Broadpeak White Paper

Giving Video Streaming Control Back to Operators

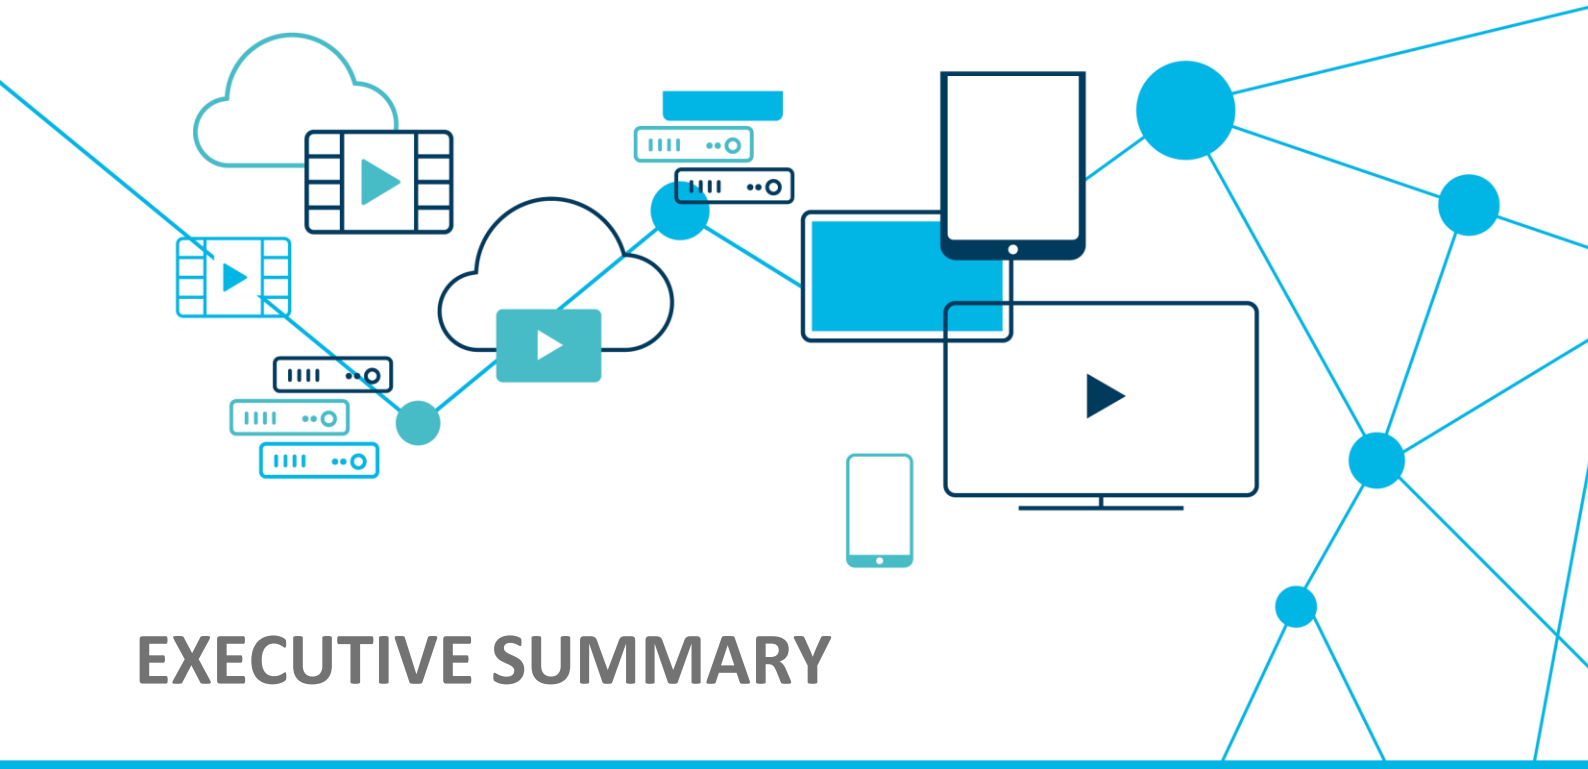

EXECUTIVE SUMMARY

Adaptive bitrate (ABR) technology has been the enabler of multiscreen video and has largely contributed to the popularity of de-linearized modes of consumption (VOD, catch-up, start-over, etc.).

The original principles of ABR propose delivering several quality streams of the same content and leaving it up to the end-device to decide which version to use. Now, what could be seen at first as an advantage in terms of adaptability is gradually becoming a limitation as ABR traffic is soaring, pushed naturally by the increasing number of mobile devices and, even more dramatically, by its progressive adoption for linear channels on big screens in place of traditional broadcast.

In order for video service providers to control the quality of their service and to secure a rational usage of their infrastructure, this paper considers the option of having the network participate in the stream selection process and regulate the amount of bandwidth distributed to each connection individually.

Moreover, it details how server-side selection provides additional benefits, such as low-latency streaming and fair distribution of the available bandwidth between the devices of a same home network.

At the end of the day, using the S4Streaming approach will decrease transport network expenses and significantly improve the quality of experience (QoE) and loyalty of end-users.

