UmbrellaCDN™ or CDN selector for live

&

on-demand content

A CTO Consulting white paper authored by Benjamin Schwarz and Philip Hunter, co-sponsored by broadpeak

White Paper contents

1 Executive summary

CDNs evolved to provide delivery services for a range of content providers and a few large global players such as Akamai, Level 3 and Limelight still dominated the field. They have worldwide reach but do not fully cater for the increasingly diverse range of delivery requirements and cannot often reach into local access infrastructures provided by ISPs (Internet Service Providers). Local CDNs have been set up by telcos and cable operators over their own networks, to provide services better tuned to their own customers. These CDNs were first deployed to deliver operators' own content for their IPTV or digital cable services, but since then have been opened up to third party content providers. This has led to operator CDN approaches becoming new sources of revenue for broadband service providers.

Meanwhile OTT video services are growing rapidly as content providers increasingly go direct to end-users. This is happening at a time when more and more channels over satellite, cable TV and terrestrial are in HD, raising expectations and demand for better QoS over the Internet as well. This in turn has driven demand for scalable CDN services with global reach but that at the same time are sensitive to varying local requirements. The result is demand for a new generation of CDN services combining the reach of the global one stop shop with the quality control of the local operator CDN.

The idea of the Federated CDN then emerged, whereby multiple local or regional CDNs participate in single end-to-end services through standards that enable interoperability. The concept sounded good but has not turned out so well in practice because, although the different CDN components support common standards, internal technical differences prevent individual CDN providers from using their own specific techniques for optimizing delivery over their own networks.

Moreover, Federated CDNs are not designed to meet diverging content requirements on the basis of business rules. For example an operator might want to deliver free movie trailers at a lower resolution than full premium movie content and therefore at a lower price, which cannot be achieved over a Federated CDN.

More recently still some CDNs have come along to resolve the problems of Federated CDNs, quite simply by putting content providers in full charge of end-to-end delivery. One CDN vendor coined the term Confederated CDN for this new model. This enables CDNs to be selected on the basis of business rules that take account of changing content requirements and network conditions, optimizing price and quality. With a CDN confederation the content owner can select a CDN for each part of an end-to-end path, rather than as with a Federated CDN delegating the decision to the infrastructure. In this white paper we set out the benefits of the CDN confederation and show how a "CDN selector" operating as an umbrella CDN can harness them.

2 Why do content providers need a CDN?

CDNs originated in the mid-1990s in the early days of public Internet usage, as web pages started to incorporate graphics that meant delivery was no longer a question of bytes, but a matter of kilobytes. Today's well known global CDN providers such as Akamai established their roots around that time under the label of web acceleration, exploiting caching to distribute content closer to the points of consumption. The motive of caching then was to reduce latency more than to cut bandwidth costs, but established the key CDN principles of enabling some data to be prioritized for express delivery, along with caching hierarchies to stage content at various points in the network according to its consumption profile.

This model has carried across to the OTT video era but with some important extensions. The proliferation of video data has meant that cost and also capacity have become important factors driving CDN usage and development, on top of the desire to improve performance and reduce latency. Streaming from the edge costs less than from a central point, but consumes more storage, implying a trade-off between cache and bandwidth. The balance of this trade-off depends on the popularity of content more than ever. While it pays to locate blockbuster movie or premium sports content near the network edge to save consuming bandwidth distributing large numbers of unicast streams, niche material is still more cost effectively stored in just one location on a central server.

Another factor is that content varies in availability between locations according to rights or business factors, with geo-blocking sometimes used to restrict access in particular regions. Such factors need to be taken into account by CDNs, increasing overall complexity. As we will see the umbrellaCDN approach enables content owners or OTT service providers to take account of these factors across a confederated CDN, without having to create the geo-blocking rule on every single CDN.

3 What are the options?

Given that CDNs are now essential for video delivery at most scales to ensure QoS and contain costs, content owners and OTT service providers now have two broad options. Firstly there is the one stop-shop from the likes of Akamai, Level 3 and Limelight, with global coverage, while the alternative is the operator CDN serving a local market. A third option is the D.I.Y approach but this is only viable for the very largest players such as Netflix with the resources to establish their own CDN infrastructure and so will not be considered any more in this paper.

