



**Digital Video Broadcasting (DVB);
Dynamic substitution of content in
linear broadcast: interfacing an advert
decisioning service and optimal
preparation**

DVB Document A178-1

November 2019



This page is left intentionally blank

Contents

Interfacing to an advert decisioning service and optimal preparation of media	1
1 Scope	5
2 References	5
2.1 Normative references	5
2.2 Informative references	5
3 Definition of terms, symbols and abbreviations	6
3.1 Terms	6
3.2 Symbols	7
3.3 Abbreviations	7
4 Background	7
5 Request, Response & Measurement	11
5.1 Requesting an advert	11
5.2 Targeting an Advert	11
5.3 Direct response	11
5.4 Nested response	12
5.5 Tracking an advert	12
5.6 Substituting adverts in broadcast	12
5.7 Extending the response capability	14
6 Advert Encoding and Delivery	14
6.1 Advert encoding	14
6.1.1 Format for receiver	14
6.1.2 Format from Advert Producer	15
6.2 Advert Delivery	15
6.2.1 Secure Advert Delivery	15
7 Trade-off between reach and quality	16
7.1 DAS app approach to reach/quality trade off	17
7.2 Advert Server approach to reach/quality trade off	17
8 General Guidance for broadcast streams and receivers	18
8.1 Broadcaster Responsibility for Advert Substitution Preparation	18
8.2 Broadcast Stream Conditioning	18
8.3 Receivers with Dual Decode capability	18
8.4 Receivers with Single Decoders and Precise Switching Timing	18
8.5 Receivers without Precise Switching Timing	19
9 Annex: Change History	20
10 History	21

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (<https://ipr.etsi.org>).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the [ETSI Drafting Rules](#) (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Executive summary

Introduction

For broadcasters who wish to dynamically substitute advertising in a linear broadcast or for platform operators who wish to enable the functionality for broadcasters, this document provides guidance relating to the interface with prevailing digital advert decisioning systems and advice on the preparation of media for an optimised viewer experience. The techniques described in this document may also be applied to dynamic substitution of programme content.

1 Scope

This document provides guidance and recommended practice for delivering and measuring targeted adverts in horizontal and vertical market deployments using in-market advert technology. The focus of this document pertains to the context of linear broadcast and the substitution with IP delivered adverts. Inasmuch as information on video and audio encoding/packaging or receiver implementations impacts the advert media served, it is discussed. In-depth technical audio/visual guidance is not given, neither is in-depth information on receiver implementations.

2 References

2.1 Normative references

Normative references are not applicable in the present document.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] IAB VAST - 4,1 – Video Ad Serving Template

NOTE: available at: <https://iabtechlab.com/wp-content/uploads/2018/11/VAST4.1-final-Nov-8-2018.pdf>

[i.2] IAB VMAP Video Multiple Ad Playlist 1.0

NOTE: available at: https://www.iab.com/wp-content/uploads/2015/06/VMAPv1_0.pdf

[i.3] TS 103 285 Digital Video Broadcasting (DVB); MPEG-DASH Profile for Transport of ISO BMFF Based DVB Services over IP Based Networks

NOTE: available at: https://www.dvb.org/resources/public/standards/a168_dvb_mpeg-dash_oct_2019.pdf

[i.4] TS 101 154 Digital Video Broadcasting (DVB); Specification for the use of Video and Audio Coding in Broadcast and Broadband Applications

NOTE: available at: https://www.etsi.org/deliver/etsi_ts/101100_101199/101154/02.04.01_60/ts_101154v020401p.pdf

[i.5] IETF RFC6381 The 'Codecs' and 'Profiles' Parameters for "Bucket" Media Types

NOTE: available at: <https://tools.ietf.org/html/rfc6381>

[i.6] W3C MSE Media Source Extensions

NOTE: available at: <https://www.w3.org/TR/media-source/>

[i.7] EBU Recommendation R128 Loudness normalisation and permitted maximum level of audio signals

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

Advert: An audio-visual presentation aimed at communicating the benefits of a product or brand.

Advert media: A particular combination of audio-visual encoding and packaging used by a receiver to render an advert. An advert response may reference several versions of advert media.

Advert producer: The organisation responsible for supplying source media for an advertising campaign to an Advert Server. Often this organisation is a creative agency.

Advert request: A resource locator, which when resolved communicates the context of a placement opportunity to an advert server.

Advert response: Delivered in reply to an advert request, advert response encapsulates all information related to a particular instance of an advert e.g. identifiers and resource locators for the advert's media and tracking.

Advert server: A trusted web service that is used to decision advert opportunities. It receives advert requests and returns advert responses.

Break: An inter-programme or intra-programme gap provided in a linear broadcast stream to present one or more adverts.

Broadcaster: An entity responsible for compiling and disseminating audio-visual programme content as a TV channel on a linear broadcast stream.

Bumper: A short audio-visual presentation, inserted between programme and break for the purpose of visually separating commercial content from programming content. It usually takes the form of promoting the broadcaster's channel brand and may be present going into a break and again coming out of it. Also referred to as a 'sting', 'ident', 'opener' or 'closer'.

DAS app: Software application to perform DAS on the receiver via the API exposed by the receiver's capabilities.