CONTENT

HOW ABR HAS BECOME THE NATURAL FUTURE OF VIDEO DISTRIBUTION 4

…AND WHY IT SHALL EVOLVE TODAY? 5

LOW-LATENCY STREAMING NEEDS A FIX 7

GIVING CONTROL TO THE NETWORK 9

PRACTICAL BENEFITS AND USE CASES 10

Controlled Quality of Experience 10

Controlled network resources usage 11

Cellular networks particular use case 12

Resolve multiresolution limitation in low-latency 13

Orchestrate the home network 14

OPTIONAL OPTIMIZATIONS ON THE PLAYER 15

BROADPEAK'S S4STREAMING™ IMPLEMENTATION 16

Quality streams control 17

Centralized management 17

Analytics and monitoring 18

CONCLUSION 18

ABOUT BROADPEAK 19

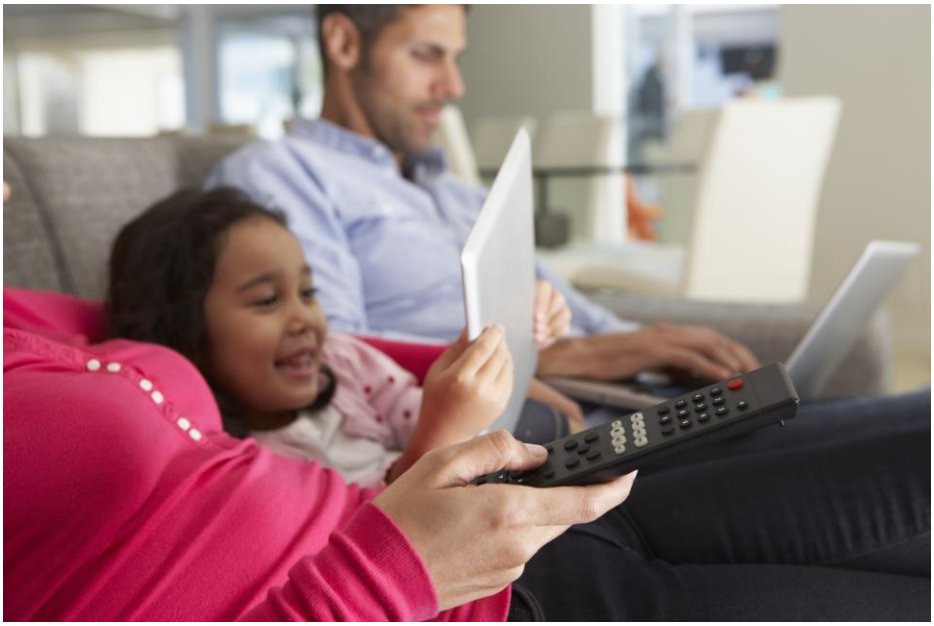

HOW ABR HAS BECOME THE NATURAL FUTURE OF VIDEO DISTRIBUTION

ABR has been one of the most disruptive evolutions in the history of video delivery. The difference in how video content has traditionally been delivered is essentially based on three characteristics: ABR streams are multi-resolution, pre-segmented, and referenced.

MULTI-RESOLUTION

Maybe the most obvious benefit, ABR can adapt to any type of screen and therefore to any type of device and, quite logically, started becoming popular at the same time as widescreen mobile devices — smartphones first, and then tablets — began appearing on the market.

PRE-SEGMENTED AND REFERENCED

An ABR stream is composed of video segments that are technically autonomous from each other and that are displayed successively at the client level following a list given by a manifest. This implies that content can easily be manipulated by simply modifying not the stream itself, but the references to that stream in the manifest that each user receives individually.

Such manipulations can be applied on the timeline of an asset, thus more easily enabling applications such as trick play (i.e., pause, fast forward, rewind), start-over, catch-up, or network PVR. The content itself can also be modified, which allows content personalization at any level of granularity, from program regionalization to individualized streams.

Major benefits brought to video distribution with ABR streaming

The possibilities that ABR streaming offers have already been widely adopted, and viewing a program on a mobile device or browsing through its timeline are standard today. Personalization, in particular targeted ad insertion, is also a critical feature because it increases the monetization of the product, and it is essential that television technology manages to reach the same level of evolution as what is

already available on the web in order to better compete. Looking forward, personalization also has an important role to play in what promises to be the next big step of TV evolution: making the viewing experience more immersive, a movement that technologies such as multi-viewpoint and virtual reality have already started. For all these reasons, **ABR is much better adapted to the context and challenges of our time; it is very certainly a necessary evolution for pay TV**, and it is progressively replacing traditional broadcast in more and more distribution networks and for more and more applications.

…AND WHY IT SHALL EVOLVE TODAY?

The advantages of using ABR instead of traditional broadcast can be summarized in its ability to propose a number of variants of the same content — different resolutions or different sequences of video segments — and adapt them depending on the user's context.

This naturally comes with a technical challenge, which is the necessity of multiplying the number of physical streams where there used to be, by definition, one-for-all when using traditional broadcast.

To some extent, it can be said that ABR has been working very well so far because its adoption is still in the phase of ramping up. Generalizing it to all video consumption modes today would definitely represent an unprecedented extra load on the networks, which could easily exceed their current capacities in many cases. This is, in particular, the case for linear channels that are mostly still received today via dedicated broadcast means of distribution and, even if declining, represent about two-thirds of the time spent by users when consuming video.

Including them in the ABR distribution basically means tripling the current amount of ABR traffic (probably much more in reality since network dimensioning is based on peaks rather than average usage).

U.S. time spent on traditional TV and online video

This is where **ABR streaming is reaching a glass ceiling in its development: It remains largely left out of popular and high-concurrency content.** The huge and unpredictable scaling efforts it would imply incite service providers to continue distributing linear channels and the highest-value events via traditional broadcast methods. Moreover, this is only considering today's context. Adding the development of some of the future ABR applications mentioned earlier, and despite all technical evolutions and the tremendous investments made to increase networks' capacity, this scalability limitation will become even more obvious. This is why, as ABR is reaching its level of maturity, control and containment of this expected surge in video delivery traffic will be one of the main topics to address in the years to come.

How can ABR traffic better be controlled?

Of course, there are certain optimizations that can be made to the content delivery network (CDN), for example efficient caching deep in the network or multicast ABR for the most popular content, but these can only optimize the overall traffic, not control it.

To ensure that bandwidth consumption always remains within a determined limit and manage how this limited resource is then shared between users, the first and most obvious option is to further manage the stream selection process — that is the choice of the video quality stream, defined by its resolution and bitrate, to be displayed on each end device.

In ABR principles, this choice is typically left to the player only. This can seem logical considering that several video layers are made available, and who better than the client device to know what best fits its screen size or its current available bitrate? But there is also one essential flaw with this approach: each device acts as if it were alone on the network, with no consideration for other devices sharing the same resource or the state of the network. This leads to poor optimization of the distribution in general and barely any control for the network service provider.

As a simple example, we observe that a lot of players will actually always use the highest quality stream available on the network, regardless of the screen resolution of its hosting device, which leads to a lot of a wasted bandwidth on a small-display phone where the perceived video quality would have been the same with a much lower bitrate.

Management of load peaks will become an increasingly critical topic as traffic grows. One advantage of TCP, and therefore ABR, is that whenever the amount of bandwidth on the network is reaching a limit, the players will detect that the delivery is slowing down and should automatically request lower quality streams. In theory, this should ensure that the service is always operational.

There are, however, two limitations to this mechanism. Naturally, it can only work up to certain limit as quality cannot be lowered indefinitely and new connections will have to be rejected at some point.

A second issue is that, whenever devices are competing for a common resource, there's currently no mediator or rules to determine how the resource shall be split between the different requestors. This generally results in a quite erratic fight and the distribution of the resource ends up being unfair, unpredictable, inhomogeneous, and unstable. It's unfair because the players better served will be the greediest and not the ones that should logically get served more bandwidth. Some players, for example, will always open numerous TCP connections in parallel for one single session, giving a false sense of priority. This can result in a smartphone, for instance, receiving more bandwidth than even a TV set.

