The Vital Role of Unified Distribution To Ongoing Success in the TV Business

Part I of 3

End-to-End HTTP Infrastructure Is Key To Expanding Revenue and Cutting Costs

Technology Advances Enable Broadcasters and MVPDs To Maximize ROI in the Multiscreen Services Environment

Introduction

As disruption continues to roil the premium video marketplace, content producers and distributors alike are under mounting pressure to find a uniform approach to reaching all screens that can minimize impediments to cost containment, strategic flexibility and revenue growth.

Providers of every stripe have been contending with audience fragmentation by investing heavily in IP-optimized transcoding, streaming and other infrastructure elements to ensure their content is accessible on connected devices even as they continue to rely on the modes of processing and distribution that undergird legacy pay TV services. There may have been some justification at the outset of the multiscreen era to tailor infrastructure support for IP video streaming as an adjunct to the core pay TV infrastructure, forcing video providers to operate two parallel platforms or to deal with suboptimal "bolt-on" platforms for OTT. But today the dual-silo approach stands in the way of realizing the full ROI potential in both domains.

Spending on components performing duplicative processes along with continued reliance on proprietary hardware to support traditional pay TV functionalities drives costs ever higher while limiting responsiveness to consumer demand and new revenue opportunities. Consequently, there's nothing more important to the success of all players than implementation of a Unified Distribution architecture that eliminates unnecessary costs and opens the door to seamless execution of advanced advertising and content strategies across all end devices.

So far, notwithstanding the appeal of the concept, any wide-scale move in this direction has been blocked by business realities, which is to say, the need to avoid replacing legacy infrastructure elements, especially set-top boxes and the associated local pay TV transport mechanisms, without assurance of revenue that will justify the expense. Simply collapsing everything end to end onto the IP infrastructure is not a viable option as long as non-IP set-tops dominate the subscriber landscape.

Yet it's essential that broadcasters and distributors be empowered to bring the advantages of IP technology into the legacy domain, including monetization through dynamic advertising; personalization and localization of content; virtualization of processes through use of software running on commodity servers; efficient aggregation and analysis of data, and much else. The challenge, then, is to create a Unified Distribution architecture that can be cost effectively deployed to eradicate as much of the legacy infrastructure as possible while extending the benefits of IP technology over QAM and IPTV links to set-top boxes.

Imagine Communications has met this challenge with development of a distribution architecture that takes an innovative approach to utilizing the mechanisms that have made HTTP (Hypertext Transfer Protocol)-based streaming technology a viable means of delivering TV-quality video to connected devices. This new Unified Distribution architecture, now in the early stages of deployment by a growing number of major providers, is built on a simple premise:

It is now possible to rely on HTTP with fragmented encoding of content as the foundation for delivering revenue-generating TV content in HD and even UHD formats to TV sets as well as personal devices, eliminating complications that result from layering the HTTP infrastructure onto legacy transport stream-centric infrastructures.

By making fragmentation of content central to the encoding process at core distribution facilities, broadcasters and MVPDs (multichannel video programming distributors) can rely on one mode of distribution for every piece of content, eliminating the need at various secondary and tertiary staging points for the transcoding and packaging processes currently used to format video for HTTP delivery to connected devices. This, of course, requires a break with the way providers utilize terrestrial backbones, where premium content is encoded and encapsulated in IP packets at core facilities for multicasting in UDP (User Datagram Protocol)-based transport streams to regional and local distribution facilities.

To make this possible while ensuring support for TV distribution in the legacy domain,

Table of Contents

Introduction

The Growing Reliance of the Pay TV Business on OTT Strategies

- Trend Lines in Consumer Viewing Behavior
- Broadcaster Responses
- MVPD Responses

The Emergence of a TV-Caliber ABR Infrastructure

- The Technical Foundation for Robust Streaming
- Advances Enabling Transformations in Monetization, Services and Operations
- The Expanding Role of CDNs

The Mounting Case Against Operating in a Two-Silo Environment

- The Need to Support UHD in the Legacy and IP Domains
- The Need to Execute New Monetization and UX Capabilities in Legacy Pay TV

Realizing the Potential of a Unified Distribution Architecture

- Technical Requirements
- The Cost-Saving Benefits
- The Time-Shifting Advantage
- Enabling DevOps Agility
- The Monetization and UX Benefits of Leveraging Just-in-Time Packaging

Conclusion

Imagine has created the first software-based platform designed to convert HTTP-delivered fragments into the continuous UDP-encapsulated MPEG-2 Transport Streams that are required to accommodate reception on legacy set-tops. Known as the Selenio Video Delivery Edge (VDE), this advanced component runs on the same type of COTS (commercial off-the-shelf) hardware that's commonly used to support generic compute and networking functionality, such as today's HTTP content delivery networks (CDNs).

