BROADCAST
THE BRIDGE
Connecting IT to Broadcast

Intelligent CDN For Broadcasters

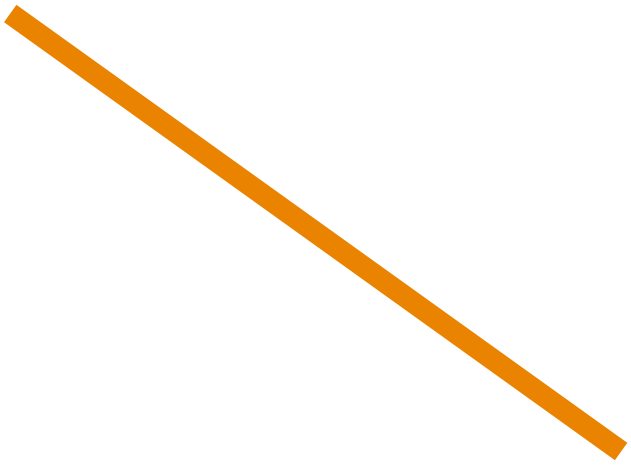

Essential Guide

EG

ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

Media streaming over the internet is unique. Packet switched networks were never designed to deliver continuous and long streams of media but instead were built to efficiently process transactional and short bursts of data. The long streams of video and audio data are relentless in their network demands and to distribute them effectively requires the adoption of specialist CDNs.

Networks assume the data will be bursty and to a certain extent they rely on this transactional demand to make systems more efficient. The synchronous nature of broadcasting may keep latency very low, but it does so at the expense of flexibility. Packet switched networks have incredible flexibility, but we always need to keep an eye on the latency.

Web server technology is built on layers of underlying protocols. HTTP provides a data structure for web pages, TCP guarantees delivery of packets, and IP is the lowest level protocol until we reach the data link layer at the physical network interconnect, such as ethernet, WiFi, or HDLC. The combination of HTTP/TCP/IP describes the protocols needed along with their hierarchy to transport data across the internet.

Broadcasters are used to maintaining backwards compatibility and we've spent the past seventy years making sure the latest technology is put to good use while guaranteeing the viewers using the previous generation of technology can still watch their programs. However, media streaming over the internet is different as broadcasters have been forced to adopt a technology that wasn't designed for them and has led to a lot of shoehorning and compromise.

Latency is an inherent characteristic of the internet and any TCP/IP network. From a web browsing experience, users are generally interested in the response to mouse clicks and other isolated events. Any variable latency isn't usually noticed and has little effect on the user experience. If the response times are too long, then this will have an obvious detrimental consequence for the user, but the recovery time will often be very quick. However, the continuous data streams created by video and audio streaming is another matter and doesn't lend itself well to variable and long latency, in part due to the widespread use of large video buffers.

CDNs have been available for many years and seem promising for media streaming, but they are often built to optimize web server type applications. Due to the complex nature of media streams, they often demand special consideration to maintain high data throughput while keeping latency low.

A new breed of CDN is now emerging that optimizes media streaming by taking into consideration the specific needs of video and audio data. By focusing on the QoE and QoS associated with video and audio, the specialist CDNs help keep data throughput high, and latency low.

Machine learning (ML) is playing an increasingly important role in these specialist CDNs as a wealth of metadata is available from distributed monitoring probes throughout the internet. Combined with the monitoring information and datasets available for video and audio streaming, ML is improving reliability and resilience while keeping flexibility high.

Tony Orme.

Although IP networks and COTS infrastructures are providing unprecedented opportunities for broadcasters, we should always remember that video and audio is different to web server traffic, and this demands more focused and specialist solutions to be provided, especially for CDN.

Tony Orme
Editor, The Broadcast Bridge

Intelligent CDN For Broadcasters

By Tony Orme, Editor at The Broadcast Bridge

CDN is a term that is used regularly and many hope it will solve the challenges of distributing IP media over the internet. However, the CDNs needed for broadcast applications are unique and specific technological challenges must be overcome.

