White Paper

Challenges of Dynamic Ad Insertion in the OTT Environment

Saurabh Gulati
Interra Systems, Inc.

1. INTRODUCTION

We are living in a technology-centric world where each day a new innovation attempts to make our lives easier, more comfortable and enjoyable. OTT is one such technology that has brought conventional television to the palm of our hands. Today, with HTTP adaptive streaming, anyone can watch the content of their choice anytime and anywhere using smartphones, tablets, and PCs.

With the drastic growth in consumer demand, OTT service providers and broadcasters are looking for a profitable business model to generate revenue from OTT streaming services. Subscription-based, pay-per-view, and ad-based models are common revenue models for OTT. Now ad-based models are increasingly popular and being used due to their potential to generate high revenues. YouTube is one example of the advertising video-on-demand model. Here is where dynamic ad insertion (DAI) finds its place in the OTT workflow. DAI enables broadcasters to insert ads, specifically targeted to end users, into the content on the fly. This paper talks about the challenges of DAI in an already complex OTT workflow.

2. WHAT IS DAI?

To understand DAI, it's important to first comprehend static ad insertion (SAI). Traditionally, SAI burns the commercial spot into the video asset. This static advertising does not take demographics into consideration, and it means the same advertisement is shown across all states of a country. Due to drawbacks of this approach, and enabled with technological advances, most online services have moved to a client-side model that inserts specific, relevant advertising into the video stream at the point of playback. This ensures that the commercials are appropriately targeted and timely, because the commercial can be changed for each viewing.

DAI is a technology that allows advertisers to target ad creatives in linear, live, or video-on-demand content to each specific viewer by leveraging the audience insights that are available through video ad servers. This allows more relevant ads on an impression-by-impression basis, increasing the likelihood that consumers will engage with the message and

ultimately make purchases. Today, DAI is in use across platforms for linear broadcast, VOD, mobile, and OTT, enabling targeting by platform, device, audience, and geography.

Media owners rely on ad revenue through higher Cost per Mile/Thousand (CPMs). Thus, delivering well-targeted ads provides media owners with more revenue.

How DAI Works

DAI for linear streaming relies upon SCTE-35 cue messages embedded in MPEG-2 transport streams from the linear network or SCTE-104 cue messages from the SDI content to signal ad opportunities. These cue messages are timed metadata, as they are used to mark and signal information related to a certain timestamp or time range in a stream. As depicted in Figure 1 below, upon receiving these markers or cue messages the splicer connects to an ad server to retrieve the ad to be inserted in the live stream. The output stream generally carries the national ads, along with the cue messages, which can be replaced by local ads as and when required.

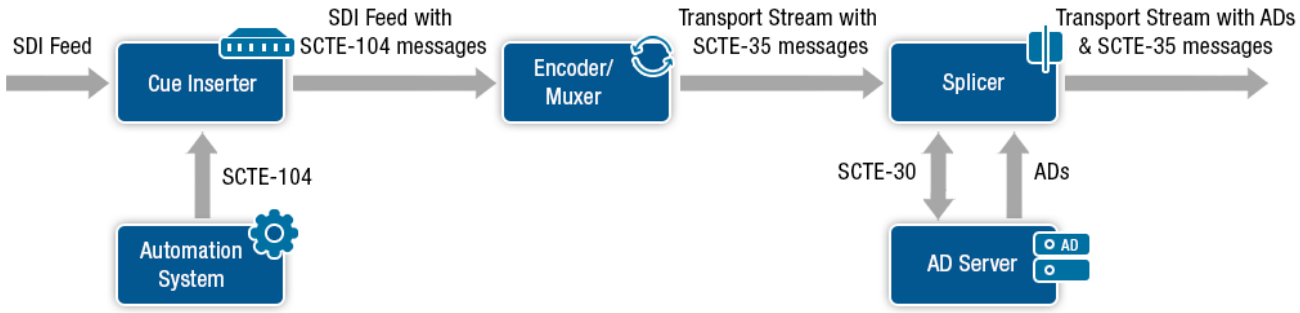

Figure 1. DAI linear streaming workflow

The live-linear stream with cue messages, as shown in Figure 2 below, is fed to a transcoder and packager that coverts it into HLS or MPEG-DASH specific ad markers in OTT manifests.

