Xumo HLS Specification - October 2018

Table of Contents

Overview

Where Content Partners supply Live Simulcast or Live Event streams to Xumo, it is essential to ensure compatibility between the encoded streams and Xumo's target devices. The HTTP Live Streaming specification leaves many factors, such as bit rates, segment duration, codecs etc to the implementer. In practice, such flexibility can cause compatibility problems where a given stream does not play correctly or optimally on all devices, and certain players cope only with streams encoded in very specific ways. This document removes some freedom from the HLS specification by mandating certain encoding parameters ensuring compatibility between streams and devices. This document also outlines some best practices for ensuring a high quality of experience for audiences on Xumo.

Summary

- [ ] Test the streams using Apple Media Stream Validator
- [ ] Streams must be HLSv3
- [ ] Files must be served over HTTPS for both M3U8 and TS files
- [ ] The MIME type must be correct on the origin and CDN, for both M3U8 and TS files
- [ ] All files must be CORS compliant on the origin and CDN
- [ ] Streams should use a consistent 6 second segment duration
- [ ] Live playlists should have at least 6 segments in the sliding window
- [ ]

- [ ] Transport streams should not be deleted from the CDN for at least 2 minutes
- [ ] The keyframe interval should be a consistent 2 seconds
- [ ] Streams must not contain discontinuities
- [ ] Audio must be 2 channel stereo, with audio loudness normalized to ATSC A/85
- [ ] At least 5 video profiles for 24/7 channels, suitable for mobile through to smart TVs
- [ ] At least 3 video profiles for one-off events
- [ ] H.264, high profile video
- [ ] Consistent video frame rate
- [ ] Content should not be interlaced
- [ ] Closed captions for accessibility

Validation & Verification Tools

Apple Media Stream Validator

Streams should be validated against Apple's Media Stream Validator tool, available through Apple's Developer Portal, as referenced in Apple Technical Note 2224 Best Practices for Creating and Deploying HTTP Live Streaming Media for Apple Devices. The HLS Authoring Specification for Apple Devices includes detailed technical explanations of the results. The Apple Media Stream Validator tool does not test compliance against one particular version of the HLS specification, but tests against Apple's best practice for their device footprint. Successful validation with the Apple Media Stream Validator does not guarantee playback across all Xumo devices, but is an excellent analysis tool.

Below is the sample usage of the Media Stream Validator tool for macOS, analyzing a 300 second sample of content.

```
$ mediastreamvalidator -t 300 --verbose --device=atv
--validation-data-path="./validationdata_data.json"
'https://example.com/hls/master.m3u8'
$ hlsreport.py --id --pl --verbose --output "./validationdata_data.html"
"./validationdata_data.json"
$ open -a "Safari" "./validationdata_data.html"
```

Apple Technical Note 2224 Best Practices for Creating and Deploying HTTP Live Streaming Media for Apple Devices	https://developer.apple.com/library/ios/technotes/tn2224
HLS Authoring Specification for Apple Devices	https://developer.apple.com/library/content/documentation/General/Reference/HLSAuthoringSpec/Requirements.html

HLS Specification

HLS Version

Although the HLS Specification was standardized under RFC 8216 as protocol version 7 in August 2017, the lowest common denominator of Xumo's target hardware (Smart TV, Connected TVs, OTT devices) is limited to protocol version 3.

In order to provide compatibility with the widest range of devices, HLS content streams must adhere to Version 3 of the HLS protocol, as defined in Draft 6 https://tools.ietf.org/html/draft-pantos-http-live-streaming-06. HLS Version 3 has the most ubiquitous support among consumer electronics devices.

Media Playlist order within the master playlist

Most devices and players will commence playback of the first media playlist within the master playlist. The first media playlist in the master playlist is considered the 'default' playlist.

