ANEVIA
INNOVATING THE VIDEO EDGE

WHITE PAPER

EDGE-PACKAGING WILL BE KEY FOR GUARANTEED LIVE OTT DISTRIBUTION WITH CDNS

Three takeaways from this white paper:

- **Retail devices are rapidly changing, often requiring adapted or new content formats.**
- **Operators are struggling to produce flexible and scalable "future-proof content".**
- **Today, processing at the network edge is the most cost-effective way to achieve this.**

Intended Readers

This White Paper is intended for technology readers seeking some in- depth business analysis and vice-versa for business-focussed readers wanting to learn something about OTT streaming technology.

Subject Matter

The industry has been stumbling in the dark trying to enable high quality OTT delivery, especially for live streams. This paper will show that there is light at the end of the tunnel and provide a blueprint for how OTT distribution will eventually extend the supremacy of broadcasting for live content. We will explore the history of OTT content distribution and then, after looking at the big picture, will drill down into the critical video content packaging function – including its risks. Then we will demonstrate how positioning this function correctly within the CDN will prove the most cost-effective way of getting to that light by 2015. While new technologies may appear to undermine existing broadcasting incumbents, by embracing emerging CDN and delivery technologies wholeheartedly, they can convert these threats into opportunities. Many hybrid broadcast solutions have already shown this.

Contents

1. Executive summary

After some years of uncertainty, OTT content is clearly here to stay. But as it takes a growing share of the distribution pie, we can tell neither how fast nor how far this market penetration will go. Therefore, it is vital for operators to deploy flexible solutions that can be scaled up to cope with future levels of demand that can as yet only be guessed at.

This White Paper explains the technical and business context of OTT distribution. We look into the drivers of OTT from different stakeholder perspectives, arguing that the exploding variety and quantity of retail devices is a common challenge and opportunity. The paper contends that wider adoption of OTT for live TV consumption will bring disruption to the first generation of centralized OTT distribution architectures. Core video head ends can still be centralized, gaining economies of scale across IP and different DVB broadcast technologies. But only edge processing and in particular edge packaging will sustain anticipated growth of live content while bringing exciting on demand features such as session based watermarking that will enable new business models. Watermarking has been used some years in the broadcast environment and is now mandated by Hollywood studios to protect content from illegal redistribution over the Internet. It involves insertion of indelible code into video frames that enables individual streams to be traced back to their source, without any discernable impact on the user experience. Watermarks can be inserted at the client during decoding and decryption, or at the server end, which has the advantage of being more secure and identifying content streams at source.

With a case study from the French Pay TV operator CANAL+ we take a high-level view of the business case for the French pay TV market. Going beyond the description of the technical value chain of OTT distribution, we explore different architectures for network and content operators focusing on scalability and security issues these new approaches bring. We look at the different locations a packager can have within the value chain, explaining why it should ultimately reside behind the Origin server, at the network's edge. This paper wraps up looking at the unique advantages of edge packaging and concludes that the key one is the promise of reliable OTT distribution of live content.

About the author

Damien Lucas
CTO & co-founder Anevia

Co-founder of Anevia with Tristan Leteurtre, Damien Lucas is an expert in video streaming technologies for fixed and mobile networks. Currently CTO of Anevia, Damien also comes from the VLC Media Player development team and graduated from Ecole Centrale Paris, France.

2. The origins of OTT content distribution

It has now become apparent that OTT distribution is here to stay. To better grasp where it might be going let's briefly look back to where it has so very recently come from.
No pun is intended in the title here, we will simply explore where OTT content distribution came from. In chapter 4 we will describe the technical setup and explain what an "origin server" is.

The meaning of the phrase OTT

"Over The Top" dates from World War I, when it was a reference to soldiers having to go over the top of the trench walls. It's probably no coincidence that it was then associated with danger and getting into the line of fire. Our definition today dates back only to 2007, when the first iPhone was launched, with the phrase only gaining widespread usage across the industry in 2009. Even in 2007, when using an app such as YouTube on the iPhone, video was delivered across network resources that were provided neither by Apple nor YouTube (which had just been acquired by Google at that time). Skype is another early example of an OTT service from that era, while Netflix is today showing that this model is far more than a flash in the pan with over 40 million clients choosing to pay about $100 a year to subscribe to that service.

Anevia, author of this paper, takes pride in having published one of the very first White Papers on OTT streaming back in 2010 (ref. 8.2).

Until recently, Telcos and ISPs attempted to confine direct support solely to services that exclusively ran on their own networks. This was for two primary reasons. They could more easily guarantee high quality services over limited bandwidth and they did not want to effectively subsidise delivery of third party content over their own network infrastructures.

However the industry is changing fundamentally, partly for business reasons and partly through maturation of technologies related to adaptive bit rate streaming. Offering one's own services to competitors' clients has become the Telco's dream. As one example, French Telco SFR's Home security package is available to any broadband subscriber in France. Croatian incumbent Telco T-Hrvatski Telekom launched a similar service as early as 2009. In the same vein, IPTV operators are providing content apps that run on connected TVs and so work not just on their own networks, but also their competitors'.

