A Converged Content Delivery Platform for IP and QAM Video

James Barkley, Weidong Mao
Comcast Cable

Abstract

HTTP based Adaptive Bit Rate (ABR) video delivery to IP enabled CPEs via Content Delivery Network (CDN) for both primary and secondary screens (IP Set-Top Box, mobile devices, etc.) is becoming a common feature of service providers. However, the large installed base of QAM STBs will require QAM based linear and Video On Demand (VOD) delivery for the foreseeable future. Traditional approaches to QAM linear delivery include statistical multiplexing and Multiple Program Transport Stream (MPTS) multicasts. A dedicated content library and purposely-built proprietary CDN have been used for QAM VOD delivery. As video migrates to H.264, separate encoding systems are dedicated to produce non-ABR assets (for QAM delivery), sharing little with the production of ABR assets (for IP delivery). Other video features such as advertising insertion, content encryption, data collection and reporting are also enabled through separate silos for IP and QAM delivery. This will not only increase the overall infrastructure cost, but also require separate operational resources.

A converged CDN platform along with common encoding, packaging and content origin are proposed to enable a single common content delivery platform for both IP and QAM video delivery, supporting all use cases, including linear and VOD. The proposed platform uses an MPEG DASH-TS based Common Intermediate Format (CIF) for packaging, ingest/storage, and delivery. It leverages the common ABR video and audio encoding system for H.264 video and various audio codecs. Legacy QAM video encoding is transitioned from MPTS, statistically multiplexing and multicasting to Single Program Transport Stream (SPTS) CIF segments with HTTP delivery through CDN. A CIF Media Presentation Document (MPD) is developed to describe all the codec and bitrate / resolution profiles for the CIF segments. Common encryption on CIF can be used to enable content security at rest in the storage and through the CDN delivery path. Common reporting and data collection schemes can be applied to the infrastructure. A common advertising system can be enabled for both IP and QAM video. At the edge of the network, the CIF MPD and segments can be converted to the appropriate MPEG-2 MPTS format for QAM and player ABR format (HLS, DASH ISO-BMFF, etc.) in conjunction with Digital Rights Management (DRM) for IP.

This converged content delivery infrastructure allows operators to increase the number of IP video clients while still supporting existing QAM STBs. This approach will result in lower overall capital and operational expenses. The approach will also enable seamless incorporation of the latest generation of ABR and codec technologies into the content delivery platform.

INTRODUCTION

With a large amount of QAM based digital set-top boxes deployed in the field, cable operators have been delivering linear video and services over backbone and regional IP networks through Gigabit Ethernet interface edge QAMs.

Traditional linear video delivery has been using techniques such as MPEG-2 Transport Stream over IP multicast, statistical multiplexing, and local Ad splicing. In particular, dedicated MPEG-2 encoders encode linear channels as SPTSs. Statistical multiplexers combine multiple SPTSs into an MPTS by taking advantages of dynamic allocation of bits to each SPTS, depending on the variability of the content. The statistical multiplexing technique enables bandwidth efficiency by packing more SPTSs into a single 6MHz QAM channel than would be supported without it. Local Ad insertion can be performed on each resulting MPTS by using a dedicated Ad Splicer with Ad Servers before sending to the edge QAMs.

Traditional video on demand delivery has been using techniques such as MPEG-2 SPTS over UDP/IP unicast. Typically, each VOD asset is encoded as SPTS using CableLabs ADI (Asset Distribution Interface) and CEP (Content Encoding

Profile) specifications. A purposely-built proprietary CDN is then used to deliver VOD assets from the content library to the headend locations, where streaming servers are typically located.

Given the large installed base of H.264-capable HD QAM digital set-top boxes in the field, significant bandwidth and storage savings can be achieved by transitioning linear and VOD delivery from MPEG-2 to H.264. On the other hand, IP video for primary and secondary screen is becoming a reality. The ABR-encoded H.264 streams are delivered through HTTP based CDNs and various IP enabled access networks to the IP players. This includes all the linear, VOD, and cloud DVR services.

