

SEPTEMBER 2021
VOL. 18, NO. 6



streaming media

THE BUSINESS &
TECHNOLOGY OF
ONLINE VIDEO

*ENHANCED
SPORTS AND
ESPORTS
EXPERIENCES*

BEYOND SOFTWARE

*WELCOME
TO THE
AD TECH
JUNGLE*



SM

streaming media **west**

the business and technology of online video

Registration is Open!

NOVEMBER 2-3, 2021

Preconference Workshops | November 1

**HYATT REGENCY
HUNTINGTON BEACH**
Huntington Beach, California

- 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.

SAVE \$100 WITH CODE SMMAG21 [STREAMINGMEDIA.COM/WEST](https://streamingmedia.com/west)

Featuring These Dedicated Tracks & Special Events:



DIAMOND SPONSOR



PLATINUM SPONSORS



GOLD SPONSORS



BRONZE SPONSOR



READERS' CHOICE AWARDS
SILVER SPONSOR



MEDIA SPONSORS



ORGANIZED AND PRODUCED BY



PUBLISHERS OF

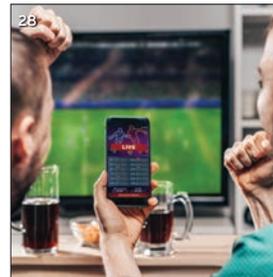


Table of contents

September 2021

FEATURES

Welcome to the Ad Tech Jungle NADINE KREFETZ	16
Beyond Software TIM SIGLIN	22
Enhanced Sports and Esports Experiences STEVE NATHANS-KELLY	28



COLUMNS

The Streaming Gold Rush of 2021 EDITOR'S NOTE ERIC SCHUMACHER-RASMUSSEN	2
Pandemic Mergers and Acquisitions CLASS ACT LIAM MORAN	4
Collecting 'Junk' for Better Streaming STREAMS OF THOUGHT TIM SIGLIN	6
The Future Is Real Time, and It Starts Now THE VIDEO DOCTOR ROBERT REINHARDT	8
Apple's Advanced Video Quality Tool THE PRODUCER'S VIEW JAN OZER	10

DEPARTMENT

Vistex, Agora, and Transmit STREAMING TOOLBOX NADINE KREFETZ	33
--	----

REVIEW

The Roland V-160HD Hybrid Event Video Switcher REVIEW ROBERT REINHARDT	37
---	----

SPONSORED CONTENT

Innovation Series: Content Delivery & Edge Compute TELESTREAM	11
--	----



The Streaming Gold Rush of 2021

It's a cliché to say that the last year has been unprecedented, but for those of us in the online video industry, it's been especially so. And not just in terms of the sheer number of streaming service providers and platforms entering the market or the spectacular growth in viewing numbers, but also in the amount of investment and consolidation among technology providers. I could list them all here, but I'd need a full page!

Perhaps it's not entirely unprecedented, but you'd have to go back 15 years or so to see this much activity in the market. I asked Brett Sappington, VP of consumer insights firm Interpret (interpret.la), for his thoughts on the great streaming gold rush of 2021.

"Companies are realizing that many of the changes brought about by the pandemic are likely to be long term, and they are adding pieces that better position them for the future," Sappington says. Multi-party video conferencing is here to stay for both business and personal use, he adds, and live streaming, remote production, efficient online workflows, and analytics are key technologies for now and the future.

As I've spoken with vendors in the space, one consistent theme has emerged. Over and over, I hear them say that the streaming video industry has crammed 5 years' worth of growth and technology advances into 18 months. To keep up with demand, streaming tech providers have needed to grow quickly. "In some cases, companies are acquiring proven partners to enhance their current offerings," Sappington says. "In others, they want to bring key expertise in-house as they consider future product design."

So the big question is whether or not the activity will slow down. Sappington doesn't think so.

"I think the activity will continue to be brisk in the short term," he says. "Overall M&A activity is high due to a number of factors, including available capital and a desire by companies to be more competitive coming out of the pandemic. In the video tech space, as companies make these acquisitions, rivals want to be sure to keep pace."

Eric Schumacher-Rasmussen (erics@streamingmedia.com) is VP and editor of Streaming Media and chairs Streaming Media conferences. Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

Vice President & Editor-in-Chief Eric Schumacher-Rasmussen
Vice President & Publisher Joel Unickow
Streaming Media Producer Editor Steve Nathans-Kelly
Senior Graphic Designer Lisa Boccaduttre
Editorial Services Manager Lauree Padgett
Copy Editor Terri Koenig
Contributing Editors Nadine Krefetz, Liam Moran, Jan Ozer, Robert Reinhardt, Dom Robinson, Tim Sighin, Jason Thibeault

CORPORATE

Chairman Roger Bilboul
President & CEO Thomas H. Hogan
Production Manager Tiffany Chamenko
Group Publisher Bob Fernekees
CFO Mike Flaherty
Vice President of Marketing Tom Hogan Jr.
Video Publishing Director Steve Nathans-Kelly
Ad Trafficking Coordinator Jacqueline Walter Crawford
Circulation Manager Janeen Welsh

List Services

Joel Unickow, VP/Publisher
joel@streamingmedia.com
 Direct: (250) 933-1111; Fax: (844) 267-1471

Direct Advertising Inquiries to:

Joel Unickow, VP/Publisher
joel@streamingmedia.com
 Direct: (250) 933-1111; Fax: (844) 267-1471

Streaming Media (ISSN: 1559-8039; USPS: 25466) is published monthly except combined issues in January/February, April/May, July/August, November/December by Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055 USA; Phone: (609) 654-6266; Fax: (609) 654-4309; Internet: infotoday.com. Registered in U.S. Patent & Trademark Office.

Periodicals postage paid at Vincentown, NJ and additional mailing offices.

SUBSCRIPTION INFORMATION

Subscriptions are available free to qualified recipients in the U.S. only. Nonqualified subscribers in the U.S. may purchase a subscription for \$79.95 per year. Nonqualified subscribers in Canada and Mexico may purchase a subscription for \$89.95 per year. Delivery outside North America is \$146.00 via surface mail per year. All rates to be prepaid in U.S. funds.

Subscribe online at <https://infotodaytoday.omega.com/stm/login.do> or write Information Today, Inc., 143 Old Marlton Pike, Medford, NJ 08055-8755.

Back issues: \$10.50 per copy, U.S.; \$15.00 per copy, Canada and Mexico; \$22.00 per copy outside North America; prepaid only. Missed issues within the U.S. must be claimed within 45 days of publication date.

Change of Address/Cancellations: Mail requests, including a copy of the current address label from a recent issue and indicating the new address, to P.O. Box 3006, Northbrook, IL 60065-3006; call (800) 248-0588; or email us at stm@omega.com.

Reprints: For quality reprints of 500 copies or more, contact Tiffany Chamenko, *Streaming Media* Reprints, 143 Old Marlton Pike, Medford, NJ 08055-8750; (800) 300-9868 x 122; email: tchamenko@infotoday.com.

© Copyright 2021, Information Today, Inc. All rights reserved.

No part of this publication may be reproduced in whole or in part in any medium without the express permission of the publisher.

PRINTED IN USA

POSTMASTER: Send address changes to *Streaming Media* Magazine, P.O. Box 3006, Northbrook, IL 60065-3006.

Rights and Permissions

Permission to photocopy items is granted by Information Today, Inc. provided that a base fee of \$3.50 plus \$0.50 per page is paid directly to Copyright Clearance Center (CCC), or provided that your organization maintains an appropriate license with CCC.

Visit copyright.com to obtain permission to use these materials in academic coursepacks or for library reserves, interlibrary loans, document delivery services, or as classroom handouts; for permission to send copies via email or post copies on a corporate intranet or extranet; or for permission to republish materials in books, textbooks, and newsletters.

Contact CCC at 222 Rosewood Drive, Danvers, MA 01923; (978) 750-8400; Fax: (978) 646-8600; copyright.com. If you live outside the USA, request permission from your local Reproduction Rights Organization. (For a list of international agencies, consult ifro.org.)

For all other requests, including making copies for use as commercial reprints or for other sales, marketing, promotional, and publicity uses, contact the publisher in advance of using the material. For a copy of our Rights and Permissions Request form, contact Lauree Padgett, lpadgett@infotoday.com.

Online Access

Visit our website at streamingmedia.com.

Contents also available online under direct licensing arrangements with EBSCO, NewsBank, ProQuest, Gale, and H.W. Wilson and through redistribution arrangements with information service providers including Dow Jones Factiva, LexisNexis, OCLC, STN International, and Westlaw.

Disclaimers

Acceptance of an advertisement does not imply an endorsement by the publisher.

While best efforts to ensure editorial accuracy of the content are exercised, publisher assumes no liability for any information contained in this publication.

The publisher can accept no responsibility for the return of unsolicited manuscripts or the loss of photos.

The views in this publication are those of the authors and do not necessarily reflect the views of Information Today, Inc. (ITI) or the editors.

Writing for Streaming Media magazine

Contact Eric Schumacher-Rasmussen, N6865 Hillview Dr., Beaver Dam, WI 53916; phone: (920) 342-6263; email: erics@streamingmedia.com

Streaming Media magazine is printed and bound at St. Croix Press, New Richmond, WI.

Streaming Media magazine is a trademark of Information Today, Inc.



ALL VOTERS
ENTERED TO WIN!



VOTE

YOUR INDUSTRY LEADERS TODAY!

VOTING CLOSSES
September 20



bit.ly/SMRCA21

#StreamingRCA

SILVER SPONSOR



PRIZES INCLUDE



MANTIS HD VIDEO
ENCODER
(VALUE IS \$995)

WINNERS ANNOUNCED AT STREAMING MEDIA WEST IN HUNTINGTON BEACH, CALIF.,
AND IN THE NOVEMBER/DECEMBER ISSUE OF *STREAMING MEDIA*.

Pandemic Mergers and Acquisitions

We've entered a new era in which well-capitalized companies gobble up less-successful competitors that have something to offer. The music has stopped, and as schools return to mostly normal operations, there will be fewer chairs available for educational technology providers. Companies that failed to gain sustainable market share during the past year of unprecedented mandatory school closures face uncertainty in a future in which remote instruction decreases to somewhere closer to pre-pandemic levels.

Each of these ill-fated companies was started with a dream and driven along by smart people putting their best efforts into realizing those dreams. Many of them were created with bright ideas that became useful features, and those companies provide value in an acquisition through intellectual property (IP). Others may have offered run-of-the-mill products that were marketed well but not quite well enough; They provide a buyer with additional customers. All offer talented personnel.

Three transactions impacting instructional media delivery have already been initiated in the past two quarters. These transactions serve as exemplars for capitalization methods that are well-suited for different phases of a company's growth and as datapoints for trends in streaming media in education.

The first is the partnership between K1 Investment Management, LLC and Panopto to accelerate Panopto's growth in a conventional (for tech, at least) private equity deal. This capitalization immediately resulted in the acquisition of Ensemble Video, the company that provided the Ensemble video management system (VMS) in 2005. To my knowledge, this was the first VMS designed for the education sector just before the explosion between 2006 and 2007 that saw iTunesU, Kaltura, and Panopto hit the market for managing the delivery of video in schools. This acquisition seems to be a good fit from an IP/personnel standpoint, since both platforms are built on Windows server technologies.

The second major transaction was Kaltura's initial public offering (IPO), which raised \$150 million for the company and its venture capital partners. The original IPO plan was set for earlier in the year, but an unexpectedly soft investor market delayed the share sale. The prospectus summary from Kaltura's IP filing with the SEC ([go2sm.com/kalturasec](https://www.sec.gov/edgar/disclosure/primary/efiling/14/000119312514000001/kaltura-20140923.htm)) is required reading for anyone following this industry, as it is a strikingly clear-eyed take on the state of the educational video business as we exit the pandemic.

The final exemplar transaction is 2U's purchase of the edX platform for \$800 million, which was announced at the end of June and set off a flurry of controversy and summertime hot takes. A pair of well-informed and complementary analyses are Dhawal Shah's, from The Report by Class Central ([go2sm.com/2uedx1](https://www.classcentral.com/report/2u-acquires-edx/)) and Michael Feldstein's, from eLiterate ([go2sm.com/2uedx2](https://www.eliterate.com/2020/07/20/2u-acquires-edx/)). 2U is a company that primarily serves as an online program market (OPM), recruiting potential students to online post-baccalaureate degree programs and boot camps that are run mostly by universities in the U.S., U.K., and Australia.

The acquisition of edX is about building out a customer base and an international brand: Aside from branding elements, student data, and existing websites, no edX IP changed hands. To finance this purchase, 2U paid almost half of the value with cash on hand and the rest with a very favorable Term Loan B. (The Term Loan B is the institutional analogue of a balloon mortgage, in which small payments are made over a period of a few years with the full balance plus interest due at the end of that term. The major difference is that a single-family property doesn't typically generate \$80 million of annual revenue.) The \$800 million sale price is being paid into the as yet unnamed new nonprofit corporation being created by edX's founding institutions, Harvard and MIT.

That nonprofit's mission will be to support, among other things, the development of next-



generation learning experience platforms, including the Open edX open source learning management system that edX is built on and that also powers HarvardX; a new Open edX instance at MIT called MITx; and manufacturingworkforce.org, a partnership between MIT and the U.S. Department of Defense that was announced on the same day as the 2U acquisition. Its funding is included in a \$3.2 million grant.

With video proving itself among the most critical and impactful educational technol-

ogies over the past year, it stands to reason that these next-generation learning experience platforms supported by MIT and Harvard's nonprofit will show a rejuvenated focus on video technology.

Liam Moran (liammoran@gmail.com) produces curricular digital media and manages the servers to store, process, and deliver it at the Center for Innovation in Teaching and Learning at the University of Illinois–Urbana-Champaign. Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

A Perfect Pair for SRT Streaming

Ultra Encode
Universal Live Media Encoder

Pro Convert
Multi-Format IP Decoders

**SRT
READY**

Internet
WAN

Live Streaming • Remote Production • Multi-Site Distribution

Magewell's feature-rich **Ultra Encode** live encoders and **Pro Convert** decoders provide reliable, low-latency, end-to-end H.264 or H.265 media transport using the SRT protocol. The free **Magewell Cloud** software adds multi-device management, SRT relay, and protocol conversion. Visit our website to learn more.