Content owners, like pay TV providers, are moving to OTT models but with even greater conviction. Content originally broadcast with DVB technologies over satellite, cable or terrestrial networks is now also becoming available to apps in consumer devices that only have IP connectivity. Examples abound, including US cable operator Comcast's Xfinity Streampix, or French Digital Terrestrial operator TDF's catch-up TV service Salto.

From an operator's perspective, OTT can be Inbound or Outbound. Inbound OTT allows the operator's subscribers to access an external service provided by third parties, with or without the operator's blessing. Many operators are still sitting on the fence, happy to enable access to a YouTube service for User Generated Content (UGC) like funny cat videos, but determined, for example, to block pirated premium HD content. Outbound OTT is the more recent phenomenon where content and service providers alike are engaged in a race to make their own offering available as widely as possible, beyond the reach of their traditional distribution network. The technologies used for Outbound OTT, originating mainly from the web, also give greater reach to operators' own services within subscriber homes. So, for example, when a TV operator develops a mobile app to offer its live and on-demand content, for example on a tablet over Wi-Fi, subscribers not only get to use the service on the move, but also in their gardens, or parts of the home where there is no access to broadcast TV.

3. How OTT has become so central to the Video Business

3.1 The drivers pushing OTT today

Explosion of devices

Attempts at launching connected entertainment devices predate even the Internet. But apart from a few remarkable but isolated technological achievements, such as the Danish Kiss Technology connected DVD player in 2003, Steve Job's launch of the iPhone in 2007 marks the debut of the modern era of OTT devices.

The iPhone 6 should be launched within a few months of this paper's publication, yet all devices back as far as the iPhone 3GS must still be supported by operators, amounting to six different models in total. On the iOS tablet side there are already four major variants of the iPad to support. These 10 different platforms must all be tested when any new service or even a minor upgrade is released.

Since Apple's ecosystem is closed, this testing remains relatively easy, given a limited number of variants. But there is now a plethora of connected TV platforms to test as well. Furthermore, when considering the Android market, the situation is much more complex because of the more open nature of the OS. This openness has led to Android's success in overhauling Apple in the smartphone market but it has made testing for operators more complex because each manufacturer tweaks the operating system, both with specific drivers at a low hardware level and usually their own enhancements to the user interface at the top-level. For example Samsung, the largest Android licensee of all, has put its own stamp on the Android look and feel, having developed its own UI called Magazine UX sitting on top of the operating system. With a dynamic dashboard and its own app short cuts, this represented a dramatic departure from the Android vision at a time when Google was attempting to reduce the operating system's fragmentation. So for operators, we have reached the point where a quick software bug correction that could take just an hour to remedy might take a whole quality assurance team weeks for all the required non-regression testing.

All these OTT devices can be categorized in various ways including screen size, processing power and memory. This is not the place to discuss these in detail, other than to highlight the need for operators to adopt flexible streaming infrastructures that can cater for all of them.
A number of important device platforms were developed just for proprietary environments, like the Microsoft Xbox game consoles, and only work with one of the several widely deployed ABR streaming protocols (Smooth Streaming in this case). Although Microsoft has been committed to MPEG-DASH for some time, earlier Xbox 360 consoles still run in 10s of millions of living rooms and only support up to version 2.0 of Smooth Streaming Media Element (SSME 2.0). Although the newer Xbox One doesn't have this limitation, it illustrates the legacy streaming challenge already faced in this still-young OTT video field. It has become impracticable and unaffordable for operators to continue supporting all these variants on a case-by-case basis, involving preparation and storage of all the required versions and formats.

The hopes for a common format throughout the market have so far been dashed. Microsoft's Smooth Stream and Apple's HLS are still being installed [in earnest] with new deployments throughout the world and even if it has lost some ground, the Adobe solution is still quite prevalent in the field. MPEG-DASH might still save the day and become the Lingua Franca that the industry needs for OTT encoding and delivery. Yet this isn't happening fast enough for operators faced with deployment decisions in 2014 and 2015.

The solution is for operators to make use of CDNs to ensure adequate end to end quality of service. As currently deployed, CDNs lack the ability to distinguish between traffic in order to give premium video such as a live football match

priority over say an OS update. Fundamentally, CDN technology works by delivering lots of small data files with the best compromise between processing cost, bandwidth efficiency and quality. We will explore in greater detail in the next chapter exactly how Edge Packaging takes that forward and how the positioning of this key component in the technology chain is critical. Packaging in the CDN alleviates much of the processing load and, by being executed in real-time with only an 8ms delay, adds little to the overall video processing time of typically 5 seconds. Although measurable, this extra delay is unnoticeable in most use cases.

In addition Edge packaging can already distinguish between traffic types. An edge packager's main role is precisely to package managed-video, so almost by definition has the ability to distinguish between this traffic and the rest. Once identified the traffic can be managed appropriately. From a technical perspective this could mean adapting the ABR manifest file in the edge server in case of congestion or edge-processor overload at that part of the network. From an end-user perspective this would mean that users could be prevented from accessing higher profiles. Damien Lucas, one of Anevia's co-founders, states that "we can't actually block say a gigabyte iOS update occurring during a key football match because of net neutrality rules, but we can ensure that the football match isn't interrupted by bringing users down to a sustainable bitrate."