Due to the large installed base of QAM STBs, migration to IP STBs and devices is slow. Having completely separate platform/component infrastructures dedicated to QAM and IP (encoder, packager, origin, delivery, security, ad insertion, and operation management tools) is not cost effective.

This paper describes a converged delivery platform for both IP and QAM video, supporting all use cases, including linear and VOD.

- Common encoding: It leverages the common ABR video and audio encoding system for H.264 video and various audio codecs.
- Common packaging: The proposed platform uses the MPEG DASH-TS based CIF.
- Common CDN delivery: In addition to IP video ABR HTTP delivery via CDN, the legacy QAM linear video delivery is also transitioned from MPTS, statistically multiplexing and multicasting to the SPTS CIF segment with HTTP delivery through the same CDN. The CDN also delivers CIF segments for QAM VOD.
- Common encryption: Common encryption on CIF can be used to enable content security at rest in storage and through the CDN delivery path.
- Common advertising systems can be enabled for both IP and QAM video. At the edge of the network, the CIF MPD and segments will be converted to the appropriate MPEG-2 MPTS format for QAM and player ABR format (HLS, DASH ISO-BMFF, etc.), in addition to with DRM for IP.
- Common reporting and data collection schemes can be applied to the infrastructure.

This converged content delivery platform allows operators to migrate to an all IP infrastructure while still serving millions of QAM STBs in the field. This approach will result in lower overall capital and operational expenses. The approach will also enable efficiency in storage and network bandwidth resources in the evolution of content delivery platforms.

COMMON CDN DELIVERY

When QAM VOD was first introduced, Streaming Servers were deployed regionally. Each region was loaded individually with the titles in the library. As the number of titles and usage grew, CDNs were introduced allowing the use of central libraries to hold the library. These CDNs were proprietary systems built around the specific characteristics necessary to provide assets to the Streaming Servers. The QAM VOD delivery is shown in **Figure 1**.

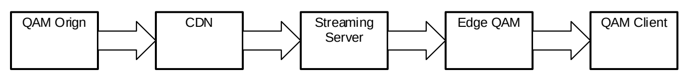

Figure 1 - QAM VOD Delivery

When IP VOD was introduced, it used a central library. However, the new IP library differed greatly from the QAM library because the QAM library contained MPEG-2 titles and the IP library contained H.264 titles. Furthermore, the CDN was different between QAM and IP. For IP VOD a title is transferred as a series of small files. QAM VOD views a title as a single, much larger file and is very sensitive to latency in the delivery. The IP VOD delivery is shown in **Figure 2**.

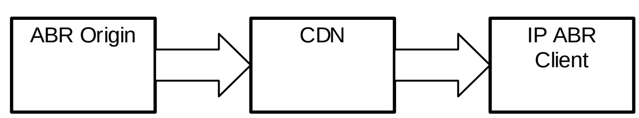

Figure 2 - IP VOD Delivery

In their 2015 paper, "Leveraging IP Video Delivery for Linear QAM TV" Weidong Mao and David Brouda introduced a HTTP ABR to UDP Gateway (QAM Linear Gateway) as an IP ABR client [1]. The gateway retrieves designated ABR profile(s) from a CDN, performs the necessary HTTP ad insertion and streams the result to the edge QAM over UDP/IP as a transport stream. Using the gateway allows a CDN to be shared in the distribution of linear content with both IP and QAM clients. A converged linear architecture is shown in **Figure 3**.

