Delivering the IP Future

Enabled by LUMEN

Forward

The changing nature of video delivery is a topic that no media executive can afford to ignore.

As we all know, the media industry is in the midst of vast transformation – one which is fundamentally rooted in data; and one which spans every phase of the workflow, touching every team. Perhaps the most crucial change is happening in how audiences are choosing to consume their content – with the flexibility of streaming services enabling new viewing experiences. With audiences moving online at an accelerating pace, we ask whether we are truly ready for a totally IP-delivered world. *Delivering the IP Future* looks at challenges and considerations for realising this vision.

Our customers around the world are delivering video in unprecedented volumes, and exponential growth in bandwidth consumption fuelled by IP video delivery is set to continue. The result is an internet in which media delivery is arguably the primary use case, which represents something of a departure from its design origins and architectural assumptions. We – the broadcast and media community – need to look carefully at the potential challenges posed by this and work together to deliver the online experiences viewers demand, at broadcast scale.

This new research provides an industry-wide view on the evolving landscape of media delivery and helps us all to direct our attention where it matters most.

It also provides insight as to how the scale challenge can be conquered, identifying key technologies which will play a role in enabling more efficient use of network resources as online audiences skyrocket. The evaluated technologies include:

- Content Delivery Networks
- Peer to Peer
- Edge Computing
- IP Multicast
- 5G Broadcasting

Lumen empowers Media and Entertainment customers to deliver against their IP-enabled vision: from acquisition, through data centre and cloud connectivity, and on to global content distribution. We are delighted to support industry forums which play an active role in driving positive change and help bring about new ideas which enable us all to realise our collective vision for the best customer viewing experience imaginable.

Together, we can build a better internet for media customers and their audiences.

Lumen Technologies, EMEA

LUMEN

The Context

In May 2019, experts from content and technology companies gathered at a DPP AT HOME event to discuss if – or perhaps when – we will arrive at a fully IP future for distribution of content to consumers. While opinions differ about whether traditional broadcast networks will cease to be used altogether, it is undeniable that content consumption over the internet is growing fast, and that growth is accelerating.

Every major broadcaster now offers some kind of VOD service, whether catch-up or library, AVOD or SVOD. Most stream their linear channels online too. Huge new global players have been born digital, while longer-established giants are pulling their content from other services – including broadcast channels – to bolster their own online offerings.

Generational changes in viewing habits are accelerating this change, whether broadcasters want it or not. In the UK, 16–34 year olds are around twice as likely to watch SVOD services or YouTube, and 25% more likely to watch broadcaster's VOD services, compared with the general population. And they watch only half as much broadcast TV[1] .

Recently the shift has become more dramatic, as traditional live broadcast events are now moving online. Facebook has been streaming La Liga in some markets, and Amazon snapped up rights to US Open Tennis and Premier League football in the UK. And from 2021, the Champions League rights in Germany will leave broadcast channels to be split between two streaming services, Amazon and DAZN. Where the shift to online viewing has thus-far been relatively gradual, the fact that some of these huge events are now available *exclusively* via an OTT service in some countries generates a significant and sudden increase in online viewing.

Meanwhile the BBC have used their online player to deliver Wimbledon and the World Cup in UHD, delivering new services at higher bit rates than ever before.

Today, many of the world's biggest sporting events – such as the Super Bowl – are dominated by broadcast television viewing. Indeed, even though the Super Bowl in 2019 saw a record 2.6m concurrent streams, there were 98m broadcast TV viewers. But it would be a mistake to take this as a sign that broadcast TV will continue to dominate; only that such events are one of the big challenges for a fully IP future.

So if that fully IP future is heading our way, will we be ready?

1 https://www.statista.com/statistics/898353/daily-audiovisual-viewing-time-by-device-uk/

The Approach

96% of the attendees at our DPP AT HOME event felt that the transition to IP would be at least somewhat problematic, with 59% seeing it as *very* or *extremely* problematic.

One of the top challenges they identified was *peak load scaling*, showing concern at the ability of our technical infrastructure to meet the demands that will come from the highest volumes of content viewing.

So in this report we ask: **how can we rise to the challenge? Which technologies will help us to meet the demands of scale, and over what time period do we need to invest in them?**

Through in-depth interviews with experts at DPP member companies, we have been able to build a picture of the current technology landscape and its direction of travel. We've gathered and compared the views of companies across the distribution chain – from broadcasters and online content platforms, to CDNs and networks, to specialists in emergent technologies. This has allowed us to form a balanced view on the technologies that matter, and the considerations that should be on the minds of leaders across the industry as we forge a path into the IP future.

In this report, we will begin by exploring the standout challenge: large live events. Consecutive sections will then explain each of the key technologies that might help us to significantly scale up delivery over IP. Finally, we draw key conclusions about the direction of the industry, and what it means for us all.

The Experts

Akamai
JAMES TAYLOR
Director of Media Services, EMEA

BBC
RICHARD COOPER
Controller of Digital Distribution
PHIL LAYTON
Head of Broadcast and Connected Systems

BT
SIMON JONES
Chief IPTV Architect

DAZN
ALASDAIR KENNEDY
Head of Video Architecture

DTG
RICHARD LINDSAY-DAVIES
CEO
SIMON FELL
Independent Chair

eluv.io
MICHELLE MUNSON
CEO

Equinix
MATT GEORGE
Senior Manager, Enterprise Solution Marketing, EMEA

fastly
LEE CHEN
Head of Streaming Media and Cloud

Lumen
PIERRE-LOUIS THERON
VP Product Content Delivery Services

Telstra
BERNHARD PUSCH
Head of Global Internet Strategy
ANNA LOCKWOOD
Head of Global Sales, Telstra Broadcast Services

touchstream
BRENTON OUGH
CEO

ViacomCBS
DOUGLAS WILLOCK
Senior Director, Distribution Operations and Support

The Live Conundrum

It's hard to disagree that the internet will become the dominant mechanism for content distribution over time. As an interviewee from one broadcaster put it, "We absolutely consider a world where broadcast doesn't exist. I believe there's an inevitability to that."

> **We absolutely consider a world where broadcast doesn't exist. I believe there's an inevitability to that.**

As an industry, we're well underway on that journey, with online viewing having increased by a factor of 1,000 in the last decade. One content provider explained that their current growth in internet-delivered data is 40% year-on-year, factoring in both increased viewership and increased bitrates. They expect that growth rate to continue for the foreseeable future.

There isn't universal agreement on the timescale for the transition ahead, but a range of 10 to 15 years was discussed by many. That's a long time in media; after all, 15 years ago, Netflix was just a DVD rental service, and neither the iPhone nor the BBC iPlayer would launch for another three years. We can be sure that the timeframe won't be consistent around the globe, with many contributors noting that each country has its own profile of viewership across devices (TVs, tablets, mobile phones) and networks (broadcast, fixed line, cellular).

A timeframe of well over a decade means that as the solutions evolve, the problem will too. Many predict that linear channels will decrease in importance (though few foresee their total demise), so it seems logical that on-demand viewing will continue to increase. Indeed, one CDN told us that they serve four times more VOD traffic than live. Yet large live events certainly won't go away. These are the moments at which the country, or even the world, comes together around great content.

The challenges created by these huge live events were hotly debated. Many feel that these major spikes in live viewing will cause a problem, or indeed, already do. Infrastructure must be scaled for these maxima, which may not be practical or cost effective with current approaches. As one contributor explained, the problems of VOD are largely solved, whereas live is a different problem set, and people are still struggling with it.

Others felt that peaks in live viewing will naturally displace viewers from VOD, or other services like gaming, resulting in only a modest increase in overall bandwidth use. There are, they argued, significant spikes in VOD usage too, such as when a new season of a high profile drama becomes available. Nonetheless, it's telling that those who deal with large live events every day were much more concerned about the ability of networks and CDNs to cope, compared with those whose primary output is on-demand.

Live content is also often some of the highest value content, with the highest Quality of Experience (QoE) expectations from viewers. A viewer of a drama series may be irritated by some buffering or a dropout; but if it only happens once or twice they can replay the problematic content without too much trouble. Whereas if a sports streaming customer experiences even one dropout at a crucial moment in a football game, they may very quickly become an ex-customer.

The criticality of delivering live events well – and the challenge of doing so – has been demonstrated through a number of high profile failures. When Australian mobile network operator Optus became a rightsholder for the 2018 FIFA World Cup, they might have expected to achieve a great deal of publicity. That they did, as their service failed and a number of their games had to be broadcast by SBS, an over the air broadcaster. Even technical giants suffer challenges, with Amazon's launch of US Open Tennis suffering from errors, quality issues, and latency problems for some users. Of course each of these problems was overcome, but should media technology professionals lie awake at night worrying about the risk of further problems, as more live events move online?

