IP Video Status Report

By Wes Simpson, Telecom Product Consulting

Internet protocol (IP) applications for video services are growing rapidly, as networks continue to expand in scope and size around the world. More consumers are being connected with high-speed circuits, and the available connection speeds continue to increase. Fiber-to-the-home initiatives such as Verizon's FIOS project and the recent certification of the first round of high-speed Docsis 3.0 equipment by CableLabs both point to the increasing amount of bandwidth available online for consumers. This growth is not limited to the U.S., as countries like Japan routinely offer 100 Mbit/sec service to consumers and France boasts a huge penetration of DSL services. With all of this increased bandwidth, it becomes possible to deliver more video at higher quality to more viewers using IP networks.

This article focuses on four areas in which IP video has made significant progress over the past year. They include Internet protocol television (IPTV), Internet video, videoconferencing/surveillance, and professional broadcast/contribution video. Before going into the details, a quick review of these easily confused terms will help make the rest of this article easier to understand.

Types of IP Video

First, it is important to recognize that many different types of video run over IP networks, but not all of them are IPTV. Although there are some differences in technologies between the different types of IP video, there are also major differences in the user experiences and the business models.

IPTV is characterized by the following characteristics:

- Continuous streams of professionally produced content (such as a TV network feed)
- Hundreds of 24 x 7 channels
- Uniform content format (all channels typically share one compression method and use roughly the same bit rate)
- Delivered over a private network, such as a telco digital subscriber line (DSL)
- Viewed on consumer televisions by way of a set-top box
- Paid service via monthly subscription

Internet Video differs from IPTV in a number of aspects, including:

- Discrete content elements, ranging from clips lasting a handful of seconds to full-length movies
- Millions of content offerings
- Widely varying content formats, including dozens of different types of video compression, rights management technologies, and image resolutions
- Delivered over the public Internet
- Viewed on PCs with special software or on portable video players
- Free or advertising supported service

Videoconferencing and video surveillance are increasingly done through the use of IP network technologies. These two services share several attributes. They operate live, in realtime; are highly compressed, use a variety of compression formats; have one-to-one or many-to-one system architecture; have private, not public audi-

WES SIMPSON is president and founder of Telecom Product Consulting, an independent consulting firm that focuses on helping companies develop and market video and telecommunications products. He is a frequent speaker and analyst for the video transport marketplace, a columnist for *TV Technology* magazine and author of the popular books *Video over IP*, published by Focal Press in 2005 and *IPTV and Internet Video*, published by Focal Press in 2007. The second edition of *Video over IP* was recently published. Simpson has over 25 years experience in the design, development, and marketing of products for telecommunication applications. He was a founding member of the Video Services Forum, and was a member of its Board of Directors from 1997 to 2001. He can be reached at wes.simpson@gmail.com.

ences; and have widespread adoption of new technologies, such as MPEG-4 and H.264.

Professional video signals are used by broadcasters to produce programming for a viewing audience. These signals can be transported over IP networks and share a number of characteristics, including uncompressed or lightly compressed video; intra-frame only compression (MPEG or JPEG); high bit rate signals (over 50 Mbits/sec); one-to-one or many-to-one system architecture; and are transported over private networks or encrypted channels on public networks.

IPTV

In delivery systems to viewers, IPTV is still a minor player compared to the huge numbers of viewers for cable, satellite, and terrestrial broadcast TV. However, growth is accelerating as new systems are launched and IPTV deployments continue to roll out around the world. AT&T announced a total of 549,000 U.S. IPTV subscribers as of the end of June 2008, representing an increase of 170,000, or 45%, in the preceding three-month period. Globally, the installed base is much larger, with PCCW (Pacific Century CyberWorks Limited) in Hong Kong reporting 882,000 subscribers, Orange/France Telecom reporting 1.3 million; the Internet service provider (ISP) called Free in France reporting over 3 million TV-enabled broadband subscribers; and well over 2 million other IPTV subscribers elsewhere in Europe; and another 2 million in Asia. Overall, this market seems to be well positioned to add new subscribers at an accelerating pace through the end of 2008 and for the next few years.

