Targeted Advertising In Linear Television

A Technical Paper Prepared for SCTE•ISBE by

Srinath V Ramaswamy
Principal Solution Architect
Comcast
1717 Arch Street, Philadelphia, PA 19103
(215) 286 5444
srinath_ramaswamy@cable.comcast.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Abstract

Advertising in video is an important business segment of the cable industry, and especially targeted advertising. In this paper, targeted advertising architectures deployed for linear television are discussed. First, we discuss the architecture successfully deployed for the legacy devices followed by the architecture for IP-based devices. In the case of the solution for legacy devices, linear video and targeted ads were delivered over QAM whereas in the case of IP-based and hybrid QAM-IP devices, linear video was delivered over QAM and targeted ads over IP. This is followed by the architecture for targeting advertising for linear television over IP. Finally, detailed are challenges faced with seamlessly splicing MPEG transport streams and their possible resolutions. It is important to note that when service providers use set-top box data to deliver relevant programming and advertising to customers, they do so without disclosing personally identifiable information about customers to programmers or advertisers. Comcast collects and uses data in accordance with our privacy policies and all relevant laws and regulations.

1. Introduction

Targeted advertising in digital cable television is an important feature required to offset growing programming costs, increase ad relevancy and reduce ad load, for a better viewer experience. This paper will focus on targeted advertising on linear television. Introduction of this feature into existing digital cable systems for both primary and second screen devices is faced with numerous challenges. Some of the challenges are the need to support legacy systems, seamless ad insertion, integrate into existing video infrastructure, evolving digital video architectures and ad decision systems.

2. Targeted Advertising Architectures

The targeted advertising architectures for various linear usage scenarios are detailed in the next few sections.

2.1. Legacy Devices on Linear QAM Channels

The architecture described in this section is for legacy devices on linear QAM channels. Legacy devices in scope for this architecture include the non-IP Arris STBs with linear network video served over a 256 QAM channel. This solution is currently live on millions of legacy devices and is based on Addressable Advertising solution from Invidi.

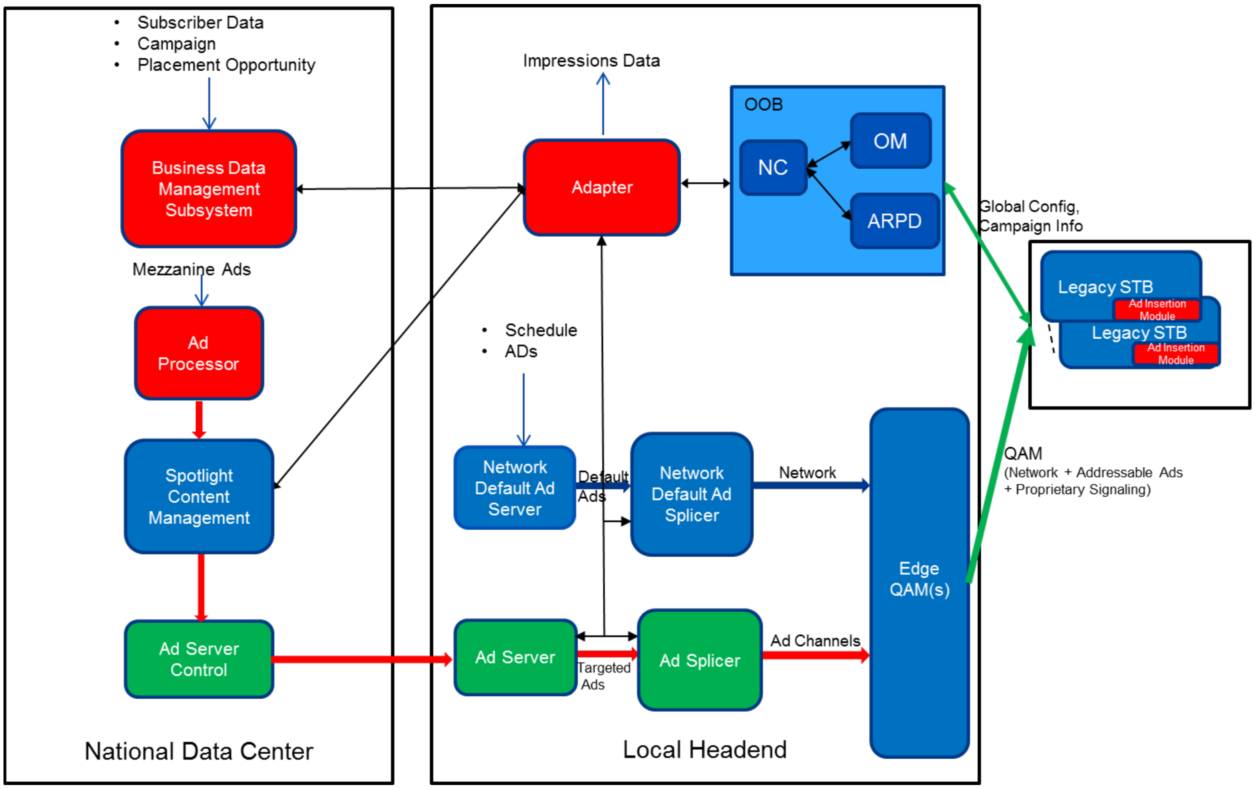

Figure 1 - Targeted Advertising Architecture for Legacy STBs (Linear QAM)

In the architecture shown in Figure 1, linear network video is over a QAM channel, and the targeted ad variants are served on dedicated QAM Ad channels that the legacy STB would tune to and play an ad that was targeted to it.

The solution components are deployed in the national data centers and local headends. Local headend components are for a single ad zone, therefore these components are mostly duplicated for each ad zone implementing this solution.

At the National Data Center we use components to manage campaign, subscriber data and placement opportunities. Addressable campaigns are created on the Invidi Business Data Management System (BDMS) system setting the ad variants to be used for each campaign and attributes of each ad variant: Maximum viewings, frequency, CPM, desired impressions, start and end date etc. Schedule files that identify breaks/spots per linear channel where targeted advertising can be exercised are also ingested into this system. Anonymized subscriber data that provides campaign-specific targeting information is provided in a proprietary format to this system.

Several targeted advertising models are available when setting up campaigns in this system. In one of the models, subscribers currently tuned to a linear channel and the ads appropriate to them are determined prior to the ad break. And the ads appropriate to that current audience are scheduled on the dedicated QAM Ad channels. In another model, a fixed set of ad variants is selected for a particular spot, dedicated to a single advertiser.