MAGEWELL[®]
www.magewell.com

Distributed by Mobile Video Devices, Inc.
www.mobilevideodevices.com
sales@mobilevideodevices.com
610.921.5720

MVD

© 2021 Nanjing Magewell Electronics Co. Ltd. All rights reserved. All trademarks are property of their respective holders.

Collecting 'Junk' for Better Streaming

A bit over a year ago, Dom Robinson and I wrote an article about the need to consider the impact that streaming at a global scale has on power consumption. The article, called “The Greening of Streaming,” argued that there are three P’s that should be focused on: adding power (consumption) to the traditional price and performance considerations most streaming companies use to sell or buy products and services.

As you can read in Dom’s follow-on article, “The Greening of Streaming, One Year On” (go2sm.com/greening2), we’ve not just pitched the idea to the industry; instead, we’ve actively worked to raise awareness of the importance of greening by launching a formal initiative of the same name (greeningofstreaming.org). By the time you read this, we’ll have had our first Greening of Streaming virtual roundtable and will be moving on toward several member-only events in the last quarter of 2021.

To say that I advocate for anyone involved in the business of streaming—including cameras, content acquisition, encoders, content delivery networks, and end-user playback peripherals—to lend their voice to the Greening of Streaming initiative would be an understatement. Do it, and do it now, for both environmental and social reasons.

There’s a second initiative that I’d also like the industry to get behind: collecting “junk” as a way to research better streaming approaches.

Almost 3 years before I got involved with the “greening” discussion, a group of industry veterans got together to discuss ideas on how to close up the digital divide. Our initial approach, and the driving force behind starting up a not-for-profit charitable organization—now a 501(c)3 called Help Me Stream Research Foundation—was to focus on the digital divide in emerging economies.

To do that, we needed to gather electronic junk that’s sitting in drawers and closets across the U.S. (and Canada and Europe, for that matter) and repurpose all of it into efficient and very inexpensive streaming solutions that could run on intermittent power.

Beyond emerging economies and the non-governmental organizations (NGOs) that serve them, we’ve also been challenged to provide solutions for disaster areas and—with the advent of COVID-related lockdowns last year, which laid bare the rural-urban digital divide in the U.S.—to consider solutions for rural communities.



Help Me Stream Research Foundation opened a test lab at East Tennessee State University’s Innovation Laboratory to evaluate new gear and turn old gear into usable tech in emerging markets.

I’m happy to report that we’ve caught enough attention—both from key influencers and forward-looking companies in the streaming space—that Help Me Stream is opening its first test lab this month, where we’ll not only validate claims made about newer gear in the industry, but also work to turn that donated junk into streaming systems to distribute to NGOs in key markets across the globe.

We’re doing so in partnership with a nationally known public health program at East Tennessee State University. This will give us the opportunity to turn that junk into robust streaming solutions and also to work with medical students—many of whom may work in areas of the U.S. or the world where resources are limited to what one has on hand. The goal will be to help them identify components from a variety of abandoned computers to build a working device that could power not only an electronic medical records solution, but also a basic, on-demand playback system to offer education to those who may not otherwise have access to it.

In turn, we’ll gain from the experience this public health program already has in training



its students to use bicycle parts to create small power generators that are currently used to run small water pumps for providing better hygiene and crop irrigation.

The common thread between the Greening of Streaming and Help Me Stream Research Foundation initiatives is to simultaneously grow the streaming industry's innovations for at-scale delivery, while also reducing electronic landfill through the intentional repurposing of older desktops, laptops, servers, smartphones, and tablets.

How can you help? After considering a membership in Greening of Streaming, spend a few

minutes digging in your closets, sock drawer, or garage to see if there's any junk that you'd like to donate. As a 501(c)3, we're able to give donation letters—either for this junk or any monetary donation to the mission—and we'd love to have your insights into projects that would benefit from the junk you've found. Visit hmsrf.org/donate or email donations@hmsrf.org for more details.

Tim Siglin is a streaming industry veteran and longtime contributing editor to *Streaming Media* magazine. Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.



TRANSCODE YOUR STREAMS ANYWHERE



Hotels



Cruise Lines



BRUTUS GPU:TELCO

Visit us at NAB Booth:C1942!

- Adapts MPEG-2 streams to H.264 for Hotel, Cruise Line, University, and Resort feeds
- Grooming for OTT and Mobile
- 1 RU with 300 watts for excellent performance
- IPTV Grooming for differentiated services at different rates
- Down converter ability to convert all HD Services to SD

sales@dveo.com | +1 858 613-1818 | www.dveo.com

The Future Is Real Time, and It Starts Now

We're on the precipice (aren't we always, in this streaming world?) of advanced real-time collaboration and the melding of traditional video workflows with new, real-time tools. Last night, I was in a hotel room in Seattle using my phone's hotspot over LTE to tech-direct a live-streaming event for one of my clients. All of the tech hurdles were anticipated, and contingencies were planned. I wasn't going to do a traditional video switch on my laptop or use dedicated video switcher hardware in my studio office. We set up an Amazon Web Services (AWS) instance running Windows 11 and vMix, controlling the instance via Microsoft Remote Desktop. Thankfully, I didn't need more than a couple of megabits up and down, as the AWS instance was doing all the heavy lifting, and the 2-hour event was a success.

So where and how does "real time" fit into your workflows in this "COVID and beyond" era? Similar to global wars accelerating advances in medicine such as plastic surgery, COVID has pushed all of us working in streaming media to rethink, innovate, optimize, and rebuild many components of our day-to-day responsibilities. And the big players are recognizing it too. Adobe just announced a \$1.275 billion acquisition of [Frame.io](#), a video editing collaboration service. Last year, Verizon acquired BlueJeans Network, a real-time conferencing platform, for \$400 million. This summer, Zoom added NDI support to its desktop conferencing apps, allowing remote contributor feeds into any NDI-enabled video switcher to be more easily shared. Real-streaming CDNs such as Millicast, nanocosmos, Frozen Mountain, and Wowza Streaming Cloud are growing at a faster pace than ever and are constantly adding new features to address the needs of their customers. The list goes on and on. It's hard for me to keep track of how many people have approached me during the last 18 months for advice on building out new real-time collaboration tools or enhancing existing systems with the injection of real-time data.

As you and your company's services evolve, keep in mind that you don't need to focus on a quick sellout of your company as so many

have attempted to do since the inception of the "dot-com" business world. There are real dollars to be made without raising millions and pouring nearly all of it into R&D. If you're reading this column, chances are, you have existing clients and services that could benefit from some aspect of real-time collaboration or data services. While I've already written other columns that discuss project planning and "build-or-buy" approaches to tech development, here's my short list of important factors to consider for real-time development:

- 1) Open source isn't "free."** While there are many open source codebases to enable real-time applications, such as FFmpeg and Janus WebRTC Server, you'll need to have the developer chops to make those tools work with your own services. You won't likely launch a new service from scratch within 30 days using open source if your dev teams aren't familiar with the trials and tribulations of WebRTC, video and audio codec transcoding, and so on.
- 2) Utilize existing real-time services from other vendors.** Most real-time streaming CDNs have intentionally built their platforms to address a wide range of business requirements for their customers. While some price tags for monthly enterprise services may seem steep to small or startup businesses, there's a lot to be said for everything you get with those monthly subscriptions: software development kits, application programming interface documentation, support services, and tech stack management.
- 3) The cloud is diverse and not a singular entity.** Admittedly, the vast majority of the cloud development my team and I perform is on AWS. But there are many companies that may offer cheaper hosted services and servers, such as Google Compute Engine, Microsoft Azure, DigitalOcean, Liquid Web, Vultr, and more. Be sure to explore



the best “bang for the buck” to stretch your development dollars.

4) Develop a long-term road map. I can't emphasize this last point enough. Too many of my single stakeholder clients feel they need to own the IP of everything they want from Day 1, and I firmly believe that this is largely impractical and nearly impossible for most small businesses and startups to achieve. Get your proof of concept (PoC) for new services up and running as the first phase in the road map, utilizing whatever resources are readily available, and have as many “knowns” (vs. “unknowns”) as possible. After you've got the PoC done, head into Phase 2 and define your MVP (minimally viable product) to get to market as soon as possible, again utilizing and piggybacking

other vendors' services. Then, if necessary, build a plan to take over and assume the responsibilities that other services were providing in Phase 2.

I hope you're as excited as I am about the near- and long-term potential of our real-time streaming media tech. Be sure to check out upcoming sessions at the next Streaming Media West conference (streamingmedia.com/west), in person or remotely, this November to learn more about new tools and services that can help your business thrive now and in the future.

Robert Reinhardt (robert@videorx.com) is the founder of VideoRx and is internationally regarded as an expert on multimedia application development and online video, particularly in HTML5, iOS, Flash, AVC/H.264, and HEVC/H.265. Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.



STREAMING FORUM CONNECT

October 5, 2021 | www.streaming-forum.com

FREE webinars delivering expert-led panels and commentary to help you resolve streaming challenges, gain strategic insights, and make great technology decisions

<p>Low Latency, Ultra-Low Latency, & You</p> <p>SPONSORED BY:</p>  <p>Nearly every live stream can benefit from lower latency. But how low can you go? What other strategic and tech decisions do you need to make to get there? Our expert panel discusses, live.</p>	<p>How to Maximise Efficiencies & Minimize Costs in Your OTT Workflow</p> <p>SPONSORED BY:</p>  <p>Your OTT workflows are in place, now it's time to focus on making them more productive, more efficient, and more cost-effective from glass to glass. This panel identifies the improvements you'll need to make.</p>	<p>Live Streaming in 2021 & Beyond</p> <p>SPONSORED BY:</p>  <p>Audiences, performers, and teams are back in venues and stadiums. Live-streaming producers debate efficiencies in remote production workflows without sacrificing quality, both now and ahead.</p>
--	---	--

Join us on Tuesday October 5 from 13.00 BST / 14.00 CET / 08.00 EST

www.streaming-forum.com

Brought to you by 



Apple's Advanced Video Quality Tool

Apple launched the Advanced Video Quality Tool (AVQT; go2sm.com/avqt) at its Worldwide Developers' Conference 2021. Here's an overview of what I learned about it.

In a post (go2sm.com/slcavqt), I point out that AVQT is both a quality metric, like VMAF and SSIMPLUS, and a tool to produce that score, like the Moscow State University Video Quality Measurement Tool (VQMT) or SSIMPLUS VOD Monitor. As a tool, AVQT is free, Mac-only, and command-line-driven. To produce the AVQT score, it compares the encoded file to the source and reports the score in either JSON or CSV format. You can't visualize the frames or score within the tool.

The tool has a number of features not typically seen in first-gen quality tools. For example, it can compare low-res videos to the source without pre-scaling. It can measure quality at different display resolutions, so you can measure the quality of a 640x360 video played in a 640x360 window, as well as fullscreen. You can also choose the viewing distance and the pooling method.

AVQT supports Dolby Vision 5 if the source and reference videos are in that format, and it is extremely fast. On an 8-core M1-based Mac mini, processing a 2-minute file took 15 seconds. On an 8-core Intel Xeon CPU E3-1505M Windows-based computer, computing VMAF on the same file took 8:40, about 35 times longer.

Strangely, however, the tool computes the score on a segment-by-segment basis, not on the file as a whole. You can set the segment size up to 60 seconds and get a single score, but if you're measuring files longer than 60 seconds, you'll get back a segment-by-segment score that you'll have to convert to a total score if needed.

I asked Apple about this and got this response: "We agree there is value in reporting an overall score for the whole video. However, this is quite challenging as it requires subjective data on long duration videos to design and evaluate an aggregation model. The aggregation model would need to mathematically model several memory related aspects in the human visual system such as first and last impressions, sudden quality drops and the length of low-quality periods."

This "perfect is the enemy of the good" response ignores the fact that every other metric reports a single score and that most of the items referenced have more impact in a short file than a long file. Five bad seconds in a 10-second file could destroy my overall opinion; 5 bad seconds in a 90-minute file is irrelevant.

In another post (go2sm.com/slcavqt2), I tested the metric and discovered that AVQT will fail if the encoded durations or frame rates vary from the source by even an irrelevant difference. In one test, in which the encoded file was 10.0767 seconds long and the source was 10.07, AVQT failed, even though both files had 300 frames. In contrast, FFmpeg, VQMT, and the SSIMPLUS VOD Monitor had no problems.

You can force AVQT to ignore the difference and produce a score, but you run the risk that the score is invalid, with no way to determine if there's a misalignment, as you can with VQMT or the SSIMPLUS VOD Monitor.

Many metric applications involve short, pristine YUV files encoded by precise reference encoders, so neither issue would be a limitation. If you work with longer, real-world videos encoded via commercial tools, AVQT might be a problem.

A final post (go2sm.com/slcavqt3) compared the AVQT metric to VMAF and SSIMPLUS, with some subjective evaluations from Subjectify.us. Here, although AVQT showed some bright spots, I didn't test enough data to conclude anything other than that AVQT, VMAF, and SSIMPLUS gauge quality differently, so you can't use VQMT as a faster VMAF or a cheaper SSIMPLUS.

As a tool, assuming the duration and sync issues aren't showstoppers, AVQT's speed and JSON output make it ideal for production. For experimentation-oriented practices like mine, the inability to visualize frames is a significant limitation. As a metric, AVQT showed some bright spots, but it's hard to see it bumping VMAF or SSIMPLUS from real-world workflows without a lot more verification.

Jan Ozer is a contributing editor to *Streaming Media* magazine and runs the Streaming Learning Center, where he produces courses like [Streaming Media 101: Technical Onboarding for Streaming Professionals](#).



INNOVATION
series

content delivery & edge compute

achieve the perfect balance of
reach, reliability, and immediacy

telestream



VOLUME 1, NO. 5 of a series of guides
to streaming products and services



Worried About the Last Mile? Don't Forget the First!

BY KEN HAREN, DIRECTOR OF PRODUCT MANAGEMENT, TELESTREAM

There is an ongoing revolution in media centered on how dynamic video routes can be leveraged over unmanaged internet and public cloud infrastructure. The timeline for growth in IP video delivery has only accelerated as COVID has mandated vastly reduced on-prem staff and production elements. From 5G and mobile edge compute (MEC) platforms, contribution and remote production (REMI) workflows, business-to-business (B2B) and affiliate distribution to direct-to-consumer (D2C) applications, these enhanced IP delivery flows have evolved from science projects to essential production and distribution services. And while these emerging platforms are enabling unprecedented flexibility and new viewer experiences, they may be doing so at the cost of visibility.

