Maximizing ROI for OTT

Addressing scaling challenges for live television streaming

A WHITE PAPER BY LTN GLOBAL

MAY 2020

Maximizing ROI for OTT

Table of contents

Abstract

Media companies in the business of streaming live television over the internet can get a return on their investments in programming in two ways:

1. Selling subscriptions to the linear programming
2. Selling advertising during gaps between content segments in the linear programming

In both cases, the revenue brought in over any given period of time is directly related to the number of people who subscribe to and view the programming during that time period.

To maximize profitability, media companies must focus on reaching as many eyeballs as possible in as efficient a manner as possible. However, many of these companies are finding over-the-top (OTT) delivery of live television programming and dynamic ad insertion to be much more complicated to execute at scale than they had anticipated. Given the level of investment they have made in the past several years, and the missed opportunity to monetize viewership across digital platforms over this time period, media companies have an urgent need to start showing a financial return.

This white paper focuses on the key barriers to scaling OTT delivery of live television programming and dynamic ad insertion — and specifically how the LTN Network and associated technologies remove those barriers to enable revenue optimization for media companies.

Introduction

Traditional linear TV has shown dramatic ratings declines in the past several years as baby boomers and Gen Xers (the TV generations) respectively have retired and entered their 50s; millennials (the computer generation) have cut the cord; and Gen Zers (the internet generation) have grown up without a cord. This has led media companies to invest massively in distributing their existing linear networks over the internet to the devices and platforms the vital millennial and Gen Z audiences are now using to watch programming. At the same time, advertisers have leveraged the much richer viewer data available on the internet to target advertising in those linear networks in much the same way websites have done for nearly two decades now.

Despite media companies' investments in OTT streaming, however, they still see a relatively small number of viewers for streaming programming compared to viewership for traditional TV delivered over cable, satellite, or over the air (OTA) broadcast. This is true for a variety of reasons, including the fact that TV was almost universally available in the early 1980s, when the millennial generation began, so there is still a very large base of TV viewers. OTT is growing rapidly today, of course, because it fits into the millennial and Gen Z expectation of content anywhere on any device. Nevertheless, because it started from zero just a few years ago, OTT viewership is still relatively small compared to traditional TV viewing.

Several challenges associated with the delivery of live video over the internet will impede further growth if they remain unresolved. These challenges are set out in detail in the following section. The final section covers how these issues can be resolved in a practical manner at the scale needed to ensure that the business becomes profitable at a material level in a reasonable timeframe.

Key challenges of live streaming over the internet

Reliable distribution of live television content through any medium is a fundamental requirement if a service is to generate any revenue at all. The **distribution** of live television over the internet presents significant challenges that are not typically issues for traditional TV distribution, where end-to-end distribution is under the control of a single operator (e.g., an MVPD). The key challenges with internet distribution are:

- **The spoiler** — Normally when a linear channel is originated, it takes a few seconds for video to appear on a TV screen in the home. This latency is caused mainly by the video compression and statistical multiplexing processes that occur before distribution via satellite, cable or OTA. However, with OTT delivery there are additional processing steps required by the streaming protocols and CDNs that add tens of seconds or more to the overall latency. Obviously, this is undesirable, especially for live sports. Viewers don't like being alerted that a goal has been scored several seconds before it happens on their screens.

- **The spinning wheel** — Unlike traditional distribution, where the end-to-end delivery is fully managed by a single operator, OTT is delivered via unmanaged networks (e.g., the public cloud and internet) and often involves passing the video and audio content through multiple entities. Monitoring the health of the end-to-end distribution is very difficult because it is not under the control of a single entity. The net result is a fundamentally less satisfactory viewer experience. Viewers quickly become frustrated if the screen freezes (even momentarily) during a live game.

- **The home movies effect** — The amount of processing applied to video and audio during distribution (compression/decompression) because of the much larger range of devices that need to be supported in OTT has the side effect of degrading video quality. This effect is not limited to the visual quality; other problems, such as lip sync error, occur because the audio and video are processed separately, and the processing introduces different delays for each component. If not carefully managed, the addition of delay leads to the audio being out of sync with the video. Viewers find the resulting lip-sync error to be highly objectionable.

