solution series

ADVANCED ENCODING & TRANSCODING

codecs
DRM
encoding
formats
captions
ABR technologies

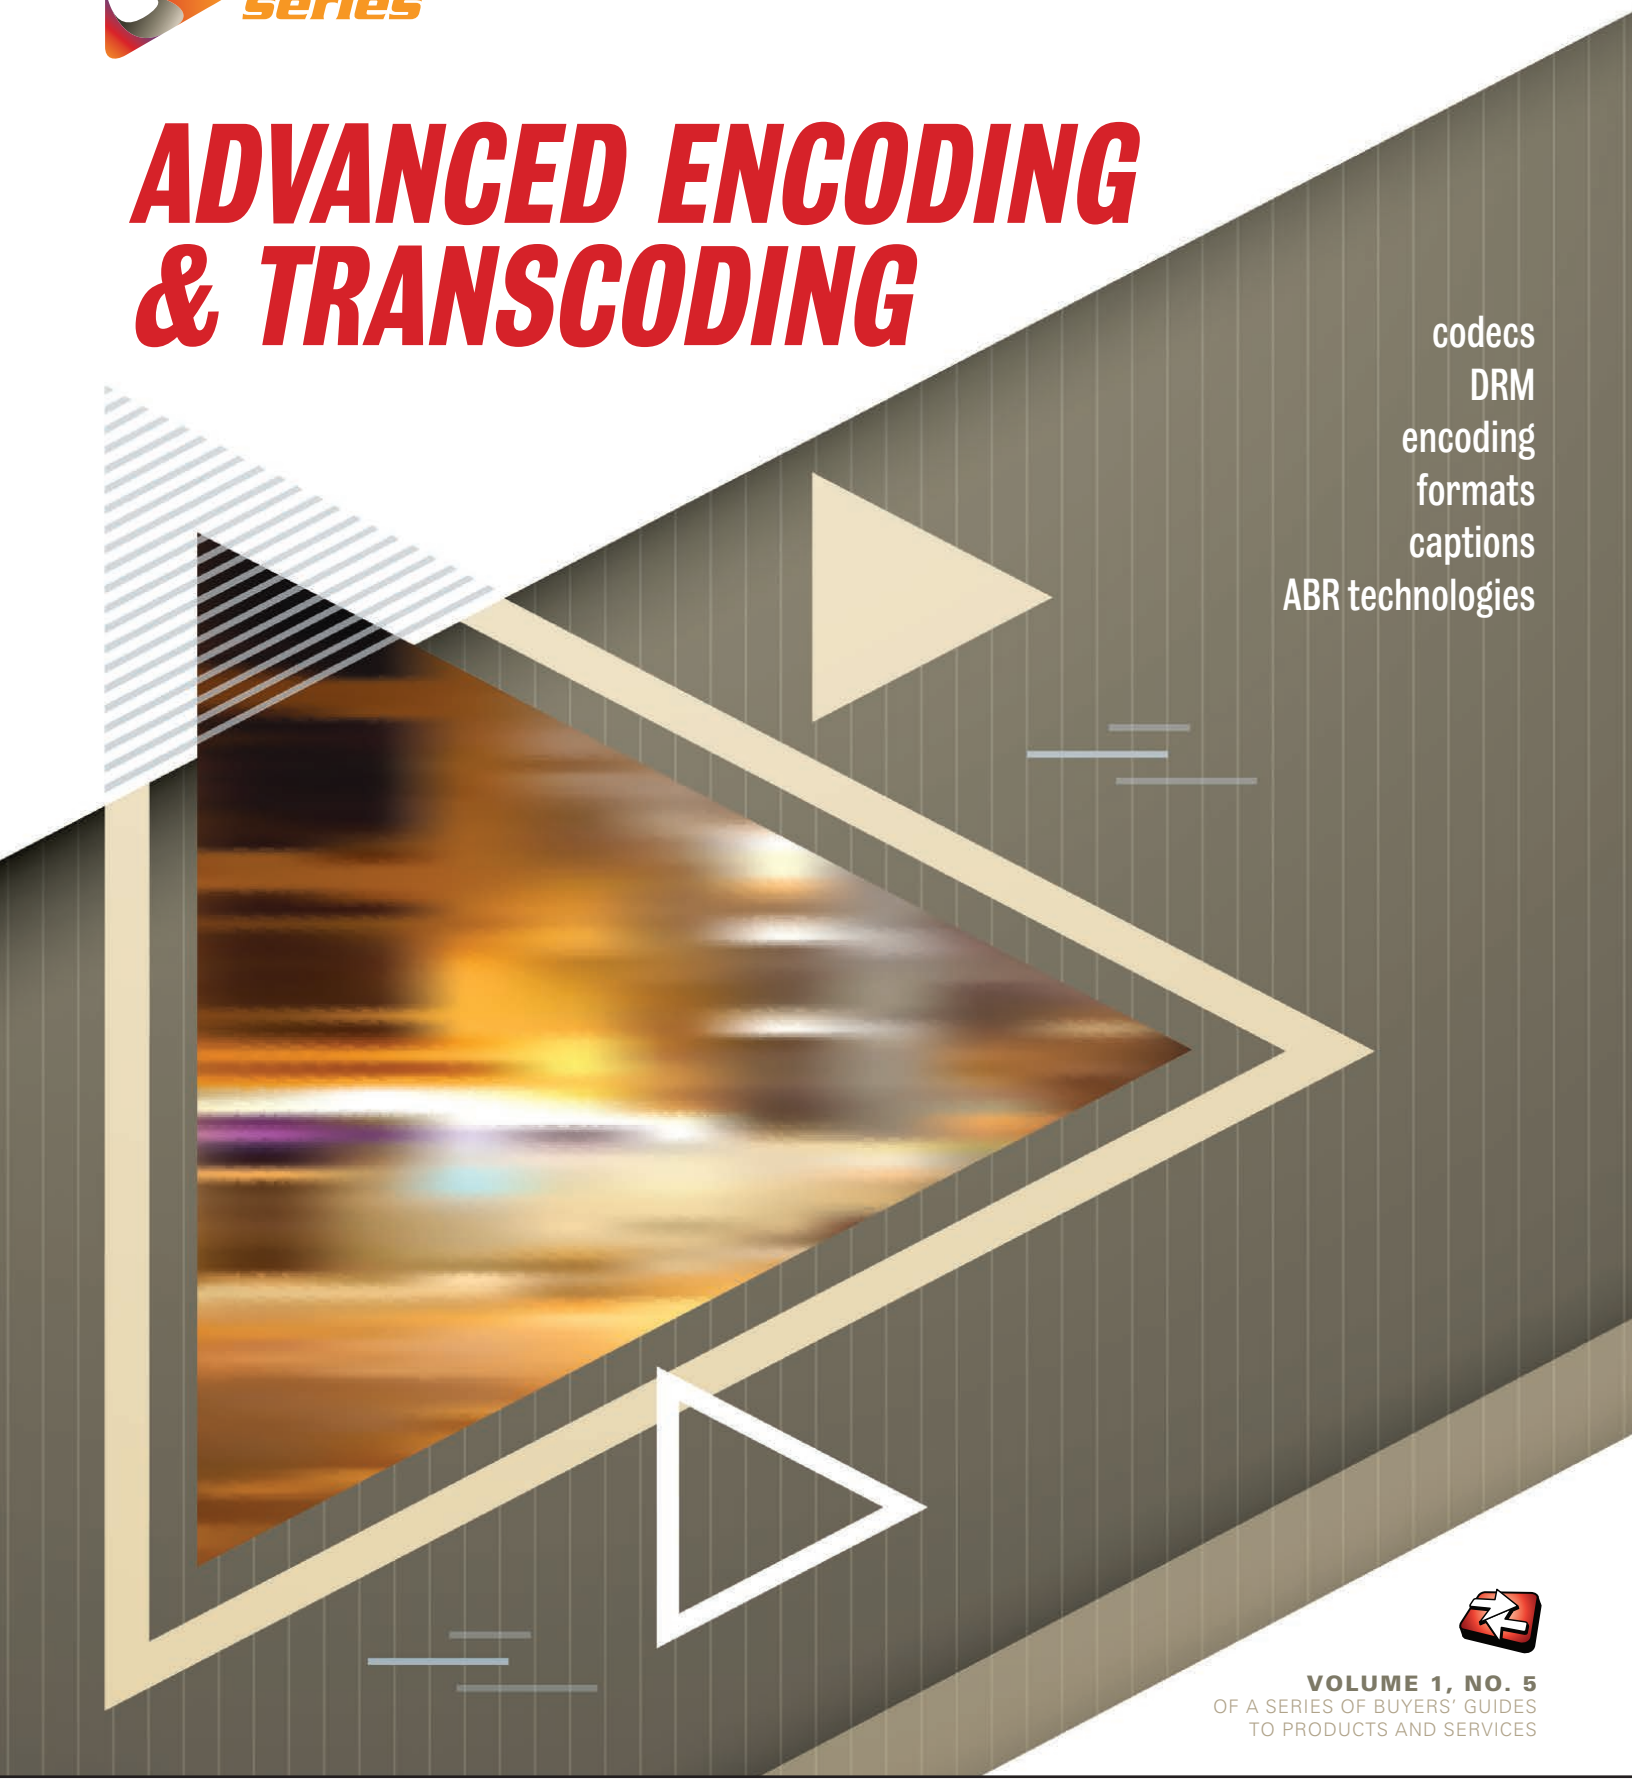

VOLUME 1, NO. 5
OF A SERIES OF BUYERS' GUIDES
TO PRODUCTS AND SERVICES

TABLE OF CONTENTS

streaming media west

the business and technology of online video

November 19-20, 2019

PRECONFERENCE WORKSHOPS
MONDAY, NOVEMBER 18

- Learn cutting-edge streaming tech and techniques—from content capture and creation to distribution and management.
- Get in-depth training on online video technology and applications in intensive preconference workshops.
- Get up-to-speed on the latest business models and distribution strategies for online video and OTT.
- Gain access to the Exhibit Hall featuring major streaming and online video vendors.
- Network with your colleagues, peers, and experts in streaming media.

STREAMINGMEDIA.COM/WEST

FEATURING THESE SPECIAL EVENTS:

live streaming SUMMIT | OTT LEADERSHIP SUMMIT | Video Engineering Summit | ESPORTS & SPORTS STREAMING SUMMIT | STREAMING MEDIA UNIVERSITY

Diamond Sponsor

MediaKind

Platinum Sponsors

applicaster | THEO

Gold Sponsors

BITMOVIN | SSIMWAVE | UI CENTRIC | XILINX

Silver Sponsors

interra systems | verizon media

On-Demand Video Sponsor

brightcove

Association Sponsor

MoCA Multimedia over Coax Alliance

Media Sponsors

Church Production | ProductionHUB | streaming media magazine | streaming media producer | TECHNOLOGIES FOR WORSHIP | Videomaker

ORGANIZED AND PRODUCED BY

Information Today, Inc.

PUBLISHERS OF

streaming media

CONNECT: #StreamingWest

NETINT

Edge Encoding for the Video Cloud

As the industry adapts to the ongoing growth in video streaming volumes and applications, cloud computing and video distribution companies need to reconsider the locations for expanding video encoding capabilities in their networks. Historically, video encoding was most often processed inside centralized data centers to consolidate operations and to achieve economies of scale. However, the costs associated with video distribution, along with the low latency encoding requirements for emerging interactive video applications, is motivating video distribution network architects to reevaluate their edge encoding options.

VIDEO ENCODING ESSENTIAL FOR STREAMING SERVICES

Video streaming is usually delivered to end user devices using adaptive bitrate (ABR) technologies. For ABR to work, a set of video encoding ladders needs to be prepared from each high-bitrate source video input. As a general benchmark, an encoding ladder is typically a set of ~6 to 10 encodings of incremental sub-bitrates and resolutions of the exact same video content, resulting in a total bitrate for each encoding ladder set equal to ~2.5x the original high-bitrate source video input.

ALTERNATIVE LOCATIONS FOR ENCODING IN THE VIDEO CLOUD

