

# The Impact of Adaptive Rate Streaming:

A Primer for Digital TV Executives

Authored by Benjamin Schwarz



  
**verimatrix**<sup>®</sup>  
harmonic<sup>®</sup>

# The Impact of Adaptive Rate Streaming: A Primer for Digital TV Executives

Authored by Benjamin Schwarz

## What This Paper is About

Adaptive rate streaming is a breakthrough approach that enables a video's bitrate to adapt dynamically to resources available in the network and on playback devices. The goal of dynamic adaption is to provide a smooth presentation of media and a "TV like" consumer friendly experience even when bandwidth availability is unmanaged. Appearing commercially in 2007, this approach was first deployed on a mass scale during the 2008 summer Olympics. Contrary to previous unsuccessful attempts, it takes advantage of existing Internet infrastructure.

This is a key enabling technology for over-the-top (OTT) video that will profoundly impact content distribution even in walled garden environments. With improved media presentation comes greater consumer enthusiasm and enhanced prospects for commercialization. Monetized services can therefore be enabled for more devices, from high-end 3D TVs to phones over varied network conditions.

This paper describes what adaptive rate streaming is, as well as its potential as a disruptive technology. We will often use the term Apple has submitted for standardizing adaptive rate streaming, an architecture known as HTTP Live Streaming (HLS).

We start with a description that is as non-technical as possible, covering the key stakeholders that are Apple, Microsoft and Adobe and illustrate some deployments. We then dive into the technology to explain what technical tradeoffs are required when deploying this technology today. The third main section covers game-changing business issues from the perspective of different types of operators and vendors. Before wrapping up, we further explain what the main issues are and take a peek into the future of adaptive rate streaming technologies. A list of references will be found at the end of the paper.

## Adaptive Rate Streaming for My Aunt Who Still Doesn't Own a PC.

Imagine a movie as a puzzle broken into 1,000 pieces. You're at home and the puzzle is in the store. Each piece of the puzzle exists in high-resolution pristine versions, but there are also smaller and simpler versions that have lower resolution and are lighter and quicker to sort and move around. If you buy the puzzle online, each piece needs to get to your house for you to build the whole picture.



If you go ahead and buy, you first get a plan (the playlist), then pieces themselves start arriving immediately for you to start assembling and having fun.

Every piece can be downloaded at its own pace. So if you have a good connection, you can ask for, and get, bigger pieces arriving at the same rate you are putting the puzzle together. Those parts of the puzzle will be of dazzling quality. If the Internet is not working well for part of the time, maybe because your geeky nieces are doing things you don't want to know about on their PCs in their rooms, you will get lower resolution pieces during that period. But as long as the Internet connection is there at all, you won't be stalled waiting for pieces to arrive, and all parts of the puzzle will still look fine even if some are a little less glossy and high-resolution than others.

What is certain is that you'll get your whole puzzle with no interruption to your enjoyment, on the TV, the tablet or the phone, complete in as good a quality as is possible.

Before this technology was invented, you'd either have had gaps in an otherwise uniform puzzle, or had to wait for the whole puzzle to be delivered before you started to play.

*As with most new technology, there is not yet a consensus on its name. So we'll navigate between adaptive rate streaming, Dynamic Adaptive Streaming on HTTP (DASH) aka Multiple Bit Rate Streaming (MBR) or Microsoft's term Smooth Streaming aka in Apple's lingo http Live Streaming (HLS).*

## A Non-technical Description of Adaptive Rate Streaming

At the apex of the Internet bubble, Microsoft, Real Networks and others launched commercial products to attempt to adapt a video stream's bandwidth to available network conditions. These unsuccessful solutions required expensive dedicated infrastructure.

Adaptive rate streaming is a relatively new technology for delivering video to a variety of user devices. The video resolution varies depending on the resources available from the network and the device on which you play it back, but as long as there are at least minimal resources, the "show will go on." There is no need for dedicated network infrastructure, as video is transported through regular Internet data protocols.

In contrast to the technology known as IPTV, which uses dedicated networks, this is often called Internet TV, over the top or OTT TV.

Among its many advantages, adaptive rate streaming offers rapid channel-change or fast start-up, as a low bitrate can often be used for the first video frames displayed. Once the video starts decoding, the device can move up to the highest supportable bitrate and quickly improve the displayed resolution. Smoothly changing the resolution or compression rate of displayed video will typically go unnoticed by consumers, in contrast to their obvious frustration at slow startup or stalled playback during re-buffering.

### **Streaming vs. Broadcast & Multicast vs. Unicast**

Streaming occurs between a server and a client device both connected to the same network. At the end of the 1990s it promised to make the Web a more interesting place where static pages were enhanced with sound and video. A few years ago, the processing power available in devices (except PCs) was too low to support high resolution decode and the management of multiple streams. In traditional streaming the server keeps track of what a client is doing; in HLS and other adaptive models the client manages the relationship. But in any kind of streaming there is a connection between the server and the device.