Naturally, all these problems are augmented when the last mile is a wireless medium, like in Wi-Fi and cellular networks. Different tricks can be implemented at different levels, and the result is quite unpredictable, making it impossible for the content distributor to know in advance the general level of perceived quality of its own service, even more so that this quality can be very different from one user to another.

Moreover, this fight for resources can often lead to unstable oscillations for who's getting more bandwidth at a given time and significantly impact end users.

Finally, having streaming congesting the network doesn't impact video applications only, it slows down all data traffic and produces negative consequences on other services using the same resources, such as telephony, web browsing, or other internet applications.

In summary, ABR is a fantastic technology, with only one flaw, which is to consume much more resources than traditional video distribution. It works well, but only as long as the extra load doesn't make the overall traffic reach networks' limits. In which case, leaving all control to the players individually doesn't allow any rational management of resources and consequently induces random problematic behaviors that can neither be anticipated nor be fixed whenever they are detected. This leads to what we believe is a necessary evolution of ABR video distribution: giving back some control to the network and allowing stream selection to be driven on the server side is a way to achieve this goal.

LOW-LATENCY STREAMING NEEDS A FIX

On top of the aforementioned network considerations, there is a huge limitation that ABR critically needs to fix as it is technically unable, in its standard form today, to stream properly in a low-latency context.

Low latency, which is the ability to reduce the time between the moment an image is captured live and the moment it is displayed on an end device (also called "glass-to-glass" latency), has probably been the hottest topic of these past years in video distribution and certainly one of the most awaited ABR features.

In reality, what is called low latency in ABR is no more than what traditional broadcast can already achieve today, and this is what makes the feature so critical. As seen previously, today live linear

content remains the most popular way of consuming video and, as such, is a priority target to convert its distribution from legacy broadcast to ABR. But adding around 20 to 40 seconds of latency, as it does in the standard ABR streaming form, would induce an obvious regression. Low latency is a mandatory step for ABR to complete before becoming the overall distribution reference format.

Now, the issue is that low latency in ABR relies mainly on video segment chunking and that it makes the job of stream selection much more complex for players. Segment chunking means that each video segment is sub-divided into smaller entities called chunks and can thus be downloaded, chunk after chunk, as it is being created without having to wait for the whole segment to be available. Without chunking, this waiting time for segment availability adds up for all components along the delivery chain and accounts for a great part of the intrinsic delay that ABR traditionally adds.

But chunking comes in conflict with how players are usually designed to assess the available bitrate and therefore how they select the adequate quality to stream. As illustrated in the figure below, the typical algorithm used by the player is to take the size of a segment and divide it by the download time to deduce the available bitrate on the line. As long as the segment is entirely available at the moment of the player's request this works well. If the segment is built progressively chunk by chunk, the time of download is no longer constrained by the available bitrate but rather by the time it takes to create the segment, and the measure gets flawed.

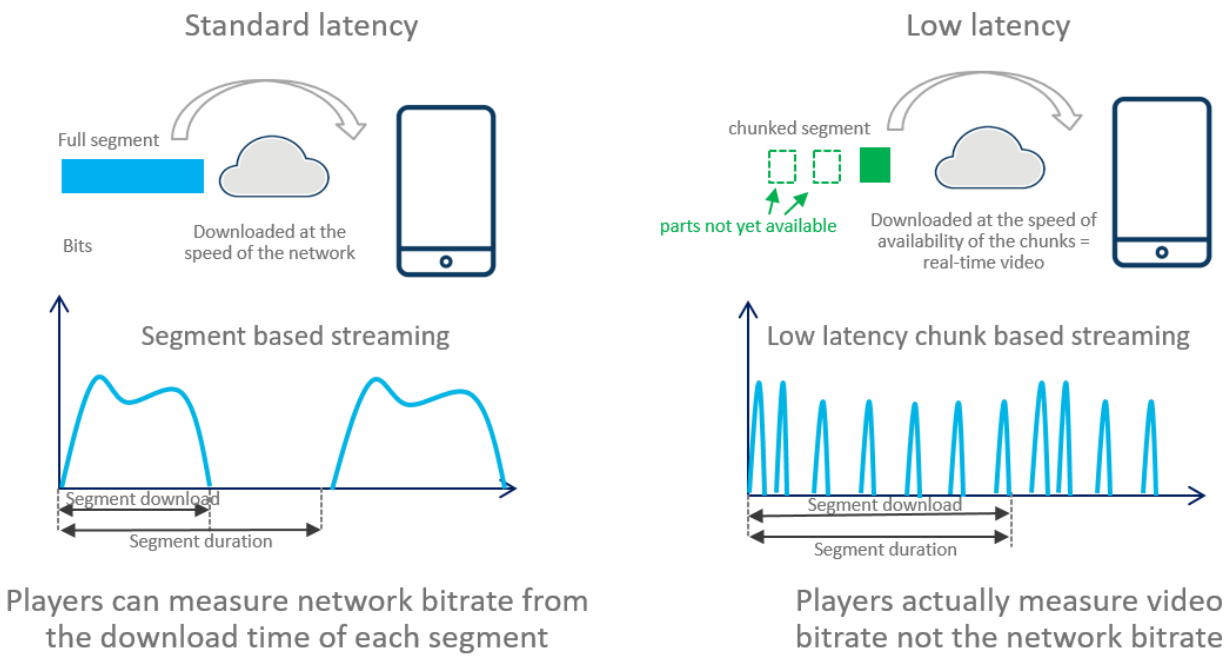

Measuring available bitrate in low latency

Some players are naturally trying to fix this antinomy between ABR low latency and the essential multiresolution feature, but this is an ongoing process that still bears too much uncertainty to be

considered for production. As detailed further on, driving the measure and the server-side stream selection is not subject to the same constraints and can be a good option to resolve this issue.

GIVING CONTROL TO THE NETWORK

Service providers should have more control over how their network is used by ABR devices for at least two basic reasons. First, because it's only fair that the entity investing in the infrastructure would want to ensure proper usage. Second, because it's the only way to ensure that the high-value content that is video streaming content is delivered with a minimum guaranteed quality to all users. Else said, **service providers should have control on their network usage and on the delivered QoE**.

Giving control to the network with server-side stream selection process requires two essential functions:

- The streaming servers must be able to assess the available bitrate for every connection established with an end device.
- The streaming servers must be able to dynamically alter the resolution choice initially made by the player (downward as well as upward). It can be based on the measured bitrate above but also on external parameters, such as pre-defined rules or commands, in a way that doesn't affect the ongoing playback on the devices.