> **The criticality of delivering live events well – and the challenge of doing so – has been demonstrated through a number of high profile failures.**

This question, of live distribution, sits at the heart of the choices we make around technology investment. Multicast, 5G broadcast, and peer to peer all have particular advantages when many viewers are watching the same content at the same time. Most felt that unicast streaming over fixed line or mobile broadband would continue to work well for on-demand content distribution, provided the CDNs and networks continue to invest. But is a more radical approach required to serve live content?

The Technologies

Just which technologies will help with the transition to IP?

With the rapid growth of mobile data services, radio spectrum is increasingly valuable. Already we see the 700MHz spectrum being cleared across Europe, taking frequencies previously used for TV broadcast and repurposing them for 5G mobile services. So will 5G broadcasting be the answer?

Peer to peer delivery still has theoretical benefits in reducing the load on centralised infrastructure, despite its turbulent history and its associations with content piracy. So can it make a resurgence, and alleviate the challenges of large scale IP delivery?

Or will IP multicast be a saviour? Long written off by many as impossible to support over the open internet, it's now in use within the closed networks of many Internet Service Providers (ISPs) to deliver their own IPTV services.

First we had mainframes. Then the personal computer. Along came the cloud. The newest computing paradigm is edge; the concept of performing server-side computing functions closer to the end user, even in the ISP's network or the cellular network. Will it change the game for performance and efficiency?

Perhaps there is no silver bullet. No new technology on the horizon to dramatically change the landscape. Instead, perhaps the Content Delivery Networks (CDNs) and ISPs will grow and expand their infrastructures, keeping up with the scale required by evolving their architectures. After all, today's technologies allowed Indian streaming platform Hotstar to deliver 25.3m concurrent streams of the 2019 ICC Men's Cricket World Cup semi-final[2].

In the following sections, each of these technologies will be explored in more detail.

2 https://www.icc-cricket.com/media-releases/1277987

Content Delivery Networks

Over the last year, Lumen has seen peak bandwidth demand on its Content Delivery Network grow by 43%. Other CDNs report similar levels of growth on their networks. This dramatic growth in content delivery requires specialist infrastructure, and that's the role of a CDN.

Content delivery networks are geographically distributed networks of proxies, caches, and computing resources, that allow scalable and performant distribution of content. Using a CDN means that your content can be served more quickly, to more users, with less load on your own serving infrastructure. They also relieve the load on the congested internet core by reducing the number of requests having to traverse that part of the network. They are a crucial part of any modern internet video distribution system, but their use and their architecture are evolving.

CDN capability is sometimes now considered a commodity, and it is certainly true to say that the costs of CDN services have dropped dramatically as demand has grown. But there are considerable differences in the architecture of different vendors, as outlined below. In conjunction with the physical infrastructure, CDNs also offer software capabilities and other services which may differentiate them. One CDN customer explained, "We see CDNs starting to focus on their unique selling points, whether that's security, functionality, or price".

But ultimately, a CDN exists to deliver content. So whether or not an all-IP future can be achieved is a question inextricably linked to the CDNs' ability to scale. Many of our contributors were bullish, commenting that, "Scale is absolutely not a problem on the CDNs. They want more of our traffic." However, others told of scaling problems, especially when launching services in new countries. "Whenever we go to a new country there's shocks, and we shake up the infrastructure. We have to understand CDN and ISP capacity, then negotiate investment and improvement in capacity and peering."

We have to understand CDN and ISP capacity, then negotiate investment and improvement in capacity and peering.

CDN ARCHITECTURE

A key objective of the CDN is to serve content to viewers with high performance, while reducing the number of requests back to the content origin, enabling greater scalability. Performance can be achieved in a number of different ways, with each CDN vendor balancing factors such as the number of nodes, the compute performance of those nodes, the connectivity between them, and the type of storage used.

Different CDNs use different architectures to distribute content. Some take a hierarchical approach, using separate edge and mid-tiers. Others use a mesh topology, enabling direct communication between all nodes.

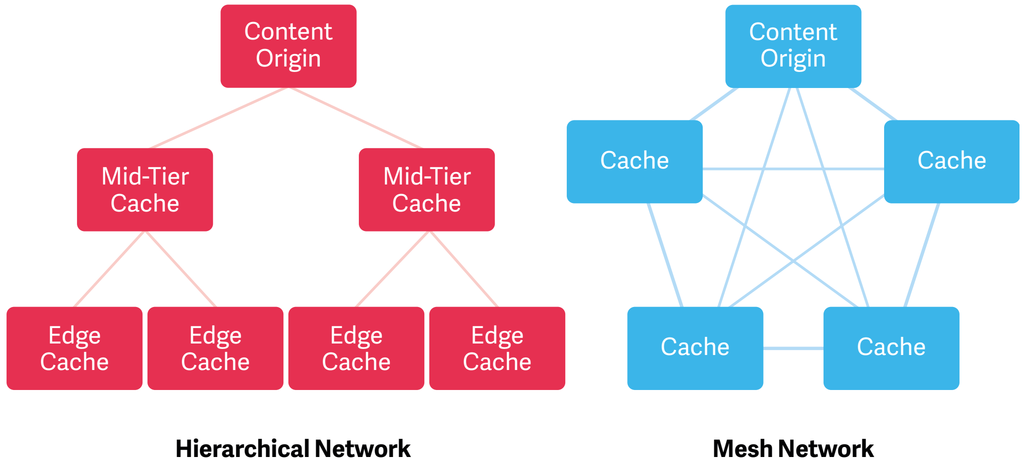

Figure 1: Simplified CDN Architectures

Some CDNs use tens of thousands of servers spread across many points of presence (PoPs) to achieve large scale, while others focus on deploying high performance infrastructure in fewer locations. Some use public or shared networks to interconnect their locations, where others also own private interconnection networks.

In order to avoid flooding the origin with requests, a number of techniques can be used. A hierarchical CDN uses mid-tier caches to avoid each edge cache requesting the same content from the origin. Other vendors use software co-ordination between their nodes to collapse down requests, serving content to one CDN node from another node which already has a copy, rather than going to the origin.

In CDNs, ISPs, and private network providers, Software Defined Network (SDN) capabilities are starting to have a significant impact. One CDN provider explained that it helps them reduce complexity, managing thousands of servers as a common entity rather than dividing up the intelligence in a more autonomous way. When a customer wants to reach new territories, explained a network provider, their software defined networking capability allow them to rapidly configure a "virtual network appliance" in the required locations. "A company with physical points of presence in the UK and the Netherlands may want to deliver to Sweden suddenly. They can set up virtual connections to networks there very rapidly." This kind of flexible use of networks was discussed by others too. "Most companies don't need their own point to point connections; public cloud has encouraged investment in reusable network connectivity and layer 2 peering is widely available at relatively low cost."

CAPACITY PLANNING

Cost is of course an important factor in managing any distribution channel, including CDNs. One broadcaster told us that their CDN distribution cost is now "in the same ballpark" as their satellite distribution cost. But unlike the relatively fixed costs of satellite or terrestrial transmission, CDN costs (in most cases) scale with the number of viewers. "I know how much it costs to deliver a programme to a viewer [through our CDNs], and we comfortably get that back even by running a small number of adverts," explained a broadcaster.

While some have experimented with fixed-price deals, the majority of CDNs charge based on usage, with significant discounts for up-front commitments. This creates an economic incentive for content providers to accurately predict their capacity needs, ensuring they maximise their commitment while minimising overages.

Alongside that financial incentive is of course a technical incentive to ensure that CDN partners are prepared for any planned peaks in demand, to avoid any outages or service degradation. As a result, forecasting capacity has become an important function for any large volume content provider.

CDN capacity planning is a new discipline for us.

"CDN capacity planning is a new discipline for us," explained one contributor. "We have to manage our commercial agreements and predict our usage based on viewer growth projections, and new service launches." Another explained that they've grown a significant internal team to manage CDN capacity. For those with peaky usage patterns due to content such as live sports events, it's important to proactively work with CDNs to plan for the peaks. "We have hypercare relationships with our CDNs. We're on the phone with them days before each big event."

BUILDING RELATIONSHIPS

This proactivity relies on a close relationship between the content providers, CDNs, and network operators. As one network ISP put it, "Scaling the network isn't that hard... when it's planned!"

In a broadcast distribution chain, the broadcaster's responsibility ends in the air – when the radio waves leave the masts. In an OTT path, one broadcaster explained, "It ends at the network port of the last cache we've contracted". So working with those contracted partners, including CDNs and network providers, is just as important as working with the DTT network operator or satellite platform.

The complexity has certainly increased compared with broadcast networks. "Putting a distribution network together requires putting multiple vendors together. You can't simply go to one vendor. You need to control the entire distribution stack."

Putting a distribution network together requires putting multiple vendors together. You can't simply go to one vendor.