In addition to **distribution** challenges, there are also challenges with inserting or replacing **advertising** in live content delivered over the internet. The main challenges are:

- **Bad timing** — By definition, targeted advertising has to be inserted at the edge, close to the viewer. Bad timing is a direct result of the insertion system being unable to accurately determine the exact frame upon which to begin inserting. Too often, though, the insertion occurs too early, cutting off the end of the programming, or too late, displaying the underlying content — a "peek through" — before the ad rolls. Both of these are bad experiences for the viewer, with negative results for advertisers, as well.

- **Leaving money on the table** — If a broadcast ad break is 2 minutes and 30 seconds long, the replacement break will need to be at least 2 minutes and 30 seconds long —and should really be exactly 2 minutes and 30 seconds long! With dynamic ad insertion, this can be a problem because if the ad decision system doesn't deliver enough ads to fill the space, there is no option but to put up a static slate to fill the gap. Some operators even leave a black screen with no audio. Regardless, it is obviously a poor experience for the viewer and is hardly likely to attract premium advertisers. This problem only gets worse as the number of simultaneous viewers increases; the ad decision system must work even harder, and fill rates decline rapidly as the number of viewers increases. When it comes to generating revenue, this problem can result in a lot of money being left on the table.

- **Déjà vu** — Linear scheduled TV advertising has become very sophisticated over the years. Advertisers pick their position in a break, they demand that they aren't next to the competition (or even exclude the competition from the same ad break) and they require that their ads are not shown too many times in a given period (frequency capping). All for good reasons. Dynamic ad insertion really hasn't caught up with this level of sophistication, and so the aforementioned conditions aren't currently being implemented for OTT. The net result of this approach (known as "scatter") is that the same ad is shown over and over, and competitors are placed next to one another. This hugely diminishes the value for advertisers and so consequently reduces the revenue potential.

The combined impact of the aforementioned challenges is a reduction in efficiency. The delivery of live television content via the internet presents significant costs, and media companies naturally need to generate a return on this investment at some point.

By viewing the internet as another **distribution** opportunity for their valuable content and carefully architecting the end-to-end architecture, media companies can eliminate the key challenges and deliver an optimized experience for viewers while making internet-based distribution a profitable enterprise. The rest of this white paper will discuss how a fully managed network — like the LTN Network — and advanced monetization tools – like the ones LTN provides - make this possible.

Ensuring efficient, reliable distribution

LTN's core technology and service model

were developed to ensure low-latency, yet highly reliable, delivery of live video over the internet.

The internet was never designed to carry live video. In fact, up until just a few years ago, most OTT was on-demand, and most people considered OTT to be a video-on-demand (VOD) service. The internet is actually quite well suited to this model. Users request video files in much the same way as they request web pages. The streaming protocols were even designed to make VOD streams look like a series of web pages being requested by a browser using the HTTP protocol. (The "H" in HLS, Apple's widely used streaming protocol, stands for "HTTP.") The view at the time was that linear TV was dead.

Once the media industry realized that linear was still needed — people still want to watch live sports and news, and they still sometimes want content curated for them the way traditional TV has for decades — streaming protocols were adapted to enable continuous streaming of linear programming rather than streaming a file with a start and end. This is why the LTN Network has been designed and developed to manage the inherent limitations of the internet for supporting reliable live video delivery. The LTN Network is optimized for ultra-low latency (i.e., live), broadcast quality, and security to ensure a great viewing experience while supporting effective monetization.

Latency

LTN's core technology and service model were developed to ensure low-latency, yet highly reliable, delivery of live video over the internet. The LTN Network does this by continuously measuring the performance of each "hop" on the internet, as well as by ensuring that errors are detected as soon as they occur, not after a "time out" at the other end. LTN can therefore (and does in a growing number of situations) deliver live content reliably — at least as well as current satellite and managed fiber with comparable latency — to the last-mile distributors (i.e., vMVPDs).

