ADVANTAGES AND CHALLENGES OF A VAST SERVER-SIDE VIDEO ADVERTISING SOLUTION

T. LEVY

Cisco, Service Provider Video Software & Solutions, France

ABSTRACT

Video advertising is still one of the leading approaches with which content and service providers can monetize video content. Combine that with the common-sense approach of using open standards and surprisingly you find yourself facing a problem.

This paper presents the two main open standards for providing dynamic video advertising are SCTE-130 and IAB VAST. But, SCTE was developed with Service Providers and traditional North American broadcast cable in mind and IAB VAST was developed for web-based advertising. Neither of these advertising standards options is ideal for the entire range of devices.

This paper presents a novel VAST server-side solution where a video streaming server acts as a VAST client fronting the client device, and merges the main video and the advertising returned from a VAST server into a single video stream. This solution combines the advantages from both options above and thus allows ad and content providers to address the full set of devices with a single platform to operate.

INTRODUCTION

Video advertising is still one of the leading approaches with which content and service providers can monetize video content. Combine that with the common-sense approach of using open standards and surprisingly you find yourself facing a problem.

The two main open standards for providing dynamic video advertising are:

1. SCTE (Society of Cable Telecommunications Engineers) – which allows a video streaming server to exchange video advertising with an ad decision server through SCTE-130, and
2. IAB (Interactive Advertising Bureau) VAST – that enables a video player to retrieve video advertising from an ad decision server to be played before, after and in the middle of a video stream.

At the heart of the issue is the fact that SCTE was developed with Service Providers and traditional North American broadcast cable in mind and IAB VAST was developed for web-based advertising. Neither of these advertising standards options is ideal for the entire range of devices.

This paper presents a novel VAST server-side solution where a video streaming server acts as a VAST client fronting the client device. This server-side solution then merges the

main video and the advertising returned from a VAST server into a single video stream. This solution combines the advantages of both of the above options and thus allows advert and content providers to use a single platform to address the full set of devices.

First, the paper describes the two main video advertising standards mentioned above. Then, it presents the architecture and the advantages of a VAST server side solution. Finally, based on our ongoing implementation and deployment experience here in Cisco, it discusses the challenges and next steps of such an approach.

SCTE APPROACH

SCTE 130 standard [1] provides a reference logical architecture and standardized interfaces allowing an Ad Management Service (ADM) to perform Dynamic Ad Insertion where the ad decisions are provided from an Ad Decision Service (ADS). Although SCTE 130 can cover several use cases (including Video On-Demand, Linear services, Digital Video Recorders), this paper will mainly focus on the Video On Demand (VOD) case. However, the solution presented could be extended to these other use cases.