As workflows increasingly leverage dynamically allocated video delivery routes, often over unmanaged internet and public cloud provider infrastructure, a monitoring challenge has emerged. With these new models for supporting contribution, production, and distribution workflows, designing where monitoring is deployed and how operators are alerted to issues has become a much more complex issue.

The ability to rapidly build dynamic, low-latency, secure, and resilient video delivery routes from venue to production facilities or between distribution points has emerged as a leading technology innovation over the last couple of years. Platforms like AWS MediaConnect, Zixi ZEN Master, and Haivision Hub have simplified the orchestration of dynamic video routing and policy enforcement while integrating the video transport directly into production environments. Now it is easier than ever to directly view the status of video routes. What is less clear is what is being carried within those video routes and how impacts to one part of a distribution may be impacting other parts.

VIDEO IMPAIRMENTS/QUALITY ANALYSIS

We're all familiar with the phrase "garbage in/garbage out." Programmatically building video distribution routes only works when we're confident that clean, high-quality signals are being driven to production and distribution environments. Automating the detection and, importantly, responding to the presence of impairments within the content is important.



Silent audio, frozen video, slate, etc., are all commonly reported as conditions that are difficult to identify. The problem is compounded by remote production and D2C offerings that may include multiple camera angles, multilingual commentary, and a dynamic content schedule.

There may also be different points along the transport where the content is changed. For example, a D2C application will typically include a mezzanine contribution feed that must be transcoded to an ABR ladder and then packaged for OTT streaming delivery. Monitoring the live contribution feed and the derivative live streaming assets is essential to validating that a healthy signal is present and propagating. We recently saw a problem where a cloud-based transcoder was receiving a 4K mezzanine feed from the contribution encoder. Due to a last mile issue in the contribution environment and a need for ultra-low latency delivery, there were intermittent dropped/lost packets in the mezzanine distribution. While these were periodically causing a transport alarm on the delivery protocol, what was not immediately apparent was that the infrequent but persistent lost packets were causing significant artifacts being introduced in derivative streaming variants—an issue which was not immediately apparent to streaming ops teams but was being alarmed in the monitoring environment and was causing an impact in audience engagement.



Operations teams are looking for dynamic monitoring solutions that can be spun up on demand, pointed at media streams at various points in the distribution, and assess the audio/visual perceptual quality as the content load changes during the broadcast schedule. Settings that may be ideal for certain kinds of content may be inadequate for others.

ADVERTISING & POLICY ENFORCEMENT

Another area where monitoring has become essential is the validation that ads being produced and delivered are compliant with the rights and monetization policies mandated by the business. Creating high-quality content is expensive, and ensuring license rights are enforced and that monetization strategy is supported is foundational. However, it is not always a simple task to continuously monitor dynamically constructed delivery routes and audience-facing streams to validate the presence or absence of appropriate signaling and safeguards.

While the contributing team may be confident that their streams are decorated with SCTE-35 markers at appropriate program boundary points, these may not be propagating through to the audience facing streams, resulting in missing ad avails and viewers receiving ad blanking materials instead of personalized ads.

If a feed is not properly conditioned for ad insertion, or when the ad creatives themselves are of poor quality, this can lead to audience engagement issues. While last mile “client analytics” tools can provide a snapshot of audience engagement, it can be difficult to correlate this with root cause. This is another area where merging data produced in both the last-mile client analytics and first-mile monitoring KPIs becomes essential. Client analytics can only tell you about the streams that the audience is able to interact with. An example that highlights the need for both last- and first-mile monitoring involved audiences being served dynamic ads that were forcing their players to opt into a reduced resolution variant when the program content resumed from ad break. Because the caching network was not optimized to support the highest resolution variants, viewer sessions would report a healthy (if diminished) playback session. Because the sessions were all healthy, no last-mile alarms were typically being seen, but the first-mile synthetic clients were able to identify continuous performance issues related to the higher



quality streams. Resolving the cache configurations not only resulted in improved playback performance, but also generated a noticeable improvement in audience engagement, especially after ad breaks.

CONCLUSIONS

There is an ongoing revolution in media centered on how dynamic video routes can be leveraged over unmanaged internet and public cloud infrastructure. This has created an opportunity for flexible live video production and distribution models. Dynamic, flexible, high-performance video routing requires monitoring services that can be spun up and positioned when and where needed to ensure that what’s being carried within those routes is what’s intended. Insights from a dynamically provisioned end-to-end monitoring solution are essential to ensuring a quality distribution.

For more on comprehensive monitoring from contribution to the last mile, see <https://www.telestream.net/telestream-cloud/stream-monitor.htm>

ABOUT TELESTREAM

For over 20 years, Telestream® has been at the forefront of innovation in the digital video industry. The company develops products for media processing and workflow orchestration; live capture, streaming, production, video quality assurance, archive and content management; and video and audio test solutions that make it possible to reliably get video content to any audience regardless of how it is created, distributed or viewed. Telestream solutions are available on premises or in the cloud as well as in hybrid combinations. Telestream is privately held with corporate headquarters located in Nevada City, California and Westwood, Massachusetts.

For company and product information, visit www.telestream.net.

SPONSORED BY



TELESTREAM

848 Gold Flat Road
Nevada City, CA 95959
USA
(877) 681-2088

www.telestream.net

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

DIGITAL DOWNLOAD

DOWNLOAD THE  **streaming media**

2021 SOURCEBOOK

The World's Most Comprehensive
Video Resource for 18 years.



[STREAMINGMEDIA.COM/SOURCEBOOK](https://streamingmedia.com/sourcebook)



Brought to you by

edgecast



WELCOME TO THE AD TECH JUNGLE

Advertising technology is still scary, but here are some tools to help you make it out alive.

BY NADINE KREFETZ

The advertising technology workflow contains many moving parts, and ad tech solutions providers are constantly updating their offerings. The sector has lots of room for innovation, so I thought I'd check in with some of the leading vendors in the market to see what customer pain points they're addressing.

More Fill(ing)

First up is ad fill. Typically, anywhere from 20% to 30% of ad placement spots go unfilled, because content owners and service providers can't sell them, according to various industry observers. Wurl has introduced Ad Pool, a service to increase fill rates by monetizing unfilled inventory. "We saw an opportunity to bring in extra demand," says Yuval Fisher, Wurl's SVP of technology. "What's interesting is it's a free service, so we don't take any cut of the CPM at all. Now, the fine print is [that] not everybody does as well. Some companies pay us money, but to other companies, we actually pay them more than we take in."

Wurl brings in demand and passes along the opportunity to bid on unused spots. "We're not in a position where we would compete with the traditional DSPs [demand-side platforms] or SSPs [supply-side platforms], but in an imagined generic workflow, we would be positioned as an SSP because we've aggregated a bunch of supply," says Fisher. Wurl handles programmatic demand and direct deals. It also enriches this supply with contextual metadata (which could include episode names, topics, genres, etc.).

"One of our chief concerns is not to compete with demand sold by the publishers. When a publisher uses Ad Pool, they set a floor price they're selling their inventory for," says Fisher. "Wurl ensures [that] ad buyers can't try and circumvent publisher's offerings and go around them and buy more cheaply from us. We ensure that we don't compete with publishers."

Publishers participating in Ad Pool can put in specific exclusions. "Sometimes they don't want their branded inventory available," says Fisher. The result is that buyers can purchase

for a specific set of channels, but they don't necessarily know what channel. "Suddenly, they [have] a route for monetizing this unsold inventory, but nobody knows that it's their inventory."

This all happens within a few seconds in each individual ad break. "We have about a 30% uplift in impressions," says Fisher. So that's one way to not leave money on the table.

Welcome Multicasting

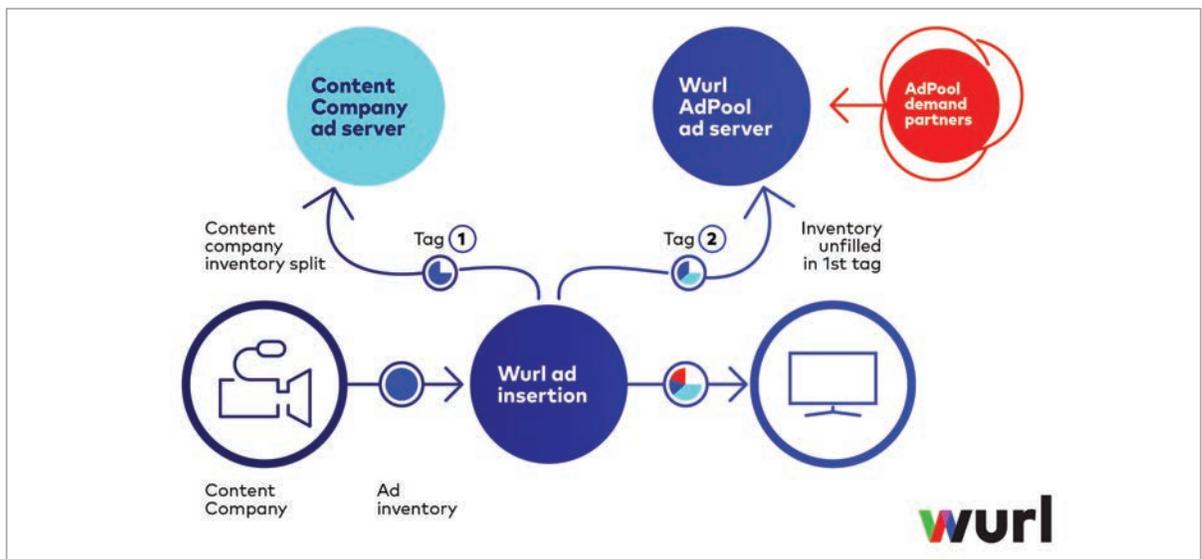
"Something that's very specific to our solution is the distribution through multicast," says Damien Sterkers, video solutions marketing director at Broadpeak. "We're doing multicast ABR [adaptive bitrate] for normal programs, which means that instead of distributing content on a one-to-one kind of transaction, ...

[we] push the same thing: one single copy to everyone."

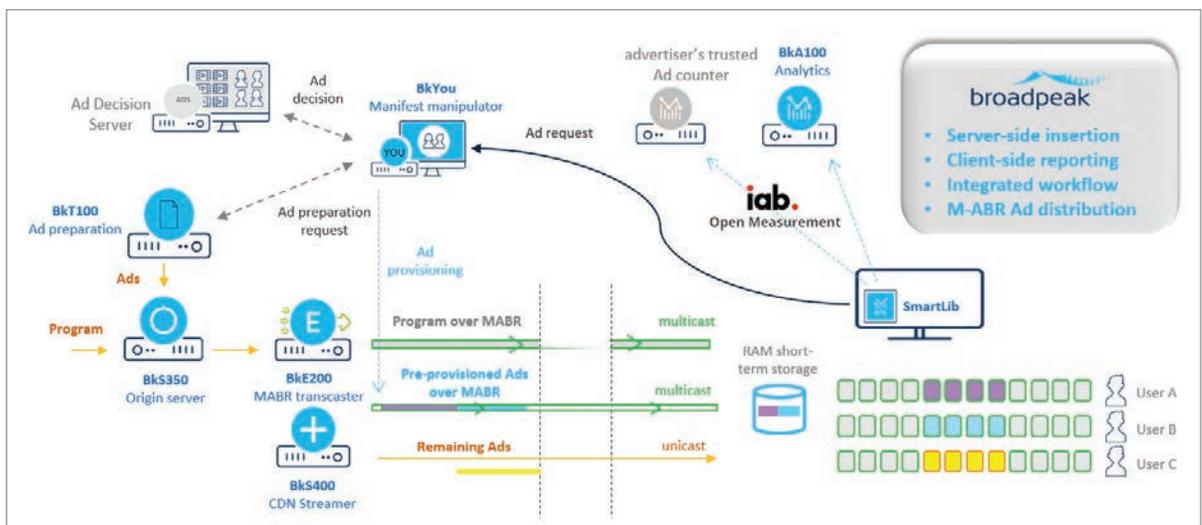
The result is unlimited capacity. "This is a copy that you can prioritize, and it is very minimal in terms of capacity on the network," says Sterkers. Broadpeak can use multicast for not only the program but for some of the advertising as well, and the ad decision could be from programmatic or direct sale.

"Everybody's watching the same thing at the same time, then comes an ad break, and, providing that the ad system is providing targeted ads, it means that not everyone will have the same content at the same time. Here, we've had to do some adaptation. The manifest manipulator knows which ads will be requested by most people," says Sterkers. "You

Wurl's Ad Pool service brings in demand from partners and presents the opportunity for publishers to bid on unused ad spots, increasing fill rates.



Broadpeak's multicast ad insertion is designed to make the ad insertion process more efficient.



can inform this multicast ABR server to say, 'These ads I know will be used a lot of times, you need to provision them on the multicast as well.' This way, they can be distributed through the multicast to the end users, and they can benefit from ... scalability and quality for these ads."

The principle of targeted advertising is that there is a theoretically infinite number of ads, but, actually, the number of possible ads that you will have during this time is limited, and some ads are more important than the others, according to Sterkers. "So the game here is to provision the most important ads so that they can be distributed in multicast."

Some viewers will have ads pre-provisioned and distributed through multicast. Other viewers get a different ad load that could not be pre-provisioned, so their broadcast goes unicast at this point. "The main thing is that you've secured that when you get to the ad break; you don't have everyone going to unicast requesting their own private copy. Most of them will get something from the multicast, and just a limited amount of people will have to do a private connection to get their ad," says Sterkers.

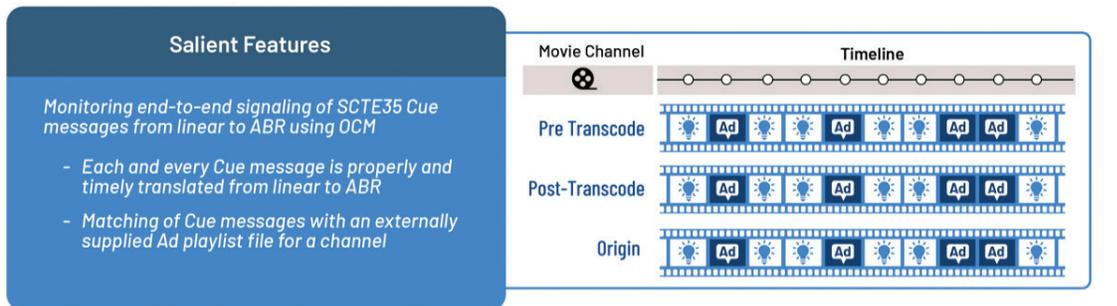
How do you get multicast? Broadpeak gets customers who've started with a manifest manipulation from Google. However, Google assumes you're getting the ad from its servers.

