MPEG-4 VIDEO-ON-DEMAND FOR CABLE SYSTEMS: AN OVERVIEW

By R. Jordan Greenhall
DivXNetworks, Inc.

Abstract

It is the purpose of this document to discuss the features and advantages of MPEG-4 video-on-demand solutions for cable systems. In so doing, the document will provide an introduction to the MPEG-4 ISO standard, review different video-on-demand products and different MPEG-4 video-on-demand solutions for those products.

INTRODUCTION TO MPEG-4

MPEG-4 is a new international ISO multimedia standard designed to be a complete and comprehensive standard for all multimedia. Unlike the MPEG-1 and MPEG-2 standards that were targeted at relatively narrow applications, MPEG-4 has been specifically designed to support a very broad array of applications across a large number of media and multimedia requirements. As a consequence, the standard is large and has a wide array of applicable tools that distinguish it from MPEG-2. However, the principal foci of MPEG-4 are:

- Superior coding efficiency. MPEG-4 is designed to provide video and audio quality indistinguishable from MPEG-2 at 1/3rd to 1/8th the bitrate.
- Interactivity. MPEG-4 natively supports object-based video (e.g., video "hotspots") and data back-channels, allowing it to provide the foundation for comprehensive interactivity.
- Multiple-platform. MPEG-4 is designed to deliver multimedia content across virtually all delivery media and to virtually any form of device. MPEG-4 can be delivered over wireless or wireline, or on physical media; it can ride on top of MPEG-2 Transport Streams, ATM, IP, and other transport protocols. Moreover, MPEG-4 is designed for playback on devices ranging from wireless handsets to digital cinema projectors.
- Massively scalable bitrate. MPEG-4 is designed to provide high levels of audiovisual quality at bitrates scaling from the micro-scale (sub 20Kbps) all the way to lossless Digital Cinema. This scalability is reflected both in the video and audio codecs, as well as in the file format and systems layer that has specific tools to allow for efficient dynamic scaling across bitrates.

VIDEO-ON-DEMAND PRODUCTS

To understand the impact of MPEG-4 on Video On Demand economics, it is important recognize that the current business environment is deploying three different kinds of "Video On Demand" – each of which has different properties and requirements.

Pay-per-view (Near VOD)

Pay-per-view, or Near Video On Demand is not, properly speaking, a Video On Demand product at all. Rather, pay-per-view combines the traditional broadcast video model where a single channel is addressed universally to all homes on a network with selective access technologies that limit the viewability of content from the client end. As a consequence, NVOD is able to restrict viewing to the granular "pay-per-movie" level and, by pushing out content on a staggered basis over multiple channels, generate an experience that is similar to "true" Video On Demand. However, while NVOD has proven attractive to subscribers, its relatively narrow content profile and, more importantly, its restrictions on subscriber control over the viewing experience (starting and stopping a movie, for example), have limited its subscriber appeal.

Pay Video-on-Demand (PVOD)

Thanks in large-part to MPEG-2 digital video technologies, service providers have recently been able to replace Pay-per-view systems with true Video-On-Demand solutions. With Pay Video-On-Demand, a subscriber is able to access a content library, select which content he wants to watch, pay for that content on a "per-view" basis and then watch it entirely under his own terms (i.e., begin watching when he wants, stop, rewind and fast forward, etc.). Because of these advantages, PVOD has proven substantially more popular than pay-per-view. Early trials of Pay VOD have shown buy-rates of 3 – 4X that of PPV, driving $16 - $22 monthly revenue per subscriber.

The principal weakness of Pay VOD is cost. Current PVOD video servers and storage facilities are expensive to deploy. Moreover, because each discrete PVOD session requires an entire dedicated MPEG-2 digital channel, the opportunity costs of PVOD are relatively high. (This is particularly exacerbated by the fact that most PVOD demand is currently focused in a relatively narrow time window during the week – requiring a large allocation of system bandwidth to meet peak demand which ends-up underutilized during the rest of the week.)

As a consequence, the utility of PVOD has been restricted to content that can command a relatively high price per view – premium events and blockbuster movies.

Subscription Video-on-Demand (SVOD)

Subscription Video-on-Demand (SVOD) is a relatively new VOD model that has arisen in response to the weaknesses of PVOD. With SVOD, a subscriber pays a fixed monthly price for access to a certain content library and is then able to view content from that library as much (or as little) as he wants during the month. The predicted utility of SVOD is fourfold. First, that bundling content libraries under a single monthly fee will provide the kind of value that will attract subscribers to different kinds of VOD content – archival movies, television content, educational content, etc. Second, that SVOD subscription fees will enable content providers to more effectively rationalize their production and revenue risks, thereby increasing ROI; Third, that SVOD usage will be less focused on nights and weekends than PVOD, thereby utilizing allocated VOD infrastructure more efficiently. Finally, that content usage and subscription fees can be appropriately calibrated to earn the system Operator a positive ROI on this alternative.