Dynamic advert substitution (DAS): The operation by which a certain broadcast advert is substituted by a targeted advert (delivered over Internet Protocol in most cases), chosen specifically for a given individual receiver. The substitution is managed by an advert server, which will effectuate the DAS operation in compliance with the business rules set by the broadcaster or platform operator.

Frame-accurate advert substitution: Where the act of substituting a linear advert is perfectly spliced into a linear broadcast stream to exacting broadcast standards.

GDPR: General Data Protection Regulation, a single set of rules by which all companies operating within the European Union are required to abide by.

Horizontal Deployment: Where one or more broadcasters output TV channels independently of each other and this is done independently of the suppliers of the receiver hardware needed to view the TV channels.

Macro substitution: A method available to advert servers enabling run-time substitution of variables, typically URL parameters, within an advert response.

Network Operator: An entity responsible for the distribution of TV channels as linear broadcast streams.

Platform Operator: An entity responsible for packaging and distributing TV channels as linear broadcast streams as well as the hardware receivers needed to view them.

Placement opportunity: The point in time in a broadcast stream at which a substitute advert could be inserted.

Receiver: A device capable of receiving and rendering the audio-visual content of a linear broadcast stream, usually a television or set top box.

Seamless advert substitution: Where the act of substituting a linear broadcast advert has no discernible visual or audible degradation to the viewer i.e. the viewer, without a comparative reference cannot notice that the advert has been substituted. This differs from frame-accurate substitution in that frames of black or freeze frames may have been introduced to visually improve a transition to the substitute advert.

Spot: A time-bound section of a break dedicated to a single advert.

Substitute advert: The advert that is presented instead of the underlying broadcast advert.

Substituted advert: The underlying broadcast advert that is omitted for the substitute advert.

Time shift: A mode of viewing a linear broadcast stream so that it is presented to a viewer delayed from the live broadcast.

Tracking: A resource locator that is resolved at run-time in order to count and measure advert delivery.

Trick play: The ability to pause, rewind or fast-forward a linear broadcast stream.

Trusted advert gateway: Server that provides receivers with substitution adverts based on decisions made elsewhere, but may provide restrictions or constraints for the decisioning process.

Vertical Deployment: Where the end-to-end chain of broadcast-to-receiver infrastructure is the responsibility of a sole entity, e.g. A platform operator. The packages of TV channels are broadcast by the sole-entity and the receiver hardware and/or software needed to view the packages are supplied into the market by the same entity.

Targeted advertising: The substitution of an advert in a linear broadcast with an advert chosen specifically for a receiver based on viewer or viewing context.

3.2 Symbols

3.3 Abbreviations

4 Background

Linear broadcast television has been established over many years as a reliable marketing channel for delivering brand awareness to a mass audience. Mass market penetration of internet connected TVs presents the opportunity for broadcast TV to support features commonly available in digital advertising.

In response to this, there is a desire from broadcasters and platform operators to evolve the capabilities available in a connected linear broadcast TV environment to include features commonly found in online/digital advertising as this is increasingly required by advertisers and agencies.

