Akamai

AKAMAI.COM

The Guide to Best Practices in PREMIUM ONLINE VIDEO STREAMING

PART 2: PREPARING VIDEO CONTENT FOR DELIVERY

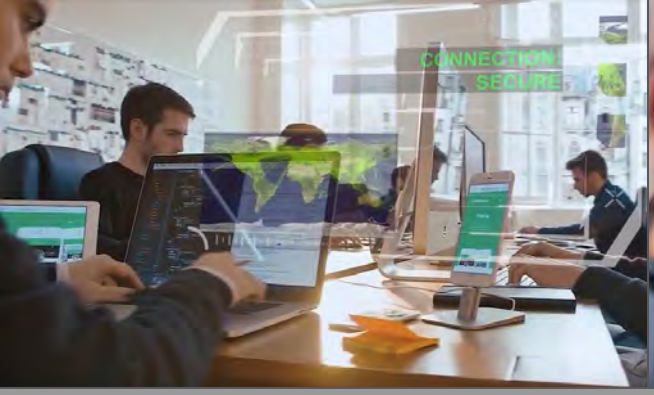

The second installment in a series looking at best practices for delivering premium video content, this piece explores preparing OTT video for delivery over content delivery networks.

Online audiences demand flawless playback experiences, particularly when it comes to premium content. This requires online content providers to carefully plan and execute during the content preparation and ingest phases. The ramifications of any missteps early in the OTT video distribution chain can result in countless issues during playback. With the cost of acquiring content rights on the rise, and audience expectations for quality of experience increasing, the stakes are higher than ever.

Preparing video content for delivery in a cost-efficient manner at the highest quality levels requires adherence to best practices for transcoding, packaging, manifest preparation, ingest, and storage. It's also essential for implementation of personalization, navigation, dynamic advertising, security, and other functions that contribute to meeting audience expectations and driving revenue.

In this second installment of "Best Practices for Premium Video Streaming," we'll explore content preparation, the interplay of origins with CDNs, and steps to maximize quality of experience.

Let's explore phase two: **Preparing video content** FOR DELIVERY.

Transcoding and Bitrate Profile
Creation for OTT Video

Fundamentally, the transcoding process requires that content be prepared at the highest quality levels and optimized prior to delivery. This entails preparation of streaming segmentation, bitrates, bit depth, and framing policies for live as well as on-demand use cases across all Internet-connected viewing devices. The quality of the video renditions produced during transcoding are critical to maintaining quality downstream during delivery and playback. However, it's not enough to create the highest-quality video renditions in all cases. Depending on the situation, video quality can be too high—resulting in wasted bandwidth, or it can be too low—resulting in pixelated playback.

Optimizing Bitrate Ladders
for Optimal Video Playback

To meet playback quality expectations of online viewers, great care should be taken during the transcoding process to select the optimal bitrate ladder for a given piece of content. Some content providers opt for a one-size-fits-all approach, creating a similar bitrate ladder for their entire VOD catalog; however, this can result in unnecessary storage and delivery costs, and result in suboptimal playback quality.

The best practice for video on demand calls for context-aware encoding, which establishes an optimal bitrate ladder for each title in a catalog. Each scene is produced in multiple quality levels and then adapts as needed. Using this approach, the perceived quality during playback is the same, but requires less bandwidth. Increasingly, content providers are using machine learning to derive optimal bitrate ladder selections.

The best practice for live streaming scenarios remains the creation of a single bitrate ladder for everyone, which can be trimmed for certain viewing scenarios. Trimming the number of available renditions can be handled further down the line via manifest manipulation or through advanced player logic.

Addressing the Next-Generation
Codec Conundrum

While AVC has long been the dominant codec used in video streaming, the advent of 4K ultra-high-definition (UHD) has encouraged content providers to look at emerging, alternative codecs. That includes HEVC, which consistently delivers a peak 50% improvement in compression efficiency over AVC.