One broadcaster, however, felt that the relationships still aren't as open as they could be. "We don't hear a lot from the CDNs about their technical architecture." For an organisation used to working closely with their partners on the design of broadcast networks, this lack of transparency creates friction. But the conversation has to be two-way, of course. A network operator told us that some more experienced OTT platforms share a lot of data about their usage plans, to help predict the required network scale, where others are less transparent.

Still, all parties in the distribution chain are incentivised to deliver a good service to the customer, and they can only do it through partnership. Following technical challenges with its streaming of the US Open, *The Times* reported that Amazon Prime Video was in talks with UK ISPs well in advance of its first Premier League games. Initial impressions are that this paid off. One content provider explained that they reserve guaranteed capacity with their CDNs for high profile events, adding to the cost but giving them a high confidence in the quality of experience they can deliver to viewers.

If customer delight is the carrot, then public shame can be the stick. Netflix publish rankings of ISPs based on their own monitoring data. While not all content providers would choose to take such a bold public stance, it is certainly true that a poor rating on such a leaderboard will have a negative reputational effect on a service provider.

So how can CDNs and ISPs deliver more reliable performance? Many think it's by working more closely together. One ISP told us that, "We maintain close relationships with over ten CDNs, including some run by broadcasters", while another supplier was clear; "CDNs need to become more deeply embedded into ISPs". The relationships are generally on a one-to-one basis though, with one contributor hoping that the future would involve more group collaboration through industry forums.

Techniques to improve performance include more interconnection – or peering – between CDNs and ISPs, and pushing CDN caching into the edge of the ISP's network. One content provider hypothesised that global giants might take a lead in this area. "Are Amazon going to start putting edge compute into ISPs? Who knows?" Just a few days after that interview, Amazon Web Services announced a new product in partnership with mobile network operators, that places AWS services at the edge of 5G networks.

MULTI-CDN

Probably the most significant IP distribution trend of recent years is multi-CDN. While in the early days it may have been simpler for a content provider to choose a single CDN partner, it is increasingly an accepted wisdom that any modern distribution strategy involves multiple CDNs.

The most significant IP distribution trend of recent years is multi-CDN.

For content providers, multi-CDN allows them to balance cost, availability, and quality. For global reach, there is no single partner who can meet all demands. "If you're truly global, the complexity of the market is a challenge in itself," said an ISP. Even within a given country, one content provider explained, "Different CDNs have regional performance differences. In particular, we consider CDNs and ISPs together – some CDNs have better relationships with specific ISPs."

If you're truly global, the complexity of the market is a challenge in itself.

According to Citrix[3], there are two important principles to building a multi-CDN plan:

1. Own your origin
2. Know your traffic

In the first principle, they recommend that companies build their content origin (i.e. the source of the content used by the CDN) outside of any given CDN, and minimise the number of provider-specific CDN features that they use. Some content providers we spoke to had indeed chosen to perform functions like transcoding and packaging in-house rather than using CDN-provided equivalents. Yet there is a trend towards placing dynamic services such as personalisation and advert insertion within the CDN edge (see Edge Computing, below). So abstracting from each CDN's features may not be so easy.

The second principle is more well agreed upon. It's important to understand the key performance indicators that matter to your business, and choose partners who can perform best against these. For live sports, low latency may be a driving factor; while for VOD, high volume at low cost may be more important. Capacity planning becomes critical for optimising performance and cost, and this is especially true when balancing multiple CDNs.

Most of the CDN providers also agree that multi-CDN is the best option for their customers. However, some did caution against the *lowest common denominator* effect. If a content provider uses multiple CDNs, are they excluded from using any value-add services offered by one provider, since they're not available on all CDNs? This is the reason Citrix recommended keeping intelligence at the origin, but there are other options emerging. One CDN told us that "We try to make our intelligent services independent of our content delivery", so that they can still be deployed in multi-CDN architectures.

There are also CDN aggregation platforms available, such as Eurovision Flow. These offer an abstraction layer that allows the content provider to take advantage of multiple CDNs without having to manage each relationship individually. Other CDN vendors offer their own multi-CDN aggregation functionality, such as origin shielding. This allows one CDN to provide an origin to other CDNs, avoiding each CDN sending requests back to the content provider's own origin servers.

3 https://www.citrix.com/content/dam/citrix/en_us/documents/white-paper/best-practices-for-evaluating-and-implementing-a-multi-cdn-strategy.pdf

BUILD YOUR OWN

If CDN capability is so important, should the larger content providers build their own capability? Netflix certainly think so, as they've invested significantly in their own CDN, Open Connect. Much of the reasoning for this was to allow them to control their caching rules in a highly proactive way. They push content into caches based on expected usage, as opposed to using the demand-driven pull model that is common elsewhere.

Netflix have installed custom server appliances in internet exchange points in their most significant markets, and also make them available within some ISP's networks. As a VOD-only platform, they are able to pre-push content to these cache devices, often before it becomes available to users. This optimises the network use for high-demand content.

The BBC have also created their own CDN for their domestic market. Working with some of the UK's largest ISPs, they are able to place caches within the ISP networks in order to reduce costs and optimise efficiency. "It's very much part of a multi-CDN strategy", they explained, working alongside commercial CDNs to serve content to specific device types that they have optimised for.

Other content providers have invested in streaming specialist companies instead. The highest profile example is Disney's acquisition of BAMTech, the streaming services company that was originally spun out of MLB Advanced Media.

So will other content providers follow suit, and build their own CDNs? The investment required means that it's only viable for companies delivering content at large scale, and the opportunity also varies regionally. One global content provider told us they had no plans to build their own, preferring to use the best partners for each country in which they operate. But another explained that they are considering it, in countries "where market CDNs don't meet our needs".

When you launch an OTT service, you don't necessarily know where your viewers will be.

This kind of global planning sits at the heart of investment decisions around CDN. "When you launch an OTT service, you don't necessarily know where your viewers will be," said one contributor. "At the start, market CDNs give you flexibility. As you grow, you may want to build your own."

ARE CDNS THE ANSWER?

Content delivery networks will undoubtedly continue to sit at the heart of video distribution. However, tomorrow's CDNs may look different to today's. With new entrants in the market, and new capabilities being offered by existing players, CDNs will continue to evolve. Over time, we expect to see more CDN nodes within ISP networks, pushing caches closer to the end users.

The largest global and regional content providers may build their own CDNs, but this won't be necessary for most; and indeed those that do will continue to use third party CDNs too.

Perhaps the most significant evolution of the CDN will be the greater availability of edge computing capability, as discussed in the next section.

Edge Computing

Modern streaming services do more than just serve video content to end users. They may need to authenticate those users, adapt the stream to the receiving device, insert personalised advertising, adjust the bitrate to account for changing network conditions, or otherwise tailor the stream for the end user. In some circumstances, some of the logic to enable these features is placed on the client device, but often it is advantageous to make these adaptations at the server side.

Therefore, one way to reduce latency and improve the streaming experience is to perform this complex server-side processing closer to the user, rather than at the origin. This requires distribution of computing capabilities around the network; commonly referred to as Edge Computing.

Figure 2: Simplified Network Architecture

It is important to realise that there is no single definition of where the 'edge' is. For many, it means the edge of the CDN; in this case the capability is located in the CDN node. These nodes are often located at internet exchange points, though in some cases they are within the ISP network, at their core sites or even the Metro sites.

There is no single definition of where the 'edge' is.

Many of the CDN providers are working with the ISPs to deploy more nodes within their networks, in order to reduce distribution costs and reduce latency between the end user and the CDN. One CDN explained that, "The connection between the ISP and the internet backbone is a bottleneck." Putting a cache inside the ISP will help alleviate that.

As telcos modernise their infrastructure, many are removing legacy telephony equipment from telephone exchanges. Some CDNs have seized this opportunity to deploy edge equipment closer to the user, but it's not always a simple swap. Computing resources require more power and cooling than telephony equipment.

Some networks push the 'edge' even closer to the user. 5G networks can be deployed with computing capability located at the cellular base station, for example. Of course there is a trade-off to be made between the complexity and cost of deployment, and how close to the end user the capability needs to be. One contributor cautioned that like all good things, edge computing should be used in moderation. "The more complex you make the delivery process, the more inefficient it becomes."

FUNCTIONALITY AT THE EDGE

Why run computing functions closer to the end user? The most talked about use case is personalised advertising. The viewer may not always realise it, but every advertising break they see when watching content online involves a complex set of calculations and transactions that result in delivering the most relevant adverts, while maximising the value of the content provider's advertising inventory.

But targeted advertising is just one example of the ways in which a stream can be personalised to the individual viewer. Different devices require different video packaging formats, and different network conditions require different media streams. Different users also have different rights to view content, requiring functions like subscriber authentication and Digital Rights Management (DRM).

Performing these operations at the CDN edge "saves you a few hundred milliseconds for every operation, compared to backhauling the traffic to the central cloud," said one CDN. The aggregate effect of that can be a significant performance increase. It can enable use-cases that simply aren't practical in a centralised infrastructure.