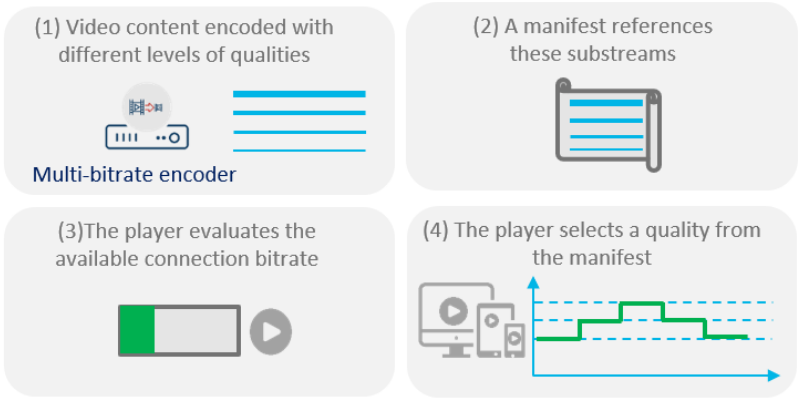

Client-side ABR principle

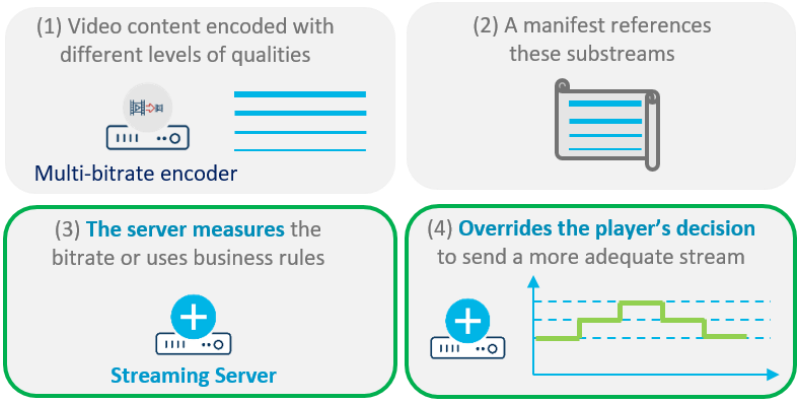

New server-side approach

Naturally, the streams selected and sent from the server to the end devices shall remain fully compliant to streaming standards, namely Apple HLS, MPEG-DASH, or Microsoft Smooth Streaming, and work seamlessly on any player with no specific adaptation. Otherwise the functionality could be extremely difficult to implement and wouldn't be applicable to some devices' native players such as Apple AVPlayer or Android Exoplayer.

The most important challenge in doing this is to make the player accept a video resolution that is not the one it has requested in the first place. This doesn't come by default, but it can reliably be achieved

by adding, on the stream itself, some extra technical information that will then automatically be used by the player to adapt.

Regarding the bitrate assessment aspect, measuring the available connection speed on the server side is actually not something new. There are a number of innovative algorithms that have already been implemented on standard servers. BBR (Bottleneck Bandwidth and Round-trip) is probably the most popular at the moment, but other options are also available. The great advantage they offer is that they are implemented in the lower layers of the system, which means that the impact on the server's CPU is much less than if the process had been implemented at an application level.

Overall, having streaming servers measuring the bitrate of each of their connections doesn't significantly impact scalability. This is a critical aspect, as it would certainly be an issue if such implementation was required to reconsider the dimensioning and the number of servers to be deployed in a network.

PRACTICAL BENEFITS AND USE CASES

CONTROLLED QUALITY OF EXPERIENCE

BETTER OVERALL QoE

Some network conditions can impact the overall delivered video quality so, simply put, it will always be better to take them into account rather than ignoring them. By adding network-aware criteria, decisions of which quality stream to deliver to each connection are more consistent all across the network, more stable, and they can focus the available resource where it matters most, leading to an improved average QoE for the overall population of users.

PREVENT ANTICIPATED DEGRADATIONS

In standard ABR distribution systems, when monitoring indicates a probable video degradation to come, typically when getting close to a resource capacity limit for a given area of the network or even when the problem is already occurring, there's little that network operators can do about it since they don't manage the different end-devices streaming from their network. Conversely, placing the network bitrate distribution on the server side is an opportunity to act on the issue and resolve it before it has any significant impact and before it affects more subscribers.

OPTIMIZE QUALITY/COST TRADE-OFF

For video distribution, as in most domains, the curve of video quality against the cost it represents on the infrastructure has an elbow shape. Investments can have a significant impact on quality only up to a certain point, then the cost becomes rapidly unreasonable for little benefit.

A server-side approach offers both more precise tools to monitor the streaming at any point in the network and the ability to adjust this volume of streaming, making it possible to identify the sweet spot and maintain the operation within that range.

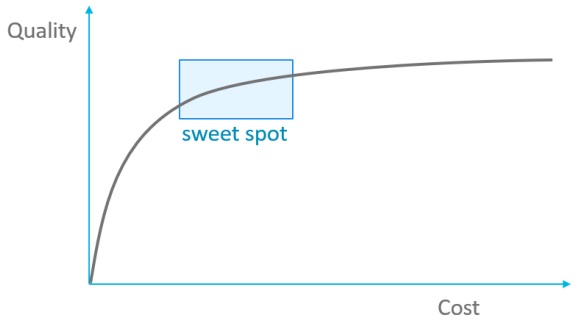

Typical quality vs. cost trade-off

DIFFERENTIATED QUALITY

When all video sessions cannot be served at the best available quality, for example during peak hours, the best way service providers can secure the most optimized overall quality is to maintain control over which sessions may be affected by this resource limitation. Server-side stream selection can gracefully and dynamically help by preserving the bitrate of higher value content such as live events, streams serving the main screen rather than mobile devices, or of premium rather than regular customers.

CONTROLLED NETWORK RESOURCES USAGE

BANDWIDTH MEASURE THROUGHOUT THE NETWORK

Controlling network resources implies first having a clear view on how they are used. Implementing a mechanism that allows all servers on the network to dynamically measure the available bandwidth ensures very accurate and valuable information on the state of the network. As streaming servers are typically tiered, this information becomes available all along the delivery chain.