Centralized DC—Video encoding requires a great deal of computing hardware or specialized encoder equipment, especially for live video streaming, making centralized data centers (DCs) the natural choice to build encoding capacity with economies of scale. Operators can consolidate encoding operations to just a few centralized locations, leading to the lowest encoding infrastructure investment, but not necessarily the lowest overall operating costs. Meanwhile, ABR streaming servers are typically hosted closer to the users in regional data centers or points of presence (POPs) to minimize latency and maximize quality of experience for a concentration of video consumers in each region. The operational challenge is that distributing full encoding ladders from centralized encoding farms to the regional locations requires 2.5x the bandwidth and ongoing networking cost, compared to distributing only the original high-bitrate source video to each regional location.

Regional DC—In today's core networks, central and regional data centers are often connected with dedicated, high bandwidth point-to-point links which are expensive to operate - especially during peak video streaming periods. A regional encoding strategy would increase encoding processing CAPEX across a larger number of regional data center or POP

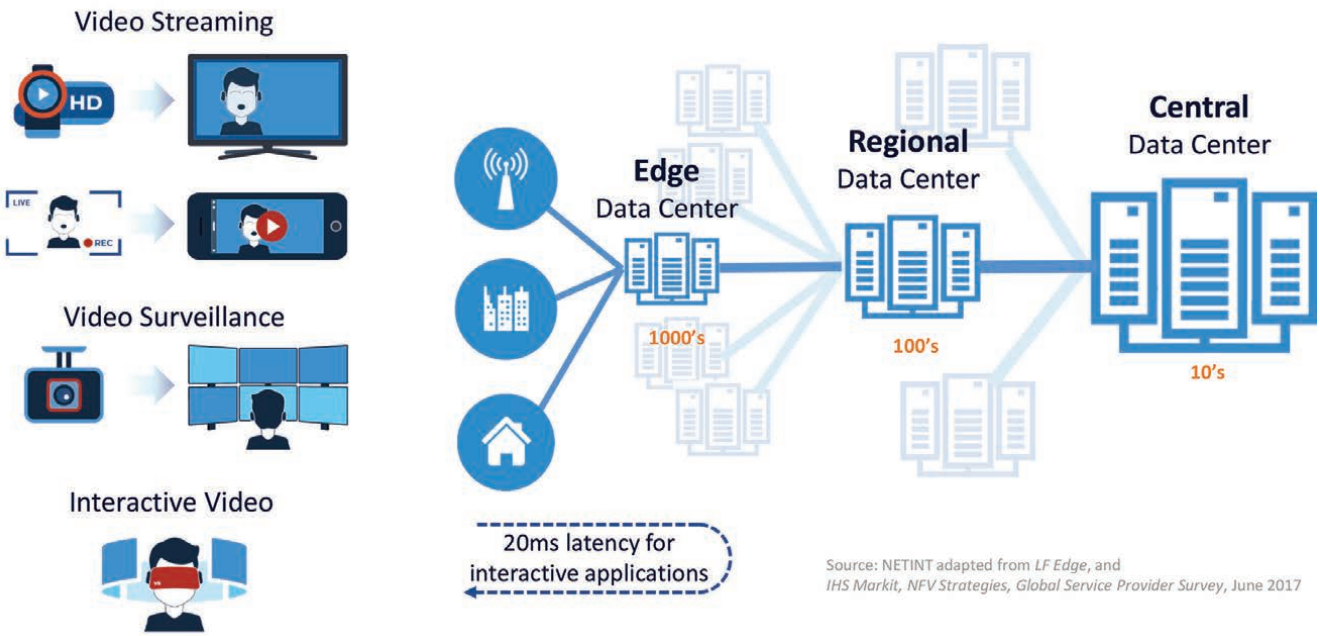

Source: NETINT adapted from *LF Edge*, and *IHS Markit, NFV Strategies, Global Service Provider Survey*, June 2017

NETINT

Encoding in regional data centers achieves the lowest TCO due to an ongoing monthly reduction in video distribution bandwidth and associated costs.

locations, but the benefit is only the original high bitrate source video is distributed to each region. Total cost of ownership (TCO) analysis over a multi-month planning period would show that a regional encoding strategy, when both video distribution and encoding costs are considered, would be more economical.

Edge DC—Many emerging interactive cloud-based video applications, such as cloud gaming, AR, VR, or 360° video, will require very low latency 20 ms feedback loops. Supporting these low latency requirements from centralized data centers will be difficult to achieve due to network propagation delay. Instead, edge data centers will be the preferred hosting locations of these next-generation applications and associated video encoding services. With the rollout of 5G services, mobile network operators are also deploying edge data centers close to 5G radio equipment to support the next generation of interactive or latency-sensitive mobile video applications and services.

Consider the low-latency processing required for a typical interactive video application, starting with a game controller or VR headset movement command sent into the cloud gaming engine and rendering for each player in edge gaming server, encoding the next video frame for each player, before sending the next frame back to each player's video display. These applications will demand high video encoding performance requirements, with solutions expected to generate UHD resolution for gaming or VR headset displays with ultra-low encoding latency.