At a high level, adaptive IP streaming works by making available a combination of a *manifest* file, which instructs the player where to find the content, and many *video chunks*, which each contain a small amount of content (usually a few seconds). With edge computing in the CDN, each user's manifest file can be personalised; dynamically edited for them at the moment they request it.

With edge computing in the CDN, each user's manifest file can be personalised.

Some modern networks even cache the media in a componentised form – with separate audio, video, and subtitles that can be combined into the relevant version for any given user. These components are then dynamically packaged in real-time for delivery to the user. This means that fewer individual representations of the content need to be delivered to each cache, optimising the use of the network.

MANAGING DISTRIBUTED COMPUTING

To achieve this vision of computing at the edge of the network, we need both hardware to be deployed to perform the computing functions, and software to manage it. In order for a content provider to deploy their functions at the edge of a CDN, they need a standard mechanism for doing so.

A number of approaches exist for this today. Some CDNs offer the use of containerisation tools such as Docker, and container orchestration systems such as Kubernetes. This essentially offers public cloud like services, closer to the edge of the network. It gives a high degree of flexibility, at the expense of a certain amount of overhead.

Other CDNs use a language such as Varnish Configuration Language (VCL) to dynamically load code into a shared infrastructure, essentially offering serverless code execution. This creates a highly dynamic system in which the code author (e.g. the content provider) is abstracted from the execution platform, offering simplicity at the expense of placing some constraints on the code that can be run.

Optimising the efficiency of this edge computing is a huge focus for many of the CDNs. One is building out a new edge computing capability using Web Assembly. This new language became a World Wide Web Consortium (W3C) recommendation in December 2019, and enables the creation of portable binary code that can be run in web browsers. By running it on a modern edge compute platform, they are hoping to achieve significant performance improvements, saying that "Our design goal is for the response to any request to leave our network within 1ms."

Our design goal is for the response to any request to leave our network within 1ms.

IS EDGE COMPUTING THE ANSWER?

There will continue to be a number of edge computing options available, both from different providers, and increasingly from each individual provider. Just as Amazon, Microsoft and Google all offer both virtual machines and serverless functions, so CDNs and network providers will offer edge computing options to suit a range of needs.

The growth of edge computing will be significant in the coming years. It may serve to improve distribution efficiency, and therefore help with scaling. But the main benefits will be improving latency and enabling new personalisation features.

Multicast

Multicast has existed for a long time, yet its use in content delivery has always been niche. As one interviewee put it, “Multicast hasn’t delivered on its promise yet because we haven’t really needed it.”

Multicast hasn’t delivered on its promise yet because we haven’t really needed it.

Multicast is a technique for one-to-many data distribution over a network, as opposed to the one-to-one model of unicast. With unicast, a server fulfils requests from an individual client by delivering data packets back to that client. If multiple clients request the same data, it is delivered multiple times. In a multicast model by contrast, one data packet may be received by multiple clients, reducing data duplication.

This clearly has advantages in bandwidth usage and cost. With multicast, these metrics scale with the number of different video streams (e.g. TV channels), rather than the number of viewers. One provider estimates that for an ISP, the cost-per-viewer of multicast is lower than unicast, provided that there are above around 500 viewers per channel.

Multicast is not without complexity, however. Network routing devices must direct the multicast packets to the parts of the network where the intended clients reside, without flooding the rest of the network with unnecessary data. In a network the size and complexity of the internet, there is potential for failures in multicast routing to cause considerable congestion. Meanwhile, the requirement to have multicast supported, enabled, and properly managed across all routing devices on the network has long been seen as a blocker in the adoption of multicast over the open internet.

To achieve end-to-end multicast distribution would also require support not just in the professionally-managed equipment at content providers, ISPs, and CDNs; but also in consumer devices. Generally speaking, it cannot be assumed that TVs, mobile phones, tablets, and other consumer devices support multicast video.

CLOSED IPTV SYSTEMS

Nonetheless, some ISPs use multicast to great effect within their own, controlled networks. BT, for example, offers its customers a hybrid broadcast and IP service based on the *YouView* platform. The set-top box's electronic programme guide (EPG) combines channels delivered by Digital Terrestrial Television (DTT) with those delivered by IP. Multicast is used to deliver those IP channels, ensuring that bandwidth is used very efficiently in BT's network.

Some ISPs use multicast to great effect within their own, controlled networks.

But of course, the same content is also made available to mobile devices via an app, which uses unicast distribution via CDNs. The comparison in bandwidth usage is stark. For a single sporting event, BT's network carries tens of megabits per second of data to their set-top boxes, whereas the unicast clients create a load of terabits per second.

As more ISPs launch content offerings, we may see more of these kind of deployments. Commercial multicast products aimed at IPTV operators are now available from CDN vendors. These offer the server-side software to generate and manage multicast streams, and Software Development Kits (SDKs) for client devices, such as set-top boxes.

Should all ISPs implement such solutions? There is, after all, considerable investment required to deploy multicast video services. They may need new routing capabilities, and possibly to build separate, traffic-managed VLANs (virtual local area networks) to carry the multicast traffic. And they will need to ensure that all consumer gateway devices support multicast. For an ISP with millions of customers using legacy home gateways, that could be very costly indeed.

There are additional capabilities required at the application layer too. Multicast uses UDP rather than TCP, so many of TCP's traffic management and error protection features must be implemented and managed at a higher layer. Capabilities like retransmission must be implemented by the TV application, using unicast.

So for an ISP with an IPTV product, the ability to generate revenue from this offering may make the investment worthwhile. But for those who don't have content at the heart of their customer offering, the business case simply may not stack up.

MULTICAST ADAPTIVE BITRATE

In 2018, the Digital Video Broadcasting Project (DVB) released a draft document entitled *A176 Adaptive Media Streaming Over IP Multicast*[4]. It seeks to address the demands of delivering the same content simultaneously to many viewers, but without multicast support being required in every receiving device.

It uses multicast packet replication at Layer 3, paired with the same adaptive bitrate (ABR) media encoding and packaging formats as used in unicast delivery. Thus it attempts to offer the best of both worlds.

In order to achieve compatibility with the greatest number of receiving devices, the DVB architecture uses a *multicast gateway*, which receives multicast content and serves it via unicast to end devices. The gateway could be implemented in a device such as an ISP's home gateway, or a television set-top box; alternatively it may be deployed at an edge point within the network, making the architecture essentially invisible to devices in the consumer realm.

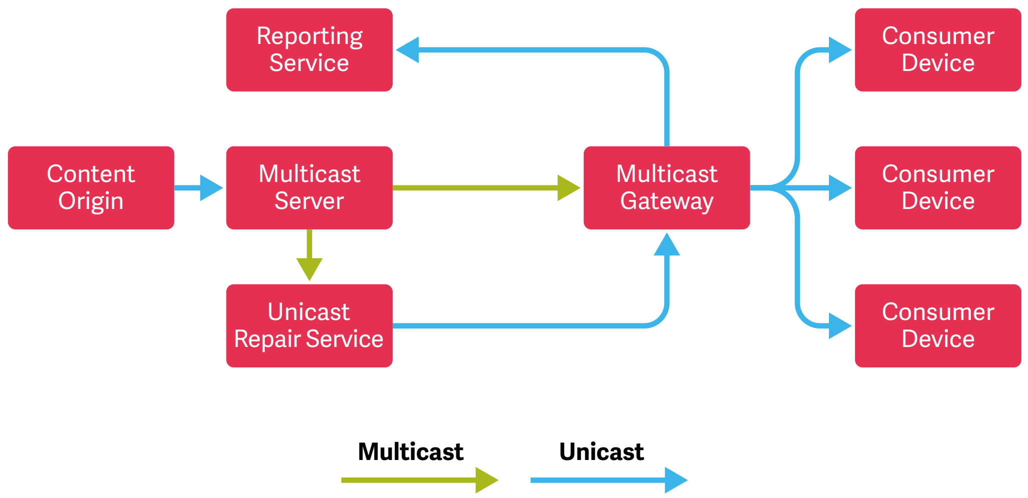

Figure 3: Simplified Multicast ABR Architecture

The multicast gateway can handle complexities such as *unicast repair*, a process by which any missing or damaged packets from the multicast stream can be recovered by using a unicast request back to the content source. It also manages service reporting, providing metrics and reporting back to the content provider. This is critical to building an economically viable content delivery proposition; as one broadcaster put it, "We wouldn't invest in any platform where we couldn't understand a lot of detail about viewers' behaviour".

4 https://www.dvb.org/resources/public/standards/a176_adaptive_media_streaming_over_ip_multi-cast_2018-02-16_draft_bluebook.pdf

IS MULTICAST THE ANSWER?