The default playlist at the top of the master playlist file should be one with a *"reasonable bitrate for a broadband connection"*, ensuring fast channel change while maintaining a suitable video quality. Apple recommends that *"the default video variant(s) SHOULD be the 2,000kb/s variant."*. In practice, Xumo has found that 2,000kbps is somewhat optimistic and a default bitrate of 1,000kbps is more appropriate.

Hosting & CDN Requirements

Secure Hosting

Xumo requires that both M3U8 and TS files are served over HTTPS. Xumo's apps run in secure mode. Modern web browsers will block mixed-mode HTTPS-HTTP content, therefore HTTPS is an absolute requirement. Master manifest files, media playlist manifest files and transport stream segments must all be available via HTTPS.

TLS v1.0 is the minimum supported security level. Some older devices are limited to HTTPS using TLS version 1.0.

Relative URLs are supported. HTTPS 302 Temporary Redirects are also supported.

MIME Type declaration

The origin server and content delivery network must both declare the MIME type.

File Extension	MIME Type
`.m3u8`	The preferred MIME Type is `application/x-mpegURL`. `vnd.apple.mpegURL` is also supported.
`.ts`	`video/MP2T`

The MIME Type can be inspected using cURL.

```
$ curl --silent --head 'https://example.com/hls/master.m3u8'

HTTP/1.1 200 OK
Date: Fri, 02 Feb 2018 21:40:24 GMT
Content-Type: application/x-mpegURL
```

Cross-Origin Resource Sharing (CORS) Headers

The origin server and content delivery network must both be fully CORS compliant for both M3U8 and TS files. This is required by modern browsers and Google Chromecast devices.

The hosting server should either (a) return a wildcard or (b) dynamically allow the domain of the client which makes the request:

```
Access-Control-Allow-Origin: *
Access-Control-Allow-Methods: GET
Access-Control-Max-Age: 3000
```

CORS Headers can be inspected using cURL.

```
$ curl --silent --head --header 'origin: https://xumo.tv' --request GET
'https://example.com/hls/master.m3u8

HTTP/1.1 200 OK
Date: Fri, 02 Feb 2018 21:40:24 GMT
Access-Control-Allow-Origin: *
Access-Control-Allow-Methods: GET
Access-Control-Max-Age: 3000
```

Note: Some modern browsers do not allow an `Access-Control-Allow-Origin: *` wildcard to be used in combination with `Access-Control-Allow-Credentials: true`. If using `Access-Control-Allow-Credentials: true`, in order to be CORS compliant, the server must not return `Access-Control-Allow-Origin: *` and should instead parse the value of the origin header in the request and dynamically populate return the particular origin, eg `Access-Control-Allow-Origin: https://xumo.tv`.

```
$ curl --silent --head --header 'origin: https://xumo.tv' --request GET
'https://example.com/hls/master.m3u8'

HTTP/1.1 200 OK
Date: Fri, 02 Feb 2018 21:40:24 GMT
Access-Control-Allow-Origin: https://xumo.tv
Access-Control-Allow-Methods: GET
Access-Control-Max-Age: 3000
```

HTTPS 302 and 307 Temporary Redirects

Some players do not support a combination of HTTPS 302 or 307 Temporary Redirects and relative URLs.

- If the master playlist is hosted behind a HTTPS 302 or 307 Temporary Redirect, media playlists must be referenced by absolute (and not relative) URLs.
- If the media playlist is hosted behind a HTTPS 302 or 307 Temporary Redirect, transport stream segments must be referenced by absolute (and not relative) URLs.

One-time URLs

If the hosting server or content delivery network uses one-time URLs for security or dynamic ad insertion, please work with your Xumo representative to ensure compatibility.

Geo-fencing

It is acceptable, but not a requirement, that CDN geo-fencing is applied. Xumo requests that a small number of individual IP addresses can be whitelisted so that Xumo's manufacturer partners can test streams prior to new models being launched to the market.

Maximum URL Length

Some Smart TVs are unable to load very long URLs. The maximum length of the URL is 512 characters.

