ateme
Transforming Video Delivery

GOING GREEN:

HOW TO REDUCE THE CARBON FOOTPRINT OF STREAMING & SERVICE PROVIDERS

White Paper - September 2021

Table of contents

I. INTRODUCTION

There is growing concern over how video streaming contributes to global warming.

As governments, scientific institutions and UN agencies warn us of the dire consequences of excess greenhouse gas emissions, an all-hands-on-deck attitude toward the need for remedial action on the part of all sectors has taken hold, no matter how small their contributions to the overall problem might be.

According to the International Energy Agency, 72% of the energy consumed in streaming is attributable to user devices, 23% to transmission and 5% to datacenter-based processing[1].

This means a big share of the streaming energy footprint, that of devices, is beyond the control of service providers. But that still leaves a big 28% chunk – just over 0.6% of global greenhouse gas emissions – that service providers can do a lot about. Moreover, the non-device elements are sure to become an ever-larger share of the total as new formats and usage trends demand ever more energy for encoding and storage in core facilities and CDN caches.

According to one summation of recent findings produced by the BBC, the Internet – including the energy consumed by connected devices, storage and distribution infrastructure, and processing systems – accounts for about 3.7% of global greenhouse gas emissions. This is similar to the contributions of the global airline industry. Video streaming represents about 60% of all Internet traffic, so it is a big contributor to the total.

So what will the future bring for OTT providers?

We at ATEME are confident providers can continue riding the streaming tidal wave well into the future, with greater control than ever over their carbon emissions. This is because there are more efficient transcoding and streaming delivery solutions that minimize energy consumption in video delivery.

This upbeat perspective could not have come at a better time. While many other power-consuming activities contribute to the media and entertainment industry's carbon footprint, the carbon impact of the global surge in video streaming has gained media attention. As research conducted by Netflix shows, the need to act is clear:

"One hour of streaming on Netflix in 2020…[is] equivalent to driving a gas-powered passenger vehicle a quarter mile (or 400 meters). These results are consistent with our peers and validated by our independent advisory group."

Emma Stewart, Ph.D., Sustainability Officer, Netflix[2]

[1] https://www.iea.org/commentaries/the-carbon-footprint-of-streaming-video-fact-checking-the-headlines

[2] https://about.netflix.com/en/news/net-zero-nature-our-climate-commitment

Interestingly, these findings do not include the energy consumed by the content delivery networks (CDNs) Netflix uses.

While the findings disprove earlier assertions of much higher gas-consumption equivalency, they still point to significant energy consumption. With 204 million subscribers spending an average of two hours per day watching the service[3], the metric suggests that the annual global viewing of Netflix is equal to the carbon footprint of 37.2 billion passenger car miles.

[3] https://variety.com/2019/tv/news/netflix-cindy-holland-subscribers-watch-average-two-hours-day-1203159868/

Choosing transcoding and streaming technology that can reduce the carbon footprint of content distribution is a no-brainer for any company that streams video, whether for audiences at a global or regional scale. The benefits are the same: reduce carbon emissions while significantly saving on power, hardware and bandwidth costs.

It's easy to see how emission control can be directly paired with financial health in the streaming business.

The bottom line: every innovation that drives more efficient use of processing, storage and network resources must be the top priority – from both the bottom-line and the green credentials that all businesses will be measured on.

And such innovations are even more powerful when combined. Incremental gains in efficiency at different points of the video-delivery process quickly add up. This is what service providers can expect when they deploy the integrated cloud-based software components that make up the ATEME Green Delivery solution. As illustrated in Figure 1, combining ATEME's TITAN transcoders with ATEME's NEA platform gives end-to-end energy savings exceeding 60% against previous industry norms.

Figure 1:

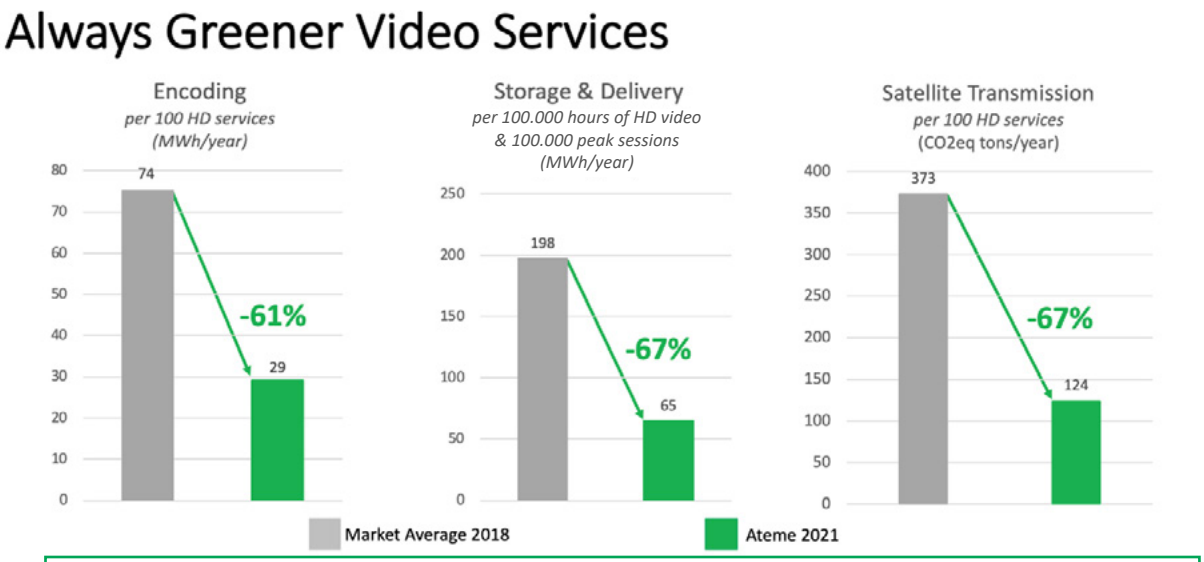

TITAN & NEA innovations: a huge energy reduction (two-thirds) in 3 years for OTT delivery, while improving the viewing experience

How does this happen? This paper explores the approaches to transcoding, packaging, and CDN infrastructure that enable these huge energy savings – and the trends that are making it essential to adopt these approaches.

II. TRENDS: CHANGING CONSUMER HABITS

VIEWERS SWITCH TO STREAMING SERVICES

The most important trend is the ongoing shift from traditional pay-TV to streaming services. A recent research report from Digital TV Research projects annual pay-TV revenues will drop from $173 billion in 2020 to $143 billion in 2026[4]. On the streaming side, the same research predicts revenue generated by SVOD and AVOD services will reach $210 billion in 2026, nearly doubling the $106 billion total reached in 2020[5].

Figure 2 shows the OTT vs. pay TV subscription trend in the U.S., which leads the world in cord cutting. Even there, judging by Digital TV Research's numbers, the OTT market is nowhere near maturity, with revenues expected to go from $42 billion in 2020 to $88 billion in 2026.

Figure 2:

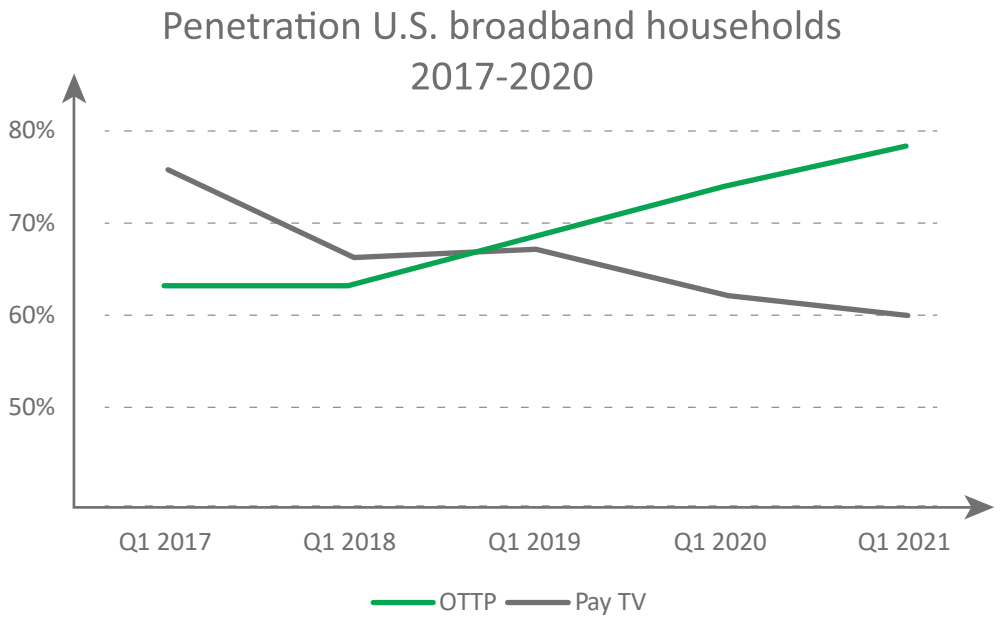

Source: Parks Associates[6]

Globally, the amount of time viewers spend watching streamed video, including everything from clips to movies, is on the rise as well, jumping by 36% from Q1 2020 to Q1 2021, as tracked by the online metrics supplier Conviva[7]. The increase was astronomical in some regions, including South America (240%), Africa (149%) and Europe (122%).

[4] https://www.rapidtvnews.com/2021051760476/global-pay-tv-revs-to-take-a-tumble-of-30bn-by-2026.html?utm_campaign=global-pay-tv-revs-to-take-a-tumble-of-30bn-by-2026&utm_medium=e-mail&utm_source=newsletter_2696#axzz6vDHNcAhI

[5] https://www.digitaltvresearch.com/products/product?id=321

[6] http://www.parksassociates.com/services/ott-tracker

[7] https://pages.conviva.com/rs/138-XJA-134/images/RPT_Conviva_State_of_Streaming_Q1_2021.pdf

As consumers switch to online video services, the connected TV (CTV) is emerging as the dominant viewing platform. According to Conviva, viewing on CTVs, including smart TVs and TVs connected to IP set-top boxes, accounted for 73% of streamed video viewing time in Q1 2021, up from 71% a year earlier. Cisco predicts that by 2023 there will be 3.2 billion CTVs in operation worldwide[8].

Trends in viewership are pushing bandwidth requirements to levels that were unimaginable not very long ago. Along with the bitrates required to serve CTVs, increases in the amount of time multiple users in broadband households are watching streaming video are profoundly impacting downstream bandwidth usage. (Figure 3).

