail programming or even medium tail programming if there are not many viewers. Unicast usually uses TCP transport, but UDP is also sometimes used.

Adaptive Bitrate Video

IP Unicast ABR Video

Adaptive bitrate (ABR) technology requires the encoding of a single program into multiple streams of decreasing resolution (and therefore average bitrate). Each stream is broken into chunks of constant time that align to the same quanta. All of the chunk

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

files corresponding to the same instant in time are listed in a manifest file. All of these chunks of each of the resolutions and their manifest file are stored in files on the video server network. When an ABR client starts playing a program, it requests a manifest file and usually starts by requesting the file (usually an HTTP request) with the lowest resolution chunk (although some systems start with the highest resolution chunk). It looks at the next manifest file and immediately requests the file for the second chunk. If the 2nd chunk file arrives long before it is needed, then the player may decide to request a slightly higher resolution for the 3rd chunk file. The player will play progressively higher resolutions per chunk until it either reached the maximum resolution or the next chunk file takes longer than expected; in which case the player may step down in resolution for the next chunk.

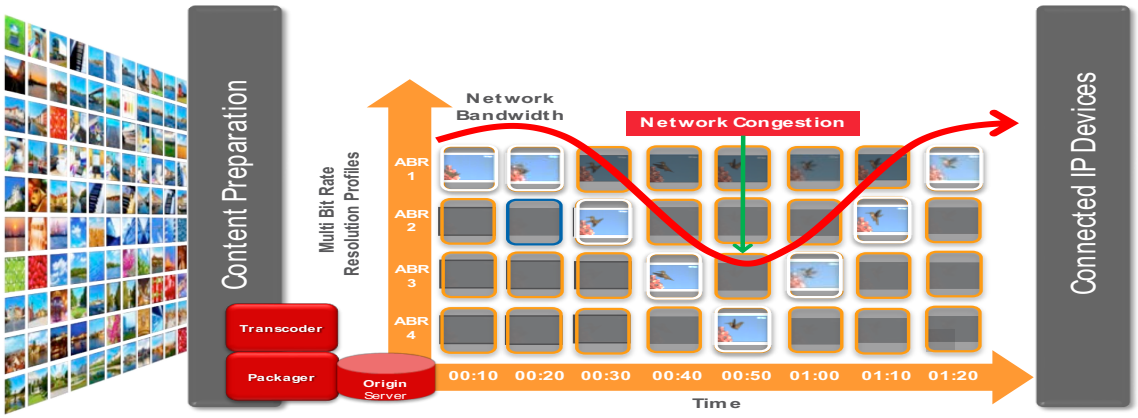

Figure 3: Unicast Adaptive Bitrate Transmission

Multicast-assisted ABR Video

Multicast-assisted ABR (M-ABR) was created to leverage IP unicast ABR advantages in conjunction with the bandwidth-savings of IP multicast. Standard ABR video players make HTTP requests to their home video IP gateway for ABR video content. The home video IP gateway is a device which is placed between the DOCSIS network and the home IP network. The home IP gateway may have a local data cache (flash disk or similar) whereby it may store content that it receives over a multicast distribution service. This content may be provided via unicast HTTP protocols to standard ABR players. Meanwhile, the home video IP gateway joins any available multicast groups for the requested video content so that the content can be prepositioned into the local data cache in anticipation of a request for the next video chunks. Ideally, the multicast prepositioning of the video chunks in the home video IP gateway local data cache can be performed at a reasonably high resolution with good QoS so as to provide a good display for the ABR clients in the home.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Home Video IP Gateway Cache Miss

If the home video IP gateway does not have the requested ABR chunk file, then it behaves very much like the unicast ABR case and forwards the HTTP ABR request towards the caches and servers in the network who then provide the chunk file back to the home video IP gateway. The home video IP gateway then forwards the chunk file to the ABR video client.

Home Video IP Gateway Multicast JOIN and Cache Hit

When a home video IP gateway notices that one of the program streams that it is serving becomes available via a multicast group, then it makes a request to JOIN the correct multicast group at its group address.

Once the multicast group packets begin arriving at the home video IP gateway, it can stop forwarding unicast ABR requests into the network and simply serve the requested chunks out of its own cache memory.

Multicast Controller

The multicast controller (running in a server north of the CCAP) decides which program streams are made available to the home video IP gateways via multicast and upon which group address. It publishes this list to the home video IP gateways.

Home Equipment Clarification

There is sometimes an ambiguity about which device in a home is being used for managed IP video. The primary use case would use home equipment as shown in Figure 4 but there are some managed video systems that deviate from this reference architecture.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

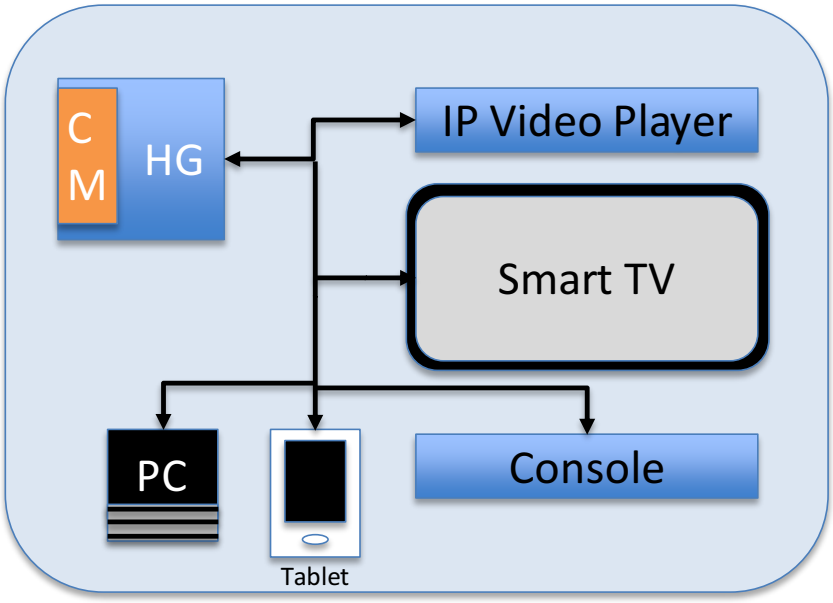

Figure 4: Home network with home gateway

Cable Modem

The cable modem provides a bridging interface to the DOCSIS network. Cable modems may be either embedded (as in the case of the Home Gateway) or they may be stand-alone. Either way, the CM behaves the same way for the purposes of its DOCSIS configuration. All of the forthcoming recommendations may be applied to embedded as well as stand-alone CMs.

Home Gateway

The Home Gateway (HG) is a whole-home appliance which provides multiple services over DOCSIS via the embedded multi-channel DOCSIS cable modem. These services may include but are not limited to managed IP video, telephony, high-speed data (HSD), home alarm and Wi-Fi hotspot. The HG is usually placed in an inconspicuous location as it usually does not include a video player (i.e. it is considered to be “headless”). If the HG does contain a video player then it would normally be placed near the main television of the home. The HG should contain some amount of buffering space for short-term video storage.

IP Video Player

The video player is an IP device that typically uses HTTP to request video chunks from a server (in this case, the HG) to render the video for a display device (like a television set). The video player may have buffering space that is used to smooth out jitter in the home network.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Smart TV

A smart TV is one which is capable of receiving over-the-top (OTT) IP video and other services via the internet (HSD service). It usually does not perform the role of video player by requesting chunks from the home gateway.

Console

Many modern gaming consoles are capable of serving as clients to OTT video services. If used this way they operate in a way similar to smart TVs.

Tablet & PCs

Tablets, mobile phones, and PCs contain video player applications. These applications may work with the managed IP video service (in which case they behave like an IP video player to request video chunks from the HG) or they may work with OTT video.

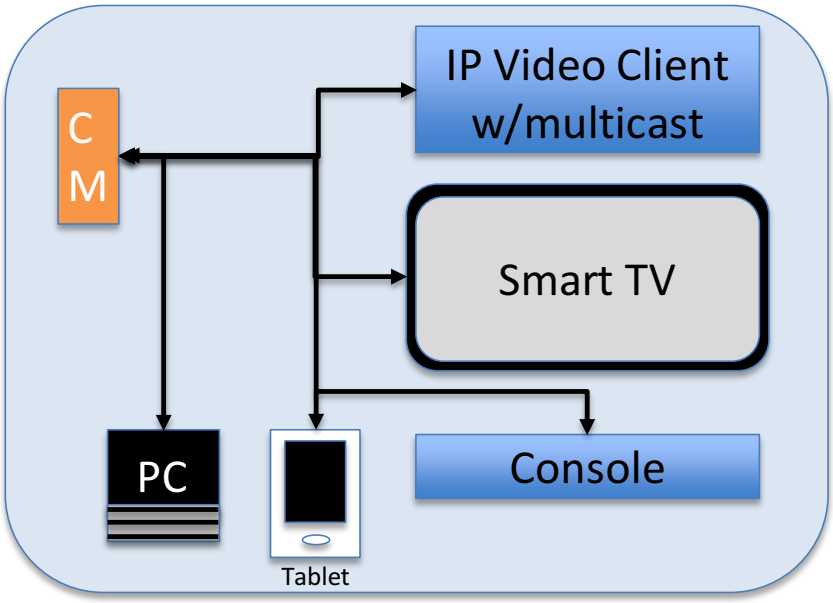

Figure 5: Home network without a home gateway

IP Video Client w/multicast

Some cable providers may choose to not deploy home gateways and instead choose to deploy more sophisticated IP video clients. In this case, multicast sessions are initiated by the video client. Furthermore, the multicast IP sessions flow all of the way to the IP video client. The video client itself would be responsible for any and all program buffering.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Data over Cable System Interface Specification

Version I01 of the Data over Cable System Interface Specification (DOCSIS) was first published on March 26, 1997 with the intent of deploying high-speed data communications over cable television systems. Cable Television Laboratories, Inc. (CableLabs), and its member and vendor companies, prepared a series of specifications that permit the definition, design, development, and deployment of data-over-cable systems on a uniform, consistent, open, non-proprietary, multi-vendor interoperable basis. The service allows transparent bi-directional transfer of IP traffic between the cable system headend and customer locations, over an all-coaxial or hybrid fiber-coax HFC cable network.