While opinions varied somewhat, and the situation differs in different countries, overall we must conclude that it is unlikely to happen at a significant global scale. It is safe to conclude that it will continue to be valuable for ISP's own IPTV solutions within their networks. But will it become economically advantageous to use it in a more open way? While opinions varied somewhat, overall we must conclude that it is unlikely.

One content provider told us that they believe multicast should be on the agenda for anyone delivering large scale live content. And multiple contributors told us they are planning multicast proof of concept trials over the next twelve months, so there does continue to be interest. Yet one broadcaster explained that they've approached ISPs to discuss multicast trials, and had little interest from them.

Multicast delivers benefit when multiple clients want to access the same content at the same time. So content providers who focus on VOD have little to gain. As one interviewee explained, "I've been a big proponent of multicast in the past. If you believe linear channels are the predominant method of viewing, then it makes sense. But linear channels originally exist because of finite bandwidth in the broadcast spectrum. We have a new generation of viewers with different habits. Building out multicast networks to support legacy delivery makes no sense."

> **We have a new generation of viewers with different habits. Building out multicast networks to support legacy delivery makes no sense.**

5G

The next generation of cellular networks, 5G, has created considerable excitement in the media industry. In content production, it has potential applications for video contribution, wireless cameras, low-latency control systems, and more. However, in this report we'll examine the potential of 5G for distribution.

> **5G is not one technology.**

It's important to understand that 5G is not one technology. It's a range of technologies to build and manage mobile networks, and indeed the standards are still being actively developed. It offers higher data rates, lower latency, and new functionality. The 3GPP (3rd Generation Partnership Project) publishes the 5G specifications, and the latest releases target use cases such as low latency emergency services communications, as well as media distribution.

We can consider two primary effects of 5G, as relevant to video delivery. The first is in providing a 'bigger pipe' for regular OTT delivery. That is to say, more bandwidth is available for general-purpose internet connectivity. The second is the option of broadcast or multicast functionality, which could even be used as a replacement for traditional TV broadcasting via mechanisms such as DTT.

FASTER INTERNET

The potential benefits of 5G have been well reported; gigabit speeds and millisecond latency. In principle of course, the service offered by 4G is perfectly sufficient to deliver good quality HD video to mobile devices, and even UHD when using modern codecs. However, with highly contended networks delivering more and more high bandwidth content (whether video, gaming, or other applications), we often see that in densely populated areas, connection speeds don't live up to expectations. As a result, the additional bandwidth of 5G will be very welcome indeed.

It is in these densely populated areas where millimetre wave cells will most likely be deployed. Millimetre wave communications offer the highest bandwidth and lowest latency, but have a much shorter range than 4G signals, and do not penetrate walls and objects so well.

As a result, millimetre waves will be most widely used in areas where many small antennas can be deployed, such as in cities. Mid- and low-band 5G frequencies offer a greater range, albeit with lower speeds. They will be used in combination with millimetre wave to make up most 5G networks.

In countries like the UK and USA, where fixed-line broadband is widely available, 5G will be used first and foremost for mobile connectivity in the coming years. Yet already there are home broadband options available using 5G, and it's expected that this will become more common. Many predict that 5G wireless connectivity will be a cost-effective way to provide broadband to those who don't have good fixed-line connections today. Governments around the world are investing in ensuring that their populations have access to good quality internet connections, and 5G may play a significant part in that. As one contributor put it, "By 2030, not having access to broadband will be as unacceptable as not having access to water."

By 2030, not having access to broadband will be as unacceptable as not having access to water.

Already in countries such as India, where the historic fixed-line network may not be as strong, cellular data connections make up an important part of the broadband infrastructure. In these countries, 5G has the potential to bring viable online video streaming to millions of new users.

BROADCAST

As well as simply offering faster data connectivity, 5G offers the potential for broadcasting or multicasting video content.

Broadcasting over cellular networks is not a new idea. 3GPP's *Evolved Multimedia Broadcast Multicast Service* (eMBMS), also known as *LTE Broadcast*, has offered broadcasting functionality on cellular networks for almost a decade. Yet while there are some deployments – one of our CDN contributors has a large-scale trial ongoing in India – it's never taken off in a big way.

The technology continues to evolve, however. More recent 3GPP specification releases include *Further Enhanced Multimedia Broadcast Multicast Service* (FeMBMS). There's also considerable industry effort around promoting what are more generally referred to as *point to multipoint* (PTM) distribution systems over 5G. The European 5G Public Private Partnership (5GPPP), for example, runs the 5G-Xcast project, promoting the development and deployment of broadcast and multicast PTM capabilities over 5G.

So why has cellular network broadcasting not been more widely used? First, the mobile network operator (MNO) has to make significant investment, with compatible equipment installed at every radio transmitter. Second, and perhaps more crucially, is that the receiving device must also support the standard. As is so often the case, the economics are more important than the technology here. We must ask – what is the incentive for device manufacturers to add support to their products?

One of our experts was blunt in their assessment; "Apple don't support LTE Broadcast, so it's dead in the water." That perspective certainly holds weight given Apple's strong position in mobile devices. And there's clear precedent; for many years, the iPhone had the technical capability to receive FM radio, due to its inclusion in the multi-purpose communication chips installed in the devices. Yet no iPhone has ever enabled support for FM radio reception.

5G NETWORK ARCHITECTURE

There's a large unsung benefit of 5G, and it is not related to the wireless technology at all. It's the chance to build a new and more architecturally advanced network, both in the way it interconnects to the internet, and in the deployment of edge infrastructure.

The interesting thing is how 5G changes the last mile architecture.

"The interesting thing is how 5G changes the last mile architecture," explained one CDN. 5G networks generally have far more backhaul connection points compared to 4G networks. To understand this, let's consider a user in Provence, making a connection to the internet via 4G. They will have their traffic routed via Paris before it leaves the mobile network, because the 4G network likely only has one or two *egress points*, where data leaves the MNO's network and passes onto the internet core. By contrast, it's expected that 5G networks will have connection points in every major metropolitan area, rather than only one or two globally.

5G and edge computing also have the potential to go hand-in-hand. Five different contributors – including CDNs, network operators, and content providers – said that the most significant aspect of 5G for video delivery is the ability to place caches or computing resources within the mobile network. The benefits of this are discussed further in the *Edge Computing* section, but in short this enables reduced latency video delivery, and opens up new opportunities for scaling and personalisation.

One telco told us that they're working with CDNs to investigate putting caches in the cellular base stations, as well as upgrading the *backhaul network* (that is, the connectivity from the cellular base stations to the rest of the network, and the internet).

IS 5G THE ANSWER?

For 5G broadcasting to become practical would need cooperation between broadcasters and MNOs. There is a broad range of views on whether this will happen. The European Broadcasting Union (EBU) has launched the *5G Media Action Group* (5G-MAG) to promote the adoption of 5G broadcasting features in networks and devices. However, our contributors were less optimistic, with one suggesting that "5G broadcasting will only happen if there's regulatory intervention."

5G broadcasting will only happen if there's regulatory intervention.

Others questioned not just whether 5G broadcasting *will* happen, but whether it *should*. "We're assuming that we're saying goodbye to broadcast and hello to IP." So why replace existing broadcast networks with new broadcasting systems? "5G is more about new consumption opportunities than replacing existing methods."

Those new opportunities include data connections to vehicles – which stand to become new venues for content consumption as they become driverless – and low latency gaming experiences. These undoubtedly have the potential to be significant.

We cannot confidently predict widespread adoption of 5G broadcasting, but we can confidently predict the growth of video viewing over 5G data networks. And there is great cause for optimism in 5G networks being a key enabler for other technologies like edge computing.

Peer to Peer

Peer to peer (P2P) delivery split the opinions of our experts more than any other technology. Some were skeptical, while others felt that it could be a useful way to moderate the level of investment required in scaling the ISPs and CDNs.

In theory, the technology removes load from the CDNs, so the CDN providers could naturally be expected to have mixed views. One has recently invested in acquiring a peer to peer specialist, while another has placed investment in the technology on hold. A third told us that, "Peer to peer could be useful. But historically it's created more problems than it's solved."

In a peer to peer distribution model, each client device can share the data it has with other client devices in the network (its peers). In other words, receiving devices become not just consumers of the content, but also suppliers of it. Peer to peer was popularised by file-sharing systems like *BitTorrent*, though there are now commercially available P2P systems for distributing video content based on WebRTC, an open source system for real time communications over the internet.