Insofar as there is nearly universal reliance on CDN facilities positioned at network edges to support caching of popular on-demand content, there's already a vast installed global infrastructure available for implementation of the capabilities enabled by VDE software. Critically, these capabilities not only include conversion of HTTP fragments to legacy transport streams but also implementation of the manifest manipulation capabilities intrinsic to HTTP-based adaptive bitrate (ABR) streaming that enable dynamic changes in both on-demand and live IP content, such as on-the-fly insertion of ad spots, replacement of regionally blacked-out programming and personalization of programming features.

With VDE at the network edge these processes can be applied to customize content as it is converted for local distribution over legacy pay or free-to-air TV links. In this way, the set-top-connected or OTA TV sets are able to deliver a viewing and advertising experience that is more closely aligned with what viewers experience when accessing content on IP-connected devices. Consolidation onto the HTTP infrastructure also adds efficiencies to time-shifting in the legacy domain through use of content fragmentation in cloud DVR, catch-up and trick-play applications.

From a monetization standpoint, the incorporation of the TV set into the dynamic advertising paradigm creates opportunities for both broadcasters and service providers to enable converged multiscreen ad campaigns. Such convergence, enabling localized and personalized ad targeting across all screens that has long been sought by ad agencies and their clients, has the potential to drive higher CPMs and, with them, higher total revenues than can be attained under current conditions, where linear advertising over legacy channels is treated as a business apart from dynamic advertising in the digital realm.

These benefits can be achieved by broadcasters and network service providers working independently of each other in the sale, respectively, of national and local avails. But there's also an opportunity to grow revenue even farther in both segments through collaboration in the use of ABR technology with HTTP-to-UDP gateways as embodied in the VDE platform. Service providers, by making such CDN capabilities available to broadcasters on a wholesale basis, could generate new revenue for themselves while expanding the higher-priced targeted ad reach for broadcasters' national ad campaigns.

In the three-part series that follows we will explore in depth these and other monetization opportunities along with the cost savings and other benefits to be realized by broadcasters and service providers with implementation of this Unified Distribution architecture. In Part 1 we explore market trends that underscore the growing need for ecosystem adoption of this architecture, the emergence of HTTP as a TV-caliber mode of distribution and how the proposed Unified Distribution architecture answers these needs with respect to technological underpinnings and key benefits.

In Part 2 we will focus on specific needs of broadcasters and how they can exploit the Unified Distribution architecture to maximum advantage. Part 3 will look at what a shift to the Unified Distribution architecture as enabled by Imagine's end-to-end solutions means to network service providers.

The Growing Reliance of the Pay TV Business on OTT Strategies

As things stand today, broadcasters' and distributors' efforts to ensure their offerings reach the broadest possible audience by bringing IP-connected devices into the viewing matrix have grown increasingly costly without a corresponding increase in revenue. Judging from trends in consumer behavior and the emergence of new video formats, these costs will continue to mount so long as providers continue to rely on separate infrastructures to deliver content in the legacy and IP-connected domains.

The pay TV market is undergoing major disruption as providers move beyond the original TV Everywhere (TVE) paradigm where IP-streamed offerings of live and on-demand premium video were targeted to pay TV subscribers as a way to enhance the appeal of the legacy services. Now growing numbers of broadcasters and distributors see an urgent need to use Internet streaming technology to engage consumers who aren't drawn to traditional premium packages – the so-called cord nevers and a growing legion of cord cutters.

Trends Lines in Consumer Viewing Behavior

In the U.S., according to Parks Associates, as of Q3 2016 63 percent of broadband households were subscribing to at least one OTT video service, and 31 percent were subscribing to two or more (Figure 1).[1] Parks reported that all of the top ten OTT services increased their subscriber bases in 2016. While, traditionally, such services were positioned as providers of subscription video-on-demand (SVOD) content, today streaming of linear TV channels, including live sports, has become a key component of most business models.

Measuring actual network usage in 2016, network traffic tracker Sandvine said the average North American household had at least seven active devices in use daily, with video streaming accounting for 65 percent of usage.[2] By 2019 the average number of connected media-enabled devices per North American household will climb to ten, according to IHS projections.[3]

The TV, of course, is now among the devices playing a big role in the growth of online video consumption, thanks to the proliferation of smart TVs and a wide range of IP streaming media players. In the U.S., 34 percent of all broadband households own a smart TV and 26 percent own streaming media players supplied by Roku, Apple TV, Amazon Fire, Google Chromecast and others, according to Parks Associates.[4]

By the end of 2017, streaming media player penetration of U.S. broadband households will reach 40 percent, according to NPD Group, marking a 150 percent increase since the beginning of 2014.[5] Overall, counting media

Figure 1

OTT Video Viewing Trends

U.S. Broadband HHs*				Average NA HH	
OTT Video Subscriptions		Smart TV Ownership	TV Media Player Ownership	Media-Enabled Devices in Daily Use	
1 or More	2 or More			2016	2019
63%	31%	34%	26%	7**	10***

** Source: Parks Associates ** Source: Sandvine *** Source: IHS*

players, smart TVs, game consoles and Blu-ray players, 211 million devices will be in place to enable OTT viewing on TV sets in U.S. households by YE 2017, NPD predicts.