One interesting new IPTV development is specifically targeted at smaller telephone companies and other local providers. For these companies, equipping a digital head-end to receive and process hundreds of television channels needed for a competitive IPTV package can be a very expensive proposition. To solve this problem, SES Americom (www.ses-americom.com) has developed a product called IP-Prime that delivers pre-compressed and formatted video signals to IPTV providers, who can then distribute those signals directly to subscribers with very little additional processing. This technology provides several benefits for telephone companies, including lower capital costs to establish video head-end; pre-negotiated content contracts reduce time-to-market; established middleware platform helps to limit costs of system integration, although billing system compatibility needs to be implemented and tested. Of course, monthly operating costs may be higher because of the cash flow necessary to cover the costs of SES-Americon's facilities and staffing. Limits are also placed on the types of software and set-top boxes that telcos can deploy. But overall, given the high costs of building a digital head-end, IP-Prime appears to be a good solution for smaller telephone companies.

In another new twist in video delivery to the home, Sezmi (www.sezmi.com) is offering a combined RF broadcast and IPTV service called TV 2.0. The idea is that the popular network channels are broadcast over the air for free and can be viewed live or recorded and played back on demand at a later time. RF spectrum can also be leased from DTV broadcasters to allow distribution of private licensed content, for channels such as ESPN or Discovery that do not have local broadcast outlets. Programs that are not broadcast over the airwaves are streamed over a broadband connection and can be stored for later playback. To make this work, Sezmi supplies a high-performance DTV receiver that is built into a powerful personal video recorder (PVR) that also has a broadband network connection. This technology is just entering technical trials; however, it may reach commercial status by the end of 2008. Of course, success depends heavily on Sezmi's ability to get enough bandwidth (both for RF broadcast and for consumer broadband) to make the system appealing to consumers. Plus, there is always the challenge of acquiring licenses for all the content.

Internet Video

Figure 1 shows a simplified diagram of how video signals are delivered over the Internet. Beginning at the source, video content is captured and ingested into a workstation for editing and compression by the author/filmmaker. The compressed file is then uploaded to a service provider, where the video can simply be stored in its uploaded format, or, in the more frequent case, the video is transformed by changing the compression format and placing the content into a standard container format along with metadata related to the video. The video can then be published on the service provider's website. Users who visit the website will be given thumbnails of one frame of the video along with a title/description, and will be invited to click on the video to view it. When a video is selected, the server delivers a file to the user's browser in the form of a stream or one or more files. As soon as the video arrives at the user's device, the user's browser software reads the container format, loads the appropriate software plug-in to decode the video, and then feeds the video file to the plug-in for display to the user.

Internet video traffic continues to grow rapidly in the U.S. According to comScore (www.comscore.com), over 12 billion videos were viewed online in the U.S. in May 2008, representing a 45% gain over May 2007. More than 142 million viewers watched an average of 85 videos in March, or about 2.7 videos per user per day. Online video viewing is not limited to a single country; in fact, according to comScore, 27 million U.K. viewers each watched 4.1 videos per day in March, and 30 million German viewers each watched 3.7 videos per day in May.

Although YouTube is still the dominant force in Internet video delivery, new providers are arriving at a rapid clip, many of which are seeking to offer creative alternatives to short clips played with Flash video. For example, Hulu (www.hulu.com) has done a lot for long-form Internet video delivery supported by advertising. One impor-