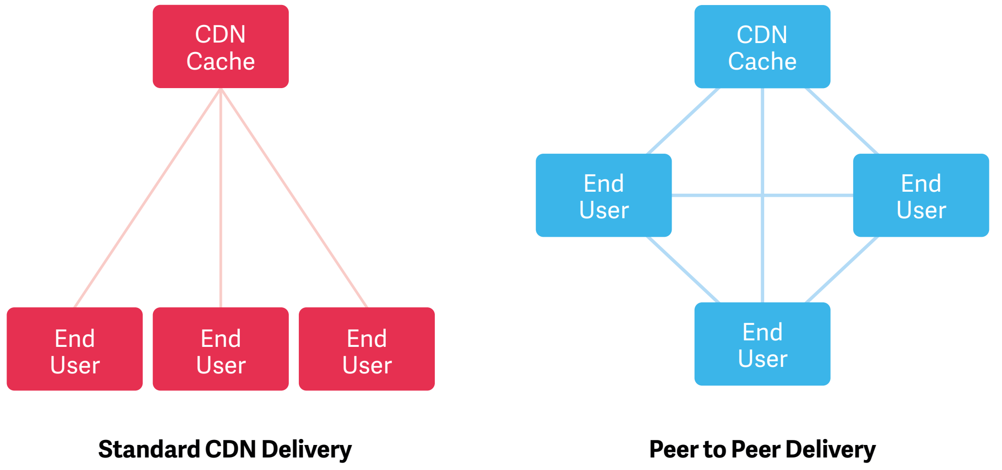

Figure 4: Simplified Peer to Peer Architecture

For delivery of premium video, peer to peer architectures are deployed in conjunction with one or more traditional CDNs. A CDN provider explained, "It acts as an overlay to CDN delivery. It helps offload the CDN, but it doesn't remove the need to have a well scaled CDN and network."

It's this *offload*, removing data load from the CDN, that offers potential gains. The amount of data that can be offloaded is very sensitive to factors such as which content is being viewed by which users, and where they're geographically located. But if those factors align, realistic figures of 50% offload can be achieved.

To implement P2P, there are no changes required to the CDN, or to networks over which the content is delivered. From the point of view of the client device, the P2P network appears much like any other CDN. A software development kit (SDK) is used in the player to communicate with the P2P system's back-end servers, and together they manage decisions about which source to retrieve content from – a CDN node or a peer. The decisions can, in fact, include not only peers and a CDN, but multiple different CDNs, supporting a multi-CDN strategy.

For the content provider, there is some additional overhead of including a third party SDK in their player, but little else to change.

Peer to peer networks have some very particular characteristics. Notably, they actually become more efficient as the number of peers increases. This is because the number of *seeds* (i.e. peers which can serve content) goes up as the number of clients goes up. This seems like it ought to be the holy grail for dealing with increasing scale, so why does it provoke mixed responses?

LATENCY AND UPLOADS

Peer to peer technology is particularly of interest in managing distribution of large live events, because these cause a spike of viewers watching the same content at the same time. Yet such live events are often extremely sensitive to latency. Nobody wants to hear their neighbours cheering for the crucial goal in a football match seconds before they see it on their screen. So one of the primary concerns expressed by our contributors was that peer to peer can introduce additional latency.

The packet path from me to my neighbour is actually quite long.

Of course, before a chunk of video can be delivered from one peer to another, the first peer has to have downloaded it. Additionally, the delivery from one peer to another can add further latency. An ISP explained, "The asymmetry of most consumer connections is a problem. The packet path from me to my neighbour

is actually quite long." In most cases, the network traffic from one household to another would be routed a long way back into the ISP's network to the Broadband Remote Access Server (B-RAS), before being directed onwards to a neighbour. "That's only one step away from where the CDN cache is, so the roundtrip time will be worse than CDN delivery."

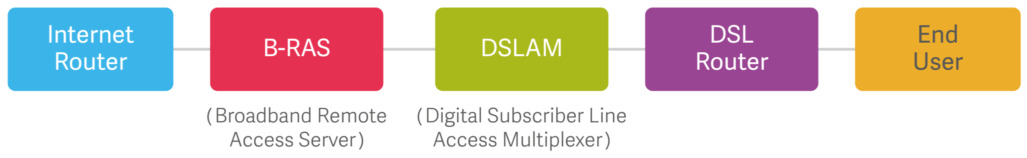

Figure 5: Simplified ISP Architecture

Another key characteristic of P2P is that clients not only download content, but also upload it to other clients. This means that the speed and cost of users' data uploads are important, unlike in regular CDN delivery.

In many countries, most fixed-line broadband connections are asymmetrical, meaning that their upload bandwidth is lower than the download bandwidth. While Fibre To The Home (FTTH) connections generally offer symmetrical data transfer, those running over copper telephony cabling are asymmetrical, including Asymmetric Digital Subscriber Line (ADSL) and Fibre To The Cabinet (FTTC). Cable connections are also generally asymmetrical.

In areas where data use is metred, uploading data over the users' connection can also cost them money. Some content providers are concerned that if a user receives a bill for the upstream bandwidth use, they may blame the content provider, causing reputational damage. This all depends on the prevailing internet charging models in the countries of operation, but it certainly adds complexity to the issues that must be considered.

Available peer to peer systems offer controls to attempt to manage these issues. The P2P network will adapt to each user's upload speeds, with the SDK on their device monitoring this in real time. Content providers can often configure default settings around things like upload speed caps, with users able to further tune these if required. An example is that uploads are often disabled when a cellular connection is in use, but enabled when connected via Wi-Fi.

Of course, many of these factors vary geographically. Availability of CDN capacity, dominant types of connectivity, and pricing all vary. In countries with lower CDN availability, or higher prevalence of symmetrical internet connections, P2P has greater benefit.

SHARING AND SECURITY

As ever, technology is only half the story. Like it or not, peer to peer technology has many cultural associations with illegal file sharing. As one contributor put it, "The main challenge with P2P is the legacy of the terminology."

The main challenge with P2P is the legacy of the terminology.

But it isn't just a challenge of individuals' perceptions. Restrictions on the use of peer to peer distribution exist in a number of content licensing agreements. "The biggest challenge is rights holders' terms and conditions," explained one content provider, while a CDN had analysed the terms used by major US studios and found that while some had no issues, others explicitly prohibited peer to peer sharing. These problems are not insurmountable, of course – contracts can be changed – but they are a factor when considering deploying P2P today.

Historical concern about peer to peer sharing has also resulted in some ISPs "traffic shaping" this content. That means that P2P traffic would be deprioritised on the network compared with other traffic. Implementations based on open technologies like WebRTC may alleviate this, of course. Nonetheless, one contributor with experience in P2P solutions was clear that as with most technologies, it's important to ensure the right relationships and dialogues are in place with ISPs.

Some of our contributors also expressed concern at the security implications of users having the content *at rest*, which is to say stored on their devices. This introduces additional risk points for piracy. However, the commercially available P2P solutions take security extremely seriously. Generally, content is only stored in RAM, and interoperability with Digital Rights Management (DRM) solutions is carefully tested to ensure these work as expected. Additional security features like geographic security can also be implemented.

IS P2P THE ANSWER?

There have been successful uses of P2P technology at scale. French broadcaster TF1 used peer to peer to help them scale their streaming delivery of the 2018 World Cup. According to Streamroot, up to 4m viewers were provided with an average 3Mbps stream by the system, which amounts to considerable offload of the CDNs.

Multiple content providers told us that they were planning or considering trials of P2P delivery, so there continues to be some interest, though none expressed any urgency. "It's on our roadmap, but there's no imperative," explained one, while another told us that they would trial it after looking into multicast, as they felt that has a bigger potential impact.

Another contributor told us that "Peer to peer lacks a killer app. It improves things slightly, but it's not a game changer." Interestingly, ISPs were especially unexcited by P2P. One stated that, "Peer to peer doesn't add any value to the ISP", while another does not see it being used in the consumer domain. However, one of our CDNs pointed out that their mesh architecture (in which each CDN cache can retrieve content from any other CDN cache) is an application of P2P on the server side.

Peer to peer lacks a killer app. It improves things slightly, but it's not a game changer.

P2P is also being consididered outside of video delivery, for example for large scale gaming software downloads. After all, software updates often cause bigger peaks than streaming video. If P2P delivery models can be successfully developed and deployed in this space, there will be additional capacity available for video.

In video delivery, the range of technical and political concerns mean that peer to peer still has a difficult road ahead. As advances in technology solve some of the historic issues, mindshare may gradually change. Indeed, there's some evidence of this happening in recent PSB tenders looking at P2P specifically, in France and Spain. So we do expect that P2P will continue to grow as an overlay for CDN distribution, and it will be used for some particular use-cases, such as certain large scale live events.

Nonetheless, we expect P2P to remain relatively niche. As one contributor quipped, "It's a useful insurance policy." Niche technologies certainly have a meaningful impact when they're used in the right way, but we do not expect P2P to become the dominant delivery mechanism in the coming years.

ATSC 3.0

When considering the transition from broadcast to IP delivery, one tends to think of delivering content over general purpose internet connections, also known as over the top (OTT). Or one might consider dedicated IPTV services, such as those discussed in the *Multicast* section.

So it may seem strange to consider an over the air (OTA) broadcast technology in a report about IP delivery. But when a new broadcast system uses IP as part of its over the air data stream, it merits a brief exploration.