3.1 . The one stop shop

Until quite recently the one stop shop was the only option and has the one big advantage of global reach. This can be a good starting point for an operator initiating a Greenfield project in a geographical area covered by one of the big global CDN providers. The business models vary, but as there is almost invariably some usage component to the billing based on bandwidth and storage, such a one stop shop may be cost effective for non-premium content.

A single global CDN can also take care of the content processing, including encoding, packaging and encryption. Another benefit can be availability of analytics to track usage and QoS, although it has to rely on agents in the end device or the application to provide feedback about a session. This can prove very expensive and complex to administer, because of the number and diversity of client devices. An embedded agent would have to be compliant with every single device, with a possible re-qualification required whenever a new version of an OS is released.

The one stop shop works less well for premium content where QoS has to be as close as possible to traditional broadcast or pay TV services delivered via cable, satellite, IPTV, or digital terrestrial. The issue here is that in many parts of the world, the global CDN operator will only have a very high-level peering agreement with local networks, often with only one or two points of presence per market. As soon as a service gains commercial success in any given market, service providers using global operators will need to rethink their CDN strategy regularly as it will always be a key part of their cost base. We believe a flexible approach is then most cost effective, making use of the second category, the operator CDN.

3.2 . The operator CDN

The major disadvantage of a global CDN is the lack of harmony with local networks to optimize both the end to end QoS and the economics of delivery. A global CDN can only work effectively when integrated well with the local operators in that given market. In practice this has not happened so far, which may be one reason Netflix is now working directly with local ISPs to optimize delivery of its own content, rather than as before dealing with global CDN providers.

Operator CDNs are deployed directly over the operator's existing infrastructure, involving direct contracts with either content suppliers or pay TV providers. Operator CDNs are deployed most often by telcos but also sometimes by cable operators, given that the latter are increasingly competing in broadband services as ISPs.

Both telcos and cable operators can exploit the intimate control they have over their infrastructure to offer a higher quality of service than a one stop global shop. They can prioritize video traffic and guarantee end to end QoS for all local services streamed entirely over their infrastructure. They are in full control of the path and can provision whatever bandwidth is needed to guarantee a given level of QoS, whether for premium or best effort services.

The ability to deploy specific technologies on their networks for traffic optimization means furthermore that operators can meet their SLAs (Service Level Agreements) efficiently and therefore be very competitive on price. It is now possible for operators to extend their CDNs as far as the home to achieve further gains in

efficiency and performance. Some vendors now provide software that turns the home gateway into a CDN end node supporting caching so that linear or live content to be transmitted to the home in a single stream, even if it is being accessed simultaneously by a number of users who are not all connected to the same box. Similarly on demand content can be sent at off peak times and cached within the home, utilising storage in the set top box or gateway. This optimizes use of the operator's access bandwidth and can postpone need for an upgrade to deliver a given number of channels at a specified quality.

Operator CDNs can also score on their ability to deliver timely and detailed analytics relating to performance, service usage and content consumption. This can be done at lower cost and with greater visibility down to the level of user devices without need for separate agent software.

	Global CDN	Operator CDN
Reach	Usually **international** through own servers or peering/transit contracts	Main **focus on own network** but can be extended through peering/transit contracts
Cost	Gets **lower with higher volume** commitments, can be high for small volumes	Can be **very competitive** if the operator deploys **specific technologies** on its network for optimizing traffic (eg nanoCDN)
Quality of Service	Uncertain according to the regions	Guaranteed on operator's network
Analytics	**Global information** excluding distribution network in the standard offer, detailed information if an agent is available but very expensive	**Detailed information** in the operator's network at a better cost
Support	Can **lack reactivity** with big companies	Usually **focused** and local

Chart comparing global one stop shop with local operator CDN.

3.3 . The Federated CDN

The one advantage for the one stop shop over the operator CDN, global reach, can be neutralized by interconnecting multiple operator networks to cover a much larger region. This is the idea behind the Federated CDN, designed to couple multiple operator CDNs into a coherent global content delivery service via common standards.