With industry confidence in the superiority of 4K UHD stream quality and built-in support from the latest Apple smartphones and TVs, HEVC has gained traction in advance of wide-scale UHD availability. HEVC is also gaining on the encoder side as content providers and distributors replace or expand encoding assets.

With that said, the emergence of other next-generation codecs indicate we may not reach the industry consensus with HEVC that we had for AVC — at least not for some time. Royalty costs remain an obstacle for HEVC, while high-performing options like Google's royalty-free VP9 and The Alliance for Open Media's AV1 are both supported by a range of devices, web browsers, and industry leaders, giving HEVC some stiff competition.

All of this uncertainty points to a need to pay closer attention when choosing codecs. Since not all codecs support all devices, it is advised to consider multiple codecs when the economics make sense. The rule of thumb is to consider implementing new codecs at the point the savings on delivery costs reach the incremental cost of storage and encoding. The economics of any use case should always be the deciding factor in choosing the right mixture of codecs to support optimized delivery to all target devices.

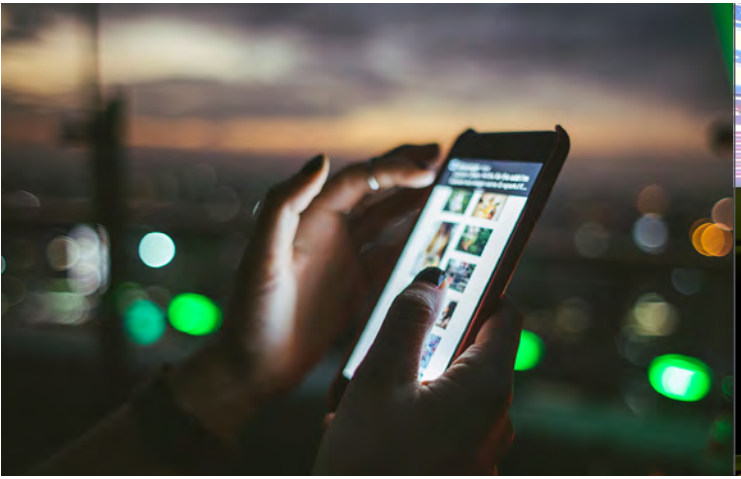

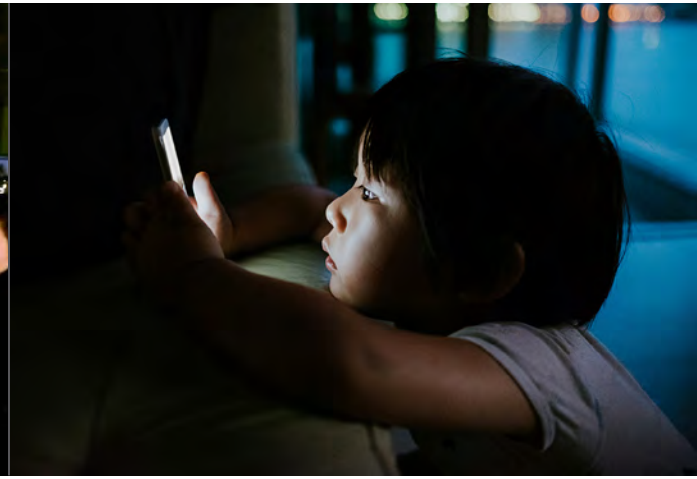

Manifest Preparation and Packaging for OTT Video:
HLS, DASH, and CMAF

For a given piece of content, each set of encoded renditions must be packaged with a manifest file that allows targeted clients to acquire and render the content using the recommended streaming formats that work with their players.

The process must also accommodate the complexity that characterizes preparation of adaptive bitrate (ABR) manifest files. This includes creating subsets of manifest files that utilize compilations of metadata for directing audio track captioning and linguistic subtitling; selecting digital rights management modes; descriptors associating advanced features with specific content segments; and placement options for dynamic advertising.