Transport Streams

Segmented TS vs Single-file Byte Range TS/MP4

Streams must include segmented transport streams, in accordance with version 3 of the HLS protocol. Single-file, byte-range HLS was introduced in version 4 of the HLS protocol and is therefore not supported across all target devices.

Transport Stream Segment Duration

The transport stream segment duration must be between 3 and 10 seconds. Some target devices do not support a segment duration of less than 3 seconds.

Xumo strongly recommends a 6 second transport stream segment duration, since this provides an excellent compromise between cold-start time and network efficiency.

The transport stream segment duration should not exceed the declaration within the `EXT-X-TARGETDURATION` tag. The Apple Media Stream Validator Tool can be used to confirm that the target duration is not being exceeded.

Live and Live event media playlists should contain a minimum of 6 transport stream segments to ensure performance under higher latency network conditions.

Video, Audio and Data Tracks

Some platforms do not support playback if the Transport Stream contains data tracks such as ID3 or SCTE-35 data tracks. Transport streams should contain one video and one audio stream.

Encoding Profiles for Maximum Compatibility

The following are recommended bitrates and bandwidths, however these are guidelines rather than absolutes. Xumo will work with Content Partners to test compatibility across target devices.

If source content is only available in Standard Definition, please contact your Xumo representative.

> These recommendations are subject to change as device compatibility improves and broadband availability increases.

	Video Codec	Video Resolution	Video frames per second (FPS)	Keyframe Interval	HLS Segment Duration	Video maximum bitrate (kbps)	Video color space	Audio Codec	Audio bitrate (kbps)	Audio Sample Rate (kHz)	Audio Channels
Smart TV	H.264 High Profile	**768x432 ****	23.976, 25, 29.970	2 seconds	6 seconds	**1000 ****	YUV 4:2:0	AAC	96	As per source, i.e. 44.1 or 48.0	2
Cellular 3G	H.264 High Profile	256x144	23.976, 25, 29.970	2 seconds	6 seconds	96	YUV 4:2:0	AAC	32	As per source, i.e. 44.1 or 48.0.	1 or 2
Smart TV	H.264 High Profile	640x360	23.976, 25, 29.970	2 seconds	6 seconds	700	YUV 4:2:0	AAC	64	As per source, i.e. 44.1 or 48.0	2
Smart TV	H.264 High Profile	1280x720	23.976, 25, 29.970	2 seconds	6 seconds	2000	YUV 4:2:0	AAC	96	As per source, i.e. 44.1 or 48.0	2
Smart TV	H.264 High Profile	1280x720	23.976, 25, 29.970	2 seconds	6 seconds	3000	YUV 4:2:0	AAC	96	As per source, i.e. 44.1 or 48.0	2
Smart TV	H.264 High Profile	1920x1080	23.976, 25, 29.970	2 seconds	6 seconds	4500	YUV 4:2:0	AAC	128	As per source, i.e. 44.1 or 48.0	2

** Note: Xumo recommends that the media playlist with the bitrate closest to 1000kbps is declared at the head of the master playlist. This is know as the default media playlist. If a low resolution media playlist is declared as the default media playlist, the visual experience is sub-optimal on a large screen device. If a high bitrate media playlist is declared as the default media playlist, the load time on cellular devices can cause unnecessary buffering.

Special Considerations for Cellular Devices

Xumo supports multiple platforms, including smart TVs, desktop and mobile. In Apple's latest versions of the AppStore Review Guidelines https://developer.apple.com/app-store/review/guidelines/#media-content, Apple state that HLS content should *"include a cellular profile of 192 kbps or lower"*. Prior to 2015-05-04, Apple AppStore limit was 64kbps.

The peak bitrate of the cellular profile should not exceed 192kbps. The Apple Media Stream Validator can be used to measure and report the peak and average bitrates.