The explosion of OTT offerings

As we saw in section 2, OTT was initially just free IP telephony from a Telco perspective and YouTube from a content operator perspective, representing only a small threat or opportunity and not much to get worried about. It has now evolved into a wide range of offerings from both Telcos and content providers. Video is core to any OTT offering and because unmanaged networks are used, Quality of Experience is a key differentiator.

While ABR may meet the challenge of content delivery over the last mile, it doesn't prevent edge-processing overload or connectivity issues between the network core and edge. With Edge packaging, operators can change rules to avoid overloading either CPUs or the pipe, e.g. by moving down from say 9 profiles to 3 during a major live event like a national team football match. This can be done automatically in real time, e.g. as soon as 70% capacity has been reached, or alternatively it can be manually scheduled.

The strategic importance of OTT

For network operators

The industry is rife with stories about how network operators like Telcos and ISPs are forced to carry ever increasing amounts of traffic for services like Netflix that are based on video streaming. The business risk is that they derive no revenue opportunities from subscriptions, transactions or advertising, but on the other hand they face losing broadband subscribers if they attempt to block or slow down such services. This issue also evokes the net neutrality debate, which is painful for operators and beyond the scope of this white paper.

There are also corresponding business opportunities through negotiation between OTT providers and traditional pay TV operators. Netflix, villain of the previous paragraph, is the hero of this one. The OTT streaming company is paving the way for a new business model through a watershed interconnection deal with Comcast, the largest US pay TV operator. Although the precise terms of the deal had not been revealed at time of writing, it is clear that it involved Netflix locating storage caches inside the Comcast network, reducing peering and improving the user experience for their subscribers. The service will resemble a halfway house between OTT and IPTV in effect, with Netflix having visibility over the whole end-to-end path, while Comcast derives some extra revenue. Netflix followed up by striking a similar interconnection deal with

Verizon in April 2014, suggesting that there would be more to come.

For TV operators

For operators of pay TV services, the main risk of not getting onto the OTT bandwagon is being cut out of the content loop. They risk being dis-intermediated by content owners going directly to their subscribers using their own OTT services. Numerous examples exist already such as the BBC's iPlayer or HULU in the US. There is also an increased risk of piracy through content redistribution over the Internet. While all forms of pirate operation threaten pay TV operators, the more recent breed of pirate OTT services presents a greater risk to the TV operators' own OTT efforts.

Success stories like Netflix illustrate that simplicity and convenience are key but moreover, as the deal described with Comcast illustrates, OTT streaming needs to work all the time for continued success. OTT is also a fantastic opportunity for TV operators. Indeed "TV everywhere" used to be on TV operators' long-term roadmap. OTT technology makes TV Everywhere readily accessible as the underlying ABR technology makes it feasible to deliver TV streams to many more devices over unmanaged networks with varying levels of performance and congestion. This "TV Everywhere" opportunity is synonymous with multiscreen.

In the next parts of this White Paper we will explore what content packaging is from a business and technical perspective with a discussion of pipes, content, devices, ABR technology and the central role of CDNs.

3.2 High-level business case of the French pay TV market

French TV operators like CANAL+ currently pay CDN providers such as Smart Jog, Akamai or Level3 to deliver their IP content via 5 different large-scale IP networks, each of which has from 1 to almost 10M subs. There are also as many operators with less than1M subscribers that still have pay TV potential.

However the IP networks still bear the traffic cost. We are convinced that business logic will eventually change the market dynamics so that the TV providers will reduce dependency on big CDN providers and also bring their content directly to the operators, probably for half the cost.

Net Neutrality is an issue that must be addressed here, as we can see by the forceful reactions to the Comcast Netflix deal in the US (see ref 8.1.3). The big question to be addressed by all stakeholders, including the regulatory bodies, is how to guarantee OTT traffic delivery. For an operator like CANAL+ to fully cover their home market of France, fewer than 10 operator deals would be required.

The ease of use of the technology and the deployment challenges depend on many factors, of which the number of channels and users are key to determine the optimal architecture. In the case of nPVR Opex can be three times lower in the cloud, as edge-caching doesn't make any sense for content destined for a single user, so that everything can be managed on the Origin server. Scalability for peak usage remains cheaper in the cloud as one operator's peak requirements will differ in timing from another, allowing the common infrastructure to achieve economies of scale. Nicolas Carr first made the case for cloud as the future of IT persuasively in his 2009 book The Big Switch (Rewiring the world, from Edison to Google). It draws the parallel between the commoditization of IT infrastructure and its transformation into a utility, as happened with electricity a century ago. OTT TV distribution uses IT technology, which is migrating to the cloud anyway.

Two of the more critical challenges faced by operators needing to deliver best-in-class services quickly are PVR and multiscreen. These are naturally implemented with a cloud

infrastructure. In the case of the nPVR, the cloud easily addresses the twin current issues whereby a minority of heavy users find their PVR capacity far too small while a majority of hard disks are underutilized, standing largely empty in users' homes.