Cable Modem Capabilities

The DOCSIS protocol has published its fifth generation (including DOCSIS 1.0, 1.1, 2.0, 3.0, and 3.1) of specifications in nineteen years. As the protocol has evolved, so has the hardware that runs it. However, due to very careful planning by both the MSOs and vendors, each generation of the DOCSIS protocol has always maintained the ability to coexist with earlier generations of DOCSIS equipment. One of the primary tools to preserving that ability is the notion of a capabilities negotiation between the cable modems (CMs) and the network equipment.

Due to traffic engineering concerns, the authors do not recommend that CMs/HGs with fewer than 24 SC-QAM downstream channel receivers be used for IP video services.

Number of CM Receivers and Transmitters

Pre-3.0 DOCSIS CMs

The very first commercial cable modems supported one downstream (DS) single-channel QAM (SC-QAM) and one upstream (US) single-channel QAM (SC-QAM). This model prevailed thru DOCSIS 1.0, DOCSIS 1.1, and DOCSIS 2.0.

- 1 DS SC-QAM receiver x 1 US SC-QAM transmitter

DOCSIS 3.0 CMs

Cable modems with each of these combinations were supported in the DOCSIS 3.0 specifications where multiple downstream SC-QAM and multiple upstream SC-QAM channel support was coupled with channel bonding in both the downstream and upstream directions.

- 4 DS SC-QAM channel receivers x 4 US SC-QAM channel transmitters
- 8 DS SC-QAM channel receivers x 4 US SC-QAM channel transmitters
- 16 DS SC-QAM channel receivers x 4 US SC-QAM channel transmitters

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

- 24 DS SC-QAM channel receivers x 8 US SC-QAM channel transmitters
- 32 DS SC-QAM channel receivers x 8 US SC-QAM channel transmitters

DOCSIS 3.1 CMs

The DOCSIS 3.1 specifications built upon the multiple SC-QAM channel and channel bonding capabilities of DOCSIS 3.0 by adding OFDM downstream and OFDMA upstream channels. When configured for their maximum settings, these OFDM and OFDMA channels are capable of supporting much more bandwidth than their SC-QAM counterparts.

- 32 DS SC-QAM + 2 DS OFDM channel receivers x 8 US SC-QAM + 2 US OFDMA channel transmitters

Number of Downstream Service Identifiers

DOCSIS 3.0 introduced the downstream service identifier (DSIDs) as a context tag that is inserted into the DOCSIS downstream extended header of all packets that belong to a downstream bonded service flow and/or a multicast group service flow. A DOCSIS 3.0 CM is required to support at least sixteen resequencing DSIDs (and associated bonding contexts) and at least sixteen multicast DSIDs (and associated multicast contexts). DOCSIS 3.1 CMs are expected to support more than this minimum number of DSIDs but the specification minimum remains the same.

Channel and Flow Assignment

DOCSIS 3.0 introduced the concept of attribute-based resource assignment. A service flow creation request – whether from a CM configuration file or some form of dynamic service flow signaling – can indicate that the resources used to service the request need to have a specific combination of attributes. This specific combination of resource attributes is compared by the cable modem termination system or converged cable access platform (CMTS/CCAP) to the resource masks of all of the available single channels and bonding groups.

Attribute Mask Bit Definition

The high-order half of the mask (16 bits) is reserved for DOCSIS specification use; 3 bits are currently defined in DOCSIS 3.0 and DOCSIS 3.1:

- Bonded – Highest order bit; in a service flow creation request attribute mask, this bit indicates whether the service flow should be bonded. In a resource attribute mask, this bit indicates whether the resource is a single channel or a bonding group.
- Low Latency – 2nd highest order bit; in a service flow creation request attribute mask, this bit indicates whether the service flow should use resources that allow for a low latency path. In a resource attribute mask, this bit indicates whether

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

the resource allows for a (relatively) low latency path. The definition of exactly what low latency means in this context is left to the CMTS/CCAP vendor.

- High Availability – 3rd highest order bit; In a service flow creation request attribute mask, this bit indicates whether the service flow should use resources that allow for high-availability path. In a resource attribute mask, this bit indicates whether the resource allow for a (relatively) high-availability path. The definition of exactly what high-availability means in this context is left to the CMTS/CCAP vendor.

The low-order 16 bits of attribute masks are available for MSO definition and use.

Some possible suggested assignments for MSO-defined attribute bits:

- IP Video service
- IP Video multicast (might be a subset of channels that support IP Video service above)
- Telephony

Service Flow Request Attribute Masks

Each service flow creation request is associated with three Attribute Masks:

- Required Attribute Mask - The CMTS only uses channels or bonding groups with these bits set in the attribute mask
- Forbidden Attribute Mask - The CMTS only uses channels or bonding groups with these bits NOT set in the attribute mask
- Attribute Aggregation Rule Mask - If the CMTS/CCAP is capable and configured to do so, the CMTS/CCAP creates a bonding group that is a logical combination of the attributes that correspond to these bit positions
    - A ‘1' in this mask for an attribute means that each channel must have this attribute
    - A ‘0' in this mask for an attribute means that at least one channel must have this attribute

These service flow creation request attribute masks may be either directly signaled in the request itself (as they might be in a PacketCable Multimedia gate) or they may be defined as part of a provisioned service class on the CMTS/CCAP (as they might be for a signaled multicast JOIN).

Resource Attribute Masks

Every configured resource (i.e. channels and bonding groups) in the system has a resource attribute mask denoting the capabilities of that resource. This resource attribute mask is compared to the attribute masks of the service flow request. The channel and bonding group attributes are not connected; the attribute mask of a single

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

channel and the attribute mask of a configured bonding group which includes the channel may not have matching attribute bits.

Resource Mask Matching for Assignment

The CMTS/CCAP attempts to assign service flows to channels or bonding groups such that all required attributes are present and no forbidden attributes are present. In this way, resources may be configured such that certain services (like IP Video) may be steered toward the resource or away from the resource.

Service Flow Types

A service flow is a MAC-layer transport service that provides unidirectional transport of packets either upstream to the CMTS/CCAP from the CM or downstream to the CM from the CMTS/CCAP. The following headings denote different ways in which service flows might be described. These types are orthogonal in nature and thus can be combine to describe a particular service flow (e.g. an unbonded dynamic service flow might be used for telephony whereas a bonded static service flow might be used to carry unicast IP video).

Unbonded vs. Bonded

A bonded service flow is one in which the data of the service flow is transported across the channels of a bonding group. A bonding group, in turn, is a set of single channels associated with a single MAC Domain which reaches a common set of cable modems and which transports at least one bonded service flow. Bonded service flows may either be sequenced, in which case the packets are guaranteed to not arrive out of order (within some sequencing interval); or they may be unsequenced, in which case the application must not consider order important. Most bonded service flows are sequenced. Bonding groups may be configured or they may be dynamically created when the request is received by the CMTS/CCAP as guided by the service flow request Attribute Aggregation Rule Mask.

An unbonded service flow is one in which the data of the service flow is transported on only a single channel associated with a MAC Domain which reaches a common set of cable modems. Unbonded service flows may have less latency than bonded and sequenced service flows due to no need for resequencing.

Static vs. Signaled

Service flows may be either statically created via provisioning or dynamically created during runtime. Service flows may be statically created at CM registration time by the back office provisioning server placing service flow definitions into a CM configuration file which is downloaded by the CM as part of the initialization process. A service flow

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

might also be dynamically-created as a result of application signaling (Ex. IGMP, PacketCable Multimedia, etc.) during runtime.

MAC Management Dynamic Services Transactions

The DOCSIS protocol has gotten ever more complex over its almost 20-year history. The process that started as a simple QoS engine controlled via a static configuration file has morphed into a complicated dynamic services engine with multiple transactions per service request. The result is that service requests can take a long time to process and they have many more ways to fail than they once did.

The DOCSIS protocol uses a 3-way transaction for critical signaling operations between an initiator and a target. In many cases, the CMTS/CCAP MAC Domain is the initiator and the CM is the target but there are some transactions in which the CM might be the initiator. A 3-way transaction typically takes the form of the initiator transmitting a request (with parameters) to the target (and running a timer while waiting for a response), the target performing an action and transmitting a response (with results) back to the initiator (and running a timer while waiting for an acknowledgement), and then the initiator transmitting a final acknowledgement (ideally with no important parameters) to the target to inform the target that the response was received. If one of the timers expires then the corresponding message transmission is retried (up to a certain number of total transmissions) and the timer is run again. Typically there may be as many as three retransmissions of either the request or the response. The acknowledgement has no critical data so no timer is run.