These trends are replicated worldwide. U.K. analyst Ovum in a recent report predicted online video subscriptions, which topped 100 million worldwide at the end of 2015, will increase to 177 million by 2019.[6] In another report, Digital TV Research projected that by 2020 online subscription revenues will reach $21.6 billion, three times the 2014 total, with penetration exceeding 33 percent in ten countries.[7]

Broadcaster Responses

In response to these trends, broadcasters are licensing content to ever more OTT subscription services, including not only competitors to traditional MVPDs but also a growing number of MVPDs who are moving beyond the TV Everywhere model to offer standalone OTT services. At the same time, many if not most broadcasters have implemented direct-to-consumer (DTC) strategies that leverage the versatility of IP-based production, post-production and distribution processes to deliver unique blends of their branded content to consumers, often without requiring them to be pay TV subscribers.

DTC strategies have become especially important to building an international presence for many U.S.-based TV and cable networks. Discovery Communications, for example, reaches some three billion subscribers in more than 220 countries and territories through Discovery Channel and multiple other network brands. With its launch of DPlay in 2015, Discovery entered the rDTC subscription market, starting with an $8 monthly service combining live and on-demand content in Norway, Denmark, Italy and Sweden. And the company has become a major sports network for Europe and Asia-Pacific regions through its recent acquisition of Eurosport.

Another ambitious undertaking is NBCUniversal's GlobalNetworks, which manages a variety of acquired properties in Europe and Asia as well as internationally formatted iterations of branded NBC cable networks. Fox, too, has an aggressive global OTT initiative in place through Fox International, which partners with leading regional platforms to make programming available online, sometimes licensing content to third-party aggregators, sometimes offering a DTC package on its own. Disney/ABC, Viacom, AMC Networks, Scripps Networks and Time Warner's HBO and Turner Broadcasting System are other examples of U.S. media companies with notable international expansion strategies in play.

MVPD Responses

Meanwhile, the ranks of MVPDs offering standalone OTT services now include all the Tier 1 providers in the U.S. and many more abroad. Dish Network's Sling TV, the first MVPD offering to launch in the U.S. is now ranked sixth in Parks' top ten behind Netflix, Amazon Prime, Hulu, MLB.TV and WWE Network and ahead of HBO Now, Crunchyroll, Showtime and CBS All Access.

Like Sling TV, Verizon's more recently launched Go90 OTT subscription service and AT&T's DirecTV Now are available nationwide. DirecTV Now, which launched in late 2016 with packages that include a 100-channel lineup, currently priced at $60 monthly following an introductory offering at $35 per month, attracted over 200,000 subscribers in its first month of operation.

The number three telco, CenturyLink, having abandoned its legacy Prism IPTV service, has begun rolling out an OTT TV service as well. Meanwhile, taking a somewhat less aggressive approach, cable leaders Comcast and Charter Communications are limiting access to their offerings, Xfinity Stream and Spectrum TV Stream, respectively, to their own broadband subscribers.

As in the U.S., MVPDs elsewhere are reacting by going beyond TVE enhancements to legacy pay TV services with their own standalone OTT services. Some are in the SVOD mode, such as Netherlands-based Altice Group, which in late 2015 launched its €9.99 Zive SVOD service in France followed by expansion in 2016 to six other countries. The MVPD says the service, now offering 15,000 HD and 400 UHD titles, will be rolled out by Altice USA, the third-ranked stateside MSO, at an unannounced date.[8]

Other European operators are offering both live and VOD programming with their new OTT services. Major players in this space include Sky with NOW TV in the U.K. and other versions offered through Sky Deutschland and Sky Italia, Belgium cable operator VOO with Be tv Go, Orange in France and Vodafone in Germany and Spain.

Increasingly, MVPDs are consolidating video processing across legacy and OTT outlets as they upgrade facilities to support more advanced multiscreen services. For example, Vodafone Deutschland is offering its new UHD-enabled GigaTV service both as a traditional pay TV service and as an OTT option, which at €9.99 per month includes 120 TV channels, 55 catch-up portals and pay-per-view access to over 3,000 VOD movie titles.[9]

The Emergence of a TV-Caliber HTTP Infrastructure

All these strategies underscore the growing confidence of MVPDs and broadcasters in the viability of IP-based streaming as a means of delivering services that measure up to the performance and quality standards set for legacy pay TV. ABR streaming technology, originally used primarily for delivering content at relatively low resolutions to personal devices over the public Internet, has emerged as a viable means of broadcasting live sports, news and other TV programming in 1080p HD or even 4K resolution at quality levels required for viewing on the largest displays, with QoS and QoE that are equal to or better than traditional delivery systems.