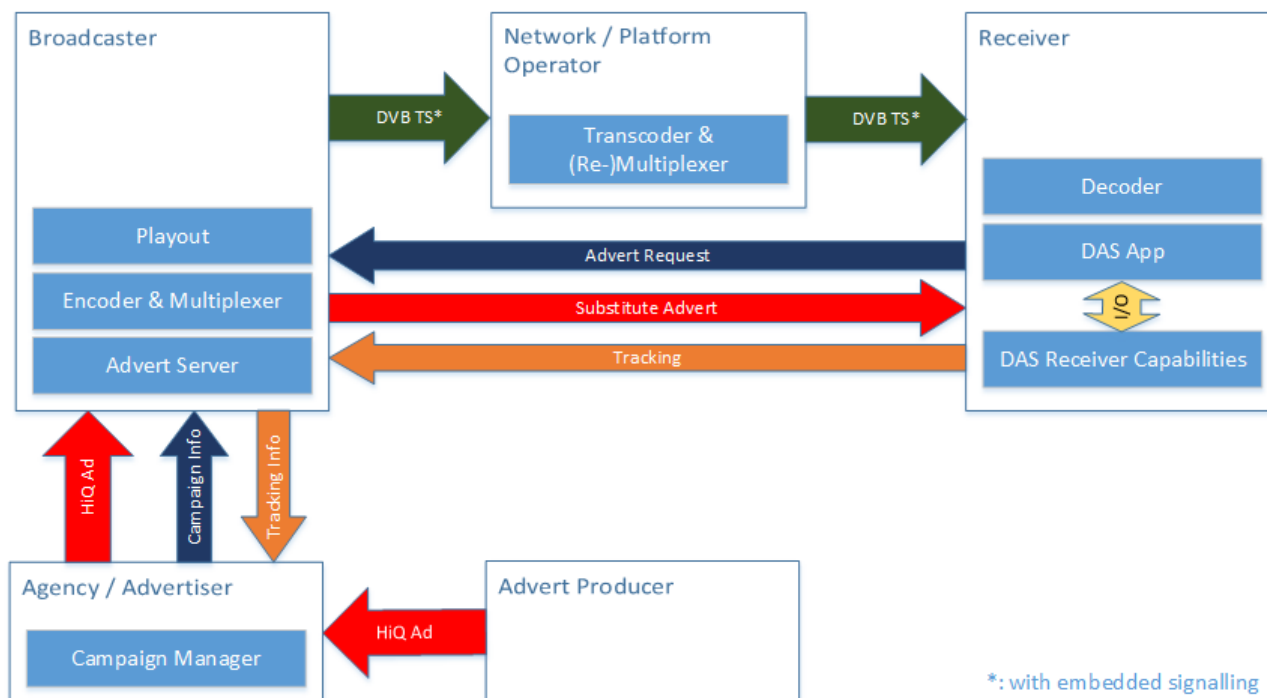


Figure 1: Logical components of the DVB-TA system

Figure 1 above illustrates the logical components found in the technology framework proposed by DVB to enable Targeted Advertising (DVB-TA), which means the substitution of an advert from a linear broadcast with an advert chosen by an advert decisioning service based on criteria such as viewer profile, viewing behaviour, environmental or contextual factors.

The DVB-TA technology framework addresses the following four areas:

- Signalling
- Seamless advert Splicing at the Connected Receiver
- Measurement & Reporting
- Integration with Existing Advert-Decisioning Systems

This document provides guidance for points 3 and 4 and gives advice for the preparation of streams and media to aid point 2.

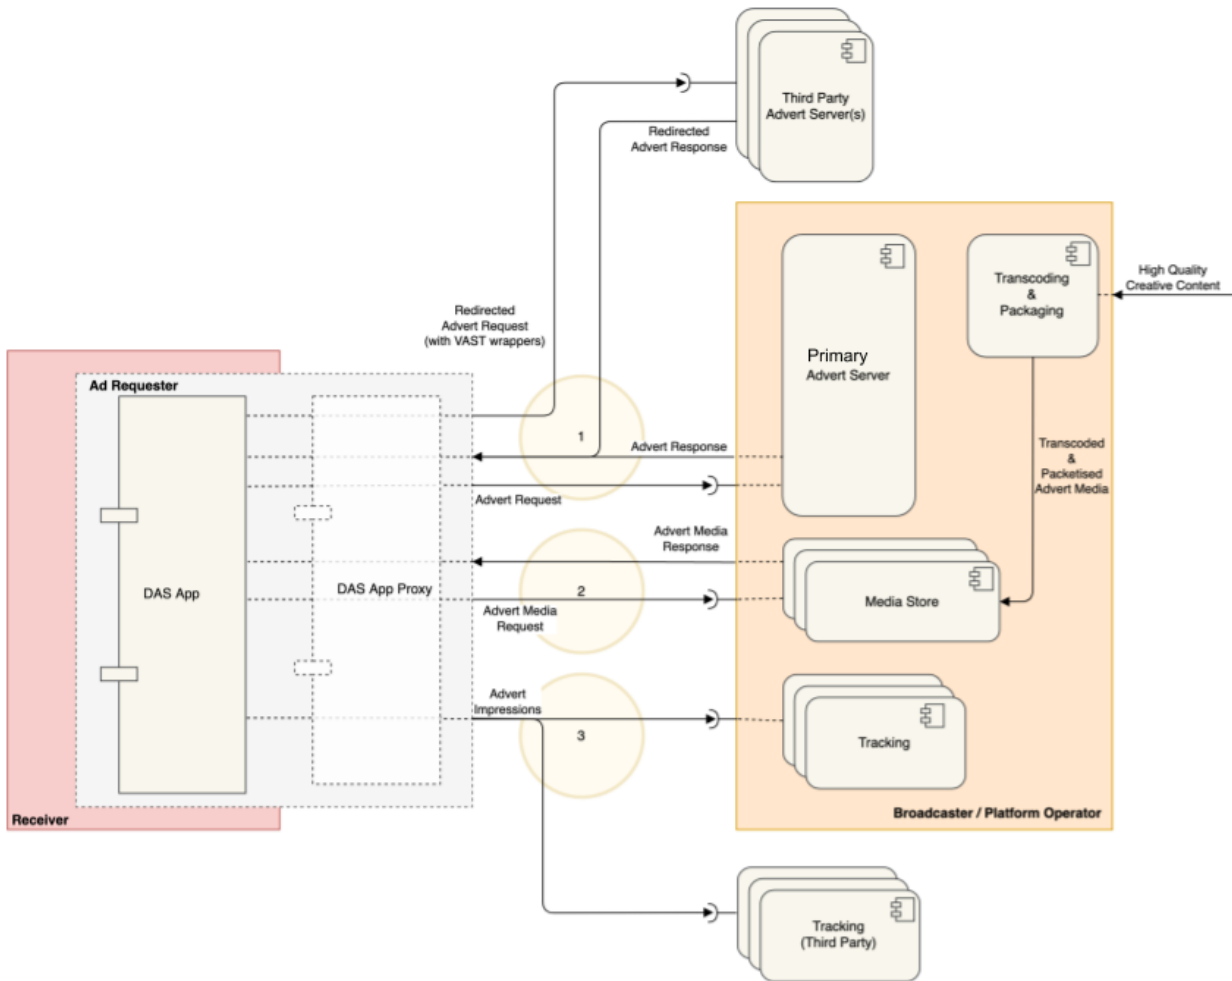


Figure 2: A reference architecture for DVB-TA interfaces considered in this document

Figure 2 above shows an example of how a typical integration between a DAS app and existing advert system might look. Interfaces 1, 2 and 3 are discussed in this document. Interface 1 and 3 are discussed in the Request, Response & Measurement section, and interface 2 is discussed in the Advert Encoding and Delivery Section.