The idea of the Federated CDN initially created much excitement among Telcos when it was first launched around 2009. It promised to let them compete with the big one stop shops we have mentioned and generate added value revenue from their infrastructures by delivering premium content, rather than being relegated to the status of dumb pipes. This excitement crystallized into a group called the Operator Carrier Exchange, founded in June 2011 with a remit to share ideas, then to come up with coherent standards enabling members to interconnect their existing local CDNs into a global service.

Since then progress has been slower than expected and the promised standards have mostly yet to materialize, as has the precise business model. The underlying problem seems to be the realization that no standard can effectively overcome fundamental incompatibilities between the respective operator CDNs in the federation. Interconnectivity is not the same as full seamless interoperability. There are for example three different flavours of Adaptive Bit Rate Streaming (ABRS) that have been commonly deployed. Then if one CDN supports just say Apple's HLS while another uses only Microsoft's Smooth Streaming, they are not going to dovetail effectively as an end to end CDN. For this reason the early excitement has died down as it becomes clear that the Federated CDN concept is unworkable, at least in the foreseeable future.

The other big issue that the Federated CDN fails to address is the overall business model. It fails to determine how much each CDN should be paid on the basis of the amount of content it carries. It does not resolve where the billing should be done, how revenue should be shared, or how the volume of content should be measured.

3.4 . The Confederated CDN

A new model has already evolved that overcomes the deficiencies of the Federated CDN. The Confederated CDN gives up on the idea of working to a common palate of standards to establish compatibility between multiple local CDN components. Instead it is based on the view that for a given situation a particular CDN will be most suitable, best placed perhaps to provide the right QoS or the lowest possible cost.

It may be for example that a particular approach, such as extending the CDN right to the home gateway, may be most appropriate for a live session, even though it does not adhere to a specified group of Federated CDN standards. Under a Confederated CDN arrangement, a network supporting such an extended could be selected just on the basis that it is deemed best at the time according to the given business rules. The difference is a subtle one in that while a federation of CDNs can only work if each CDN is compatible with every other one, within a confederation of CDNs the focus shifts to coordination rather than integration. The coordination is achieved via a central server, which determines which CDN to use at a given time, with other CDNs then standing back. This is done via some CDN selector.

4 The criteria for selecting a CDN

The great advantage of a selector function within a confederated arrangement is that now CDNs can be chosen according to a wide range of criteria that can be precisely defined and controlled by the content owner or pay TV operator. Here we have identified five specific criteria for selecting a particular CDN, plus three additional use cases that exploit the fact multiple CDNs are available for content distribution at any given time.

4.1 . End user location

There are great regional variations both over content availability, as governed by rights or other business factors, and over the service available from the local CDNs themselves. CDNs vary in their footprint within a region, as well as in their capacity and performance.

Typically video content distribution is skewed, with the biggest volume concentrated in a given location but shifting over time, so that optimal choice of CDN will vary. This could be determined by real time analytics feedback and involve switching at a given time between one CDN and another to cater for a change in the traffic profile. This may happen when a large amount of traffic encroaches into an area where the CDN currently being used has poor coverage or is short of capacity, while another is better placed.

4.2 . Type of content

The content profile is a major factor in choice of CDN because making the right decision here can have a big impact both on cost and QoS. Some CDNs perform best with live content, because they are optimized for streaming with a real time end to end path spanning the streaming source, the CDN ingest server, appropriate content servers on the path and finally the end device. This will typically be via a flat network infrastructure with layers designed specifically for live streaming, imposing the least possible latency on video packets.

Other CDNs though may be optimized for VOD content, which has very different requirements, since there is no urgency to get video out across the network but with potential for cost savings by caching popular assets close to the network edge. It therefore makes sense to optimize live streaming and VOD delivery independently, with many CDNs better at one than the other. CDN selection should therefore take account of this.

Quality considerations should also come into play, given that some CDNs may be best for delivering high resolution content while others are perfectly capable of handling say free trailers at a lower cost. Another factor can be the type of streaming required. As we mentioned earlier, one problem with Federated CDNs is that constituent operator CDNs may support different streaming technologies, making it harder to establish efficient end to end content delivery. With a CDN Selector, CDNs can be chosen purely for the type of

streaming they support. Then a content owner can select one CDN say for HLS and another for Smooth Streaming.