The internet is built from just over 100,000 privately owned networks known as Autonomous Systems (AS) that collaborate to connect servers and ISPs to give the impression of one single network.

When a user browses a web page, the IP datagrams forming the messages and server responses traverse through the connected networks. Each AS has a commercial relationship with the ASs it connects to so seamless routing can be achieved.

Internet Operation

One of the challenges we have with the individual networks that comprise the internet is that they were originally designed to transport HTML pages between web servers and browsers, not long media streams. HTML provides a simple text-based markup language that is easily edited by users, a system that is the core of web browsing today.

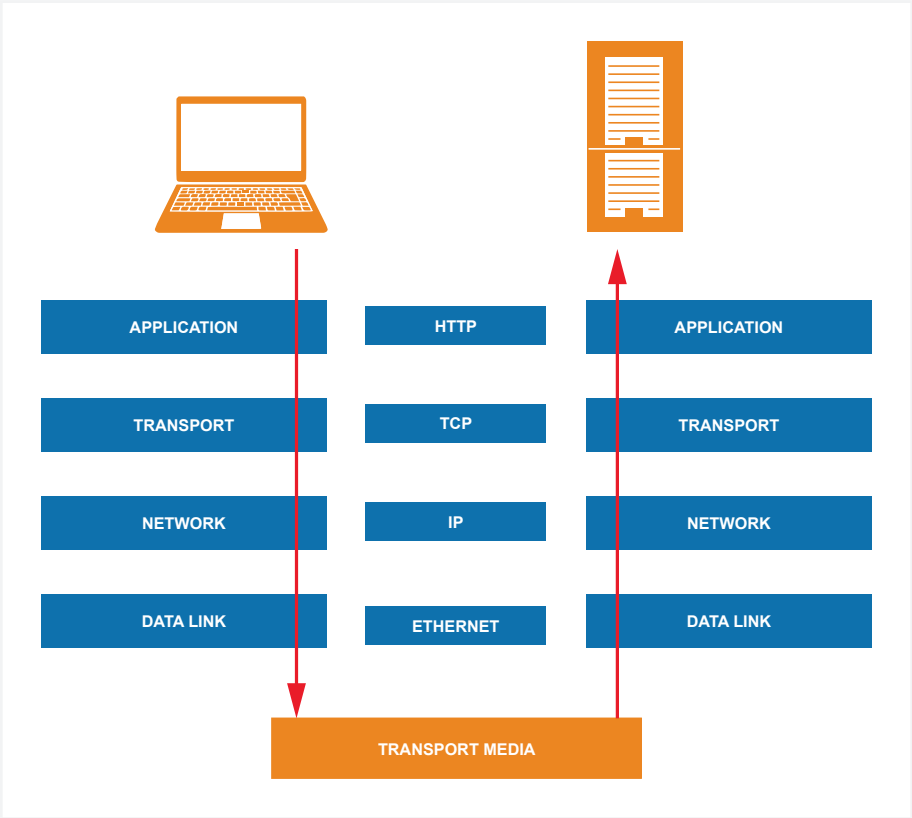

Figure 1 – A browser requesting a webpage update must transfer from HTTP to TCP to IP and then across the physical networks, the response from the web server is the opposite. This mechanism provides the fundamental operation of the internet and has been fine-tuned over many years for short messages. Streaming media uses long files resulting in many compromises when delivering over the internet.

HTTP provides a specialist transport mechanism that encapsulates the HTML files so that compliant browsers and webpage servers can easily exchange data. Control messages and data responses must have their own protocol otherwise the server would just receive meaningless data and not know what to do with it. Adding the HTTP layer brings structure to the message and data exchanges.

TCP is needed to guarantee delivery of the IP packets as the underlying IP protocol only manages best-effort packet distribution. Although this may sound like an oversight, it is a great strength as it keeps latency incredibly low and provides much greater flexibility for system designers when building protocols to work with IP.