In order to switch near seamlessly between linear network content and targeted ads, the ads are processed and filler black frames are added to the beginning and end, with silent audio. The filler black frames reduce artifacts when switching from network content to ad and vice versa. This processing of ads is handled by the Ad Processor shown in Figure 1.

An Adapter sends configuration, targeting, cue, and break descriptions to STBs, either Out of Band (OOB) via the OOB components shown in Figure 1, or In band via the MPEG stream. To signal ad breaks for use by a distributor, content providers typically insert a SCTE 35 message into their streams few seconds before (minimum four seconds) the start of the ad break. This SCTE 35 message is referred to in the remaining sections of the paper. The Adapter reacts to SCTE 35 messages in the network stream, based on the schedule and campaigns applicable to that network determines which targeted ads need to played out over the dedicated ad channels for an SCTE 35-triggered break.

A proprietary component is integrated with the legacy STB firmware to support this client-side ad insertion solution. This component listens to incoming proprietary cue messages from Adapter, and based on the targeting metadata it receives the STB determines if it needs to force tune to an Ad channel streaming the targeted ad during the targeted ad spot. If none of the targeted ads streaming on the ad channels is applicable to this STB, then it does not force tune to any of the targeted ad channel.

In Figure 1, the Network Default Ad Splicer is responsible for inserting local zone-based ads during ad avails. The network default ad splicer listens to SCTE 35 messages in the stream, and upon reception of an SCTE 35 message, it communicates the cue tone information to the Adapter, via the SCTE 30 protocol. The targeted ads are played over these default ads if a STB chooses to do so. The “Ad Splicer” and “Ad Server” in this Figure 1 retrieve and inject Targeted Ads into the dedicated QAM ad channels based on instructions from the Adapter at the appropriate times. Deployment tools were developed for automated configuration of Adapter and Network Default Ad Splicer that paved way for a faster deployment of this solution.

One of the challenges of this elegant and scalable solution, deployed on millions of legacy devices, was the need to acquire a dedicated QAM channel for playing out targeted ads.

2.2. Hybrid QAM-IP based RDK devices on Linear QAM Channels

In the architecture shown in Figure 2, linear network video is received over QAM channels, but the targeted ads are retrieved over IP. This solution is currently live on several million hybrid QAM-IP based RDK devices.

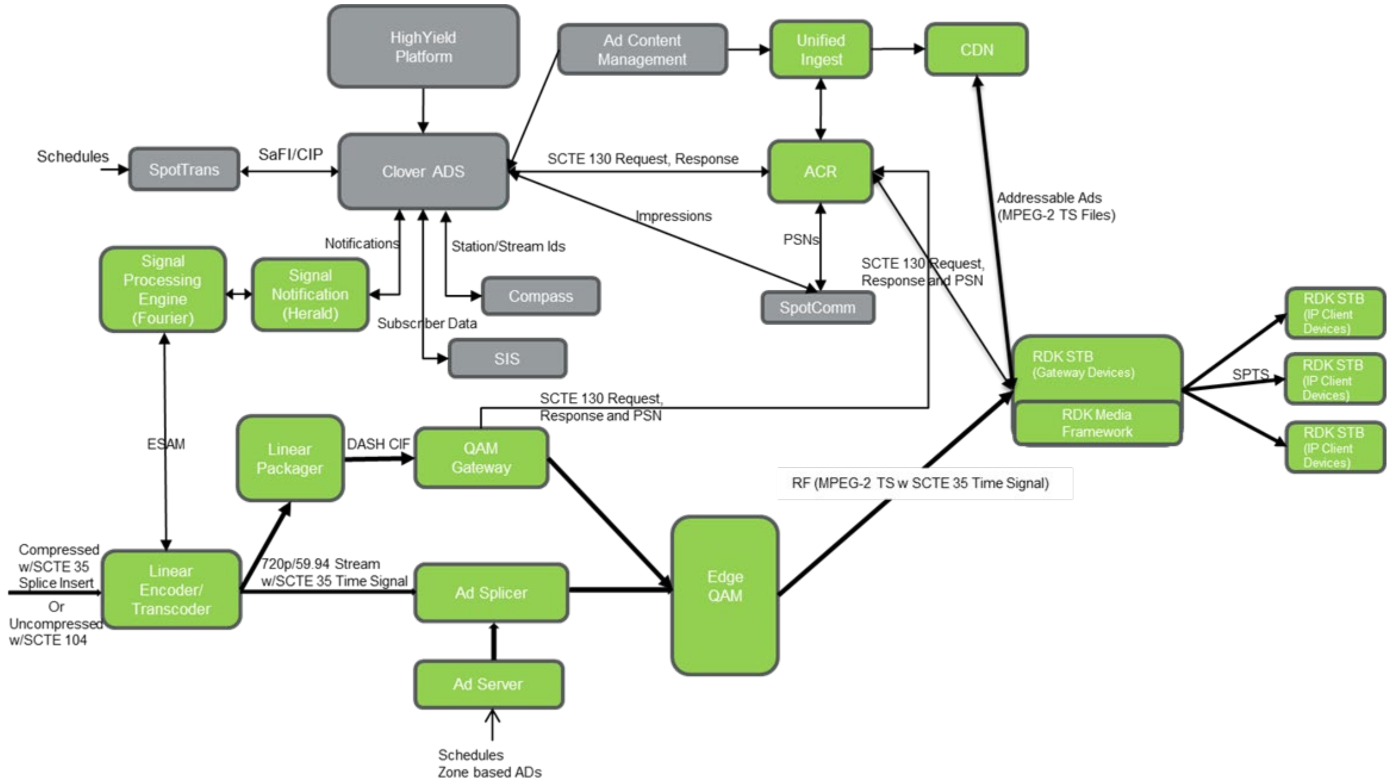

Figure 2 - Targeted Advertising Architecture for Hybrid QAM-IP based STBs (Linear QAM)

In Figure 2, a compressed MPEG video stream with SCTE 35 splice insert messages or uncompressed stream (SDI) with SCTE 104 (in the VANC) is input into Adaptive Bitrate encoders/transcoders. The encoder/transcoder transforms the SCTE 35 splice insert message for the "Out Point" in the incoming streams into an SCTE 35 time signal command with a segmentation descriptor. The transformation of this signal is done in close coordination with the Event Signaling and Management (ESAM) signal processing component Fourier. In the SCTE 35 segmentation descriptor segmentation_upid_type is set to a value of 0x9 (ADI) and the ID value set to SIGNAL:<identifier>, identifier is set to a Signal ID value generated by Fourier. The Segmentation_type_id is set to a value of 0x36 for Distributor Placement Opportunity Start message. The Fourier component also requests injections of SCTE 35 "In Point" time signal messages occurring at suitable times. The Herald component notifies Clover Ad Decision Server (ADS) of the signals injected into the incoming linear video stream; the notification info includes the signal IDs and acquisitionInfo (transcoder, channels, time of acquisition). The transcoder also conditions video at the splice points by ensuring that an Instantaneous Decoder Refresh (IDR) frame is present right after the splice points.