The following sequence diagram (Figure 3) illustrates the typical component interactions present during dynamic substitution of a broadcast advert.

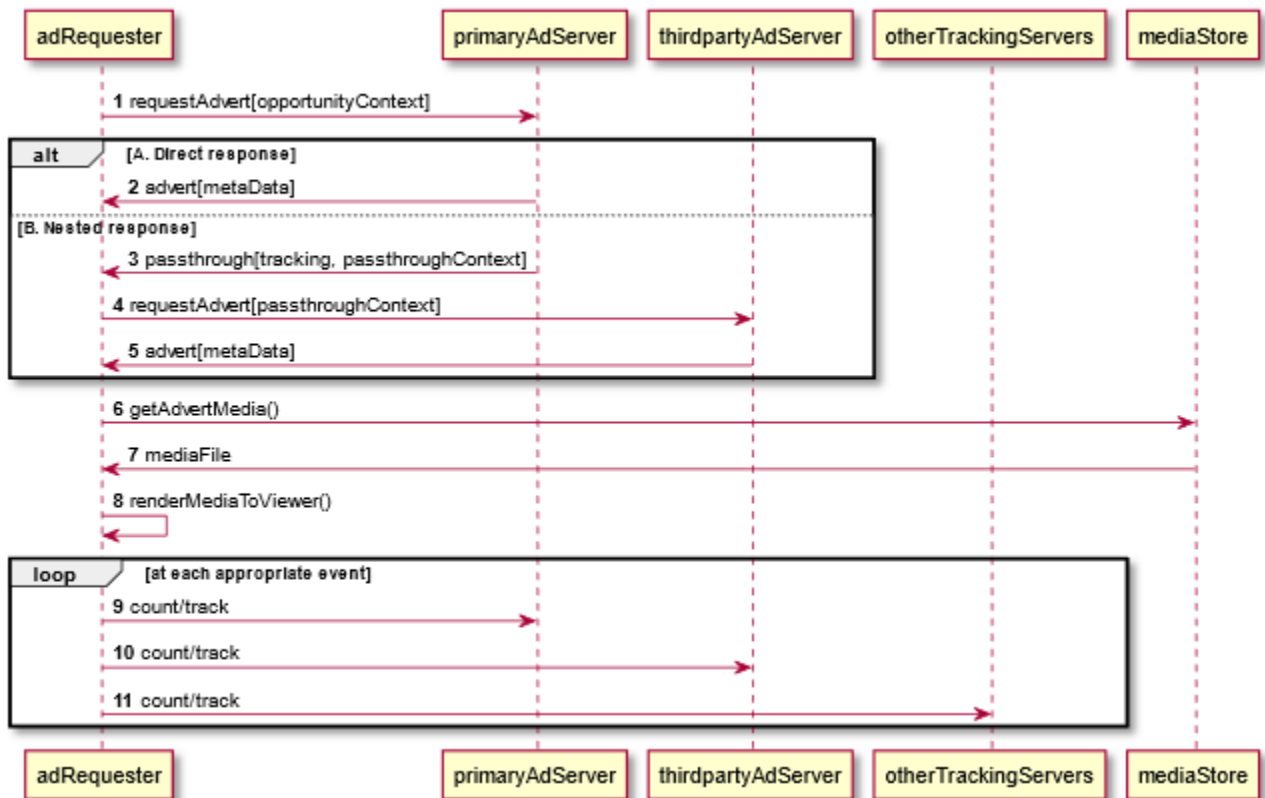


Figure 3: Typical interactions for an adRequester during dynamic advert substitution

The actors in Figure 3 are defined as:

adRequester: a system actor that is responsible for requesting an advert, interpreting a returned advert response and displaying the advert to a viewer. The adRequester is typically a DAS application hosted on a DVB-TA receiver, but could equally be hosted in a platform's head-end or other proxy-server (e.g. trusted advert gateway) acting on behalf of a DVB-TA receiver.

primaryAdServer: a system actor that initially decisions an advert request and returns one or more adverts in response to that request. The primaryAdServer is typically owned by the broadcaster, but could equally be owned by a platform operator or another third-party e.g. an advert exchange.

thirdpartyAdServer: primaryAdServer is able to defer an advert decision to a thirdPartyAdServer. The thirdpartyAdServer in turn is able to defer an advert decision to subsequent thirdPartyAdservers. thirdpartyAdServer may belong to any group separate from the owner of the primaryAdServer, but is generally a separate sales house or advert exchange.

otherTrackingServers: otherTrackingServers is a system actor that represent one or more servers that do not necessarily decision an advert request, but may receive signals during playback of an advert. These are often owned by advertising agencies and allow them to track advert delivery independently of the decisioning advert servers.

mediaStore: mediaStore is a system actor that hosts media files for a particular advert. It is typically a CDN (content delivery network), but could equally be storage in a platform's head-end or even storage on a viewer's DVB-TA receiver.

primaryAdServer, **thirdpartyAdServer**, **otherTrackingServers** and **mediaStore** actors together represent the Advert Server shown in Figure 1.