It may seem strange to consider an over the air (OTA) broadcast technology in a report about IP delivery.

ATSC 3.0 is the next generation over the air TV standard for countries including the USA and South Korea. Its OTA broadcast mechanism in fact uses IP as its foundational protocol, delivering video in chunks rather than as a continuous bitstream, such as was used in previous broadcast standards. ATSC 3.0 further allows for hybrid content delivery, mixing OTA and internet delivered content on one receiving device.

By offering more robust support for reception of OTA signals on mobile devices, ATSC 3.0 also aims to alleviate the demands on mobile data networks, by enabling mobile viewers to watch live content using a broadcast signal.

Proponents of ATSC 3.0 aim to see the standard integrated into general purpose communications chips, much like wi-fi has been. However, the same concerns exist here as were expressed for 5G broadcast; will device manufacturers integrate receivers and make their functionality available at a software level? As was discussed in the 5G section, key manufacturers like Apple have a history of favouring internet delivery over broadcast for consuming content on their devices.

One possibility to aid adoption in the home is that, much like with IP multicast, there is the possibility of a gateway to support devices which are not compatible. LG have demonstrated a 'smart antenna' that can receive ATSC 3.0 signals and offer unicast IP streams to devices on the same network.

IS ATSC 3.0 THE ANSWER?

There is considerable support for ATSC 3.0 in the broadcast markets in which it will be deployed, and it will certainly have some success in modernising broadcast services in those countries. But it is much less certain that it will gain widespread adoption for mobile viewing, or that it will meaningfully change the trend of viewer behaviour towards watching content via internet connected devices instead of broadcast receivers.

It has been evident for some time that industry trends are increasingly driven by consumer behaviour, rather than by content providers. One of our contributors felt that this particularly applies here, noting that ATSC 3.0 is a development of the broadcast industry. As they explained, "There's a fundamental shift from broadcasters defining the technology, towards consumer electronics manufacturers setting the direction." Another contributor was more direct, saying that when it comes to keeping broadcast relevant as viewers move online, it's "too little, too late".

There's a fundamental shift from broadcasters defining the technology, towards consumer electronics manufacturers setting the direction.

Quality of Experience

Whatever technologies are being used to deliver video, a crucial aspect of any distribution mechanism is effective management of Quality of Experience (QoE). As we've seen, content delivery over IP is complex, with many different technologies and organisations in the path between the content provider and the viewer. So understanding, monitoring, and optimising the delivery is crucial.

It's not just the content providers who care deeply about this challenge. One ISP explained, "People don't measure us by our download speed. They are more interested in whether they can do what they want to do. Customer satisfaction is impacted by whether streaming services work well."

People don't measure us by our download speed. Customer satisfaction is impacted by whether streaming services work well.

In a broadcast distribution chain, there is generally parity in quality for all viewers (provided their aerial, satellite dish, or receiving equipment is correctly installed and configured). In IP, outlined one CDN expert, "There's much more variability in the network congestion, utilisation, and architectural characteristics at every point of the delivery path." This can make it difficult to isolate problems; one contributor posed the question, "If a user is having buffering problems in Leeds, but it's fine in London, where's the problem?"

A number of companies offer specialist solutions for video QoE monitoring and management, and there are a small number of recognised leaders. However, across the industry there is a lack of commonality in the metrics used. One content provider explained that they've defined their own set of metrics and targets, which they monitor across each of the CDNs they use. "We're trying to define what an acceptable viewer experience is, and then any CDN that meets that can be considered on price." Some – including two of the CDNs we spoke to – have called for common industry agreement on which are the most important QoE metrics and targets.

It's easy to see why. Even defining what QoE means can be complex. Slow playback start, rebuffering, jitter, unexpected disconnects, latency, reduced bitrate due to congestion; all these issues and more will affect the viewer's experience. Which are the most important may also depend on the content provider or content type. In particular, the need to reduce latency may be critical in some cases (e.g. live sport) but not in others (e.g. drama); meanwhile buffering or disconnection will not be popular for any content!

As well as deciding what to measure, one must consider how and where to measure it. There are of course many points in the chain at which monitoring can take place. Popular solutions use client-side monitoring, in which the player monitors the QoE, adapts playback accordingly, and reports back to a monitoring server as required. In unicast delivery, the server knows every request that is fulfilled to every client, so server-side monitoring is also common. Indeed, the availability of real-time data feeds or logs from the CDN is a selling point in some cases. One of our contributors, Touchstream, also offers cloud-based services to proactively monitor streams across different CDNs, giving the content owner a balanced view of how their different CDNs are performing.

Whatever is monitored, this data must be acted upon. When architecting an IP streaming service, it's important to build redundancy into the solution, and real-time monitoring can allow for automatic failover should an error occur. But it can also be used to adapt and improve performance in real-time. Many video players now use mid-stream switching between different CDNs (including P2P networks) to continually optimise delivery based on the playback conditions.

Many video players now use mid-stream switching between different CDNs to continually optimise delivery based on the playback conditions.

In the coming years, it will become increasingly important that vendors across the distribution chain offer clear metrics and logging in real-time, and that these can be combined together into an end-to-end view. Being able to pinpoint issues and respond to them rapidly will be key to achieving a QoE that is equal to or greater than broadcast.

More Pieces to the Puzzle

The next decade will continue to see rapid evolution of the technologies used to deliver video over the internet. The basic building blocks we have today – ISPs, CDNs, adaptive bitrate streaming, and others – will still exist, but they will themselves evolve, and in some cases be augmented.

But as we've identified, there are many complexities to delivering video over IP. While this report has examined technologies used for transporting video to the user, there are many other factors which must also be considered when planning an IP video service.

ENCODING AND PACKAGING

Before it can be distributed to the viewer, content must be encoded and packaged. Codec technology is perhaps one of the areas of most rapid technological progression. An exploration of the newest video coding options – and the licensing situation that accompanies them – would be a large enough topic for its own report. For this paper we'll settle for saying that in the next three to five years we expect to see the next generation of codecs gain widespread adoption, offering the same quality at lower bitrates.

> **An exploration of the newest video coding options – and the licensing situation that accompanies them – would be a large enough topic for its own report.**

This will help to alleviate congestion in the networks of course, though history tells us that just as we make bandwidth savings, greater availability of higher quality services will offset those savings. While reduction in bitrates for SD and HD video will be helpful, new codecs will also be an enabler for more 4K and even 8K content.

A large focus in 2019 was reducing the latency of live IP distribution. Optimisations in network architecture and caching policies certainly help this, but the greatest change has been the availability of new packaging formats like Low Latency HLS and DASH.

SUBSCRIBER MANAGEMENT

Of course, it's not just video and audio that make up a streaming service. Billing, user management, advert insertion, and other transactional services need to be able to scale for peak demand too. A major sporting event, for example, may generate sudden peaks in a streaming service's signups in the 24 hours before the event, and logins in the minutes leading up to the start.

Despite being a huge success overall, the launch of flagship streaming service Disney+ suffered much-publicised problems on its launch day. According to web monitoring site DownDetector, over a quarter of these were related to simply logging in, rather than streaming video.

And after all, it's no use having perfect video delivery if the user can't sign in to the service in the first place.

PLANNING FOR FAILURE

It's also critical to carefully plan the way in which all of these technology components work together to create a reliable and resilient distribution medium. "It's a lot more complex with IP," said one contributor. "You need to test, test, test."

> **It's a lot more complex with IP – you need to test, test, test.**

When building a broadcast distribution chain, we needed to consider the impact of any link in that chain failing. But when distributing video over the internet, a problem may occur even if every component is operating as expected. That's because in a shared network, issues like network congestion can come from sources that have nothing to do with video delivery.

For many countries, Christmas Day and Boxing Day (the day after Christmas Day) pose huge challenges, for example. TV viewing generally reaches high levels – including notable live sporting events. Yet at the same time, millions of lucky gift recipients are turning on their new phones, tablets, smart speakers, and other devices; all reaching out to download the latest software updates. Game updates are gaining a particular reputation as drivers of network congestion. As one ISP told us, "Some game providers stagger their updates. But some of the biggest ones don't, and that can overwhelm our network."

Meanwhile, others will be sitting down to play a game streamed live from a cloud gaming service such as Google Stadia. The bandwidth used for this can dwarf those used by most video streaming, with current profiles using up to 35Mbps to deliver a game in 4K HDR 60fps[5].

Some contributors spoke about the changing skills base required to design great IP distribution solutions. But at least one CDN found themselves impressed with the expertise on show in broadcasting, compared to other industries. "Broadcasters are very good at managing resilience and redundancy, and they plan for all eventualities." Reflecting on one broadcaster's team, who they'd worked with on a high profile sporting event, they commented, "I'm struck by the professionalism. How deeply they think about redundancy, fault tolerance, and failover."