TECHNOLOGY OPTIONS FOR VIDEO ENCODING

Today, video encoding services are often deployed as software processes running on compute resources. Variations on this architecture include cloud-native software that can be deployed on virtualized or containerized instances in the cloud or enterprise data centers. While software encoding on CPU is your most flexible option, it is also likely your most expensive option requiring the most servers, rackspace, power, and cooling, making this encoding strategy undesirable, if not impractical, for edge data center environments. Other options to improve encoding densities and reduce power consumption include software encoding solutions that run on graphic processor units (GPUs) or field programmable gate arrays (FPGAs). However, to achieve the density and low power targets of edge data center environments, application-specific integrated circuit (ASIC) encoding solutions are your best option.

EDGE ENCODING AT SCALE

The **Codensity T408 Video Transcoder** (go2sm.com/codensity) is a compact ASIC-based module delivering a staggering 4K/UHDp60 UHD encoding throughput. T408 modules are based on the standard 2.5" U.2 form factor, for plug-in provisioning into standard NVMe servers, delivering the highest density, high encoding quality, lowest power, and reduced rackspace requirements. FFmpeg can be used for integration with existing encoding ladder workflows, or an API can be used for specifying ultra-low latency configurations for cloud-based interactive video applications.

For applications requiring both encoding and storage in space-constrained edge data centers, the new **Codensity EdgeFusion E408** (go2sm.com/edgefusion) offers similar transcoding capabilities with integrated 3.5TB SSD storage in the same form factor.

Codensity™ EdgeFusion E408

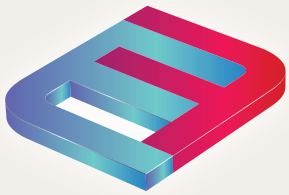

- Highest Density Encoding
- High Encoding Quality
- Least Rack Space
- Lowest Power
- Ultra-low latency
- FFmpeg or API integration
- Integrated 3.5TB NVMe SSD

In summary, the versatile Codensity T408 Video Transcoder and EdgeFusion E408 deliver the following capabilities and benefits for cloud computing and video distribution companies:

- Encoding at scale for centralized encoding,
- Lowest TCO with regional encoding,
- Interactive video encoding at the network edge, and
- Video storage at the most economical location.

To learn more, please download our latest whitepaper **Edge Encoding for the Video Cloud** (go2sm.com/netintwp), or visit **www.netint.ca**.

ABOUT NETINT TECHNOLOGIES

NETINT Technologies is a pioneer of innovative SoC solutions intersecting computational storage and video processing. Its Codensity portfolio enables cloud data centers, edge computing companies, and content providers to deploy scalable high performance applications, while minimizing their data storage and video processing costs. NETINT, founded by an experienced team of storage SoC veterans, is a Canadian venture-funded high-tech company with R&D facilities in Vancouver, Toronto, and Shanghai, China.

telestream

Today's Content Distribution Challenges Go Far Beyond Transcoding and Exporting Files

CHRIS OSIKA, CHIEF MARKETING OFFICER, TELESTREAM

The media & entertainment supply chain is under increased pressure as the number of devices and locations from which consumers access content continue to rise. According to a recent Cisco whitepaper (go2sm.com/ciscovni), one million minutes (17,000 hours) of video content will cross global IP networks every second by 2021. While new distribution agreements enhance the revenue of content owners and originators, there hasn't been a corresponding allocation of resources to fulfill these agreements. Transcoding and packaging of media is a substantial part of the media & entertainment supply chain process, particularly when prepping for a new outlet. However, it is far from the only issue. In any on-prem workflow, graphics are added/removed/altered, dialog is replaced, legislative advisories are added/removed/replaced, promos are inserted, and much more. Platform requirements further complicate the issue with adaptive bitrate (ABR) variants being optimized for different devices and closed captioning, dynamic ad insertion, sidecar content, and loudness standards to be adhered to. Working with international distributors creates still another set of demands focused mainly on localization of programs and movies. In many cases, these additions and alterations are performed by an automated "bag and tag" edit function running largely autonomously. (We certainly don't want people doing such repetitive, mundane work.)

KEEP SILOS ON THE FARM

Many organizations have set up distinct workflows for each distribution destination as the need has arisen. Frequently, these workflows will be comprised of different technologies and service providers that were available at the time to get the job done. This is not surprising, as each of the workflows comes with unique demands that content originators may not be equipped internally to provide. With these custom tools in place and independently deployed for each use case, organizations become challenged when moving to new, centralized workflows that can scale as required. These resultant siloed workflows are inefficient, difficult to maintain, and costly in terms of resources (machine and human).

IS THE CLOUD THE ANSWER? YES AND NO

The need to scale up media processing resources has never been greater as distribution contracts get signed, back catalogs are refreshed, and programs find new avenues to viewers. For a growing number of companies, though not all, the logical solution is to move the media supply chain up into "the cloud" and make use of the well-understood advantages that such architectures promise. The scalability of cloud processing, along with the idea of consumption pricing for the infrastructure and processing make this an easy decision in the minds of many CFOs, as moving more expenses to an op-ex budget can be very attractive.

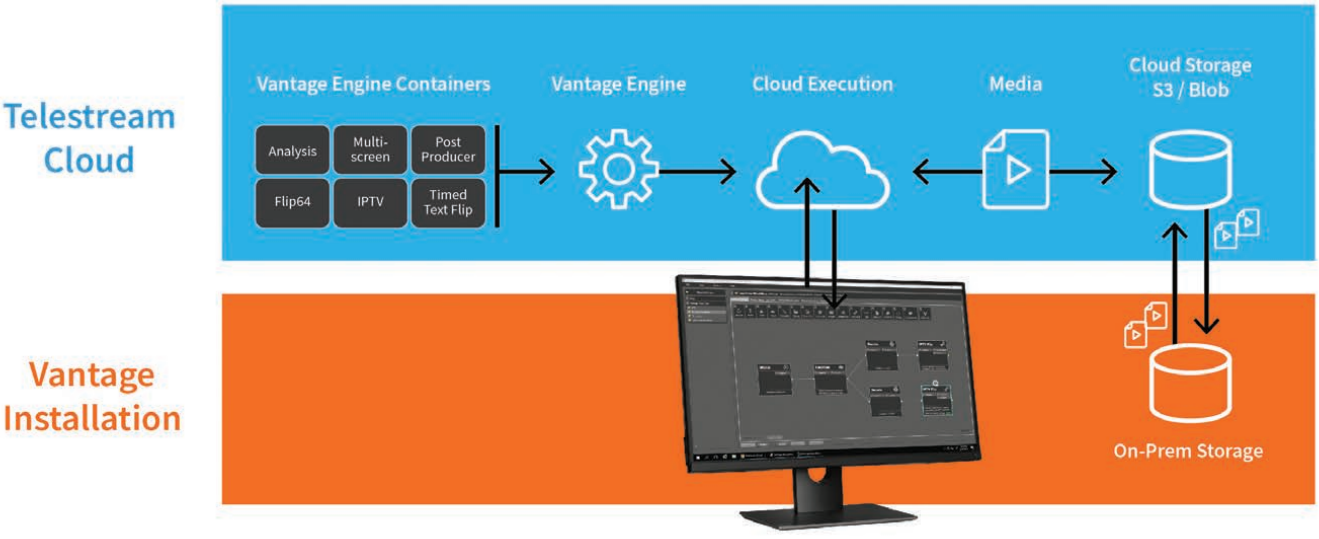

telestream

However, when considering moving myriad processes into the cloud, companies are quick to recognize the necessity of replicating in the cloud all of the tools that they rely on locally. This includes the workflow automation engine that ties all of these processes together into an efficient, cohesive whole. Transcoding alone simply does not suffice. For example, more and more outlets require IMF packages as the mechanism of delivery. These are not simply transcoded copies of the original master (which may have been made many years ago), but program segments that require significant processing in order to create the multiple components that make up an IMF deliverable. Depending on the number and variety of silos that have sprung up to manage and prep content, moving everything lock, stock, and barrel into the cloud could be a tall order.

USING THE CLOUD WITHOUT GETTING RAINED ON

When it comes to processing media for distribution, the first rule of the "Cloud Club" is not that nobody talks about it, because everybody is talking about it incessantly. Instead, the first rule is to "follow the media." If your media is in the cloud or going to be in the cloud, then by all means you should process it there. If it is "on the ground," then it's going to be very expensive to push it to the cloud for some processing, only to have to pay handsomely to bring it back down. One of my colleagues refers to this as the "Hotel California scenario," as businesses can check in their data for free but can't ever leave without paying extortionate egress charges. The lesson here is that there are many instances where on-premise workflows will still rule when it comes to efficiency, particularly when it comes to specialized, high-resolution, high-bandwidth content workflows.