4.3 . Time of Day

CDNs can vary their prices by time of day or for that matter week, month or year in an attempt to balance loads or extract maximum revenues when demand is at its peak. Content owners or OTT service providers can deploy a selector function readily to optimize costs by exploiting these price changes. There is potential for combining forward planning based on known prices with changing traffic conditions. It is worth mentioning that these selection criteria are not mutually exclusive and may overlap, with one sometimes trumping another. For example the need to switch to a CDN supporting a particular live streaming technology may outweigh cost savings that might be made by not doing so.

4.4 . End-user NSP

The contracts negotiated between NSPs (Network Service Providers) and CDNs within a confederation will usually not be all the same, so this needs to be taken account of in CDN selection. Contracts will typically be based on the number of open ports rented by a CDN Service Provider with each NSP. Moreover, if an operator's own CDN is in the CDN list, its subscribers will be best served by that CDN.

4.5 . Dynamic quality of service

Quality of service is obviously another key criterion for CDN selection, although one that again needs to be balanced against some of the other factors such as cost. It is very important that selection on QoS be based on real time feedback from the heart of the network and not just from the device itself. In the next chapter, we look in more detail at the role of analytics in measuring QoS.

4.6 . Three additional use cases

We have identified three use cases that exploit the availability of multiple CDNs to improve the overall business case of the confederated CDN model. Load balancing

4.6.1 Load balancing

We have identified three use cases that exploit the availability of multiple CDNs to improve the overall business case of the confederated CDN model.

The ability to balance loads across several CDNs can enable content owners or service providers to improve quality of service, reduce costs and ensure maximum availability. It can ensure continuity of service in the event that a single CDN goes down. There is also the ability to assess the relative performance of each CDN for a given type of service, such as live streaming, without at that stage relying totally on it. In basic operation,

load balancing mode would simply assign specified percentages of requests to given CDNs, determined by set rules.

4.6.2 Cost control with quotas

In some cases CDNs impose quotas on customers, rather like consumer data caps from broadband or mobile service providers. The additional bandwidth can also be more expensive. By modeling their contract with CDN Service Providers, content owners can ensure they do not exceed their quotas and as a result control costs.

4.6.3 Fail over / Disaster recovery

The ability to route between multiple CDNs gives content owners or service providers some fail over or disaster recovery capability. CDNs should have their own internal fail over procedures but the confederated model enables an additional level of protection through being able to switch from one CDN to another.

The CDN approach of a leading European TV broadcaster

In a recent public presentation, the chief architect of a leading European TV station's OTT streaming platform explained how they came to look to CDN selectors as a way to meet their distribution challenges.

"Clients pay for our content and they don't care if it is carried on OTT networks or not. We are simply expected to deliver a live TV experience in line with what satellite TV can offer. OTT head-ends can have trouble keeping up with the growing number of devices, channels and rising security requirements. User experience is then more and more under threat especially whenever a live football match aired."

The architect then went on to explain how they face these challenges.

"Our early experience showed that encoders should just encode, packagers should just package and encrypt and that Origin servers should also deliver DVR services and playlist manipulation."

The physical location OTT head-ends turned out also to be critical and they should be where connectivity is best and cheapest in a Net-Center.

The architect went on saying "We became agnostic to CDN providers and now believe in choosing the best one at any point in time. Indeed a major iOS release can impact the service as much as a live football match and different problems can require different solutions." Beyond merely caching video content, the TV operator found they also had to pre-fetch and cache licencing data otherwise the licence servers would be overloaded

5 Measure the quality of service for video delivery

As video services have grown in sophistication, it has become more important to measure QoS as perceived by the end user. But since end user QoS is determined by factors beyond the control of CDNs, it is important to base CDN selection on hard QoS data obtained from probes deployed in the network.

5.1 . What can be measured

The aim is to collect a basket of key metrics that between them cover all the principle aspects of a service, including video and audio quality, along with responsiveness and availability. Key measures are:

- Start-up or channel-change time
- Services' availability and continuity
- Latency (or responsiveness)
- Video or audio impairments

1 Video resolution and frame-rate

5.2 . What needs to be deployed for measurement

Monitoring information comes from two sources, the client device and the network, which need to be combined to provide a full picture of QoS.