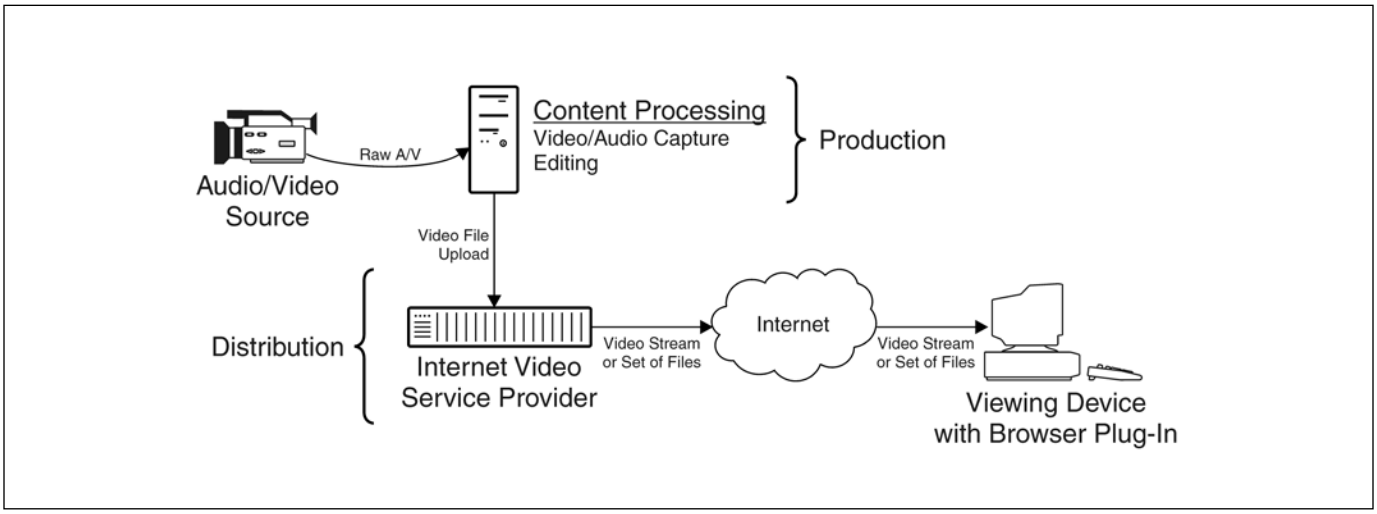

Figure 1. Simplified diagram of Internet video delivery.

tant item is that viewer sign-up is not required, which means that even casual viewers can easily get access. Also, there is no need to download a special player, because the videos are distributed in Flash format. Like a lot of other sites, Hulu has clips that have been extracted from movies and television shows; however, full episodes of some shows and movies are also available, and the list is growing. A few choice shows have every episode available, such as "Arrested Development" and "Firefly." Hulu has also made instream commercials somewhat tolerable by indicating their location on the player status timeline with a small white dot. Although viewers cannot fast-forward through the commercials, it is possible to tell when the commercials are going to appear. Overall, watching a few commercials certainly beats paying to download each episode, and Hulu is a great implementation.

America Online (AOL) Television (television.aol.com/in2tv) also offers full episodes of some television series online, especially favorites from the 1960s and 1970s. Unfortunately, the user interface has a significant amount of display advertising surrounding the playback window. On the plus side, a large amount of in2tv's content is shown with only a single pre-roll advertisement per program, so that there are no interruptions once the content begins to play.

Videoconferencing and Surveillance

It is amazing how completely the closed-circuit television (CCTV) industry has embraced video over IP technology. This adoption has been motivated by three strong forces:

- Video services can use existing datacom infrastructure.
- Multiple video signals can be carried on the same data cable.
- A single network can carry video, audio, voice, and data services, which are often needed in large surveillance and perimeter control systems.

Currently, hundreds of models of security cameras are available from major manufacturers that have multimegapixel IP video outputs at full 30 Hz frame rates. These cameras can be connected directly to Ethernet connections, and the video streams can be delivered to a variety of display and recording devices. For security applications, one very popular viewing device is a PC with a software video decoder and multi-image viewing software. There is also a variety of IP-network-enabled digital video recording equipment, some based on dedicated hardware and some in the form of software that will run on a standard PC platform. Based on the wide array of products available, it is clear that camera manufacturers have mastered the technologies needed for delivering IP video streams directly from a camera. Thus, it is somewhat surprising that more manufacturers have not taken the step of adding direct IP connections to consumer and prosumer camcorders.

Fox Television presented an interesting broadcast videoconferencing application during the Video Services Forum (VSF) meeting on May 13. For its Big 10 network, Fox has installed an HD videoconferencing endpoint in the athletic department of each of the Big 10 schools. These systems use a combination of H.323 conference signaling and H.264 video compression to bring live HD video back from each school, but at the relatively low bit rate of 2.5 Mbits/sec. At this rate, the compressed video contains a lot of image detail, provided that there is little or no motion in the scene (making it very useful for videoconferencing, but not for live sports). These signals are sometimes used live to air; therefore, the public Internet was not deemed to be reliable enough, so a dedicated IP network was constructed using leased telephone company circuits. Each camera supports pan/tilt/zoom remote control that was implemented over the IP network; this capability gives the staff at the Fox studio in Chicago the ability to control the remote cameras. Overall, this system has improved the availability of players and coaches for interviews, giving the broadcasts a

wider array of perspectives and making them livelier.