For broadcast, one typically refers to a transmitter rather than a server. There is not necessarily any kind of physical or logical connection between transmitter and receiving devices, so broadcasting a signal takes the same resources whether there is one person listening or several million.

When sending a TV show in a stream to a single user, typically in an on-demand use case, the stream is called unicast. But when several viewers are watching the same stream at the same time, multicast streaming can be used to optimize network usage. Telcos think of this rather like broadcasting within their controlled network. If several ADSL users connected to the same telephone exchange are watching the same live broadcast, then that exchange only requests one stream from the head-end server and copies it to each of the viewers. Multicasting in today's telco IPTV deployments is complex; it requires special infrastructure and bandwidth management in the network which is not easy to duplicate in the multi-hop connections of the Internet.

## The Move to HTTP Unleashed the Beast

Early PC implementations of adaptive rate streaming were unsuccessful because they relied on specialized protocols (like RTSP). They could therefore not use equipment designed to optimize ordinary Web traffic that used the more standard HTTP protocol.

As the H part of the acronym suggests, HLS uses the Internet's basic HTTP protocol. It's counter-intuitive that a protocol dating from the very beginning of the Internet, originally designed to transport small text files, should end up being the enabling technology for mass deployment of video delivery.

One key reason for this is that sophisticated CDNs from the likes of Akamai can be used without much adaptation at all. Even a cheap stand-alone HTTP web cache helps deliver HLS.

Velocix provides service providers with infrastructure to build their own CDN. Paul Larbey, the company's CEO, told us that HLS simplifies deployment and typically reduces cost by 30%. Velocix infrastructure is currently used by Verizon to deliver "HBO Go" and by Talk Talk to prepare for YouView delivery. Velocix further found that the nature of adaptive streaming makes CDNs even more appropriate for streaming to mobile device with a better Quality of Experience.

By using the existing infrastructure, HLS creates a mechanism akin to a "multicast" architecture that almost comes for "free". Video content divided into small files is cached at various points in the network and saved for a short time. Another viewer can request the same content directly from the cache if it's still there. There's no need to travel any further up the network to the head end.

By avoiding a silo architecture, HTTP live streaming is also a good enabler of TV-Everywhere.

The HTTP protocol is about moving files around, not streams. The technical description below explains how video streams are converted into "chunks" that are nothing other than files.

A valuable benefit that comes with using the HTTP protocol is that it carries traffic seamlessly through all parts of the network including the home and through firewalls and other security mechanisms. It uses the standard “port number” (80) of all ordinary web traffic, where other protocols use specific ports (554 or 7070 for TCP and 6790 and higher for UDP) that are often blocked.

### **A New Client Server Architecture, with an Active Client**

Broadcast existed long before streaming, so it was only natural for streaming video to mimic its predecessor. The server originally had more resources to control the streaming than the client. In IPTV, video playing software on thin client architectures typically joins or leaves an existing stream that runs at a fixed constant bitrate. Until recently, streaming on the Internet has mainly used Real-Time Streaming Protocol (RTSP), which is a stateful protocol developed by Real Networks & Netscape in the 90s. That means that the server keeps track of the client’s state at all times and tries to help optimize delivery and deal with errors in real-time.

HLS is a stateless protocol and it is now the software running in the end device that completely takes over the setting up and optimizing of video, wholly from the point of view of optimizing the consumers’ experience, while the head end server becomes dedicated to creation and management of a fixed range of chunked video sources. Encoding and preparing the video chunks takes place with no reference to the streams that are actually being consumed, so no client state is required or maintained. The encoding process, like everything else at the video head end, has to be set up for the type of service being offered, and for the types of content that must be monetized, this includes appropriate mechanisms to encrypt video and manage the delivery of keys to authenticated users.

### **The key rate-adaptive technology suppliers**

This section mentions five companies notable for their role in introducing and promoting adaptive rate streaming solutions, focusing on the three main contenders for mass market adoption - Apple, Microsoft and Adobe.

Drew Major started Move Networks in 2006 and was also the co-founder of Novell. Move Networks was the first company to bring a working adaptive rate solution to market. Somehow, though, Move Networks has not been able to harness this first-

mover advantage and in mid 2010 went through a major restructuring, shedding a large proportion of its workforce. In September 2010, Move Networks was, however, granted a patent for some of the pioneering work it did on adaptive rate technology.

### **Apple**

Apple's public promotion of HTTP Live Streaming dates from March 2009, when the standard was introduced as a part of the iOS 3 and associated iPhone and iPod launch. It has since been enhanced and extended in iOS 4 and, of course is central to streaming video services support on the iPad and now the Apple TV.