Moving to the cloud doesn't have to be an all-or-nothing proposition, and it's important to be wary of any opinion to the contrary. There are a number of scenarios where cloud-based processing—and particularly processing hosted by a public cloud provider—is not preferable or even feasible. There are data ownership provisos in many source agreements that prevent material from being housed on a public cloud platform. There are also some scenarios in which an on-prem platform can actually be more cost-effective than a cloud approach—mainly those where the "run rate" business is well known, and there is less need for "bursts" of processing.

SECRETS TO SUCCESS

For many media organizations, the solution is to go for a hybrid approach that covers the run-rate business, or a significant portion of it, with on-prem processing but with the capability to process in the cloud where it makes the most sense (follow the media). Or whenever the company has a burst of work which cannot be fulfilled with on-prem processing within a given time frame. The secret here, though, is to ensure that both on-prem and cloud-based workflows can offer all of the same capabilities with no exceptions and hopefully with the same interface. Once again, it's so much more than transcoding. All of the processing steps and options need to be available in both scenarios for this approach to be successful.

Telestream released Vantage Cloud Port earlier this year to enable Vantage users* to select encoding actions within their workflows that would normally consume on-premises resources and to export those actions to the cloud. Vantage continues to control and orchestrate those actions, but it is now consuming public cloud resources via Telestream Cloud running in a public cloud platform of choice. Customers pay for these actions on a per-content-minute basis. It is a hybrid model that enables a customer-managed on-prem installation to be augmented from an infrastructure perspective and a licensing perspective by Telestream.

Vantage Cloud Port provides a very cost-effective entry point into Vantage, enabling enterprises that could not previously afford Vantage to make use of its rich media processing feature set. The system is differentiated from other systems in its access to all the workflow automation capabilities that are intrinsic to Vantage. These new cloud-enabled media processing actions enable Vantage to utilize the benefits of a public cloud and provides consumption-based pricing in a streamlined and straightforward fashion.

Vantage Cloud Port is a fully featured media processing platform with most of the characteristic Vantage transcoding functionality including Telestream's specialty formats and packaging.

Learn how to collapse distribution silos so you can scale, optimize, manage content via centralized resources, and perhaps make your job more productive. Listen to the on-demand webinar titled "Are Your Video Distribution Workflows Siloed?" at https://pages.telestream.net/dist_sm

**Telestream has thousands of Vantage users*

ABOUT TELESTREAM

Telestream provides world-class live and on-demand digital video tools, workflow solutions and quality monitoring capabilities that allow consumers and businesses to transform video on the desktop and across the enterprise. Many of the world's most demanding media and entertainment companies, and service providers, as well as a growing number of users in a broad range of business environments, rely on Telestream products to streamline operations, reach broader audiences, generate more revenue from and ensure the quality of their media. Telestream products span the entire digital media lifecycle, including video capture and ingest; live and on-demand encoding and transcoding; captioning; playback and inspection, delivery, and live streaming; automation and orchestration. With its iQ product line, Telestream enables the monitoring and management of quality service and experience over any network.

Telestream's corporate headquarters are located in Nevada City, California and Westwood, Massachusetts. The company is privately held. For company and product information, visit www.telestream.net.

MEDIAEXCEL

Multiscreen Encoding/Transcoding for Ad-Supported Services

How the encoding/transcoding platform impacts the performance and value of ad-supported services

Multiscreen video content and service monetization models revolve around two main pillars: subscriptions and advertisements. Subscription-based services may range from short-term/pay-per-view to longer-term monthly/annual commitment for the viewer. Ad-supported services may range from a spray-and-pray, mass targeting to a razor-thin, personalized approach. Those monetization models are widely used today across live/linear as well as VOD services and offerings.

Additionally, several operators (most prominently Hulu and CBS All Access in the U.S.) now employ a hybrid "subscription with limited adverts" monetization model in order to better position their offerings across the consumer segments and maximize their revenue potential. This also improves ad targeting opportunities, as the viewer profile is better known/tracked through the user's subscription account.

Before pay TV operators and content providers can harvest the full monetization potential of ad-supported services, they need to put in place a complex ecosystem, from acquisition to signaling to targeting to tracking. An essential component of this value chain is multiscreen encoding/transcoding, and that is the focus of this article.

ACQUISITION

In the live workflows, ad opportunities are generally ad-hoc (e.g. a timeout in a basketball game). The encoder platform is expected to process those opportunities in a frame-accurate manner through the following:

- in-band ad signaling, such as SCTE-35 or SCTE-104
- out-of-band API-based ad signaling, such as CableLabs ESAM
- out-of-band API-based timed metadata insertion/injection, such as ID3 tags (e.g. for sidebar, overlay based ads, or VPAID-based ad tracking)
- fingerprinting or AI-based automation for ad detection and replacement

In the VOD workflows, ad opportunities are generally pre-decided (e.g. pre and mid rolls). The encoder platform is expected to process those opportunities in a frame-accurate manner through the following:

- sidecar edit lists, such as CableLabs ESAM XML
- automation/rules (e.g. pre-roll, the first scene change after 15 mins of content)

Most high-end encoder platforms today support one or more of the above practices, and the main challenges are time-accuracy and interoperability. A best-effort approach in ad signaling, instead of a frame-accurate one, severely affects the viewer experience (e.g. a movie scene is cut mid-sentence for an advert).

CONTENT PREPARATION

After the ad opportunities are marked, the content itself needs to be prepared for seamless ad insertion. The encoder platform is expected to do the following:

- Condition the GOP structure of the encoded feed and insert IDR frames at the boundaries of the ad opportunities.
- Adjust the segment duration of adaptive HTTP outputs so that the boundaries of the ad opportunities coincide with the segment boundaries.
- Decorate the corresponding manifests to signal the ad opportunities for downstream DAI components.
- Ensure frame alignment across all adaptive variants and across primary and backup feeds.

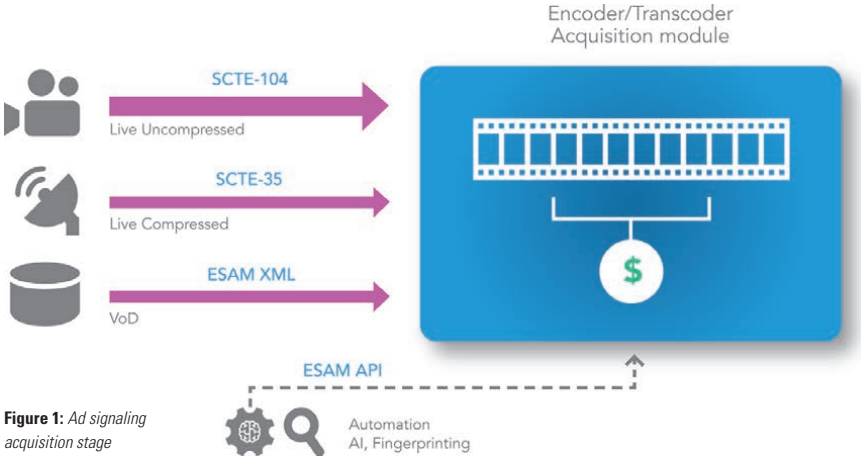

Figure 1: *Ad signaling acquisition stage*

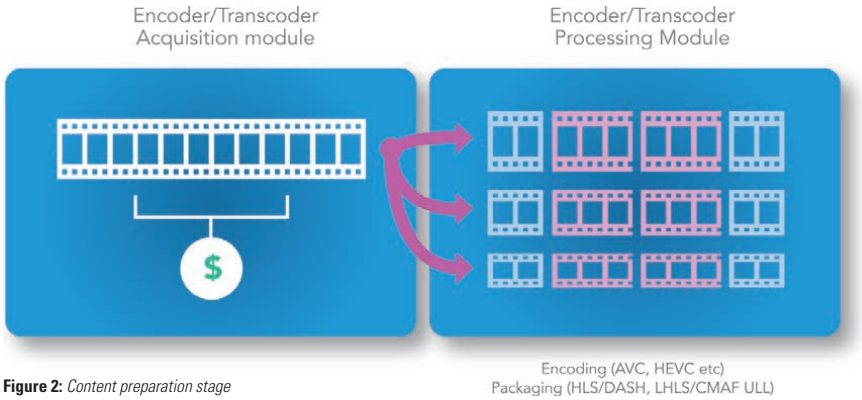

Figure 2: *Content preparation stage*

- Enable features such as overlays, scrolling text, etc., for sidebar-type ad insertion.
- Enable recording automation practices such as ad-free archiving or proof of ad delivery.

End-to-end interoperability is the key challenge in this stage, as the encoder needs to be standards-compliant across all adaptive delivery technologies and resilient to timestamp drifts and failover events. Further, for live workflows where low latency is a key competitive advantage, the encoder's challenge is to enable dynamic ad insertion (DAI) even in low latency mode (e.g. LHLS or CMAF ULL over HTTP/2).