Figure 3:

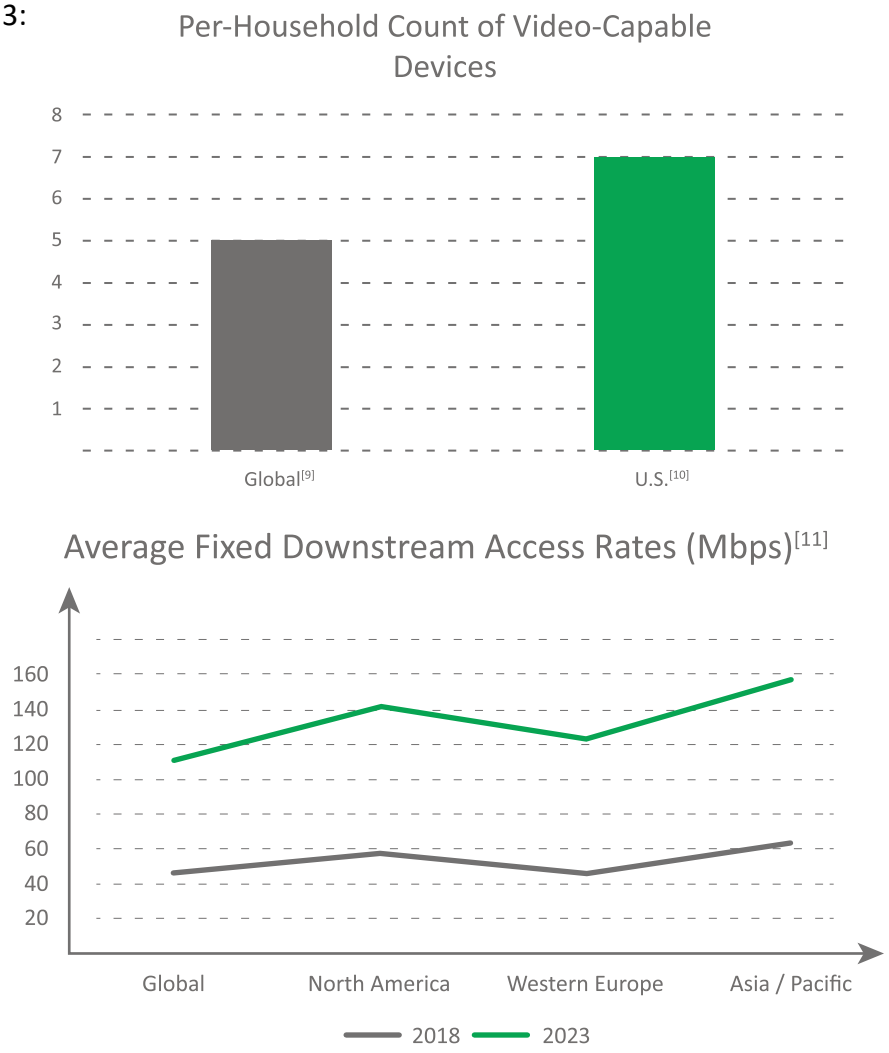

[8] https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html

[9] https://www.broadband-tvnews.com/2015/09/04/five-connected-media-devices-per-home-by-2019/

[10] https://variety.com/2019/digital/news/u-s-households-have-an-average-of-11-connected-devices-and-5g-should-push-that-even-higher-1203431225/

[11] https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html

III. THE ATEME GREEN DELIVERY SOLUTION: REDUCING ENERGY CONSUMPTION BY 67%

Clearly, video service providers have much to gain by adopting green strategies, from both a good citizenship and a public relations standpoint. But there's also a major cost benefit that comes with taking the steps to minimize the emissions impact of content distribution.

Just what those steps are can best be seen in the energy-reducing measures embodied in the ATEME Green Delivery solution. As illustrated in Figure 1 in the Introduction, ATEME's encoders used to process HD content in 2021 consume energy at a rate of 29 megawatt hours (MWh) per 100 services annually, representing a 61% reduction from the market average in 2018.

Over three years, ATEME's encoders achieved a 61% reduction in energy consumption from the market average.

These efficiencies, combined with those enabled by ATEME's NEA delivery solution, have reduced energy consumption in storage and delivery by 67% over the same time-frame. And energy consumption in satellite transmission has been mitigated by the same percentage.

ATEME's NEA delivery solution has reduced energy consumption in storage and delivery by 67%.

But how did ATEME manage to achieve such huge savings in the energy consumption of its encoders? There are many aspects to this:

1. Thanks to its expertise in codec development
2. By leveraging microprocessor evolutions
3. Through accelerated distributed parallel processing
4. Through content-aware encoding
5. Using Film Grain Synthesis
6. By leveraging cloud-native microservices efficiencies.

EXPERTISE IN CODEC DEVELOPMENT

With more and more video being consumed online, there is a need for new codecs offering better compression to address the greater bandwidth requirements. Better compression also reduces VOD storage requirements, and cascades to reducing cache storage requirements in CDN distribution.

Indeed, there's a multiplier effect to storage capacity savings, since every piece of efficiently encoded content moves into multiple edge caches across the service provider's primary and backup CDNs. And as each content file transits Internet backbones and peering points to enter CDNs, service providers reap savings in energy consumption and costs there as well. These workflows will become even more important as 5G adoption takes off.