4. The technical value chain of OTT distribution

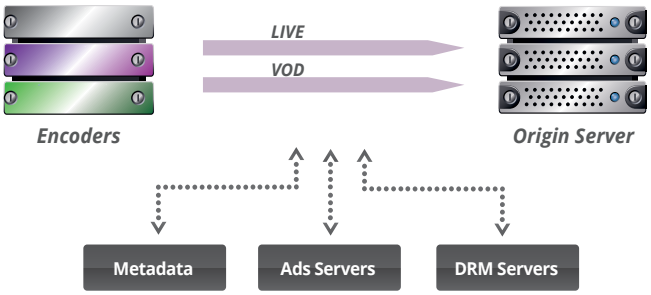

The above diagram shows live an on demand video content being sent to an origin server that will enable its distribution using ABR technology. The content packager is not represented.

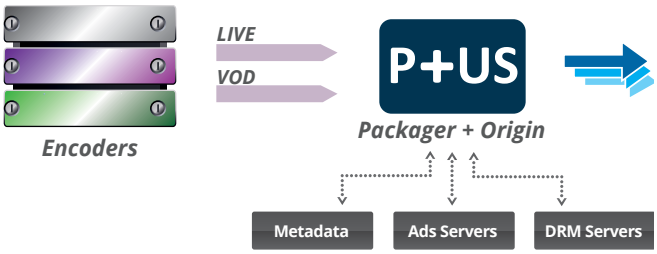

In this diagram we see the content packager component appear as part of the Origin server. Before honing in on the critical packager component, we take a look at what the OTT distribution concepts really are. In Anevia's first 2010 White Paper, OTT streaming was defined as "the delivery of video streams to different types of device via the Internet. Unlike traditional IPTV, there is no need for a dedicated network or infrastructure provided by the operator, as OTT is transported through regular Internet data protocols and uses the open Internet over unmanaged networks."

This definition assumed that existing CDN and other caching infrastructure would carry ABR streams without any extra effort. This is still largely true, although it does not take into account the explosion of the number of devices and hence formats, as well as the sheer number of live TV channels operators want to deliver. In the case study of CANAL+ in chapter 6.1, the French operator wants to deliver 100

live channels already. So the situation is now changing with the growing role of CDNs coupled with partnerships between ISPs and OTT operators, as witnessed by the Netflix/Comcast partnership discussed above. The key change in architecture under the hood since that 2010 paper is the increased importance of the Packager component.

What is a Packager and how did it become so critical to success?

A packager in an OTT infrastructure should require just a single copy of the video stream to be carried. Then, by repackaging that stream at the edge of the network, the appropriate video format can be reproduced to fit any device. "Packaging on the edge" massively offloads the CDN and thus reduces bandwidth requirements.

A state-of-the-art packager must work for both online and offline content in real-time. It should turn a pivot format into the desired OTT formats to the origin servers. A pivot format, often called a mezzanine format, is an innovation that has evolved largely for OTT distribution. It is an intermediate master format at a resolution between that of the contributed content and that required by the end devices, avoiding the complexity of encoding at multiple resolutions and bit rates. The pivot format is usually encoded at 5 to 10 times the bit rate of the highest resolution target device.

Some key criteria to differentiate Packagers and choose one that is future-proof

To ensure it stands the test of time, a packager should be a modular part of a complete OTT ecosystem. It is important that it remain interoperable with multiple security systems. Indeed in the real world example of CANAL+ described below, three different DRM solutions are supported (i.e. FairPlay, Nagra & PlayReady). Different encoders must also be supported, so if an operator is seeking this best-of-breed goal they should ensure different vendors are used for the packager component and the encoder component even if a single integrator is used. In the same spirit, pre-integration of different CMS systems is preferable.

A track record of innovation is also a key attribute to look for here, especially while the principle standards like MPEG-DASH are still evolving. Alongside this capacity for innovation, the OTT infrastructure must be able to scale up to accommodate the growth of TV delinearization. Pundits mostly agree that while 86% of TV is still consumed at the time of scheduling (see ref 8.1.2), on demand is growing fast, although with disparate views over how quickly this will happen or how far it will reach. This is why flexibility is so critical for operators who must make decisions today. For nPVR to work in a multiscreen environment service providers must ensure that a recording made today remains viewable in say two years time. Currently the best way to attain this objective is by using a Pivot file format and On-the-Fly to achieve "future-proof content". Multiscreen is actually multi-device requiring multiple formats and protocols. A Pivot format is the only cost effective way to address this. In this way, future formats that may be required for devices yet to hit the market will be compatible with today's content.

Ecosystems that use a standards based approach can enable the next generation of OTT content distribution to scale further still with a CDN-Federation. Here operators will be able to exchange CDN data in the same way that voice minutes have been exchanged in the past. CDN technology must optimize these costs that will increasingly affect the bottom line. As an example, the recent deal between Comcast and Netflix is said to involve annual payments of several million dollars (see ref 8.1.3).

A brief history of packaging, from the head end to the edge

The earliest OTT encoding solutions were proprietary closed formats such as Move Networks. They showed promising results but yielded little commercial success. From 2008 the first dedicated OTT encoders that didn't require dedicated players started taking hold in the market. They were each dedicated to specific formats such as Apple's HLS or Microsoft's Smooth Streaming and used MPEG Transport just like IPTV.

In 2010, common encoding created some economies of scale, but hit scalability issues with the ever-growing number of profiles needed. Packaging was still concentrated in the video Head End.

In the fall of 2011 external dedicated packagers were launched, notably by RGB with their Just-In-Time or JIT packager. Externalising the packager in this way enables federated encoding for different distribution channels, i.e. Satellite, Terrestrial, Windows/Mac, iOS, xBox, etc.