As part of the optimization process, Xumo may request separate master playlists, each optimized for mobile and smart TV platforms, where the smart TV version does not include very-low bitrate streams. This is only required in exceptional circumstances where low bitrates and long segment lengths create an undesirable user experience.

Discontinuity

Although the HLS specification includes support for the `EXT-X-DISCONTINUITY` tag to indicate that the characteristics of the video may have changed, devices are very sensitive to variation in either the transport, video or audio stream. Many Smart TV and OTT devices have significant playback issues if the characteristics of the media changes within a stream, even if the change is indicated with an `EXT-X-DISCONTINUITY` tag.

The PMT, PID, sequence, segment duration, frame-rate, resolution, aspect ratio and codecs must remain constant within any single HLS profile.

If the channel requires monetization through server-side ad insertion (SSAI) aka dynamic ad insertion (DAI), please contact your Xumo representative to discuss compatibility. Maintaining a seamless stream without discontinuities presents significant challenge for SSAI.

Video Frame Rates

Video must be progressively scanned. Interlaced sources (eg 480i, 576i and 1080i) must be deinterlaced to progressive format.

Video should have a frame rate of either 23.976p, 24p, 25p, 29.970p or 30p.

Frame rate	Exact Frame rate	
23.976	=24*1000/1001	NTSC Film
24.000		Film
25.000		PAL. European 'shot for broadcast television'
29.970	=30*1000/1001	NTSC. US 'shot for broadcast television'
30.000		NTSC. US 'shot for broadcast television'

Content should be encoded at either the original framerate of the source or if the framerate is greater than 30fps, a factor of the original framerate. For example:

- 60fps progressive source should be converted to 30fps

Frame-rate conversion (including as Telecine and 3:2 pull-down) should be avoided.

- 23.976 and 24fps progressive source should not be converted to 29.97/30fps and should remain as either 23.976 or 24fps
- 25fps progressive source should not be converted to 29.97/30fps and should remain 25fps

Streams must never contain mixed frame-rates, even if separated by `EXT-X-DISCONTINUITY` tag, since support for the HLS discontinuity tag cannot be relied upon across all devices.

The framerate should remain consistent between profiles.

Video Keyframes & I-frames

All transport stream segments must always begin with an Instantaneous Decoder Refresh (IDR) I-frame or keyframe.

Additional I-frames should be placed every 2 seconds, so that subsequent HLS segmenters can produce a consistent segment duration.

Many encoders also place I-frames at scene changes.

Xumo recommends a 2 second keyframe interval with a 6 second TS segment duration.

Video Codecs

Video must be encoded using H.264/MPEG-4 AVC compression. H.265 is not currently supported across all target devices, therefore H.264/MPEG-4 AVC remains the required codec.

All variants should use H.264 High Profile.

The video codec must be declared within the `EXT-X-STREAM-INF:CODECS` tag. The `CODECS` declaration must match that of the encoded video. The Apple Media Stream Validator can be used to confirm that the encoded video matches the `CODECS` declaration.

Profile, Level	Numerical representation	EXT-X-STREAM-INF:CODECS
H.264 High Profile, Level 3.0	High = 100 = 0x66 Constraint = 00 = 0x00 Level 3.0 = 30 = 0x1e	`"avc1.64001e"`
H.264 High Profile, Level 3.1	High = 100 = 0x66 Constraint = 00 = 0x00 Level 3.1 = 31 = 0x1f	`"avc1.64001f"`
H.264 High Profile, Level 3.2	High = 100 = 0x66 Constraint = 00 = 0x00 Level 3.2 = 32 = 0x20	`"avc1.640020"`
H.264 High Profile, Level 4.0	High = 100 = 0x66 Constraint = 00 = 0x00 Level 4.0 = 40 = 0x28	`"avc1.640028"`
H.264 High Profile, Level 4.1	High = 100 = 0x66 Constraint = 00 = 0x00 Level 4.1 = 41 = 0x29	`"avc1.640029"`

Video Aspect Ratio

Recently produced video should be 16:9 aspect ratio and encoded in 16:9 compatible resolutions. However, it is appreciated that cinematic content may have a wider aspect ratio and retro content may have a 4:3 aspect ratio.