A DOCSIS MAC Management Dynamic Services transaction is necessary for any of the following actions:

- Dynamic Services Addition (DSA): Create a service flow (3-way transaction)
- Dynamic Services Change (DSC): Change an existing service flow (3-way transaction)
- Dynamic Services Deletion (DSD): Delete an existing service flow (2-way transaction only)
- Dynamic Channel Change (DCC): reassign the channel set of a CM with reboot (3-way transaction)
- Dynamic Bonding Change (DBC): modify channel or bonding parameters (3-way transaction)
    - Modifications to CM Transmit Channel Configuration (TCC)
    - Modifications to Service Identifier (SID) Cluster parameters for upstream bonding
    - Modifications to CM Receive Channel Configuration (RCC)

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

  - Modifications to downstream service identifiers (DSIDs) parameters for downstream multicast filtering and/or bonding
  - Modifications of Security Associations (SA)

Some types of service requests may require the chaining of multiple types of these transactions. For example, inherent in the process of creating/modifying/deleting a unicast service flow in one direction is a DOCSIS MAC Management Dynamic Bonding Change transaction to create/change/delete the parameters necessary for a bonded or service flow followed by a DOCSIS MAC Management Dynamic Services Add/Change/Delete transaction. Per the DOCSIS protocol, the CM can only handle only one Dynamic Services MAC Management transaction at a time so these two transactions must be done serially.

Each of these DOCSIS MAC Management transactions involve signaling over a possibly noisy and congested RF Plant whereby up to four message transmissions (at 1-second timeout intervals) may be required for each message.

Quality of Service

Classifiers associate packets into exactly one Service Flow. The Service Flow Encodings provide the QoS Parameters for treatment of those packets on the RF interface.

Classifier

A Classifier is a set of matching criteria applied to each packet entering the cable network which consists of some packet matching criteria (destination IP address, for example) and a classifier priority. A QoS Classifier additionally consists of a reference to a service flow. If a packet matches the specified packet matching criteria of a QoS Classifier, it is then delivered on the referenced service flow [8].

The following are all options for use in classifiers:

- Priority: determines the search order for the table. If a packet matches more than one classifier, the classifier with the highest priority is used to determine which service flow will carry the packet.
- Source
  - MAC Address
  - IP address
  - IP address mask
  - TCP/UDP Source Port Start, TCP/UDP Source Port End
- Destination
  - MAC Address

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

    - IP address
    - IP address mask
    - TCP/UDP Destination Port Start, TCP/UDP Destination Port End
- Ethertype/Service Access Point (SAP)
- Protocol
- IP TOS/DSCP Range/Mask
- 802.1P Priority Range
- 802.1Q VLAN ID

QoS Parameters

The QoS Parameters of a service flow consists of three sets of Service Flow Encodings – one for either the Configured/Authorized Set, one for the Admitted set and one for the Active set. Service Flow Encodings which are defined in a CM configuration file or which are signaled as authorized are considered the Configured/Authorized Set. The Service Flow Encoding which corresponds to resources that have been reserved for use by a service flow are considered to be the Admitted set. The Service Flow Encoding which corresponds to the resources that have been activated for the service flow is considered the active set.

A service flow encoding consists of the following parameters

- Service Class Name – used as a key to identify a predefined service flow encoding on the CMTS/CCAP. If used in a configured or signaled service flow request, then the definition of the predefined service flow encoding on the CMTS/CCAP is used as the service flow encoding of the requested service flow. Some parameters of the predefined service flow encoding may be overridden in the service flow request
- Service Flow Required Attribute Mask – see section: Service Flow Request Attribute Masks
- Service Flow Forbidden Attribute Mask– see section: Service Flow Request Attribute Masks
- Service Flow Attribute Aggregation Rule Mask– see section: Service Flow Request Attribute Masks
- Traffic Priority - 0 to 7 — higher numbers indicate higher priority. The default priority is 0 (lowest)
- Maximum Sustained Traffic Rate – (a.k.a. T_{max}) a token-bucket-based rate limit for packets; does not limit the instantaneous rate of the Service Flow. If this parameter is set to zero, then there is no explicitly-enforced traffic rate maximum

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

- Maximum Traffic Burst - token bucket size (in bytes) for this Service Flow. The minimum value is 1522 bytes. The default value is 3044 bytes. This parameter has no effect unless a non-zero value has been provided for the Maximum Sustained Traffic Rate parameter
- Minimum Reserved Traffic Rate - (a.k.a. T_{min}) a CMTS/CCAP should schedule forwarding of all service flows' traffic such that each receives at least its Minimum Reserved Traffic Rate when transmitting packets with the Assumed Minimum Reserved Rate Packet Size. This parameter defaults to a value of 0 bits/sec (i.e., no bandwidth is reserved for the flow by default)
- Assumed Minimum Reserved Rate Packet Size - If the service flow sends packets of a size smaller than the Assumed Minimum Reserved Rate Packet Size, such packets will be treated as being of the Assumed Minimum Reserved Rate Packet Size for calculating the rate forwarded from the service flow for purposes of meeting the Minimum Reserved Traffic Rate
- Peak Traffic Rate - (a.k.a. T_{peak}) rate parameter P of a token-bucket-based peak rate limiter for packets of a service flow. Configuring this peak rate parameter permits an operator to define a Maximum Traffic Burst value for the Maximum Sustained Traffic Rate much larger than a maximum packet size, but still limit the burst of packets consecutively transmitted for a service flow. If this parameter has a value of 0, peak traffic rate is not limited by the QoS engine and is limited only by the line rate of the channel or bonding group

Admission Control

The DOCSIS protocol supports a two-phase activation model in which authorized resources for an application are first "admitted," and then (perhaps after an end-to-end service negotiation is completed) the resources are "activated", although these two phases may be and often are performed simultaneously. This model serves the purposes of:

- conserving network resources for other uses until a complete end-to-end service negotiation has been established
- performing policy checks and identifying available resources as quickly as possible before committing any resources for the service

At a high level, admission control is a CMTS/CCAP function that is used to determine when adequate resources are available to permit a new service flow to be established. If adequate resources are not available, the new service flow cannot be admitted (or activated).

Primary-Capable Downstream Channels

A DOCSIS cable modem must receive certain pieces of information from its CMTS/CCAP MAC Domain in order to enter and remain in service. This information is primarily timing

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

information (SYNC messages from SC-QAM DS channels, PHY-Link Channel (PLC) timestamp from OFDM DS channels) but it also includes downstream and upstream topology ambiguity resolution information from MAC domain descriptor (MDD) messages (helps to identify which channels a CM is expected to be able to use), as well as upstream channel descriptor (UCD) and Upstream Bandwidth Allocation (MAP) messages for at least one upstream channel in each of the service groups that the downstream channel reaches. A downstream channel (whether SC-QAM or OFDM) that carries this information is known as primary capable because the channel is capable of being used by the CM as its primary channel timing and information source.

Different generations of CMs have slightly different requirements with regard to their downstream primary channel requirements. Pre-3.0 DOCSIS CMs have only one downstream channel receiver so the downstream channel that is assigned to them must be primary capable. A multiple-downstream-channel DOCSIS 3.0 or later CM need only be assigned one primary capable downstream although often they receive multiple primary capable downstream channels and are told which channel to use as the CM's primary channel. DOCSIS 3.1 CMs allow the option of being provided a prioritized list of primary capable downstream channels such that the CM may take fault recovery actions if the CM's primary downstream were to ever fail. Judicious (i.e. sparse) assignment of primary capability can save bandwidth across the service group.

Upstream Supervision

The Upstream Channel Descriptor (UCD) and Upstream Bandwidth Allocation (MAP) DOCSIS MAC Management messages for a single upstream channel are sometimes referred to as upstream channel supervision. As indicated in section 0, upstream supervision for at least one upstream channel must be carried on each primary capable channel.

While supervision for every upstream channel must be carried on at least one downstream channel, not every upstream channel's supervision need be carried on a primary downstream for a DOCSIS 3.0 or later CM to use the upstream. In this case, each CM must be able to receive supervision for each assigned upstream channel on one of its downstream channels. Judicious assignment of upstream channel supervision can save a significant amount (0.5 to 1 Mbps per US channel supervision) of downstream bandwidth that may be used for services such as IP Video.

Energy Management

DOCSIS 3.0 and later CMs support an energy management 1x1 mode whereby the CM monitors traffic flow and, upon sensing a low utilization, requests that the CMTS/CCAP reduce its active channel set to one upstream and one downstream. Obviously, such a reduction would preclude the reception of a bonded unicast or multicast IP video flow.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

DOCSIS 3.1 CMs support an energy management DOCSIS Light Sleep (DLS) mode whereby the CM monitors traffic flow and, upon sensing a low utilization, requests that the CMTS/CCAP place the CM into a temporary sleep mode. Obviously, this mode would preclude the reception of a bonded unicast or multicast IP video flow.

In order to prevent a CM from requesting energy management 1x1 or DLS operation, the type, length, value (TLV) parameter of the energy management Downstream Entry Bitrate Threshold (TLV 74.2.1.1) in the CM's configuration file must be either set below the minimum bitrate of a single IP Video program or set to a value of zero to disable energy management activity detection.

IP Multicast

Static

See section: Unicast Flows.

Signaled

Multicast client protocol signaling (eg. IGMPv3 for an IPv4 network or MLDv3 for an IPv6 network) is the usual method used to join clients to multicast group streams. See section: In-band Multicast Protocol Signaling.

The PacketCable Multimedia (PCMM) protocol has extensions for managing IP multicast group service flows. See section: PacketCable Multimedia.

QoS Control Plane Signaling

Cable Modem Configuration File

During initialization, each CM must download a cable modem configuration file. That CM configuration file may contain many pieces of information that determine how the CM behaves. Many of these pieces of information are sent to the CMTS/CCAP as part of the registration process. Unicast packet classifiers and service flow definitions are the main pieces used to control QoS. Once installed as part of the CM registration process, these definitions cannot be changed without a CM reinitialization.