This need for more codecs continues to be met, which explains the multitude of codecs currently available: besides the legacy MPEG-2 and the current mainstream standard HEVC, we have H264, AV1, and the upcoming successors of HEVC – VVC, EVC and LC-EVC, to name a few. This presents a challenge to the companies developing streaming solutions: how can they not only stay up to speed on the latest codecs, but also be able to implement them quickly into their products and solutions?

ATEME is well equipped to face this challenge. It has an unmatched record in encoding advances that have allowed service providers to consistently achieve the highest levels of video quality at the lowest bitrates. These achievements rest not only on the company's laser-focused dedication to encoding technology, as reflected in the expertise of its personnel and in its commitment to engaging in the activities of standards bodies (see figure 4), but also on the speed at which the company can implement new codecs into its solutions as it leverages its full control on its modular software stack, which prevents it from relying on third parties.

By staying closely attuned to which ideas in each succeeding version of existing codecs as well as new codecs are being adopted, ATEME quickly allocates internal resources to deliver the advances that matter most to customers. Every codec supported by ATEME encoders uses standards-compliant algorithms that ATEME builds from scratch to achieve the highest possible efficiencies in on-demand and live content applications.

This is a major difference between ATEME and many other suppliers of cloud-based encoders. While a codec standard is a language enabling the encoder to describe video in as few bits as possible, the ability to use that language to its maximum effect, thereby maximizing the bit reduction with the most efficient use of compute power, depends on designing the most efficient algorithms.

Figure 4:

ATEME's current involvement in codec and standards bodies

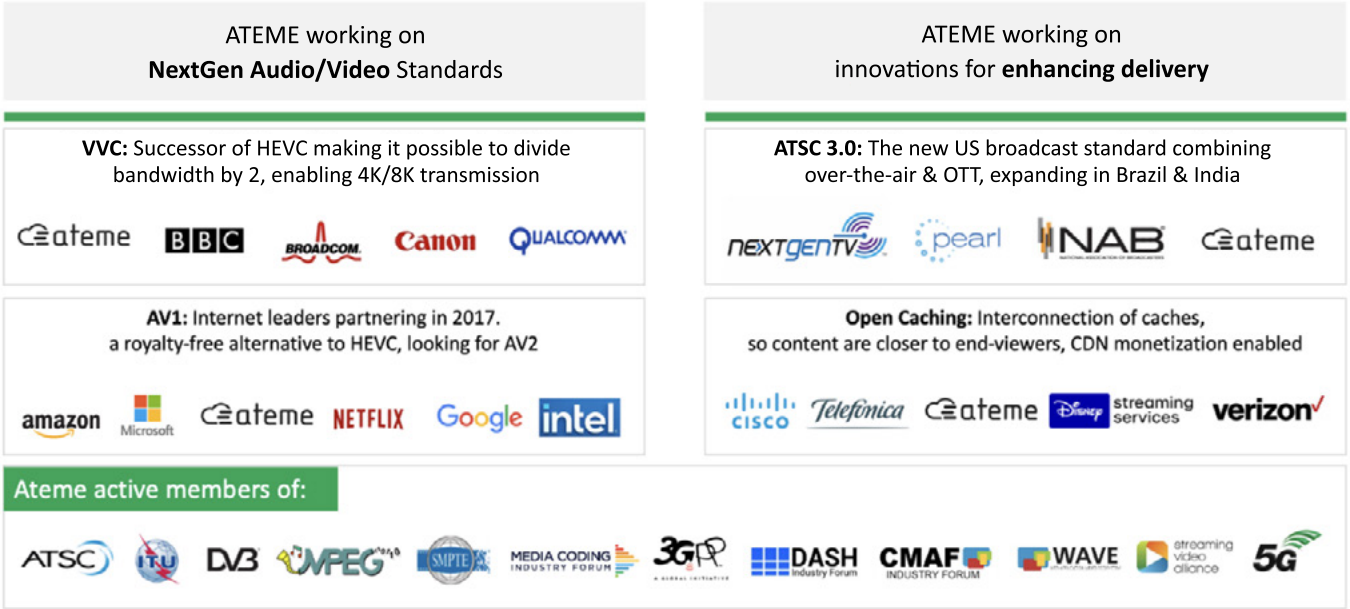

Customers can expect further energy reduction as ATEME supports the AV1 codec today and prepares to support VVC, enabling additional bitrate savings of 20% and 39%, respectively, compared to HEVC. Of course, these gains are measured against HEVC in the context of all the advances ATEME has applied in making TITAN the efficiency leader in HEVC encoding.

LEVERAGING MICROPROCESSOR EVOLUTIONS

Another way that ATEME increases the efficiency of its encoders is by leveraging microprocessor evolutions.

Over the past 15 years, chips have repeatedly surpassed predicted processing limits. As predicted by Moore's Law, CPU integrated circuit (IC) processing capacities have soared with the biennial doubling of transistor counts per area of silicon, going from hundreds of millions per IC used in video processing 15 years ago to billions today. So we expect that commodity processors will continue increasing their densities. As a matter of proof, the table below demonstrates how ATEME has reduced power consumption per channel at the chipset level by 80% in 8 years.