- Video must never be window-boxed, where black bars appear top, bottom, left and right.
- Ideally, widescreen video should be encoded at the original aspect ratio, without letter-boxing. The player will adapt accordingly. Widescreen video may also be letter-boxed in a 16:9 frame (eg Cinemascope with black bars top and bottom). If video is letter-boxed, it must only ever be letter-boxed to a 16:9 aspect ratio.
- Ideally, 4:3 or Academy Ratio video should be encoded at the original aspect ratio, without pillar-boxing. The player will adapt accordingly. 4:3 or Academy Ratio video may also be letter-boxed in a 16:9 frame, with black bars left and right. If video is pillar-boxed, it must only ever be pillar-boxed to a 16:9 aspect ratio.

Square aspect-ratio 'social media' content and portrait user-generated content is not generally acceptable, except where it is contextually relevant and appropriate. If in doubt, please contact your Xumo representative for clarification.

All video aspect ratio, display aspect ratio and pixel aspect ratio should be consistent throughout all HLS variants to ensure a seamless user experience during rate-adaption. By definition, 1080p and 720p HD standards use a square (1:1) pixel aspect ratio, therefore all profiles should also use a square pixel aspect ratio.

Video Resolution

Xumo is a high definition video service and expects the video quality to be suitable for consumption on large Smart TVs.

- Recently produced content must be available up to a resolution of at least 1280x720.
- Exceptions can be made for older and classic content which was originally produced and distributed at 480p and 576p standard definition resolutions.
- Low resolution 'web' or 'mobile' content is not accepted.
- Content should not be upscaled to a resolution higher than than the source video, since upscaling is best handled by the video player.
- `EXT-X-STREAM-INF:RESOLUTION` must be declared in the master playlist.

Video Bitrates

For any particular video resolution, an encoder needs enough bandwidth to represent the pixels, otherwise playback can appear blocky,

especially during fast-motion content. Measures of quality are subjective and cannot be strictly defined. However, a good guide is to use a video bits per pixel of between 0.075 and 0.15.

Apple recommends that the bitrate between adjacent variants should be a factor of 1.5 to 2 apart.

EXT-X-STREAM-INF:BANDWIDTH

The HLS specification requires that the bitrate for any individual segment must not exceed the declared `#EXT-X-STREAM-INF:BANDWIDTH` within the M3U8 file, measured across the duration of the segment.

Apple's Media Stream Validator can be used to check that the bitrate of the segments do not exceed the declared `#EXT-X-STREAM-INF:BANDWIDTH`.

The popular x264 codec includes a HRD/VBV Max Rate and Buffer Size rate-control mechanism, which can be used to constrain the bitrate across a buffer of approximately the same duration as the HLS segment duration.

Audio Streams

The HLS protocol version 3 requires that video and audio must be multiplexed together into an MPEG Transport Stream.

Although more modern versions of HLS support separate audio and video elementary streams, this configuration is not supported across all target devices.

Audio Codecs

Audio streams must be encoded with either the AAC Low Complexity (AAC-LC) or High Efficiency AACv1 (HE-AACv1) audio codec. Support for HE-AACv2 cannot be guaranteed. Although the bitrate can change between variants, the codec must be consistent across all variants.

The audio must should be declared within the `CODECS` attribute if the EXT-X-STREAM-INF tag.

Codec name	EXT-X-STREAM-INF:CODEC
AAC-LC	`"mp4a.40.2"`
HE-AAC	`"mp4a.40.5"`

Audio Volume Levels

In order to ensure that content from different sources (including ads) have similar levels of volume, audio must be normalized in line with ATSC A/85 https://www.atsc.org/recommended-practice/a85-techniques-for-establishing-and-maintaining-audio-loudness-for-digital-television/ for North America or EBU-R128 https://tech.ebu.ch/loudness for European distribution. Most modern encoders have settings to enable ATSC A/85 (North America), EBU R128 (the European standard) or BS.1770 (the underlying loudness measurement algorithm).

