CASE STUDY

Streaming Workflow Assessment
for a Regional Streaming Service

A confidential engagement involving codec optimization, packaging architecture, DRM workflows, and codec licensing strategy.

Contents

Executive Summary

A regional subscription streaming service with both VOD and live components engaged the Streaming Learning Center to evaluate its video encoding, packaging, DRM, and delivery workflow, and the work expanded to incorporate the potential business impact of emerging codec royalty programs. The objective was to identify opportunities to improve streaming efficiency while balancing video quality, encoding cost, CDN delivery cost, operational complexity, and long-term maintainability.

The engagement began with a straightforward observation: that all outputs (two VOD/two live) should share similar output quality to deliver a uniform experience to subscribers. Existing H.264, HEVC, VP9, and AV1 encoding ladders had evolved independently over time and targeted different quality levels. Before bandwidth savings or encoding costs could be compared, we reconfigured each codec to produce similar perceptual quality. Then the engagement turned to exploring alternative encoding optimizations to enable similar quality transmissions at lower bandwidths.

The project ultimately generated more than 6,500 test encodes and produced approximately 250 pages of technical analysis across three reports covering encoding optimization, packaging and DRM architecture, and codec licensing strategy. These reports were accompanied by a 12-page executive summary that prioritized recommendations for each encoding workflow and estimated the operational and financial impact of each recommendation using the client's encoding costs, CDN pricing, and actual viewing distribution.

Engagement at a Glance	
Duration	~4 months
Test encodes	6,500+
Codecs evaluated	4
Technical reports	3
Technical documentation	~250 pages
Executive summary	12 pages
Areas analyzed	Encoding, Packaging, DRM, Licensing

Scope of the Engagement

The work was organized as a progression, with each phase building on the previous one.

1. **Establish a common quality target** by tuning each codec to comparable perceptual quality.
2. **Optimize encoding parameters** for each codec using objective quality metrics and bitrate analysis.
3. **Model operational impact** by incorporating encoding cost, CDN pricing, and the client's viewing distribution into each recommendation.
4. **Review packaging and DRM architecture** to identify opportunities to reduce complexity while maintaining compatibility.
5. **Evaluate codec licensing considerations** to understand how emerging royalty programs could affect future codec strategy.

Recommendations were developed only after completing each phase, ensuring that technical conclusions were supported by both engineering measurements and operational analysis.

Establishing a Common Quality Target

The first engagement phase established a common quality target across all four codecs so that viewers receive a uniform viewing experience irrespective of codec or whether the stream was live or VOD. This was necessary because the client's H.264, HEVC, VP9, and AV1 encoding ladders had been developed independently over time. Each codec targeted a different quality level, making direct bitrate comparisons misleading.

Figure 1 shows the client's original encoding ladders, with a VMAF differential approaching five VMAF points between the highest and lowest quality streams.