Figure 5:

Power Consumption Reduction with ATEME Technology

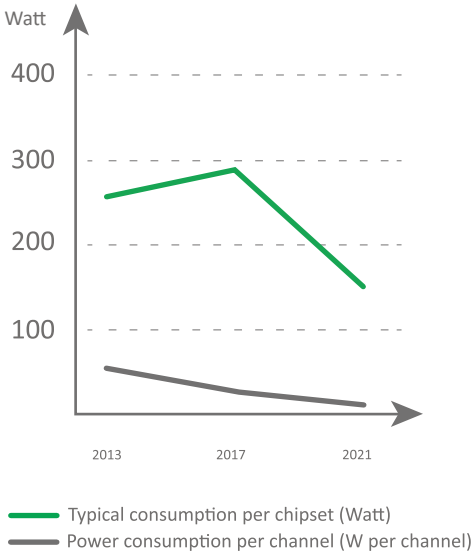

Figure 6:

Transcoding Density Increase with ATEME Technology (number of channels)

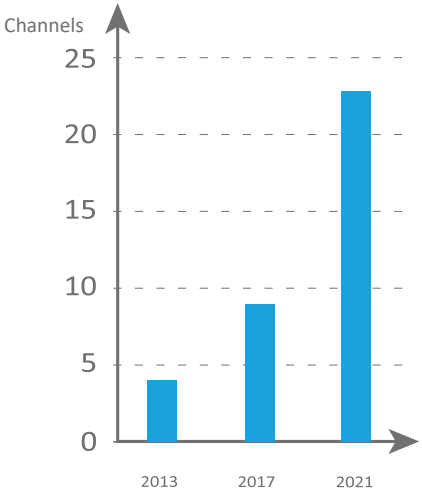

Of course, raw processing capacity is just one of many factors that go into choosing processors. ATEME capitalizes on advances in elastic computing technology by remaining vendor-neutral when choosing the processors to be used with succeeding enhancements to the TITAN encoders.

ACCELERATED DISTRIBUTED PARALLEL PROCESSING

ATEME's encoding efficiency depends on the elastic computing architectures of multi-core ICs that have made cloud-based encoding the go-to option in today's streaming market. This has enabled accelerated parallel processing techniques to achieve efficiency gains.

For example, TITAN encoders use algorithms that minimize the number of computations needed to predict patterns in future frames based on analysis of patterns in immediately processed frames. The better the encoder is at these functions, the better it is at minimizing the bitrates for delivering content at a given level of quality (or maximizing the quality level that can be achieved at an established bitrate). And this, of course, means less bandwidth and storage requirements – so ultimately, a lower carbon footprint.

Other techniques contribute to reducing encoding time. If a library of media assets needs to be encoded faster than real time, each asset can be divided into multiple chunks, and all the chunks can be encoded in parallel with no compromise in the quality of output.

CONTENT-AWARE ENCODING TO HALVE STORAGE SPACE

Through content-aware encoding, TITAN can achieve up to 50% reduction in storage requirements. This is one of three of ATEME's technologies that won a Technology & Engineering Emmy® award this year.

Indeed, another major factor in the efficiencies achieved with ATEME encoding involves the use of AI with the ATEME Quality Index (AQI) perceptual quality metric, which is an automated means of objectively measuring video quality that closely mimics how humans perceive quality. AQI is directly incorporated into TITAN.

This is a far more meaningful way of assessing encoding quality than the conventional Peak Signal-to-Noise Ratio (PSNR) and Structure Similarity Index (SSIM) measures. AQI is tuned against a large, proprietary database of tens of thousands of video-quality test encodings applicable across all codecs and resolutions supported by TITAN encoders. Machine learning (ML) techniques were used to train the AQI algorithms to match the encoder outputs to the visual quality evaluations of test participants, many of whom are highly skilled Golden Eye video engineers.

As illustrated in Figure 6, the AQI metrics provide quality assessments across five vectors, including:

- Encoding Quality, which is an objective metric that reflects raw encoder efficiency before other factors are taken into consideration.
- The Spatial Index, which describes the expected quality loss incurred by resolution down-sampling in relation to viewing parameters (e.g. viewing distance) and content spatial features.
- The Temporal Index, which leverages content motion analysis to describe the impact of framerate reduction, given certain viewing parameters.
- The Dynamic Range Index, which calculates ratios between brightest and darkest parts of the picture, and the impact on perceived quality of tone-mapping HDR content to SDR formats.
- The Color Gamut Index, which describes how transformations relating to the content's color characteristics will impact the perceived quality of the output.