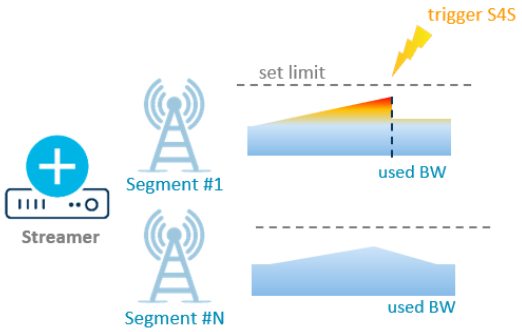

Per-segment stream selection control

DYNAMIC MANAGEMENT AT CAPACITY LIMITS

Control also means that it is possible to act on the information reported by the network. One simple example is when the streaming of one given segment of the network is getting near its capacity limit, server-side stream selection technology provides the tools to see this situation approaching and the ability to gracefully and dynamically reduce the overall delivered bitrate on that segment. The trigger sent to the streamers can either be an automated rule, pre-provisioned in the server itself, or a command sent from a central management point via API. The trigger only applies to the relevant network segments, leaving the other unchanged. Once the trigger is received, the streaming bitrate can be reduced smoothly, with managed policies preserving a certain category of devices, users, or programs.

CONSISTENT AND FAIR REPARTITION

Driving the streaming on the server side means that all users are served consistently, according to the policies defined by the service provider. It ensures that players don't have to worry about anarchical competition between each other, often biased by opaque implementation tricks that some are using, and can focus on playing their role fairly.

CELLULAR NETWORKS PARTICULAR USE CASE

Similar to wireline distribution modes, cellular networks have seen their data traffic dramatically increasing these past years as a consequence of the appeal for video consumption on mobile devices. Adapting the network capacity to that surge is particularly challenging for the last-mile radio access, and this segment has generally been the limiting factor for video streaming. 4G has enabled video on-the-go on smaller mobile devices but, in most cases, has fallen short in generalizing and extending the service to in-home consumption, for example on large screens.

A lot is expected from the technology upgrade to 5G, but its deployment will take time, and it is likely that, by then, new requirements will have appeared, rapidly filling up the significant additional amount of capacity 5G will offer. This will certainly include supporting the peaks of popular live events, streaming on larger displays, more individualization, more bandwidth-hungry video formats such as Ultra HD and Virtual Reality, and a number of changes to come that will consistently ask for more resources.

More generally, radio network technology will improve but so will wired networks and currently, since not constrained by the extra complexity of wireless connections, are more likely to set the standard of how much bitrate the video applications of the future can reasonably consume. In other words, video traffic optimization is even more relevant today, and probably tomorrow as well, for cellular networks.

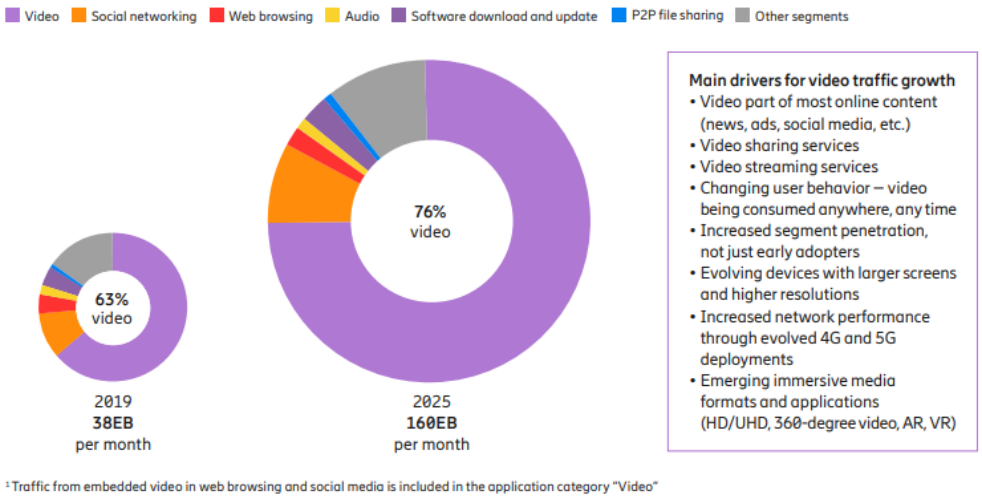

Mobile traffic by application category per month (percent)
Source: Ericsson Mobility Report, Nov 2019

A second specificity of cellular networks is that the bandwidth available to the end user through a radio connection is particularly variable, notably because of interferences, noise, and the user's mobility. Typical client-side estimation algorithms are, in many occurrences, not precise enough to accurately follow these variations, leading to recurring image freezes. In comparison, server-side measures are by nature more reactive and can already better follow these fluctuations, but they can also benefit from network information at the cell level, from a calculation of its servers' total throughput and by exchanging with the bandwidth scheduler of the cell. A radio link is always precisely scheduled, and synchronizing this information with the stream selection process can be very beneficial to predict future variations.

Moreover, as mentioned earlier, whenever a resource is shared among concurrent users, managing its fair distribution is important to avoid erratic competition between players and the instabilities it implies. Radio channels are indeed shared, and the notion of fairness in the distribution is even more relevant considering that these instabilities would add up to the intrinsic variability of wireless transmissions.

For these reasons, the ability to optimize and stabilize the stream selection process on the server side offers interesting benefits in comparison to the exclusive client-side approach available today.

RESOLVE MULTIRESOLUTION LIMITATION IN LOW LATENCY

As mentioned previously, streaming live content without adding latency requires breaking up the several-second video segments into smaller chunks. This causes issues on the players deployed today, as they rely on the existence of these segments to choose a quality stream. Server-side bandwidth measurements are not using the same principle, rather they use lower-layer TCP information provided by the server's operating system and, consequently, remain immune to that limitation.

Actually, streaming only one video layer in low latency may appear acceptable if we consider that it is a must only on the main screen, where it will be compared to traditional broadcast latency, and not on mobile devices, where people are used to a much longer latency anyway. There wouldn't be any obvious regression compared with the situation today and, providing that the set-top box feeding the main screen is wired, it is quite doable to ensure that there will always be enough bitrate capacity on the link to consume the highest quality layer in all cases. In order to evolve toward a system that could also offer low-latency streaming on wireless set-top boxes or on other mobile devices, it is necessary to have a solution that lifts this limitation and enables multiresolution even in a low-latency context.

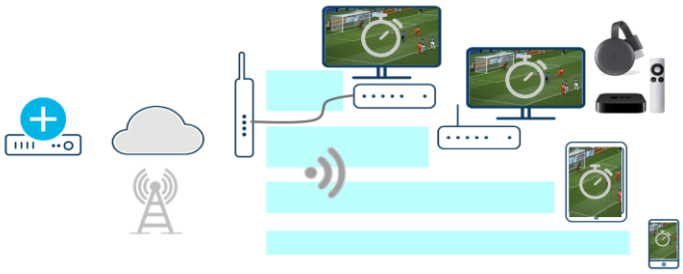

Multiscreen low-latency streaming and synchronization

One other benefit of enabling multiresolution in low latency is that, once all the devices of the home can play at the closest possible to the live edge, it becomes possible to synchronize them and make sure that two users can view the same program without being disturbed by the interferences caused

by the other. A slight video delay between two devices in the same visual field is already annoying, but an audio delay easily becomes unbearable. This feature can greatly improve the experience when proposing multiroom TV services and is often a must when proposing video in public venues.