AD ASSETS PREPARATION

Apart from the original content, the ad assets also need to be transcoded in order to ensure a smooth user experience. The transcoder platform therefore should achieve the following:

- Match the encoding properties of the original content (e.g. codec, resolution, bitrate, aspect ratio, etc.) across all adaptive variants. This needs to be dynamic and adjustable so that any changes in the content are also replicated for the ad assets by invalidating the corresponding caches.

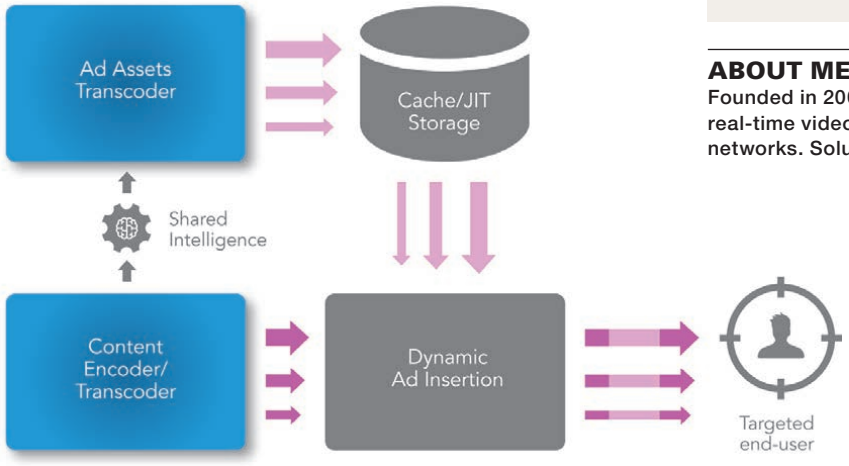

Figure 3: *Ad assets preparation stage*

- Deliver exceptional performance and scalability, in terms of turnaround time, in order to enable programmatic ad insertion, where ad decisions are made within milliseconds.
- Enable audio normalization and loudness management in order to enhance user experience without jeopardizing compliance.
- Allow for payload-level timed metadata insertion/injection (e.g. ID3 tags), to aid ad tracking regardless of the packaging/delivery protocol.

Most encoder platforms fail to deliver a unified solution across live and VOD workflows, thus the original content and ad assets are often produced by separate independent encoder platforms. This practice introduces artifacts (buffering, flickering, rescaling, audio volume fluctuation, etc.) and severely impacts the user experience, the operator's brand equity and consequently the ROI.

MEDIA EXCEL'S VALUE PROPOSITION

- ✔ Media Excel's HERO product portfolio covers all aspects of multiscreen encoding/transcoding. The HERO software can be deployed as a turn-key appliance, virtualized solution or cloud instance, enabling performance, scalability and orchestration across all form factors.
- ✔ Pristine quality of experience with adaptive encoding for H.264/AVC, H.265/HEVC up to 8K/HDR and ultra-low latency CMAF packaging.
- ✔ Unified and intrinsic support for live-to-live (streaming), file-to-live (playout), live-to-file (catchup, archive), file-to-file (VoD), enabling knowledge sharing across content and ad preparation workflows.
- ✔ Native in/out-of-band ad signaling, timed metadata injection, GOP/segment/manifest conditioning and ad insertion.
- ✔ Media Excel today supports Fortune 500 clients and global brands for their 24/7 and event-based mission critical deployments in broadcasting, cable, satellite, IPTV, pay TV, OTT/TVE, mobile, and government/DoD with five 9's reliability.

ABOUT MEDIA EXCEL

Founded in 2000, the company has been the industry leader of real-time video processing software to distribute video over IP networks. Solutions from Media Excel provide the reliability, scalability and performance required to deliver high-quality video via appliance and cloud deployment models. Powering more than 500 million multiscreen subscribers globally, and with the largest market share in multiscreen encoding, Media Excel assists pay tv operators, content providers, broadcasters, and telcos worldwide.

To learn more, please visit www.mediaexcel.com or contact us at info@mediaexcel.com

NewTek™

Swimming Upstream

Encoding and Streaming for Real-Time Production

Much of the energy and attention about encoding and streaming focuses on delivery and distribution—all the way to the viewer watching at home or somewhere on a mobile device.

All of this is possible through the use of IP as the transport mechanism. IP can now be considered the last media standard, because IP has the capability to grow and scale in bandwidth, complexity, and installation type.

In media production, transport revolves around SDI, a proven methodology going back in time over decades. Consequently, encoding and transcoding has revolved around hardware to translate SDI signals and the supported codecs such as MPEG-2 to IP-compliant packetized streams with codecs such as H.264 and HEVC.

For capture, acquisition, processing, and transport, the challenge has been in how to get enough performance out of mainstream IT practices and infrastructures using IP to make it practical and reliable enough to use for producing live digital media content. Production has desired the same benefits of downstream distribution. If media production could use IP, then the burden of confronting a new standard again goes away. This is not true of SDI or HDMI.

At the same time, powerful currents have dramatically re-sculpted the production landscape: Inexpensive cameras with incredible features are everywhere, producing an explosion of valuable content. Storage media for content acquisition is faster and more affordable than ever, and economical high-quality file formats abound. Non-linear editing is quicker, cheaper, and more fun than anyone could have imagined, and capabilities that formerly cost millions of dollars are now a few taps away on a mobile device.

NewTek believes in making media inherent to common IT environments. By developing Network Device Interface (NDI®), NewTek has changed the paradigm—giving virtually

everyone access to easy-to-use, high-quality video over IP. NDI is a software standard that *requires no special hardware of any kind*. NDI provides connectivity with existing software applications, computer platforms, and network infrastructure. The lightweight, low overhead implementation of NDI makes it as simple as downloading an NDI tool at no charge, and seeing live video moving across the network with Windows, Mac, and Linux software applications. NDI includes multiple source encoding of signals for viewing, hearing, mixing, recording, and editing, while retaining visual quality, frame accuracy, and source synchronization.

NDI is not just an SDI cable replacement. It enhances workflows with new capabilities such as resolution, aspect ratio, and frame-rate-independent encoding up to UHD formats. Floating-point audio processing starts at 16 channels and scales to virtually unlimited numbers. 16-bit video color computations ensure image precision and quality for HDR standards. Video, proxy, key, audio, control, tally, custom metadata, and precision timestamp data are all included in NDI streams. Production systems can easily produce program outputs and steaming outputs simultaneously.

Any NDI source can be recorded without recompression and no need to transcode. Media streams are directly converted into files, and support for "growing files" permits access to content even while media capture continues. Recording can be made from sources in any combination of video aspect ratio, frame rate, or resolution. The recordings include multi-channel audio, embedded alpha for keying, and video proxies, and are time-stamped during capture,

to make multi-camera synchronization simple. All of this is great for live production, but the multi-channel synchronization across inputs, systems, and locations also make post-production editing fundamentally easier as well. Recording functionality extends NDI beyond the demands of live production where it originated into affiliated media ecosystems such as archive, video on demand, and post-production. An NDI tool from NewTek will import NDI files directly into Adobe Creative Cloud applications, and NDI devices and connectivity reduce the cost of complexity of moving editing into cloud-based implementations.

NDI uses a unique codec with the benefit of enabling transmission of multiple video streams across existing infrastructures. NDI employs discrete cosine transform compression, which is the method of compression commonly used in encoding formats and mezzanine codecs around the media industry. NDI is the first-ever codec to provide multi-generational stability. Once a video signal is compressed, there is no further loss. Generation 2 and generation 1000 of a decode-to-encode sequence are identical. As NDI produces IP-friendly streams and files, transcoding into some other format for downstream use is low-overhead exercise for software or hardware systems.

NDI technology enables identification of and communication between IP-based production components connected to the same local area network and operates bi-directionally so every source is also a destination. Further, it facilitates interconnected production workflows by enabling encoding, transmission, and reception of multiple streams of high-quality, low-latency, frame-accurate video, audio, and data in real time by compatible networked components. Devices can include video mixers, graphics systems, capture cards, and other commonly used production equipment—all through software, and across existing network infrastructures.

Streamlined management of individual components means that any NDI-enabled system or device connected to the network is seen through IP. Extremely low latency permits sources to be synced and switched in real time for live productions with only a single field of latency, so no buffering is required. High-quality, resolution-independent video provides an image virtually as clean as if it were viewed directly in the camera. NDI implementations include video access rights, grouping, and IP commands.