Figure 7:

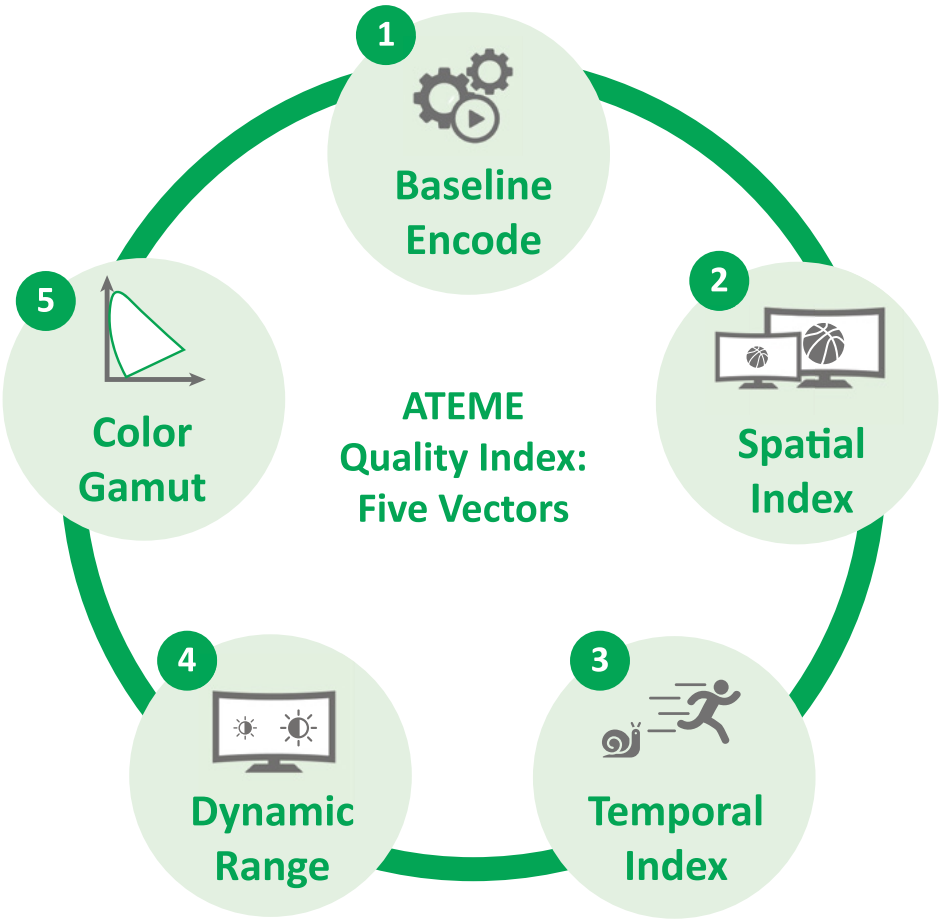

ATEME's AI platform uses AQI metrics for greater efficiency through content-aware encoding, which makes it possible to minimize the scene-by-scene bitrates assigned to each ABR profile. At each scene transition in the bitstream, the AI platform assigns an encoding bitrate precisely matched to the service provider's intended quality parameters as measured by the AQI.
Additional efficiencies are achieved by aligning the intelligent scene-change detection algorithms with the ABR fragment boundaries, which minimizes the use of I-frames. Together these techniques result in fewer variants and/or lower bitrates for each variant, meaning up to a 50% reduction in storage space at the core and at each CDN cache.

By combining artificial intelligence with the ATEME Quality Index (AQI) and intelligent scene-change detection, TITAN encoders achieve up to a 50% reduction in storage space.

FILM GRAIN SYNTHESIS FOR 90% SAVINGS

TITAN encoders deliver film grain content with 90% savings in processing energy.

The question of how to deal with film grain, which gives productions their distinctive look and artistry, has long been a nettlesome issue for streaming service providers. While there is value in preserving the effect, the random noise constituting film grain cannot be compressed in the encoding process. By nature it's temporally unpredictable, so it disrupts the predictive process used to maximize bitrate efficiency.

The process known as Film Grain Synthesis (FGS) puts this issue to rest. The long-ignored support for FGS built into the AVC and HEVC standards but never executed in players is gaining attention. Service providers using the AV1 codec will automatically benefit from FGS insofar as support is mandated in both the encoders and players used with that platform.

With AV1 still in the nascent stage of adoption, ATEME is working with SVOD and AVOD providers to support the technology in their players with its activation in TITAN encoders. Rather than encode the noise, TITAN analyzes its distribution across the frames and its intensity relative to the primary signal, then replicates it algorithmically. With this technology, ATEME won a Technology & Engineering Emmy Award in January 2021.

Players equipped with the necessary algorithmic function add the noise in close approximation to the original film during the decoding process. The result is a movie-viewing experience that can deliver a 90% reduction in the processing that would have been used to encode the noise.

With Film Grain Synthesis, TITAN delivers a 90% reduction in processing power.

LEVERAGING CLOUD-NATIVE MICROSERVICES EFFICIENCIES

Adding to all the encoding efficiencies mentioned above is the significant efficiency gained when using datacenter resources intrinsic to container-based cloud-native virtualization with the cluster management system known as Kubernetes. This allows the myriad of TITAN processes to be executed on servers by the semi-autonomous OS modules known as containers.

All the microservices associated with a particular application can boot up or be taken down in a fraction of a second. This means hardware resources can be allocated to containers supporting different applications as needs change. This scaling flexibility translates to a far more efficient use of datacenter resources. According to 451 Research, container-based virtualization enables double-digit savings in resource use compared to the approach based on virtual machines, even with relatively simple implementations[12].

[12] https://robin.io/resources/451-research-containers-economically-better/

Automated load balancing extends to all the containers in the cluster to ensure uninterrupted operations whenever failures occur. And ATEME enhances the power consumption efficiencies by enabling customer supervision of multi-server transcoding farms from their TITAN operations portals.