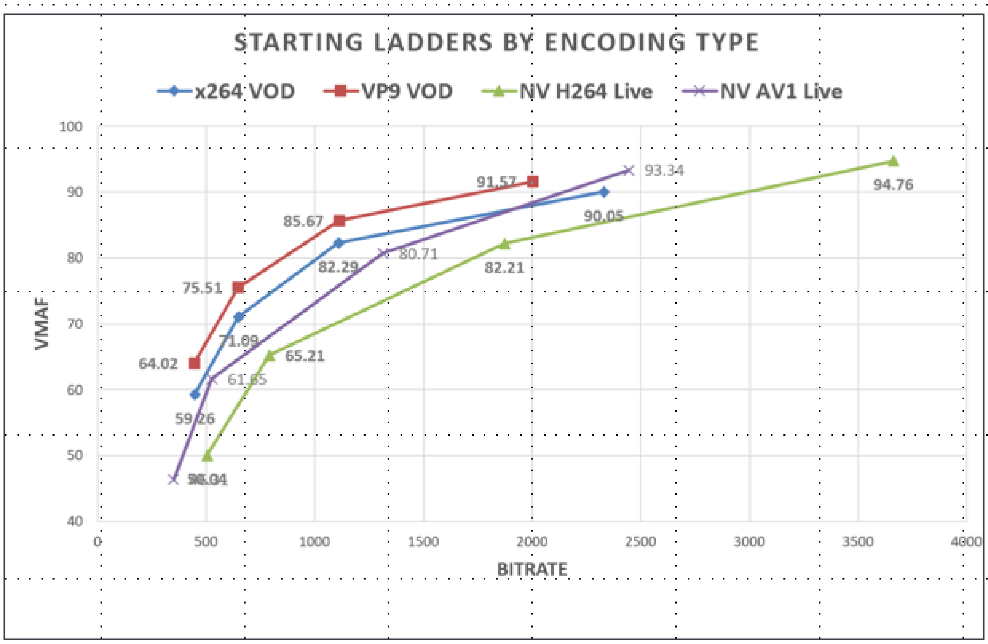

Figure 1. The original encoding ladders showed significant differences in viewer experience based on the codec and whether the stream was VOD or live.

To establish a common baseline, each encoding ladder was adjusted to produce a target VMAF of 93 for the top rung. Objective quality metrics guided the tuning process, with visual inspection used to verify the results. VMAF was the primary metric deployed with key findings verified via spot checks with PSNR and SSIM.

Figure 2 shows the normalized ladders after tuning, with the VMAF differential under 0.75 VMAF points, producing a much more uniform viewing experience.

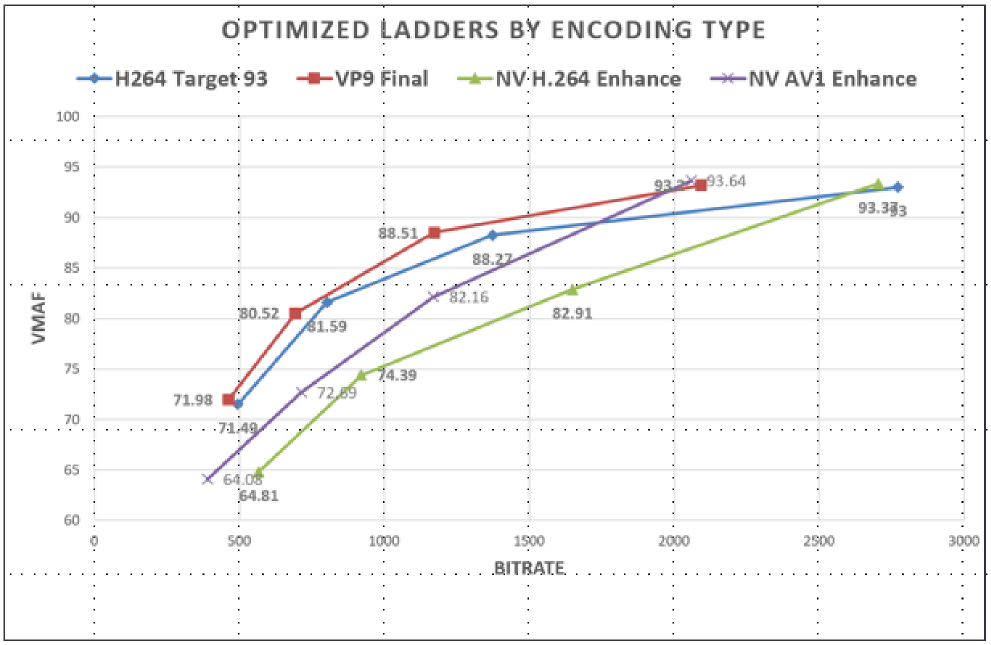

Figure 2. Normalized encoding ladders. After tuning, all four codecs produced comparable perceptual quality, providing a consistent baseline for subsequent optimization.

Normalizing the ladders affected every stage of the analysis that followed. Codec optimization, operating cost estimates, CDN savings, and breakeven calculations all depended on comparing equivalent quality. Establishing that baseline ensured that the recommendations presented later in the project were based on valid technical comparisons rather than differences in target quality.