The 720p output at a frame rate of 59.94fps from this linear transcoder is then sent to a Linear Packager and then to a QAM gateway or Ad splicer which is then modulated using edge QAMs for an RF output that contains the MPEG-2 transport stream with SCTE 35 time signal messages.

Inserting local ads applicable to an Ad Zone is done either in the QAM gateway or the Ad splicer, per the schedule defined by advertising back office and breaks allocated to the distributor. The SCTE 35 outpoint message received from the transcoder triggers the possible replacement of the national, network ad with local, zone-based ads.

Hybrid QAM-IP based Gateway STBs receive the RF signal containing the MPEG-2 TS with SCTE 35 time signal messages, and in the STB the Receiver Based Insertion (RBI) component detects this SCTE 35 message and creates an SCTE 130 placement request, which it then sends to Alternate Content Router (ACR). The ACR routes this request to the Clover ADS. The Clover ADS decides which targeted ads to play during this break, and sends this back in the form of an SCTE 130 placement response to the Hybrid QAM-IP based STB. If the ADS response indicates that a Network Ad needs to be replaced with a targeted ad, the Hybrid QAM-IP based STB, at the appropriate time, starts downloading the targeted ad, replacing the underlying ad with the targeted ad. The replacement involves replacing the underlying network ad MPEG-2 transport packets with targeted ad transport packets. The targeted ad is downloaded from Content Delivery Network (CDN) over an IP connection. The Hybrid QAM-IP based STB provides notifications to upstream components when a targeted ad starts/stops playing, or when an error occurs during the targeted ad insertion process.

For IP-based client devices connected to Gateway STBs via MOCA, targeted ad insertion for streams played out on IP-based client devices is actually done in the Gateway STB. Therefore, the incoming Single Program Transport Stream (SPTS) into IP-based client device already has the network ad replaced by the targeted ad.

The Clover ADS works with various components to make an ad decision. Specifically, it retrieves anonymized subscriber info from the Subscriber Information System (SIS), schedule information from SpotTrans, and network stream info from Compass. The Clover ADS managed by the HighYield Platform works with Ad Content Management systems to determine the targeted ads available. The Unified Ingest transcodes ads from a mezzanine format to a ready-to-air format.

The Alternate Content Router (ACR) routes the placement request and ad playout status to the appropriate components.

2.3. Title VI devices on Linear IP Channels

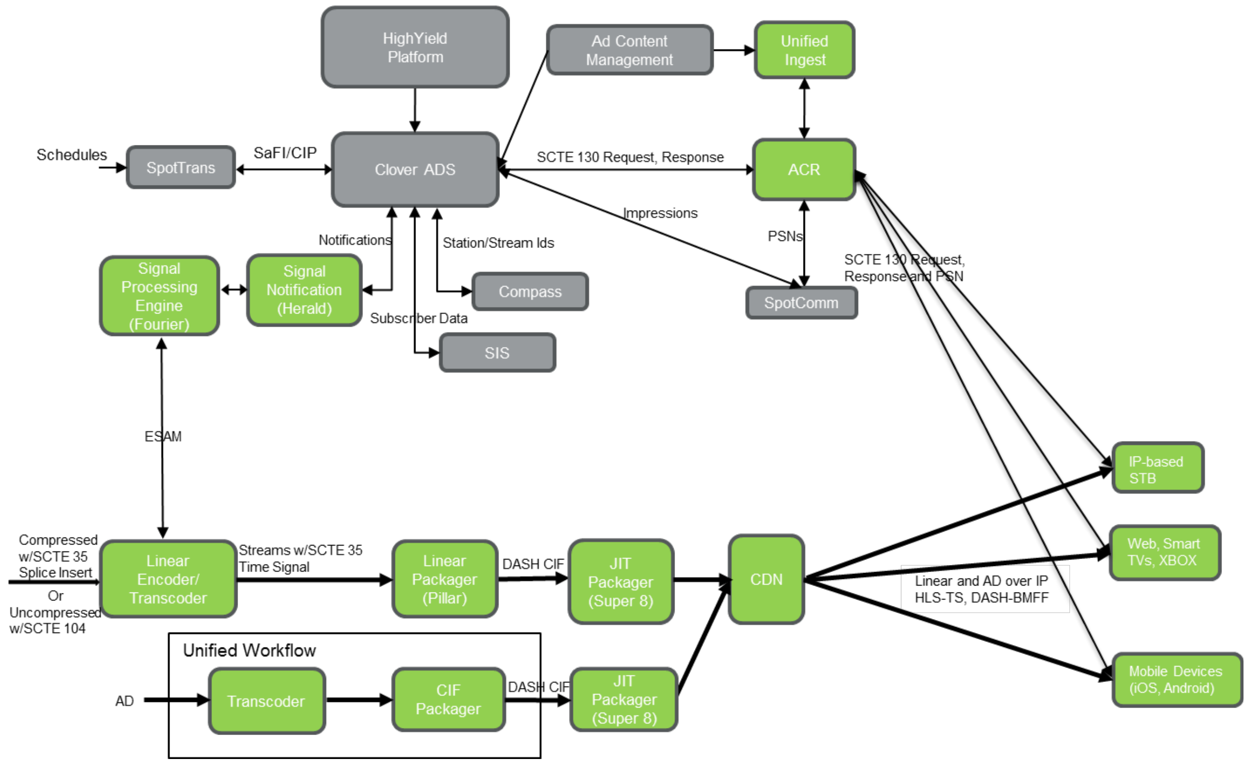

Figure 3 - Targeted Advertising Architecture for Title VI Devices (Linear IP)

There are several similarities between this architecture and the architecture for targeted advertising on Hybrid QAM-IP based STBs over linear QAM channels (Section 2.2). The Ad decision, Fourier/Herald signal processing and Encoder/Transcoder components are the same between the two. Linear IP video in either HTTP Live Streaming (HLS) or Dynamic Adaptive Streaming over HTTP (DASH) formats is created by Linear Packager (Pillar) and Just In Time (JIT) Packager (Super 8) and delivered to CDN. In the case of HLS, the segments are carried in MPEG transport stream packets and in the case DASH, the fragments are carried in fragmented MP4 (ISO-BMFF format) files. In DASH, SCTE 35 messages included in the DASH MPD file are typically carried in the Event child element under EventStream element, see SCTE standard 214-1 for additional details. In HLS, SCTE 35 messages are either carried in SCTE 35 defined HLS Tag EXT-X-SCTE35 or HLS defined tag EXT-X-DATERANGE, see HLS Internet draft and SCTE 35 standard for additional details.