The continuing rapid uptake of these Apple devices, and the introduction of QuickTime X on the Mac is certainly helping HLS gain traction as a key enabler of high quality video delivery to this ecosystem, but uncharacteristically Apple has also published the HLS specification as an Internet Engineering Task Force (IETF) draft and is pushing for broader adoption as a fully ratified standard. This openness has enabled a range of other device types to adopt the same standard infrastructure – notably a wide range of STBs and other popular mobile devices. Until recently, a notable exception to the range of device support was the PC browser and plug-in marketplace, even when using Safari. As one aspect of the growing HTML 5 support initiative, HLS is very likely to rapidly grow in market significance and breadth of application. The ecosystem supporting head end components and other aspects of the video delivery value chain is also growing rapidly.

Apple seems to have the most pragmatic technical approach using the well-known MPEG2-TS structure for video encapsulation where other solutions use the newer and less mature MPEG-4 structure. Each chunk is stored as a file with a '.ts' extension and then an index or playlist file is generated to link the different chunks. Being based on well proven technology Apple's solution will have to be updated to scale up to future requirements as can be seen with the 1.6mbps limit on some streams. By leveraging the highly standardized and highly developed HTTP server and web delivery marketplace, rather than focus on a proprietary server product, important practical issues such as service Failover Protection mechanism (potentially implementing a load balancer throughout the delivery system) helps the HLS perspective goes way beyond simple bandwidth optimization.

Another critical aspect of the HLS architecture is the definition of a standardized stream encryption mechanism, describing the encryption algorithm and formatting of encrypted chunks. The standard also defines how references to the key file required for stream decryption should be included in the playlist or manifest file. This should be seen as a DRM-independent approach, because even with this level of detail on protecting content, the standard is deliberately silent about how to authenticate devices or control which devices can obtain which stream keys. Security companies like Verimatrix therefore have the opportunity to provide significant value to the ecosystem by overlaying the HLS protocol with a mechanism for registering and tracking individual devices, together with a way to manage entitlement for content viewing. This type of security opens up the potential for flexible business models that can help up-sell OTT content for premium services and cross-sell over multi-layer, multi-device distribution.

HLS has been deployed successfully to Apple's iPhone, iPod, iPad and MAC (OS X) and is also being used for OTT set-top box (STB)/TV deployments.

### **Microsoft**

Smooth Streaming is Microsoft's IIS Media Services extension to enable HTTP adaptive rate streaming to Silverlight and other client devices. It commonly uses H.264 video encoding but can also use Microsoft's VC-1 codec.

From a workflow and architecture point of view, the Microsoft Smooth Streaming solution is similar to HLS although initial feedback is that the use of MP4 files makes it harder to implement in a broadcast environment due to the difficulty of preserving meta data. Microsoft has already deployed Smooth Streaming to PCs and has high hopes that this video capability will help make its new Windows Phone 7 a success. Note that Nokia has already announced its support for Smooth Streaming. STB and connected TV support from Smooth Streaming is also planned with PlayReady DRM.

Microsoft, together with Apple and Cisco are all taking an active part in the MPEG Dynamic Adaptive Streaming over HTTP or MPEG (DASH) standardization effort that aims to converge the different industry initiatives.

**Adobe**

Most recently, Adobe has implemented a variant of adaptive rate streaming within its Flash player 10. The feature is called HTTP Dynamic Streaming and requires Flash Media Server 3.5. The initial release used RTMP to deliver the video, but since Flash player 10.1, HTTP delivery is also supported. Current Flash deployments use RTMP, and it is expected that in 2011, Flash Dynamic Streaming will be available on PCs and Mobile phones. Android has already announced support of HTTP Dynamic Streaming. In his technical white paper (see references) Andrew Fechey-Lippens states that Adobe’s Flash player uses significantly more device resources than Apple’s solution. But multiple parameters like “chunk size,” which is described below, differ considerably between Apple, Microsoft and Adobe solutions making comparisons very hard to make.

For completeness, it is worth noting that Google is backing WebM, an open-source initiative based on the VP8 and Vorbis codecs. WebM is a relatively untried technology as yet. It uses a different approach for sending the video in a single file, encompassing variable bitrates rather than chunks. It is much harder to benefit from caching with this approach and head-end scalability might become an area of concern.