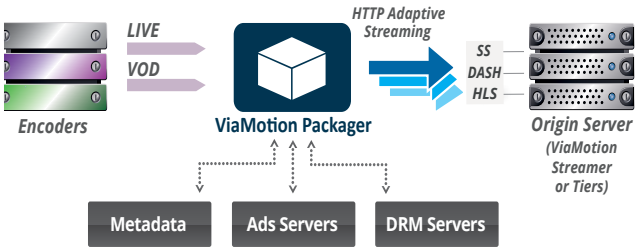

The packager then gradually moved from being positioned in front of the origin server to being integrated into the Origin server and finally to being positioned after the Origin server. The major advantage of the latter position is that the range of devices types can carry on increasing without needing to upgrade the whole head end all the time.

The current trend that we have started seeing in 2014 is for the packager to move into the CDN. This brings the key advantage of the ability to guarantee QoS for live streaming.

For operators that still have to go through this evolution, the watershed moment comes when the packager moves from being in front of the Origin server to being positioned behind it.

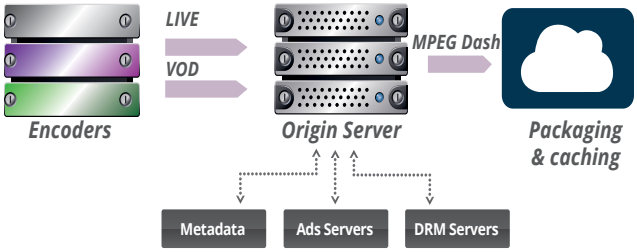

Benchmark data from an Anevia client on the scaling of OTT distribution

In this case, the operator supports ~200 live TV channels with a 30- minute catch up TV (CUTV) buffer and 10 different formats. The 10 formats are actually 5 formats * 2 DRMs. Using a Pivot format brings a factor of 10 saving in storage and bandwidth requirements.

This means using 128 Gigabytes of live memory instead of 1.2 Terabytes, so here two 64 Gigabyte servers are enough to power the entire origin server.

Storing CUTV over 10 days requires 1.5 TB per Channel. Considerable savings are achieved in storage even in the theoretical case limit where all users are watching delinearized content. Note that in such a theoretical extreme case, CPU requirements would remain the same.

For the on demand part, if there are already over 20k video assets, and when a new device or format comes out, it is no longer feasible to reprocess and store a new version for each. The pivot format associated with edge packaging means that the operator simply updates parameters, or in some cases upgrades the packager software.

A telling example of operational hassle avoided with a pivot format is illustrated by the need to maintain Smooth Streaming 2.0 support to ensure Xbox compatibility, as the Xbox 360 doesn't support the more recent Smooth Streaming 2.2 version.

5. Different architectures for network and content operators

This paper contends that edge-processing is the future for OTT delivery of live services. In the real world, many operators will need to build a straightforward setup to start with, positioning all their OTT infrastructure in their video head-end, which is connected to the network core. However, packaging is a relatively complex part of the ecosystem so it is advisable to start with a flexible technology from the outset in order to adapt processing and configurations as the service matures and needs to scale.

5.1 Scalability of OTT architectures

The benefits of HTTP based streaming are clearly laid out in many other White Papers (including Anevia OTT Streaming paper Ref 8.3), so here we will focus solely on scalability. The streaming protocols are all based on Adaptive Bit Rate (ABR). This adaptation to the available bandwidth and processing capabilities already enables a first level of scalability, as services can be maintained even as resources get scarcer. What really enables OTT live TV streaming to scale dramatically compared to earlier RTSP based approaches comes from the way HTTP streaming mimics multicast solutions.

Multicast is an IP network's way of broadcasting, where a stream is only carried once on any link in the network. This Multicast mimicry is achieved through various features. The two most important come from avoiding the need to maintain a central server session for all active streams and from the chunking and packaging of streams that allows cache and CDN infrastructure (already in place for Internet data distribution) to be reused for OTT video with only minor adaptations, if any.

With session based encryption, it is the edge-server that performs the actual stream encryption, while the key server remains an external component. In this respect, a packager in the OTT delivery world can be similar to a MUX in a DVB broadcast world. Both devices receive multiple elements and let some pass through while others are encrypted or filtered out. The scalability challenge introduced at the encryption level by edge processing is simply that key servers will need to interact with hundreds of edge servers instead of tens of MUX or encryption servers, as they would in a centralized processing architecture.

The real scalability challenge of the key servers and Subscriber Management System (SMS) remains the same. It is to manage a database of millions of subscribers with potentially multiple keys being used on different devices. Stephen Christian of Verimatrix confirmed, "the high-performance key management and delivery systems within our products fully support on-the-fly encryption solutions, including edge-packaging approaches. The same architecture that can scale to support security on multi-million client deployments also scales straightforwardly to manage keys for a local mux network or hundreds of packaging servers. Bear in mind that a distributed encryption solution implies putting a lot of trust in the delivery network and edge servers, which are then just as much part of the security envelope as the origin".