Short Duration Messages

The combination of HTML, HTTP, TCP, and IP leads to a message structure that encourages short bursty data exchanges. Intuitively this is correct as anybody surfing a non-media type webpage will be clicking on menu selections and hyperlinks, and then receiving data from the server in response to these requests.

From a user's point of view, short messages are preferred as latency will be low, and hence response time will be fast. In webpage surfing, one of the key quality checks is the response time of the webpage from the users input.

The need to keep user-response times as fast as possible has led to the internet being developed and streamlined to exchange short messages. Generalized CDNs used throughout the internet embrace this philosophy to maintain high levels of user experience. However, this type of operation is completely the opposite to the type of delivery required by media.

Streaming video and audio, whether within an embedded webpage or a full blow OTT system uses very long continuous files, and for live broadcasts the live stream doesn't even have a file length as it carries on for the duration of the transmission. But to transfer a media file over the internet, we must be compliant with the fundamental transport stream and the webpage server and browser architecture, that is HTTP/TCP, and these require short transactional bursts of data. Consequently, media streams are subdivided into smaller blocks of data through the process of chunking, but this leads to quality compromises.

Also, streaming media suffers from RTT (Round Trip Time) delays which are a consequence of the TCP protocol and are particularly problematic for long distances. When a server sends a block of data to the web browser it waits for the browser to acknowledge the data by sending a message back to the server. Only when the server receives this ack-message can it send the next sequence of the media stream. The objective of this exchange is to guarantee packet delivery and resend any packets that are lost in the network, a consequence of this is increased and unpredictable latency.

Media Special Requirements

Media streaming relies on video and audio compression. If an unrecoverable packet loss occurs in the GOP sequence, the video disturbance is evident for many seconds. One method of reducing this sort of risk is to increase the media players buffer size, however, this often results in long latencies.

Many scenarios occur within a network, such as lost packets and buffer management issues in switches and routers, that result in the TCP packets sent from server "timing out", so it must resend the data, which in turn leads to the potential for duplicated packets wasting valuable bandwidth, thus exasperating a poor user experience due to increased latency.

CDNs succeed by moving the media files closer to the user by caching the data. In principle, this sounds like the perfect solution as the RTT times are significantly reduced leading to a decrease in latency, and congestion can be better managed. However, many of the CDNs are designed for general internet traffic consisting of short messages from servers delivering static content, not long media streaming files.

Making Media Fit

Although the media file or stream may be divided into smaller chunks of data to allow it to progress through the internet, the consequences of the media being an incredibly long file cannot be understated, even when it's separated into short chunks. Webpage technology relies on the exchange of data being stateless. In essence, this means each data transaction is treated in isolation as an individual mouse click, or data entry, creates a unique message-response action. We may have worked around this by splitting the video and audio into small chunks of data, but they cannot be treated in isolation as they are assembled in the player to form a continuous streaming experience.

Another consequence of IP delivery is that there are no guarantees that packets will arrive in the correct order. This is particularly evident if subsequent packets take different routes through the ASs to the ISPs. Not having the constraint of making sure subsequent packets take the same route does make the internet more efficient for short term messages as TCP re-orders packets to put them in the right sequence, but it does so at the expense of latency. Furthermore, broadcasters have no control over how packets route through the internet.

Another issue within generalized CDNs is that their inner workings generally lack visibility. Also, broadcast traffic comprises a relatively small amount of the overall network traffic leading to limited focus from service providers. Companies such as YouTube, Netflix and Facebook all employ their own private CDNs to overcome this.