The ad content is transcoded to various Adaptive Bitrate (ABR) formats and made available in either HLS or DASH using a DASH Common Intermediate Format (CIF) packager and Just In Time (JIT) packager.

In this architecture as well, the client detects SCTE 35 and makes a call to the ACR to determine the ads to play during the break. The client is provided with Uniform Resource Locater (URL) of the manifest for the targeted ad that it needs to play out.

This architecture is applicable to mobile devices, IP-based STBs, Web, Smart TVs etc that may be operating in the Title 6 environment.

2.4. TV Everywhere devices on Linear IP Channels

In the architecture shown in Figure 4, encoding/transcoding steps for linear and ad content is the same as that for Title VI. The key differences here are that the ad decisions are made by FreeWheel Monetization Rights Management (MRM) and the devices in scope for this architecture integrate the FW Plugin Ad Manager for interacting with FreeWheel MRM. The protocol used between the devices and FreeWheel is the proprietary SmartXML protocol.

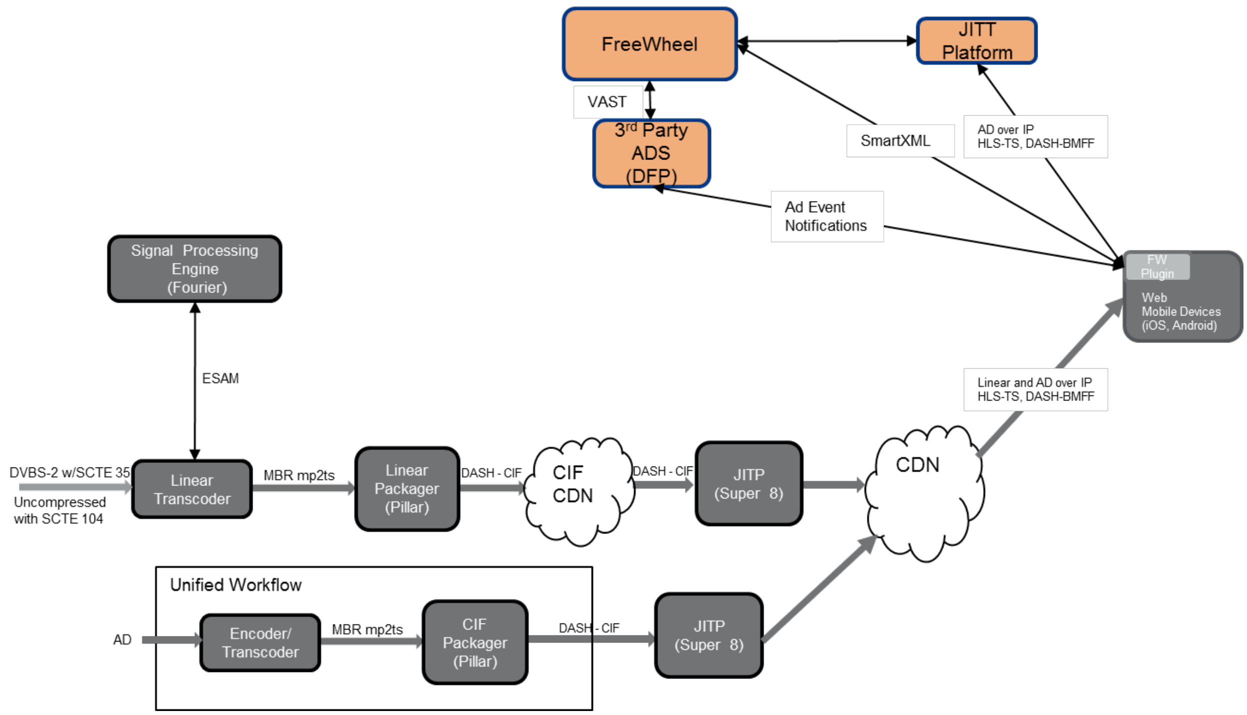

Figure 4 - Targeted Advertising Architecture for TV Everywhere Devices (Linear IP)

There are scenarios where the ads are not encoded by the Unified workflow but instead by the JITT platform.

In this architecture too, ad insertion is done in the clients, triggered by SCTE 35 messages in the linear video (HLS or DASH) streams.

This architecture is applicable to mobile devices (iOS, Android), Web etc., in the TV Everywhere environment.

2.5. Programmer Targeted Advertising on Linear QAM/IP Channels

The targeted advertising architectures for Linear Channels discussed in the previous three sections are not only capable of supporting targeted ad insertion into Distributor ad breaks but also into Programmer ad breaks with a few enhancements to the architectures. In the case of Programmer ad breaks, the programmer signals a Programmer Ad break for targeted ad insertion with an SCTE 35 time signal message with the Segmentation_type_id set to a value of 0x34 for Provider Placement Opportunity Start.

3. Hybrid QAM-IP based STB Targeted Advertising Internal Architecture for QAM/IP Channels

The Reference Design Kit (RDK)'s STB targeted advertising internal architecture is shown in Figure 5. In this architecture, the RDK Media Framework (RMF) is the component that processes input video streams from different sources and outputs them to the MPEG decoder. The component RBIFilter parses a decrypted input video stream for SCTE 35 messages, and inserts targeted ads into the stream that is then sent to the MPEG decoder. There are other components between RBIFilter and MPEG decoder that are not shown in the figure here for simplicity. Once RBIFilter detects SCTE 35 out message, it makes a call to the ACR via RBIDaemon for ad decisions. The received ad decision is parsed by RBIDaemon to determine the ads to play, and the corresponding time slots to play them. The ad download via HTTP starts four seconds ahead of the spot start time.