It's especially noteworthy that the biggest challenge for ABR streaming of TV-caliber live programming, namely coverage of fast-action sports, has been met with resounding success. OTT delivery of professional and college-level football, baseball, basketball, soccer and other sports to mass audiences in HD is now a routine feature of broadcast operations worldwide with 4K transmissions entering the picture as well.

Perhaps nothing better illustrates what can be accomplished with IP technology utilizing ABR streaming than NBC's live OTT delivery of all events from the 2016 Summer Olympics (Figure 2). One hundred million unique users spent an aggregate time of over 45 million hours watching live coverage on connected devices during the 17-day schedule, according to NBC.[10] IP-based production and post-production

Figure 2

Digital Viewing Summer Olympics

	Total Hours of Live Streammed Programming	Unique Digital Viewers	Viewers of Live Streaming under Age 35
2016	4,500	100 Million	50%+
2012	3,500	77.5 Million	

Source: NBC Sports

processes facilitated storage and preparation of live-captured content for time-shifted viewing with enhanced features offering statistics, athletic profiles and much else relevant to each event and to the games as a whole.

The Technical Foundation for Robust Streaming

Many technological advances have come into play to enable such capabilities over the open Internet, starting with enhancements in the functionalities intrinsic to HTTP. ABR was designed for use with HTTP servers as a way to utilize the packet-loss recovery mechanisms of TCP (Transmission Control Protocol) in support of maintaining continuity of the A/V stream by adjusting bitrates to accommodate fluctuations in available bandwidth during a user session.

Fundamental to today's ABR streaming capabilities are software-based platforms that utilize high-density COTS compute to transcode content from original sources at bitrates suited to reaching multiple types of devices at various resolutions and frame rates across both QoS and best-effort networks. These platforms are designed to perform all this processing, which might include use of more than one codec for each bitrate profile, in real time for live feeds and at accelerated speeds for content stored for on-demand distribution.

State-of-the-art premium video transcoders execute many other tasks as well, such as de-interlacing of NTSC files to progressive mode; adding IDR (instantaneous decoder refresh) frames to enable SCTE 35-based ad insertion; performing GOP (group of pictures) alignment to facilitate smooth output in the ABR streaming process; making automatic loudness adjustments, and processing and synchronizing ancillary feeds such as closed captioning, picture-in-picture displays and foreign language subtitles.

ABR streaming packagers, software components that can be co-located with transcoders or positioned remotely, utilize a communications framework that allows HTTP servers to send to each client device a "manifest" file of information pertinent to the content the user is accessing, starting most fundamentally with a list of the various bitrates at which the content has been transcoded for distribution. The packagers direct processors to fragment each transcoded version of the content into "chunks" of a few seconds duration, the length of which depends on which of several ABR formats is used with a given streaming session.

While the multiplicity of ABR formats, most prominently including HTTP Live Streaming (HLS), Adobe HTTP Dynamic Streaming (HDS) and Microsoft Smooth Streaming, complicated use of the technology through its formative years, HLS has emerged as the dominant mode with close to universal support on recent vintage devices while HDS and Smooth have been made interoperable through the standardized format known as MPEG-DASH. Whichever format is used, throughout the streaming session the client device calculates its available network capacity and requests from the video origin the video resolution, called a profile, that's best suited for the available bandwidth.

Moreover, the client can also take into account video format or CPU capacity when requesting profiles. Thus, for example, even if enough bandwidth is available to send a chunk at 60 frames per second with 1080p resolution, a device equipped to render resolution no greater than 480p at 30 fps will ask for fragments transcoded at the appropriate resolution, thereby avoiding over use of bandwidth and buffering delays that result from over saturation of the CPU.

Advances Enabling Transformations in Monetization, Services and Operations

Beyond the robust performance capabilities of multi-bitrate ABR streaming, several other factors are contributing to the utility of ABR in pay TV operations. At the top of the list are the functionalities introduced through what is known as "manifest manipulation."

Over time the information communicated in manifest files has been expanded to include data telling clients where to find and exactly when to pull content such as advertising, special features or alternative programming that the distributor wants any given user to receive at any point during the session. These adjustments through manifest manipulation are implemented in tandem with use of HTTP applications servers, ad decision systems and other elements, all of which are managed by orchestration platforms that employ data analysis and policy servers to call for assets relevant to a particular transaction based on the location and profile of the end user.