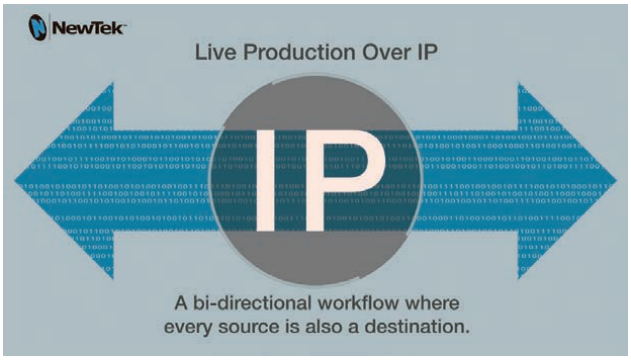

NDI uses minimal bandwidth even for multi-source productions. GigE LANs are usable for small workflows, and only 10Gb connectivity is needed for robust and higher resolution ones. Extremely efficient video packaging permits multiple streams to be stacked together even on a single Gigabit network.

NDI does not include an inherent clocking mechanism, but does include a 100ns UTC timestamp on every video frame and associated audio stream. In addition to using UTC, high-precision timing schemes such at NTP can also be utilized, synchronizing streams arriving from different locations on the network. NDI also supports integration with ASPEN, SMPTE 2022, SMPTE 2110, and other emerging IP standards.

As NDI is IP-based, it supports a variety of IT conventions to work across various network environments including mDNS (multicast Domain Name System) discovery server functionality, UDP multicast, forward error correction, and multi-TCP.

NDI has been designed to permit all type of suppliers to work together. NDI is easy for manufacturers of production components to develop compliant products. This has led to NDI achieving the widest support of any IP protocol for media by far. There are now thousands of companies with software and hardware products supporting NDI, and millions of users incorporating it into their environments.

NDI, NDI Tools, NDI Applications, and third-party NDI systems form a powerful combination to make creative imagination the only digital media production limitation. Reducing the costs and complexities of encoding, transcoding, and steaming media perfectly matches the need for collaboratively creating and delivering more content, more quickly, for more platforms, with less time and resources.

Find out more at **www.newtek.com/ndi**.

About NewTek

As the leader in IP video technology, NewTek is transforming the way people create network-style television content and share it with the world. From sporting events, Web-based talk shows, live entertainment, classrooms, and corporate communications, to virtually any place people want to capture and publish live video, we give our customers the power to grow their audiences, brands and businesses faster than ever before.

Clients include: New York Giants, NBA Development League, Fox News, BBC, NHL, Nickelodeon, CBS Radio, ESPN Radio, Fox Sports, MTV, TWiT.TV, USA TODAY, Department of Homeland Security (DHS), the National Aeronautics and Space Administration (NASA), and more than 80% of the U.S. Fortune 100.

NewTek is based in San Antonio, Texas. For more information on NewTek please visit: **www.newtek.com**, Twitter, YouTube, Flickr or connect with us on Facebook.

verizon media

Encoding for High Quality Live Sports Streaming

BY NATHAN BURR, SENIOR ENGINEERING MANAGER

Many things must be considered to ensure success when streaming a live sporting event, including encoding costs, venue bandwidth, and bitrates required to optimize your workflow. And that's just for starters. Let's take a deeper look at factors that impact meeting your quality and budget requirements.

SIGNAL ACQUISITION

Venue bandwidth can be a weak link in your workflow and restrict your ability to deliver great viewing experiences reliably. You need to preserve as much of the original feed as possible, with an eye toward even greater adoption of higher-quality formats and 4K. However, you also need to optimize the stream so it can be delivered efficiently without getting bogged down by additional overhead such as personalized advertising. Finding the right balance during this step of the video workflow is critical.

Backhaul costs, which can add up quickly, must be factored into your budget. A satellite uplink rents for around $2,000 per day, and bandwidth costs about $400 an hour. Given the inconsistent and bandwidth-constrained conditions at some venues, it's always a good idea to reduce upload bandwidth requirements as much as possible, which can help you deliver a broadcast-like experience to your viewers.

ENCODING AND BITRATES

Once the live video feed leaves the venue, the next step in the workflow is encoding. Here, a video encoder creates multiple versions or variants of the audio and video at different bitrates, resolutions, and quality levels. It then segments the versions into a series of small files or media segments. There are several additional steps that must be performed as well, such as creating a media playlist for each version containing a list of URLs pointing to the version's media segments. The resulting master playlist is used by the player to decide the most appropriate version for the device and the currently measured or available bandwidth.

Live feeds need to be encoded into any number of adaptive bitrate formats and protocols, such as Apple HLS and MPEG-DASH, while minimizing latency to cover a myriad of playback devices. You may also need to add support for Microsoft Smooth Streaming for gaming consoles.

DRM can also complicate encoding by requiring its own set of multiple formats to support large audience requirements. For example,

verizon media

older players that don't support DRM need HLS and AES-128. Older iOS devices require HLS and FairPlay. Newer iOS devices support HLS, FairPlay, and CMAF CBC. Older versions of Windows and Android only support CMAF CTR. Newer releases of Android, Windows, and iOS should support all CMAF formats. This means your content must be packaged in multiple different formats to allow playback on all of those devices.

As resolutions increase and codecs become more complex, it becomes harder to encode a complete ABR encoding ladder on a single machine, whether that machine is located in the cloud or on-prem. If your encoding hardware can't keep up with the live feed, you may need to look at reducing the number of rungs on the encoding ladder, a move that ultimately could impact your audience's experience.

CONCLUSION

Configuring your live stream encoding workflow often means striking a compromise that balances desired quality with cost and system constraints. A cloud-based managed service can simplify the workflow and reduce bandwidth requirements on the front end. It can also dramatically reduce upfront and ongoing expenditures while delivering the high-quality low-latency streams your viewers expect.

To learn more about streaming live sporting events at scale, please see our series on Medium.com.

PTZ OPTICS

Encode and Stream Right From Your Camera

Did you know that the second-generation PTZOptics cameras include a built-in video encoder? The SDI, USB, and ZCam models can encode and stream multiple high-definition video streams simultaneously while still outputting video via HDMI, 3G SDI, NDI, and/or USB 3.0. The combination of the camera's advanced encoding and robotic control capabilities have a lot to do with why broadcast professionals are falling in love with these cameras. And don't just take our word for it: The second-gen PTZOptics 20X-SDI camera was voted "Best PTZ or Network-Controlled Camera" in the 2018 Streaming Media Magazine Readers' Choice Awards.

By building IP video encoding capabilities into each PTZ camera model, PTZOptics reduces the cost of video production for many workflows in an industry that is quickly adopting video over IP. Because PTZOptics cameras have supported IP video connectivity from the release of the very first models in 2015, most users do not need to purchase an additional video encoder to stream content as they would with a standard camcorder. IP video production is obviously important to engineers at PTZOptics, and they have worked with NewTek to make sure every PTZOptics camera model has an option to support NDI. NDI is a new low-latency and high-quality video production protocol now supported by the world's largest video production companies.

PTZOptics' built-in camera encoders are useful for streaming media to platforms like Facebook, but also for streaming audio and video directly into video production applications such as vMix, Wirecast, Livestream Studio, or Open Broadcaster Software (OBS). PTZOptics' built-in video encoders are commonly used with simple video applications like VLC media player and OBS. But the cameras are also deployed to operate with advanced video management servers like Wowza and GoEasyLive.

The PTZOptics built-in camera encoders can encode video and audio into popular streaming formats such as H.264, H.265, MJPEG, and NDI HX (High Efficiency). Using the built-in encoder inside your PTZOptics camera is a handy way to deliver video over a local or a wide area network. Using a single ethernet cable you can power your PTZOptics camera, deliver video, gain robotic PTZ positioning controls, and configure your encoding protocols with precision.

Many users like to control PTZOptics cameras using an IP joystick, which is also available with built-in PoE (Power over Ethernet), but others like to use their iOS or Android smartphones over Wi-Fi. Perhaps the most popular way to control PTZOptics cameras is already built-in to the video production system you are already using. PTZOptics now supports direct PTZ camera integrations with vMix, Wirecast, OBS, MimoLive, TriCaster, Livestream Studio, Zoom Video Conferencing, and more.

In esports, the popular Twitch live streaming platform features an extension allowing broadcasters the ability to monetize camera control. The PTZOptics Twitch extension allows broadcasters to give their audience control of a camera inside a location in their broadcast. Viewers can pay bits for access to the pan, tilt, zoom camera controls kept secure by an intermediary server. With two easy operational modes (Que and Code) broadcasters can choose to protect camera controls with a passcode for a trusted friend for remote operation, or choose how much time each viewer can have operating the PTZ camera for a fee.