MICROSERVICE-BASED EFFICIENCIES OF AN ELASTIC CDN

ATEME's cloud-based NEA CDN platform extends the use of Kubernetes technology all the way to the network edges. This creates more datacenter resource savings, and the efficiencies quickly add up. Whether service providers rely exclusively on public cloud or on in-house or hybrid infrastructures, they can orchestrate many functions when putting this highly elastic CDN as near as possible to the end user.

Figure 8:

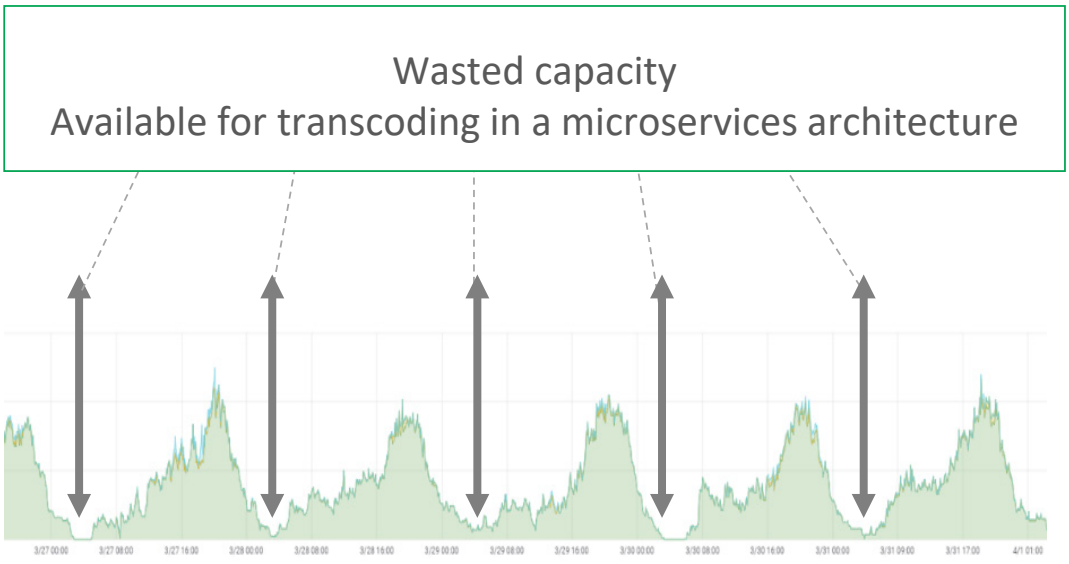

Service providers thereby avoid the unused capacity costs and functional limitations of traditional CDN technology. They can also reduce power consumption and storage and bandwidth requirements by:

- Orchestrating time-shifting with analytics for efficient use of cache and storage resources for trick play, catch-up and VOD. For example, the use of analytics can help to predict which content will be required where, enabling more efficient and targeted distribution of content across cache nodes. So instead of provisioning the same content (almost) randomly across nodes, it is provisioned only where it is likely to be needed. This greater caching efficiency leads to reduced storage and bandwidth requirements, while also increasing the quality of the viewing experience without increasing power consumption;
- Using just-in-time (JIT) packaging. This reduces the number of profiles to be stored and sent, as only profiles actually requested by viewers will be packaged;
- Maximizing performance by using virtualized resources for load balancing and redundancy across edge points. Again, this leads to a more efficient use of resources, thereby reducing both storage and bandwidth requirements and also hardware resources;

HOW THE COMBINATION OF NEA AND TITAN DELIVERS REAL EFFICIENCIES

These capabilities can put a big dent in service providers' carbon footprints. But the efficiencies enabled by TITAN encoding and NEA origin, packager and CDN go much further.

They begin with the NEA's energy-reducing contributions to the encoding process itself.
The CDN platform's JIT packaging and encapsulation of encoded content in the Common Media Application Format (CMAF) enable substantial savings in the use of encoding resources.

THE PROMISE OF CMAF AND JUST-IN-TIME PACKAGING: HALVING STORAGE & TRAFFIC REQUIREMENTS

CMAF has the potential of bringing huge energy savings by halving both traffic and storage requirements. It specifies a fragmented MP4 (fMP4) container that can be used by the two dominant adaptive streaming formats: Apple's HTTP Live Services (HLS) and MPEG-Dynamic Adaptive Streaming over HTTP (DASH). This eliminates the need to maintain separate storage and processing silos to enable content streaming via these two modes.

For content distributors supporting CMAF, that means the volume of traffic and cache storage required to deliver that content to those devices can potentially be cut in half.

In the case of encrypted content, CMAF works with MPEG Common Encryption (CENC) to enable DRM-interoperable delivery with ad signaling, closed captioning, subtitles and other enhancements. CENC can assign encryption keys for Apple's FairPlay DRM and any keys supported by CENC in DASH.

This is possible because CMAF encapsulation is compatible with HTML5 Encrypted Media Extensions (EME), which defines a common API that can be used with CENC to discover, select and interact with DRMs. Rather than defining DRM functionality, EME standardizes the discovery hooks, moving the responsibility for such interactions from plugins or third-party applications to HTML5-compatible browsers.