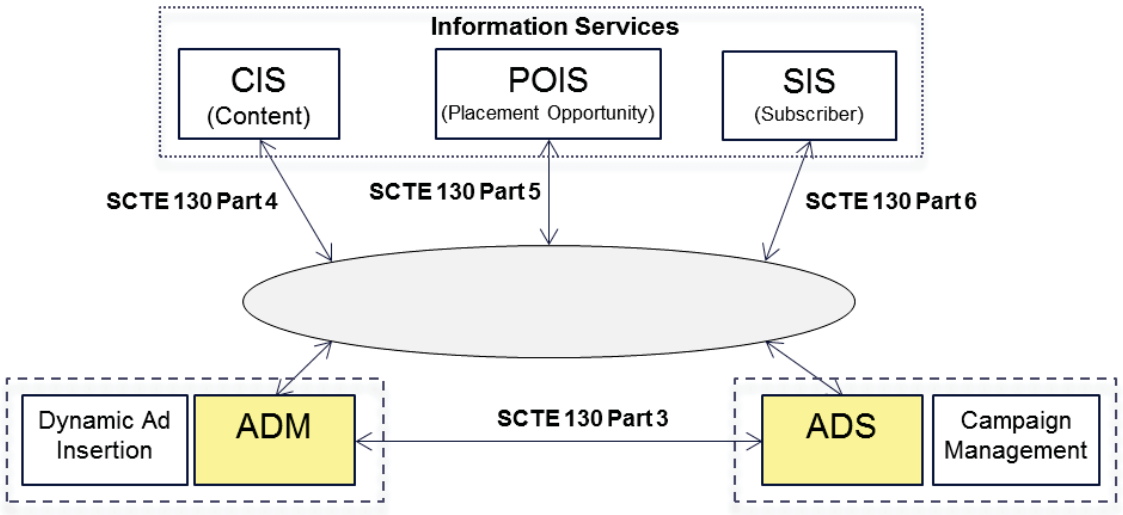

Figure 1 – Dynamic Ad Insertion with SCTE 130 - Reference Architecture

Logical Elements

The SCTE 130 reference architecture shown in figure 1 above includes the following elements:

- **ADM**: The Ad Management Service is usually located within the network element that performs dynamic ad insertion. The ADM queries the ADS to determine the ads to place inside the content. The decisions are conveyed using the PlacementRequest and PlacementResponse messages defined in SCTE 130 Part 3 [3].
- **ADS**: The Ad Decision Service determines what advertising content needs to be inserted inside non-advertising (i.e. entertainment) content. It is typically linked with campaign management where a service provider defines how a particular set of advertising needs to be mapped to content depending on various criteria (such as subscriber data, time of day, type of content …).

- **Information Services**: A grouping of several logical services providing information to the ADS and ADM elements described above. This includes:
    - **CIS**: The Content Information Service manages metadata describing all the assets (both advertising and entertainment) available to the other SCTE 130 logical services. The CIS interfaces and messages are defined in SCTE 130 Part 4 [4].
    - **POIS**: The Placement Opportunity Information Service manages the descriptions of placement opportunities in the entertainment assets. The POIS interfaces and messages are defined in SCTE 130 Part 5 [5].
    - **SIS**: The Subscriber Information Service manages per-subscriber information relevant to ad placement decisions. The SIS interfaces and messages are defined in SCTE 130 Part 6 [6].

SCTE 130 protocol

SCTE 130 is defined as an extensible XML protocol with core elements defined in SCTE 130 Part 2 [2] that can be extended with other more specific interfaces.

Although defined initially in a cable network context, SCTE 130 can be applied to any video advertising environment relying on an IP capable network.

In particular, in a pure IP environment, it can be used for doing dynamic ad insertion within a content delivered in Adaptive Bitrate streaming over HTTP (ABR) – such as HTTP Live Streaming (HLS) or Dynamic Adaptive Streaming over HTTP (DASH).

Analysis

SCTE 130 is a mature standard that has been successfully deployed by many content providers in the delivery of live and Video on Demand services to users with managed equipment (such as: Customer Premise Equipments (CPE) or Set Top Boxes (STB)).

However, it relies on the assumption that the ad provider providing the ADS manages the campaign with advertisement assets that are shared with the service provider that delivers the video services. Typically, the ads need to be ingested in advance by the video service provider. Thus, it requires a strong coupling between the video service provider and the ad provider that may restrict the exposure of a new ad campaign and its related revenue.

VAST APPROACH

Video Ad Serving Template (VAST) is a specification released by the Interactive Advertising Bureau (IAB) that allows a video player to retrieve ads from one (or more) ad servers. VAST target multiple types of devices (but web clients in particular) meaning that it supports multiple type of ads. These types cover:

- linear ads that are video, and are typically displayed as a video at the beginning/in the middle/at the end of the main video content, or
- nonlinear ads that are typically images displayed over video content, or
- companion ads that are served as linear or nonlinear ads, but are displayed outside the video player.

VAST Protocol

VAST Protocol version 3 [7] is currently the most widely deployed VAST version. Like SCTE 130, it is defined as an XML schema for serving ads to the video player. However, unlike SCTE 130, it directly includes the media files corresponding to the ads selected by the ad server. Thus, the client itself retrieves the ads and displays them as independent elements. Consequently, the video player typically displays a linear video ad as a different video compared to that used for the main video content.