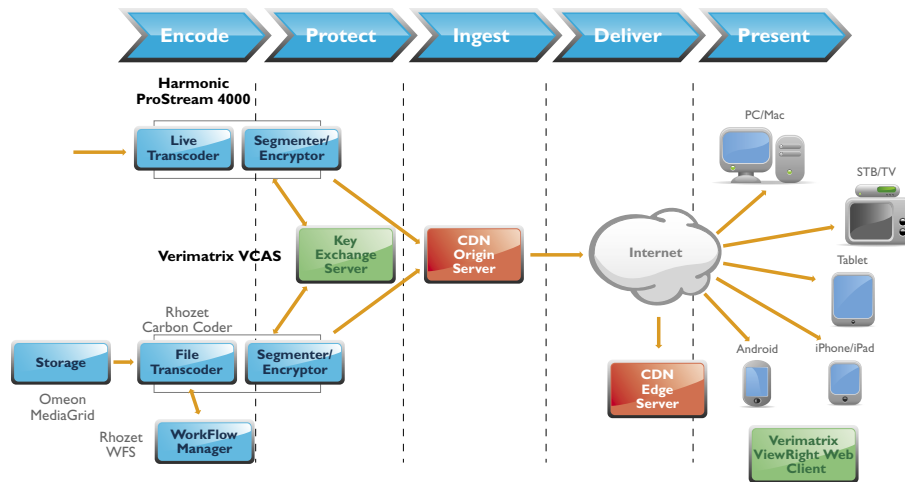
**Examples of Some Live Services**

Operator	Service	Type	Technology Used	Devices
NetFlix	Internet VoD	VoD	Smooth Streaming	PC, Roku STB, connected TV
BSkyB	Sky Player	VoD	Smooth Streaming	PC
AT&T	U-Verse Mobile	VoD	Smooth Streaming	
Apple HLS	Window Mobile Phone			
iPhone				
Canal+	Canal + Foot	VoD	Smooth Streaming	PC
Swisscom	TV Air	Live / VoD	Smooth Streaming	
Apple HLS	PC, iPhone, iPad			
DirecTV	NFL Ticket on the Go	Live	Apple HLS	iPhone
Major US TV network	Live events	Live	Apple HLS, Akamai	iPad
Global OTT content platform	TV Shows	VoD	Apple HLS	PC, iPad, iPhone, Android

## Implementing & Fine-Tuning HLS (Getting a Bit More Technical)

So this is the chapter to skip if you are disconcerted by the difference between a kilobyte and a kilobit.

Producing adaptive rate streaming requires a similar workflow, whoever the supplier: encode, segment (except for Google), encrypt, store, wait and deliver.



## Getting a head-end together

In a typical solution like the Verimatrix / Harmonic architecture illustrated here, both live linear content and on-demand file based content can be managed with the same delivery architecture. The video source is encoded in one pass for all different bitrates and then segmented into chunks. The playlist/manifest file is created to tie all stream information together for the client devices. While the streams are being segmented, a closely integrated key exchange process communicates with the security sub-system to provide a steady stream of encryption keys and reference identifiers. It's worth noting that keys and keyfile references do not need to point to the same network resource as the URLs for the media chunks.

Encoding and segmentation are somewhat linked processes and it is often a benefit to have these tightly integrated within a single hardware device. The segmenter must be able to begin and end each bitrate chunk for each time period to an accuracy of just a few frames across all bitrates, and must ensure that all chunks contain a close sequence

of compressed video information. Any failure to create these accurate splices will undermine the smooth presentation of video as a device moves between bitrates.

The number of bitrates generated depends on the business case, as does the frequency of key rotation and any differences between keys applied to specific bitrate versions of streams.

A greater number of bitrates enable a finer adaptation to precisely the available resources that are bandwidth and CPU. But all bitrates have to be generated for the length of the content and subsequently stored. In some large-scale deployments, each different bitrate version of a stream might need to be distributed in parallel over a core network to a set of edge servers. Too many bitrates profiles obviously have the potential to consume excessive resources at the head-end or in the network, even while they may improve the end user experience or the reach of the content to different devices. Improvements in live transcoding, segmentation/encryption and storage architectures will tend to alleviate this problem.

In most implementations, content is encrypted using the same key for each chunk that represents a specific time period, so protecting content doesn't explode the amount of data involved or prevent caching. For a real-time stream, the so called "cryptoperiod" is adjusted so that a given key is used for encryption of a number of chunk periods, which reduces the need for the client devices to frequently fetch new keys. Setting an appropriate cryptoperiod is the same as with any pay-TV environment – a compromise between key management overhead and appropriate protection of content from various types of attack. Typically a key will be used to protect a few dozen chunks before being renewed. For on-demand content, the strategy is often to use a small number of keys, as the length of the playlist/manifest file, read before any video is displayed, can grow very rapidly, if it contains a large number of key references.

Integration of security is also important at the head-end when chunking and encrypting streams. The key exchange protocols are not yet fully standardized for this type of streaming in the same way they are for DVB delivery systems. The security of the exchange process and the flexibility to configure different protection mechanisms and client signaling schemes becomes more important as the OTT market matures.

These mechanisms are not always fully defined in the different specifications and so it is important to work with a set of partners who have an appropriate level of experience.