The principal weaknesses of SVOD are threefold. First, that SVOD content libraries can be very large – orders of magnitude larger than PVOD libraries, thus presenting substantial storage challenges. Second, that each SVOD session requires a dedicated VOD server channel. Thus while SVOD will likely fill VOD capacity troughs during non-peak hours, SVOD will also add to peak VOD consumption, thereby requiring expensive additional capacity. Similarly, the third weakness of SVOD is that under current MPEG-2 technology, each SVOD session requires an entire dedicated digital channel. As a consequence, a single subscriber watching an SVOD episode of Seinfeld would require as much bandwidth as 100,000 subscribers watching NBC.

DIFFERENT METHODS OF DELIVERING VIDEO-ON-DEMAND

MPEG-4 technologies can be used to deliver Video-on-Demand more efficiently and flexibly than current MPEG-2 based VOD systems. Cost savings using MPEG-4 can be found in storage, server infrastructure, bandwidth utilization and even in customer-premise equipment. Because of its flexibility, MPEG-4 can be delivered both via a cable system's digital video infrastructure as well its IP data infrastructure. In both of these modes, MPEG-4 is significantly more efficient than MPEG-2 for VOD. However, as discussed below, delivery of MPEG-4 VOD over cable IP data infrastructure is particularly exciting and presents the most compelling argument for MPEG-4.

The principal weakness of MPEG-4 is simply that it is a new technology and must compete with legacy systems and legacy investments in the older MPEG-2. Although some of this legacy investment is found in the headend (storage and servers), these technologies can be largely repurposed to MPEG-4 use and in any event represent the smallest portion of legacy investment. Rather, the largest hurdle faced by MPEG-4 is in the customer premise: digital set-top boxes. None of the currently deployed digital set-top boxes support MPEG-4. As a consequence, any decision to deploy MPEG-4 within a cable infrastructure must present a compelling argument to replace digital set-tops, or to supplement or upgrade them with collateral equipment (e.g., more functional cable modems, gateway devices, set-top upgrades, etc.).

MPEG-4 Digital Video

As discussed above, MPEG-4 is designed for delivery over many different protocols – including the MPEG-2 Transport stream. As a consequence, MPEG-4 technologies can be used within the MPEG-2-based digital video infrastructure to leverage as much of the current infrastructure as possible while taking advantage of some of the benefits of MPEG-4. The key advantages of such an approach are found in reduced storage requirements and reduced network bandwidth burdens allowed by MPEG-4's superior compression.

By compressing VOD content by a factor of 3 - 5X without any loss of quality, MPEG-4 gives cable operators significant flexibility in their VOD strategy. A cable operator can choose to simply use less storage for their current VOD content; provide a larger library of content over VOD; distribute their content more widely over their network, thereby reducing overhead on their costly ATM backbone; etc. Similarly, because MPEG-4's superior compression means 3 – 5X as many simultaneous streams on a given portion of allocated VOD bandwidth, a cable operators opportunity cost to scale VOD services is considerably reduced.

However, it is very important to note that in order to take advantage of these efficiencies, a cable operator must have in-place customer premise equipment that is capable of decoding MPEG-4 video. Although current MPEG-2 set-tops will be able to receive and decode the MPEG-2 Transport Stream carrier, the "supercompressed" MPEG-4 video signals within that Transport Stream will be inaccessible to the set-top without a dedicated MPEG-4 decoder.

The advantages of MPEG-4, including its native support for interactivity and ability to carry compression efficiencies to high definition and beyond, are compelling for any MSO that is serious about VOD (and interactivity) in its value-added strategy. However, in the event of such an upgrade to CPE, a service provider would be well-disposed to pursue more multi-functional "gateway" devices which, in addition to providing a platform for multiple bundled services in addition to cable and VOD, provide a pathway for the delivery of MPEG-4 video over an a cable system's IP data path.

MPEG-4 Video Over Data Networks

MPEG-4's ability to supercompress video, combined with its ability to delivery video over multiple transport protocols (specifically IP), allows an MPEG-4 enabled MSO to exploit their IP data infrastructure as truly effective video delivery channel.