Because the ad and its related media files are directly served by the ad server, when a client requests an advert, the primary ad server can redirect the request to a secondary ad server (for instance: if it doesn't own the appropriate ad campaign for the content). In such a case, the secondary server will return the ads and the related media files in the final response. Upon triggering specific events (like the start) for the ad, both ad servers are notified using a tracking URI provided in their VAST responses. This process is shown in figure 2.

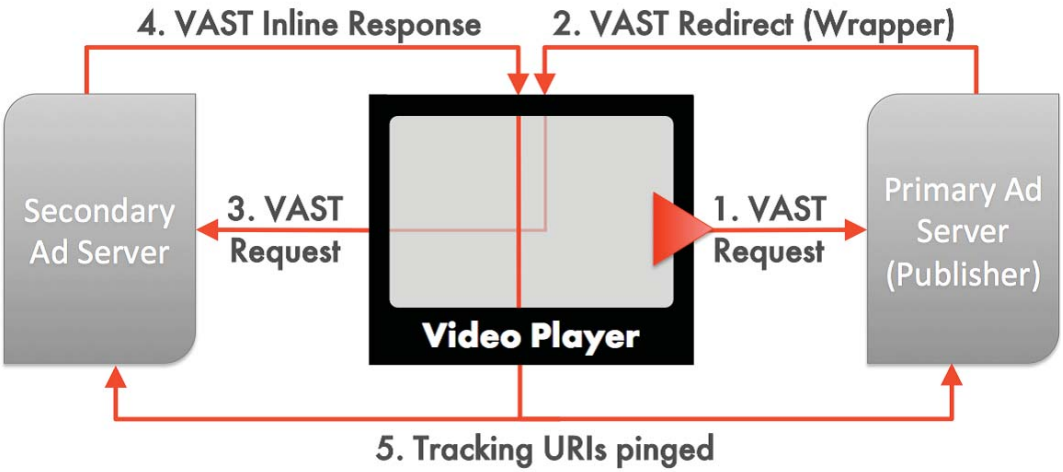

Figure 2 – VAST with redirection to a secondary ad server

This ad-serving scenario can be extended beyond just two ad servers. For instance, the secondary ad server may respond with a VAST wrapper that points to yet another ad server.

Analysis

The VAST Redirection shows how a relationship between a video service provider and its primary ad provider can be extended to other ad providers where the primary ad provider doesn't have enough (or the appropriate) ad campaign. This aspect is attractive to video service providers with the possibility to expand their revenue by filling their ad placement opportunities with multiple ad providers.

Unfortunately, the fact that the video player directly retrieves the ads can generate a user experience that suffers from delay and buffering during the playback of the video content. This may be widely accepted when playing OTT services on web clients (like Youtube or Dailymotion), however on managed devices (such as CPE or STB), such user experience may be more difficult to accept. This is particularly the case when comparing with the existing experience with pre-inserted ads inside video content.

Finally, the rise in Ad Blockers installed on web clients represents a threat to the VAST model. The ease with which requests handled at the video player level can be filtered threatens the basis of making decisions at the video player level that in turn is a serious challenge to the basis of the VAST model.

VAST SERVER SIDE MODEL

To mitigate the limitations of the current VAST model, we have developed a VAST server side model. In this approach a server side component that acts as a VAST client fronting the client device and merges the main video and the advertising returned from a VAST server into a single video stream.

This approach combined with ABR delivery (HLS or DASH) is particularly interesting because the VAST server side component can manipulate video manifest to insert ads without having to manipulate the video segments, thus allowing the VAST server side model to be introduced in a very smooth way.

VAST Ad Manager

The solution that provides VAST server side combined with ABR relies on a new service called VAST Ad Manager (VADM). This service can be also viewed as an extension to an SCTE ADM component that manages the VAST server side ad management in addition to the normal SCTE-130 ad management. The architecture of this solution is shown in Figure 3.