ORCHESTRATE THE HOME NETWORK

Adaptive bitrate streaming is often called “multiscreen,” since it is extending video consumption from what used to only be set-top boxes to any type of device. A logical consequence is that video distribution systems today often have to serve more than one session per home, as several persons, each connected on their own device, can now consume a different program. The bitrate available for one given home is generally constrained, so the topic discussed previously of sharing this limited resource in a consistent and fair manner in the network also applies at the home level.

The principle of home orchestration is that each device informs the server, in its first streaming request, of which user and which device does the request for which program. Then, depending on this information, the server applies the network provider’s policies to select the appropriate quality stream. On the same model when managed at network level, the selection is dynamic and can vary depending on the conditions of the home network.

As an example, let’s imagine that one given home has 10 Mbps internet access and that the children in their room are consuming 8 Mbps on their smartphones. The parents come home and want to watch an important event on the main TV screen, then the orchestration process can decide to reserve 4 Mbps to the main screen, ensuring a minimum QoE, while seamlessly lowering the bitrate delivered to the smartphones to a level that is still very acceptable for such screens.

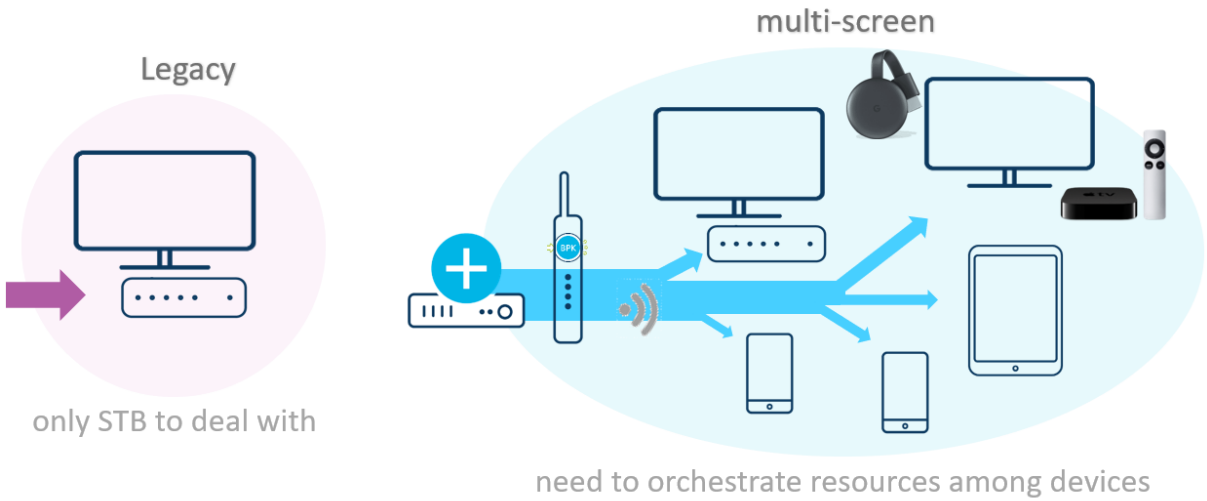

Multiscreen home orchestration

The information given in the first request can quite easily be available without having to customize anything on the player. It can come from connection details such as the user-agent that most mobile devices already include by default, or it can also be enriched by the TV application, in general provided

by the network operator. That said, specific developments on the player can optionally be added if some of the data it can share is relevant for a particular use case. This option is further detailed in the next section.

Note that local video streaming servers can be embedded directly in-home gateways. It would, for example, come by default when using multicast-ABR technology, since the multicast-to-unicast function can indeed be considered an ABR streamer downstream and therefore can include such technology. Stream selection from a local streamer can make sense for managing home orchestration, as the intelligence and decision process can remain local within the home network without having to rely on an external system to gather, sort, and process individual data from many households in a centralized point.

OPTIONAL OPTIMIZATIONS ON THE PLAYER

As mentioned, server-side stream selection doesn't require any modification on the client, but there are some extra benefits that could optionally be obtained by allowing information exchanges between the server and the video application or the player. Selecting the video quality stream ultimately remains the responsibility of the server, but its decision can be enriched by information that can potentially be easier to get from the client than from the network. These could include, for example, technical information about the hosting device, synchronization data to optimize the player's buffer usage, specific information about the nature of the program being played, or further user information.

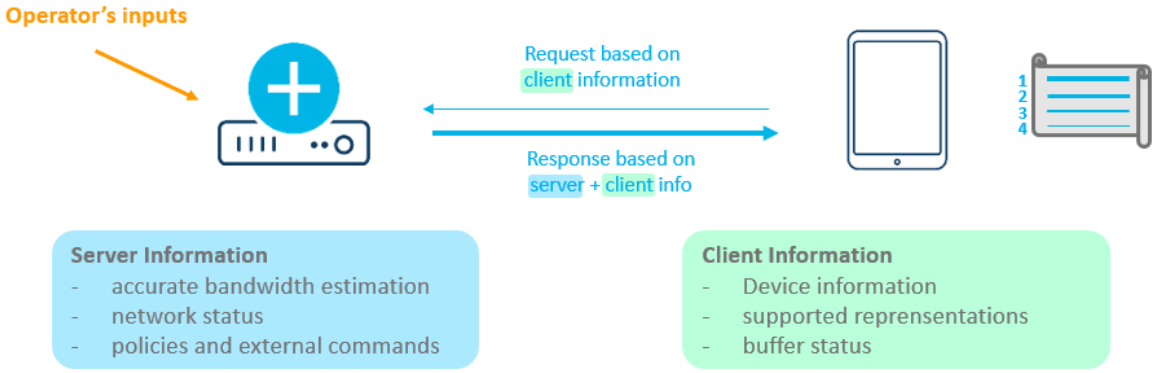

Information exchanges between server and client for stream selection

The ability for the client to pass network information for one particular connection and one particular customer is something that is already largely used today. One simple implementation consists of adding these parameters as query tags appended to the HTTP request, the model that is generally used, for example, for manifest manipulation at the origin server level for targeted ad insertion. Most applications and players on the market today already have implemented such mechanisms, and their adaptation to server-side mode of stream selection is quite straightforward.