With support in Apple iOS and Android devices since 2017, CMAF has reached

With the NEA low-latency JIT packager, content distributors can expect a further 50% savings in the energy consumed for delivering video.

significant device penetration worldwide. This enables a unitary streaming approach to delivering unencrypted video. Service providers who immediately adopt CMAF in their streaming operations will be able to exploit the efficiencies as the market advances. For content distributors supporting CMAF, that means the volume of traffic and cache storage required to deliver that content to those devices can potentially be cut in half.

As streamed services become more feature-rich and geographically diverse, and support ever more display formats, we expect CMAF to deliver even more efficiencies.

That's what CMAF brings when 100% of the workflows have been standardized. Of course, this may take a while and we might have to live with several formats (including CMAF) for some time. Does that mean that for now, we have to carry on encoding and storing all profiles in all formats?

The answer is no. Encoding and storage efficiencies can be greatly enhanced by using ATEME's NEA Just-In-Time (JIT) packaging solution. Operating in pull mode, this solution prepares content for delivery to each requesting device on an as-needed basis. This eliminates the need to encode and store each content file in all packaging formats.

GOING EVEN LEANER WITH ANALYTICS

While clever design can bring significant efficiencies that reduce the carbon footprint of video delivery, knowing how video is consumed can multiply these efficiencies.

In a cloud-native, elastic CDN, the size of the cache can fluctuate to meet consumer demand. But this will only work if there is a fairly reliable prediction of what

consumer demand will be. So by combining an elastic CDN with analytics using artificial intelligence that can predict consumer demand, you can tap into the full cost-saving (and energy-saving) potential of the elastic CDN.

Moreover, while JIT packaging eliminates the need for packaging in multiple profiles, if you go further upstream, the content still needs to be transcoded in all the profiles.

By combining ATEME's clever design – better compression, JIT packaging, single workflow for both live and time-shifted video, and elastic CDN- with analytics, service providers eliminate waste in video delivery. They become lean and green.

Imagine now if the service provider could predict which resolutions and bitrates would actually be requested. With this knowledge, it could encode content only in the required profiles.

Furthermore, while cache only stores content once it has been requested, the ATEME CDN design contains multiple origin-server clusters. These clusters require some redundancy to make the system robust. But with precise data on which nodes will need which content in which profiles, the service provider can significantly reduce the amount of redundancy, streamlining video delivery so that each cluster has exactly what it needs. Again, this brings significant energy efficiencies by reducing the number of servers and bandwidth consumption in the core, while maintaining a high quality of experience for viewers.

All of the above scenarios are in fact possible thanks to ATEME's own server-side analytics, and the integration of ATEME's solutions with a broad range of industry-leading client-side analytics that can provide the precious data on the behaviors of viewers.

VII. CONCLUSION

OTT video service providers can significantly limit their impact on climate change by adopting energy-saving approaches to encoding and distributing their content, which amount to three strategies:

- Reducing storage requirements
- Reducing bandwidth consumption
- Reduced hardware requirements.

ATEME has made this possible by integrating its cloud-based TITAN and NEA platform into a holistically managed end-to-end Green Delivery solution, which reduces storage, bandwidth, and hardware requirements through:

- Better compression, for sending and storing the same quality in fewer bits;
- Just-in-Time packaging, for sending and storing fewer profiles;
- A single workflow for both live and time-shifted video, for sending and storing the same profile only once instead of twice;
- An elastic CDN, to drastically reduce the CDN platform's hardware requirements;
- Analytics, to transcode only what is really needed.

Now service providers can orchestrate all the elements of these best-of-breed components to minimize their carbon footprints.

What's more, energy savings also contribute to the bottom line – both by reducing consumption of encoding resources, and by reducing bitrate contributing to lower storage costs and greater bandwidth efficiency. So going green is by far the best business strategy. Benefits to both agendas will increase rapidly as new codecs and growing service consumption call for ever greater use of cloud resources.

ATEME leads the market in driving encoding efficiency through accelerated parallel processing, content-aware encoding, and other innovations. With these, service providers can leverage new codecs and keep bitrates in check with the least possible increases in CPU power consumption.

And in encoding, storage, packaging and CDN distribution, service providers can optimize their use of datacenter resources by using ATEME's container-based microservices.

The green way forward is clear: drive maximum efficiency in processing to minimize storage and bandwidth resources by implementing the end-to-end distribution infrastructure that is designed to do just that.

VIII. REFERENCES

[1] IEA, The Carbon Footprint of Streaming Video, December 2020

[2] Netflix, Net Zero + Nature: Our Commitment to the Environment, March 2021

[3] Variety, Netflix's Cindy Holland Says Subscribers Watch an Average of Two Hours a Day, March 2019

[4] Rapid TV News, Global Pay TV Revs To Take a Tumble of $30 BN by 2026, May 2021

[5] Digital TV Research, Global OTT TV and Video Forecasts, May 2021

[6] Parks Associates, OTT Video Market Tracker, November 2020

[7] Conviva, State of Streaming Report, April 2021

[8] Cisco, Annual Internet Report, March 2020

[9] Broadband TV News, Five connected Media Devices per Home by 2019, September 2015

[10] Variety, Households Have an Average of 11 Connected Devices, December 2019

[11] Cisco, Annual Internet Report (2018-2023), March 2020

[12] 451 Research, Containers, December 2018

Authored by

Remi Beaudouin, Chief Strategy Officer, ATEME

Jean-Louis Lods, VP Media Supply Chain Solutions, ATEME