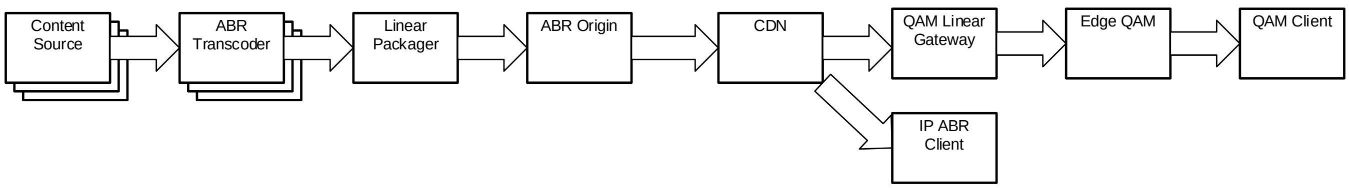

Figure 3 - Converged CDN for QAM and IP Linear

Taking the general concept of an HTTP ABR to UDP Gateway, we create a QAM VOD Gateway that converts from the large file based requests that are made from existing Streaming Servers into the ABR segment requests of the CDN and origin that are deployed for IP VOD. The QAM VOD Gateway satisfies the latency requirements of the Streaming Servers. This makes any H.264 title available to all QAM clients that are capable of decoding H.264.

All ABR delivery begins with reading a manifest that identifies the profiles and the files (segments) that make up the profiles. Manifest are one of two types dynamic, or static. The QAM Linear Gateway works with dynamic manifests, in which segments are constantly being added and removed. The QAM VOD Gateway, however, uses static manifests. Working with static manifests is, at the same time, easier and harder. A static manifest needs only to be read once, rather than at each update. Once read, it is only the most recent segment that requires processing for a dynamic manifest. Conversely, for a static manifest all segments require processing.

In contrast to QAM Linear, which traditionally supports statistical multiplexing, QAM VOD is constant bit rate delivery. This means that unlike other ABR clients a QAM VOD Gateway will select a single bit rate and never change. A gateway could change bit rate if there are issues with the network through the CDN back to the origin, however, there is no way to reclaim the unused QAM capacity.

Caching has long been use used in QAM VOD Streaming Servers to manage the impact of network through the CDN back to the origin. Having a gateway change the bit rate would not combine well with this caching. The change of bit rate would affect multiple sessions (until the stream ages out of the cache). Transition from caching in the Streaming Server to adjusting bit rate in the gateway would significant increase the network traffic into the Streaming Server.

A converged VOD architecture for H.264 delivery utilizing the QAM VOD Gateway is shown in **Figure 4**.

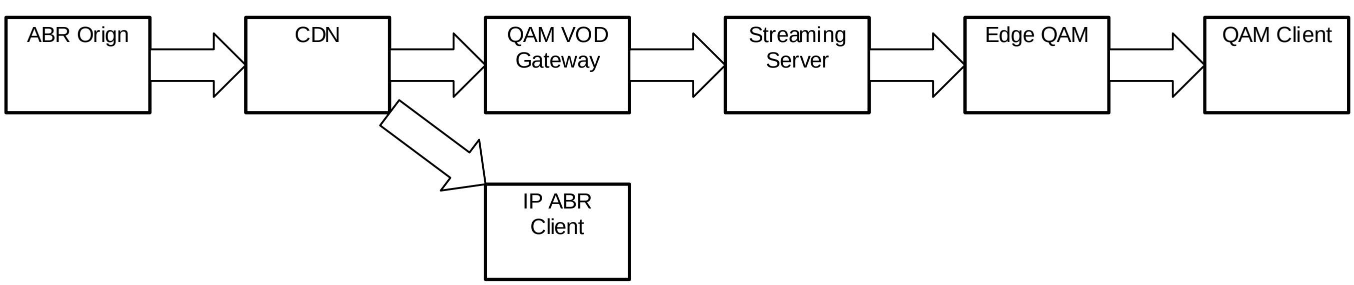

Figure 4 – Converged CDN for QAM and IP VOD

This converged CDN provides H.264 delivery to QAM VOD clients without incurring additional origin costs by reusing ABR origin. However, this does not reduce the number of CDNs, as there is still a separate CDN for MPEG-2. Additional savings is possible by transition to a single CDN for all VOD. This does require that the existing MPEG-2 library be transformed from large files to ABR files. It is not necessary to create multiple different bit rate profiles since the network has already been scaled to support the single MPEG-2 bit rate.