In the client, QoS information is gathered by agents, or embedded software, measuring various parameters relating to end user experience and reporting them back to a central monitoring system. These parameters may include channel change time and responsiveness to user requests, as well as increasingly some measure of the subjective quality.

Historically end user video quality was expressed as a MOS (Mean Opinion Score) determined by user panels, but this is no good for assessing how QoS varies over time. Now some specialist monitoring vendors have developed algorithms that attempt to assess the subjective quality and provide a rolling MOS score. One such vendor provides video and audio MOS in real time via its client robots, which detect artifacts and combine these with a subjective factor derived from tests with humans, to provide continuous QoS feedback at the device level. This is based on the ITU-R Rec. BT.500 standard.

While client agents provide a view of QoS at the individual user level, probes in the network measure the underlying issues that may cause impairments in the first place. Key parameters measured by network probes

include jitter and packet delay variation, as well as packets lost, reordered or duplicated. These can pinpoint problem points such as areas suffering from congestion or link failure.

The key point is that while client based measurements may identify a service failure, they require integration with data from network probes to locate faults and so make correct CDN routing decisions. It may be that a fault is traced to a particular operator CDN that is part of a CDN Confederation. In that case the indicated action for customers of that CDN confederation would be to switch to a different CDN.

5.3 . What is relevant for selecting a CDN

The performance of a CDN is only one factor determining the QoS experienced by the end user. Other key factors are the state of the last mile, the home network and the client device itself, which are beyond the control of any CDN within a confederation. For example several people in a home may be streaming video content at the same time and therefore getting reduced resolution as access bandwidth is squeezed, which has nothing to do with any CDN. For this reason measurements taken by agents in the end user device are totally irrelevant for selecting the best CDN at a given time, despite their importance from the customer's perspective.

A good end to end QoS monitoring solution therefore measures the QoS in the heart of the network and not only on the device. Basing selection of the CDN on the QoS measured in the device alone would be no better than throwing a dice. So while the monitoring solution can provide valuable information about service consumption and general QoS trends at the level of the end user, this is not what is used for selecting the CDN.

At the same time the monitoring solution should be independent of the selection tool, since the choice of CDN will be based on other factors we have already described, as well as QoS feedback. The monitoring solution should though provide a variety of information that can be applied for both real time service management and trend analysis, such as type of content and location of users, along with status of ISPs and CDNs.

Example of an active Telecom Operator: SFR

> *We were fortunate in speaking in the fall of 2013 to SFR's Bruno Beaugrand, who is Marking Director for Cloud & CDN at SFR Business Team that serves over 165 business clients.*
>
> *Beaugrand told us that SFR's Operator CDN initiatives are part of the larger IAS (Infrastructure as a Service) trend and more specifically part of the Storage As A Service (STAS). SFR originally used Akamai, but early business planning showed that the business model was unsustainable.*
>
> *"In 2008 the decision to start building our own CDN infrastructure was taken. Merely caching our website http://www.sfr.fr was getting too expensive as we were already spending tens of thousands of Euros per month on that alone." SFR's CDN was initially built for internal use and just after that demand for OTT multi-channel, multiscreen traffic appeared.*
>
> *In September 2012 Bruno Beaugrand took over CDN business development and launched a packaged commercial operator-CDN Service.*

"We use open-source Varnish caching technology for this project where we have 12 heavily meshed major PoPs in France. SFR is a powerhouse in France as the ZDNet's published Meteo du CDN shows. We are convinced that a multi-CDN approach is needed internationally. Our clients, like Cedexis for example, would benefit from using a selector such as that offered by Broadpeak. CDN selection can be made on behalf of the content provider, or we can bring the tools for them to make the selection themselves". Other typical clients of SFR are CDiscount, Leroy Merlin and M6.

"SFR is a challenger in the global CDN space and so must differentiate with open APIs and better local meshing. We believe agile development will let us follow market demands rather than in having a heavyweight long-term roadmap."