Two issues arise: Ads can be far away, and Google charges for each time the ad is pulled, according to Sterkers. Caching the commercials closer to the edge and provisioning for ABR multicast solves those problems. "We have an added value, because Broadpeak is providing the distribution system," says Sterkers.

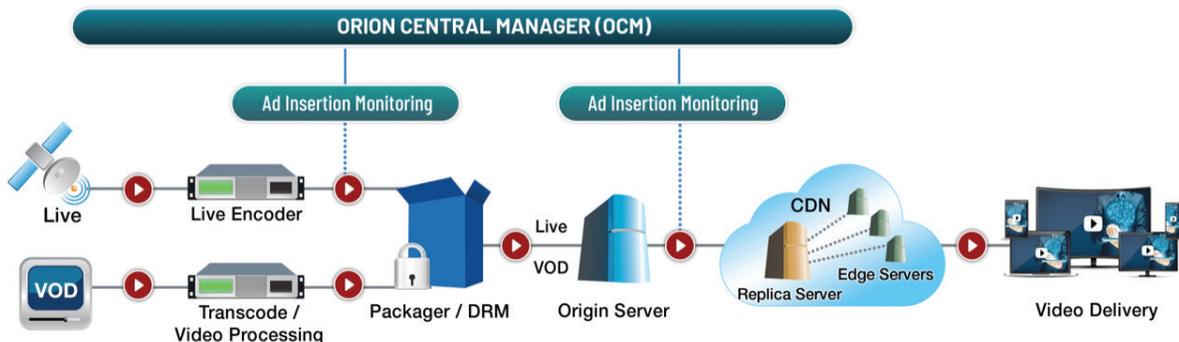
Trust and Verify

"Ad insertion monitoring is very important, and I would say in the last two years, there is an increased interest because for some of our customers, the primary monetization is from ads," says Anupama Anantharaman, VP of product management at Interra Systems. "In the ABR world, ad insertion monitoring is particularly complex."

Interra Systems' Orion now features detailed ad insertion monitoring reports, which include video thumbnails for visual proof that the ad was in fact inserted. The software is installed on a server, and there may be multiple probes running. For example, one probe may be running at ingest and another at a CDN location. Orion monitors the incoming streams in real time and can perform manifest validation for different profiles. "If there are any problems, then you can locate where the problem is. You can troubleshoot with all the built-in tools that we have in the monitor," says Anantharaman.



Interra Systems' Orion monitors ad insertion on both live and video-on-demand content, giving brands confidence that their ads are being delivered.



“We can compare both at pre-insertion and post-insertion to make sure that the ad in fact went into the same spot where it was supposed to go. We also have information about resolution of the ad and bitrate, and we can also perform [quality control] on that ad to make sure that audio and video quality meets customer requirements,” says Anantharaman. Orion also looks at how many ad slots were available and how many went unfilled.

In an OTT scenario in which content is being consumed in many countries, probes can be deployed in the cloud nearest to the customer. “Dynamic ad insertion enables streaming services to reach a much wider audience across so many different geographies, but that also means many different viewing platforms,” says Anantharaman. “It definitely adds to the complexity.”

Context and AI

To get even more complicated, maybe we should invent some new ad formats? That’s exactly what CatapultX is doing with its On-Stream platform, which inserts display-like ads into contextually relevant moments in video that’s adjacent to the main video. It can identify activities, logos, and objects, as well as significant moments that brands may want to be associated with.

“Our AI identifies different brands and logos and finds one we’re programmatically connected to, and then delivers [it] into the stream at that moment,” says Zack Rosenberg, CEO of CatapultX. One example might be that a soccer goal being scored could trigger an ad that’s selling the player’s jersey. This works for live or on-demand content. It can also be used in

ad-supported video on demand and more subtly in subscription video on demand. The premise is that this type of ad gives a better experience than showing a traditional 30-second ad in a subscription service in which viewers don’t traditionally see ads, according to Rosenberg.

“84% of people will leave as soon as pre-roll starts, 65% will hit skip the moment they can, and 25% of people will block ads,” says Rosenberg. “We invented a completely new set of inventory. We provide publishers with effectively three options: They buy it [based] on time, on percentage, or high-impact moments. If we know the video is three minutes long and they only want 10% of the video monetized, we have 30 seconds.

“We’re connected to DSPs. Everything can be bought programmatically. So while they’re not kind of your ... standard formats, they use existing ads,” says Rosenberg. “We took the approach to partner with video players so that we get access to any imaginable content. We’re compatible with everything from JW Player and Brightcove [to] AirTV and Dailymotion.”

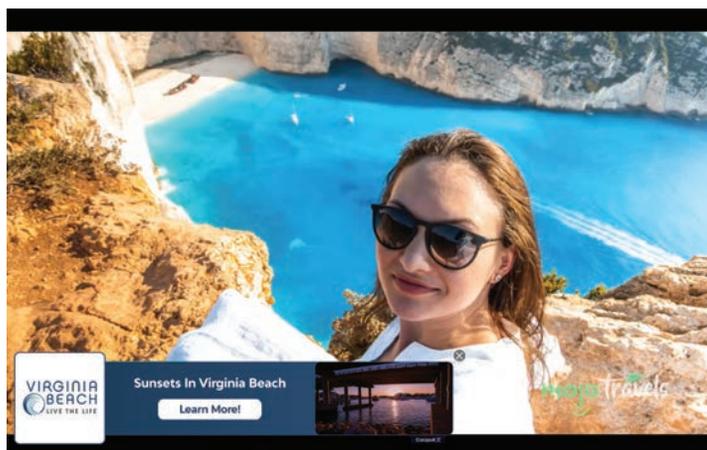
Sports has gravitated toward this type of advertising, and it didn’t hurt that CatapultX recently worked within a sports technology incubator. Its first ads went live in December 2020 and are contextually targeted. If an advertiser’s spokesperson is Serena Williams, CatapultX can deliver top CPM around content featuring her, then lower CPM around content featuring Venus Williams, and lower CPM still for other tennis players. The viewer has the option of clicking on the ad to pause the video or to continue consuming the content with the inline, or what the company calls on-stream, ads.

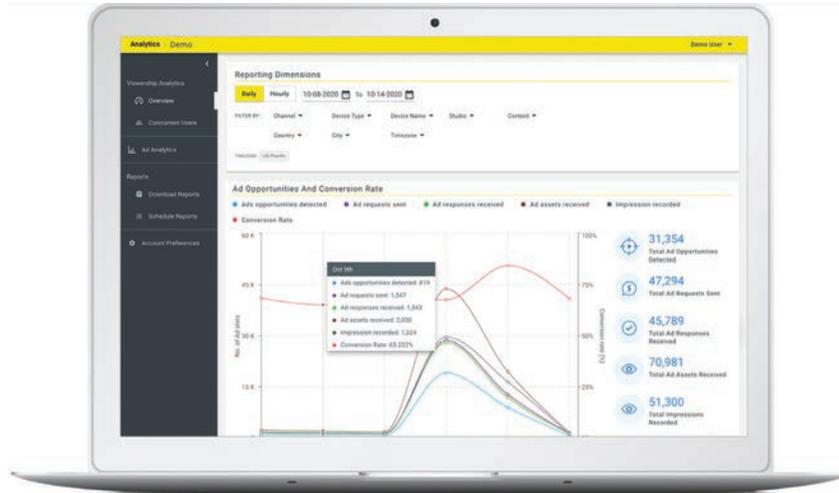
Ad Intelligence

Amagi’s SSAI platform Thunderstorm has been on the market for about 6 years. Some of the features that customers appreciate are the analytics capabilities that accompany it.

“I was talking with a customer who said that they completely revamped their programming based on the analytics that they actually got,” notes Amagi co-founder Srinivasan KA. He says the customer was able to get answers to a number of questions: Should I have 24/7 unique programming? How many viewers fall off if I

CatapultX’s On-Stream delivers contextually relevant video ads on top of the main content, and the new format has been more effective than traditional pre-rolls.





Amagi brings granular analytics to SSAI with its Thunderstorm platform.

do repeats? What sort of ad loads should I have? What's the trade-off between monetization and viewership drop? What sort of stickiness is there in terms of average time spent across, and when do the viewers fall off?

Amagi's platform has three common use cases. "The content partner has cue points or SCTE markers within the stream. When we see the SCTE markers, we make the ad request. We transcode the ad in real time and serve in the format that the platform needs based on the specific profile," says Srinivasan KA. Amagi also does real-time header bidding, in which it provides the request to multiple partners at the same time. "Based on the highest bid or the second highest bid, we would insert the appropriate ad in real time." The third use is for live streams that do not have SCTE markers. "We can recognize the ads in real time using [machine learning] and can replace those ads in those linear streams without any markers."

The New DSP

The Trade Desk's Solimar DSP provides advertisers with data-driven tools that enable them to leverage the power of their first-party data and still abide by the latest privacy requirements.

Solimar uses AI capabilities, so marketers can optimize their campaigns and tie their ad spend to real business outcomes. "The result of more than two years of product development, Solimar addresses key concerns for today's marketers, including easy and secure onboarding of first-party data; the need to connect marketing performance to business growth goals;

an increasingly cross channel digital media environment including the fast growing world of CTV; and a rising focus on digital identity," according to a company press release.

Marketers are grappling with seismic changes—including data privacy requirements and shifts in TV consumption—that make delivering ROI and performance measurement even more important. Solimar shows how decisions impact the success or failure of a campaign, and The Trade Desk's advanced AI tools are there to help ensure the campaign is a success.

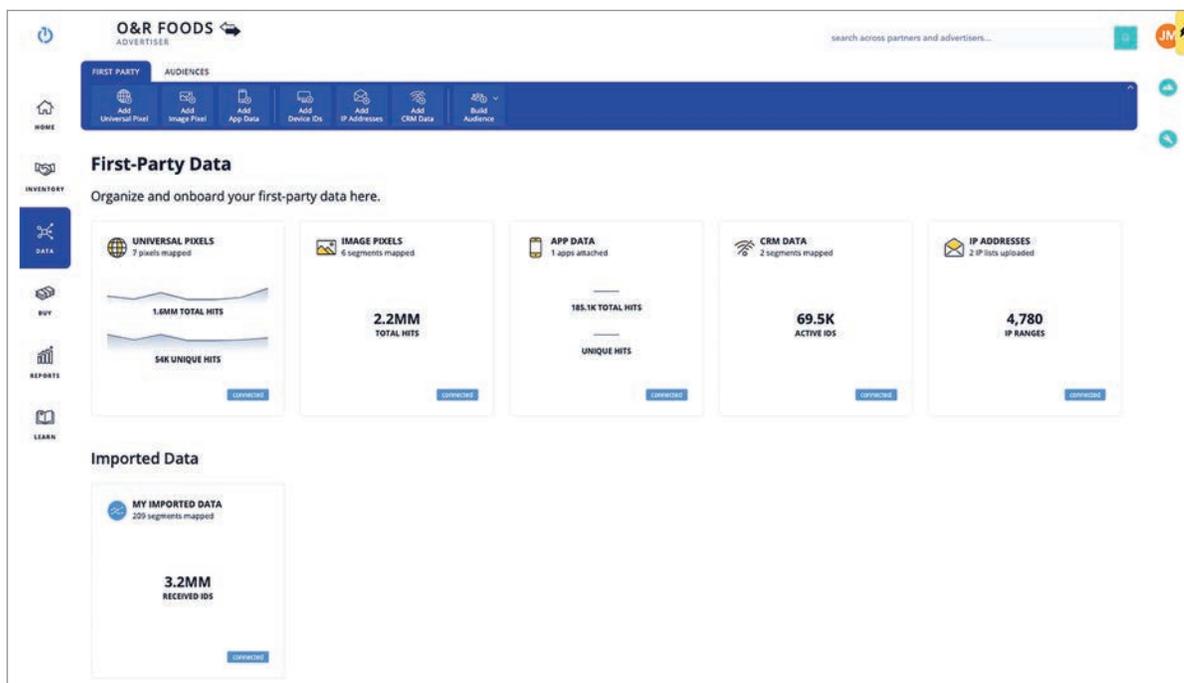
Marketers can onboard their CRM and Unified ID 2.0 data (the data identification standard that has been promoted by The Trade Desk). "Solimar connects this first-party data to pseudonymized data, such as data generated by [Unified ID 2.0]. This allows brands to find new customers who share the same characteristics as their most loyal consumers," according to the press release.

The New SSP

If one of your biggest complaints about ad tech is that there are too many moving parts, FreeWheel feels your pain and has created a one-stop solution with its supply management platform Monetization Rights Management (MRM).

"[This is the platform used by the] majority of the big TV media companies for their digital and dynamic operations, as well as the big cable and satellite distributors. It helps them to manage and monetize their inventory across both linear and digital," says David Dworin, VP of advisory services and trust at FreeWheel.

The Trade Desk's Solimar demand-side platform uses AI to let marketers leverage first-party data without running afoul of privacy requirements.



FreeWheel built capabilities to coordinate direct-sold and programmatic advertising into one source. According to a company press release from 2020, “By combining an ad server and SSP into a single system, FreeWheel has the functionality to consider all factors that impact both a publisher’s advertising clients and viewers, including:

- Programmatic demand across multiple sources, including guaranteed deals and private marketplaces
- Rules, objectives and requirements of the buyer
- TV-level compliance, such as category and competitive separation requirements
- High-quality user experience considerations, such as ad repetition and relevance.”

SSPs allow publishers to set pricing floors as well as other controls, like rejecting certain types of ads or assigning the amount of inventory to sell programmatically. MRM connects different types of demand coming in from agencies, programmatic partners and resellers, and publishers. “If you think about a world where video is more fragmented on all these different devices, you have both linear and TV-style inventory,” Dworin says. “You might have multi-

ple brands that are operating separately. You might have multiple platforms. We help them to bring that all together into a single place.”