Another function crucial to using ABR in premium video distributions is automated support for time shifting. Applications servers linked to the packaging platforms enable implementation of multiple time-shifting modes under the control of end users, including trick-play functions, catch-up viewing in limited time windows and cloud-based DVR options utilizing long-term storage facilities.

In addition, it's now much easier than it once was to apply whatever mode of content protection is suited to the type of content and device associated with a given session. In today's premium service operations the execution of encryption, device authentication, user certification and DRM management is performed under the guidance of multi-platform systems that support on-the-fly association of policies and DRMs specific to user and device profiles for each live and on-demand session. Some of these systems also incorporate support for applying traditional conditional access (CA) protection to content destined for viewing over legacy Pay TV streams.

Rounding out the advances that have made ABR a reliable mode of delivering TV-caliber performance to connected devices, including large-screen displays, are quality assurance (QA) platforms that orchestrate both the QoS (quality of service) functions that enable fast, often proactive measures against network-based threats to expected service performance and the QoE (quality of experience) functions that track and analyze data from every viewing session in support of various business models. Utilizing advanced analytics engines that tap into data flows from network elements and user devices, these QA platforms provide content producers and distributors the means to ensure the caliber of performance they're looking for is adhered to in the production, post-production and distribution processes as well as in the application of dynamic ad placement and other functionalities tied to manifest manipulation.

The Expanding Role of CDNs

A major factor underlying the viability and flexibility of today's IP video infrastructure is the fact that all the foregoing functions as well as all the functions performed by IP-optimized production and post-production platforms are now executed by software systems running on COTS facilities. As a result, all the advances supporting shared use of resources through standards-based virtualization technologies can be leveraged by every element in the end-to-end HTTP infrastructure. This means that broadcasters and distributors will be able to exploit whatever cost benefits emerge with ongoing advances in the virtualization domain to enable adjustments to ever-changing market needs in the most cost-efficient ways possible.

Anchoring the implementation of all these capabilities is the fact that all providers in the global premium video ecosystem rely on the processing capabilities of public and private CDN facilities positioned at the edges of metropolitan areas and often at points closer to end users. Originally designed to provide

caching support for web pages, COTS-based CDN facilities in growing numbers have been expanded and enhanced to support the more advanced capabilities of HTTP technology as discussed above for both on-demand and live feeds via the caching of video fragments.

Now these facilities can be further enhanced to make it possible to integrate legacy pay TV distribution with the advanced IP streaming infrastructure in support of the Unified Distribution architecture described in the following section. Critically, these enhancements can be readily implemented in software utilizing COTS appliances, thereby avoiding the need for investments in purpose-built hardware.

Thus, one of the major advantages of the Unified Distribution architecture is its ability re-use existing protocols on public and private CDNs for video delivery. The resiliency and scale-induced cost effectiveness of those networks translates directly to both a cost and QoE benefit for content providers. And, when required, the addition of Imagine Communications' VDE device can convert the fragmented video stream back to a legacy transport stream while conveniently utilizing the same COTS hardware used as part of the CDN.

The Mounting Case Against Operating in a Two-Silo Environment

Notwithstanding the emergence of HTTP as a technology supporting TV-caliber distribution, the premium video industry continues to rely on the legacy mode of end-to-end content distribution as the core architecture while investing increasing amounts in the overlaid HTTP infrastructure. This results in duplicative spending on transcoding platforms and other elements as well as inflated operational costs and complexity as providers attempt to keep up with evolving requirements in the IP and legacy pay TV domains.

The Need to Support UHD in the Legacy and IP Domains

Ultra HD (UHD) has been slow to get off the ground, much as was the case in the early going with HD. But whatever the time frame for introducing UHD services on a mass scale turns out to be, it makes no sense to have to invest in two infrastructures to support implementation of those services.

The prospects for market adaptation to UHD have been greatly improved with the inclusion of HDR (High Dynamic Range) technology as a true differentiator in user experience. As a result, supporting UHD has come to mean investment in not only the encoding and CPE infrastructure required for transmitting programming at 4K resolution but also the processing mechanisms that are needed to add HDR enhancements to 4K content and, likely, HD content as well. And, of course, the same types of dual-silo cost barriers loom when it comes to planning for future investments in other formats, including virtual reality (VR) and, eventually, 8K UHD.

Moreover, the onset of UHD brings another major infrastructure adjustment into play in conjunction with meeting new security requirements for licensing the highest-value content, including HD as well as UHD programming. With the need to thwart video piracy emerging as a major priority across the motion picture and TV programming sectors, the Enhanced Content Protection (ECP) specifications issued in 2015 by the motion picture studios' tech consortium MovieLabs are now widely viewed as a template for protecting all types of premium content.