When we asked Bruno how a multi CDN world might work, he told us "hardware makers like Cisco have tried to offer Operator-CDN federations before, but lack of standards hindered those initiatives. It might have worked technically, but the provisioning, reporting and billing were not standardized. So for example if a service provider wanted to implement a policy lowering the price during a specific period, there was no way to propagate that down the whole value chain. The way forward must be based on simplicity and pragmatism. Our content provider clients are not asking for overly complex CDN features, but they must have access to simple decision trees that ensures they keep their freedom. In this light, we believe less in a CDN federation than in a CDN selector approach, that makes it easier to change CDN provider".

6 Customer focused architectures leveraging retail devices

The winning strategy is to put customers in control of their service delivery, whether these are IPTV, Cable, OTT operators, or content providers. This means firstly being able to select among multiple CDNs over the end to end path beyond the operator's own network. Secondly it means extending the CDN over the operator's access infrastructure right to the service end point, the gateway in the home.

Then IPTV, cable or OTT service providers can propose a cost effective and high quality operator CDN solution to their content provider customers. It allows operators to harness the control they already have over their own network with their multicasting capability and extend this to the home gateway by implementing innovative technologies there. This enables operators to both optimize cost and quality.

6.1 . Extending the CDN to the home

Broadpeak launched their nanoCDN™ approach over a year ago. Many other vendors are also starting to talk about extending the CDN to the home gateway to optimize the quality of video delivery on operator networks with better performance than any CDN service provider. Users experience better QoS because the adaptation based on ABR technology now occurs between the client device and the network edge, rather than all the way

to the central server. This brings multicasting right down to the network edge, with ABR adaptation only occurring from there out to the device. This ensures that issues in the heart of the network that may reduce streaming performance less affect users.

Akamai's Julien Privé said in Paris recently that the CDN giant would need to leverage multicast infrastructure to follow the demand for OTT streaming. The company's CTO was recently quoted as having said that Akamai must also "turn consumer devices into content nodes".

Operators can contain costs with this extended CDN approach that scales much better as the number of devices grows, by ensuring that the amount of bandwidth consumed in the core network is constant no matter how many users are watching it. It makes live OTT video delivery to any device truly scalable by extending the CDN infrastructure as far as the user's home gateway. There is no need to add more servers or to dimension the system for peak capacity. Hence costs are controlled and the operator can enable its customers to benefit from very competitive CDN prices.