Thus far we have discussed linear, or VOD, separately. However, when looking at the CDN, the capabilities are the same, whether the files being delivered are linear or VOD. Files used for linear can have a much shorter relevance, but this difference is straightforward to express within the HTTP standard. Rather than having two CDNs, (one for VOD and one for linear), additional savings are possible by sharing the CDN between linear and VOD.

COMMON ENCODING

The introduction of the gateways allows a single encoding platform to be used without consideration of the IP or QAM delivery format. An asset (linear channel or VOD title) is encoded and packaged into ABR segments. IP clients, using the manifest, will select the specific encodings that they are capable of decoding. QAM clients will have access to the same assets via gateways.

Common video encoding is H.264. Some QAM devices are only capable of supporting MPEG-2 video encoding. However, that encoding will not be destined for IP devices. Similarly, there are experiments in delivering H.265 video encodings to IP devices, but little interest in creating QAM devices capable of such a video decode. Since we are using ABR, separate segments could be selectively produced for those assets that require alternative video encodings.

It may not be necessary to have multiple video encodings on the origin. Gateways already are transforming from ABR to MPEG-2 transport streams. Depending on the tradeoff of origin storage, CDN cache and gateway computing performance, it is possible that the gateway may transcode from H.264 to MPEG-2 on the fly.

COMMON PACKAGING

DASH-TS based Common Intermediate Format is introduced as a common packaging format for both linear and VOD content at the origin and CDN. This introduces a CIF conversion between the CDN and IP client, but allows for more efficient use of the origin and CDN cache. The savings of resources in the origin and CDN is partially offset by the resources at the edge, converting on-demand to the IP streaming format. However, having the conversion at the edge and not the origin further reduces the burden when multiple streaming formats are to be supported.

In order to support ABR delivery multiple encodings are produced at differing bit rates. Video not only requires more bits to encode than audio, but also is easier to change the bit rate through changes in resolution. The result is that only one bit rate will be created for audio. Rather than multiplex the audio into all of the video profiles the audio is packaged separately from video.

Although QAM has traditionally relied on a single audio, encoding IP has not. It is common to have some devices capable of decoding Dolby Digital Plus (DD+ also known as EAC-3), others limited to Dolby Digital (DD also known as AC-3 and still others limited to AAC. Creating unique segments for each of these audio encodings and identifying them in the manifest allows the client to select the encoding they are capable of supporting.

Separate audio segments also simplify accessibility support within VOD. QAM settled on a limit of two audios minimizing the fixed bandwidth cost for audio. As a result, QAM VOD selects either a Secondary Audio Program (SAP) or Descriptive Video Service (DVS) for the second audio – never both. With separate audio segments, a title can be created with any number of audio tracks. The gateway will select the two that are to be delivered.

COMMON ENCYPTION

For QAM delivery, encryption is applied at the QAM device itself. Traditionally the distribution network has not required encryption because it has been a closed system. Only specific, qualified devices are allowed to access the distribution network. Careful monitoring is required to ensure that no unqualified device is granted access and periodic audits ensure that qualified devices are not compromised.

For IP delivery, encryption is tied with the DRM system associated with adaptive streaming format.

Thus far, it has been stated that the gateways function as IP clients for the converged system. Therefore, common encryption can be achieved by when gateways utilize the same streaming format as all other IP clients. Unfortunately, not only is there no single streaming format, innovation changes the existing protocols while introducing new ones.

Instead, the aforementioned CIF is used. This format supports encryption of the data as it is loaded onto the origin. Data stays encrypted as it passes through the CDN.

For QAM, the data is decrypted in the gateway. Depending on the exact security desired, data can be encrypted leaving the gateway, or it can be encrypted at the QAM.