- ATSC A/85 requires that the Integrated loudness should be -24 LKFS/LUFS +/-2 LU
- EBU R-128 requires that the Integrated loudness for live streams should be -23 LUFS +/-1 LU

In practice, the difference between ATSC A/85 and EBU R128 is minimal. All content on Xumo must be normalized to either ATSC A/85 or EBU R128.

Testing Audio Volume Levels

The open source FFmpeg tool includes an ebur128 library which can be used to measure the Integrated Loudness for both ATSC A/85 and EBU R128. The following FFmpeg command will sample 3 minutes of content and analyze the Integrated loudness.

```
$ ffmpeg -i 'https://mystream.example.com/1280x720.m3u8' -t
"00:03:00.000" -filter:audio "ebur128" -vn -f "null" /dev/null

Summary:
  Integrated loudness:
    I:         -23.4 LUFS
    Threshold: -33.5 LUFS
  Loudness range:
    LRA:         2.7 LU
    Threshold: -43.5 LUFS
    LRA low:   -24.9 LUFS
    LRA high:  -22.2 LUFS
```

Closed Captioning Support

Content Partners must comply to the appropriate regulations for the target market. In the US, the recently updated FCC guidelines are available from https://www.fcc.gov/guides/captioning-internet-video-programming.

Xumo requires that if closed captioning or subtitles are available, they must be included in the stream. This is especially important for movies and full-episode television content which has previously appeared on broadcast television.

There are several methods of encoding closed captioning (in-band CEA-608/708 or a timed text sidecar file such as SRT or WebVTT). This document does not attempt to cover all the methods of delivering closed captions, since this is usually limited by the capabilities of the encoding platform. Please work with your Xumo representative to discuss closed captioning options.

Example of a simple HLSv3 playlist

Example Master Playlist.

```
#EXTM3U
#EXT-X-VERSION:3
#EXT-X-INDEPENDENT-SEGMENTS
#EXT-X-STREAM-INF:BANDWIDTH=1119342,CODECS="avc1.64001f,mp4a.40.2",RESOL
UTION=768x432
a_432p30.m3u8
#EXT-X-STREAM-INF:BANDWIDTH=2155321,CODECS="avc1.64001f,mp4a.40.2",RESOL
UTION=1280x720
a_720p30.m3u8
#EXT-X-STREAM-INF:BANDWIDTH=1493322,CODECS="avc1.64001f,mp4a.40.2",RESOL
UTION=960x540
a_540p30.m3u8
#EXT-X-STREAM-INF:BANDWIDTH=797268,CODECS="avc1.64001f,mp4a.40.2",RESOLU
TION=640x360
a_360p30.m3u8
#EXT-X-STREAM-INF:BANDWIDTH=223621,CODECS="avc1.64000c,mp4a.40.2",RESOLU
TION=256x144
a_144p30.m3u8
```

Example Media Playlist for an event.

```
#EXTM3U
#EXT-X-VERSION:3
#EXT-X-PLAYLIST-TYPE:EVENT
#EXT-X-TARGETDURATION:6
#EXT-X-INDEPENDENT-SEGMENTS
#EXT-X-MEDIA-SEQUENCE:0
#EXTINF:6.000,
a_540p30_000.ts
#EXTINF:6.000,
a_540p25_001.ts
#EXTINF:6.000,
a_540p30_002.ts
#EXTINF:6.000,
a_540p30_003.ts
#EXTINF:6.000,
a_540p30_004.ts
#EXTINF:6.000,
a_540p30_005.ts
...
#EXTINF:6.000,
a_540p30_754.ts
#EXT-X-ENDLIST
```