5 Request, Response & Measurement

The digital advertising industry has established VAST [i.1], as the standard response format from an advert decisioning component. This standard is used by both primary and secondary advert servers. DVB-TA recommends its use for all advert decision responses.

VAST supports several features, which are summarised below.

5.1 Requesting an advert

In action 1 of Figure 3 the adRequester issues an advert request to the primaryAdServer so that it can decide which advert is suitable for substituting the broadcast advert. An advert request typically conveys the context of the placement opportunity that is to be substituted. This may include, but is not limited to programme, receiver and user context.

The method of submission and format of this message is specific to the primaryAdServer and context that the broadcaster or platform would like decisioned. It usually consists of key-value pairs describing the opportunity and is submitted via HTTPS GET request, but could just as well be HTTPS POST or PUT.

DVB-TA recommends using HTTPS as a method for advert requests so that the content of messages between adRequester, primaryAdServer and all subsequent thirdPartyAdServers are privy to the participants.

DVB-TA does not specify a mechanism for the submission or formatting of an advert request. These are implementation specific to a broadcaster/operator's internal business rules. However, to increase interoperability with downstream advert servers, DVB-TA recommends that macro substitution defined in reference [i.1] 'Macros' be supported by adRequester.

5.2 Targeting an Advert

The primary benefit of DVB-TA over existing broadcast advertising is the ability to target one or more adverts with finer granularity than normally possible. Targeting at a geographical level, while regularly used within broadcast advertising in some markets, can be achieved with DVB-TA without the need to transmit differing broadcast streams at the regional level. However, device or even viewer level targeting is also possible to meet advertiser demand. In order to achieve this, unique identifiers for device or viewer would need to be generated. How device or viewer identifiers are derived is outside of the scope of this document, but they would need to be present in the adRequester in order to transmit them with the advert request. Broadcasters or other DAS app providers are reminded that often national and/or international legislation (e.g. GDPR) applies to the use of identifiers and implementations should take into account such legislation.

5.3 Direct response

In action 2 of Figure 3 primaryAdServer has decided what advert to display and returns a message to adRequester. This message contains all meta-data necessary for adRequester to locate the

advert's media and later issue tracking counts and respond to any viewer interaction that may be permitted.

Reference [i.1] 'VAST Implementation - <Inline>' provides a comprehensive framework for structuring this information in an interoperable manner.

Note: Not all of VAST [i.1] will be relevant to a broadcaster's or platform's use-case so only a subset of the standard need be implemented by adRequester.

5.4 Nested response

In action 3 of Figure 3 primaryAdServer has decided to pass the substitution decision to a downstream advert server. It does this by returning a passthrough response to adRequester. This response contains the advert request for the downstream as well as any additional tracking that primaryAdServer would like to add to the final decided advert.

VAST [i.1] 'VAST Implementation - <Wrapper>' provides the format to use for passthrough responses.

Actions 4 and 5 of Figure 1 are similar to actions 1 and 2, except that the advert request of action 4 was extracted from the primaryAdServer and the destination advert server is a downstream thirdPartyAdServer. The final response in the chain takes on the same format as action 2 i.e. it will at minimum contain information on where to retrieve advert media. adRequester will be responsible for recording the tracking information that was collected in upstream responses. Too many levels of nesting may have a detrimental effect on the timeliness of the ad substitution. DVB-TA recommends no greater than 5 levels.

5.5 Tracking an advert

Actions 9 to 11 of Figure 3 are triggered by adRequester at appropriate points in playback or in response to viewer interaction.

adRequester will have extracted and accumulated the tracking end-points from actions 2, 3 and 5 of Figure 3. The recipients of these tracking events are typically primaryAdServer, thirdPartyAdServer and otherTrackingServers. Each of whom have a stake in validating the performance of the advert.

Tracking URLs can be HTTP and/or HTTPS. Using HTTPS method will provide benefits as described in the advert request and responses.

5.6 Substituting adverts in broadcast

A broadcaster or platform operator may decide to substitute an entire commercial break or individual adverts in a commercial break. It is important to note that some receivers may not support the replacement of consecutive adverts or may only support this scenario given certain advert media so it is recommended that adRequester be designed to communicate the target receiver capability relating to support of consecutive advert substitution so that primaryAdServer can decide to abandon the substitution or return links to appropriately tailored media.

The digital advertising industry refers to commercial breaks (consecutive adverts) as 'Ad Pods' and the VAST specification caters for both types of substitution under the section entitled 'Ad Pods and

Stand-Alone Ads'. Note: Several deployments within the industry may make use of IAB VMAP [i.2] to describe the advertising content of a break. However, the recommended approach for DVB-TA is to describe breaks as defined by IAB VAST [i.1].

After actions 2 or 5 of figure 3, adRequester is presented with one or more media locations within the response's <MediaFiles> element. If several media locations are presented, they may represent the same file from diverse locations or the same visual content encoded and packaged in multiple formats e.g. combinations of bitrates, resolutions, encoding profile or packaging. Each <MediaFile> element will contain attributes describing its format.

adRequester can either choose the most appropriate <MediaFile> from the list or if sufficient context was passed to primaryAdServer in action 1, it may have already narrowed down the list to only contain suitable media locations in the response. This theme is further described in section 7, Trade-off between Reach and Quality.