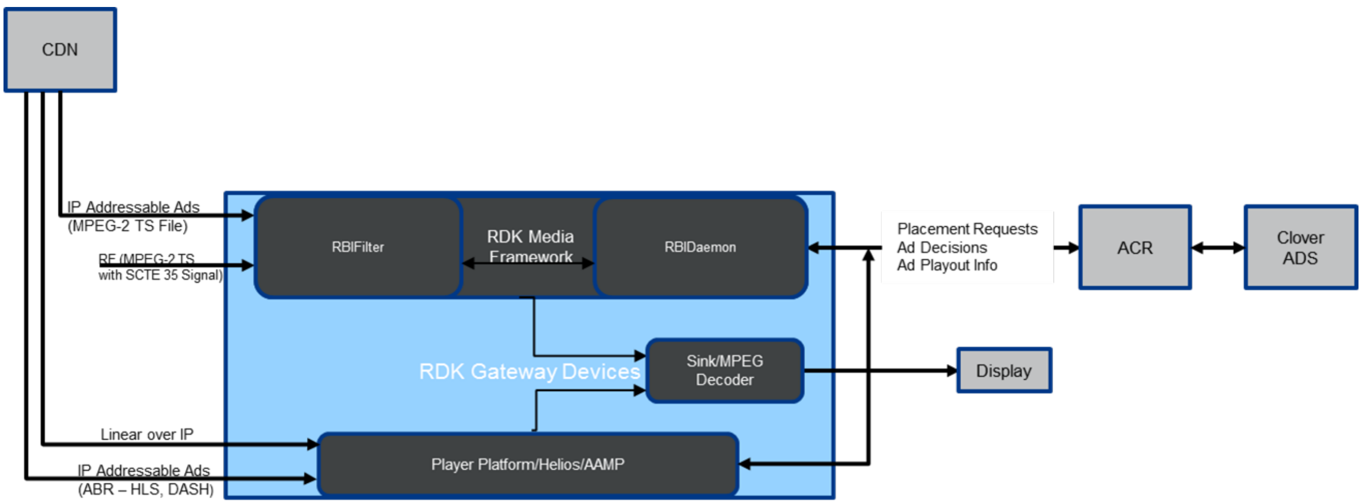

Figure 5 - RDK STB Targeted Advertising Internal Architecture (QAM/IP Linear Channels)

In addition to SCTE 35 messages, the RBIFilter continually monitors the input video stream packets for various things. When the RBIFilter detects the Out Point in the input stream, it starts replacing incoming audio/video packets after the Out Point with ad packets and sends them to the sink/MPEG decoder. This continues until it detects the In Point, at which point it stops the replacement and passes the input stream as it is to the sink/MPEG decoder. During these transition times, RBIDaemon sends notifications indicating the start and stop of ad play to Ad BackOffice (via the ACR).

The Player Platform/Helios/AAMP component would use a similar approach for replacing underlying ads in the incoming IP video stream, with targeted ads. The Player Platform/Helios interacts with the ACR to determine the ad decisions. In the IP stream, as well, the output MPEG stream with targeted ad is sent to sink/MPEG decoder on RDK STBs, or Apple AV Foundation in iOS devices, or the ExoPlayer in the case of Android devices.

This internal architecture is applicable to an RDK STB with a QAM tuner; for an RDK IP STB the RBIFilter and RBIDaemon components in the figure are not present.

Detailed next are some technical challenges faced with targeted ad insertion in linear video and their possible resolutions.

4. Technical Challenges with Targeted Advertising on Linear Channels

4.1. Seamless Audio/Video Transitions

Achieving seamless audio/video transitions when transitioning from network to targeted ad content and vice versa with compressed linear video and audio frames is not a trivial task. The Audio and Video compressed frames received by a client device either in the MPEG transport stream or in ISO BMFF formats (MP4 file) are sent to the MPEG decoder for decoding and presentation. If some of these frames are not sent to the decoder in a timely manner or not sent at all, it can result in video/audio buffer overflow or underflow, which, in turn, can produce video/audio artifacts. This is especially true in client-side targeted ad insertion, where the client device is replacing linear network content with targeted ad content. For that reason, additional steps must be taken to ensure that the compressed frames from this modified linear stream are received by the decoder and in a timely manner. The following paragraphs detail some approaches for non-seamless splicing, after which we discuss how seamless splicing is achieved, specifically in cases where linear content and targeted ads are carried in MPEG transport stream.

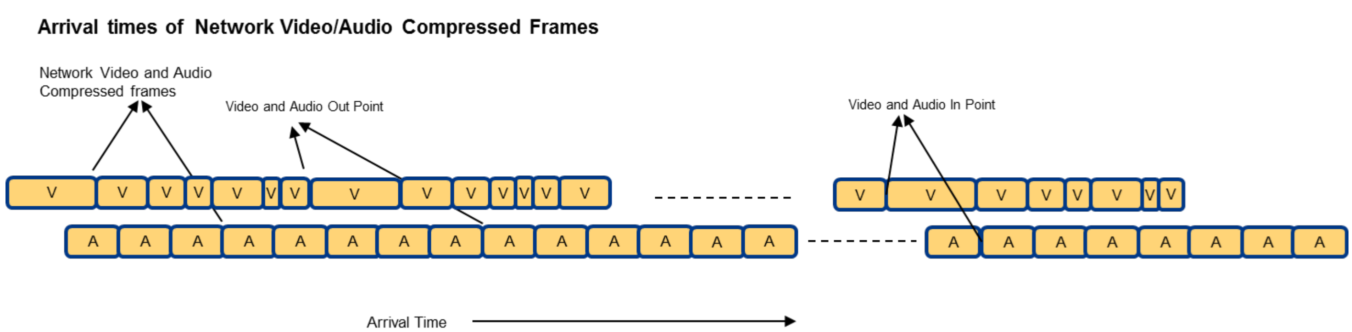

Figure 6 – Arrival times of Network Video/Audio Compressed Frames

Figure 6 shows the arrival times of Network Video/Audio compressed frames in the MPEG transport stream at the input to a decoder. The Video and Audio Out Points (a point in the stream between two presentation units, suitable for exit) are determined based on the PTS value (pts_time as modified by pts_adjustment) in the SCTE 35 message. The video frame before the Out Point is the one from which there is a clean exit (for example after the last B frame in the sequence PBB... in decode order.) Typically, the first video frame in the network content after this Out Point is an IDR frame. The Video and Audio In Points (a point in the stream between two presentation units, suitable for entry) could be determined based on duration of the targeted ad or the PTS value in the SCTE 35 message for In Point.