5.2 Edge processing protects against DDoS attacks

Relocating the initial point of contact to a service for subscriber devices from the network core to its edge enables the edge-servers to act as checkpoints for attacks. That way, the attacks can be deflected before they compromise the entire network or even a single origin server, so that they cannot bring down the services supported. CDNs need to be protected against this sort of attack in all cases.

The foremost risk is that of bringing the whole network down. Potentially just as harmful is the risk of blocking just one service by swamping an Origin server with requests in a Ddos (Distributed

Denial of Service) attack. The third risk level concerns the network edge where just one server could be brought down.
The top-level security concern must be to prevent malicious attacks from reaching the centre of the network. A more distributed architecture, where the task of handling requests for streams or chunks is pushed out to the network's edge, achieves this intrinsically, just through the network topology.

In the case of Anevia's edge packaging technology, tokens are used to sign an incoming request. This signature includes multiple elements based on date and time, the content being requested, subscriber information and the IP address of the client device.

Client software on a subscriber's device will add this signature to each request for a chunk. The Anevia edge server decompiles the token to determine if a request is legitimate. All standard DdoS attacks can thus be detected at the network edge because fake signatures will have the wrong IP address. More sophisticated DdoS attacks using IP Spoofing can also be detected by the network edge thanks to the other elements used for the signature.

Serving requests for chunks at the network edge is therefore a way to protect the CDN, limiting the worst-case attack scenario to just a few edge servers going down.

5.3 The importance of properly locating the packager within the value chain

Basic ABR, which is the cornerstone of OTT content delivery, cannot deliver Watermarking without edge processing. Edge Packaging offers the best solution for Watermarking, including user session based systems. At Anevia, we believe that as edge processing is inherently session based, it can resurrect the benefits of session-based protocols like RTSP without their scalability limitations.

Content is secured by encrypting it with a key obtained from the DRM key server. In the case illustrated here, the service platform (SPF) then issues a token to authenticate the client device so that the edge server knows whom to trust with a greater level of security than merely checking the IP address. With IP filtering alone tech savvy users can fool the system. Unless TR-111 or similar protocols are used, a server might view multiple devices from a single home network as a single client from an authentication perspective.

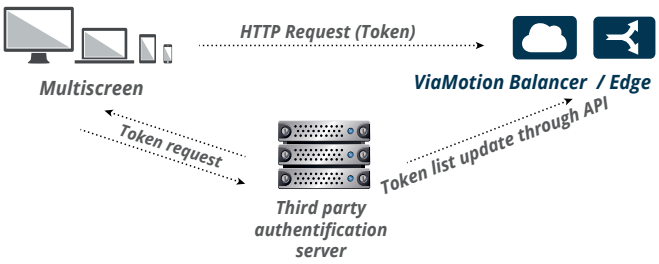

5.4 Security concerns

Concerns on the security of Edge Packaging …

Different stakeholders in the content value chain have different incentives. With content owners for example, the motivation to enhance the security of premium content supersedes most other requirements. This can lead to a desire for encryption to be handled as early as possible so that only encrypted content is distributed within the CDN. Some discussions are going on with studios to explore the possibility of having content encrypted right at the source. This extreme security concern can cause some mistrust in an edge-packaging architecture described in this paper. Indeed, the security components such as key servers need to send security keys to the edge of the network where servers may be harder to protect from malicious attack.

… are not justified in the long term

Broadcasters on the other hand need to balance security against other business issues. As we have seen, in the world of OTT video distribution, new devices are still coming to market and some of the streaming protocols are still evolving. The relative market shares of retail devices just three years from now is unpredictable and new device categories still appear regularly. A recent example of a new device category is the HDMI dongle. Each new hardware and software variation can create the requirement for the broadcaster to set up yet another profile for OTT distribution. This constant evolution is why using a Pivot format that can be encrypted, packaged and transcoded on the fly is the best way to guarantee that future requirements can be met.

In line with this movement towards standard encoding and packaging, the Common Encryption Scheme (CENC) has evolved as part of MPEG-DASH to reduce the cost and complexity of interoperating with different DRMs on target client devices. CENC specifies standard encryption and key mapping methods that can then be used by any DRM downstream to decrypt the given file. It defines a common format for the encryption related metadata transmitted to devices for decrypting the protected streams, but leaves the details of rights mappings, key acquisition and storage, and DRM compliance rules, up to the DRM system.

Meanwhile all the leading CDN vendors are introducing more and more security within their networks, so that even if an edge-processing device were in an unsecure location and sensitive data sent to it, that data would at least be lightly encrypted. The continuous monitoring that is now possible means that any security issue can be identified and remedied within hours. Nevertheless some service providers are still reluctant to decentralize the crucial key servers, out of concerns for scalability as well as security according to Christian van Boven, VP Multiscreen Product Management & Architecture at Nagra. “Service providers want to decentralize the packaging of content, possibly as far as the CDN edge node,” said van Boven. “But it does not necessarily follow that key servers should be decentralized, since these two elements do not have to be coupled.”

However van Boven does believe that key servers can be safely and successfully decentralized, providing care is taken to ensure both scalability and security. “On the scalability front, session based encryption makes it harder to build more resilient OTT infrastructures that can scale at peak times, for example major live sports events, which can entail predelivered licenses and keys,” said van Boven. “But this is solvable with the help of improving redundancy, load balancing and associated technology, although this may raise the cost.