As noted above, receivers in the target population may vary considerably in capability and viewing configuration (as determined by viewer settings).

Where possible, broadcaster or platform operator should take steps to preserve the in-home viewing experience during advertising substitution.

Where a viewer is consuming programme content in a specific viewing mode e.g. with a specified language track or parental lock activated, the broadcaster/operator may wish to ensure that substituted adverts match the particulars of the viewing mode, e.g. only adverts that support a particular language track or within the correct parental lock tier respectively are returned.

Similarly, when a viewer has activated accessibility features e.g. subtitles, audio description or signing, the broadcaster or platform operator should ensure that, if available, substitute adverts match the active accessibility features, e.g. only adverts that support subtitles, audio description or signing respectively are returned.

To enable this, advert requests can be authored to signal information pertaining to the viewer's current accessibility and other settings so that all advert servers in the decision chain are able to decide on suitable adverts. There is a small risk that the viewer may change their settings between the time of an advert request and the showing of a replacement ad. The DAS app should be authored to allow for this edge case.

Matching substitute adverts to programme content in a specific viewing mode or accessibility settings may be subject to varying degrees of regulation depending on geographic market.

Regulatory and legal compliance of advert media is out of scope for this document. It is typically managed by market specific broadcaster or platform operator processes.

Capabilities relating to trick play may also vary in the receiver target population. Where possible, a DAS App should query the receiver to determine its trick-play capability. The capability level can then be signalled to primaryAdServer, which can in turn provide appropriate media location in the advert response. adRequester can then decide to enforce a viewer experience based on any of these business rules:

- Match receiver capability: substitute advert is played back to match the viewing capability of the underlying broadcast advert e.g. on a non-PVR receiver, trick-play is disabled for the replacement advert. On a PVR receiver, pause, rewind and fast-forward are enabled for the replacement advert when viewed in a time-shift

scenario. Note: time-shifting of replacement adverts requires careful management as undesired effects may result when re-joining the broadcast programme.

- Disable on all receivers: adRequester disables trick-play on playback of all substitute adverts.
- Enforce dynamically: adRequester acts dynamically upon the advert response element <skipoffset> (see VAST [i.1]).

5.7 Extending the response capability

On occasion, a broadcaster or platform operator may need to support functionality unavailable to the VAST specification or specific to an internal use-case. Provision is made for this in the <Extension> element of VAST [i.1].

Some broadcasters/platform operators may have the capability to substitute broadcast adverts with interactive adverts. Interactive adverts allow a viewer to respond to a call to action e.g. a ‘Press Red Button’ prompt, to access an enriched advertising experience. The enriched experience might provide a viewer with a choice to view a variety of styles/models/previews relating to an advertiser’s product.

The implementation of interactive adverts is out of scope for this document, but where needed, broadcasters and platform operators should make use of ‘VAST Interactive Templates’ in VAST [i.1] to meet their requirements.

6 Advert Encoding and Delivery

6.1 Advert encoding

6.1.1 Format for receiver

It is recommended that the video and audio for substituted adverts be encoded according to the capabilities of the target receiver population, and, for the relevant formats, in conformance with DVB video and audio encoding and packaging specifications, i.e. TS 103 285 for fragmented ISO BMFF [i.3] and TS 101 154 [i.4] for Transport Stream formats respectively. Non-fragmented ISO BMFF (often referred to as MP4) files are commonly used for on-line video advert substitution, and although not specified for DVB receivers, may be used if supported by the target receiver population, and profiled for targeted advertising elsewhere.

The advert server may offer more than one <MediaFile> for an advert to the adRequester such that a suitable version may be selected according to receiver and/or network capabilities which may be unknown to the advert server. Each version should have a technical quality similar to that of the broadcast service being replaced, e.g. for the video, maintaining the same resolution, and a similar level of coding artefacts. If the same codec is being used for broadcast and advert, then the encoded bit-rate would need to be similar between broadcast and advert (although the advert encoding may benefit from non-real-time techniques allowing a higher level of compression for the same quality compared with a real-time broadcast encoding). If a later generation encoding standard is being used for the advert than for the broadcast, an advert with matching technical quality is likely to have a lower bit-rate than the broadcast.

If the advert is being streamed using Adaptive Bit-Rate (i.e. DVB-DASH [i.3] according to TS 103 285), then even the lowest bit-rate representation should provide an acceptable quality in comparison with the broadcast service.

Depending on the receiver capability, if the advert is encoded with the same codec as the broadcast, then there may be advantages in the quality of the transition between broadcast and advert. The advert server should therefore, where possible, include a <MediaFile> with the same codecs as the broadcast. However, it is noted that according to TS 101 154[i.4], standard-definition broadcasts may use MPEG-2 video and MPEG-1layer 2 audio, but the receiver might not support IP delivery of content in these formats, and TS 103 285[i.3] does not specify these codecs for packaging with fragmented ISO BMFF. In this case, any potential advantage in the quality of transitions between broadcast and advert by using the same codecs for broadcast and adverts cannot be realised for these receivers on broadcasts using these legacy codecs. See also section 7 on Trade-off between reach and quality for a discussion on approaches to dealing with this situation.