Unicast Flows

A unicast flow and the associated classifier(s) may be created in the cable modem's configuration file. This is the primary method that is used to create the unicast service flows that are used for cache misses for multicast assisted adaptive bit rate (M-ABR) IP video. If created via this method, the classifier definition must accommodate all traffic that is expected for the flow as it cannot be changed to accommodate dynamic services.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Multicast flows

When supported by the CMTS/CCAP, one or more CMTS Static Multicast Session Encoding (TLV 64) parameters may be configured in the CM configuration file to cause the home gateway to statically JOIN the most popular multicast streams at CM registration time. When the CMTS/CCAP sees the CMTS Static Multicast Session Encoding in the registration request message, it will perform all of the processing necessary to add the CM to the multicast group as if the CM had issued a multicast signaled JOIN. The registration response message will contain all of the DSID information necessary for the CM to receive the multicast group session. This way, the Home Gateway will not need to issue multicast signaling (IGMP or MLD) requests for these statically JOINed group sessions. The Home Gateway may choose to either store or discard traffic for these sessions when another program is being viewed.

The group IP multicast address would be provisioned as the Static Multicast Group Encoding (TLV 64.1) It is recommended wherever possible that source-specific multicast (SSM) be used by specifying the Static Multicast Source Encoding (TLV 64.2) for each CMTS Static Multicast Session Encoding. The Static Multicast CMIM Encoding (TLV 64.3) must be set to forward packets to the eSTB-IP interface [Bit 17 (0x00 00 40)].

All DOCSIS 3.0 or later CMs support a minimum of 16 Static Multicast MAC addresses.

PacketCable Multimedia

The PacketCable Multimedia (PacketCable Multimedia protocol is one way that is currently defined to signal the creation, modification, and deletion of dynamic unicast and group service flows and associated classifier(s) for certain applications such as telephony.

The PacketCable Multimedia protocol consists of an application manager sending an “application request gate” containing information about CM identity, flow classifiers, QoS settings, etc. to a policy server who then “authorizes” (according to some local policy) and “normalizes” (modifies to properly prioritize amongst possibly multiple applications) the request before sending the gate to the CMTS/CCAP as a directive.

If the gate from the policy server is deemed to be acceptable by the CMTS/CCAP (i.e. passed access control checks, etc.), then the CMTS/CCAP must install the gate by creating/modifying/deleting a service flow in one direction and then replying to the policy service with a transaction status and possibly identifying information about the service flow. This reply is forwarded back to the application manager.

Inherent in the process of “creating/modifying/deleting a service flow in one direction” is the DOCSIS MAC Management Dynamic Services transaction processing between the CMTS and the CM that is described in section 0.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

The fastest possible successful “Gate-Set” transaction with all its underlying DOCSIS transactions may require fifty milliseconds. If DOCSIS MAC management request or response messages are lost (perhaps due to noise?) then they will typically be retransmitted three more times (each with a timeout) until the transmitting device receives the response from the other device. If the DOCSIS MAC Management transactions each catch only the last retransmission of each of the messages then the successful gate set could take as long as seven seconds for the successful case. It would be safest to assume that at least one message gets lost during each Gate-Set transaction so we should assume a “normal“ channel change time of about 1.25 seconds. An unsuccessful Gate-Set would likely take about 5 seconds since the DOCSIS Dynamic Services transactions cannot be interrupted.

Meanwhile, if the subscriber tunes to another channel while waiting somewhere between fifty milliseconds and five seconds for the previous channel change to occur, since the DOCSIS transactions cannot be interrupted, they would need to complete before handling a new Gate-Set for a new channel for the same CM. If the Gate-Set gets triggered by a trio (because most multicast IP Video applications assume that three requests are always safer than one) of multicast protocol LEAVEs for the old program channel followed by three JOINs for the new program channel, then each individual request will look like a separate channel change to the CMTS/CCAP device and will result in a DOCSIS transaction logjam.

Due to the possibility (likelihood?) of rapid channel change behavior causing this DOCSIS transaction logjam, the PCMM architecture is **not** recommended for an IP video system.

In-band Multicast Protocol Signaling

An IP multicast signaling protocol such as IGMP (for IPv4) or MLDv2 (for IPv6) may be used by a home gateway to dynamically JOIN and LEAVE IP multicast group sessions. The DOCSIS multicast session establishment is initiated by the first home gateway to JOIN the session and the multicast session is terminated sometime after the last home gateway has issued a LEAVE for the session.

When an IP video player switches TV channels, there is a good possibility that new TV channel is available via a multicast session as well. This results in the home gateway issuing a LEAVE for the old channel followed by a JOIN for the new channel. Each JOIN or LEAVE will result in a 3-way DOCSIS MAC Management Dynamic Bonding Change transaction to create/change/delete the parameters necessary for a bonded group service flow (but, unlike what would be necessary for a unicast flow, there is no need for a 3-way DOCSIS MAC Management Dynamic Services Add/Change/Delete transaction). It is necessary for the home gateway to always issue the LEAVE because the CM has a limited number of bonded multicast DSID contexts [see section: Number of Downstream Service Identifiers].

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

As might be expected, when television programs change every half hour, there may be large number of program channel changes encountered by an IP video system. In kind, this would create a flood of LEAVE and JOIN multicast signaling for the IP video system to process. Worse yet, with the M-ABR architecture, each time that a home gateway is asked for a IP video chunk that it does not have, it will issue a unicast HTTP request to the IP video servers to retrieve the missing segment file. This action will cause temporary blooms in unicast IP activity on the DOCSIS and backbone networks until the home gateway can start filling requests with video segments that have been received over an IP multicast group session. Some techniques (such as the use of static multicast for short tail TV streams) may be used to mitigate this flood of signaling.

NACK-Oriented Reliable Multicast (NORM)

The NACK-Oriented Reliable Multicast (NORM) was developed by United States Naval Research Laboratory (NRL) Protocol Engineering Advanced Networking (PROTEAN) Research Group to serve multicast segments to gateways. NORM utilizes content-level forward error correction (FEC) to minimize the potential for loss and has the ability to retransmit over multicast in case multiple receivers missed the same content transmission - at the cost of latency.

NORM has been identified by IP video industry experts working thru CableLabs as the protocol of choice for delivery of M-ABR video. The primary reason for this is that NORM is extremely flexible. For example, it can be utilized to transfer data streams or files, with or without FEC and with just error detection or also multicast retransmission. Operators who want maximal multicast efficiency may want to retransmit over multicast at the cost of latency, but other operators may be willing to accept reduced multicast efficiency to minimize latency.

The optimal configuration may not only vary between operators, but might even vary across regions within the same operator.

See reference 3 in conjunction with reference 6 for more details.

RECOMMENDATIONS - TYING IT ALL TOGETHER

A DOCSIS IP video distribution system should be deployed in such a way as to maximize stability and efficiency and minimize latency and perturbations to the system. All resources should be configured once as early as possible (ideally, at initialization) and service request activity such as a TV program channel change should, whenever and wherever possible, not cause the reconfiguration of resources (channel assignments, bonding group changes, etc.) in the system. These types of operations, while supported

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

(or at least not disallowed) by the DOCSIS protocol, can result in excessive latency and significant opportunity for system error when ideal conditions are not present. Furthermore, system configuration should leverage the advantages of the available technologies as much as possible to minimize the resources used to provide service to the subscriber base.

This overarching philosophy is a strong motivator of many of the following recommendations.

CM Capabilities

The authors recommend that only CMs and home gateways that are capable of receiving at least twenty-four (preferably 32) SC-QAM downstream channels be used for IP video service. This recommendation (see Figure 6) is based on the relatively low penetration of OFDM-capable CM and home-gateway devices at the present time (mid-2016). In addition, a CM device of this size allows for a multicast IP bonding group, an encompassing unicast-ABR bonding group, and an encompassing high-speed data bonding group.

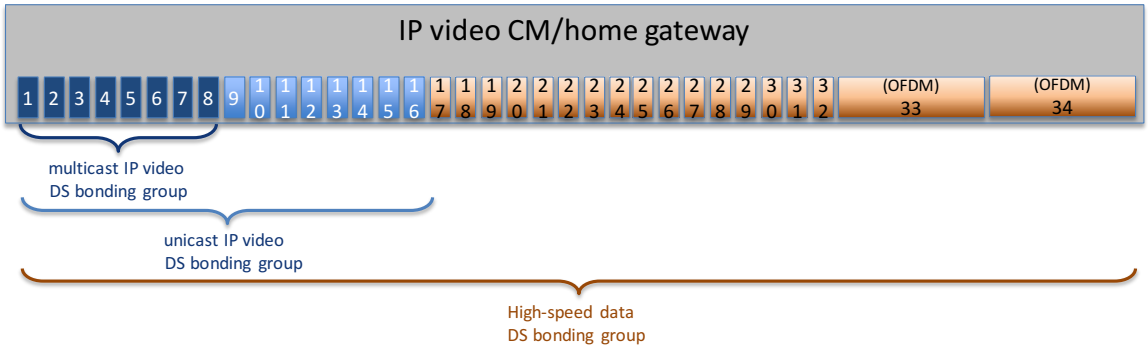

Figure 6: Example downstream channel assignment for 32 DS device

Channel Assignment to CMs

Consolidate Multicast Program Streams

While the authors do not recommend that CMs with fewer than twenty-four SC-QAM downstream receivers be used for IP video services, some MSOs might ignore this recommendation and use CMs with a smaller number of downstream receivers. Similarly, in the future, CMs are expected to increase in capability so today's twenty-four SC-QAM downstream CMs may become the least capable CMs in the future. In either case, in order to maximize the multicast benefit (and minimize bandwidth) without having to retune CM receiver channel configuration and downstream bonding group configuration during TV channel-zapping time, the authors recommend that all IP video multicast group sessions be consolidated onto a number of DOCSIS SC-QAM

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

downstream channels that is (ideally, much) smaller than the number of downstream SC-QAM receivers of the least capable CM that will be used for IP video services. Furthermore, a bonding group for all multicast sessions can be created for this entire set of channels to allow for maximum bandwidth with minimum latency. All CMs which are to be used for IP video service should be assigned to this same IP video downstream bonding group.