Optimizing Codec Performance

With comparable quality established, the next phase focused on identifying encoding parameters that delivered the target quality at the lowest bitrate. Each codec was optimized independently, and tested options included codec presets, rate control modes, GOP structure, CQ/CRF settings, and other codec-specific parameters.

Figure 3 illustrates the results of the optimization process for AV1. The Baseline configuration was the original configuration adjusted to achieve VMAF 93. The Enhanced ladder was produced using a configuration that incorporated all recommended optimizations.

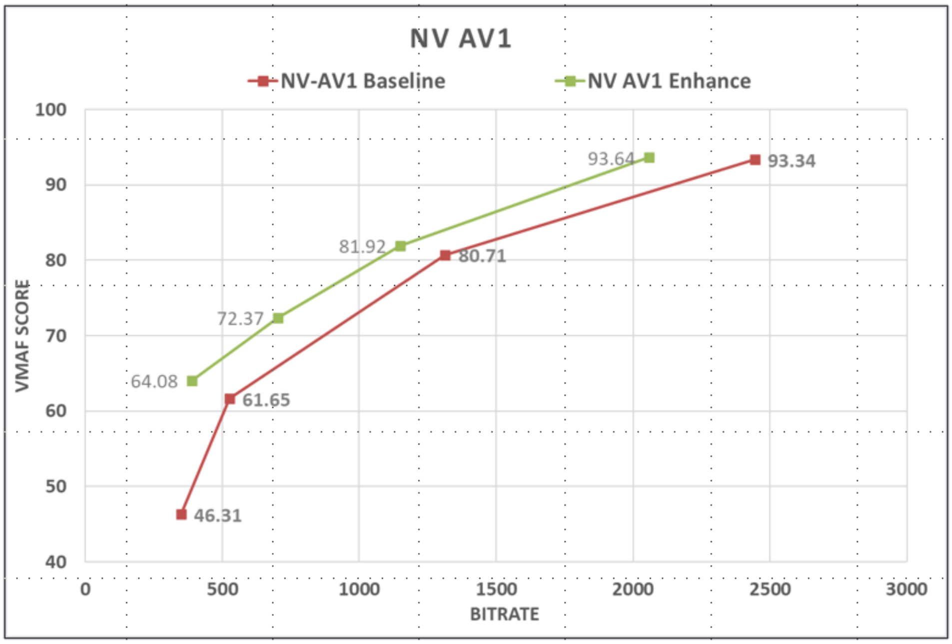

Figure 3. *AV1 optimization results. The Enhanced configuration reduced bitrate while maintaining the target quality established during the normalization phase.*

As you can see, the Enhanced ladder produced higher VMAF scores at lower bitrates through all ladder rungs, particularly the lower rungs typically delivered to mobile viewers. Similar optimization procedures were performed for H.264, HEVC, and VP9.

Modeling Deployment Costs

The optimized encoding configurations established the bitrate and throughput for each codec. The next phase estimated the cost of deploying each configuration in production.

The analysis combined three client-specific inputs. The first was encoding performance, including bitrate and throughput measured during the optimization phase. The second was the client's encoding and CDN costs. The third was the distribution of viewing across the encoding ladder. Rather than assuming a generic viewing profile, the analysis used the client's actual viewing distribution so that bandwidth savings reflected how their content was consumed.

The analysis differed for VOD and live workflows. For VOD, the additional cost of encoding each asset was compared with the projected bandwidth savings during playback to determine the number of viewing hours required to recover the additional encoding cost.

Live streaming required a different approach. Reduced throughput increases the number of encoding servers required to support the same channel count, creating an additional capital expense. The measured throughput of each configuration was used to estimate the required hardware, and the resulting capital cost was amortized over the expected service life of the equipment. The breakeven calculation then compared this additional cost with the projected bandwidth savings based on the client's viewing distribution.

Figure 4 shows an example from the live analysis.