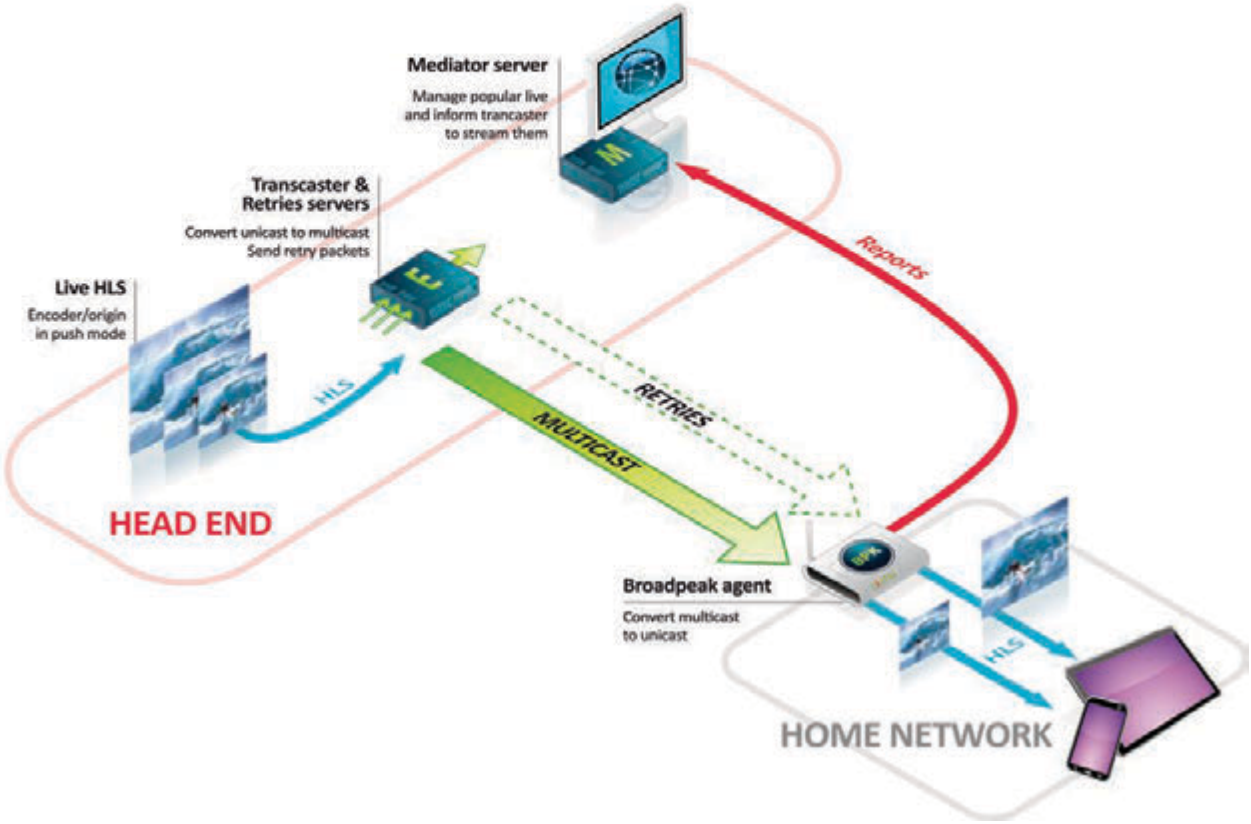

Broadpeak's nanoCDN™ Architecture

6.2 . Using the umbrellaCDN concept

An umbrella CDN approach allows content owners to determine the CDNs they want to use under specified conditions. They can select operator CDNs that use advanced technologies in their region of focus.

With an umbrellaCDN approach, content providers and pay TV operators can select the CDN that best addresses each request for live or VOD content. This is done using rules for assigning a CDN to a session that are easily defined through a simple graphical interface. The process can then be tracked on a monitoring dashboard.

Such rules can be based on a wide variety of factors, including:

- End-user geo location
- End-user Internet Service Provider (ISP)
- Content category (e.g., live/VOD, pay/free, premium/trailer...)
- Time of day
- Quality of service, through feedback provided by an external probe on a location/CDN/NSP combination

Any combination of these parameters can be used to create rules that meet particular needs