Thanks to accelerating market adoption of Dynamic Adaptive Streaming over HTTP (MPEG-DASH) and Apple's HTTP Live Streaming (HLS), along with vanishing support for Microsoft's Smooth or Adobe's HTTP Dynamic Streaming (HDS), best practice for distributors is to utilize HLS and/or DASH streaming formats for the vast majority of use cases. However, there are still some corner cases that require the use of Smooth.

The master manifest file—identified as the Master Playlist in HLS and the Media Presentation Description (MPD) in DASH—provides the player information about the audio and video codec, available bitrate profiles, segment sizes and sequencing, details relating to captioning, subtitles, and advertising.

All of these elements must be presented and synchronized to ensure precise, smooth playback on client devices. Best practice, stemming from common use of fMP4 containers with HLS and DASH and the need to maximize CDN efficiency, is to utilize the emerging common media file format (CMAF).

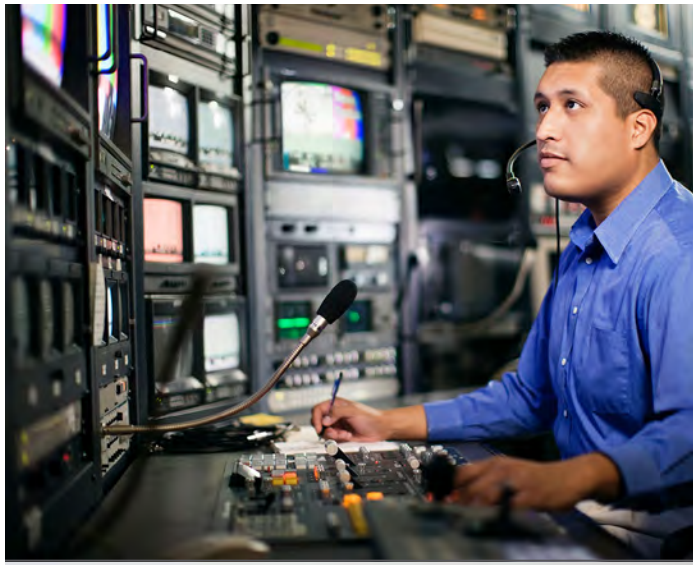

The Emergence
of CMAF

The Common Media Application Framework (CMAF) makes it possible to encode a video in multiple bitrate profiles with uniform segmentation utilizing the fMP4 container for streaming over either HLS or DASH. As an Application Format (AF) in the ISO lexicon, CMAF provides a lightweight framework that does not introduce new processes, but rather serves to combine existing formats and standards in a new way. With formal standardization achieved in 2017, CMAF needs to be taken into account as distributors adapt their workflows to accommodate best practices for 2018 and beyond.

For on-demand scenarios, best practice entails using DASH and/or HLS, fMP4 containers, and CMAF. Content providers can utilize one set of audio and video files packaged in CMAF along with two manifests (one for HLS and another for DASH) that reference the file. This can help lower content preparation and storage costs while providing better CDN efficiency, thanks to increased cache hit ratios.

In the case of live streaming, best practice again calls for the use of DASH and/or HLS, fMP4 containers, and CMAF. Similar to on-demand, CMAF enables content providers to utilize a single set of live audio/video files with two manifests to reference the file.

However, for live streaming, CMAF can help content providers achieve lower content preparation and ingest costs. Best practices include consolidating live ingest feeds into origin servers to ISO only, instead of both TS and ISO, which can cut bandwidth in half.

CMAF also has an optional chunked encoding mode for live streaming that, when combined with chunked transfer support through an origin and CDN, can cut end-to-end latency down to a few seconds and can also allow latency to be decoupled from segment duration. This option provides an avenue for content providers to achieve lower latency without decreasing segment duration, which comes with increasing quality and overhead penalties.