The advantages of doing so are multiple:

1. Commoditized Backend Infrastructure – delivering VOD over IP Data means that the cable MSO can utilize its backend data infrastructure to distribute content to headends. This backend infrastructure can be built on-top of commoditized standard IT hardware such as Gigabit Ethernet, rather than dedicated and much more expensive ATM.
2. Commoditized Server and Storage – similarly, using MPEG-4 over IP enables VOD systems that utilize standard commodity IT hardware for both storage and video servers. This means storage and per stream costs that are a fraction of current VOD systems (over and above the savings driven by smaller files and reduced bandwidth).
3. Repurposed Hardware – because MPEG-4 over IP is treated by an MSO's infrastructure as "just more packets" a great deal of the infrastructure required for distributing and delivering MPEG-4 video content is the same as is used for all other IP data: email, files, web browsing, etc. As a consequence, capital investment in that infrastructure is paid for by multiple services, not just VOD.
4. Repurposed Bandwidth – similarly, and perhaps more importantly, MPEG-4 over IP uses the same bandwidth within the cable plant as all other IP data. Thus, the cable operator using MPEG-4 over IP as their VOD delivery medium is able to **better utilize** their existing IP data bandwidth, rather than attempt to figure out how to deal with underutilized dedicated MPEG-2 VOD bandwidth.
5. Flexible Delivery Approaches – as discussed above, there are many different kinds of VOD products, each of which presents its own tribulations for efficiency and ROI-conscious network operators. Delivery of MPEG-4 over IP can be done in three different ways, each of which has strengths that give network operators tremendous flexibility in efficiently and effectively delivering content to their subscribers.

MPEG-4 IP Data Streaming

Streaming is a technique used in data networks to provide an end-user experience that is identical to MPEG-2 based VOD. Using streaming, a video is delivered on-demand, in real-time and has complete "VCR"-style control (stop, rewind, fast-forward, etc.). Within a private network with adequate bandwidth, IP-streaming can provide quality of service indistinguishable from MPEG-2 VOD.

The principal drawbacks of streaming are that it is more server-intensive than other techniques (described below), thereby requiring more server infrastructure per simultaneous stream, and that streaming creates a virtual "channel" of dedicated bandwidth that lasts throughout the playing of the video content. As a consequence, streaming cannot take advantage of IP data's ability to "burst" delivery and provide more effective bandwidth shaping (described below).

The biggest advantage of streaming is found in multi-user events (particularly sporting events) where many of the features of VOD (such as when the content will be viewed, and fast-forward) are not applicable. In this case, network operators can take advantage of a variety of "IP Multicast" approaches to deliver an effective VOD experience to multiple simultaneous subscribers while dramatically reducing both server and network bandwidth overhead.

MPEG-4 IP Data Downloading

For many kinds of VOD content, the most efficient mode of delivery is "downloading" rather than "streaming." This is because downloading separates the act of viewing from the act of delivery and enables the network to deliver VOD content in the most efficient way for the network. Thus, where a piece of content is delivered at an average of 800 Kbps and the network is capable of delivering 3 Mbps to the subscriber, the network can "choose" to burst a 120 minute movie to the subscriber in just over 30 minutes. As a consequence, rather than tying down a server session and 800Kps of network bandwidth for the entire viewing, network resources are rapidly freed-up for other uses within the IP Data pipe (other VOD sessions, e-mail, web-browsing, etc.). Moreover, because the content resides at the client after the download is complete, reviews of the content take place entirely on the client and impose no additional burden on the network infrastructure.

The principal drawbacks of Data Downloading are:

- Storage Requirement – for content to be pushed to the client, the client must have some significant storage capacity (in the range of 350 MB per hour of content). In a gateway device equipped for PVR (and datacasting as described below), this is a non-issue, but this can be a serious hurdle for very inexpensive set-top devices.
- Limitations to Trick-Play – data downloading is fully capable of all of the features of trick-play with the single exception that fast-forward cannot go beyond where the content has been downloaded. Thus, if a subscriber begins viewing a movie and wants to fast-forward to the end, he will have to wait some amount of time (30 minutes in the above example) before he can do this. Because of the way that consumers typically use trick-play features, this limitation is usually unimportant, but should be considered when choosing which delivery approach to take.

MPEG-4 Datacasting

One of the more unique and compelling applications enabled by MPEG-4 over IP is "datacasting". With datacasting, the network operator "pre-loads" certain content on the subscriber's CPE storage by downloading that content during network down times before the content is available for viewing. A typical application, for example, would be to datacast a VOD version of a very popular blockbuster movie before the first day of the VOD window. When the content is available for viewing, all requests actually come from the pre-loaded content on the subscriber's client – meaning that the operator's VOD infrastructure takes no load whatsoever from the request.