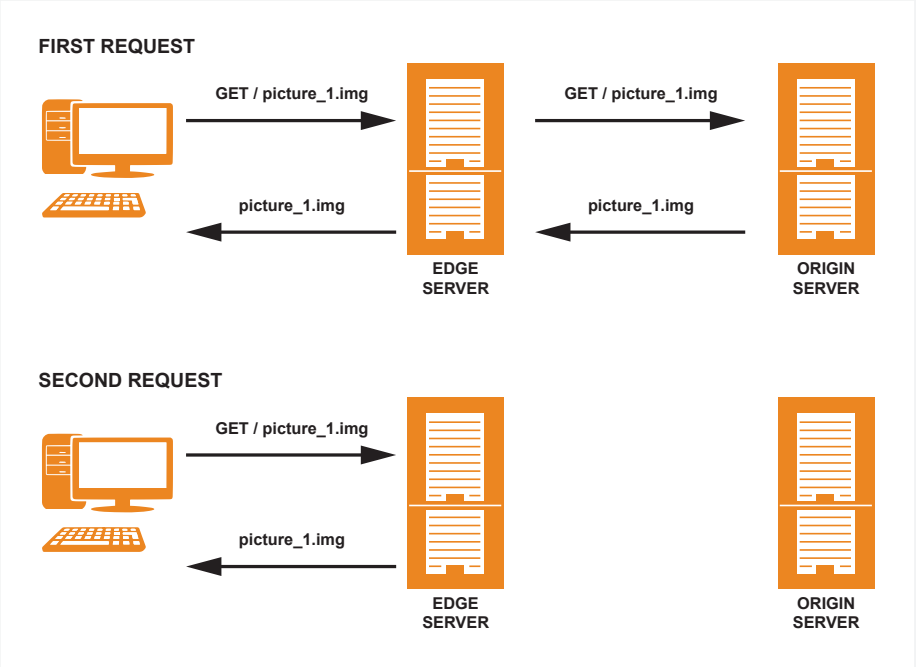

Figure 2 – Generalized CDNs move the data closer to the user to reduce latency. However, this model assumes the data is static so very few requests are made to the origin server. Although streaming uses this model, it results in very high latency, especially during live broadcasts. Specialized CDNs fix this challenge.

The cost to broadcasters of building their own CDNs by themselves is prohibitive. Anybody building a CDN does not only have to consider the deep knowledge needed to create multi-layer software, the cost of the infrastructure, physical network, and connections, but also how they form the contractual relationships to integrate into the ASs and ISPs that make up the internet. These are often reciprocal relationships that rely on the trading of connectivity into homes, something broadcasters probably will not have.

Specialist Knowledge

The historical broadcast television model is essentially a fire and forget system, that is, the video and audio is presented to the television transmitter and then sent to all homeowners in the region. There is no reverse link to verify the signal arrived intact other than confidence off-air monitors in the vicinity of the broadcast center. Consequently, building a huge regional or national private CDN system would be beyond the knowledge and experience of most, if not all broadcasters.

A broadcaster can easily use a generic CDN, but they do so at the expense of a reduction in the amount of control they have over the allocation of data bandwidth and latency guarantee within the network, leading to the potential loss of viewers through a poor viewer experience.

IP is delivering unprecedented opportunities for broadcasters and seeing the number of facilities transitioning to IP is testament to this. As an industry, we have been able to ride on the back of the massive amount of innovation that has occurred in IT and other industries such as finance and medical. However, due to the uniqueness of streamed media, generic IT and COTS solutions need optimization.

Streaming media for OTT using CDNs is one area where specialization is needed. Specialist CDNs solve many of the challenges already discussed as the proprietors have an in depth understanding of how media is constructed as well as understanding the expectations of the viewers.

Specialized CDN

For a CDN to be successful it needs to have massive interconnect to ASs and ISPs, provide dedicated capacity and low latency, and accurate network measurements. CDN suppliers who focus on media data capacity, latency and network measure are critical for broadcasters looking to provide their viewers with the highest possible QoE.