The first video compressed frame in the network content after this In Point is an IDR frame. As shown in Figure 7, with seamless transition between network content and targeted ad, the expectation is that the video and audio frames across the transitions are presented at the appropriate rate without any interruptions (such as frozen frames, dropped, distorted frames etc.) To achieve seamless transitions, it is also required that both network and ad video is compressed using the same video compression format, and audio using the same audio compression format (Dolby Digital or Advanced Audio Coding (AAC)). Further, any video/audio attribute differences between the two does not require a decoder reset at the transition points. Resetting the decoder would result in observable artifacts.

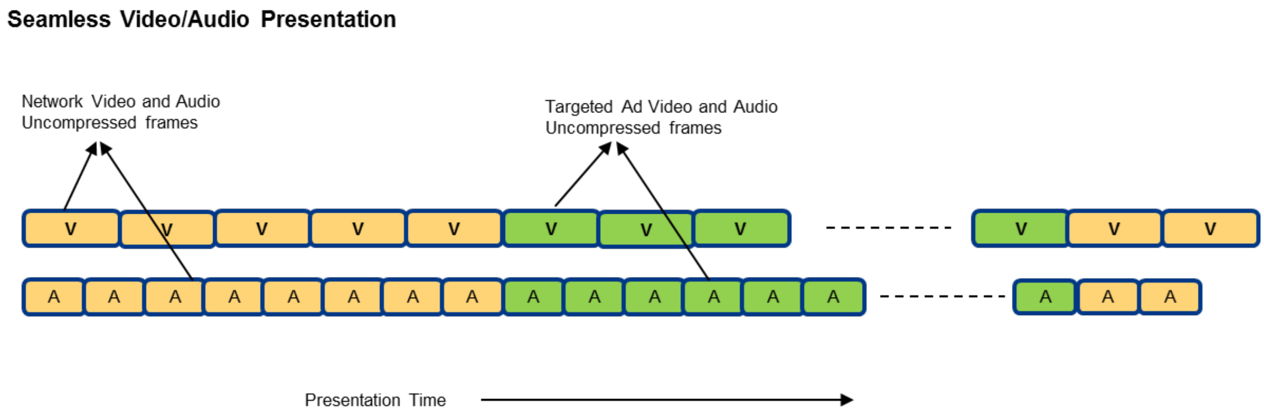

Figure 7 – Seamless Video/Audio Presentation

For seamless splicing with uncompressed video/audio frames where the arrival times of audio/video frames with the same presentation times are time-aligned, we exit the stream right after the last network video frame preceding the splice point, and enter the network stream right after the last ad video frame. Taking a similar approach when splicing with compressed video/audio frames, as shown in Figure 8, network transport stream packets are replaced with targeted ad transport stream packets right after the last packet of the network video frame, which is immediately before the Video Out Point. As a result, as shown in the figure impacted network audio frames indicated by a pattern fill are dropped, and also an audio frame overlapping the Video Out Point is partially sent. In this scenario, audio pop can be heard around the transition, and network audio clipping occurs (which can manifest as dropped words or phrases). The same issues are seen at the transition back to the network feed where the targeted ad audio frames are clipped.

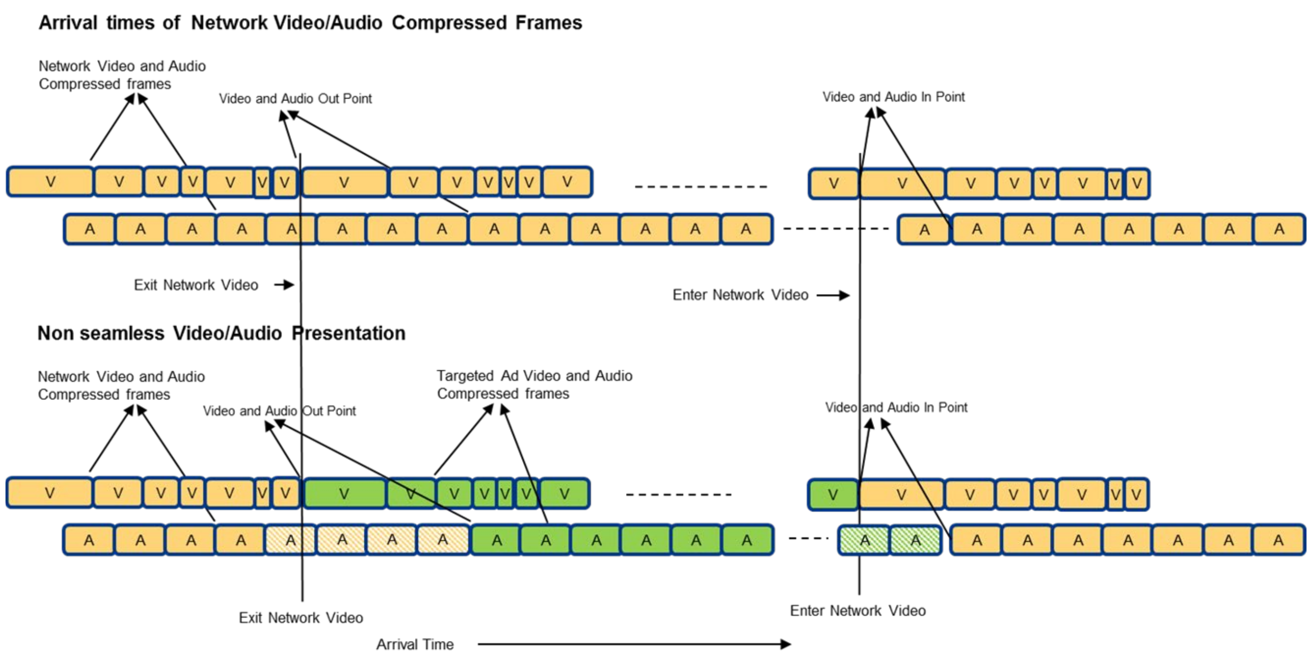

Figure 8 – Non Seamless Video/Audio Splicing (Network Video Out Point Exit)

One other approach as shown in Figure 9 might be to replace network transport stream packets with targeted ad transport stream packets, right after the last packet of the network audio compressed frame that immediately precedes the Audio Out Point. With this approach, the duration available to insert the targeted ad packets is reduced by dt (time interval between the arrival of the last packet of the last compressed network audio frame and the last packet of the last compressed network video frame) as shown in the Figure 9. As a result, not all the targeted ad video/audio compressed frames are placed in the transport stream sent to the decoder. During the duration dt, any network video transport packet is converted to null packets before including them in the transport stream sent to the decoder. This results in frozen network video frames at the start of targeted ad playback, and also clipping of targeted ad video/audio, with the possibility of sending a partial ad audio compressed frame towards the end of targeted ad. This reduction in duration available for streaming ad video/audio frames could be resolved by presenting the audio/video frames at a slighter faster rate, but doing so could introduce audio/video distortions.