As an example, (see Figure 1) if a MSO chooses to use CMs with sixteen, twenty-four, and thirty-two SC-QAM receivers, the recommendation is to limit (by careful traffic planning) the total amount of IP video multicast session traffic to fit in a relatively small IP multicast downstream bonding group (DSBG), in this example, eight channels. Also, by carefully applying DOCSIS QoS parameters (priority, attribute bits, etc.), multiple unicast IP video bonding groups can be defined which includes the same channels of the IP video multicast bonding group. Furthermore multiple high-speed data bonding groups can be defined which overlaps all of these IP video bonding groups plus any OFDM channels (if available).

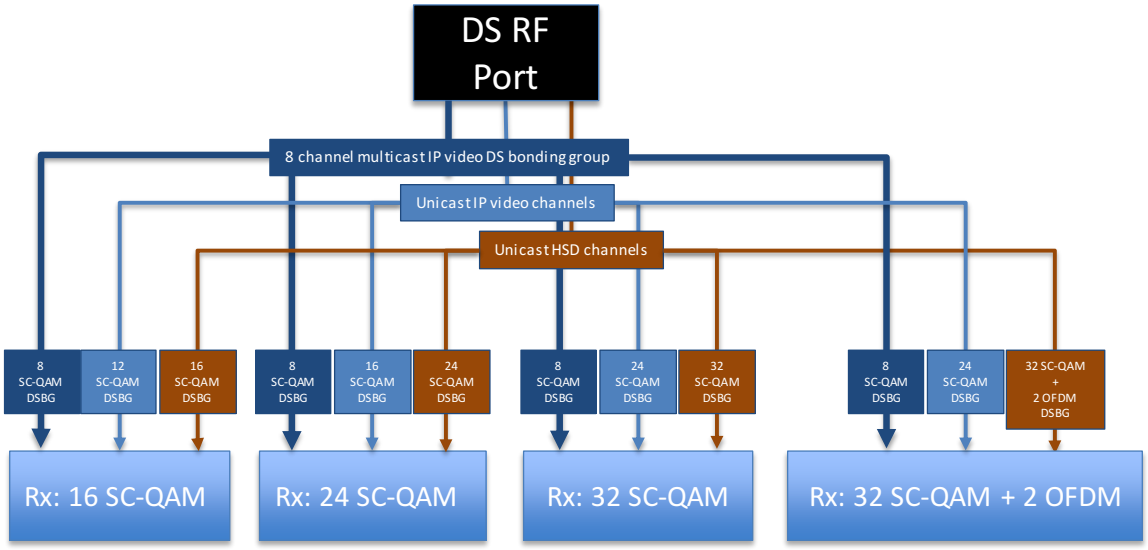

Figure 1: Example configuration of DS channels across CMs with different numbers of receivers

In Figure 1 the entire one and only multicast IP video downstream bonding group is assigned to all CMs which are participating in IP Video. Then a unicast IP video bonding group is defined for each CM size which uses the channels of the multicast IP video downstream bonding group as well as other downstream channels to carry the (lower priority) unicast IP video traffic. This way the IP video unicast traffic can use any extra capacity left on the multicast IP video downstream bonding group while not impacting the multicast IP video traffic. Finally, a large unicast bonding group is defined for each CM size for other services (such as high-speed data) at a lower priority than unicast IP

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

video traffic and these services may use all of the channels of the CM while not impacting the higher priority traffic.

Static Multicast for IP video Delivery

If some number of linear TV streams can be identified that are nearly always viewed by a relatively large number of viewers at any one time, then these TV streams are excellent candidates to be carried over statically provisioned multicast groups via the use of the DOCSIS CMTS Static Multicast Session Encoding (TLV Type 64 in a CM configuration file). Statically configuring these groups will help to alleviate the need for multicast (IGMP or MLD) signaling; thereby saving bandwidth, signaling time, and latency.

For example, in the real world observation example of Figure 1, all ten of the most popular linear TV streams might be configured as static multicast joins by using the CMTS Static Multicast Session Encoding (TLV type 64) in the CMs' configuration files. For this particular system, doing so might eliminate up to 95% of the multicast signaling that may otherwise occur with IP video channel changes.

Signaled Multicast for IP video Delivery

If TV program stream popularity is volatile, then an MSO may choose to not include it in a set of statically assigned multicast streams. However, for program streams which are popular at least part of the time, a significant bandwidth savings in the access network can be achieved by allowing these programs to be carried as a signaled group IP multicast session. This way, the home gateway of each viewer of a program stream will signal when the viewer is watching the stream so that the network resources are only assigned when they are needed. In addition, single multicast IP group sessions can be shared amongst all viewers' home gateways (when there is more than one) at once. It is certainly possible that all program streams that are not included in the set of statically assigned multicast streams could use the signaled multicast approach described in this section. An MSO may choose this approach if they wish to minimize bandwidth utilization resulting from IP video delivery, because the signaled multicast approach ensures that any program stream viewed by one or more viewers within a service group will typically be transmitted only once within that particular service group. Alternatively, the MSO could choose to use the signaled multicast approach for medium-tail programs, and then use the unicast approach (of the following section) for long-tail programs.

Unicast for IP video Delivery

If a particular set of TV program streams are of the special interest variety (i.e. not very popular across the entire subscriber base) and only viewed by a small number of viewers at any given time, then an MSO may choose to not include the specialty

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

programs in the set of statically assigned multicast streams and also not include them in the set of signaled multicast streams. This implies that the home gateways will always have to use unicast ABR requests to receive these program streams. This approach may be appealing to MSOs who do not wish to complicate their networks with all signaled multicasting. These unicast streams would typically be planned as part of the normal high-speed data channels mentioned below.

Load-Balance Remainder of Channels

Once the set of SC-QAM channels that are needed for short and medium tail multicast IP video have been allocated, all of the remaining resources should be considered unicast high-speed data (HSD) resources and may be used for load balancing all non-multicast traffic (including unicast ABR). Depending on the configuration, this load balancing may be performed at CM registration time via the CMTS/CCAP load-balancing mechanisms.

Primary downstream assignment

A DOCSIS 3.0 or later CM needs only one primary downstream channel but may be assigned many primary-capable downstream channels. Because primary capable downstream channels carry more DOCSIS signaling overhead than non-primary capable downstream channels, the downstream channels that are used to carry the IP multicast IP video program streams should not be configured to be primary capable.

As for the remaining high-speed data (HSD) channels, unless the population of pre-3.0 DOCSIS CMs is high, only a small handful of these channels should be configured to be primary-capable. Ideally, if sets of channels are created as described above, only one or two channels from each set would be configured as primary capable.

Upstream supervision

For many of the same reasons cited in section 0 about primary capability, the downstream channels that comprise the IP multicast bonding group should not be provisioned to carry supervision for any upstream channel.

As for the remaining high-speed data (HSD) channels, any pre-3.0 CM must be able to receive supervision for every upstream channel to which it is assigned. For this reason, all supervision for upstream channels to which a pre-3.0 CM is to be assigned should only be carried on primary-capable downstream channels. Supervision for upstream channels on which only DOCSIS 3.0 or later CMs are to be assigned need not be carried more than once in each downstream channel set that is assigned to a D3.0 or later CM.

CM Unicast IP Video Service Flow QoS Parameters

Modeling and simulation has resulted in a recommendation that a unicast DOCSIS service flow which is to carry an adaptive bitrate (unicast) IP video stream should be

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

constrained to no less than twice the maximum bitrate of the highest resolution version of the stream. This maximum bitrate will accommodate the majority of bursty media traffic and allow the stream to refill video player jitter buffers in the face of any transient cable plant noise.

Downstream

Unicast

A best effort downstream service flow for use by IP video unicast service should be created in the CM configuration file with a slightly higher priority than the normal high-speed data service flow. This IP video downstream unicast service flow should have a maximum sustained bitrate equal to at least twice the maximum bitrate of the highest resolution of a single stream multiplied by the integer number of unicast program streams to be simultaneously served to a household (which will all be carried within that downstream unicast service flow).

A high-priority downstream classifier should be created in the CM configuration file to reference the IP video downstream unicast service flow. The classification field specifics are different from MSO to MSO and possibly even from system to system within the same MSO. Some MSOs may dedicate a TOS/DSCP marking to all unicast IP video packets. Some MSOs may define an expected source host or subnet for all unicast IP video packets. The MSO must define the classifier so that the IP video unicast packets can be identified and classified by a DOCSIS classifier.