Some publishers will optimize for fill, showing as many ads as they can. Others will focus on filling fewer spots with premium inventory at a higher price. “For a buyer, it means that their ads are in front of eyeballs that they wouldn’t have access to before, because there’s actually not nearly enough inventory, especially at scale in streaming,” Dworin says.

The MRM platform also makes programmatic buying possible for publishers that may have wanted to buy programmatically but didn’t have the relationships needed to do so. “Normally, what happens is you get some programmatic demand, and you have to either wait for it, and you might miss out on good opportunities, or it flows between different systems, and you might lose things from a workflow perspective due to latency,” Dworin says. “By moving everything into a single system, we improve the workflow.”

Nadine Krefetz has a consulting background providing project and program management for many of the areas she writes about. She also does competitive analysis and technical marketing focused on the streaming industry. Half of her brain is unstructured data, and the other half is structured data. She can be reached at nadinek@realitysoftware.com or on [LinkedIn](#). Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

BEYOND SOFTWARE

By Tim Siglin



*Highly efficient
and cost-effective
hardware processing
for streaming at scale*

I'm going to start off with a realization that hit me while writing this article: Every codec that I've ever used over my 30-year media career started as a hardware-based codec. And even though these codecs worked, almost every one of them eventually went all "soft." And, as a result of moving to software on general-purpose computing platforms, each of those "soft" codecs is much less efficient today.

With the exception of the first few months of my career—when I programmed racks of multimedia slide projectors to a cue track on a reel-to-reel tape, before convincing my boss to get a Pro Tools 442 audio system—the co-

dec in each media device I've used started as specialized hardware. Here's a quick list that may be either foreign or nostalgic, depending on how long you've been in the industry:

- My first broadcast field camera (ENG for you news folks) was BetaSX, which used an MPEG-2 hardware encoder.
- The first small form-factor prosumer camera I used, the Sony mini DV-based VX1000, was based on an MPEG-2 encoder to fit content on a small cartridge.
- My first HD camera, using the same mini DV tape, was a Canon XL2 that utilized—



in a really odd way—an MPEG-2 encoder to encode “almost 1080p” and then a hardware decoder to extrapolate “almost 1080p” to actual 1080p for playback.

- The first videoconferencing system I used, the Rembrandt II/VP, had an MPEG-2 encoder and encryption module. (It was the MilSpec/Department of Defense version, but I assume even the enterprise version had the same.)
- All of the videoconferencing systems I consulted on, from Polycom to PictureTel to Lifesize, used hardware encoders, including H.261, H.263, and H.264 codecs, long before there was even a thing called streaming. In fact, hardware encoders were so embedded (no pun intended) into videoconferencing that my consulting recommendation to move toward soft clients (desktop apps to allow participants to join a videoconference from their desk) received a response somewhat equivalent to Bill Gates’ famous quote about not needing more than 640KB of memory. In other words, hardware was the present and future, until Apple introduced Mac-software-based QuickTime Videoconferencing and Microsoft followed suit with Windows-based NetMeeting.
- The first Internet Protocol Television (IPTV) encoders I worked with were all MPEG-2 hardware running on ASICs (more on those in a bit) for hundreds of channels that could be tuned by quality, bandwidth constraints, or even targeted geographies without requiring a reboot or software reload.
- AV integration solutions, which required zero latency, all relied on hardware for encoding, trans-rating, and scaling in order to stay within one field of video (that’s 16 milliseconds, in case you’re keeping score). Even to this day, when there’s a need to design a video matrix or synchronize multiple monitors in a given room, a combination of a master clock and hardware decoders is used to maintain consistency for devices that need to run flawlessly for weeks, months, or even years.

- Early H.264 adoption in streaming, once MPEG-1 and MPEG-2 encodes moved toward software but were still unable to generate compression efficiencies at scale, required specialized hardware.

This is by no means an exhaustive list, as there are numerous silicon-based H.265 encoders in the field in rugged, fanless cases that use just a trickle of power to do their encoding task. We’ve covered a number of these in past articles, spanning market verticals such as oil and gas exploration, remote military operations, manufacturing, and the automotive sector. But this article is geared toward understanding the benefits of purpose-built hardware encoders and decoders as they relate to real-world streaming and video transport scenarios.

It’s also exactly why Dom Robinson and I explored the “greening of streaming” in an article a year ago, in which we recommended the industry add a third P (power) to the two other P’s we’ve gauged streaming solutions by: price and performance (go2sm.com/greening).

Whether you use the cloud or a hybrid solution to encode your live or on-demand streaming content, high-efficiency encoding has come of age. This article explores the math and science behind these at-scale solutions that are not only good for business, but also good for the environment.

Software Solutions

Why is software so popular? The short answer is its flexibility in programming video-specific workflows in the field, using field-programmable compute engines.

Around the end of the first decade of streaming, the H.264 codec had been optimized enough to work with general-purpose computing chips (what we’d call CPUs) and was soon optimized to take advantage of graphics processor chips (GPUs) to output more content from a single generic server. But unlike MPEG-1 and MPEG-2, in which numerous high-density hardware solutions arose that fit the five-nines model for telecom operations (99.999% uptime) and allowed for early IPTV delivery, the move to streaming with H.264 went in the opposite direction: more software-based workflows.

It's a fairly typical pattern. In the first few years of a new codec's lifecycle, dedicated hardware is required to keep encoding times within a reasonable range. But as codecs mature and feature sets solidify, most encoding moves to a purely software solution that can run on general-purpose CPUs and, if more raw processing is needed, CPU workloads can be pushed on to a high-end GPU sitting in the same server.

Yet these software solutions aren't all that efficient, certainly not at scale, consuming more processing power—and actual power—than should be needed.

And that's where high-efficiency processors that go by strange acronyms like ASIC, DSP, or FPGA come in to play: At scale, these processors produce more quality encodes faster and at a much lower power rate than standard CPUs and GPUs.

Why Hardware Instead of a 'Generic' Server-Based Cloud?

I'll explore the acronyms later, but if you're still with me on this look at dedicated hardware being used to compress at highly efficient power rates, I'm sure one question is on your mind: Why does any of this matter if we all have access to significant cloud resources?

One reason is that the cloud isn't all that efficient. Yes, it's accessible 24/7 across the globe, which is convenient and means that instances can be spun up rapidly and torn down almost as quickly. This is perfect for services that don't require dedicated hardware, but it's not so perfect in the fact that the generic computing hardware on which these instances run needs

to be kept in an always-on operational state that consumes significant power when the generic hardware isn't being used to full capacity.

And then there's the question of exactly what full capacity entails, as most systems architects will consider generic servers to be at capacity when either the processor load or the network interface card (NIC) is above 60% utilization. That's understandable, because systems architects have to accommodate potential overhead and want to avoid peak capacities choking a generic server. But it's not efficient.

Another issue with a generic server-based cloud is that the cloud itself adds inherent delay. On-prem encoding can be used for both on-prem and internet distribution, as we'll see later, with the added benefit of synchronization among encoders, decoders, and monitoring equipment.

A Practical Use Case for Hardware

Remember at the beginning of the article when I mentioned AV integration solutions? These are the kinds of solutions that are installed in corporate boardrooms, training centers, or enterprise auditoriums. They're also used in entertainment venues, where multiple cameras and screens need to be synchronized and operating conditions are less than ideal for both local display and live streams.

You choose these hardware solutions when the video absolutely has to stay in sync among dozens of monitors that can be seen from one vantage point in a venue. (I'm looking at you, regional ice rink wiener-dog-race participants that I got to watch a few months back between hockey game periods.)

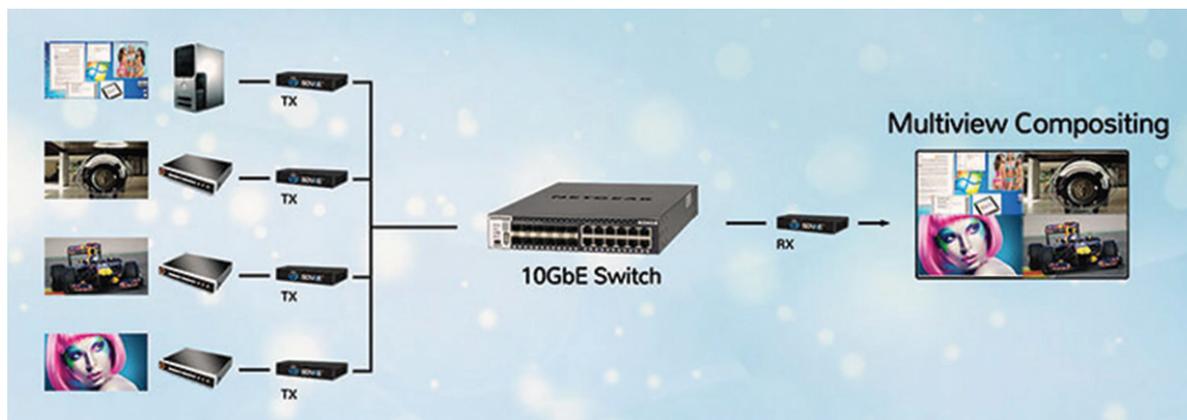


Figure 1. Software-Defined Video over Ethernet's (SDVoE) zero-frame latency encoders can downscale the incoming video image, allowing multiple encoders' video images to appear on a single screen. Known as multiview compositing, this many-to-one display scenario takes advantage of Ethernet transport to eliminate the need for an expensive matrix switch. (Image courtesy of the SDVoE Alliance)

You've probably experienced the disconcerting effects when a computer-based video mixer (e.g., TriCaster) or even software package (e.g., OBS or Wirecast) has far too many frames of video delay to use as both a video mixer and a source for image magnification (IMAG) within a given venue. The speaker at the front of the room raises her arm to emphasize a point, and a second or so later, the IMAG monitor in the room, be it a flat panel or a Jumbotron, shows her raising her arm.

How do the AV integration solutions get around this unnerving visual delay? Hardware encoding.

Sometimes, compression is added, but with many hardware solutions, the data path is wide enough—much wider than a general-purpose processor or even a GPU—that compression isn't necessary for the local display portion.

One of the best examples of this is a solution designed and licensed by the Software-Defined Video over Ethernet (SDVoE) Alliance for use in a variety of products offered by Black Box, IDK, Netgear, Semtech, ZeeVee, and others. I covered the SDVoE approach 2 years ago (go2sm.com/sdvoe), but here's a recap.

The solution offers zero latency, passing through uncompressed 4K UHD up to 4:4:2 video and adding a light compression to content to the maximum level of HDMI 2.0 (meaning full 4K 4:4:4 60Hz support). "Our compression codec, when enabled, adds 5 lines of latency," says Justin Kennington, president of the SDVoE Alliance. "At UHD, 60Hz, that's 7.5 microseconds, which blows away even I-frame-only AVC/HEVC, etc."

What Kennington didn't mention was that the solution does so while also carrying uncompressed 7.1 sounds, adding AES-128 encryption and supporting 12-bit color depth (this means it easily supports HDR10 and HDR10+, which are only 10-bit color space). And it does all of this on a prepackaged hardware encoder/10Gbps network interface that's about half the size of a pack of playing cards. The encoder runs so efficiently that it needs zero airflow across the two chips—an FPGA and a 10G PHY—so it can be shoved into small enclosures.

In addition, it can be combined with half a dozen other encoders in something that fits in a quarter of a rack unit. (Typical dimensions

of these multi-encoders run about 6 inches deep, so they can be mounted in AV racks instead of requiring the larger and deeper data center rack that most generic servers need.)

In other words, the performance and power efficiency allow for smaller, cooler devices that can also be put into hibernation while awaiting the next task. Try doing that with your generic server based on decades-old designs.

Types of Dedicated Hardware Compute Engines

Now that I've talked about a practical use case, I'll spend the rest of the article discussing the different types of dedicated hardware.

First on the list is the ASIC (application-specific integrated circuit). An ASIC has the potential to be the most efficient and most powerful compute platform for streaming solutions, but it's also the most difficult to "get right," precisely because it's purpose-built for a specific set of tasks. As such, great care is taken before the silicon is "spun" or committed to pressing/fabrication because any oversight could render the ASIC unusable for its intended purpose.

A handful of companies in the industry have designed and deployed dedicated ASICs. They spin silicon every 2–3 years, but the field life of the products based on these ASICs can easily be a decade or two.

Is there a way to get similar compute efficiency while still allowing for programmability to accommodate changing needs or emerging industry standards? Yes, and that comes in the form of two different system-on-chip (SoC) approaches: DSPs and FPGAs.

A DSP (digital signal processor) is an SoC that's often used to process baseband signals, such as audio or video elementary streams, that have been converted at some previous step in the workflow via an analog-to-digital converter (ADC). Think of the microphone on your smartphone, which feeds an analog signal into an ADC. This is then processed by a version of a DSP from Texas Instruments (TI) or Qualcomm to be recorded to a file, sent as part of a phone conversation, or both.