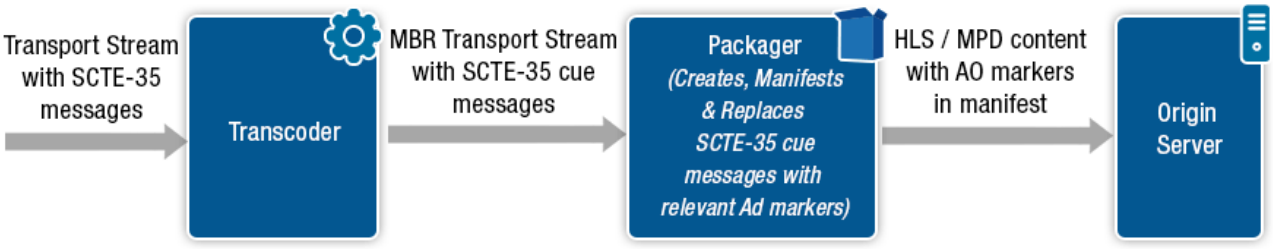

Figure 2. DAI live-linear workflow, transcoding and packaging phase

After receiving these ad markers, the player or the delivery platform connects to the Video Ad Serving Template (VAST) or Video Player Ad-serving Interface Definition (VPAID) compliant ad network, which provides demographically relevant ads that are inserted into the video stream. The IAB's VAST specification is a universal XML schema for serving digital video advertisements to players for in-stream advertising, and it describes expected player behavior when playing those ads. VPAID is a script that instructs a video player on what ad to play, the length of the ad, when to surface the ad, and where to place the actions (i.e., play, pause).

For VOD content, the broadcaster can choose the point of ad insertion. The ad insertion requires only manifest manipulation by adding ad markers and ad URLs to be played.

DAI in the OTT world

A quick word on OTT first. OTT uses Adaptive Bitrate (ABR) technique to deliver content in varying uncontrolled network conditions. ABR provides a method for delivering multiple video quality levels within a video stream. Multiple versions of a video are created, with each version encoded at a different bitrate and profile. Each version is further broken into short-duration segments, which are aligned with the same segment in other versions. Depending on the network bandwidth available on the consumer device, an appropriate segment from a specific file is sent to the user. This assures that the user receives the best quality video in an uninterrupted manner.

In order to manage these segments of content, a manifest (or playlist), is sent alongside the stream. The manifest contains URLs (pointers) to the video files contained within the stream. Additionally, the manifest contains metadata that identifies the bitrate, type of content (i.e., programming, advertising) being delivered to the client player.

In this OTT world, ad insertion is different. It can be done both at the server side as well as the client side. To enable DAI, an alternate content decision service (ACDS) parses the manifest metadata and provides an alternative manifest that identifies the targeted advertising, blackout, or geographic content to play out. In this way, targeted advertising can be delivered to client browsers and devices.

3. CLIENT SIDE AD INSERTION (CSAI)

A player like YouTube can display different types of ads during playback. Ads can be played before, during, or after a particular video is being played. Some of these ads can be skipped and some not. The advertiser has to pay when a viewer watches the ad for certain amount of time, for instance 30 seconds. Some ads are shown in a certain region of the display, above the video suggestion list. Lastly, some semi-transparent ads can be overlaid on the video being played. (See Figure 3.)

Figure 3. Types of client-side ads

The content creator/uploader is free to choose the type of ads to be displayed along with the content. The time of ad display and the ad to be shown is the choice of the player, based on viewers' interest, searches, and many other factors. Here, the player inserts the ad in the content, performing CSAI.

When content is played, the player requests the manifest file from the network and based on the available bandwidth, starts playing the relevant profile. The player also buffers the next few chunks to handle bandwidth fluctuations. Upon receiving an ad marker in the manifest, it sends a VAST request to the ad server. In response, the player receives the relevant ad to be played. (See Figure 4.)