Multicast

Each IP multicast group of an M-ABR architecture is concerned only with transferring a single resolution of the program stream in question. It is recommended that a single multicast program stream be associated with a single IP multicast group, so the transfer of a number of multicast program streams to a particular service group will require that same number of IP multicast groups be transported to that service group. Using NORM (see section: NACK-Oriented Reliable Multicast (NORM)), there may be extra data transferred in the form of FEC to recover nack'ed frames. If using proactive FEC, an IP video downstream NORM-Based multicast service flow should have a maximum sustained bitrate equal to about 150% of the maximum bitrate of the program stream to account for NORM FEC overhead. The CableLabs M-ABR document [3] recommends unicast repair using HTTP Range-Requests (an HTTP term for accessing pieces of a video chunk) for repairing missing video segment data be used in the small number of cases where proactive FEC cannot correct the segment.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Upstream

A best effort upstream service flow for use by IP video unicast service with a slightly higher priority than the normal high-speed data service flow should be created in the CM configuration file. This IP video upstream unicast service flow should have a maximum sustained bitrate equal to about 5% of the maximum rate of the corresponding downstream IP video unicast service flow. This IP video unicast service flow will carry all forwarded HTTP requests and all multicast signaling.

A high-priority upstream classifier should be created in the CM configuration file to reference the IP video upstream unicast service flow to help ensure low-latency delivery of HTTP GETs and multicast JOINs. This classifier is used at the CM to determine upon which flow to send the signaling and TCP acknowledgements. The per-field specifics for the classifier are different from MSO to MSO and possibly even from system to system within the same MSO. Some MSOs may dedicate a TOS/DSCP marking to all unicast IP video packets. Some MSOs may define an expected destination host or subnet for all upstream unicast IP video packets. The MSO must define the classifier so that the IP video unicast packets can be identified.

QoS Settings across Different Services

The IP video multicast streams will be the primary transport for video to the majority of the subscribers so the integrity of these service flows is considered very important.

Priority 5: VoIP

Voice over IP (VoIP) traffic tends to be a low-bandwidth service that is extremely sensitive to both latency and packet loss. The default priority for PacketCable VoIP services tends to be pretty high at a DOCSIS priority of 5. Since VoIP is a very high priority, VoIP services should not normally be carried on DOCSIS channels that are used for IP video multicast flows. If this is absolutely necessary, then admission control should be used to ensure that these two high-priority services do not interfere with one another.

Priority 3: Multicast IP video

The DOCSIS traffic priority of IP video multicast traffic should be such that no other traffic (perhaps with the exception of VoIP) should be higher so that in cases of congestion on these channels, the IP video packets will always get through. For this reason, a DOCSIS priority of 3 may be appropriate. Also, by prioritizing multicast IP video on the DOCSIS channels that carry it, any excess capacity of these channels can be used by lower priority services without impacting multicast IP video.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Priority 2: Signaling and unicast ABR IP video

Just below the IP video multicast packets in priority is the signaling and unicast ABR IP video traffic. The dynamic sizing nature of ABR will allow the video service to adapt to the congestion that might exist in the network. These service flows should be provisioned with a DOCSIS priority setting below IP video multicast service flows and above normal HSD service flows to keep them from getting dropped or queued behind HSD service flows.

Priority 1: Normal High-speed Data

The normal HSD data services (web surfing, etc.) might be carried at a default DOCSIS priority of 1.

Priority 0: Traffic management "penalty box"

Finally, some MSOs may use a DOCSIS traffic priority of 0 for purposes of managing traffic for some users.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

CONCLUSION

The DOCSIS protocol uses a shared medium but provides a rich set of tools to ensure Quality of Service (QoS) to each subscriber. This paper has described some of the issues in using DOCSIS to transport IP video.

Linear television viewership continues to follow viewership patterns that suit a one-to-many (broadcast or multicast) transmission medium. The popularity curve Figure 1 shows that the vast majority of subscribers can be satisfied with a relatively small number of program streams. This fits in well with the DOCSIS shared-medium architecture if the system is configured properly. To this end, the authors recommend a configuration whereby all CMs participating in the IP video service share a common bonding group that is dedicated to carrying popular program streams in IP multicast groups. These dedicated channels should be designed to support maximum utilization to create the smallest possible shared multicast bonding group to gain bandwidth efficiency.

The DOCSIS MAC Management Dynamic Services facility is both powerful in its capabilities and cumbersome in its realization. In the name of backward compatibility the DOCSIS protocol contains artifacts of historical development that have led to increased latency and transaction error opportunities due to serialization of necessary transactions. In fact, the effects of these artifacts have made the PacketCable Multimedia protocol ineffective in handling a high transaction rate (such as might occur during TV channel zapping) for a single subscriber device and thus, PacketCable Multimedia is not recommended as a control plane architecture for IP video.

Some examples of possible configurations for a MAC-Domain Downstream Bonding group to support CMs of differing sizes are shown in the Appendix of this document. The reader may use these recommendations directly or may tailor them to suit any special circumstances inherent in a particular deployment.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

ABBREVIATIONS & DEFINITIONS

ABR	Adaptive bitrate
bonded service flow	A service flow which is transported across the channels of a bonding group
bonding group	A set of single channels associated with a single MAC Domain which reaches a common set of cable modems and which transports at least one bonded service flow
bps	Bits per second
CableLabs	Cable Television Laboratories, Inc.
CCAP	Converged cable access platform; a device that performs both EQAM and CMTS functionality
classifier	A set of criteria used for packet matching according to TCP, UDP, IP, LLC, and/or 802.1P/Q packet fields. A classifier maps each packet to a single Service Flow.
CM	Cable modem
CMTS	Cable modem termination system
DBC	Dynamic bonding change
DOCSIS	Data over cable system interface specification
DS	Downstream
DSA	Dynamic services addition
DSBG	Downstream bonding group
DSC	Dynamic services change
DSD	Dynamic services deletion
DSG	DOCSIS set-top gateway
DSID	Downstream service identifier
DSL	Digital subscriber line
EQAM	Edge-QAM; a device that supports broadcast digital video and/or switched digital video; and/or video on demand
FEC	Forward error correction
FIFO	First-in, first-out
HFC	Hybrid fiber-coax
HD	High definition
home gateway	A DOCSIS appliance that serves as a unicast HTTP source for IP video content to clients within the home network
Hz	Hertz
HTTP	Hypertext transfer protocol
IGMP	Internet group management protocol
IP	Internet Protocol; may denote either IPv4 or IPv6 or both
IPv4	Internet Protocol, version 4; This earlier version of IP uses a 32-

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

	bit address space, typically written as four integers between 0 and 255 separated by dots. IPv4 does natively support IP broadcasting but does not natively support a multicast signaling protocol
IPv6	Internet Protocol, version 6; This later version of IP uses a 128-bit address space, typically written as sequences of up to four hexadecimal digits and separated by colons. IPv6 does not support IP broadcasting but does natively support a multicast signaling protocol
ISBE	International Society of Broadband Engineers
jitter buffer	A FIFO memory buffer used to store a certain amount (corresponding to a delay in time) of media data before the media is used in real-time; this buffer serves to isolate the real-time player from the variances in reception delay (i.e. Jitter) that are common in packet-based networks
JOIN	A signaled multicast protocol command to join a multicast group
linear (TV)	A program stream whereby the program lineup has been pre-scheduled so that an audience must tune in at a particular time to view a particular program; also known as appointment TV
long-tail programming	A set (of non-determinant size) of the special-interest (least popular) video program streams; these program streams are more likely to be unicast to each requesting user
LEAVE	A signaled multicast protocol request to leave a multicast group
M-ABR	Multicast-assisted adaptive bitrate
MAC	Media access
MAC domain	A logical subcomponent of a CMTS that provides data forwarding services to a set of CMs over a set of downstream and upstream channels
MAP	Upstream bandwidth allocation
MDD	MAC domain descriptor
medium-tail programming	A set (of non-determinant size) of the video program streams that are neither short-tail or long-tail
MLD	Multicast listener discovery
MPEG	Moving picture experts group
NORM	NACK-oriented reliable multicast
NRL	Naval research laboratory
OFDM	Orthogonal frequency division multiplexing
OFDMA	Orthogonal frequency division multiplexing with multiple access
PCMM	PacketCable Multimedia
PHY	Physical layer

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

PLC	PHY link channel of an OFDM downstream channel
PROTEAN	Protocol engineering advanced networking
QAM	Quadrature amplitude modulation
QoS	Quality of service
RCC	Receive channel configuration
RF	Radio frequency
SA	Security association
SAP	Service access point
SC-QAM	Single-channel QAM; a term coined during the DOCSIS 3.1 specification effort to denote pre-3.1 DOCSIS QAM channels and differentiation them from OFDM/OFDMA channels
SCTE	Society of Cable Telecommunications Engineers
SDV	Switched digital video
sequenced	A stream in which packets are guaranteed to not arrive out of order (within some sequencing interval)
SG	Service group
short-tail programming	A set (of non-determinate size) of the most popular video program streams; these program streams are more likely to be broadcast to all user groups
SID	Service identifier
STB	Set-top box
TCC	Transmit channel configuration
TCP	Transmission control protocol
TLV	Type, length value; a data structure used to pass operational parameters in the DOCSIS protocol
trick modes	A set of features commonly found on a home video player that allow a viewer to control a video stream; these features include but are not limited to play, pause, rewind, skip, fast-forward, stop, etc.
TV	Television
UCD	Upstream channel descriptor
UDP	User datagram protocol
unbonded service flow	A service flow in which the data is transported on only a single channel associated with a MAC Domain which reaches a common set of cable modems
unsequenced	A stream in which packets may arrive in a different order than they were originally transmitted
US	Upstream
VoIP	Voice over IP

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

RELATED READINGS

- **Migrating to IP Video over DOCSIS – The Challenges, Requirements and Technologies at the Forefront of Change** – For many service providers, IP video is on the near horizon. But getting there requires careful planning. In this white paper, Cornel Ciocirlan outlines the need for a phased spectrum evolution plan and discusses several new technologies that can help service providers smooth this transition.