BROADPEAK'S S4STREAMING™ IMPLEMENTATION

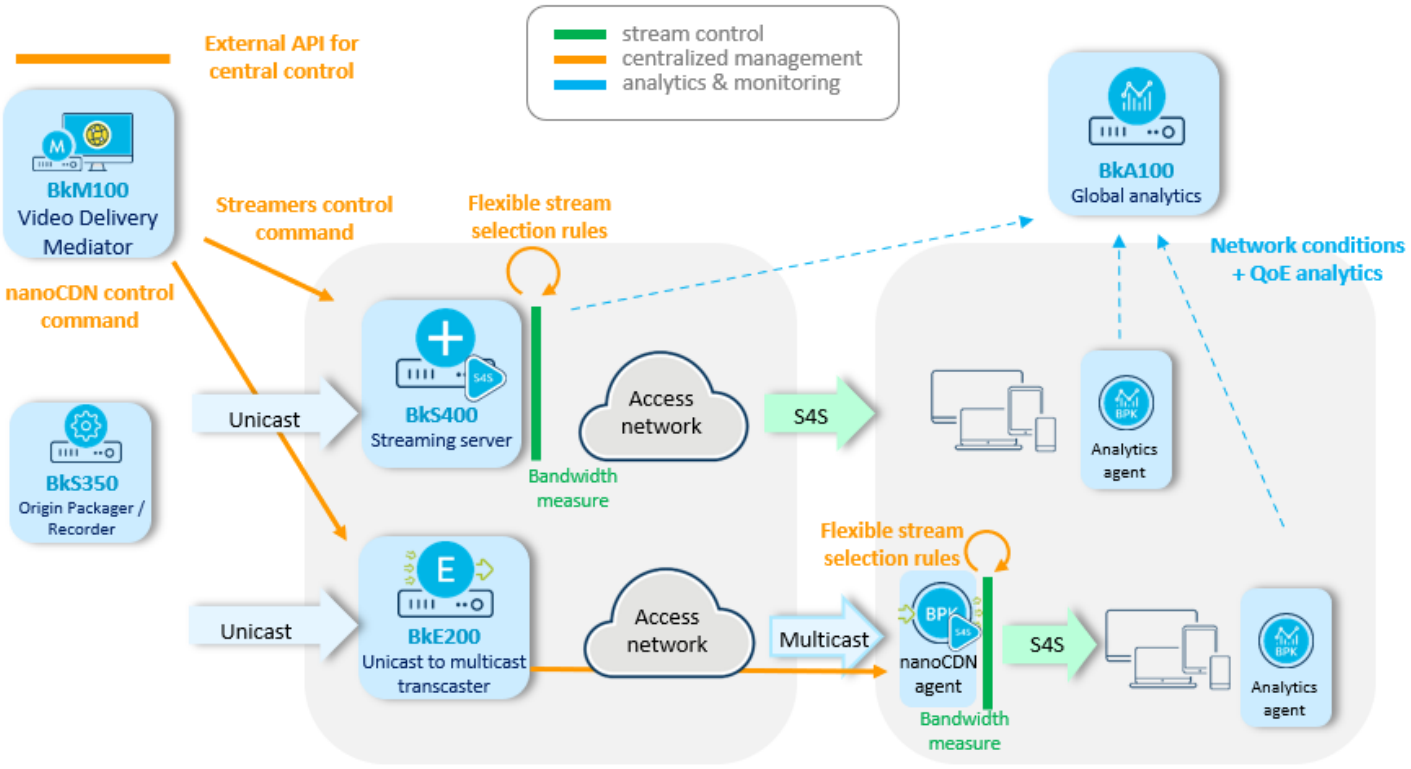

S4Streaming architecture overview

Broadpeak's S4Streaming™ solution

S4Streaming stands for "server-side stream selection for streaming" and is the complete solution proposed by Broadpeak to implement a network-driven control of video streaming.

The solution is fully compliant with ABR streaming standards, Apple HLS and MPEG-DASH, and doesn't require any adaptation to existing ecosystems.

It encompasses three different essential aspects:

- stream control itself, with per-session bitrate measure and quality selection
- commands and automated rules to drive the selection process
- analytics and monitoring, to help with decision-making and report on how they have effectively been applied

S4Streaming concerns all types of streaming servers present in the CDN, BkS400™ shield caches in the headend, BkS400 video streamers located in PoPs or deeper in the network, and the nanoCDN™ agent directly placed in a CPE in each home network.

S4Streaming technology can be implemented at one particular point of the distribution chain or at all levels, depending on the targeted use case.

> ***The solution is fully compliant with ABR streaming standards, Apple HLS and MPEG-DASH, and doesn't require any adaptation to existing ecosystems.***

QUALITY STREAMS CONTROL

There are two different aspects for stream control: the first one consists of measuring the bandwidth between the server and the clients for each and every connection, and the second one selects a video quality stream, often overriding the client's own decision. These two functions work well together, as the server can use its bandwidth assessment as the main criteria for selecting the appropriate quality, replacing or complementing the ABR client's selection, but they can also be used separately. It is possible to use bandwidth measurement just for analytics or apply streams selection based on other criteria, typically using preconfigured logical rules or commands sent from the service provider's management system.

Both processes have a very limited impact on the server's hardware resources and therefore don't imply any redimensioning when used, making them easy to activate even on a production system.

CENTRALIZED MANAGEMENT

Now that the system has a way to seamlessly impose a different quality stream than the one requested by any client, it is possible to act on these decisions from a centralized operation point. In other words, the available bandwidth on the link is no longer the only criterion for the selection, like it was for the traditional client-side approach.

In the most basic approach, commands can be sent to drive or influence these decisions like, for instance, to impose a reduction of bitrate on all the connections of a network segment getting near saturation, possibly taking into account that some users, content, or devices have to be preserved as much as possible. These commands can either be manual or driven from an external management system via API.

In a more advanced mode, the network constantly feeds the servers with technical information that can in turn trigger, by automated rules configured on the server, a predefined action.

ANALYTICS AND MONITORING

S4Streaming produces precise information on the available bandwidth of all open connections at any point in the network. It enables the computing of very valuable analytics on the usage of the network as a whole, region per region, or even down to each connection level. It also provides data on what quality stream has been provided to which user, an important metric on the perceived QoE by each subscriber of the service.

These metrics can typically be provided today by dedicated components on the client, but it's important to note that, in a server-side approach, the client doesn't necessarily know the resolution or bitrate it is actually playing, since its decision can have been overridden seamlessly by the streaming server.

Still, having a QoE analytics component in the client can be very interesting, for example to compare requested versus actually delivered streams, but also for other metrics such as zapping time, frequency of video freezes, and more.