### Packets & Chunks

Looking back at RTSP, one of the traditional streaming protocols, 1.5kb packets would represent about 10 milliseconds of video in a full SD format, so a 90-minute movie could require half a million packets.

HLS was developed much more recently and the chunk size is typically set to hold between two seconds and ten seconds of video. So a movie would require a few hundred to a few thousand chunks. Of course this is then multiplied by the number of bitrates being supported, with chunks of different bitrates having frame accurate synchronization at their start and end points. Each different bitrate version of a chunk corresponds to a separate encoding profile.

### Profiles

One aspect of a practical system to be designed carefully is the number of video profiles being simultaneously generated for the service. The more profiles, the greater range of devices that can be supported and the more those devices can maintain the smooth presentation of video in the face of changing network conditions. Apple recommends five profiles for their devices; the Move Networks solution is also usually set up with five profiles for Web delivery. Microsoft recommends just three. In some extreme cases there can be up to eight profiles. This may become more common in the future when new full HD 3D formats are supported. Here is an example of five typical Apple profiles using H.264 for the video and a fixed 16kbps HE-AAC audio stream.

Video Bitrate (kbps)	Resolution	Frame Rate
40	320x180	3
90	480x270	5
256	480x270	15
500	480x270	30
1300	1024x720	30

Profiles are defined by their bitrate. To achieve a given bitrate (i.e. compression level) the encoder will usually use video resolution, video format and frame-rate as variables. In current implementations, the audio bitrate doesn't usually vary. In sports in particular,

it is better to keep a continuous audio commentary even if the video becomes a slide show for a brief period (say on mobile networks).

iPad	iPod Touch	Google Android HTC
iPhone 3GS	Blackberry	Motorola Droid - WVGA resolution
iPhone4	PC – multiple sizes up to full res	OTT STB – multiple sizes up to full res (e.g. Netgem)

Profiles are mainly chosen so as to target specific devices. Harmonic has already implemented solutions targeting these nine different device groups:

### Tradeoffs

The chunk size will determine the speed at which a device can adapt the video bitrate. Smaller chunk sizes mean faster rate adaption and potentially faster channel changing. This implies that Apple’s out-of-the-box configuration is best suited for streaming to different devices over stable network conditions. A 2-second chunk size is better suited to playing variable network conditions. However, smaller chunks mean more chunks, which is why Microsoft (with 2 second chunks) suggests fewer profiles than Apple (with 10 second chunks) can. Smaller chunk size also requires longer playlist/manifest files, which then require more parsing and longer repeated fetches for live streams.

Here are some figures for a 90-minute movie:

Chunk Size	Number of Profiles	Number of Chunks
8s (HLS)	5	3375
2s (Msft)	3	8100
2s (Msft)	5	13500

Distributing many thousands of small files throughout a CDN can be challenging. As the protocols are state free, smaller chunk size requires more setup and tear down of TCP/IP connections for each individual chunk. This creates a higher load in the network using more traffic for setup and so relatively less traffic for delivery.

As they are made ready, chunks are encrypted and put on an HTTP server ready to be downloaded on the device’s request. Note that a 2-second duration corresponds

to a typical group of pictures (GOP) size so that video playback can start even if the client has only received a single chunk. The playback or download buffer’s parameters can be adapted to specific use cases.

The client adapts the bitrate by selecting the best possible profile given the available resources. Starting with the lowest profiles, the client may be tuned to move up to higher bitrate profiles one by one, or possibly skip some on the way up.

### Parameters for choosing a solution

Playing with all the parameters described above will lead to technical trade-offs. At this stage, it may still seem challenging to choose between the different implementation approaches. In most adaptive rate projects, the first choice must be which devices to target. The choice of delivery network will often be an external constraint to take into account in choosing the right solution. Building or renting the head-end - where companies like Harmonic or Verimatrix come into the picture – is then a key choice to make, bearing the whole ecosystem in mind. Here is a list of eight parameters that need to be taken into account:

Breadth of client implementations	Security & signaling support (e.g. output control options)
Subscriber’s user experience	Performance
Ease of future device integration	Cost (license fees, implementation cost, maintenance),
Multi-vendor support (from encoders to CDNs to clients)	Roadmap & standards track support

### New Business Models Will Emerge

#### Securing Revenue

Video consumption is becoming increasingly sophisticated. Beyond being delivered, a stream needs to be presented, searched, promoted, recommended and monetized. Enabling and securing a commercial transaction is too often regarded as an unnecessary overhead. In the new online world of OTT, as in any other, commercial viability is still key. The business environment for adaptive rate streaming is only just appearing and flexibility is paramount. Devices are used to manage rights. They are associated with individuals that can own multiple devices so rights must be valid for multiple devices.