- **Methodologies for QoE Monitoring of IP Video Services** – this paper discusses the differences between QoE and QoS and between QoE and video quality and then compares different methodologies for video quality and QoE monitoring. It also includes a review of alternatives for embedding QoE probes in the end-to-end IP Video architecture and their ability to collect true and effective QoE information.

- **Smart ABR: The Future of Managed IP Video Services** – This paper compares in detail the SABR system is to traditional unmanaged ABR delivery as well as a system with enhanced CMTS QoS. With SABR, operators can significantly increase their IP Video capacity while gracefully handling congestion and providing an improved user experience.

MEET OUR EXPERT: William T. Hanks

Bill Hanks serves as Director of Systems Engineering under the ARRIS office of CTO-Networks with a focus on broadband application solutions. Most recently, Bill has been one of the ARRIS representatives to the CableLabs DOCSIS 3.1 specification effort and has previously been on the DOCSIS 3.0, DOCSIS M-CMTS, DOCSIS DSG, PacketCable Multimedia, and PacketCable 1.5 specification teams. Bill has been with ARRIS since the Cadant acquisition circa 2002 and was with both Cadant and Motorola during the previous 15 years. Bill has seven issued patents with several more pending and is a graduate of the University of Illinois at Urbana-Champaign.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

REFERENCES

(1) Kooij, Robert et.al.; “Perceived Quality of Channel Zapping”; Proceedings of the Fifth Lasted International Conference; Communication Systems and Networks; August 28-30, 2006. https://www.nas.ewi.tudelft.nl/people/Rob/telecom/iasted.pdf

(2) B. Dekeris, L. Narbutaite; “IP video Channel Zap Time Analysis”; Electronics and Electrical Engineering. – Kaunas: Technologija; 2010. – No. 10(106). – P. 117–120. http://www.ee.ktu.lt/journal/2010/10/26__ISSN_1392-1215_IP video%20Channel%20Zap%20Time%20Analysis.pdf

(3) “Video IP Multicast IP Multicast Adaptive Bit Rate Architecture Technical Report“; OC-TR-IP-MULTI-ARCH-V01-141112; http://www.cablelabs.com/wp-content/uploads/specdocs/OC-TR-IP-MULTI-ARCH-V01-1411121.pdf

(4) White, Gerry; “Top Ten Reasons Cable Needs Adaptive Bit Rate Streaming”; ARRIS blog; June 19, 2012; http://www.arriseverywhere.com/2012/06/top-ten-reasons-cable-needs-adaptive-bit-rate-streaming/

(5) Allen, Jim et.al.; “ARRIS WhitePaper: Achieving High Channel Utilization with DOCSIS Video-over-IP”; October 26, 2010.

(6) B.Adamson, C. Bormann, M. Handley, J. Macker, “NACK-Oriented Reliable Multicast (NORM) Transport Protocol”, IETF RFC 5740, November 2009.

(7) “Data-Over-Cable Service Interface Specifications: DOCSIS 1.0 Radio Frequency Interface (RFI)”; Society of Cable Telecommunications Engineers: Data Standards Subcommittee; American National Standard SCTE 22-1 2012; http://www.scte.org/documents/pdf/Standards/ANSI_SCTE%2022-1%202012.pdf

(8) “Data-Over-Cable Service Interface Specifications: DOCSIS 3.0 MAC and Upper Layer Protocols Interface Specification”; CM-SP-MULPIv3.1-I08-151210; Cable Television Laboratories, Inc.; 2006-2015; http://www.cablelabs.com/wp-content/uploads/specdocs/CM-SP-MULPIv3.1-I08-151210.pdf

(9) “Data-Over-Cable Service Interface Specifications: DOCSIS 3.0 MAC and Upper Layer Protocols Interface Specification”; CM-SP-MULPIv3.1-I08-151210; Cable Television Laboratories, Inc.; 2006-2015; http://www.cablelabs.com/wp-content/uploads/specdocs/CM-SP-MULPIv3.1-I08-151210.pdf

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

APPENDIX: SAMPLE IP VIDEO CONFIGURATIONS

In each of the following examples, the DOCSIS MAC-Domain Downstream Service Group (MD-DS-SG) consists of 32 SC-QAM downstream channels and one OFDM downstream channel. For purposes of drawing ease, the channels are all drawn as being adjacent but this is not necessary so long as the DS channels that are assigned to each CM can be simultaneously received (fit within the capture bandwidth window) by that CM. Channel identifiers in these examples are considered to be logical in nature and the actual configured DOCSIS channel identifiers may be different from what is shown here.

All upstream channel supervision that is used by pre-3.0 DOCSIS CMs is sent on the primary-capable SC-QAM downstream channels. Supervision for upstream channels that are used only by multi-channel DOCSIS CMs may be carried on the non-primary-capable SC-QAM downstream channels. Supervision for upstream channels which are used only by DOCSIS 3.1 CMs may be carried on the OFDM DS channel.

The DS channels are configured with a McastIP video attribute (which can be designated by the MSO to be any one of the user-defined bits of the resource attribute mask – see section: Attribute Mask Bit Definition) of value 0 or 1. Downstream channels with a McastIP video value of 1 are for all services (perhaps excluding VOIP) including multicast IP video. Downstream channels with McastIP video value of 0 are for services other than multicast IP video. If telephony flows are to be steered to non-multicast-IP video channels then another attribute mask bit may be used for this purpose but this topic is not in scope for this paper.

Dynamic downstream channel assignment and dynamic downstream channel bonding is not used in these examples. Each CM is assigned the largest bonding group that it can support plus any smaller bonding groups that the MSO wishes to assign for the services to be provided. In many cases there are multiple bonding groups of the same size depicted. The CMTS/CCAP should be configured to load balance across bonding groups of the same size.

Multicast traffic is steered to DS channels using a multicast service class with a required attribute mask that includes both the bonding and McastIP video attributes set. Other services on bonded service flows (HSD, unicast IP video, signaling for services, etc.) can also be carried on the multicast DS channels but will be naturally “pushed” to non-multicast-IP video channels by the CMTS/CCAP bonded downstream packet scheduler as the higher priority multicast IP video traffic increases. The multicast group service flows must have a higher DOCSIS priority (perhaps 3) than other traffic. Unicast signaling flows and unicast ABR data flows will get the next highest priority (perhaps 2). Normal high speed data flows should get a lower DOCSIS priority (perhaps 1); thus supporting a

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

bandwidth penalty box (with priority 0) for policy enforcement purposes by MSOs. These priorities will not affect VOIP flows at their default DOCSIS priority setting of 5, thus giving low-bandwidth but latency-sensitive VOIP flows priority on the non-multicast-IP video channels.

Each of these configurations can be used as a baseline and tweaked to suit specific system needs.

Minimum 32 DS Channel CMs/HGs for IP Video

CM types participating in all services (including IP video):

- DOCSIS 3.0
  - (32 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways
- DOCSIS 3.1
  - (32 SC-QAM + 2 OFDM DS and 8 SC-QAM + 2 OFDMA US) CMs & Home Gateways

CM types not participating in IP video:

- All other CMs & Set-Top Boxes

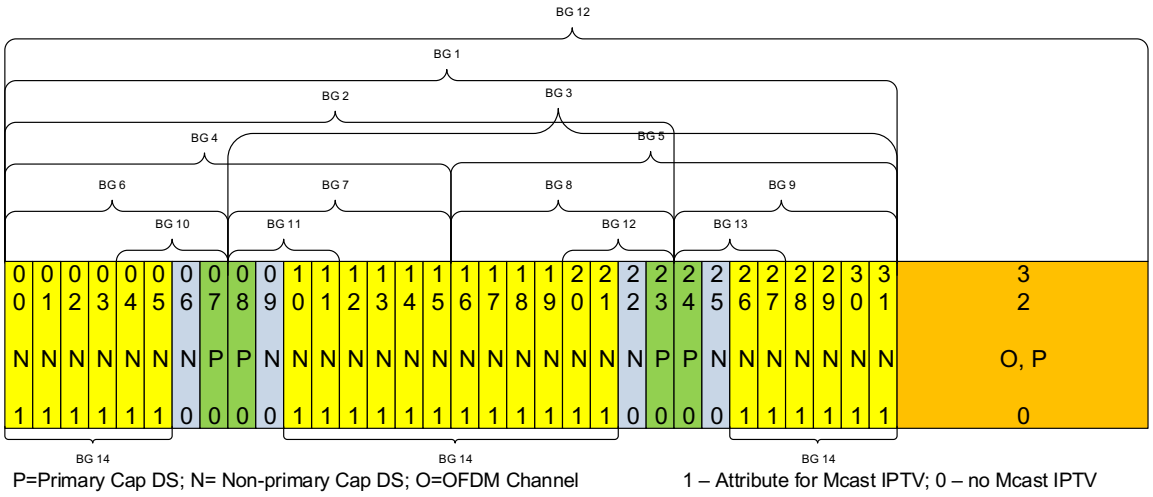

Figure 8: IP video configuration with minimum 32 DS channel CMs/HGs for IP video

Primary-capable downstream channel-ids in this example are 7, 8, 23, and 24 (more can be configured if the pre-3.0 CM population percentage is high). All DSG tunnels (if any) are configured on the SC-QAM primary-capable downstream channels. Multicast traffic is carried on multicast bonding group BG 14 (DS channels 0-5, 10-21, and 26-31).