Unicast vs. Multicast

The internet is based on a unicast (point-to-point) model. This differs from traditional TV which is point-to-multipoint — everyone gets the same content. Unicast supports the delivery of video-on-demand content well because it is point-to-point — each viewer has control over when the content starts. In order to deliver live content to multiple viewers simultaneously, the Internet has to set up multiple point-to-point connections (one for each viewer in fact). The more viewers there are, the more bandwidth the content delivery networks (CDNs) need to serve. This approach just isn't scalable enough. The capacity to deliver the same content to millions of viewers for a few minutes or hours cannot be conjured up out of thin air instantly. It must be provisioned — and paid for — even when it isn't being used.

LTN's network is inherently multicast, so it is much more efficient bandwidth-wise — and therefore cost-wise, which also means it is scalable.

LTN's Network is inherently multicast, so it is much more efficient bandwidth-wise — and therefore cost-wise, which also means it is scalable. So scalable, in fact, that there is no fundamental reason the LTN Network cannot deliver content directly to the consumer in addition to delivering to MVPDs. It would obviously take a significant expansion of LTN's current network to realize this, but the amount of infrastructure required would be far less than that required using today's CDN methodology. LTN's approach would mean a fully managed, end-to-end network running over the internet delivering live content directly to consumers. No more spinning wheels. It would also deliver a more cost-efficient (less infrastructure), low-latency, highly reliable experience for viewers that would be able to scale much more readily than today's architectures.

Quality of experience

Working with a fully managed network, media companies can address the audiovisual quality issues that are caused by multiple processing steps by multiple parties in the end-to-end delivery chain. If content — uncompressed video and audio with rich metadata — is taken directly from the source, then video and audio could be processed just once for each individual viewer, depending upon the device being used at the time.

If a viewer's device can support 4k UHD and there is sufficient bandwidth into the home, then that is what they should receive. If not, then they should get the best-quality experience possible for the device they are using at that time. This also opens up the possibility of delivering customized content or graphics to each user. Imagine graphics that change depending on the device being used.

While this all may seem a bit far-fetched, it is possible to realize this vision in stages. For example, LTN currently delivers the station output from approximately 600 broadcast stations in the U.S. to vMVPDs across its network. Because LTN takes uncompressed video and audio from each station, it would be quite straightforward to compress according to ATSC 3.0 standards (i.e., H.265/HEVC and Dolby AC4) and deliver this — as well as the current H.264 — to vMVPDs. This would allow vMVPDs (or a traditional MVPD) to use their current delivery infrastructure to offer ATSC 3.0 before it is available over the air. High dynamic range (HDR) and 4k UHD can therefore be rolled out now by broadcasters without waiting for full ATSC 3.0 deployment.

Implementing sophisticated ad insertion/replacement/enhancement

High quality, low latency and highly reliable distribution is only part of the equation. Without being able to monetize content though the use of targeted advertising, media companies cannot make OTT streaming profitable. They can't generate enough subscription revenue to support the kind of high-quality content viewers expect for the price they pay. Few media companies or MVPDs currently do a good job with ad insertion. One or more of the issues described earlier are readily apparent even in the most valuable live streaming sports today.

LTN has developed technology that specifically addresses the key challenges of monetization, thereby enabling vMVPDs and media companies to unlock the value of their content.

Timing is everything

Seamless insertion (or replacement) of advertising in an OTT environment has just a few basic requirements:

a, The precise timing of the start of the replacement must be known in advance by the system performing the insertion

b The ad creative(s) that will be displayed to each viewer must be transcoded into the same video/audio quality as the original content and be ready to play (i.e., delivered to a CDN)

If the ad insertion uses manifest manipulation in the cloud (server-side ad insertion or SSAI), which is almost always the case today, there is one more very important requirement:

c The streaming protocol's "chunks" must be perfectly aligned with the frame boundaries between programming and advertising in the original live stream. In other words, the adverts that make up a commercial break must each be packaged into an exact integer number of chunks. This is very important as manifest manipulation can only replace whole chunks. If a single chunk has part of a program and part of an ad, the replacement will not be seamless.

Without being able to monetize content though the use of targeted advertising, media companies cannot make OTT streaming profitable.

In order to achieve **a** and **c**, it is necessary to determine exactly where the boundaries between program and adverts will occur. The only system that has this information for a live TV program stream is the broadcast origination system. LTN specializes in integrating with these systems and retrieving not just the basic timing information required above but also other vital information necessary to fill the entire break and optimize the revenue.