### **The Link Between HLS, Cord-cutting and OTT**

The initial gut reaction of many stakeholders is to fret that the OTT enabled by adaptive rate streaming will shake the market, and they might lose out. But are first reactions always the right ones?

At this simplistic threat level, content operators are the most at risk. If customers cancel pay-TV and use OTT, operators could lose out all together. In many markets they are banking on a combination of exclusive rights (with a bias for live sports which are easier to control) and strong brand marketing. They are gambling that OTT services will require too much of a “lean-forward” approach. But just in case they’re wrong, they are also slowly letting some controlled OTT content into their walled gardens.

Satellite broadcasters offer parts of their line-up over the Internet. DISHOnline.com for example has extended its catalogue beyond on demand to include live TV. DISH has even committed to enabling such services for non-DISH customers.

Incumbents like AT&T or Orange also feel threatened by OTT, fearing it will push further disintermediation for them. They can at least console themselves by still providing network access, even for subscribers that get all their content OTT.

Some markets will see regulatory battles as incumbents try to use those regulatory bodies to defend their existing business.

### **The Glass is Actually Three Quarters Full**

Opportunities far outweigh threats for both network and content operators. Indeed HLS will extend the reach of services both within and beyond operator networks while content operators will achieve a global reach overnight.

### **Network Operators**

Telcos have spent much of their energies in the last decades moving away from a “bit-pipe” model, where they simply transport data as a commoditized service. Their drive into IPTV is a typical example of this movement.

Incumbents are becoming service aggregators, where most services are actually delivered by third parties. For example a telco's catch-up TV portal will provide missed episodes that are then delivered using HLS from a TV station's own portal. The value sought after here is in the aggregation. But that value created - at great cost - serves the sole purpose of selling more network access packages. HLS is a simple cost-effective way to add value to networks.

In the short term, full HD streaming can only be monetized on managed networks. Telco's are hoping to stay ahead of the game with Super HD, 3D and multiple streams so that, as with full HD today, their managed networks of tomorrow will still have a reason to exist alongside OTT services on top of a plain vanilla ISP service.

But before that, the first opportunity HLS brings for telcos is to deliver services to subscribers that were not previously eligible and widen the breadth devices served, be it in the home (as in the TV Everywhere initiatives in the US) or on the move (delivering to hotspots or 3G connected devices for example).

Network operators also use HLS to widen the breadth of services offered within their walled garden by giving selective access to OTT services at no cost to themselves.

After painfully failing to keep up with the likes of Google, MSN or Yahoo! in the portal space 10 years ago, telcos had mostly given up the dream of selling their services to other ISPs' customers. Adaptive rate technologies, by lowering the barriers to entry, could make some of them reconsider.

The technology is clearly a blessing for operators that challenge incumbents and have no access network. They see it as a key advantage to start delivering services rapidly with very little upfront cost. This is the case for many European incumbents as they move into foreign markets for example.

### Content Operators

The big U.S. networks were the first to grasp the importance of adaptive rate streaming technology. Initial deployments with Move Networks in 2007 saw companies like Fox in the U.S. and later Germany's Prosieben deliver content over the Internet directly to the customer from the content owner's own Web portal.

But just as the incumbent operators see this technology as an opportunity, newcomers threaten to also create direct relationships with end customers without needing to go through existing content operators. The barriers for entry into the video content delivery business are being drastically lowered for everyone.

### CE Manufacturers

A TV set maker with a clear strategy for OTT is Samsung. The widgets they are introducing into their connected TV platform will enable a direct relationship between content producers and customers.

Adaptive rate streaming will have an even greater impact on mobile devices within the home. Apple's iPad is set to really benefit from the technology and Samsung's Galaxy, like the other tablet makers, hopes to benefit also. Moving around the home, network conditions will vary, but thanks to adaptive rate streaming, video streaming can continue regardless.

Until recently, suppliers like Netgem were limited to offering VoD content in the paid services of their OTT devices. Progressive download was the technology of choice for this. HTTP Live Streaming is now opening a world of new business possibilities, with live content over non-managed networks becoming a real option for the first time. Yann Courqueux, Netgem's CMO, named two key advantages with this OTT enabling technology.

"Operators have greater flexibility in launching TV services. They no longer need to provision bandwidth or limit the service to a certain copper line length. The service comes as multi-screen out-of-the-box with this technology." The second advantage is for new entrants with innovative business models. "Ethnic content aggregators are a new segment for us."

Another example is consumer electronics portals like "Toshiba Places," the portal that WizTV is building that will serve TVs through Netgem's OTT boxes, PCs and android Tablets. Toshiba is putting a strong social media focus on the project.

Eric Bibolet is founder and EVP of WizTV, a French company with explosive growth in connected TV implementation. He told us, "Our customers follow the same path from a first requirement for VoD with DRM –which is what Canal+ is doing on TCL TVs – to DRM-

protected live streaming, which Eurosport is providing for €4,90 per month over the Internet. The requirement for dynamic adaptive streaming usually comes third.”