In the case of security, the issue is that centralized key servers are normally deployed in controlled environments less vulnerable to physical intrusion. However, the advent of Hardware Security Modules can mitigate this problem, while more sophisticated models are emerging that make distributed architectures less vulnerable.”

Van Boven added that for OTT players, creating partnerships with network infrastructure providers might help address both security and overall OTT scalability issues.

6. Case studies

6.1 CANAL+ group

France's leading pay TV platform, the CANAL+ Group, has over 10 million subscribers globally and the operator's IP delivery channels are growing.

Philippe Rambourg is the architect behind the OTT streaming platform that is now being converged into a single global head-end where OTT will be used as a distribution channel alongside satellite (DVB-S), terrestrial (DVB-T), cable (DVB-C) and DSL/Fibre (IPTV).

He told us how the technical architecture is being set up to execute an ambitious plan for delivery over managed and un-managed networks. A key challenge that CANAL+ had to address with its OTT streaming architecture is the multiplicity of devices including "legacy", such as the older Xbox versions that don't support the latest Smooth Streaming version. This kind of issue will only get worse over time so the operator set out to build a future-proof solution. The retail devices fully supported already include iOS and Android (both phones and tablets), the Xbox, PCs and Macs, as well of course as a multiroom service via the operator's own STBs. The certain evolution of this environment requires a flexible OTT architecture, which is best met today with an on-the-fly approach so that terabytes of stored content need not be repurposed when a new device must be supported. In this case upgrading part of the server infrastructure is much more cost effective.

To ensure scalability and resilience of the streaming architecture that already carries over 100 live TV channels, CANAL+ has set itself three primary objectives. These are (1) for the complete independence of encoding components to continue to gain economies of scale, (2) the reuse of all input streams for multiple outputs, implying some sort of Pivot format with reduction in required storage space and (3) to remain agnostic to both CDN and DRM technologies, so as always to have access to the best technology available.

As a pay TV operator, CANAL+ is still facing two key issues."*Addressing the exploding number of device types that must often also be supported in different software configurations is our first challenge,*" said Rambourg. DASH may not turn out to be the lingua Franca CANAL+ had hoped for, as there are already several versions on the market. "CANAL+ already supports three separate DRM vendors (PlayReady, Nagra & FairPlay), offering little room for reuse or economies of scale".

The second challenge CANAL+ faces is in getting all the right content streams delivered through operator networks. Rambourg continued, "our domestic French market is unique in the high penetration of IPTV with 4 major fixed IP network providers as well as a cable operator and several smaller players. Operator CDNs would fit CANAL+' purposes in providing higher commitments to video quality than the global CDN providers can achieve, and at a lower price point. As CANAL+ manages both its own head-ends and the software on customers' devices, we will be able to get much more precise measurements to fully capture the true quality of service as perceived by our paying subscribers, using a software probe approach."

6.2 TrueVisions

TrueVisions, Thailand's main cable and satellite television operator and one of the region's more innovative players, has ambitious plans to remain the leading TV provider in its market as the region deploys more IP infrastructure. Because of this technical constraint the operator must deploy a resilient solution that can cope with these multiple environments flexibly. The use of a pivot format is one way TrueVisions sought to achieve this and support current and future requirements. Vichai Sernvongsat, the operator's CTO, shared his vision of the future of OTT streaming with us.

NDS is currently used to secure content that is broadcast both by satellite and by cable. He told us that TrueVisions is implementing a session-based encryption mechanism for OTT. Vichai told us this will both maximize the security at a manageable cost and also enable advanced features like watermarking that can be used to reassure rights-owners so that they will be willing to provide their premium content for delivery over IP.

The Thai operator's foray into OTT is primarily a defensive move to retain existing customers with what will become a TV Everywhere offering. But once this vision is fully implemented, the broadcaster sees the addition of OTT streaming as an opportunity to gain new customers that might not ever have become traditional pay TV subscribers. It hopes to gain new OTT customers both at the top end, with what will probably be 4K by then, and at the lower end with people perhaps taking just some on demand content. TrueVisions has separate platforms for traditional broadcast and OTT. However, as the number of devices continues to grow, there will be an incentive to unify them. TrueVisions already supports all iOS devices (iPhone, iPad, iPod) as well as both Android smartphones and tablets. The operator's OTT platform also supports the major browsers, Internet Explorer 9 & 10, Google Chrome, Mozilla Firefox and Apple's Safari browser.

In the short term, Sernvongsat told us that monitoring the availability of OTT streams for customers is still quite challenging. Considering the longer term, he does not see OTT streaming replacing broadcast in the region, as the infrastructure will not for the foreseeable future support the 100 Mbps or so required for this to happen. But OTT is a formidable companion to broadcast TV, expanding the TV product features and enriching the user experience.

7. Why edge-packaging now and in the future?

7.1 The unique advantages of edge-packaging

As discussed above, the most cost-effective way to implement session-based encryption with the technology available today is through the use of an edge packaging architecture. This offers several opportunities, including provision of the strongest security measures possible, which might be needed for early synchronous distribution at the time of theatrical release. Indeed, at the point of encryption for each client, fingerprinting techniques can be used so that in the event of any content piracy, the source can be rapidly identified and dealt with.