This technique can be very efficient when used for content that is either likely to be extremely popular and the operator wants to avoid peak strain on his system, or for content that is uniquely targeted to the specific subscriber and the operator wants to deliver the content during network downtime (and avoid competing for resources during peak times).

The principle weaknesses of datacasting are lack of storage on the subscriber's CPE. With the CPE almost certainly serving double or triple duty as PVR, ad server and datacasting server, there is only so-much content that can be pre-loaded by datacasting.

However, this drawback points to the larger issue that each of these three methods of delivery over the IP Data infrastructure are complementary, not mutually exclusive. The best approach will require a mixture of methods appropriate to the content offering, the network and the subscriber. The key advantage of delivering MPEG-4 video over IP is that it has the flexibility to enable the operator to select the mixture of methods that is most appropriate and efficient for its profile.

(Note, it should be mentioned that the advantages of MPEG-4 compression are doubly applicable where the network operator will be providing a gateway device with PVR functionality – an MPEG-4 enabled PVR can store 3-5 times as much content as a typical MPEG-2 PVR. This means that a much smaller hard-drive can be used for the same amount of hours stored, at great savings to the MSO.)

ECONOMICS

Clearly, the actual economic footprint of a Video-on-Demand system is highly dependent upon specific network topologies and a mix of technologies used to deliver VOD. The below cost comparison between an MPEG-2 based VOD system and an MPEG-4 based VOD system provides a baseline that can be used to derive more detailed estimates based on real conditions in the cable network.

This model assumes a comparison between two cable networks with 100,000 subscribers (10,000 simultaneous streams at peak capacity). Both networks present a simplified topology of 100 nodes, each with a headend serving 1000 subscribers and a central supernode that serves the entire subscriber base through those nodes. One network is assumed to be using an MPEG-2 VOD system networked with an ATM backbone between nodes and the supernode, while the other network assumes an MPEG-4 based VOD system using Gigabit Ethernet to connect nodes with the supernode.

This basic topology assumes that each headend node stores roughly 20% of the total content library and absorbs 80% of the total VOD requests directly, while the supernode stores 100% of the content library, but handles only 20% of the total VOD requests (i.e., 2 out of 10 VOD requests aren't satisfied by the local headend and have to be delivered from the larger archive at the supernode). In the MPEG-2 system, this delivery from the supernode is via an ATM backbone, while in the MPEG-4 system, it is via a Gigabit Ethernet backbone.

MPEG-2 Based System

Total VOD Network Cost	$	5,917,139
Cost Per Simultaneous Stream	$	592
Cost Per Sub	$	59
Total Cost Less ATM Backend	$	4,717,139
Cost Per SS Less Network	$	472
Cost Per Sub	$	47

MPEG-4 Based System

Total VOD Network Cost	$	1,586,840
Cost Per Simultaneous Stream	$	159
Cost Per Sub	$	16
Total Cost Less GigE Backend	$	1,386,840
Cost Per SS Less Network	$	139
Cost Per Sub	$	14

Ratios

Total MPEG-4 vs. MPEG-2	27%
Total Less Networking Backend	29%

This model reflects cost elements associated with content storage at both the supernode and each headend node; VOD servers at both locations; and the ATM or Ethernet networking to support the VOD system. The model assumes a least-efficient MPEG-4 system that uses streaming technology and does not exploit download "bursting". The model also does not contemplate the cost savings and opportunity cost advantages associated with repurposing of IP data bandwidth contrasted with dedicated MPEG-2 VOD bandwidth. Finally, the model does not contemplate the tremendous efficiencies associated with datacasting to push the most popular or unique content to the end consumer device.

Thus, under the assumptions of this model, an MPEG-4 based IP video system is at least four times as cost-effective as an equivalent MEPG-2 based system – with tremendous flexibility to achieve even more substantial cost savings as more tools from the MPEG-4 toolbox are utilized.

Finally, it should be mentioned that this whitepaper addresses only the advantages of MPEG-4 for delivering VOD in a manner that is roughly identical to current MPEG-2 systems. This, of course, is only the tip of the iceberg for MPEG-4. The new MPEG-4 standard is a complete platform technology for a variety of next-generation applications including robust interactive and dynamic content and innovative distributed content delivery models. Consequently, upgrading to an MPEG-4 system makes sense both on a short-term ROI basis and on a longer-term strategic basis.

Contact Information:
Jordan Greenhall
DivXNetworks
10350 Science Center Drive.
Building 14, Suite 140
San Diego, CA 92121
Phone: 858.909.5300
Fax: 656.909.5301
Email: jgreenhall@divxnetworks.com
Website: http://www.divxnetworks.com