Video production professionals are always finding new ways to use the encoders and IP connectivity. Whether the cameras are being used for previewing video over a local area network, sending video over a wide area network, or integrating into an advanced IP video NDI IP video production environment, you can configure the encoders by navigating to the camera's IP address with any web browser and adjusting the advanced menu options. Each camera can support two simultaneous video streams that can be used in any resolution configuration up to 1080p at 60 frames per second. PTZOptics now offers PTZ cameras with 12X, 20X, and 30X optical zoom options. A second-generation PTZOptics 30X-SDI camera costs just $1,799, and it can zoom into a head and shoulders view of a person on stage from up to 75 feet away.

You can learn more about PTZOptics live streaming cameras at **www.ptzoptics.com**. You can also join the PTZOptics User Group to see how customers around the world are deploying these cameras at **facebook.com/groups/ptzopticspals**.

About PTZOptics

PTZOptics is a US-Based professional live streaming and broadcast camera manufacturer. Winner of the 2018 *Streaming Media* magazine's Readers' Choice Awards, PTZOptics makes some of the world's best PTZ cameras with NDI, HD-SDI, HDMI and IP Control.

PTZOptics is known for its weekly live broadcasts on Fridays at 11a.m. PST / 2p.m. EST. Show hosts, Paul Richards and Tess Protesto, interact with audience members from businesses, churches, schools and healthcare through an active online live streaming webinar available on Facebook, YouTube and Twitch. Learn more about ptzoptics at **https://ptzoptics.com**.

DVEO®
Extraordinary Products & Support

Economics and Evolution: The Two Drivers of Cloudification and Survival as a Niche Player

BY LASZLO "LES" ZOLTAN

Both economic forces and evolution are very powerful agents of change in human behavior and technology. Economic forces are like gravitational forces, quietly pulling smaller objects towards them, even though it is not evident why objects are moving in a particular direction in the cosmos. In the same way evolutionary forces are at play when multiple technologies meet in the market place and only a few ultimately prevail. According to evolution, the fittest technologies that win in the long run. Naturally, some "boutique" providers also survive to serve niche markets that the large providers ignore. In the evolutionary model the fittest technology must prevail in the long run

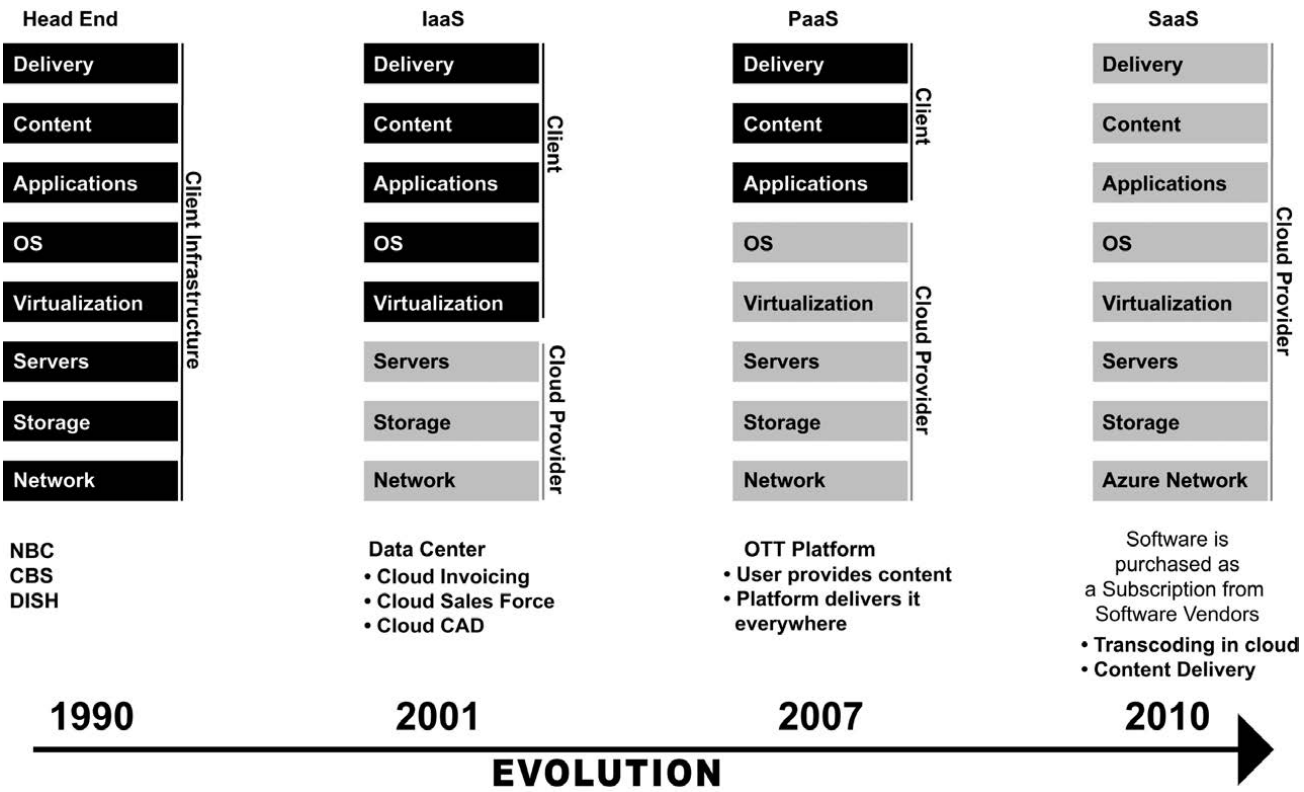

since it is better suited to the new environment. Those in the market place ultimately must choose the best fit or else they too will be displaced by more cost effective and efficient providers.

This is exactly what is happening in video consumption, and video delivery. The content delivery environment changed with the advent of the internet. At the beginning of the internet revolution, the gravitational pull exerted by slower internet speeds was rather weak. In 2019 internet speeds have improved vastly. Just as importantly, the cost of internet bandwidth has also plummeted over the years. The process curve is not linear, but rather an accelerating one. This means that the cost of delivering information as 1-'s & 0-'s is now significantly cheaper than it was 5 or 10 years ago. This is evolution at play, as "fitter" internet technology keeps displacing less-fit technologies like DSL, radiowaves, microwaves, and satellite. Companies in the infrastructure market are also in an evolutionary battle to determine who will dominate that market for the next few years. Cisco, the long standing dominant player, gets leapfrogged by younger challengers, and they must steel up to do battle against these upstarts to regain supremacy. This is the evolutionary process at play, and that is driving down the cost of connections via IP. It is no different than the process at play to select which animal, dominates the herd, but dominance is temporary. There are new agile young upstarts in technology and new challengers looking to topple the leader of the packs.

Simultaneous with the evolution of internet speed we are seeing evolution of better and faster codecs—in the space of 15 years we have moved from MPEG-2, to H.264, and now to H.265. The MPEG-2 to H. 264 evolution was peaceful, but the adoption of H.265 may not be so smooth. New, more robust, challengers are waiting in the wings to challenge the heir apparent. Codecs like VC1 and VP-9 are biding their time to challenge H.265.

So thanks to evolution, we have cheaper, more reliable internet and we have more efficient codecs. We now have a content hungry 1 billion viewers eager for content. These are the unserved masses in India and Africa who don't often have much money but are waiting their turn to get video content at a price they can afford. Economic theory dictates that this very large unmet need can actually be satisfied by massive scaling up of data centers. This is exactly what is happening. Very large data centers are being planned and built in India and Africa. This upscaling, of course, drives prices down even outside of India and Africa.

These economic forces are creating the accelerating movement towards cloudification. Cloudification is the "end stage" or "factory-building stage," massively displacing in-house and outsourced economic units that provide bandwidth and delivery. Now these super-fit giants with large financial backing and the ability to build massive factories, are displacing the smaller players. Just like consolidation in other industries, in due time a few large content delivery "factories" will dominate. Of course, smaller niche players will survive and do well. They must, however, differentiate themselves with unique offerings. Interestingly, we see that the growth of the pioneers of generic video platforms have slowed significantly. We won't name names, but several video platform pioneers are in financial trouble. Cloudification is the natural next step to the process started by home internet delivery deployment. From in house to platform to cloudification, the process is unstoppable.

DVEO saw this reality early and has chosen to be a smaller, more agile, and more specialized engine provider. We distinguish ourselves by offering non-generic encode, transcode, decode, and insertion engines. We also offer a cloud based error correction engine. We have chosen not to compete with the big guys and risk death. We chose to stay as a vital healthy niche player; coexisting with the mega sized suppliers with niche codecs like J2K, with niche capture capabilities, and niche cloud ad insertion engines that offer the most features.

We have chosen to serve the customers who need unusual features. We chose to be long-term survivors.