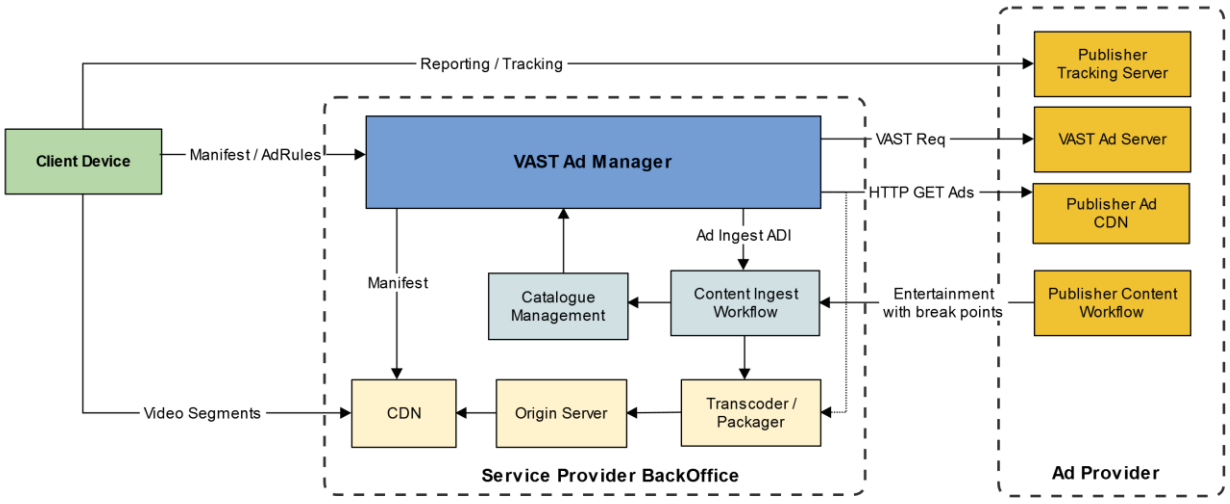

Figure 3 – VAST Ad Manager with ABR delivery

For content where a VAST ad insertion is required, the VADM will be configured as an HTTP proxy that will process the incoming manifest HTTP requests. Whether to use the VADM or request the manifest directly from the CDN can be decided based on the Content Provider or any other criteria related to the entertainment content.

When it receives a request for a Manifest, the VADM will in turn retrieve this from the CDN. However, it does not immediately return the manifest, but instead initiates a sequence of steps prior to returning a possibly modified manifest to the client that requested it in the following way:

- First, the VADM will determine whether one or more ad opportunities are present in the content the manifest corresponds to, and if they should be pre, mid or post roll.
- Second, and based on these decisions, the VADM will issue one or more VAST requests to the appropriate VAST ad server(s).
- Then, and based on the VAST response, the VADM will retrieve the linear ads (only the linear video ads can be inserted inside the video entertainment content).
- Next, it will ensure that all the adverts are available for delivery. To achieve this, for each linear ad, it will check whether a version of this video ad exists in the service provider back office at the same bitrates as the entertainment content.
- If they are not already available, the VADM will trigger the local ingest and encoding of the ad. For instance, this can be done by building a related ADI file that will be issued to the usual automatic content ingest workflow of the service provider.
- Finally, but without waiting for the previous step to complete if it was triggered, it will return the modified manifest. We do not wait for the ingest to complete as the local encoding of the ad might take few seconds depending on the ad duration, and this would unacceptably delay the response. A consequence of this is that the first time an ad is returned by a VAST ad server it will not be placed into the manifest, but it will be inserted the next time there is a VAST response with this ad.

This process will result in a modified manifest returned by the VADM that aggregates segments corresponding to the linear ads locally available into the CDN of the service provider into a single manifest corresponding to the requested content. As a result, the client will display the video and the ads as a single video content.