For IP, a device is necessary to map from CIF to the client's streaming format. If this were done prior to loading on the origin, there would be no sharing of storage, or cache, between QAM and IP delivery. In addition, new copies would be required for every new protocol. If CIF conversion were performed between the origin and the CDN, storage would be shared, but not cache. In order to share storage and cache, the conversion must take place between the CDN and the IP client.

COMMON ADVERTISING SYSTEMS

It is reasonable to start with the understanding that there are two advertising systems in content delivery; one for linear and a second for VOD. Linear delivery, both IP and QAM, requires selection of the advertisements in real time. In contrast, QAM VOD selects advertisements prior to starting any delivery. IP VOD can leverage either technique by selecting advertisements either prior to starting any delivery, or in line with the delivery of segments.

VOD advertisement selection is personalized to both the viewer and the title. The effectiveness of advertisement campaigns can be further increased by considering all advertisements presented in a single VOD session.

Linear advertisement selection is based on station, rough-time boundaries and broadcast area. Unicasting reduces the broadcast area to a single client, allowing personalization similar to that achieved in VOD. This degree of personalization can even be achieved during broadcasting when the client orchestrates the advertisement selection.

Given the progression toward personalized advertisement selection, the differences between linear and VOD advertisement selection become a lot less significant. Selecting an advertisement requires three items: identification of the advertisement opportunity; identification of the client; and identification of the entertainment.

Previous sections discussed common encoding and common packaging. The common encoding includes conditioning content for clean transitions into and out of advertisements. The manifest generated by the common packaging includes signaling of the advertisement opportunities. The packaging also includes the creation of package boundaries at these transition points. Any advertising system only has to work with this one format for advertisement opportunities.

When the client makes the selection of the advertisement, identification of the client is straightforward. For QAM Linear delivery, there exist techniques that allow the client to select the advertisements, but it is more common for the selection to be made during the broadcast. In this case, the identification is via the broadcast ad zone and all clients within that ad zone receive the same advertisement.

Linear and VOD assets have different identification strategies (station versus title). The manifest generated by the common packaging manages both strategies in a single structure.

Having aligned the inputs for selecting linear and VOD advertisements, a common advertisement system is possible. The common advertising system is called individually for each advertisement opportunity. Title or program context is provided as another input, tracking with any previous decisions that have been made.

For QAM VOD, the smallest change is to continue to select all advertisements prior to session start. Moving the selection of advertisements to the QAM VOD Gateway is possible, but has some expense; either removing caching from the Streaming Server, or additional signaling between the Streaming Server and the gateway. In addition, any support of fast-forward or fast-rewind in VOD presents challenges, predicting the timing of a request for an advertising selection to ensure all opportunities are resolved before the data for the advertisement needs to be fetched.

COMMON REPORTING AND DATA COLLECTION

The first content delivery by most service providers was QAM linear. When QAM VOD was added, the use case was sufficiently different there was no serious attempt to share any structure. Eventually IP delivery was added. However, the status quo had been formed, and there was little shared between IP Linear and IP VOD.

This has tended to result in large *diverged* content delivery systems as shown in **Figure 5**