There's a true science to working with DSPs, and the high learning curve limited their use to very specialized audio or video products. About 2 decades ago, though, TI realized that it needed to lower the barrier to programming DSPs

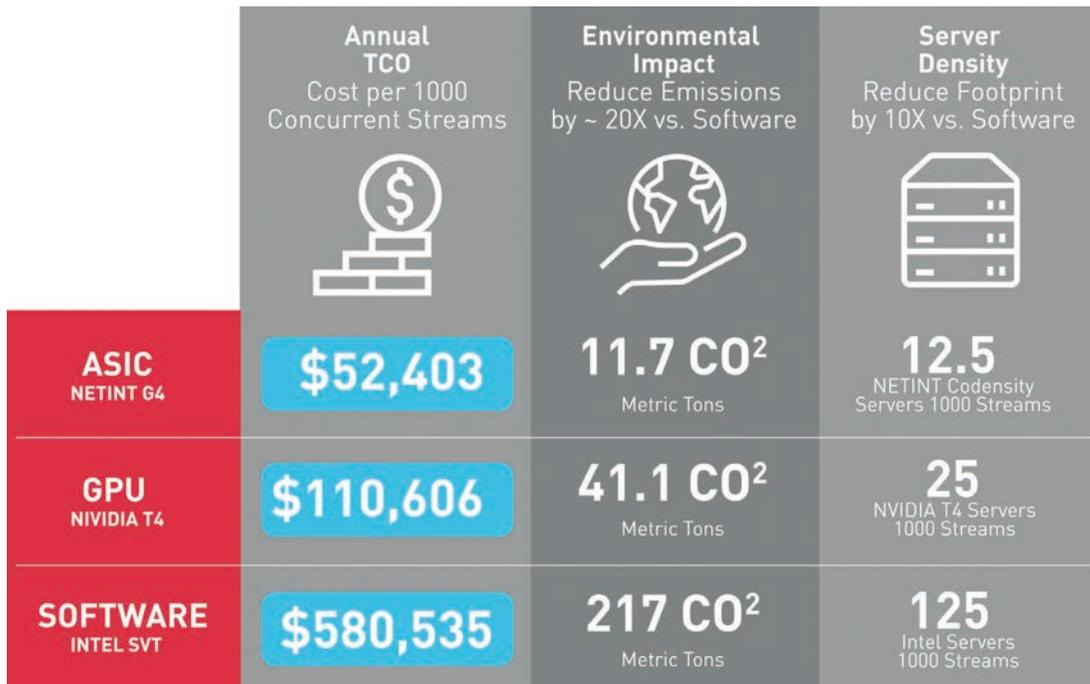


Figure 2. NETINT claims significant power reduction, at scale, over CPU and GPU encoding. (Image courtesy of NETINT)

for video compression, so it began to provide both programming interfaces (via links to integrated development environments, or IDEs) and pre-licensed audio and video codecs under the DaVinci moniker.

DaVinci allowed the programmer to focus on programming in a language environment that he or she understood to tie into video workflow tools that the video architecture team also understood. When combined with a TI DSP, the result was a much more rapid time to market, since the DaVinci software interfaces and integrated codecs eliminated the need to learn how to optimize a codec to the DSP platform.

An example of one of these chips is the Digital Media System-on-Chip (DMSoC) TMS320DM368. Its video processing subsystem was capable of 1080p at 30 fps using an integrated H.264 video processing engine. Beyond the codec, though, this particular SoC has an integrated facial detection engine, an analog front end, a hardware on-screen display, and a number of digital-to-analog converters (DACs) to allow output to a local monitor. It also contains a variety of color depth options via a 4:2:2 (8-/16-bit) interface that allows up to 16-bit YCC and 24-bit RGB888.

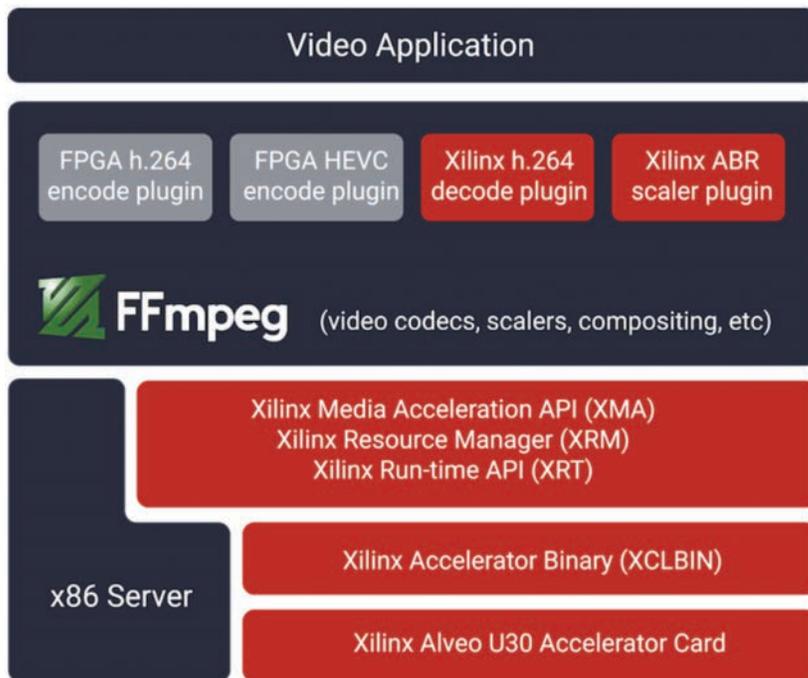
Another company that makes DSPs that are targeted at the streaming video market is NETINT, which has offices in China and Canada. NETINT has a DSP-based SoC, called the NETINTG4, which it claims is much more efficient than a GPU or even a software-based CPU solution.

As **Figure 2** shows, the SoC approach has the potential to not only scale much easier, but also to do so with significantly lower power consumption requirements. More interestingly, NETINT claims that FFmpeg workflows that were designed for CPU or GPU processing can be easily ported to use the NETINTG4 SoC.

A quick note about these claims, though. While the not-for-profit that I run, the Help Me Stream Research Foundation, normally verifies manufacturers' claims on its test bench before recommending a specific video engine, we've not yet had a chance to do so with the NETINTG4. Still, since NETINT displays this data on the main page of its website, it's worth sharing with *Streaming Media* readers.

Speaking of claims, DSPs, and DaVinci, the other approach to hardware-based encoding comes in the form of FPGA (field-programmable gate array). The solution I previously mentioned, from the SDVoE Alliance, uses an FPGA

Figure 3.
Xilinx provides two types of application programming interfaces as well as a resource manager to handle multiple FPGAs in a single server. (Image courtesy of Xilinx)



alongside the 10Gbps network interface, and the organization has explained to me in the past that it does so because of both power and performance benefits, since FPGAs provide very wide signal processing paths.

In the process of researching this article, I came across a name in the FPGA world that sounded familiar: Sean Gardner. I looked in my phone contacts and noted that there'd been a Sean Gardner at TI around the time I did some benchmark analysis of early DaVinci solutions. So I was curious if it was the same person. It was, and he's not only stayed in the video compute space, moving from DSPs to FPGAs, but he's also brought a similar make-it-easy-on-the-programmer approach to Xilinx for video solutions.

In a conversation with Gardner; his Xilinx colleague, Aaron Behman; and me, the two of them joked that they've long thought that the G in FPGA should stand for "green" based on how power-efficient FPGA solutions are. As an example, they ran me through a presentation around the Alveo U30 Media Accelerator. Again, we've not verified these claims at Help Me Stream, but if the numbers are accurate, they're impressive.

The U30 is a PCIe single-slot card that's half-height and half-length, meaning it can fit

in even small form-factor computers. The idea behind the U30 is to allow existing workflows in FFmpeg and other video tools to take advantage of an FPGA solution that offloads all video processing from the host CPU.

The card offers both hardware-integrated H.264 (AVC) and H.265 (HEVC) encoding and transcoding. At full power, it consumes 25 watts, so it's on par with power consumption of most midrange CPUs on the market. For that 25 watts, though, the U30 is capable of supporting some impressive transcoding numbers: 2 UHD 60 fps (2160p60) or 8 Full HD high frame rate

(1080p60) or 6 Full HD standard frame rate (1080p30) transcodes.

Even more interesting is the fact that Xilinx claims to have tested up to eight of the U30 cards in a single server, with the ability to generate up to 256 720p30 transcodes in a single server, as well as adaptive bitrate outputs on a scale that most GPU-equipped servers can't manage.

Gardner notes that the U30 "is deterministic in performance" from a throughput standpoint, but that the true benefit comes from lowering the host CPU to manageable levels.

"From a CPU offload perspective," says Gardner, "this reduction in CPU loading offers lower CPU cost, better thermals, and better economics at a rack level."

In other words, by increasing performance and lowering overall power output, this FPGA solution provides a way to scale streaming without adversely impacting the environment.

Tim Siglin (insights@helpmestream.com) is a streaming industry veteran and longtime contributing editor to *Streaming Media* magazine.

Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

Enhanced **Sports** and **Esports** Experiences

By Steve Nathans-Kelly

The best way for sports streaming services to build loyal followings is to deliver experiences that fans can't get in crowded arenas—or in traditional sports broadcasts either, with their action-on-the-field-centric camera views, play-by-play, and color commentary. To make online sports experiences stand out from the crowd, content creators and programmers need to offer the most immersive, interactive, and innovative experiences possible. The key to engaging today's sports streaming audiences, prominent providers say, is to satisfy fans' hunger for content and experiences that capture their attention outside the games themselves.

Only Game in Town

"If the pandemic has taught us anything," says Kristen Scott, director of digital talent production and storytelling at Fox Sports, "it's that when live sports shut down, people were grasping for alternative programming that still tangentially relates to sports. So we're creating original programming that has nothing to do with the actual game, whether that's access to the players and coaches or whether it's creating a completely made-up vote of who is the best fan base in all of college football. We saw all of that completely take off over the last 15 months."

Interest in this sort of tangential programming didn't wane when the games resumed.

It grew and pointed the way toward ongoing, enhanced sports content offerings.

“When everything shut down, we decided we were initially going to have watch parties of old games online, where we all gather and talk about them or tweet about them,” recalls Cheesehead TV co-founder Aaron Nagler. “And that extended to when the NFL got back on its feet. After the season started, whereas in the past we had done a watch party here or there, they became absolutely essential for a lot of our viewers and our regulars. Now we see fans in the stands, but a lot of the innovation that happened when things shut down is going to stick around.”

“We keep iterating upon our watch parties and other second-screen experiences, but we’re also creating complete, original programming around our marquee sports,” Scott adds. “We’re bringing in young, fresh talent who have communities that were built on the internet, who [people] are used to engaging with on a minute-by-minute basis, who are obsessed with the

data, and obsessed with how they’re performing. We’re investing in that type of talent with guys like RJ Young and Ben Verlander and Ryan Satin to create original shows that go a little bit deeper than the [NFL] RedZone broadcast on linear.”

Complete Control

One of the most innovative beyond-the-game enhanced experiences serving a growing audience today is the Fan-Controlled Football League (FCFL), which puts fans in control of the game by letting them call the plays in real time through a weighted online voting system that rewards gamer experience and demonstrated “football IQ.” The FCFL leverages the massive success of football simulation games like *Madden* and the ubiquity of fantasy football leagues by bringing fans directly into the action of live games as if they were coaching on the sidelines.

As on other channels developing enhanced sports content, FCFL is all about creating a community to enhance fan engagement. “We allow our communities to co-stream,” says FCFL co-founder and chief gaming officer Patrick Dees. “Every one of our teams is owned by a celebrity in a different vertical—people with a lot of swagger, like Renee Montgomery, Quavo, and Marshawn Lynch. As with Mark Cuban and the Mavs, we see these teams as an extension of their personalities. And we not only work with their teams, but also a lot of Twitch influencers. We’ve seen these micro-communities pop up, where they would root for their team and talk over the game. We’ve seen all of these different types of experiences, where you can just go watch it with your people wherever you are and talk trash with other fans all day long.”

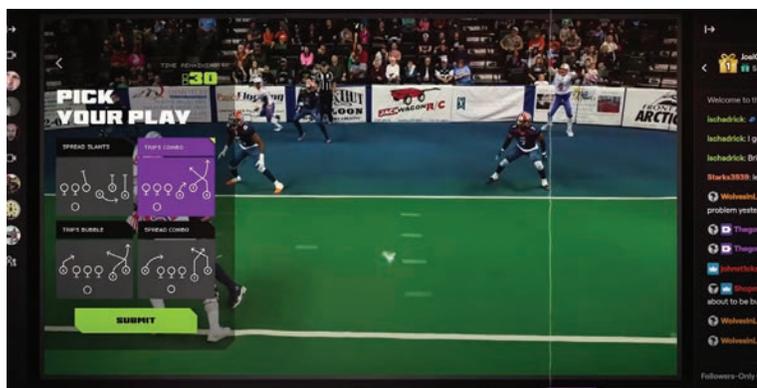
The Elephant in the Room

Gambling is arguably the elephant in the room when it comes to enhanced sports streaming

Cheesehead TV’s #CHTVDraft: 2021 NFL Draft Watch Party, Day 1



Play-calling in Fan-Controlled Football League games



experiences. Due to regulatory changes that are on the way, are underway, or have already happened—depending on where a given sporting event is occurring—betting is likely to become more integral than ever to the sports OTT experience in the coming years. The mainstreaming of sports gambling as a central element of the sports-viewing experience is having a profound effect on how sports content programmers develop new offerings. As with any other type of content, it's all about finding a niche, an angle that distinguishes your offering from others in an increasingly crowded market.

“Betting is everything right now,” Cheesehead TV’s Nagler says. “Everyone’s got a mock draft. Every content area seemingly has their betting expert. So I think that’s already a saturated marketplace. The key is gamification, extending a visit to Vegas to an online experience wrapped inside the specific content you’re creating. For us, that means tapping into whatever the betting angle is for the Green Bay Packers, whether that’s the game itself or whether Aaron Jones will put on the shades after he scores a touchdown and goes to the sideline—something that connects to your team that you can then engage with the audience about in that capacity. It’s wide open, fertile ground, but it’s hard to wrap your arms around it if you don’t come from that world. The biggest trick is finding something authentic, something that connects to your audience immediately, but doesn’t feel like it’s slapped on in a way that is going to be easily dismissed.”

“[W]e’re leaning into non-gamblers in our gambling content,” Fox Sports’ Scott says. “Maybe you’re intimidated by the concept of gambling. You don’t know what a spread means. You don’t know what an over/under means. You don’t know what a moneyline is. We want to make it approachable. We have our gambling experts who live and breathe this stuff, and we have our non-gamblers, who [say], ‘Hey, I’m just getting into this world, make a \$10 bet with me, and let’s ride it together and see what happens.’ We’re trying to make it approachable, and then we’re integrating it into all of our second-screen experiences from an overt and a non-overt standpoint, just graphically on a screen. On the more intense side, we’re creating full original programming around gambling with our top-of-the-line gambling show

that we use to specifically talk all things gambling on any given event or sport.”

During its first season, FCFL also made a conscious effort for approachability on the gambling side to ease its audience in—particularly fans who were more at home in the esports world than with traditional sports. “We focused on alternative solutions that a Twitch audience would just know and understand,” Dees says. “There’s a currency that you earn on Twitch called channel points, which are like frequent flyer miles. You watch for a few minutes, and you start earning channel points. We would let them not only bet peer-to-peer, but also bet a certain amount of their channel points on a prediction. And we saw tens of millions of channel points bet over the course of six games. I think that’s one of our key learnings: If you can make it approachable, people are going to engage with it.”

How Low Can You Go?