WizTV is working on several major projects that will go public in Q2 2011. As content owners are all keenly aware of the iPad’s potential for new revenue, WizTV pushes the Apple HLS technology on connected TVs in their implementation projects. “One major difference we’ve found over the last year is that all CE Manufacturers are now acutely aware of DRM and have implemented it in at least one flavour.”

### **Other Aggregators and Intermediaries**

Heavyweights Apple, Microsoft and Google all play multiple roles in this game and were covered above, as potential technology suppliers. They are investing heavily in adaptive rate streaming technologies because of the huge potential it represents for them to enhance the direct relationship with their end-customers.

Netflix supplies on-demand content and no live streams. Progressive download has satisfied their requirements up to now. We expect stakeholders like Netflix or LoveFilm to turn to adaptive rate technologies that offer advantages for VoD also. Customers can choose whether to prioritize continuity or video quality.

Retailers ranging from Walmart to Amazon are also moving to OTT digital content delivery for which HLS is a key enabler by lowering the barriers to entry of this segment also. Some of the arguments against adaptive rate streaming

### **Issues that Remain to be Resolved with Adaptive Streaming**

HLS isn’t a silver bullet and here are some of the main arguments one hears against the technology.

- The Quality of Experience may not be guaranteed. So how can you monetize such a solution?
  - This is a valid question, but as any IPTV operator knows all too well, there will always be weak links in the quality chain, namely the last mile for DSL based solutions and the home network for all the others. When congestion happens there, HLS is the only solution that offers at least some service. Most customers would agree that a degraded service is better than no service at all.

- With multiple bitrates, the overall bandwidth used is greater than for regular streaming.
  - Again this is a perfectly valid criticism of HLS, but it misses the point. HLS will not replace all existing streaming. When live TV is streamed over a managed multicast-enabled network, the only benefit HLS would bring would be to offer multiple profiles in the home, targeting different devices. It is likely that one day, transforming incoming broadcast streams to rate adaptive will happen within the home to create such use cases.
- The end-to-end latency of HLS is too large for live sports events.
  - In the digital TV industry, latency has been a recurring problem. First generation encoders for new codecs have always introduced more latency than the previous generation of optimized algorithms. Current worst-case scenarios for HLS are of a 30-second delay, but it can be configured to as little as 2 seconds. We expect this problem to slowly go away as the technology improves.

### What Might Lie Ahead for Adaptive Rate Streaming?

We expect the use of adaptive rate technologies to extend further from the public Internet to managed network environments.

HLS was originally designed to enable video delivery as resources become sparse. As fiber rolls out at an ever-increasing pace throughout the world, HLS can also be adapted to scale up, to full HD, 3D and beyond.

The end-user is currently unaware if adaptive rate technology is used. Use cases may appear where customers could explicitly participate in tradeoffs by perhaps accepting a lesser quality video in return for a price discount.

We also expect many new use-cases of streaming services at public hotspots and over 3G networks to appear.

Scalable Vector Coding (SVC) is an alternative approach, which has had very little market traction so far. Also new codecs like VP8 or H265 might gain some traction. In all cases, a well-built adaptive rate solution should be able to integrate new components so as to remain competitive.

In the long term, adaptive rate technologies will have an impact on the Net neutrality debate by sharing available resources more equally among users. When operators next go begging regulators for traffic prioritization rights, they may have a new point to make. In case of congestion of non-prioritized traffic, all bitrates will theoretically drop until equilibrium is reached. Resources will be shared more fairly instead of randomly, as is the case on the Internet today.

The MPEG (DASH) initiative to harmonize standards mentioned above should prevent a standards war from slowing down the take-up of this promising new technology. Apple has already initiated a normalization process as has the Open IPTV Forum (OIPF) with its standard for HTTP Adaptive Streaming.

As the current version of chunking technology is taking root in the market, researchers are already looking into improvements, which could include byte ranges instead of chunks, virtual chunks and chained chunks. So assuming standardization doesn't hinder the technology's progress, it will itself prevent any stakeholders from resting on their laurels.

### Wrap-up

You've read this far, your decision is probably no longer whether to go to an adaptive rate streaming solution, but which solution to choose. As stated previously, the first thing to get sorted is the business case – what services are to be delivered to which devices across which network? As we've illustrated here, tuning the different parameters can be quite a challenge, so make sure you get the right partners with you on the project. If you do, video will look as good as is possible, which is already quite an accomplishment, but one that only serves a purpose if someone can make some money from it. The OTT environment is still changing rapidly. Solutions put into effect to dynamically adapt video must themselves be flexible enough to adapt to the market just as rapidly. With that in hand it really is time to start adapting!