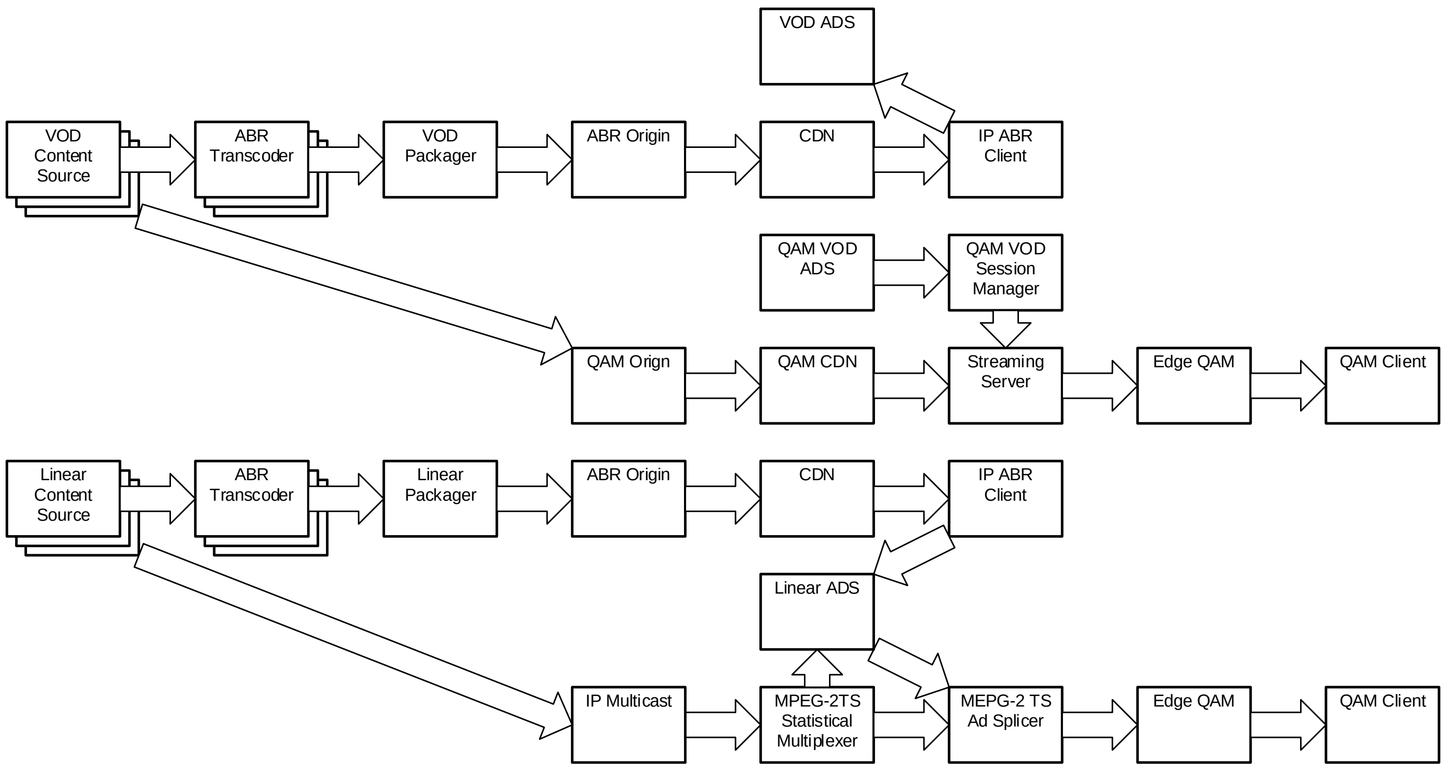

Figure 5 - *Diverged* Content Delivery

The large number of different components involved complicates Reporting and data collection.

Legacy (non-ABR) delivery makes significant use of the UDP protocol paced at the source to match the rate of video decoding. ABR delivery makes significant use of the TCP protocol. TCP allows the delivery rate to vary for short time intervals while trending to the video decoding rate over the long term. These differences further discourage any desire to converge reporting and data collection.

However, as we have shown it is possible to reduce the number of unique components in a converged content delivery system as shown in **Figure 6**.