6.1.2 Format from Advert Producer

The encoding and packaging of the advert from the advert producer, (shown as the “HiQ Ad” arrow in Figure 1), is out of scope of the current specification, but the following aspects of the advert media should be considered to enable a functional system.

The technical quality of the advert needs to match or exceed that required for the advert to be delivered to receivers.

- Ad media components, such as multiple languages, surround sound, audio description, subtitles, may be required by the broadcaster for insertion into broadcast content that includes such components.
- The audio sound level may need to conform with national regulations or codes of practice so that its subjective loudness does not exceed that of the broadcast content. EBU Recommendation R128 [1.7] provides information on normalisation of the loudness of broadcast audio signals.

6.2 Advert Delivery

Reliable delivery of advert media to receivers is critical to perform dynamic advert substitutions. It is essential to have advert media ready for consumption at the point where dynamic advert substitution starts regardless of the method of delivery.

A broadcaster’s selection of file and packaging formats for advert delivery will depend on a number of factors, including differences in implementation across market deployments, differing thresholds in advert substitution performance requirements and potentially even distinct commercial or operational practices. The overall design of the DVB-TA system defines the use of adaptive streaming with DVB-DASH [i.3], non-adaptive streaming (e.g. with TS 101 154[i.4] for Transport Stream) and simple downloading of the entire advert into memory in advance of advert substitution.

6.2.1 Secure Advert Delivery

Delivery of the ads can be achieved via HTTP or HTTPS depending on the deployment market use cases as well as authentication and authorisation requirements. HTTPS which adds SSL layer on top of HTTP provides a secure connection between the adRequester and the mediaStore.

This secure connection is recommended as it provides verification that the adRequester is talking directly to the mediaStore intended by the broadcaster or platform operator. Using the HTTPS method also ensures that only the mediaStore can read the advert request and only adRequester can read the returned response.

In most vertical and horizontal deployments adverts are not encrypted. However, in some cases adverts may be encrypted and the encryption scheme for the adverts could differ from that of the broadcast stream. To prevent preloading media files which the receiver is not capable of playing back due to encryption, the advert server should be aware of the receiver capabilities to be able to make the best decision at the time of serving the adverts. It is not recommended for advert servers to respond with media file URLs of both encrypted and unencrypted version of the media files in an advert response when both versions of the same media file available to serve. Using encrypted delivery methods and/or encrypted media files can introduce delays in advert substitutions.

7 Trade-off between reach and quality

DVB-TA targets both vertical and horizontal market deployments and could be implemented on a wide range of receivers with varying levels of hardware capabilities. As such it is a fundamental principle of DVB-TA that broadcasters are given the opportunity to trade off reach and quality.

Reach is defined as the total number of receivers a broadcaster can expect to address with targeted advertising, whereas quality is defined as the accuracy and duration of the switch to a substituted advert and back.

Given some receivers will not be able to deliver consistently seamless advert substitutions, there are mechanisms which allow a broadcaster to determine at runtime whether a receiver or an individual advert placement opportunity is capable of a seamless substitution. Some broadcasters may only want to reach users where seamless substitution is achieved, while others may want to reach the largest possible number of users, even if this could mean some users experiencing noticeable transitions between broadcast content and adverts.

In order to avoid additional costs and/or negative impact to inventory management of unnecessary advert decision or media calls broadcasters seek the ability to omit receivers they deem not capable of meeting their requirements. Hence, receivers which are capable of meeting the broadcaster quality requirements in principle, should make sufficient information available to minimise any possibility of advert substitution playback issues.

Information which could influence the success of the advert substitution may be known up-front if the particular broadcast service is identified by the adRequester when making an advert request. This includes information such as broadcast stream resolution, encoding and encryption. Further information which will factor into the performance of advert substitution is specific to the receiver design (such as memory, number of simultaneous decoders available) or to the context at the point of transition (such as resources availability and management, e.g. a second decoder is busy with picture-in-picture), or the memory currently available to pre-load adverts.

In vertical market deployments available memory may well be known at the point of advert substitution, but such information is less likely to be determinable in horizontal market receivers, not least given the large number of system and user-initiated background activities which may require memory at any time. While it is acknowledged that certain browser technologies (such as Media Source Extensions [i.5]) may provide helpful approximations to establishing available memory, it is recommended that logic within the DAS app should not rely solely upon knowing available memory at the point of advert substitution for horizontal market deployments.

While it is a requirement of DVB-TA that one advert can be fully pre-loaded, a broadcaster can also stream advert media rather than fully preload it. Streaming may be a useful approach if longer

periods of content are to be substituted. This may be helpful for a receiver identified as not meeting a broadcaster's requirements for advert substitution transitions between individual adverts. A full advert break could be replaced, relying upon the use of sacrificial material in bumpers on either side of the break. This would accommodate the transition time rather than impact on the length of the substituted adverts themselves. Substituting broadcast content with streamed material can also meet other use cases such as live longform content replacement.

Use of streaming media for advert or content substitution comes with the inherent risk that adverse network conditions may cause an interruption in playback, which would otherwise be avoided by preloading each advert entirely. Consideration should be given to the viewer experience when taking this approach.

This document identifies two main approaches a broadcaster can adopt to achieve the reach and quality trade off in practice.