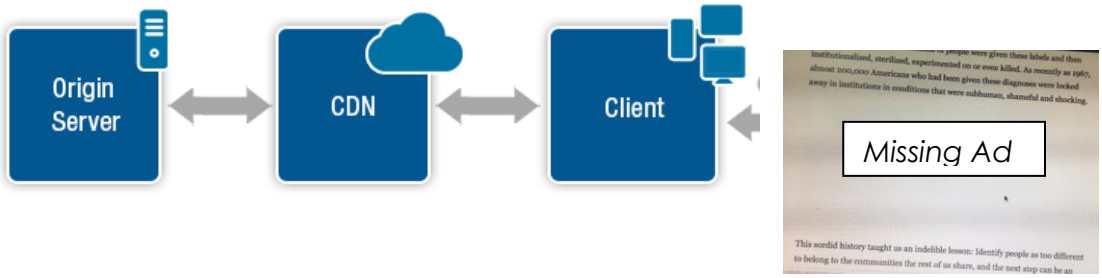

Figure 4. Client-side ad insertion workflow, highlighting interactivity between the CDN, origin server, ad server and client

A major advantage of CSAI is that it enables fairly accurate user-targeted ad insertion using client behavioral data from the device. CSAI also supports interactive and clickable ads.

While CSAI offers many advantages, it comes with some disadvantages as well. First, fetching ads from the ad server can lead to latency and buffering during playback if the ad serving system is slow or ad requests are redirected to another ad server.

This can be a major issue during live broadcasts like sports or news programs where abrupt ad breaks are frequent. It is very challenging to create an ad break and return to the live broadcast at a precise time. Second, ad blocker software can be installed on client devices

that can enable users to block or skip the advertisement, defeating the purpose of ad insertion.

4. SERVER SIDE AD INSERTION (SSAI)

SSAI is similar to linear TV advertisements where video content and ads are received in a single stream from same source. The ad is "stitched" into the video stream. SSAI relies on a video delivery network that performs manifest manipulation, communicates to the ad server (to fetch the ads) and ad transcoding (to convert the ad to the proper video content format). All of this happens on the server side.

When an ad marker is received, the splicer sends a VAST request to the ad server. In response, the splicer receives the ad in its base format (i.e., highest bitrate and quality). Then it needs to be transcoded into the format in which it is to be stitched. In a typical scenario, a splicer contacts the transcoder. If the transcoder already has a cached copy of the ad, it provides the URLs to the splicer. Only manifest manipulation is required to direct the client to the transcoded ad URLs. If the ad needs to be transcoded, the splicer serves another low priority, already transcoded ad to the client to avoid delays during playback. In the meantime, the transcoder prepares the ad so that whenever this ad is requested again, the transcoder already has cached copy available. (See Figure 5.)

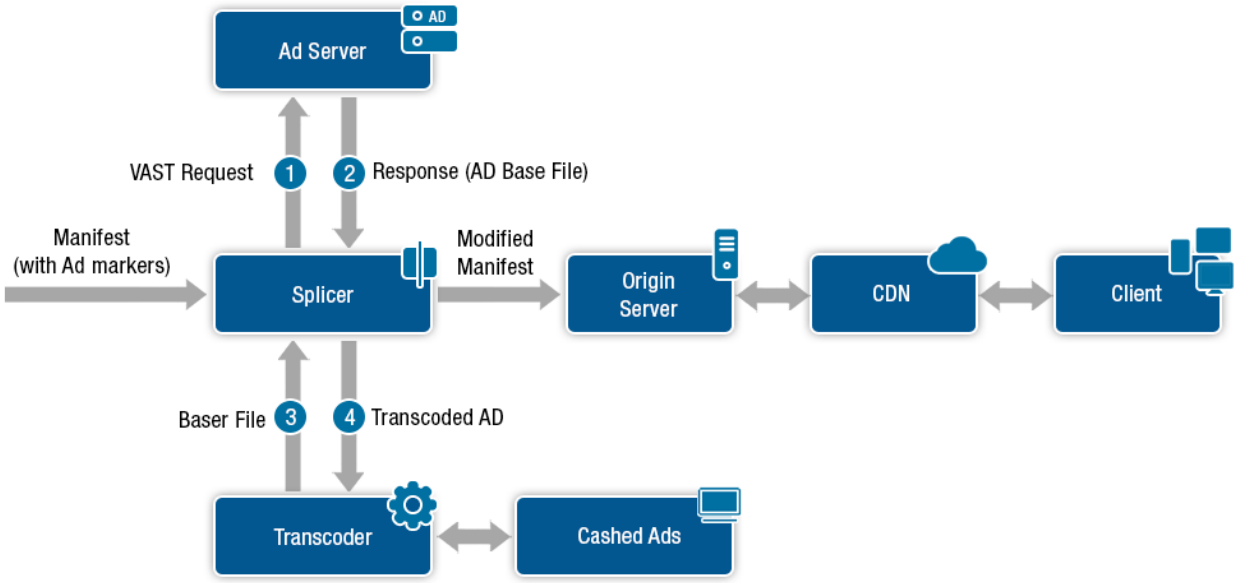

Figure 5. SSAI workflow

SSAI provides a seamless transition between the content and the ad, reducing latency during playback. It is a superior choice for live sports and news broadcasts where the duration of an ad break is not known in advance and a sudden switch is required. Moreover, video content and ads are muxed together, making it difficult for ad blockers to detect and block the ads.