Contribution Video

Professional video is also being infiltrated by IP networking technology for three main reasons. First, the costs of IP networks continue to decline both in absolute terms and relative to the costs of other types of networks. Second, organizations are very interested in combining multiple services onto a common backbone, to further reduce bandwidth costs and simplify network management and maintenance operations. Third, broadcasters are becoming increasingly comfortable with the use of IP networks for their most valuable realtime content.

Work in the area of defining interoperability specifications for contribution video over IP continues at VSF (www.videoserviceforum.org). Two main activities are under way that will likely be used to create new standards that are complementary to the existing SMPTE 2022 standards that only deal with constant bit rate streams operating at 40 Mbits/sec or less.

- The first activity focuses on support for variable bit rate streams over IP networks. The currently proposed method uses fixed-size block of data for row/column forward error correction (FEC), and a user-defined latency limit. Whenever a data block becomes full it can be sent. If, however, a data block does not get completely filled before the latency limit is reached on the oldest packet in the block, then the block is padded with null packets before sending. Note that it is not necessary to actually send the null packets—the source can simply indicate to the destination the number of null packets that were used in the FEC calculations, and the destination device can simply add them back in before recovering the data. After some remaining issues are resolved (particularly in the area of packet sequence numbering) it should be possible to submit a draft proposal to SMPTE in late 2008 or early 2009. A project for variable bit rate streams used in contribution applications has also recently started.
- The second activity focuses on support for high bit rate IP video streams. Work in this area is proceeding on two fronts—one dealing with uncompressed serial digital video at the 270 Mbit, 1.5 Gbit, and 3 Gbit rates, and another dealing with motion JPEG 2000 compressed streams. Mappings are being developed that place the video signals into realtime protocol (RTP) packets, which can then be transmitted using user datagram protocol (UDP) over IP services. Progress is being made in both uncompressed and JPEG 2000 areas, with the goal of developing proposals that can be fed into the SMPTE standardization process.

One important difference between the work being done by VSF and some of the video-over-RTP work done previously by the Internet Engineering Task Force (IETF) is in the area of audio embedding. Philosophically, the IETF prefers to have video and audio signals mapped into separate packet streams, so that participants who do not have enough bandwidth for the video signal can receive the (relatively low bandwidth) audio signal only. However, for this work, VSF is choosing to use embedded audio. The benefits for broadcasters of doing this are significant for two reasons:

- Broadcasters are highly experienced with embedded analog and digital audio, and often have the equipment needed to combine the two signals.
- By not splitting the video and audio into separate streams, one large potential area for introducing synchronization problems is eliminated.

In terms of deployed applications, uncompressed HD video over IP (operating at roughly 1.6 Gbits/sec) was used for the Union of European Football Associations (UEFA) Euro 2008 championships held in Switzerland and Austria in June. This tournament, which occurs every four years, is held between the national teams for countries from around Europe, which had to qualify for the 16 available tournament berths. There are eight stadiums, each having capacity for 30,000 spectators, located in four cities in Switzerland and four cities in Austria. Video from each stadium was routed back to the broadcast center in Vienna using an IP over SDH backbone. During each match, every stadium is equipped with a total of 30 gigabits of IP bandwidth, which is enough to handle 12 uncompressed HD video signals for multilateral feeds (fed to all broadcasters) and 18 lightly compressed (2:1) SD signals for unilateral feeds (dedicated to individual broadcasters). More details can be found at www.medialinks.com.

Conclusion

Overall, use of video over IP continues to expand in 2008, both in terms of market penetration and new applications. Clearly, broadcasters, service providers, and end users have come to recognize the advantages of IP technology in terms of cost, flexibility, and market reach. In addition, as users gain increasing experience with the reliability and performance of the technology, even more video applications will find their way onto IP networks in the coming years.