This arrangement works well if the MSO desires to multicast many television program streams at the same time. There is ample capacity for many multicast streams in this case. Meanwhile, all CMs which are not participating in IP video may use the unused capacity of the channels without impacting the IP video service.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

Minimum 24 DS Channel CMs/HGs for IP Video

CM types participating in all services (including IP video):

- DOCSIS 3.0
  - (32 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways
  - (24 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways
- DOCSIS 3.1
  - (32 SC-QAM + 2 OFDM DS and 8 SC-QAM + 2 OFDMA US) CMs & Home Gateways

CM types not participating in IP video:

- All other CMs & Set-Top Boxes

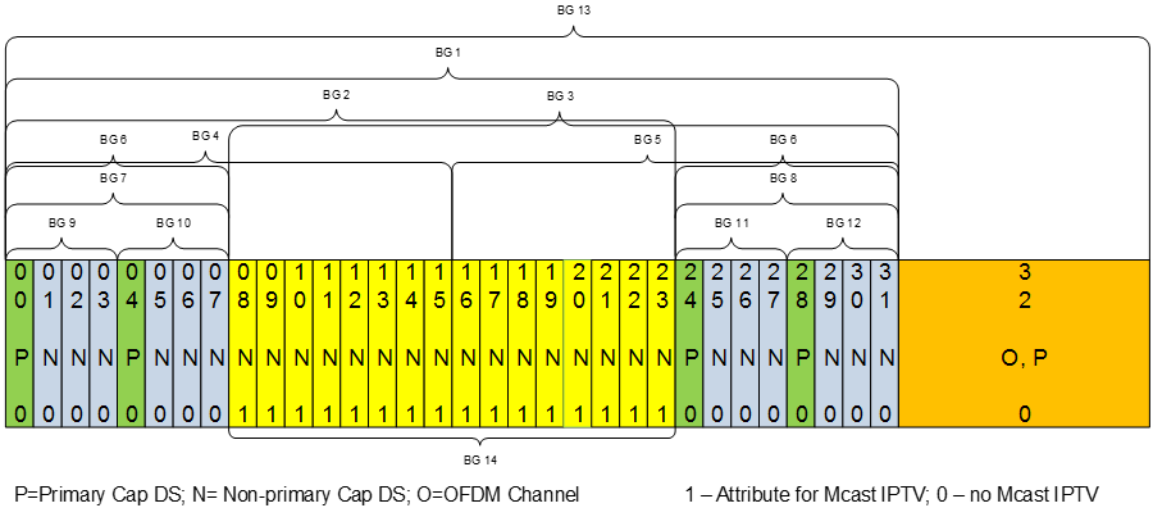

Figure 9: IP video configuration with minimum 24 DS channel CMs/HGs for IP video

Primary-capable downstream channel-ids in this example are 0, 4, 24, and 28 (more can be configured if pre-3.0 CM population percentage is high). All DSG tunnels (if any) are configured on the primary-capable downstream channels. Multicast traffic is carried on DS channels 8-23.

This arrangement works well if the MSO desires to multicast many television program streams at the same time. There is ample capacity for many streams in this case. Meanwhile, all CMs which are not participating in IP video may use the unused capacity of the channels without impacting the IP video service.

Minimum 16 DS Channel CMs/HGs for IP Video

CM types participating in all services (including IP video):

- DOCSIS 3.0

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

- (32 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways
- (24 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways
- (16 SC-QAM DS and 4 SC-QAM US) CMs & Home Gateways

- DOCSIS 3.1
  - (32 SC-QAM + 2 OFDM DS and 8 SC-QAM + 2 OFDMA US) CMs & Home Gateways

CM types not participating in IP video:

- All other CMs & Set-Top Boxes

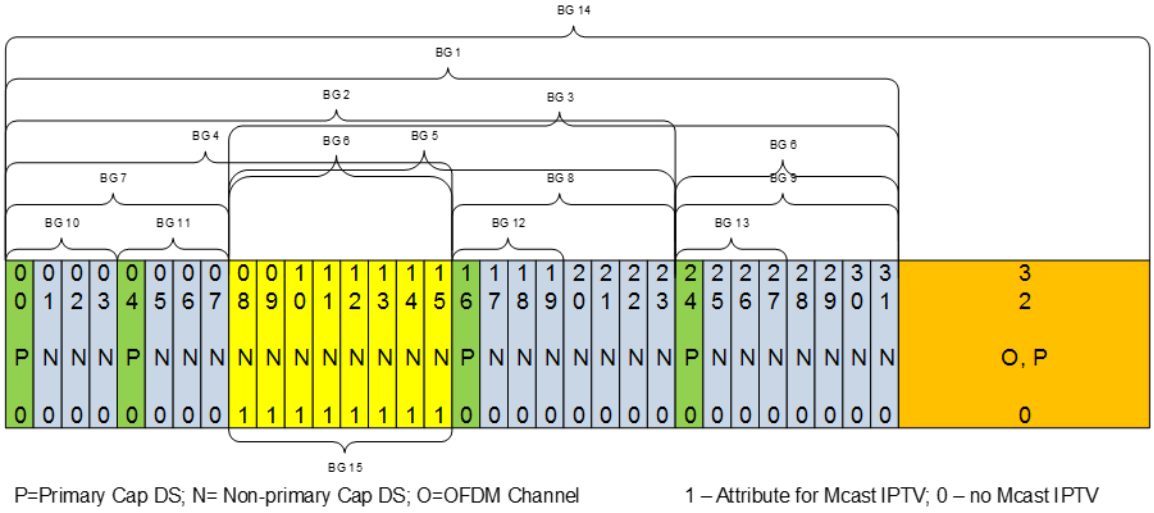

Figure 10: IP video configuration with minimum 16 DS channel CMs/HGs for IP video

Primary-capable downstream channel-ids in this example are 0, 4, 16, and 24 (more can be configured if pre-3.0 CM population percentage is high). All DSG tunnels (if any) are configured on the primary-capable downstream channels. Multicast traffic is carried on DS channels 8-15.

The use of 16 DS CMs for IP video forces a more constrictive arrangement than is possible when only larger CMs are used. In order to leave DS channel space for other services on a 16 DS channel CM, the multicast bonding group – BG 15 is only half what the example for 24 channel and up CMs could use. This may or may not be sufficient multicast bandwidth to meet the IP video goals of the MSO. Meanwhile, all CMs which are not participating in IP video may use the unused capacity of the channels without impacting the IP video service.

Minimum 8 DS Channel CMs/HGs for IP Video

CM types participating in all services (including IP video):

- DOCSIS 3.0
  - (32 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.

- (24 SC-QAM DS and 8 SC-QAM US) CMs & Home Gateways
- (16 SC-QAM DS and 4 SC-QAM US) CMs & Home Gateways
- (8 SC-QAM DS and 4 SC-QAM US) CMs & Home Gateways

- DOCSIS 3.1
  - (32 SC-QAM + 2 OFDM DS and 8 SC-QAM + 2 OFDMA US) CMs & Home Gateways

CM types not participating in IP video:

- All other CMs & Set-Top Boxes

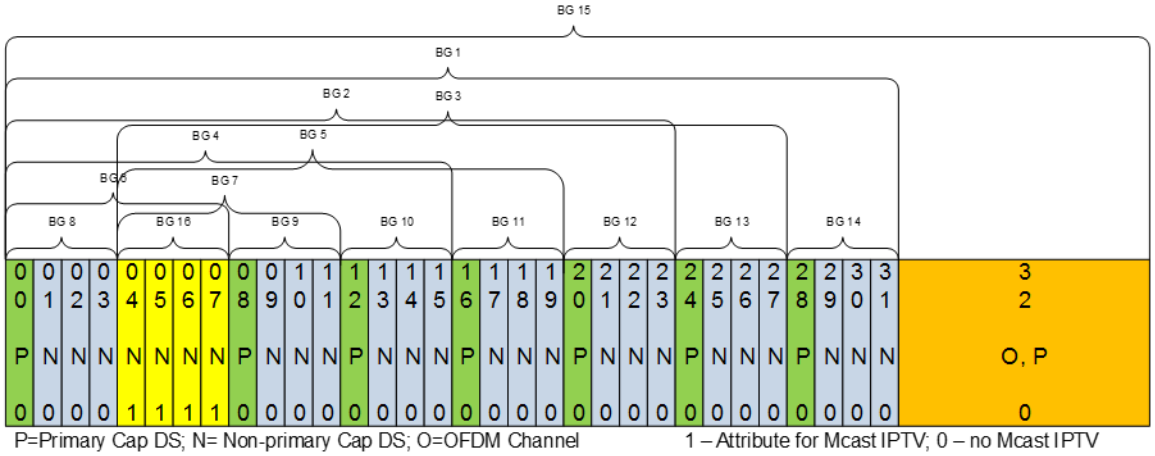

Figure 11: IP video configuration with minimum 8 DS channel CMs/HGs for IP video

Primary-capable downstream channel-ids in this example are 0, 8, 12, 16, 20, 24, and 28 (more can be configured if pre-3.0 CM population percentage is high). All DSG tunnels (if any) are configured on the primary-capable downstream channels. Multicast traffic is carried on DS channels 4-7.

This arrangement works well if the number of 8x4 DOCSIS 3.0 CMs that participate in IP video Multicast groups is relatively low as the 8x4 CMs will be constrained to use the four multicast channels plus either the four HSD channels immediately below them in the spectrum or the four HSD channels immediately above them in the spectrum. This constraint is due to the capture bandwidth constraint of the smaller DOCSIS 3.0 CMs.

Copyright 2016 – ARRIS Enterprises LLC. All rights reserved.