Extracting the precise timing and other data from the broadcast automation systems and master control is one thing. Configuring that information and delivering it in a timely fashion to the downstream system in a form that can be used readily by that system is quite another. SCTE 104 and SCTE 35 are the standards typically used to achieve this, but there are fundamental flaws with

this approach that LTN's patented temporal fingerprinting technology solves. These flaws are:

a SCTE 104/35 are inconsistently used. Every media company has a different interpretation, and this adds so much complexity for vMVPDs that many have limited its use to the local avail breaks only.

b SCTE 104/35 messages are easily corrupted or lost altogether by normal processing — and if the message isn't received, it isn't much use.

c SCTE 104/35 messages are sent in-band with the video, which means they are delayed with the video so cannot be sent far enough in advance for some systems. Live sport, for example, only provides 0.5 seconds of notice.

Filling the entire break

By using the rich metadata provided by LTN,

any ad insertion system can make smarter decisions about which ADS to use and even have multiple ADSs working in parallel.

Today, a contact closure typically is activated by the automation system — triggered from data placed in the playlist — which then generates a single SCTE 104/35 message used by all downstream systems to trigger dynamic ad insertion. Unfortunately, this clockwork mechanism doesn't meet the needs of replacing variable-length ad breaks at scale.

For an ad decision system (ADS) to be able to properly fill an entire break, it must know how long the break is. The more time the ADS is provided to make the decisions regarding which ad or ads to place, the better it can perform its job. This is especially important, for example, when it comes to large numbers of viewers simultaneously watching a live sports event, with every viewer going into an ad break at the same time and requiring a different ad (if ads are to be targeted) within the same time slot.

Though some solutions have been tried, contact closures just don't cut it here. The basic approach is to assume that all breaks will be at least, say, 3 minutes in length; to ask the ADS in advance for a 3-minute break for every viewer; and to wait for the contact closure to signal the start of the break and then start inserting. Another contact closure signals the end of the break and — with OTT only (this won't work at all for non-OTT) — delay each viewer's stream until the end of the ad that is playing when the end-break signal arrives. This may (with the emphasis on may) fill the break, but it will also result in a steadily increasing latency during the viewer's session. In other words, it compounds the latency issue. This approach

also makes the ADS work unnecessarily hard and leaves a lot of ads on the floor with orders unfulfilled. In short, it is inefficient.

LTN's approach is to provide the ADS with the break structure in advance, along with the context: information about the programming that surrounds the ad breaks, the exact length of breaks, and what's in them (in case there is no need to replace some of the ads). This information can be extracted by LTN's technology from the broadcast origination system and can be communicated well ahead of time. A contact closure-based SCTE 104/35 system simply cannot do that. By using the rich metadata provided by LTN, any ad insertion system can make smarter decisions about which ADS to use and even have multiple ADSs working in parallel. Altogether this drives better fill rates, with more contextually relevant ads targeted to each viewer, which in turn leads to better revenue optimization.

Delivering value for advertisers

A key difference between traditional scheduled linear TV and OTT is where ads are inserted. For traditional TV, they are all inserted at the origination site and are scheduled by a traffic system for display to anyone watching at the time. Ads are sold months, weeks, days and even minutes in advance, and media companies have built extremely sophisticated traffic systems to maximize the revenue from advertisers by providing placement, exclusivity and sponsorship opportunities within the schedule.

OTT, on the other hand, is today almost exclusively based on a programmatic sale for a specific viewer a few seconds at most before the insertion occurs. The ad is then dynamically inserted into the stream for that viewer. However, for live programming, this approach creates problems not only because of the limited time to make decisions for large numbers of simultaneous viewers but also because each decision is independent of every other decision made so there are no break rules. This results in scatter.

LTN's approach to delivering information about the structure of all possible breaks in advance allows the ADS to fill one or multiple breaks at a time, which means that placement in the break and exclusivity in the break can be sold too. In live sports the breaks may still be scheduled in advance, but the order they roll in depends on the events on the field. With this approach, the chances that the break will be filled (without the need for slates saying "We'll be right back!") with

ads that are contextually relevant and targeted to the specific viewer are increased massively. In the worst case scenario, the ADS cannot fill the whole break, in which case the existing ads can be selectively left in place.