This is going to require many adjustments in content protection infrastructure. New stipulations include not only the well-publicized need for mechanisms supporting per-session applications of forensic watermarking codes but also stringent requirements for end-to-end link security bearing on the design of core hardware processors, where no CPU will be allowed to perform security-related functions, and use of robust DRM technology that is superior to native DRMs employed with many types of devices.

The Need to Execute New Monetization and UX Capabilities in Legacy Pay TV

Beyond duplicative investments in new technology, there's another, potentially much greater cost incurred with continued reliance on the legacy architecture, which is the difficulty of refurbishing traditional pay TV service with the user experience (UX) and monetization advances intrinsic to the IP streaming architecture. As the UX offered with OTT services delivered to IP-connected TVs becomes ever more compelling, the competitive challenge to legacy services intensifies. To even begin to replicate the IP-platform capabilities in the legacy domain requires ever more investment in infrastructure elements specifically designed to work within the limitations of legacy mechanisms.

Consequently, as advertisers become more enthusiastic about the dynamic targeted advertising opportunity offered in the OTT space, there's no readily available way for broadcasters and distributors to capture higher CPMs for locally and demographically targeted ad placements in legacy pay TV programming. As a result, even though the industry has made some strides with regard to dynamic ad placement in VOD content, it has been stalled in building out the business and supporting infrastructure framework essential to realizing the full potential of addressable advertising in linear programming.

Realizing the Potential of a Unified Distribution Architecture

Technical Requirements

As the imperatives to eliminate the unnecessary costs of duplicate distribution architecture and to bring the benefits of IP technology into the pay TV arena intensify, the search for a way to address these issues has become mission critical. In other words, it's becoming mandatory that broadcasters and distributors define and implement a Unified Distribution architecture that eliminates duplication, enables use of cloud technology to consolidate and simplify operations with DevOps flexibility, preserves the legacy pay TV access infrastructure and extends the benefits of IP-based UX and monetization into that domain.

Clearly, as evidenced in the foregoing discussion, the streaming infrastructure has been proven sufficiently robust to play the role as the core architecture. But even as ever more distributors come to this realization, the drawback to wide-scale adjustment to this new reality is a perception that it can't be done without incurring the extraordinary costs of replacing legacy STBs and stranding investment in the supporting access infrastructure.

Consequently, three major changes in how things are currently done must be accomplished to achieve the goals of a new Unified Distribution architecture. Notably:

- The current mode of distributing premium video over terrestrial fiber backbones from core broadcaster and MVPD facilities, which entails encapsulation of traditionally encoded programming in IP packets for multicasting in continuous UDP streams to secondary points of distribution, is eliminated in favor of multi-resolution transcoding and fragmenting of content at the points of origin in support of HTTP-based delivery over those backbones with farther packaging for local ABR streaming performed at the secondary distribution points;
- All processing of the fragmented IP content at those secondary distribution points in

support of live and on-demand multiscreen service models, including dynamic advertising, personalization of user experience, adherence to local blackout policies, support for time shifting and UX feature enhancements, is performed in software running on COTS appliances;

- To accommodate delivery of the locally processed IP content to legacy set-tops and to TVs relying on OTA reception an HTTP-to-UDP gateway is employed to convert the ABR fragments to continuous MPEG-2 Transport Streams for delivery over UDP. All mechanisms essential to making this Unified Distribution architecture possible are now available in the marketplace, including the HTTP-to-UDP gateway. This innovation was introduced to the market in 2016 as part of Imagine Communications' Selenio VDE platform, which now is in the early phases of deployment in several Tier 1 MVPDs' CDN facilities and in the contribution platforms of leading broadcasters.

In Parts 2 and 3 we will explore the ways in which the capabilities of VDE and other Imagine components of the Unified Distribution architecture can be utilized by broadcasters and distributors to address their specific needs. Here we conclude Part 1 with an overview of the benefits which this architecture brings to players on all sides in the premium video marketplace.

The Cost-Saving Benefits

Fragmentation and transcoding of mezzanine-encoded content mapped to all the resolutions suited to everything from handsets to 4K TV sets at core points of distribution not only unifies downstream processing onto a single platform; it eliminates many of the transcoding steps that the current architecture requires in both the legacy and IP domains. For example, video delivered by an MVPD to connected devices goes through four transcoding processes, starting with the point of contribution, next at the point of MVPD ingestion and then in two different processing environments, one for distribution to STBs and the other for ABR streaming to connected devices.

This configuration multiplies costs as each point in the chain is adjusted to accommodate new requirements, and it introduces delays and quality degradations at each step. With implementation of the Unified Distribution architecture, once the transcoding and fragmentation are performed on the mezzanine files, the need for farther processing is greatly reduced.