For this reason, Broadpeak's solution does include this option with SmartLib™, a library directly integrated in the client. The information gathered all along the delivery chain, from the central headend, the PoPs, the CPE's embedded streamers, and down to the client can be aggregated and computed in centralized BkA100™ server to give access to powerful and business-driven analytics.

CONCLUSION

One of the main reasons for ABR streaming's success has been its adaptability to any kind of device and to different types and qualities of connections. In that context, the principle of only giving the client the choice of which resolution and bitrate to use whenever it initiates a video streaming session may have seemed logical so far.

Today, the experience shows that it also leads to a certain number of limitations on how network resources are used and how the different devices share these resources between each other. These limitations can however be lifted by not relying only on clients' decisions but by also incorporating information from the network.

S4Streaming server-side approach gives back control to the service operator so that it can better optimize its resources and ensure that they are used wisely to deliver a fair and consistent quality of experience to all users.

broadpeak

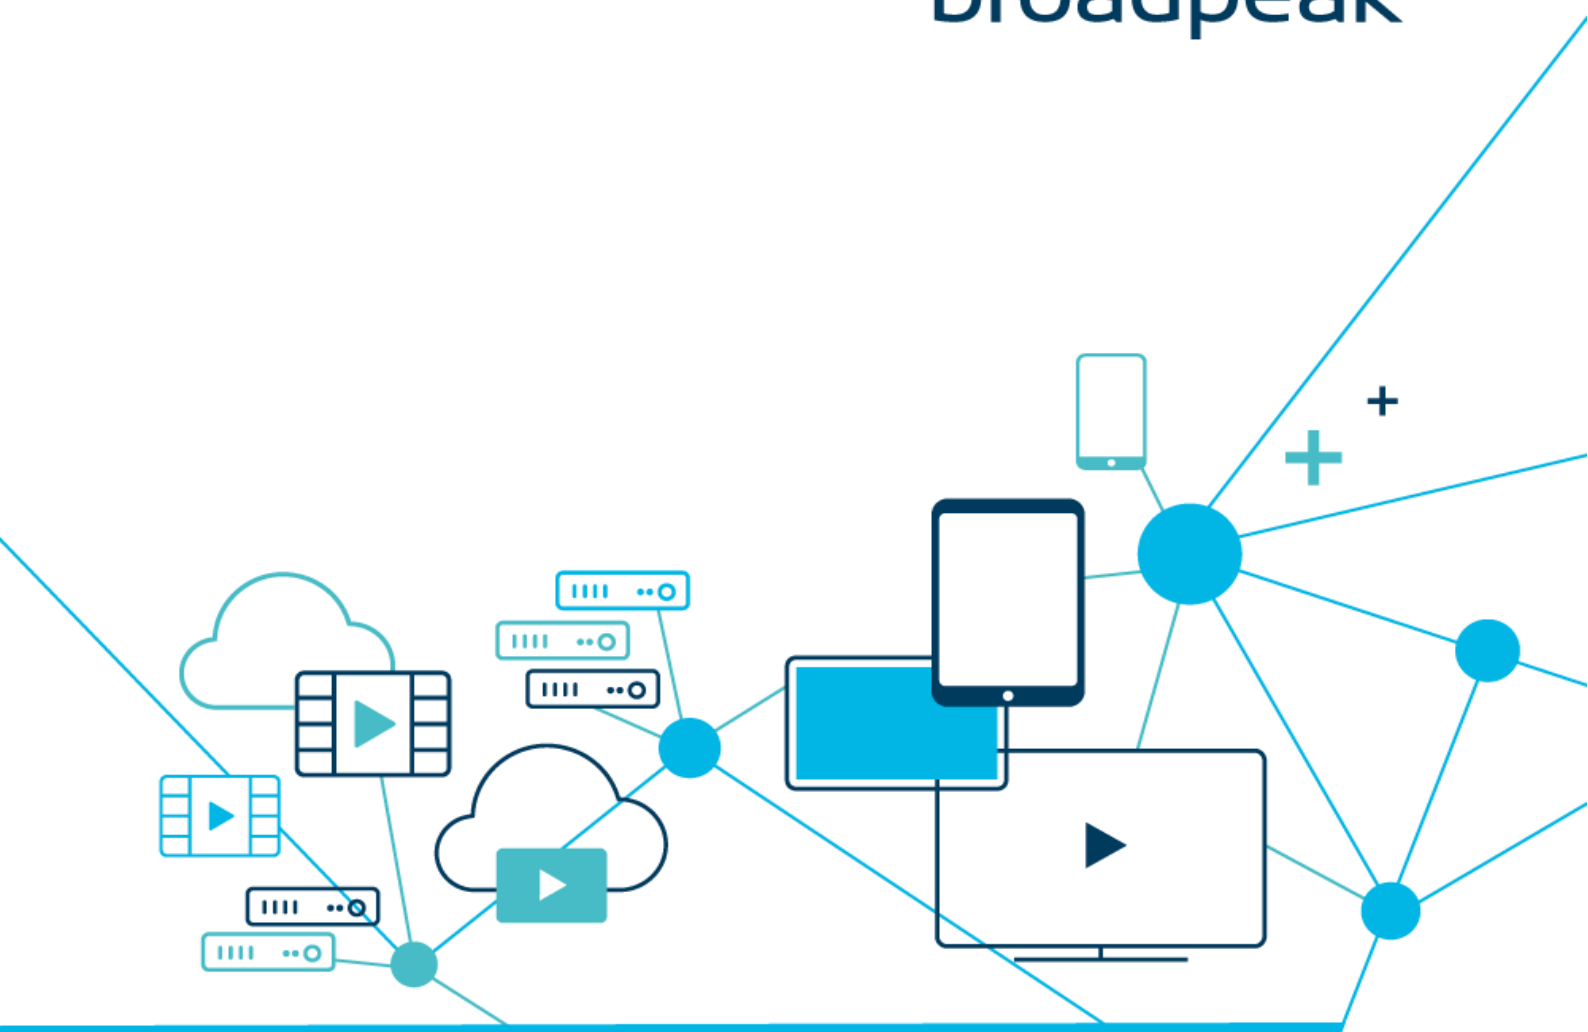

Broadpeak® designs and manufactures video delivery components for content providers and network service providers deploying IPTV, cable, OTT, and mobile services. Its portfolio of solutions and technologies powers the delivery of movies, television programming, and other video content over managed networks and the Internet for viewing on any type of device. The company's systems and services help operators increase market share and improve subscriber loyalty with superior quality of experience.

Broadpeak supports all of its customers worldwide, from simple installations to large delivery systems reaching capacities of several million of simultaneous streams. The company is headquartered in Cesson-Sevigne, France.

https://broadpeak.tv

We Deliver the Future of Television