This approach also allows for the flexible coordination of the sale of inventory between traditional broadcast and OTT. For example, a media company may choose to allow a vMVPD to sell all inventory, both national and local, for certain devices or particular demographics. This flexibility and control can only have the effect of maximizing the revenue opportunities for media companies. Another way the rich metadata provided by LTN adds value is with the advertising embedded in DVR-recorded content and live-to-VOD. The timing information LTN extracts from master control allows for frame-accurate automated live-to-VOD and cloud DVR. Today, in most cases, live-to-VOD is a separate workflow, and DVR recordings are not frame-accurate. However, LTN's metadata can help correct for these challenges to flexible ad insertion.

The LTN Network and LTN's technologies can both be used within or added to existing workflows, which makes the solutions provided much easier to deploy. The solutions are modular and utilize standards wherever possible to interconnect with other infrastructure seamlessly.

With proper marking of the exact location and identity of the programming, ads and promos in the live video stream, it is possible to automatically process the video frame accurately. This means that viewers can DVR just the program, and the decision on what to do with the ads can be made on playback. The system can leave the ads as they are for potential C3/C7 credit, replace the entire break, replace the ads in earlier episodes with the current week's ads (helpful for binge watchers), etc. None of this is possible today without either heavy and costly manual intervention, or the benefit of LTN's deep integration into master control.

Ease of implementation

With respect to video distribution, the LTN Network accepts uncompressed video and audio or IP encoded content at the source — whichever is most convenient. At each receive point, the LTN Network can hand off uncompressed video and audio or IP-encoded content independent of the source format. The network can be used just as easily for point-to-point (i.e., contributing a college sports game to a network), point-to-multipoint (i.e., distributing a Diginet to hundreds of broadcast stations) or multipoint-to-point (i.e., aggregating all affiliate stations for a big three-letter broadcaster to a single ingest point for distribution to one or more vMVPDs).

LTN's technology seamlessly integrates with the complex master control environments that already exist at the network/broadcaster's origination facility to determine exactly what is on air and, crucially, what is going to be on air and precisely when. This information can then be delivered out-of-band (i.e., via a regular unmanaged internet connect) without any impact on existing video distribution and then re-synchronized with the video downstream.

With this information, the SSAI system can understand precisely what is on air and what is coming well in advance. Sophisticated SSAI systems can use this knowledge to get one ADS (or more) to "pre-decision" ad breaks for all current viewers, transcode all the targeted ads, deliver them to the existing CDN(s) and re-write individual viewer manifests in time to ensure that all ad breaks are filled optimally. All this can be executed without changing any existing video delivery/distribution infrastructure.

Conclusion

Media companies' distribution of live linear programming over the top in a way that facilitates sales of advertising inventory in a coordinated, controlled manner will drive adoption and improve the business model. LTN's approach to solving the challenges described in this document combines technology developed specifically to enable existing broadcast infrastructure to be re-used over the internet with a fully managed service approach to the end-to-end workflow.

LTN adds a managed layer on top of the internet, enabling it to be operated in much the same way as traditional distribution. LTN also uses the rich metadata from the existing live master control and origination systems to fully enable automated and seamless targeted ad insertion at scale on any platform.

At the end of the day, the internet is just another vehicle for distributing live/linear content. With the LTN approach, media companies can overcome challenges stemming from the way the internet is architected and fully leverage the medium to reach the vital millennial and Gen Z demographics with the same quality and at the same scale at which they traditionally have served premium content to the baby boomers and Gen Xers.

Contact us

Technological innovation is driving growth and creating new opportunities for the broadcast and media industry. New competitors are vying to capture viewers, leveraging new technologies in unique and creative ways. If you are ready to begin exploring how to reach a greater global audience, connect to media partners around the world and achieve new, more efficient and effective workflows, LTN is your guide.

LTN GLOBAL

LTN provides highly reliable end-to-end transport for live video with low latency using its dedicated multicast IP network infrastructure. The Internet is used only for the first and last mile connectivity as necessary.