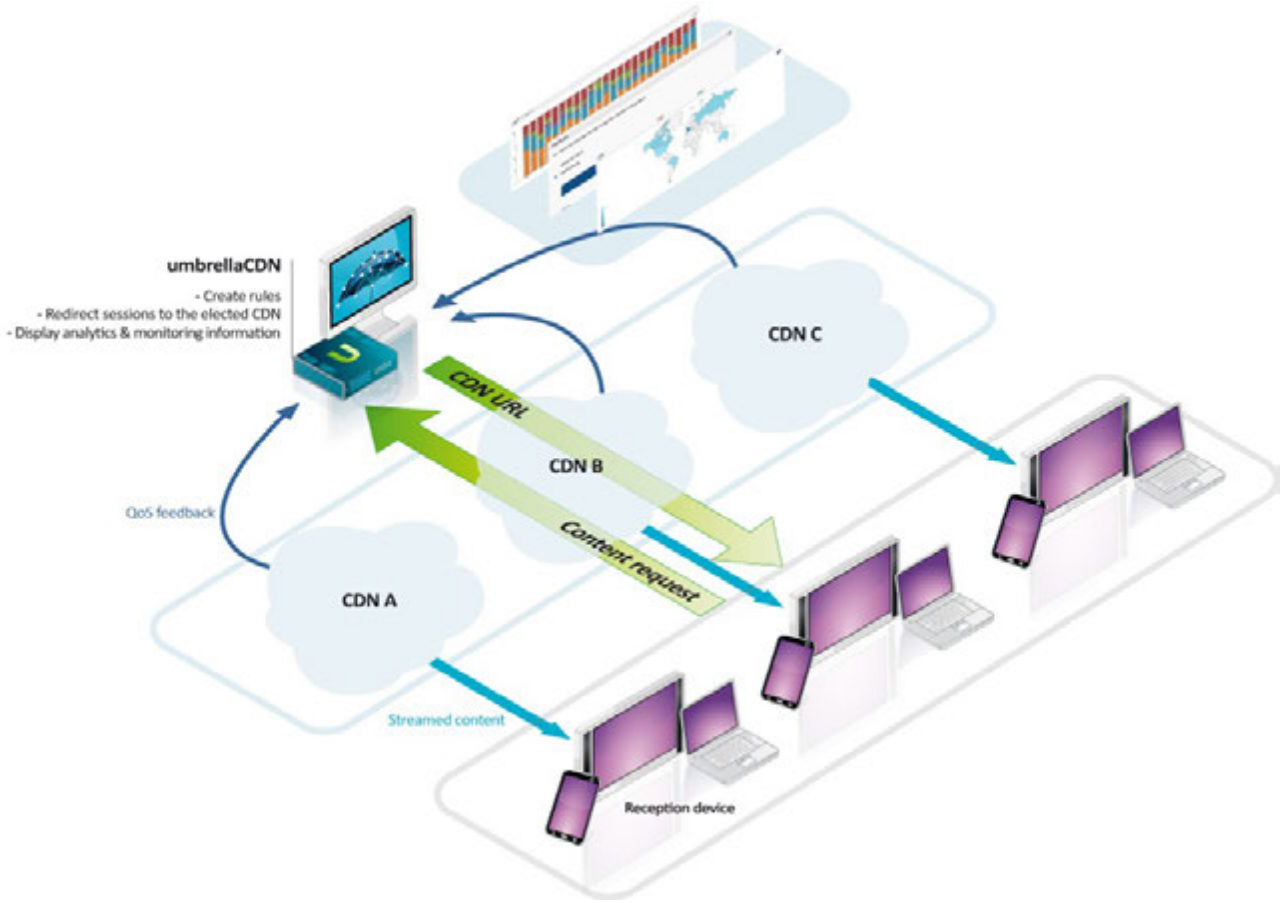

An umbrella CDN Architecture as suggested by Broadpeak

7 Typical use case

We present a scenario that can be implemented with an umbrella CDN to route either live or on-demand content within a confederation comprising five different CDNs. This use case features as content provider a French TV channel that offers live and catch-up TV content through both a web portal and a mobile application.

The catch-up TV service is available worldwide, but the live service is restricted to French Territory, including Overseas Departments and Territories referred to as the DOM-TOM. This means the channel has to deliver content to customers in three types of location, mainland France, the DOM-TOM and the rest of the world.

To meet this requirement, the channel has signed a contract with two global CDNs, which we refer to as GCDN A and GCDN B, with an operator CDN called OCDN and with two French CDNs, which we call FCDN A and FCDN B. The OCDN, FCDN A and FCDN B are all confined to mainland France. The channel also has contracts with a QoS information provider that has deployed its probes across the whole French territory, including DOM-TOM.

The channel has implemented the following rules in the umbrellaCDN for selecting CDNs, dependent on the location of the customer:

	Mainland France	DOM-TOM	Rest of the world
Catch-up TV	Choose between FCDN A and FCDN B according to which offers the best quality at the time.	Use GCDN A for 50% of the requests and GCDN B for the other 50%.	Use GCDN A for a maximum number of 10,000 sessions per day, then switch to GCDN B when the number exceeds that.
Live	- If the end-user is a subscriber from operator O, use OCDN that has deployed the extended for live content. - Otherwise, use whichever of FCDN A and FCDN B offers best quality.	Use GCDN B	Use GCDN B to send replacement content.

8 Characteristics and key benefits of the umbrella CDN approach

The benefits of an umbrella CDN approach are control, simplicity and visibility. It should be straightforward to deploy and enable centralized content distribution, monitoring and analytics.

Key characteristics to look for include:

- Deployment: Lightweight integration with CDN service providers, no specific client on end-user device.
- Open QoS/QoE criteria: interface with external vendors to assess CDN quality in a given location.
- Control over allocation rules: operators need be able to choose CDN on a session basis.
- System monitoring: operators need to check the application of rules they've defined in real time.
- Dynamic provisioning: content should only be sent to CDNs that are going to stream it.
- Centralized geo-blocking: regions with different rights are defined at the umbrella CDN level.
- Centralized content source redundancy: umbrella CDNs can switch to alternative content sources.
- Unified audience and content analytics: session requests are centralized and monitored in real-time.