Rather than requiring multiple transcoding stations and separate infrastructures for processing video, the new architecture consolidates processing for IP and legacy pay TV operations with use of just-in-time packaging at the edge distribution points to execute manifest manipulation for DAI (dynamic ad insertion), blacked-out content replacement, personalization of features and other locally oriented functions. With the addition of the software-based HTTP-to-UDP gateway the locally processed programming can be delivered to legacy STBs.

These advantages apply whether the content is delivered for traditional broadcast linear viewing or for time-shifted applications. When it comes to supporting live content distribution over the backbones, the multiple fragmented versions of the content can be delivered in traditional IP multicast mode in keeping with the low latency requirements of linear programming, including sports. Along with conserving bandwidth this offers the added benefit of eliminating the variations in latency that occur in the two-silo approach, where content delivered to consumers on their personal devices is typically out of synch with content delivered over the legacy feeds.

Another cost-saving infrastructure benefit has to do with the content recovery mechanisms intrinsic to ABR. Currently, if a UDP packet is dropped or damaged in transit, the content is rendered useless, which means providers must support alternative backup paths for each route on the backbones. HTTP, with its reliance on TCP, provides packet replacement redundancy in the stream, obviating the need for backup paths.

Consolidating video processing at MVPD facilities lowers costs of operations in many other ways as well. For example, with this consolidation comes the ability to implement content protection mechanisms, including the watermarking and advanced DRM processes required under new content protection specifications, through the new consolidated multi-platform security systems on offer from industry suppliers.

With the Unified Distribution architecture in place MVPDs can also greatly lower costs while adding flexibility in their migrations to all-IP pay TV service. At any point they can implement a cap-and-transition strategy that preserves the investment in legacy STBs while delivering the pay TV service directly from the streaming packagers to new subscribers and subscribers who upgrade to new tiers tied to the IP service paradigm.

It's also important to note that, with consolidation of processing onto a single IP platform at local distribution points, MVPDs can use ABR mechanisms to break free of encumbrances and costs imposed by the need to use traditional statistical multiplexing technology to maximize bandwidth efficiency on QAM channels. Currently, while many operators want to capitalize on the bandwidth-conserving capabilities of H.264 compression, which is now widely supported by set-tops deployed over the past few years, the greater complexities associated with H.264 compression pose problems for legacy stat muxes, which means they must be replaced with newer equipment or the transition to H.264 has to be delayed.

With reliance on ABR feeds into the HTTP-to-UDP gateway operators can avoid this problem by capitalizing on the ABR transcoding and fragmenting system to replicate the benefits of stat muxing. In this model, the client in the gateway asks the streamer to send fragments for each stream associated with a given QAM channel at the minimum bitrates required to achieve the required quality level, thereby achieving maximum bandwidth efficiency across the entire QAM channel.

Another cost component eliminated with implementation of the new architecture is splicing technology used in blacked-out content replacement and advertising. In the case of blackouts, the need to precisely splice replacement programming at local distribution points through manual operation of splicing equipment is a costly headache that goes away with the use of just-in-time packaging with manifest manipulation to perform the task.

Similarly, using a single HTTP-based ad insertion system tied to the software-managed manifest manipulation process eliminates the need for traditional ad splicers at points of content origin and secondary distribution points. For MVPDs, the shift to HTTP-based insertion has the added benefit of overcoming barriers imposed by legacy ad insertion technology on the use of H.264 compression, which requires upgrades to H.264-compatible splicers.

The Time-Shifting Advantage

Another byproduct of the Unified Distribution model is the role HTTP fragmentation can play in providing a more robust, high-availability approach to providing cloud DVR services. Pioneered by Imagine with its CloudDVR technology, reliance on storage of HTTP-fragmented content rather than moving continuous transport streams into storage has now become commonplace in cloud DVR platforms.

With this approach, operators are not vulnerable to incremental corruptions in continuous files that can render the stored content useless. With recording of fragments, in the event of corruption of a given fragment the cloud DVR system can easily replace the fragment and reconcile the interruption in the

stored sequence.

At the same time, the consolidated processing platform at the edge locations can utilize local CDN storage in combination with just-in-time packaging functionality to execute short-term time-shifting. By placing all linear streams into temporary edge storage, providers can support pause, resume and rewind trick plays on live programming along with place shifting that allows users to switch to another device without losing continuity. Storage capacity for longer retention of linear content for catch-up applications can be allocated in core as well as edge locations.

Enabling DevOps Agility

Reliance on software-defined uses of commodity datacenter hardware to manage and process video at all locations allows all players to benefit from the flexibility of a software-based DevOps environment. In this environment providers can introduce new user interfaces, on-board new devices, launch new applications, execute software upgrades and deliver other user benefits almost instantaneously whenever they wish across all screens.