Broadcasters are very good at managing resilience and redundancy, and they plan for all eventualities.

DELIVERING UNIVERSAL SERVICE

If we are ever to truly replace over the air broadcasting with IP-only delivery, then significant advances will need to be made in the universal availability of high quality internet connectivity. For Public Service Broadcasters (PSBs) especially, the ability to serve all of a country's population is an important – and often highly regulated – benchmark.

With any distribution mechanism, reaching the final 10% of the population is generally the most challenging. For example, the UK's first DTT multiplex launched in 1998 with 82 transmitters, reaching around 80% of the population. By the time the digital switchover completed in 2012, the coverage reached 98.5%, and today there are 1,150 transmitters – an increase of more than 14x.

Yet many governments around the world are committed to delivering 100% availability of broadband services. And, it could be argued, if this is achieved then it becomes much harder to justify continued maintenance of a separate broadcast network.

This is especially true when we consider that the costs of running broadcast networks will not reduce significantly, even if viewership dwindles. But this in itself causes a problem – until universal IP service is reached, broadcast networks will still need to be funded in many countries. Perhaps fewer channels or multiplexes will be maintained over the air, but the transmitter networks cannot be decommissioned until all services are removed. With even PSBs subject to increasingly rigorous value for money tests, governments will have some tough decisions to make about the right time to stop funding broadcast networks.

5 https://store.google.com/product/stadia?srp=/product/stadia_founders_edition

Looking to the Future

For established broadcasters, transition from one distribution technology to another is not new. The move from analogue to digital broadcasting was significant, yet it was in some ways simpler than the transition that faces us now. In each country, the industry generally swapped one analogue technology for one digital technology, whether ATSC, ISDB, or DVB-T. Broadcasters and their partners controlled the distribution networks, allowing them to plan for the required investment, and governments were able to lay out timetables for transition. Viewers, meanwhile, purchased new set-top boxes (STBs) or new TV sets to access the new services, but otherwise had a relatively unchanged experience.

Transition from one distribution technology to another is not new.

For IP, the situation is different. Viewer behaviour is changing; content is consumed on an array of different devices, many of which are general purpose and not specifically designed for viewing 'TV'. Even the large screens in our living rooms blend broadcast and IP technologies together. Broadcasters do not have control over the end-to-end distribution path, and bandwidth is now shared with websites, software updates, and gaming. Meanwhile, new entrants shape the market – one broadcaster told us that they keep a close watch on Amazon and Netflix's latest developments to help shape their own plans. That's because, they said, those companies can invest far more in technology development than traditional media companies can.

Perhaps most fundamental of all, however, is the pace of change. The DVB-T network in the UK is now over 20 years old, and has many years of life left in it. While it has continued to evolve, with the introduction of HD services using DVB-T2 and improvements to coverage, there has not been a need to adjust the architecture to adapt to viewing figures on a month-by-month basis. By contrast, it seems safe to assume that significant parts of this report will be outdated within 5 years. Content providers and platform operators will need to constantly adapt to changing opportunities, and they'll need to continually forecast their usage in order to deliver a truly broadcast quality service to their viewers.

Perhaps most fundamental of all is the pace of change.

One thing is for sure: no company involved in the distribution of content to consumers can afford to ignore the challenges of distributing over IP. Any who get it badly wrong will make headlines, and lose viewers. For those who get it right, the reward is that their viewers will never notice the delivery mechanism; the headlines will be driven only by the quality of the content.

10 Things You Need to Know

1 All-IP will happen. Slowly.

Technology change often follows an S-curve. New technologies take longer than we expect to take off, then adoption accelerates more rapidly than we predict. The curve finally flattens off again as reaching the last few percent takes much longer than we expect. Old technologies have a habit of staying around too; despite the prevalence of digital radio, FM is still going strong. We expect the same for broadcast TV, although an inflection point will ultimately be reached where broadcast is no longer viable.

2 Choose CDNs carefully.

Any significant OTT service will use multiple CDNs, balancing their cost and performance. But to think that all CDNs are interchangeable would be wrong. The CDNs will continue to evolve, with some investing in high-performance low-latency meshes of relatively few nodes, while others focus on deploying caches deeper into ISP networks, closer to the end user. Selecting the right partners for your use-case and the territories in which you operate is a critical decision.

3 Be on the leading Edge.

There is not a single definition of where "the edge" is, yet there does seem to be near unanimous support for doing more computing there. Better edge computing capability, creeping ever closer to the user, will define a new wave of low-latency hyper-personalised services. Whether it's in the ISP core or the 5G radio mast, the important thing is that it's a short network hop from your viewers, and much closer than a centralised cloud data centre. In fact, cloud architectures for content delivery will increasingly look old fashioned if they don't also include processing at the edge.

4 Multicast is not multi-purpose.

Multicast is an extremely useful, scalable technology for building closed IPTV solutions. As connectivity providers increasingly bundle content services, more of them will look to multicast. However, the technology will not become widely used in delivering Over The Top (OTT) streaming services via the internet. If you're delivering within your own network, multicast could be the key. But if you're delivering across the open internet, look elsewhere.

5 5G is all about more data.

It may be hard to overstate how prevalent 5G connectivity will become for accessing the internet, and that of course includes video streaming services. And its specialised functionality will unlock new opportunities in autonomous vehicles, gaming, and even content production. Nevertheless, 5G broadcasting still seems to lack a commercial incentive for Mobile Network Operators and device manufacturers, and so the use of 5G in content distribution will primarily be to provide more internet bandwidth, plain and simple.

6 Peer to peer fills a niche.

P2P technology has the potential to help reduce load on traditional CDNs in particular scenarios, so it will be useful to some content providers and will experience some growth. When lots of viewers are watching one piece of content, and the network conditions between them are favourable enough to manage latency, the results can be impressive. But with a mindshare legacy to overcome, and tight constraints on the situations in which it is most effective, growth of P2P will have limitations.

7 Better broadcast won't beat broadband.

Advances in broadcasting, and indeed in hybrid broadcast and OTT platforms, will make an impact in the coming decade, as we bridge the transition towards IP. However, adoption of broadcast standards in general purpose devices such as mobile phones will be limited, preventing these technologies from becoming ubiquitous. And with universal availability of broadband high on the agenda around the world, it will eventually become untenable to run separate broadcast networks.

8 You can't do it alone.

It was once possible for a broadcaster, or at least a consortium of broadcasters, to own and control their distribution network. For delivery over the internet, this simply isn't possible. Content providers need to partner closely with CDNs, and in some cases with ISPs. For the big players, this is happening already, but new entrants need to heed this lesson. At the same time, CDNs themselves need to build closer partnerships with ISPs, and this will increase in importance as more CDN caches and edge computing capabilities are deployed within the ISP networks.

9 Metrics matter.

In broadcast networks, we could largely assume that each viewer would receive the same quality of experience as any other. With IP, there are considerably more variables. This means that we can deliver more personalised services that adapt to the user's circumstances, but it also means that more elements can go wrong. Good QoE monitoring is crucial, and usually requires synthesising data from multiple vendors in the supply chain. Yet metrics are not common across organisations. The industry must collaborate on these metrics if it is to achieve the best results in an efficient manner.

10 Rethink the economics.

Broadcast networks run at a relatively fixed cost. Adding more channels might cost extra, but adding more viewers does not. A streaming service – at least one operating with unicast CDN delivery – has costs that scale with the number of viewers. This may seem frightening at first, but it introduces a certain discipline around the return on investment of acquiring and delivering specific content to specific users. The move to digital supply chains is beginning to bring this capability, and it will underpin a seismic shift in the economics of media.

Delivering the IP Future was researched by **Rowan de Pomerai** and **Abdul Hakim**.
This report was authored by **Rowan de Pomerai**. Design was by **Vlad Cohen**.

The DPP would like to thank all of the member companies who made this report possible by being so generous with their time and insight.

About the DPP

The DPP is the media industry's business network. It is a not-for-profit company with an international membership that spans the whole media supply chain, covering global technology companies, production companies, digital agencies, suppliers, service providers, post production facilities, online platforms, broadcasters, distributors and not-for-profit organisations. The DPP harnesses the collective intelligence of its membership to generate insight, enable change and create market opportunities. For more information, or to enquire about membership visit

thedpp.com

About Lumen

Lumen is guided by our belief that humanity is at its best when technology advances the way we live and work. With approximately 450,000 route fibre miles and serving customers in more than 60 countries, we deliver the fastest, most secure platform for applications and data to help businesses, government and communities deliver amazing experiences.

Learn more about the Lumen network, edge cloud, security, communication and collaboration solutions and our purpose to further human progress through technology at **news.lumen.com**, LinkedIn: /lumentechnologies, Twitter: @lumentechco, Facebook: /lumentechnologies, Instagram: @lumentechnologies and YouTube: /lumentechnologies.