However, as we saw above, edge processing is inherently session based so at the other extreme of the security concerns, session based encryption will in the future offer lighter weight security solutions where all security issues are encapsulated within the CDN.

7.2 What's on the roadmap if you go for an edge packaging approach?

Packaging on-the-fly offers opportunities for both technical and behavioural analytics. The inherent session based approach means that details such as usage statistics can easily be collected at the user level, paving the way for true big data analysis.

The CDN federation idea discussed in chapter 4 is promoted by most of the infrastructure vendors. Cisco's François le Faucheur has been advocating the CDN federation concept for over three years (see ref 8.1.4) as a way to take the industry forward. The idea is that network operators are empowered by using their own local operator CDN infrastructure within a global service offering.

The content industry is facing a key question over nPVR, which is whether to store a unique copy for each user or have just one shared copy for all. In this debate, which is arousing the strongest passions in the United States, session based encryption brings the advantage to operators that they do not need to record content specifically for a given subscriber in order to fulfil legal requirements specified by rights holders. In that case copies are always shared. This issue will be explored in more detail in a future White Paper. A current architectural requirement that can sometimes be seen as a limitation of existing CAS and DRM systems is the continuing requirement to have a truly centralized database, because of the lack of viable distributed alternatives. Even if high availability and the strongest security could be provided by such an architecture, no distributed data models exist as yet. This intrinsic weakness for any premium content distribution system is compounded with edge processing, since not only is the integrity of the central database critical, but a valid network path must always be maintained to it.

Vendors such as Viaccess-Orca are working on networked database models that will map more naturally to the increasingly distributed nature of OTT. “Live and on-demand service deployments, on managed and unmanaged IP networks, already rely on distributed architectures for content preparation when the content comes from multiple sources or rights management platforms that are hosted outside of the operator’s premises, as in cloud deployments,” noted David Leporini, EVP Marketing, Products and Security, Viaccess-Orca. “Until now content preparation, including content encryption, has usually taken place at multiple locations but often with a centralized rights and DRM license management platform. Provided the right security architecture is in place, distribution infrastructures can now be extended to handle content encryption capabilities deeper within a CDN, extending right to the edge servers”.

Christian van Boven VP, Multiscreen Product Management & Architecture at Nagra discussed some of the pros and cons of distributing Key Servers: “Service providers want to decentralize the packaging of content, possibly as far as the CDN edge node. This raises the question whether the key servers should also be decentralized. This does not necessarily follow, since these two elements do not have to be coupled, and in fact in the few deployments out there so far that have implemented this architecture, key servers have not been significantly decentralized. We believes there is no fundamental reason why key servers should not be decentralized, although care has to be taken to ensure scalability and security. On the scalability front, session based encryption makes it harder to build more resilient OTT infrastructures that can scale at peak times, for example major live sports events, which can entail pre-delivered licenses and keys. But this is solvable with the help of improving redundancy, load balancing and associated technology, although this may raise the cost.

In the case of security, the issue is that centralized key servers are normally deployed in controlled environments less vulnerable to physical intrusion. However the advent of Hardware Security Modules can mitigate this problem, while more sophisticated models are emerging that make distributed architectures less vulnerable. OTT players can create partnerships with network infrastructure providers to help address both these security and overall OTT scalability issues.”

One of the most exciting prospects that edge packaging brings to the table is the possibility of having truly scalable live OTT streaming, so that premium pay TV including key sporting events will one day be available with enough reliability over the Internet to become truly mainstream.

8. Annexes

8.1 References

8.1.1 Anevia 2010 White Paper on Adaptive Bit Rate streaming

http://www.anevia-group.com/wp-content/uploads/2014/06/Anevia_White-Paper_OTT-Streaming_2nd_Edition-2.pdf

8.1.2 IPSOS study on TV usage

http://www.ipsos-na.com/news-polls/pressrelease.aspx?id=6433

8.1.3 Comcast Netflix streaming agreement

http://www.nytimes.com/2014/02/24/business/media/comcast-and-netflix-reach-a-streaming-agreement.html?_r=0

8.1.4 Cisco promotion of CDN federation

http://www.cisco.com/c/dam/en/us/products/collateral/video/content-delivery-engine-series/white_paper_ibsg.pdf

8.2 Acronyms used

ABR	Adaptive Bit Rate
CDN	Content Delivery Network
CMS	Content Management System
DASH	Dynamic Adaptive Streaming over HTTP
DdoS	Distributed Denial of Service (network attack)
DRM	Digital Rights Management
JIT	Just In Time (also on-the-fly)
NPVR	Network PVR
PVR	Personal Video Recorder
QoE	Quality of Experience
QoS	Quality of Service
SMS	Subscriber Management System
SPF	Service Platform

For more information, visit: www.anevia.com

Anevia Headquarters
1 rue René Anjolvy
94250 Gentilly
France

North America
800 W El camino Real - Suite 180
Mountain View
CA 94040
USA

Latin America
Av. Dr. Chucri Zaidan, 940
16° Piso
São Paulo, SP-04583-906
Brazil

Middle East
Dubai Silicon Oasis
Office B303
PO Box 341073, Dubai
Dubai

APAC
20 Malacca Street
#4-00, malacca centre
Singapore, 048979
Singapore