Specialist CDN providers will have their own network and agreements with ASs and ISPs dedicated to media delivery, and not just generic internet traffic. This opens massive opportunities for broadcasters as data bandwidth can be guaranteed along with very low latencies. Also, as the specialist media CDN provider will have visibility of their entire network, they will have a mountain of analysis metadata available from their monitoring probes.

Network metadata is the new tool available to broadcasters. Not only does it provide the opportunity to understand how a network is performing, where the pinch points are, and where faults have occurred, but it also facilitates the new breed of data-led optimization through machine learning (ML).

Realtime monitoring provides visibility of how the low-level characteristics of the network are performing. This is far more than just checking the averaged data rates or packet loss; it provides relevance to the measure. Media is special as the individual packets of data cannot be treated in isolation. We cannot assume that just because a sequence of packets has a small latency then this can be extrapolated to the whole video stream. Delivering media over the internet requires an understanding of the link between the transport stream and media essence.

Data-Led Learning

ML is powerful as the analysts do not need prior knowledge of the patterns within the dataset. ML is data-led learning meaning that the neural network, or whatever topology the designers have chosen, finds the patterns within the training data so the model can then be applied to unseen data. Quite often, the ML model will learn trends within the dataset that human analysts were unaware of.

Key to ML is the dataset, and the more diverse the dataset, the more likely the ML engine is to learn patterns in behavior that were otherwise unknown. It's no coincidence therefore that specialist CDN suppliers have an incredible number of probes for recording metadata parameters that can be applied to this task.

Through ML learning, patterns can be detected that may point to the imminent overflowing of a switch buffer, or a potential failure in a link. Often, these patterns, or anomalies, are detected so quickly that they are used to control the network itself by changing routing or data bandwidth allocation. Because a specialist CDN provider has knowledge and visibility of their network, they have a great deal of influence of how the network performs and can react to these automated changes very quickly.

Based on the analysis of historic metadata, specialist CDN providers can provide accurate costings for services. Network utilization increases during peak viewing and understanding where and when this happens allows the accurate prediction of costings for media delivery.

Conclusions

Viewers may be switching to mobile device and smart TV viewing in their droves, but they still expect the same QoE found in traditional static television models. The viewing immersive experience is critical and the communal enjoyment the internet now delivers through social media is without compromise. Although the internet is providing untold opportunities for broadcasters, the unique requirements of streamed media demand specialist solutions, and none more so than in the provision of CDNs.

The Sponsors Perspective

Video Streaming: Strategies And Performance Challenges In A Rapidly Changing Environment

By Antonio Corrado, CEO, MainStreaming.

Entertainment over the internet has gained significant traction over the last years. For this reason, companies have developed new business models in order to retain customers, by meeting their emerging needs and studying the behavior patterns of online streaming consumption.

This new way of consuming content has been driven by the rapid adoption of smart devices, like smartphones, smart TVs, laptops, and desktops. The viewing of content on these devices has changed the habits and expectations of users. New viewers are unlikely to sit in front of their TV waiting for their favorite show to start. They want to watch their preferred programs wherever and whenever they want.

This pressure has led traditional broadcasting companies to create their own online platform offerings, as well as produce their own content in order to engage viewers. By embracing the internet approach, many broadcasters are now fully committed to the online formats transitioning.

Streaming Services Need To Deliver A Flawless Experience

A high-quality delivery is fundamental, as subscribers are used to turning on their TV and getting the same, consistent quality they've always gotten. Subscribers will be looking for that broadcast-like quality of experience when viewing content across their device, and for this reason a flawless streaming is essential to their overall experience.

In addition, with time-sensitive video content such as sports, news, or pure OTT content like esports and interactive shows, viewers expect to watch events as they unfold with no tolerance for loading time or buffering.

Content Delivery Networks, with their current infrastructure layouts, are unable to guarantee adequate quality when dealing with large numbers of users. It has emerged the need for new technological architectures capable of delivering the desired viewing experience.