5. CSAI VS SSAI

Recently, we've seen a shift toward SSAI. Still, SSAI is not the complete solution for DAI. Unlike CSAI, it does not support clickable, interactive ad insertion. Also, to provide user-relevant ads, it needs client metrics, which cannot be obtained from the server side. SSAI also requires a highly scalable architecture that can handle abrupt ad instances and creates manifests according to the inputs of the ad providers in real time. However, the stitching of ad at the server side decreases chances of delay or frame drop in SSAI. Server-side stitching also reduces probability of an ad blocker to block the ad. Both the technologies offer a unique approach to DAI, with distinct advantages and disadvantages.

6. DAI CHALLENGES IN AN OTT WORKFLOW

Compared with traditional broadcast, the OTT workflow is much more complex, as it involves delivery to multiple devices with multiple profiles and bitrates. The content coming from different media sources is re-encoded, transcoded and packaged into different profiles. DAI adds to the already complex OTT workflow. Like a traditional linear workflow, for example, in a live broadcast, the splicer needs to insert the ad and also return to the live content at a very precise moment. Unlike traditional workflows, this may need to be performed for multiple profiles at the same time for consumers watching the same content on different devices.

The following section talks more about the challenges involved with DAI in an OTT workflow.

- ### Multiplatform and multiscreen delivery issue

  OTT technology is fast growing and supports multiple formats like HTTP Live Streaming and MPEG-DASH in which the content can be delivered to consumers. Along with the content, the ads for insertion should also be available in different formats to support multiplatform playback. For this, ads need to comply to different standard formats. The Common Media Application Format (CMAF) is one way of tackling this problem. It aims to provide a single format for encoding and packaging segmented media objects for delivery and decoding on devices in adaptive multimedia presentations.

  Preparing and delivering the ads for multiscreen consumption is another cumbersome task. OTT services are also transcoded in different profiles to satisfy a wide range of viewers watching content on the go on mobile devices and on big smart TVs via a broadband connection. As for the content, each ad should also be transcoded into different profiles and available for insertion whenever an opportunity is found.

- ### Seamless ad insertion at segment boundaries

  Smooth playback requires each segment/fragment of content to be independently decodable so that whenever a switch from one profile to another takes place, no jerk in

playback is observed. For this, each segment, whether belonging to content or ad, should start with a key frame to ensure a smooth transition between content to ad and ad to content. Splicing can be a little trickier in the case of OTT, as the segments around the out and in points may not satisfy segment duration boundary conditions. (See Figure 6.)