ABOUT DVEO

DVEO® is a strong and well established supplier to leading Telco, TV/OTT, and WISP operators around the world. We provide telcoquality encoders, decoders, transcoders, ad insertion solutions, live media servers, and packet recovery technology... all with support for HLS/fMP4 plus optional H.265/HEVC.

aws

Revolutionizing High Resolution Streaming At Scale

DAN GERMAIN, SENIOR PRODUCT MANAGER COMPRESSION & VOD, AMAZON WEB SERVICES

There is no telling where VOD customers will view content; it could be anywhere from a home theatre in a dedicated projection room, an LED panel in the living room, a laptop in the office, or more likely these days on a tablet or phone during a morning commute. One thing is certain though, as the resolution and dynamic range of these displays continues to increase across all of these devices, more efficient encoding methods are being employed to try and keep the transcoding requirements for all these devices under control.

Dan Germain, Senior Product Manager Compression & VOD, Amazon Web Services

Cable networks, CDNs, and streaming content providers seem to face insurmountable challenges when optimizing their existing VOD asset libraries and trying to turn around new content to meet their delivery service level agreements (SLAs). A full-length 4K HDR feature could take traditional workflow infrastructures over a day to render out all of the adaptive bitrate (ABR) stacks required for the myriad of consumer players available today. However, with the ubiquity of cloud processing and the advent of new technologies such as AWS Elemental MediaConvert's Accelerated Transcoding, these compute-heavy workflows can be processed with drastically reduced turnaround times, allowing customers to meet studio's SLAs for transcode and QA with time to spare.

At its core, accelerated transcoding builds on parallel processing concepts like split and stitch, where media is broken down into segments for processing. However, this can be tricky business, and if it isn't done properly you can get spikes in data rates that can be problematic with different players.

Accelerated transcoding is based on incredibly powerful engineering that takes into account data rates, analyzes types of GOP structures, player compatibility, and can output a full gamut of codecs, not just for segmented outputs like HLS and Dash but also mezzanine file types like ProRes and XDCam. New engineering developments in transcoding combined with the power and price performance flexibility thanks to millions of computing cores in the cloud is truly changing the game for VOD workflows.

In practice whenever a content distributor buys rights from a content owner they want to turn that content around and make it available to their customers as soon as possible. If it takes the content distributor a day to perform the transcoding, and then even more time for quality control (QC), they may lose a substantial percentage of VOD customers to their competition. Ideally once a network secures a new piece of media from a content owner, they want the file transcoded, QC'd and available for customers on demand as soon as possible, within hours not weeks. Accelerated transcoding is a very real solution which solves these complex compute heavy workflows. Not only is the transcode sped up by an order of magnitude, but thanks to a combination of new encoding technologies like quality-defined variable bitrate (QVBR), one transcode job can produce the whole ABR stack of 60-70 different deliverables with up to a 50% reduction of file size. The CDN cost savings alone are enough to make VOD catalogue owners evaluate the benefits of re-encoding their content libraries to take advantage of these new processes—especially now, when quality, price and speed are no longer obstacles.

ABOUT AWS

For over 12 years, Amazon Web Services has been the world's most comprehensive and broadly adopted cloud platform. AWS offers over 125 fully featured services for compute, storage, databases, networking, analytics, machine learning and artificial intelligence (AI), Internet of Things (IoT), mobile, security, hybrid, virtual and augmented reality (VR and AR), media, and application development, deployment, and management from 57 Availability Zones (AZs) within 19 geographic regions around the world, spanning the US, Australia, Brazil, Canada, China, France, Germany, India, Ireland, Japan, Korea, Singapore, and the UK. AWS services are trusted by millions of active customers around the world—including the fastest-growing startups, largest enterprises, and leading government agencies—to power their infrastructure, make them more agile, and lower costs.

Recommended Reading and Viewing

solution series

For more on **advanced encoding and transcoding**, check out these articles and videos from **StreamingMedia.com**.

FEATURED ARTICLES

NAB '19 ENCODING AND QOE HIGHLIGHTS: HERE'S WHAT MATTERED

See you later, per-title encoding. There's a new system in town and it's called context-aware encoding. Learn which companies unveiled CAE solutions at NAB, as well as other streaming breakthroughs.
go2sm.com/july19a

BUYERS' GUIDE TO CONTEXT-AWARE ENCODING 2019

2018 was the year that context-aware encoding (CAE) went mainstream, but 2019 will be the year that context-aware encoding reaches critical mass.
go2sm.com/july19c

BUYERS' GUIDE TO ENCODING APPLIANCES 2019

Software encoding, especially in the cloud, is all the rage these days. So why stick with on-premises hardware? There are good reasons for keeping operations in-house.
go2sm.com/july19d

BUYERS' GUIDE TO PER-TITLE ENCODING

Per-title encoding is now a required feature on most encoding platforms. Here's an overview of the features to look for, and how to get the most out of them
go2sm.com/july19e

BUYERS' GUIDE TO ON-PREM ENCODING 2019

Despite all the hype around the cloud, plenty of use cases still call for on-prem video encoding. Here's what to look for when choosing a solution.
go2sm.com/july19f

BUYERS' GUIDE TO LIVE TRANSCODING 2019

Whether it's enterprise, ecommerce, news, or gaming, it feels like everyone is rushing to go live. Here's a guide to services that will get any company streaming now.
go2sm.com/july19g

AV1: A FIRST LOOK

FFmpeg 4.0 gives many video engineers their first chance to test the new AV1 codec against H.264, HEVC, and VP9. The results? In our tests, quality was impressive, but glacially slow encoding times make AV1 a non-starter for most publishers until hardware acceleration becomes available.
go2sm.com/july19l

FEATURED VIDEOS

STREAMING MEDIA EAST 2019: TWITCH'S TAREK AMARA TALKS MULTI-CODEC STREAMING

Streaming Media's Tim Siglin and Twitch's Tarek Amara discuss Twitch's developing the emerging codec landscape and Twitch's multi-codec content delivery strategy on the show floor at Streaming Media East 2019.
go2sm.com/july19m

NAB 2019: TELESTREAM TALKS CLOUD TRANSCODING AND HYBRID WORKFLOWS

When Telestream Vantage customers need a large video library encoded in a hurry, they can now turn to a hybrid cloud option. Streaming Media's Jan Ozer and Telestream's Ken Haren discuss Vantage Cloud Port on the show floor at NAB 2019.
go2sm.com/july19n

NAB 2019: TWITCH TALKS VP9, AV1 AND ITS FIVE-YEAR ENCODING ROADMAP

Twitch uses a head-and-tail encoding strategy where popular content (the head) is encoded one way and less popular content (the tail) another. In this interview, a Twitch engineer explains what's on the video gaming powerhouse's five-year roadmap.
go2sm.com/july19o

NAB 2019: NETINT TALKS HIGH-DENSITY H.265 ENCODING

Streaming Media's Jan Ozer and NETINT's Ray Adensamer discuss NETINT's Codensity T400, which is aimed at companies that need to do large live video encoding jobs at scale.
go2sm.com/july19p

SUPERMARKET

NEW GEAR • USED GEAR • SERVICE • JOB • STARTUPS • MORE

OUR SPONSORS

NETINT Technologies

Suite 306–3500 Gilmore Way
Burnaby, BC V5G 0B8
Canada
+1 (604) 628-8177
www.netint.ca

For more information and articles about NETINT, see go2sm.com/netint1.

Telestream

848 Gold Flat Road
Nevada City, CA 95959
USA
(530) 470-1300
www.telestream.net

For more information and articles about Telestream, see go2sm.com/telestream.

Media Excel

8834 N. Capital of TX Hwy, Ste. 230
Austin, TX 78759
USA
(512) 502-0034
www.mediaexcel.com

For more information and articles about Media Excel, see go2sm.com/mediaexcel.

NewTek

5131 Beckwith Blvd.
San Antonio, TX 78249
USA
(210) 370-8000
www.newtek.com

For more articles about NewTek, see go2sm.com/newtek.

Verizon Media

13031 W. Jefferson Blvd., Suite 900
Los Angeles, CA 90094
USA
(213) 200-3988
https://vd.ms/50SM2019

For more information and articles about Verizon Media, see go2sm.com/vdms.

PTZOptics

152 Robbins Rd.
Downingtown, PA 19335
USA
(610) 518-2200
https://ptzoptics.com

For more information and articles about PTZOptics, see go2sm.com/ptzoptics.

DVEO

11409 West Bernardo Court
San Diego, CA 92127
USA
(858) 613-1818
dveo.com

For more information and articles about DVEO, see go2sm.com/dveo.

AWS

1320 SW Broadway, Suite 400
Portland, OR 97201
USA
1 (888) 900-8104
aws.amazon.com/media

For more articles about AWS, see go2sm.com/aws.