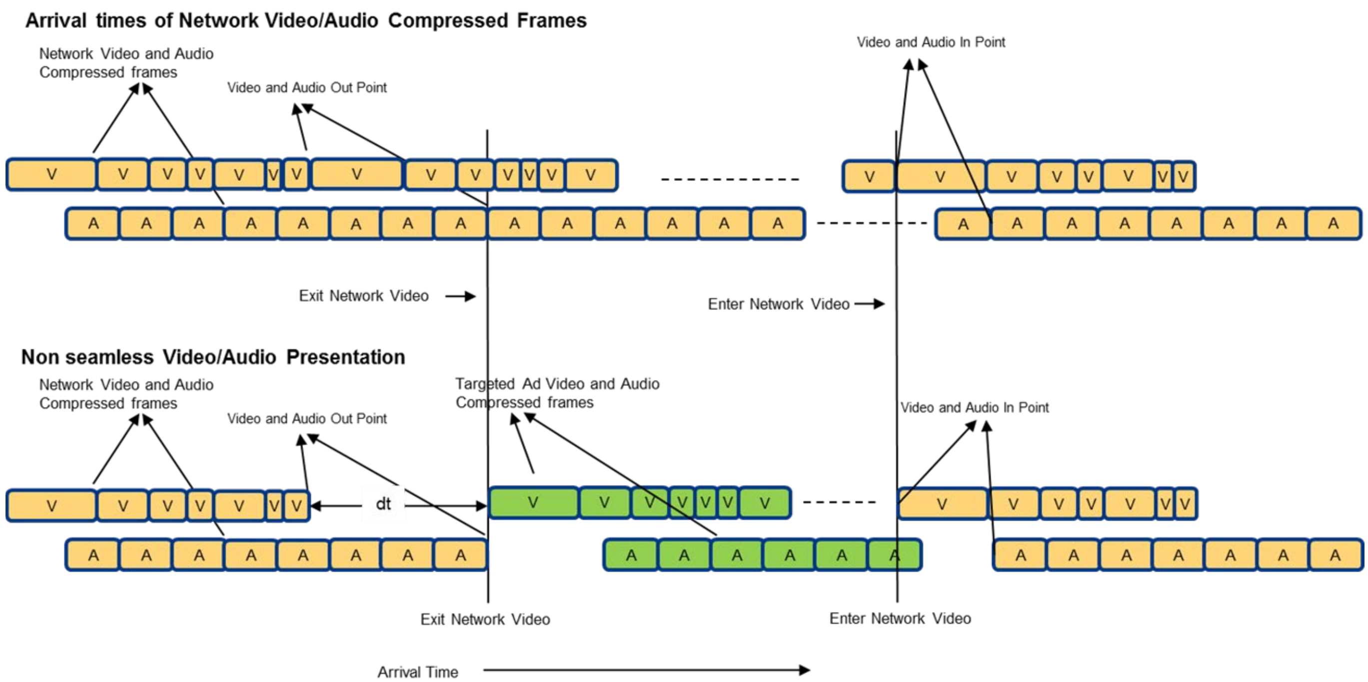

Figure 9 – Non Seamless Video/Audio Splicing (Network Audio Out Point Exit)

Figure 10 shows the arrival times of Video/Audio compressed frames with targeted ads replacing network content (typically a national ad) in the modified MPEG transport stream sent to the decoder. To achieve seamless splicing, insertion of targeted ad video frame packets should start after network Video Out Point but before the start of network Audio Out Point, and targeted ad audio frame packets should start after network Audio Out Point, as illustrated in Figure 10. The network video transport packets between the network Video Out Point and Audio Out Point are converted to MPEG TS null packets; many of these null packets are replaced with targeted ad video transport packets in this approach.

Similarly, on the transition back to network content, the insertion of network video frame packets starts at the Video In Point right after the last frame of the targeted ad, but before the start of network Audio In Point, and the network audio frame packets start right after Audio In Point. The network audio transport packets between network Video In Point and Audio In Point are converted to MPEG TS null packets; many of these null packets are replaced with targeted ad audio transport packets. Also, the arrival of the ad video/audio frames at the decoder are controlled such that the presentation of network and targeted ad video/audio compressed frames is seamless across the transition points. MPEG timestamps Presentation Time Stamp (PTS), Decode Time Stamp (DTS) and Program Clock References (PCR) in the ad audio and video transport packets are appropriately restamped for a seamless ad playout.

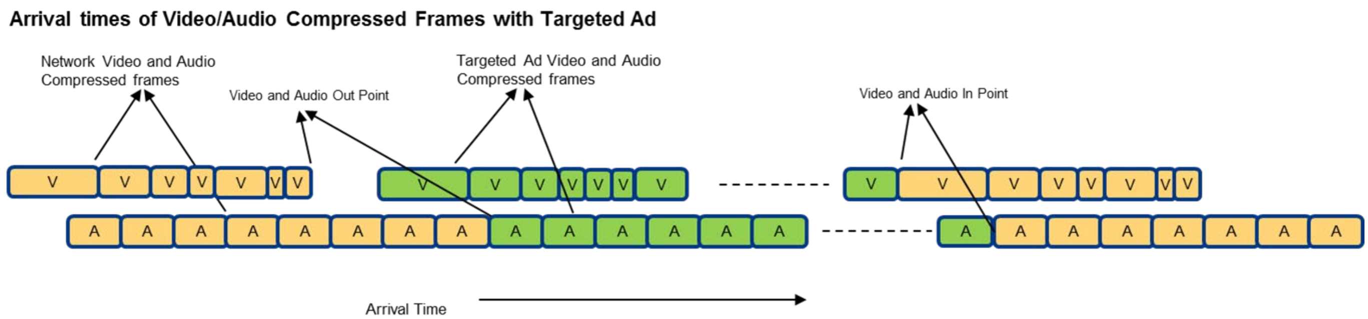

Figure 10 – Arrival Times of Video/Audio Compressed Frames for Seamless Video/Audio Presentation

4.2. SCTE 35 PTS adjustments

The PTS value (pts_time as modified by pts_adjustment) in an SCTE 35 message represents the time of the intended splice point, which is immediately prior to the presentation unit whose presentation time most closely matches this PTS value. And in the case of an Out Point it is a point in the stream, suitable for exit.