Betting’s ascendancy in the world of sports streaming also makes ultra-low latency more critical than ever on the delivery side. In-game, in-arena micro-betting is resurgent in a way that it hasn’t been in a century, since the run-up to the Black Sox scandal of 1919, when bookies and their minions blanketed ballparks throughout the major leagues, taking action on virtually every pitch. 5G has become the indispensable enabling technology of the real-time micro-betting trend. But how realistic is that degree of ultra-low latency in the near term?

“Especially in the arena, ultra-wide bandwidth promises zero latency,” says Jason Thi-beault, executive director of the Streaming Video Alliance. “When we’re talking about getting content from the camera to the handsets that are in that arena, we still have the latency that’s in the streaming workflow. But when you start combining 5G ultra-wideband with low-latency CMAF or WebRTC, you’re going to have people in the arena, watching the game, placing bets on their headset.” In that scenario, he continues, “You can’t have any latency, and 5G gives us that.”

Cheesehead TV recently tried offering a real-time trivia game for which low-latency interaction was crucial, and it quickly reaped the whirlwind of a less-than-level playing field. “We started doing live Packers trivia via YouTube, using the YouTube API pulling hashtags,”

Nagler says. “We would have four answers for each question, and you had to hashtag your answer into the chat. The computer would pull the five quickest correct answers, and then we would add the winners to a leaderboard. This was literally a no-stakes game—just people on the internet gathering at the end of the day to have a good time. And the amount of vitriol and very sternly worded emails I would get from people who said, ‘I got all the answers right and my name didn’t appear!’ I can’t even imagine having people’s money on the line. That’s only going to become more and more important as gambling turns into even bigger business, because everyone’s on a different Wi-Fi, and everyone’s got a different internet speed. Everyone’s coming technically from a different spot. How do you make that uniform and create a level playing field? That’s going to be an amazing challenge to watch as things get even faster and even more immediate.”

Feeding the Fandom

Another key aspect of enhanced sports streaming involves adding interactivity to the in-game experience, so it transcends the traditional lean-back lifestyle that’s only intermittently interrupted to yell at a ref or launch projectiles at the screen. How much interactivity are fans actually looking for, and what can sports streamers provide to satisfy that desire to engage with the game in a more active way?

“We view even our linear programming as a two-way conversation, a two-way experience,” says Fox Sports’ Scott. “Every single *Big Noon*

Saturday, [college football] fans get to choose a segment on the show. So it’s like fan-produced TV. On the digital side, they have access to all of our talent, via asking questions that ultimately lead to some of our best [video-on-demand] moments that are evergreen and live on well beyond that live experience. On the more technological side, we are producing bonus cameras so you can watch in-car with any number of NASCAR drivers while the race is going on. Fans get to vote on which drivers they want to see streaming on our app and our website. So the amount of interactivity we offer is just going up, and we continue to innovate along with the fans’ demand.”

“We have a technology we call ‘command center,’ which is your traditional drone shot—the hero cam—plus six or eight 180-degree helmet cams and a ref camera as well,” says FCFL’s Dees. “In real time, fans can bounce between any of those views. People use it mostly for replay, like after a crazy Hail Mary, where they could go to the player and watch it happen. We’ve seen a lot of great feedback on this.”

“The appetite for interactivity is insatiable,” says Cheesehead TV’s Nagler. “You need to offer a wide variety of ways for fans to consume your content or connect with you. We’ve got people who are all in and watch every single video and consume every single tweet that we put out and every Patreon or Zoom hang-out we do. And then there are people who are just passersby, and we want to make sure that we have something on offer for all of them. A couple of years ago when we first started the

Fox Sports’ in-car NASCAR camera



watch parties, we'd do our little pre-show, and we would tell people to send us pictures: 'Where are you watching the game tonight or today?' And that has grown to where I can't keep up with all the entries. So now we have a spot in our pre-game show where we've got people from Brazil, Australia, the Philippines—literally all over the world—checking in, and they're more excited about their picture being seen on our show than they are for the game. Our next iteration of that is what we call 'Carry the G,' where we're going to have a spot on the site for people to upload video from their tailgate parties. And hopefully, we expand that into little video peeks into Brazil, the Philippines, Germany, London, et cetera, just to keep connecting to people because that's what we love—sports and community. The idea is that this person might be on the other side of the world, but they're just as excited about this event as I am."

Telling Stories

One of the hallmarks of successful live event streams that create unique viewing experiences is storytelling. For traditional sports viewers, the game *is* the story. But interactive streams aren't traditional broadcasts, and the emerging generation of fans whose expectations have been shaped by different types of media experiences isn't made up of traditional fans. How important is storytelling to enhanced sports streaming, and how can sports streaming services fashion stories that don't intrude on or drown out the story of the game itself?

Nagler admits that in the early days of Cheesehead TV, he had a hard time imagining where the audience for his company's brand of ancillary content would come from. A lot of fans "are just reclining in their La-Z-Boys, just watching the game," Nagler concedes. "But there is a whole generation that doesn't even know that existence. They want their hands in the pie. There's a wide palette of engagement and storytelling and things that can augment viewers' and fans' experience of the game. When I was growing up, journalists stood on the side and told what happened. They weren't part of the story. That is so blown out of the water at this point."

FCFL's Dees says that developing additional storylines—far from being a distraction—has proven critical to the league's success. To build the league's roster of players, he says, "We had camps and combines, and then we did a cast-

ing call. We hired a casting agent to find the guys with the biggest personalities that lit up when the camera turned on them, guys that wanted to create content outside of game day and engage with fans."

What emerged, Dees says, was a sort of "virtuous circle": "As they're out there creating content on Twitch and being active on social and playing *Madden* with their fans and those types of things, we saw their number get called more on game day. So the story changes. They're actually getting more reps on game day because they've built a relationship with fans in a way that never existed before. Then, if you fast-forward two or three years, you'll look at the fan leaderboards and say, 'This fan, she's killing it.' And then you'll see not only players getting recruited to other teams, but also that fan, because she's been number one on the leaderboard for three seasons, and other fans will want her to be part of their fandom. So it creates new storylines and new opportunities for experiences."

"We're dealing with a different audience these days—a distracted audience," says Fox Sports' Scott. "They're scrolling through Twitter while they're watching the game; they're doing other things. If you're not taking a slice of that pie, you're losing. So you have to create a more interesting experience for them in order to grab that attention. If you're just providing them a base-level experience, they're going to look elsewhere."

The challenge, she says, is drawing out storylines that resonate with fans. But, often, those stories emerge organically when they provide fans with real opportunities to interact with marquee on-screen talent. "The storytelling that comes from some of the questions that fans ask of our talent and star-power guests goes well beyond the live window," she says. "On one show, Joe Montana told fans that he would get nervous on the sidelines during games and call his wife from a phone that's not even supposed to dial out. That story continues to accrue views and engagement and lives on forever on our platforms. That's super valuable for us."

Steve Nathans-Kelly (steve@streamingmedia.com) is editor of Streaming Media Producer (streamingmedia.com/producer) and video publishing director of Information Today, Inc. Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

streaming toolbox

Vistex, Agora, and Transmit

By Nadine Krefetz

This installment of Streaming Toolbox is all about tools to help run streaming businesses better. The first company manages content licensing, the next provides transport protocols and software development kits (SDKs) for interactive media app developers to ensure low-latency media delivery, and the last offers a new spin on advertising. If I were starting a new streaming service, say Nadine+, I could and would use all of the following products. Now the first question is, what kind of content should I offer?

VISTEX GTMS

The boom in streaming has made tracking royalties and licensing exponentially more difficult, and that's where Vistex comes in. Broadcasters, sales and distribution companies, studios, and direct-to-consumer platforms use Vistex's Go-To-Market Suite (GTMS) for licensing and royalties management. Vistex claims to be one of the larger companies in this space, and recently it created a cloud version of its applications—or rather, two versions. One is directly from Vistex, and the other is available via SAP (which also happens to be an investor in the company).

I'll focus on the Vistex direct product, which manages an inventory of rights. Once you have that information, you can use your content in whichever way you need, says Amos Biegun, global head of rights and royalties for Vistex. "A very common use is in sales and distribution to run availability reporting, or avails, to see what I can sell." Another use is paying out royalties to content owners.

"There isn't a typical use case, because what I've always found in media is that every use case is slightly different," says Biegun. So we talked through a hypothetical streaming service that he and I would create.

STARTING A STREAMING SERVICE

"We own nothing. We go to the market and procure content. The first thing we need is an inventory of the content that we've just procured and the rights associated with it," says Biegun. "So what tends to happen is we could go to 10, 20, 30 different suppliers of content. Each one will grant the license to us and provide us with content which has contract and rights information that differs. So the first thing we need to do is homogenize that into a single standardized repository of rights."

Then, GTMS allows us to start scheduling content for specific markets, territories, periods of time, languages, and different elements cleared for the market we'll be broadcasting in. After content is consumed, GTMS tracks any associated payments or royalties and calculates them on the contract and what is due to the parties that provided this content.

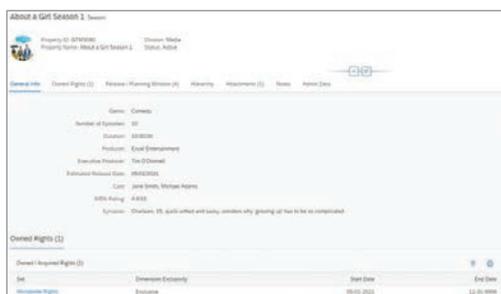
"If you're getting data that's using standards like [Entertainment Identifier Registry], that's extremely helpful for us, but even if it's non-standardized data, we have tools that will help you ingest the data, map it, embellish it, validate it, and then ultimately store it using this best-practices data model that we configured," says Biegun.

FUTURE BUSINESS MODELS

GTMS has the ability to support future business models like payment-by-viewing-session rather than upfront payments. "We've allowed customers to configure ever-evolving business models," says Biegun. "That's one of the areas where we have done a lot of forward thinking on streaming." So, if a producer decided to sell content and receive royalties on viewing, Vistex would be able to track these transactions.

Vistex offers no public demo or pricing.

Vistex's Go-To-Market Suite (GTMS) allows streaming services to more easily manage rights and royalties via clear, explicit dashboards.



streaming toolbox

Vistex, Agora, and Transmit

AGORA

Agora's platform as a service provides tools for creators to deliver low-latency, interactive, two-way video via its proprietary transport layer. The goal is eliminating the delays caused by a traditional CDN, so interaction can happen naturally without pauses. Some of the types of applications Agora helps power include interactive live streaming, real-time messaging, interactive whiteboards, video calls, and interactive classrooms. It also supports augmented reality and virtual reality applications.

"Most of what you see on Agora is about engagement, because we believe people want to engage with the content rather than sitting there and just watching it passively," says Reggie Yativ, Agora's chief revenue officer and COO. "Traditional media companies understand that engagement is going to keep them current."

TRANSPORT

Agora provides building blocks, via an SDK, to connect to the software-defined network running in its 200-plus data centers. Lately, these data centers are seeing a lot more mobile traffic and an increase in web traffic with video. "There's a lot more moving parts in video than people are aware of—the location of the audience, the location of the host, what kind of network situation will impact you, what will be acceptable latency for the use case?" says Yativ.

The self-serve platform is geared to developers. According to Yativ, "The developers will do the implementation, but there's a whole ecosystem around the developers that will make decisions, and we'll participate [with] product managers that design how the app would look, user experience experts, producers, and, of course, ... CIOs and CTOs.

"We are not involved in user interface. We do have, however, a lot of sample codes and sample apps that developers can go in

and download and customize," says Yativ. "This shortens their development cycle significantly."

USE CASES

Some sports broadcasters are choosing to deliver a more interactive game environment along with their live streaming. "You can promote an audience member to become a 'host,' and now they can jump into that conversation," says Yativ. This calls for incredibly low latency: The lip-sync needs to be perfect, and the video image should not blur or stutter. Without Agora, developers would try open source or DIY development—which can work until scaling becomes a requirement.

One customer, HP, wanted to embed live interactive streaming into its Omen gaming environment. "HP sells millions of them a year, and they want to give the Omen community the option to stream a game or hold watch parties," says Yativ.

"If you're a developer building a virtual classroom or enterprise collaboration software with video, voice, whiteboarding, or recording, we will provide you with ready-made packaged SDKs to embed into your app," says Yativ. "Your video will seamlessly reach every corner of the planet without you having to worry about quality, latency, and scalability.

"Many of our customers integrate the Agora platform into virtual tipping platforms, so

Build it with Agora
Agora provides the building blocks to enable a wide range of **real-time engagement possibilities**.

- Video Call
- Voice Call
- Real-Time Messaging BETA
- Live Interactive Audio Streaming
- Real-Time Recording
- Interactive Whiteboard

The diagram features six icons representing different real-time engagement capabilities: a video camera for Video Call, a microphone for Voice Call, a speech bubble for Real-Time Messaging BETA, a speaker for Live Interactive Audio Streaming, a recording button for Real-Time Recording, and a whiteboard for Interactive Whiteboard.

Agora offers a platform as a service that strives to eliminate latency delays for real-time communications and live streaming.

you can also reward a host or audience with virtual tips,” says Yativ.

The company went public last year, so it was a bit limited in what it wanted to promote regarding its customers, but said it has many media companies on board.

Agora offers no public demo or pricing.

TRANSMIT

Transmit is trying to break media companies away from the 2-minute ad pod by offering a contextual in-stream, picture-in-picture ad format. The company believes the advertising market has been confined by a lack of innovation, especially as more live streaming is taking place, says Rob Friedlander, Transmit’s head of corporate development. “The big technical differentiator that we’re bringing to market is the ability to insert contextually relevant ads into live streams. This could be an MLB stream. This could be a live music event. I think what we’re seeing in the market today is that the premiums on live rights have never been higher.

“Ad breaks disrupt content with the 2-minute ad pod. That has been the accepted norm since the beginning of television and video,” says Friedlander. The company’s platform in-

tegrates relevant ads in moments when the viewers are susceptible to those experiences.

“We have the capabilities to not only ingest the live video, but also ingest live data from sports data providers that are pulling in event data,” says Friedlander. “We know when someone strikes out, when someone hits a home run ... and those become relevant signals for us to put an ad in front of somebody, because that’s when attention is usually eyes off the play.”