Low Latency Has Become More Important Than Ever In Streaming Services

When streaming over the internet started to become part of the broadcast workflow, the latency of a minute or more was not uncommon, and was mostly accepted. At the time, streaming was not used for live events or online gaming competitions, so high latency was not a problem. As new habits developed, and the entertainment industries shifted business models from linear broadcasting to IP based streaming for on-demand and live events, the lowering of latency became a must.

Reducing latency has been an ever-evolving quest for the streaming industry, and achieving this would allow business models of premium live events and gaming to flourish. This is only possible with a technology specifically designed for online streaming, that provides a resilient and reliable service, and consequently a high quality of experience for viewers.

Sustainability, A Factor To Consider When Choosing A Provider

Digital video files are getting heavier as viewers today expect to watch their content in high definition, at any time and on any device, resulting in an enormous amount of energy needed to keep data centers cool and running. The current network infrastructure uses an unprecedented amount of energy that is highly inefficient and detrimental to the environment.

Data centers are essentially the factories of the information age; their 24/7 operation makes online browsing, streaming and communication possible, but delivering all this data requires a tremendous amount of electricity. Significant improvements to IT hardware and data center energy performance must be made to keep up with the market's requests without causing harm to the environment.

Antonio Corrado, CEO, MainStreaming.

To watch a video on a device, the content must be encoded. The technology used to achieve this must be capable of parallel computing to encode a video in a few seconds/minutes for a quick delivery to viewers. It is essential for broadcasting media and gaming companies to shorten the processing time for content publication, especially for sports, news and gaming, as they must be uploaded online in real-time. The process for uploading an on-demand video has to be simplified as much as possible, utilizing a parallel computing system with innovative technology that sends out packets to multiple locations, where they are encoded and then brought back together to generate the complete video.

It is important to develop technologies that can reduce energy consumption. Given the numbers that must be managed online, it is necessary to innovate, while implementing a green approach. In short, not only does the traditional approach not guarantee a smooth and perfect delivery but, because of the infrastructure currently in place, it leads to an excessive amount of energy consumption.

An Intelligent Media Delivery Platform

Flawless delivery, latency reduction, and sustainability are the main distribution goals that an intelligent platform needs to achieve.

MainStreaming's iMDP is a single integrated platform, based on proprietary technologies and developed specifically for video content, that manages the entire streaming workflow. Intelligent algorithms allow for a constant assessment of the network conditions and identify the most efficient data path, by making decisions based on a continuous flow of metrics related to the Quality of Experience of each viewer. Its features grant content owners full transparency, control and scalability. Furthermore, the iMDP minimizes its carbon footprint by providing more efficient network resource utilization.

The Challenging Demands Of Streaming Services Scalability

The iMDP provides highly-customizable streaming services that adapts itself to accommodate the most demanding live and OnDemand events. This is possible thanks to an intelligent and dynamic network infrastructure that routes the traffic to the best path to reach viewers, ensuring the highest quality of experience.

There are three service options included in the iMDP, which are important to analyze when determining what is best for your streaming workflow: Public, Private and Hybrid.

The Public option is a ready-to-use, all-inclusive, end-to-end media delivery service. It provides full control over the streaming process while supporting the growth of subscribers and meeting the needs of peak distribution and monthly volume use.

It is possible to customize the distribution network structure with a Private option, based on a decentralized resource approach. This service is created for the management of high-volume traffic for large-scale customers. It allows customers to deploy dedicated nodes within ISP infrastructures or dedicated peering and have a financially sustainable solution able to manage millions of users.

The third option available is a Hybrid service, which is a mix of the Private and Public service. One or the other can be used according to your viewers' location.

BROADCAST
THE BRIDGE
Connecting IT to Broadcast

CI
CORE INSIGHTS

EG
ESSENTIAL GUIDES


MEDIA


ARTICLES

For hundreds more high quality original articles and Essential Guides like this please visit:

thebroadcastbridge.com

11/2021

Supported by

MAINSTREAMING®
The Intelligent Media Delivery Company