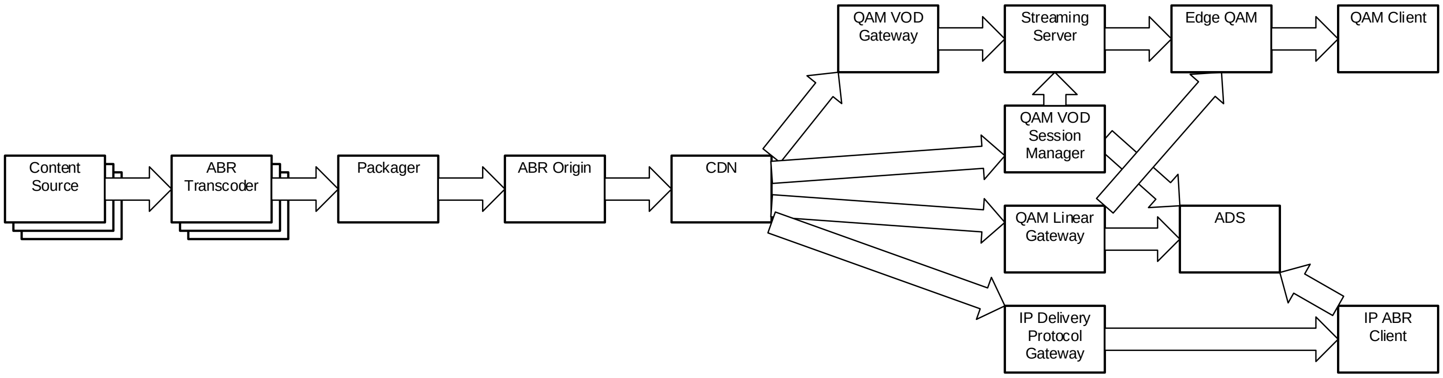

Figure 6 - Converged Content Delivery

With converged content delivery, Linear and VOD can share contents from the content source through the CDN. There are different performance characteristics for Linear and VOD, and there may be partitioning of the components to accommodate these differences. However, the default as common components is to have common reporting and data collection.

Although the components diverge after the CDN, the functionality of retrieving the manifest and segments in addition to selecting the advertisements is still common. This will continue to drive common reporting and data collection.

CONCLUSION

Content delivery has developed in a slow fashion. Different services and delivery formats (e.g., linear or VOD, QAM or IP) have been added over time. Taking a pause from the work of delivering the content allows the time to reflect and to begin to see the similarities are greater than the differences.

Although this discussion has been primarily focused on MPEG-4 delivery, as it is common to both IP and QAM content delivery, this can also be used for MPEG-2 delivery. Obviously, with MPEG-2 there is no sharing of content between QAM and IP. However, as non-MPEG-2 content becomes predominant, operations teams will be less familiar with the structure and issues of non-ABR delivery. Transitioning MPEG-2 into ABR will reduce the operations cost of continuing to carry MPEG-2.

This converged content delivery infrastructure allows operators to increase their IP video clients without needing to replace their QAM STBs. This approach will result in lower overall capital and operational expenses while enabling seamless incorporation of the latest generation of ABR and codec technologies into the content delivery platform.

REFERENCES

[1] Leveraging IP ABR Video Delivery for Linear QAM TV, Weidong Mao, David Brouda, 2015 Spring Technical Forum.

[2] R. Fielding et al., "Hypertext Transfer Protocol—HTTP/1.1: Message Syntax and Routing," Request for Comments (RFC) 7230, Internet Engineering Task Force, 2014.

[3] R. Fielding et al., "Hypertext Transfer Protocol—HTTP/1.1: Semantics and Content," Request for Comments (RFC) 7231, Internet Engineering Task Force, 2014.

[4] R. Fielding et al., "Hypertext Transfer Protocol—HTTP/1.1: Range Requests," Request for Comments (RFC) 7233, Internet Engineering Task Force, 2014.

[5] W. Mao, "Building Large VOD Libraries with Next Generation On Demand Architecture," NCTA Technical Papers, 2008.

[6] W. Mao, "Key Architecture and Interface Options for IP Video Over Cable," SCTE Conference on Emerging Technologies, 2009.

[7] S. Krishnan, W. Mao, "Open Content Delivery Networks for Managed Video Delivery to Multiple Screens," SCTE Cable-Tec Expo, 2010.

[8] ISO/IEC Draft International Standard 23009-1.2, "Information Technology – Dynamic Adaptive Streaming over HTTP (DASH) Part 1: Media presentation description and segment formats," 2011.