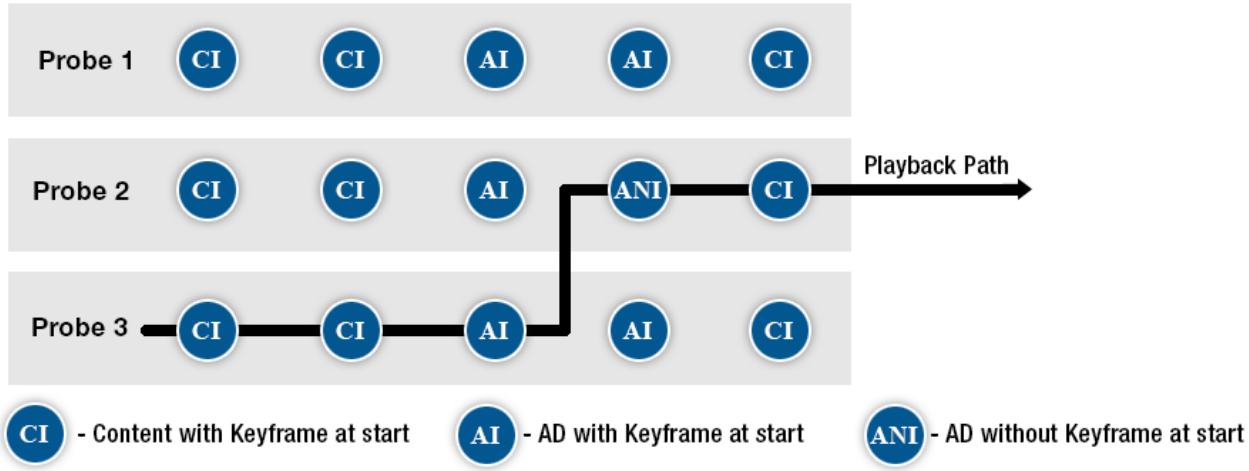

Figure 6. Seamless ad insertion example

Figure 6 shows one of the ad segments in profile 2 does not start with key frame. Either the key frame was not properly encoded or its key frame is in its previous segment due to improper segmentation. When playback path switches to profile 2 during playback of the ad, the player will not be able to properly decode and play the segment since the key frame is not found. A jerk in the playback will be observed.

- **Ensuring that the ad markers are aligned across all profiles/variants**

  The manifest ad markers must be properly aligned across all profiles. It is important to note that if an ad is about to start in one profile, the other profile should also be marked with the same ad start time. If this does not occur, and there is a switch to profile where the marker is not present, an ad playback opportunity is lost. Ensuring the alignment is quite challenging without the use of an advanced monitoring solution.

- **Consistency in ad insertion in an end-to-end workflow**

  In an end-to-end workflow, wherein a live broadcast containing SCTE-35 ad markers is being transcoded into different profiles for OTT delivery, the markers need to be converted to ad markers in manifest indicating ad placement opportunity. Ensuring that each and every ad marker in the linear streams gets correctly inserted at the right place in ABR outputs can be challenging.

- **Correct ad insertion at correct position**

  Correct ad insertion is another issue facing service providers. The ad shown to consumers must be the actual intended ad. Also, service providers need to ensure that the ads are inserted at the correct position (i.e., at logical breaks) without disrupting the continuity of the content.

7. CONTENT MONITORING: AN ANSWER TO THESE CHALLENGES

One of the biggest advantages of streaming is the potential to generate revenue from ads. In a stream, there are ad markers that tell you where to insert the ad in a stream, how long the duration should be, and other vital details.

Your monitoring system should be configured to look for these ad markers across all streams, ensuring you don't miss a single ad insertion opportunity. The monitoring system should alert you if no ads were present in a specified time duration. For example, if you expect to run one ad every 20 minutes and it doesn't do that, then that's a problem. Beyond this, you need to do a sync check for ads to make sure ads start at the same point across all variants. This is something that content providers struggle with. Hence, it should be an important part of your monitoring strategy:

- Check whether ads are inserted consistently (time and position) across all variants of a channel
- Check whether ads are inserted at the right place and at the right time