7.1 DAS app approach to reach/quality trade off

One approach to achieving the reach versus quality trade off involves the advert server responding to a VAST [i.1] request with a list of all possible technical variants available for a given placement opportunity on a given broadcast channel.

Logic in the DAS app would use its knowledge of the context and capabilities it has available to it to determine whether the receiver can make the advert substitution. The DAS app would then select the best fit media variant for the substitution or terminate the placement opportunity if no suitable media variant was available.

When taking this approach, it is important to provide as much detail as possible with respect to the file wrapper and codec information of each media variant, as such detail will allow the DAS app to make the most appropriate selection. In section 3.9.1 of the VAST specification [i.1], the <MediaFile> element is defined to express two XML attributes which help here:

1. @type for the mime-type of the file
2. @codec for information pertaining to the encoding of the audio and video of the file.

DVB-TA recommends composing @type and @codec strings as defined in IETC RFC6381 "The 'Codecs' and 'Profiles' Parameters for 'Bucket' Media Types" [i.5].

This approach can reduce the number of parameters required in the VAST request, but may lead to more advert decisions served than can be fulfilled with a successful substitution.

7.2 Advert Server approach to reach/quality trade off

Another approach to achieving the reach versus quality trade off involves the DAS app parameterising the VAST [i.1] request with a specific context revealing the relevant capabilities, such that only the most appropriate technical advert media variant is returned in the VAST [i.1] response. Failing that, no advert decision is made if the advert server determines that the available media would not result in an advert substitution which meets the broadcaster's quality requirements.

This approach enables tighter inventory management as the advert server can know immediately that the substitution did not occur for the placement opportunity on that receiver. However, when

taking this approach, broadcasters should consider that not all advert servers have the same level of configurability on capabilities and viewer preferences.

8 General Guidance for broadcast streams and receivers

8.1 Broadcaster Responsibility for Advert Substitution Preparation

In order to give the best possible chance of making a seamless advert substitution, consideration should be given to the timing of an advert request in relation to the point of substitution.

The signalling of an upcoming placement opportunity needs to be given sufficiently far in advance to allow for the request to the advert server and subsequent media download. Note, that in the case of fully preloading a single advert media asset this may require more time than pre-buffering a longer streamed media asset.

8.2 Broadcast Stream Conditioning

The requirements for broadcast stream conditioning depend on the capabilities of the receiver that will be performing DAS on the broadcast stream. The ecosystem of TA compatible receivers will have a range of capabilities, to benefit from the broadest reach broadcasters may need to condition broadcast streams.

8.3 Receivers with Dual Decode capability

If all receivers are capable of simultaneously decoding audio and video from both the broadcast and the replacement advertisement (with no resource issue, such as memory), then the broadcast stream may not require any specific conditioning, as the receivers could continue to decode the broadcast (with the substituted advert) while decoding and displaying a substitute advert.

8.4 Receivers with Single Decoders and Precise Switching Timing

In order to enable smooth transitions from advert to broadcast for receivers with a single decoder and with precise switching timing capability, some conditioning of the broadcast stream is recommended to assist the advert substitution to occur with little or no video discontinuity. The conditioning should be applied at any point where the decoder could be finishing decoding a substituted advert and returning to decoding the broadcast stream. The stream conditioning puts a Random Access Point (RAP) at the first frame of video after any advert spot intended for TA replacement. This RAP could be the first frame of another broadcast advert in the advert break, or the first frame of any bridging material ("bumper") at the end of an advert break before programme content resumes.

The PTS for the RAP point should have the value referenced (directly or indirectly) by the TA signalling for the end of the placement opportunity being signalled.

In general, audio RAPs will not precisely coincide with the video RAPs, since an encoded audio frame usually has a duration different from video frames. It may be advised as authoring practice for adverts to put a brief silence period at the start and end of adverts (e.g. for 240 milliseconds). An audio RAP should occur sufficiently soon after the video RAP described above in order that the

audio switch from substituted advert to broadcast content can occur during the authored silence, and hence with no observable discontinuity.

8.5 Receivers without Precise Switching Timing

In order to support receivers that have only approximate (not frame-accurate) stream switching capability, the single RAP described above is not sufficient, as it may be missed by the receiver, and/or the receiver may take some time to start displaying moving video from that point. In this case, the broadcast content should include some sacrificial material to accommodate the time period of uncertainty for the stream switch, otherwise incomplete advertisements may be shown. The sacrificial material could, for example, be in the form of some transition video and silent audio at the intended splicing points. If the receiver has a single decoder, then it will only be able to switch to broadcast at a RAP. The period of sacrificial material could be encoded with frequent RAPs, including one at the time indicated in the TA signalling as above, so that one of these can be selected by the receiver to restart decoding the broadcast stream. However, to do so video encoders, capable of allowing dynamic adjustment of the RAP frequency would be required.

9 Annex: Change History

Date	Version	Information about changes
16 Oct 2019	1.0	Integrating text approved by TM into ETSI format
17 Oct 2019	1.0.1	Minor editorial changes and one new text on codec and mime-type
21 Oct 2019	1.0.2	Minor editorial changes

10History

Document history		
<Version>	<Date>	<Milestone>

Latest changes made on 2019-01-29