Encoder Cost	
Base price	$5,000
Five year hosting @$25/mo	$1,500
Total	$6,500
Cost per hour (24/7/365*5)	$0.1484
Throughput - Baseline	22.2
Cost per hour - Baseline	$0.0067
Throughput - Enhanced	10.2
Cost per hour - Enhanced	$0.0145
Delta	$0.0078
Breakeven	**39.5**

Figure 4. *AV1 optimization results. The Enhanced configuration cut throughput requiring additional transcoding hardware to achieve the same throughput. Amortized over five years, the additional cost was recouped in 39.5 hours of viewing. So, if more than 40 viewers watched the average live channel simultaneously, the Enhanced configuration was more cost effective.*

Reviewing Packaging and DRM

The next engagement phase reviewed the client's packaging and DRM architecture to determine whether the recommended codec strategy could be implemented without introducing unnecessary operational complexity.

The assessment examined both live and VOD workflows, including HLS and DASH delivery, DRM and non-DRM packaging, manifest generation, segment formats, and the packaging tools used throughout the production workflow. The objective was to identify structural issues, compatibility constraints, and opportunities to simplify the packaging architecture while maintaining support for the client's existing device base.

The review found that the client's operational complexity was driven primarily by DRM requirements and device compatibility rather than by the packaging tools themselves. FFmpeg and Shaka Packager were used appropriately throughout the workflow, and the existing architecture reflected a conservative approach intended to maximize compatibility across Apple, Android, Smart TV, and legacy devices.

Several opportunities for future simplification were identified. For example, non-DRM live workflows already used a single CMAF segment set shared between HLS and DASH, demonstrating that the existing toolchain could support unified packaging where DRM constraints did not apply. For DRM-protected content, however, maintaining separate HLS and DASH packaging paths remained a reasonable design choice because it reduced compatibility risk, despite increasing storage requirements and operational complexity.

The review also showed that live latency was determined primarily by packaging, CDN, and player configuration rather than encoder settings. As a result, enabling low-latency encoder modes alone would not produce a measurable reduction in end-to-end latency without corresponding changes throughout the delivery workflow.

Evaluating Content Royalty Exposure

The final engagement phase evaluated how emerging content royalty programs could affect the codec recommendations developed during the earlier phases of the project.

The report identified the royalty pools applicable to AV1, VP9, HEVC, and VVC, summarized their published licensing terms, and documented the areas of uncertainty surrounding their enforcement and pricing. It also compared the different licensing structures used by Access Advance and Avanci and explained how those differences could affect codec deployment decisions.

The report then examined how content royalties could change the economics of deploying advanced codecs. Using representative royalty assumptions, it calculated the viewing hours required for the bandwidth savings from AV1, VP9, and HEVC to offset potential royalty costs. These calculations illustrated the scale of the savings required before royalties became economically insignificant.

Because both royalty programs are relatively new and significant uncertainties remain regarding pricing and enforcement, the report recommended repeating these calculations before making long-term codec decisions.

Conclusions

The engagement produced a set of recommendations covering encoding, packaging, DRM, and codec strategy for both live and VOD workflows. Each recommendation was supported by objective quality measurements, measured encoding performance, and cost models based on the client's encoding costs, CDN pricing, and viewing distribution.

The three technical reports documented the analysis supporting these recommendations, while the accompanying executive summary organized them into implementation priorities for each workflow. Some recommendations could be implemented immediately, while others identified architectural changes and industry developments that should be revisited as device support, packaging practices, and codec licensing evolve.

The resulting methodology provides a repeatable process for evaluating future codec and workflow decisions. By establishing comparable quality, optimizing encoding performance, quantifying deployment costs, and considering operational and licensing constraints, future technology decisions can be evaluated using the same framework as streaming requirements and business conditions change.

About the Consultant

Jan Ozer has been encoding video since the mid-1990s and is an established authority in H.264, HEVC, VP9, and AV1 compression. His books include *Video Encoding by the Numbers: Eliminate the Guesswork from Your Streaming Video* and *Learn to Produce Video with FFmpeg in 30 Minutes or Less*.

Ozer has tracked video codec licensing policies since 2008 and published the industry's first structured comparison of the Access Advance and Avanci patent pools (Access Advance vs. Avanci Video: How the Streaming Programs Calculate Royalties). Though no longer practicing in either discipline, Ozer's numbers-based approach and economic analysis are grounded in his legal and accounting background, holding both a law degree and CPA.