Moreover, operating in the IP domain from a software and cloud-based infrastructure, development teams can directly tune into operational results and make whatever adjustments are necessary with minimum disruption to services. Overseers of market trials can access operations-generated data to learn in real time how innovations are performing and, in many cases, make adjustments to the technology parameters without having to end the trial and start all over again with lab testing and a new trial. They can judge the effectiveness of an offer and adjust terms or replace it with another from dashboards that allow them to monitor results and implement changes with the click of a mouse.

The Monetization and UX Benefits of Leveraging Just-in-Time Packaging

As previously discussed, the advances tied to use of manifest manipulation in the ABR stream packaging process have enabled a wide range of capabilities, including dynamic advertising suited to localization and personalization of ads, blacked-out program replacement and personalization of UIs and other aspects of the user experience. How these benefits can be exploited in the business models of broadcasters and distributors will be explained at greater depth in Parts 2 and 3.

The important point here is to underscore the fact that, with deployment of the HTTP-to-UDP gateway, providers can deliver into the legacy pay TV streams many of the applications enabled by the ABR platform for consumption on connected devices. Obviously, the processing performed on content destined for translation through the gateway must take into account the shared viewing environment of STB-connected TV sets. Zone-based local advertising will be the same across all devices, but distributors will need to avail themselves of settings on advanced platforms that adjust the range of just-in-time applications to suit the viewing situation.

Conclusion

The pay TV industry has reached a crossroads where cost and revenue impediments attending reliance on the current distribution architecture to deliver premium content to legacy STBs and connected devices are severely limiting the ROI potential in both arenas. The need to support separate video processing infrastructures within this architectural framework makes it very hard to meet spending requirements imposed by new developments, especially when it comes to enabling the transition to next-generation TV formats, including UHD in the near term and virtual reality and 8K UHD farther into the future.

This dual-silo architecture also prevents providers from bringing into the legacy pay TV service domain the advances in dynamic advertising, service personalization and localized content management that are intrinsic to IP-based operations. This is impeding pursuit of new monetization opportunities and creating a widening gap between the pay TV UX and what consumers can expect utilizing connected TVs and other devices to receive programming over the HTTP infrastructure.

Given the proven viability of HTTP-based streaming for premium video distribution, the industry has every reason to eliminate the two-silo approach in favor of a Unified Distribution architecture that relies on encoding and fragmenting video for distribution over the HTTP infrastructure as the primary video processing framework. But in so doing providers must be able to protect legacy infrastructure investments with technology that supports delivery of pay TV services to legacy STBs.

Moreover, the Unified Distribution architecture must be able to enhance those legacy services with the capabilities that can be derived from an HTTP infrastructure equipped with the advanced functionalities enabled by manifest manipulation and various supporting components. These enhancements include addressable advertising, personalization of services, local programming adjustments related to blackout rules and much else.

All the advances required to support the Unified Distribution architecture are now commercially available, including Imagine Communications' portfolio of advanced manifest manipulation and related dynamic advertising and content substitution solutions and a new HTTP-to-UDP gateway introduced by Imagine as part of its Selenio Video Delivery Edge platform. This innovation make it possible to extend the benefits of manifest manipulation into the legacy TV domain by converting the fragmented output of ABR packagers to continuous transport streams for delivery via STBs and OTA reception.

All components of the Unified Distribution architecture are designed to run as software-based platforms on COTS appliances at the cores and edges of the broadband ecosystem. As a result, providers can operate with DevOps flexibility in pursuit of new revenue opportunities across all screens while setting the stage for migration strategies that will lead eventually to retirement of the legacy STB-based infrastructure.

In Part 2 we will explore the many ways broadcasters can exploit the technological components of the Unified Distribution architecture to drive service innovation and monetization across all their evolving business models. Part 3 will provide a thorough look into what the Unified Distribution architecture means to MVPDs, not only for meeting their internal service goals but also for their ability to develop business models that can lead to new opportunities in their dealings with broadcasters and other premium video providers. <

Footnotes

1 Parks Associates, OTT Video Market Tracker, December 2016
2 Sandvine, Global Internet Phenomena Report, August 2016
3 Broadband News, Five Connected Media Devices per Home by 2019, September 2015
4 Parks Associates, press release, January 2015
5 NPD Group, Connected Home Forecast, January 2015
6 Rapid TV News, "OTT Streaming to Hit 100 MN Subs," April 2015
7 Digital TV Research, Global OTT TV & Video Forecast, June 2015
8 Fierce Cable, Altice Remains Silent on U.S. Launch of Zive, June 2016
9 Digital TV Europe, Vodafone Germany to Launch New Advanced TV Service, February 2017
10 NBC Sports, NBC's Rio Olympics Is the Most Successful Media Event in History, August 2016