This value set initially by the programmer could be incorrect, pointing to the wrong splice point. If the incorrect splice point is before one or more presentation units of the intended presentation unit, then any network content before the start of the targeted ad is clipped, and the ending frames of network ad that is being replaced with targeted ad can be seen. If the incorrect splice point is after the intended presentation unit, then the starting frames of the network ad can be seen before the transition to the targeted ad, and the starting frames of the network content on the transition back from the targeted ad to the network content is clipped. The incorrect splice point occurs when programmers utilize master control schedules that mismatch with the programming segments, and are thus not guaranteed to be frame accurate. This is addressed by making an adjustment on the transcoder to either add/subtract an offset value from the splice point, and most effective if the splice point is off by a constant value from the intended.

Downstream devices from the transcoder could incorrectly restamp the PTS value in the SCTE 35 message, causing transition artifacts due to a premature exit from an MPEG stream say, after a P frame in IPBB sequence in decode order. Also, the transition back to network content after ad playback could be delayed causing frozen frames or joining the stream at a non-IDR frame resulting in broken pictures. This would be detected by performing an MPEG analysis of streams generated at the output of the downstream device, then comparing those against the incoming stream. Fixes would be made in the downstream device to address this issue.

4.3. Network/Targeted Ad duration mismatch

It is also possible that the duration of the network ad that is to be replaced with targeted ad might mismatch with targeted ad duration. If the network ad duration is longer than the targeted ad duration, then on the transition from targeted ad to network content, part of the network ad is shown. Plus, the entry may be delayed because of a lack of IDR frame right after the In Point, resulting in frozen frames. And if the network duration is shorter than the targeted ad duration, then on the transition from targeted ad to network content, part of the next network ad or program content is shown. Again, the entry may be delayed due to lack of IDR frame right after the In Point, resulting in frozen frames.

Checks must be made to determine duration of targeted ads provided by the ad creator, to have them fixed to avoid these artifacts. In the case of network ad duration inaccuracies, this must be addressed by the network content provider.

4.4. Network and Targeted Ad Video/Audio Attributes mismatch

For seamless splicing, the network and ad video/audio attributes need to match. For example, if the network is encoded using MPEG-4/AVC, it is expected that the targeted ad would be encoded using MPEG-4/AVC. If not, then glitches are present during the transitions, since the video decoder is reset to handle the codec change.

Bitrate differences between network and targeted Ad on the client do not pose an issue, because adjustments can be made to the rate at which compressed frames are provided to the decoder. This achieves a timely decode and presentation.

In a scenario where there is a mismatch in the number of audio streams between the network feed and the targeted ad, and the viewer is tuned to for example an SAP audio stream and the corresponding stream is not available in the targeted ad, silence is heard during the targeted ad playback. This is avoided by having the client tune to the targeted ad primary audio during the duration of the ad, or by duplicating primary audio for the missing audio track.

4.5. Invalid decisions from Ad Decision Server

The client is expected to detect errors or missing information in the ad decisions from the ADS. If mandatory information required for accurate playout is not available, the client is expected to ignore the ad decision. For example the targeted ad pointed to by the ADS might be invalid.

4.6. Scalability, Reliability and Availability

The targeted advertising architectures presented are required to be scalable, reliable and highly available.

The solutions are designed such that it is easily scalable to handle existing and upcoming targeted advertising applications. Dynamic scaling of components is desirable to handle any unanticipated increase in traffic.

To ensure high availability, redundancy is expected to be incorporated into the solution with automatic failover.

Appropriate load tests are carried out on the entire system and on the individual components in the targeted advertising architectures to confirm that they can handle the desired load. IP Bandwidth modeling is done with ad messages (For example, Placement requests, status notifications etc.), ad downloads etc. to determine peak traffic especially as a result of primetime linear events, and the solution designed to handle this peak traffic.

Conclusion

Targeted advertisements in digital cable television are an important revenue source, and simultaneously necessary, to enable a better user experience with cable television. This paper detailed the complexity involved with implementing targeted advertising across the broad spectrum of client devices and video delivery mechanisms. As indicated earlier, achieving seamless audio/video transitions when transitioning from network to targeted ad content and vice versa, with compressed linear video and audio frames, is not a trivial task. The techniques detailed in the paper help to make frame accurate seamless ad insertion achievable.

Please note that these architectures are continuously evolving, to meet the needs of Comcast and its syndication partners.

Abbreviations

ADS	Ad Decision Server
AAC	Advanced Audio Coding
ACR	Alternate Content Router
ABR	Adaptive Bitrate
BDMS	Business Data Management System
CIF	Common Intermediate Format
CDN	Content Delivery Network
DASH	Dynamic Adaptive Streaming over HTTP
ESAM	Event Signaling and Management
HLS	HTTP Live Streaming
IDR	Instantaneous decoder refresh
JIT	Just In Time packager
MRM	Monetization Rights Management
MPEG	Motion Picture Experts Group
OOO	Out Of Band
PTS	Presentation Time Stamp
RBI	Receiver Based Insertion
RDK	Reference Design Kit
RMF	RDK Media Framework
QAM	Quadrature Amplitude Modulation
SCTE	Society of Cable Telecommunications Engineers
SIS	Subscriber Information System
VANC	Vertical Ancillary Data

Bibliography & References

DASH Industry Forum, http://dashif.org/

Dynamic Adaptive Streaming over HTTP (DASH), ISO/IEC 23009-1:2014

H.262, Generic Coding of Moving Pictures and Associated Audio Information (H.262), ISO/IEC 13818-2:2013.

H.264, Coding of Audio-Visual Objects — Part 10: Advanced Video Coding (AVC), ISO/IEC 14496-10:2010.

MPEG-2 Transport Stream, ISO/IEC 13818-1.

SCTE 30, Digital Program Insertion Splicing, ANSI/SCTE 30 2009

SCTE 35, Digital Program Insertion Cueing Message for Cable, ANSI/SCTE 35 2016.

SCTE 67, Recommended Practice for Digital Program Insertion for Cable, ANSI/SCTE 67 2017

SCTE 130-3, Digital Program Insertion - Advertising Systems Interfaces - Part 3: Ad Management Service (Adm) Interface, SCTE 130-3 2013

SCTE 214-1 2016 - MPEG DASH for IP-Based Cable Services Part 1: MPD Constraints and Extensions

SCTE 224 2018r1 - Event Scheduling and Notification Interface

R. P. Pantos, editor, HTTP Live Streaming, https://tools.ietf.org/html/draft-pantoshttp-live-streaming-23