### Some References

Adaptive rate streaming is rapidly moving into the limelight. Below are some pointers to what we found useful.

1. Apple documentation:

<http://developer.apple.com/iphone/library/documentation/NetworkingInternet/Conceptual/StreamingMediaGuide>

2. Apple's (Roger Pantos) HTTP Live Streaming draft IETF proposal:

<http://tools.ietf.org/pdf/draft-pantos-http-live-streaming-02.pdf>

3. Microsoft documentation:

<http://www.iis.net/download/SmoothStreaming>

4. Adobe documentation:

[http://www.adobe.com/products/httpdynamicstreaming/pdfs/httpdynamicstreaming\\_datasheet.pdf](http://www.adobe.com/products/httpdynamicstreaming/pdfs/httpdynamicstreaming_datasheet.pdf)

5. Recent Anevia OTT streaming White Paper with good description of Google's

WebM: [http://www.anevia.com/IMG/pdf/Anevia\\_White-Paper\\_OTT-Streaming\\_19-10-10.pdf](http://www.anevia.com/IMG/pdf/Anevia_White-Paper_OTT-Streaming_19-10-10.pdf)

6. Excellent background article from Streaming Media: <http://www.streamingmediaglobal.com/Articles/ReadArticle.aspx?ArticleID=65544&PageNum=5>

7. A technical Review of HTTP Live Streaming by Andrew Fechey-Lippens:

[http://andrewsblog.org/a\\_review\\_of\\_http\\_live\\_streaming.pdf](http://andrewsblog.org/a_review_of_http_live_streaming.pdf)

8. Akamai HD network demos:

<http://wwwns.akamai.com/hdnetwork/demo/index.html> & for the iPhone:  
<http://iphone.akamai.com/>

9. Qualive comparison of adaptive streaming:

<http://www.qualive.com/adaptive-streaming>

10. Open Internet Forum (OIPF) Release 2 Specification HTTP Adaptive Streaming

[V2.0] – [2010-09-07]: [http://www.openiptvforum.org/docs/Release2/OIPF-TI-R2-Specification-Volume-2a-HTTP-Adaptive-Streaming-V2\\_0-2010-09-07.pdf](http://www.openiptvforum.org/docs/Release2/OIPF-TI-R2-Specification-Volume-2a-HTTP-Adaptive-Streaming-V2_0-2010-09-07.pdf)

### **About Verimatrix**

*Verimatrix specializes in securing and enhancing revenue for multi-screen digital TV services around the globe. The award-winning and independently audited Verimatrix Video Content Authority System (VCAS™) and ViewRight® solutions offer an innovative approach for cable, satellite, terrestrial and IPTV operators to cost-effectively extend their networks and enable new business models. As the recognized leader in software-based security solutions for premier service providers, Verimatrix has pioneered the 3-Dimensional Security approach that offers flexible layers of protection techniques to address evolving business needs and revenue threats. Maintaining close relationships with major studios, broadcasters, industry organizations, and its unmatched partner ecosystem enables Verimatrix to provide a unique perspective on digital TV business issues beyond content security as operators seek to deliver compelling new services. For more information, please visit [www.verimatrix.com](http://www.verimatrix.com) or our [Pay TV Views blog](#) to join the conversation.*

*To discuss opportunities for your digital TV network, please contact us at [info@verimatrix.com](mailto:info@verimatrix.com) or +1 858-677-7800*

### **About Harmonic**

*Harmonic Inc. offers a comprehensive, innovative and market-leading portfolio of video infrastructure solutions, spanning content production to multi-screen video delivery. Harmonic customers can efficiently create, prepare and deliver differentiated video services over broadcast, cable, Internet, mobile, satellite and telecom networks, while simplifying end-to-end asset management, reducing costs and streamlining workflows.*

*Harmonic (NASDAQ: HLIT) is headquartered in San Jose, California with R&D, sales and system integration centers worldwide. The company's customers—including each of the top 20 Fortune 2000 media companies—choose Harmonic to enable their high quality video services delivered to consumers in virtually every country. Visit [www.harmonicinc.com](http://www.harmonicinc.com) for more information.*

### About the Author

*Benjamin Schwarz has more than 20 years of international experience in consulting and Telco & Media organisations.*

*Benjamin started his career with Logica-CMG where he spent 10 years. In 2000 Benjamin was CTO of Net4Music, a digital sheet music Internet start-up. He then spent 8 years with Orange, starting as an R&D manager focusing on video & search technologies. In 2004 he joined the Content Division running international IPTV deployments.*

*In September 2008 Benjamin created “CTO innovation Consulting” bringing his unique expertise to Content owners, Technology companies and network Operators and is now involved in many projects described on his Website where he also blogs [www.ctoic.net](http://www.ctoic.net).*