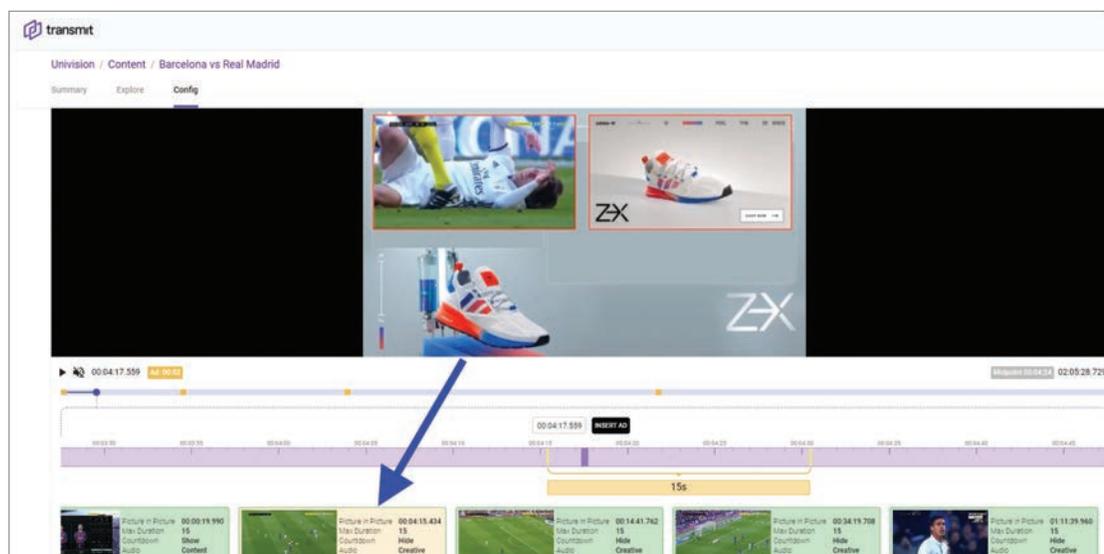
Transmit can create a two-box display within a dynamically skinned background and put an ad in one of those boxes. “Not only do we have the play continuing in the top left box, but you have the ability to dynamically frame and insert sponsor messages,” says Friedlander.

“Similar to how other SSAI vendors would work, we sit in between the original content and what the user sees,” says Michael Celona, Transmit’s CTO. “We certainly lean in with the CTV experience, but then quite a few folks are also leveraging us on the web as well. We have plug-in integration with JW Player,” says Friedlander.

GO TEAM

Transmit found that sports was an underserved OTT market. “Soccer, as an example, is a game you’ve not been able to monetize because

Transmit inserts contextual picture-in-picture ads into live streams. For instance, a sneaker company can insert its ad into a stream that includes a soccer player wearing its brand.



streaming toolbox

Vistex, Agora, and Transmit

there's really no natural break in play. So all of a sudden, we're creating value for the rights-holder and the programmer and integrating ad experiences that are additive and relevant to the programming," says Celona.

Transmit takes the video feed, which could be HTTP Live Streaming or Real-Time Messaging Protocol, then encodes the ad (usually a static MP4 file) on top of the feed and brings in assets like a logo or lower third. "Then you have some accoutrements that would sit underneath or take up the negative space," Celona says.

"We have direct-sold campaigns that Transmit is procuring through our partnerships with

the advertising brands. Sometimes our clients have their own ad tag," says Celona. "It doesn't disrupt any of the existing ad tech or workflow that exists in the agency model or any of those kind of relationships."

Transmit offers no public demo or pricing.

Nadine Krefetz has a consulting background providing project and program management for many of the areas she writes about. She also does competitive analysis and technical marketing focused on the streaming industry. Half of her brain is unstructured data, and the other half is structured data. She can be reached at nadinek@realitysoftware.com or on [LinkedIn](#).

Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

DVEO : DIGITAL PRODUCTIVITY

ELIMINATE PACKET LOSS AND CORRECT INTERNET JITTER

DOZER SRT:Box
Eliminates Packet Loss & Jitter For Any Bitstream

- COMPLETELY CORRECTS PACKET LOSS IN TRANSMISSION PATH
- SOLD IN PAIRS OR POINT-TO-MULTIPOINT
- STREAM LIVE NEWS OR SPORTS PROGRAMMING OVER CONGESTED WIRED OR WIRELESS/ LANS OR WANS
- COST EFFECTIVE ALTERNATIVE TO SATELLITE CONTENT DELIVERY

sales@dveo.com | + 1 858 613-1818 | www.DVEO.com

review

The Roland V-160HD Hybrid Event Video Switcher

By Robert Reinhardt

In July, I was given the opportunity to review a pre-shipping model of the Roland V-160HD video switcher (**Figure 1**), which is marketed as a product that “sets a new standard in hybrid event switching.”

Product-specific hype aside, hybrid event switching itself could become a new standard. As event organizers are allowed to open up venue spaces with in-person attendees, there will be an expectation to enable a virtual event experience for remote presenters and attendees.

While detailing every feature of this switcher is beyond the scope of the article, I’ll present the top 10 features that I fully appreciate on the Roland V-160HD, peppered with some constructive criticism for good measure.

Full disclosure: I’m a fan of Roland Pro A/V video switchers. My production company owns two Roland switchers, and I’ve written a review of the Roland V-60HD as well (go2sm.com/v-60hd).

1. USB Output

To enable an experience for remote attendees, the Roland V-160HD switcher is certainly well-equipped to use the switcher’s output as a USB camera source in desktop operating systems and applications that support USB video class (UVC) and USB audio class (UAC), such as Zoom, Microsoft Teams, and other conferencing applications. Many WebRTC web apps and hosting providers support capture from locally connected USB camera sources, and the Roland V-160HD can be used with those as well.

The Roland V-160HD has the newer USB-C-style port (bottom row center in **Figure 2** on page 38), and you can use any number of USB 3-to-USB-C cables on computers without USB-C ports. The beauty of this simple USB camera source feature is that you no longer need to set up a dedicated SDI or HDMI output from the switcher to use with SDI or HDMI

Figure 1.
The top console
of the Roland
V-160HD switcher



review

The Roland V-160HD Hybrid Event Video Switcher



Figure 2. Inputs and outputs on the Roland V-160HD switcher

capture hardware on your computer, as I've had to do with other video switchers.

2. Input Assign

This feature refers to two enhancements in the V-160HD that I haven't seen in previous Roland switchers I've used: eight independent HDMI and eight SDI inputs (on the right in Figure 2) and the ability to quickly assign the input source for each of the 10 cross-point buttons. For the former category, the SDI and HDMI input number on the back of the unit does not need to correlate to a specific switched cross-point button on the console.

With older switchers, the SDI inputs were dedicated to specific console buttons, as were the HDMI inputs. If you didn't have HDMI inputs, for example, those console buttons would never be used. With the V-160HD, however, you can mix and match any of the HDMI, SDI, or still-image bank resources to any of the 10 cross-point buttons.

Perhaps even more importantly, you don't need to dive into the console setup menu to hunt for the assignment of inputs. The V-160HD has a convenient Input Assign button on the console (just left of the T-bar in Figure 1). Hold this button, and use the PGM/A or PST/B buttons (lower left in Figure 1) to cycle through the various inputs for the desired cross-point button. This simple operation greatly cuts down your setup time for a live event.

3. Four PinP or Key Composition Layers

Do you want to insert four presenter headshots over a PowerPoint or Keynote slide deck

or add a green screen ASL interpreter shot to a specific corner of the program feed? The V-160HD switcher can do it, and it has more console buttons to make input position, selection, preview, and program output effortless without (again) the need to dive into nested menu options. The switcher has all four PinP (picture-in-picture)/key compositors available on the console.

Like other Roland switchers I've used, PinP has the standard options for shape mattes and border effects. The keyer can use luminance (white or black) or chromakeys, and the V-160HD adds a sampling marker feature that enables you to pick a chromakey color directly from a selected keyer input.

4. Two DSK Overlays With Alpha Channel Support

The V-160HD switcher also features two independent Downstream Keyer (DSK) compositors. While the PinP/key compositors are designed to use specific regions of the output canvas, the DSK compositors can be used for input sources that need to overlay the whole screen, such as lower-thirds for presenter names/titles, game scoreboards, and more.

My favorite enhancement to DSK overlays is the ability to use PNG image sources with 24-bit transparency or alpha channels. With older Roland switchers, knocking out the background required you to output a BMP image asset with a chromakey background (e.g., lower-thirds on a green or blue background) and adjust the keyer settings. Now, you can have super-clean alpha channels created in tools like Adobe

Photoshop to produce beautiful DSK overlays without any guesswork on chromakey settings.

Note that the demo unit supports FAT32 formatting only for USB memory. This limitation is the same on previous Roland switchers I've used. I was hoping to see exFAT support, as it's the common format now for many external USB memory sticks and drives. For now, you'll need to make sure your USB memory input is formatted as FAT32 in order for the V-160HD to recognize it.

5. Audio Inputs

The V-160HD switcher has a variety of audio inputs you can mix into your live feeds: two balanced XLR inputs, two RCA (left/right) inputs, a USB-C audio input, audio channels available in the SDI or HDMI inputs, and a new Bluetooth input (see Figure 2). The Bluetooth input enables you to connect output from just about any Bluetooth device to the V-160HD switcher. For example, you can play music from your smartphone to add as a background track to the rest of your feed.

One feature I miss from the V-60HD is the additional XLR inputs. While I usually use a professional audio mixer to pipe audio into any switcher, the ability to record independent wet/dry audio across four XLR inputs is a great asset for making postproduction tweaks to live event recordings. Like other Roland switchers, the V-160HD can output up to eight independent audio channels in the embedded audio output over HDMI/SDI, but you'll only have two XLR input channels available.

6. Output Control

The V-160HD has several views (or output busses) to push to specific output ports for a wide range of applications:

- **Program:** The final PGM/A composited feed from the switcher, including any or all active PinP/key or DSK effects. You can control it if you want PinP/key or DSK layers enabled or disabled in the program output within the menu settings.
- **Preview:** The PST(B) feed for the “on deck” input
- **Aux:** The selected input for the AUX mode on the console
- **Sub Program:** Identical to the Program view, but with separate control for PinP/key or DSK layers
- **16 Input View:** Outputs a 4x4 grid of all SDI and HDMI inputs
- **16 Still View:** Outputs a 4x4 grid of all still images stored on the V-160HD
- **Multiview:** Outputs the PVW and PGM views side by side, along with eight cross-point input views

You can assign these views to any of the three SDI outputs, any of the three HDMI outputs, or the USB-C output. For my live event workflows, the expanded number of output ports over the Roland V-60HD would reduce my reliance on an SDI distribution hub to split a program feed to a recorder and two dedicated hardware encoders. The HDMI outputs can all be used, for example, for the operator to view all inputs, multiview, and program (or preview) across several displays simultaneously or to push to the projector(s) in use at the event venue.

I also love the embedded display on the V-160HD console in the upper right in Figure 1. You can check your inputs very quickly without the need for an external HDMI or SDI display.

7. Macros and Sequencer

While I wasn't able to fully test these brand-new capabilities, the V-160HD switcher can record and store up to 100 macros, which can then be used with the new sequencer controls as well. (A macro is a list of ordered procedures you want to happen for any given effect in your program, just like you may have used macro scripts in desktop video/image-editing programs.) If you like tight control and pre-planned visual effects for your live events, you'll love these new features.

The demo unit I tested had several demo scripts pre-programmed, animating one or

review

The Roland V-160HD Hybrid Event Video Switcher

more PinP (or key) compositors from one area of the screen to the other. The new sequencer feature set allows you to arrange multiple memory and macro presets and manually step through them in an orderly fashion during your live event. It also lets you automate the procession between macros. In addition, you can load and save macros and sequences from USB memory.

As a quick aside, the demo unit uses a proprietary binary format for the saved macros or sequences when saved to USB memory. I was hoping to see a text-based format for these settings so that any text editor could be used to make new macros or sequences or modify them.

8. Enhanced Control With External Apps

Roland continues its tradition of enhanced control and display of its switcher settings with the Remote Control Software (RCS) desktop applications for Windows/macOS, as well as a new iPad app. I've used the RCS desktop apps with older Roland switchers, and I expect to see the same fine-tuned controls in the V-160HD version of the RCS app. The desktop RCS apps communicate with the V-160HD over the Ethernet connection, and the iPad app (Figure 3) requires a Wi-Fi connection to the same IP subnet that the V-160HD is connected to with the Ethernet port.

9. Console Button Highlight Colors

I've already written about my excitement over the extended console operations, which don't require you to go menu hunting in the V-160HD, as you need to do in other switchers I've used. I'll mention another feature that helps any switcher operator quickly learn and understand what's happening and when: button highlight colors. The MODE button below



Figure 3. The Roland V-160HD iPad remote control utility

the PinP/key 3 controls enables the operator to quickly switch between the aux, memory, and macro features and have specific coloring.

The PinP/key and DSK controls have unique highlight colors across the console as well. You should never be in the position of not knowing which setting is active on the console.

10. Physical Dimensions

I'm amazed that all of these features are packed into a unit that's just barely bigger than the V-60HD switcher. The V-160HD would be a tight fit in the Pelican Air 1485 case that I use for the V-60HD, but I'd probably go with a larger case to accommodate the power supply as well.

There are so many more features in the V-160HD, such as PTZ camera control, remote tally support, user-defined buttons, and the 40-channel digital audio mixer, just to name a few. When you're considering a new hardware video switcher, as hybrid events become the norm in this COVID era and beyond, I highly encourage you to check out the V-160HD.

Robert Reinhardt (robert@videorx.com) is the founder of VideoRx and is internationally regarded as an expert on multimedia application development and online video, particularly in HTML5, iOS, Flash, AVC/H.264, and HEVC/H.265. Comments? Email us at letters@streamingmedia.com, or check the masthead for other ways to contact us.

**You've got questions,
we've got answers.**



Join our Facebook group today!

bit.ly/AskSMExperts





excellence in
transcoding

ADVENT OF A NEW ERA: HERO 6000

Low latency | 8K UHD

hero

Multiscreen

Contribution/
Point-to-point

Management

Cloud

4K

Live

Contribution

File

Intelligence

Event

info@mediaexcel.com
mediaexcel.com

Visit us at
NAB SHOW 2021
October 10-13
Booth #C3331