Hosted vs. In-House

Origin Infrastructures

Following transcoding and packaging, ingestion of the packaged content onto origin servers that can be accessed by CDNs is a critical next step. To support live or 24/7 linear video services reaching mass audiences, best practice includes using an encoder that pushes content to origin equipped to handle massive call volumes for content over CDN infrastructures. Best practices dictate that the origin service be able to dynamically find the best points of entry based on network conditions, audience locations, and other factors, and also support optimal modes of transport to minimize latency with no loss in quality. For on-demand scenarios, it is advised to utilize origin services that offer highly scalable infrastructures working in tandem with distributor workflows to optimize storage for high-performance video use cases.

Providers using their own origin infrastructures must have sufficient capacity operating in pull mode to handle all calls from all CDNs, as well as separate backup facilities in the event of an origin failure at the primary origin facilities. Switchover to backup also requires automated response capabilities, 24/7 monitoring, and high-performance transport connectivity.

Maximizing Quality Assurance
for OTT Video

Encoding and packaging for OTT distribution require absolute assurance that those processes are consistently performing as required. This is not only fundamental to achieving end-to-end performance goals—it's also essential to holding CDN operators accountable when it comes to meeting SLA commitments.

Commercial-grade reliability in the ingestion process begins with having sufficient resources to handle peak loads in primary and backup facilities that have no mutual dependency on the same power resources. Each set of facilities should have secondary routes into targeted origin locations to prevent network-related interruptions. Completely automated failover mechanisms are also essential to ensuring seamless continuity in the event of failure.

Performance monitoring and analytics tools must be employed to provide the comprehensive visibility needed to identify issues before they cause disruption; compare input and output quality; confirm latency and quality expectations are met for each video program's rendition sets; and verify workloads are properly apportioned to avoid overloading on transcoders.

Sustaining High Quality Video Delivery
with Akamai's Origin Solutions

Once live video content is prepared for delivery at the highest quality levels, content providers can rely on Akamai's Media Services Live for reliable and consistent ingest, from encoder to origin using Akamai's Ingest Acceleration. This capability utilizes a UDP transport protocol to provide improved consistency in throughput of higher-resolution streams and mitigation of poor network connectivity. Ingest reliability is also maintained using a highly distributed ingest network complemented by advanced dynamic encoder-to-entry-point mapping (DEEM) technology that replicates and routes content by dynamically matching encoders to an optimal entry point. This creates a self-healing aspect to Ingest that automatically reroutes encoders after the instant there's a disruption in the chosen path, ensuring reliability and consistency of streams in the face of varying network conditions.

Media Services Live also serves as a hosted origin service that enables unmatched reliability, consistency, and scale for live and live linear streams, which can also be used independent of Akamai's CDN. The service offers purpose-built capabilities including low latency support that cuts end-to-end hand wave latency from traditional 30- to 60--second durations to just 10 seconds, addressing the need for online and broadcast content to provide near-simultaneous playback for viewers. With support for CMAF and its low latency chunked transfer encoding profile, distributors can cut latency even farther.

Another benefit available with MSL stems from near real-time monitoring and reporting of signal performance in first-mile distribution that gives customers visibility into ingestion bitrates, fluctuations over a given time period, packet loss, latency and errors.

For on demand video, content providers can utilize standard SFTP file-transfer or Aspera's Fast Adaptive Secure Protocol technology to upload prepared content to Akamai's NetStorage cloud-based storage. Whether used for long-duration DVR or for SVOD services, NetStorage provides a scalable infrastructure that works in tandem with distributor workflows to optimize storage for high-performance use cases. The high availability solution replicates content immediately to two or more geographically diverse locations.

A growing OTT video market with ongoing technical advances that continually raise the performance benchmark make sustaining competitive edge a challenge. High-value video providers can remain competitive by adhering to these content preparation best practices inclusive of the right expertise, technologies, and facilities to meet increasing performance requirements and quality expectations.