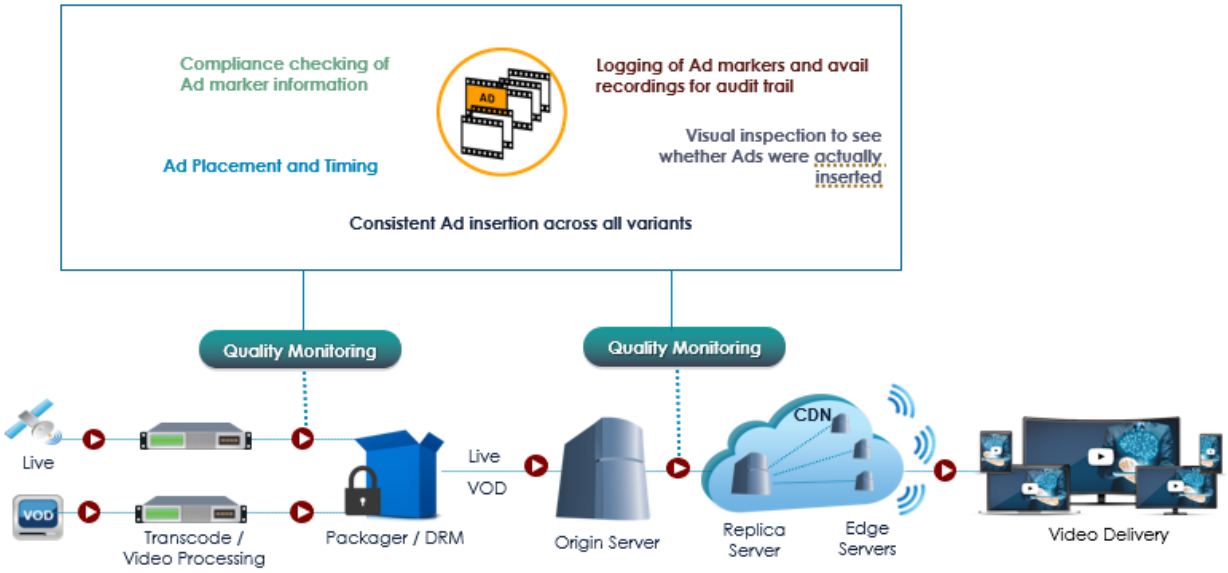

The aim of monitoring OTT content is to guarantee good quality of service and quality of experience to the consumer. A monitoring system running 24/7 is able to monitor the complete end-to-end workflow starting with linear streams and extending to content being watched by end consumers in any format. It is able to provide the correct, meaningful alarms in real time along with the location of where the issue occurred in the chain. This helps service providers to quickly resolve the issue before the end customer is affected. It is essential to monitor for all possible issues like over-compression, frame misalignment across profiles, syntax errors, and compliance slips, as these issues can negatively impact overall user experience. Monitoring the transition of ad markers from linear to OTT delivery in an end-to-end workflow will ensure all ad opportunities are properly utilized. A few crucial points that require monitoring are ad markers being properly aligned and corresponding to ads of same duration across all profiles, as well as ads starting with key frames. Monitoring content will give a competitive edge to service providers and broadcasters vs. those service providers that do not monitor content.

8. INTERRA SYSTEMS' SOLUTION

Interra Systems' ORION-OTT is a software-based OTT monitoring solution that enables you to monitor Adaptive Bitrate (ABR) content integrity and related network performance in an OTT environment.

Leveraging industry-proven audio and video quality analysis technologies, ORION-OTT enables OTT service and equipment providers, as well as content delivery network (CDN) providers, to seamlessly monitor online video delivery for quality and compliance of mainstream content, along with ads.

ORION-OTT monitors ABR content, both VOD and live, leveraging user-defined automated or manual schedules. It checks for inconsistencies pertaining to ABR package compliance, manifest and playlist syntax, download errors, content quality, and more.
Specific to ad monitoring, it offers the following monitoring checks:

- The presence of respective ad markers in MPD or HLS playlists and out time, in time, ad duration, etc.
- Inconsistent ad start segments across variants
- Duration of ads
- Minimum and maximum duration of content between ads
- Download failure or delay for ad segments
- Stale manifests in ad slots
- ABR compliance issues in ad chunks
- Audio and video quality issues within an ad stream

9. CONCLUSION

The ever-evolving and fast-growing nature of OTT technology and its viewership has attracted many players to enter the video streaming domain. For broadcasters and video content providers, the challenge is not only limited to delivering a seamless OTT experience with good quality. Another major challenge is monetizing the streaming service. The quality can be ensured by implementing content monitoring at each phase of content transformation. Revenue generation requires advertisers to provide the relevant ads to individual customers in non-disruptive way.

DAI provides a solution to dynamically insert user relevant ads with seamless playback. The two approaches CSAI and SSAI have their own pros and cons, which can be debated. A hybrid solution uses SSAI technology and client-side SDK to gather user metrics and determine user interest in ads. It can also integrate some custom SDKs in player for providing interactive and clickable ad services. Broadcasters and service providers are currently exploring this method of ad insertion.

DAI is challenging in terms of assuring correctness and seamless ad insertion at a precise time. It also requires ads to comply with different OTT content formats and profiles. A good end-to-end workflow monitoring solution is critical.

It will be interesting to see if broadcasters can come up with a solution that uses the best of both CSAI and SSAI technologies to generate revenue from the vast and expanding OTT domain while keeping the content quality high. Maybe we are heading toward something new and a different type of DAI. Nevertheless, DAI is here to stay.

Seamless ad insertion and assuring quality on the inserted ads help with better QoE and monetization. Ad monitoring should be part of OTT video service providers' core strategy.