There are two options to track events related to the ads (like the start of viewing):

- The tracking is performed at server side by the VADM, or
- The tracking is performed at client side (as for the traditional client side model).

The issue with tracking at server side is that with ABR delivery the network has limited knowledge about the client viewing. For a limited number of types of tracking events (like start viewing and completed viewing), a server side session control element can determine the start (and end) of viewing of a given ad, and so trigger the appropriate tracking event. However, in general, we recommend that the client triggers the tracking at the exact ad viewing event. This latter option requires that the tracking event URLs are passed to the client. This can be done by passing a part of VAST response back to the client in a form that it uses to trigger tracking events.

ADVANTAGES AND CHALLENGES

We have developed the VADM solution that works with a VAST server. This solution has shown that a VAST server side solution can be delivered with an existing VAST platform, and that this has several important advantages:

- There is no delay and buffering at the client side when retrieving the ads
- The ads are played with the main content as a single content (with a single timeline for the whole content).

- The video segments corresponding to the ads are delivered from the same CDN as the entrainment content. Thus, it is far harder for an Ad Blocker to identify and block the ads.

However, even though much of the VAST client logic is moved at server side VADM component (making the client easier to support dynamic ad insertion), the parts related to tracking are not, and cannot be, implemented by VADM. Thus it is necessary that some parts of an overall advert substitution solution (in particular for triggering the tracking events) are present either in the client, or if a subset of tracking events, in the session control or network elements.

Next steps

We believe that this VAST server solution is promising because it allows addressing both managed devices and web devices with the same VAST platform. We are currently deploying this solution for a European customer and we expect to have further results and experiential information after the deployment is completed.

IAB also recently released a version 4 of VAST standard [8] that includes some improvements to better support a VAST server side solution. In particular, version 4:

- Allows a server to return a Mezzanine media file associated with a linear ad. This Mezzanine is a high quality video file that is more appropriate for transcoding the ad to a better resolution and delivering a better video quality when playing ads on clients with high-resolution screens.
- Adds a Universal Ad Identifier that identifies an ad uniquely across multiple ad providers. This identifier allows the VADM to identify an ad across different ad servers that would have used different Creative Identifiers, and so ensure that ingest and encode occurs only once for every ad.

VAST servers implementing VAST version 4 are not yet widely deployed but we expect to be able to test these improvements in the near future.

CONCLUSION

Video service providers are both willing to, and actively want to, unify their platforms that offer video services across multiple screens. For many of these providers, supporting advert insertion is key to their business model and so that, too, must be unified. We have described a solution that combines the interest of the VAST model with a server side approach that makes it applicable for the different platforms, and outlined the benefits of this architecture. As a result, we believe that this VAST server side approach, and our VADM solution, will become increasingly popular for video service providers in the coming years.

REFERENCES

1. SCTE 130-1 2013: Digital Program Insertion—Advertising Systems Interfaces Part 1—Overview
2. SCTE 130-2 2014: Digital Program Insertion—Advertising Systems Interfaces Part 2 - Core Data Elements

3. SCTE 130-3 2013: Digital Program Insertion—Advertising Systems Interfaces Part 3 – Ad Management Service (ADM) Interface
4. SCTE 130-4 2011: Digital Program Insertion—Advertising Systems Interfaces Part 4 – Content Information Service
5. SCTE 130-5 2010: Digital Program Insertion—Advertising Systems Interfaces Part 5 – Placement Opportunity Information Service
6. SCTE 130-6 2010: Digital Program Insertion—Advertising Systems Interfaces Part 6 – Subscriber Information Service
7. Video Ad Serving Template (VAST) Version 3.0 – Published by IAB in July 2012
8. Video Ad Serving Template (VAST) Version 4.0 – Published by IAB in January 2016

ACKNOWLEDGEMENTS

I would like to thank the AdSuite team who have contributed to this interesting work and Kevin and David for their helpful comments on the paper. I would also like to thank the International Broadcasting Convention for giving the chance to publish and present this paper.