CABLE IP VIDEO DISTRIBUTION

Timothy O'Keefe, Robert Sayko and C.J. Liu; Edited by Sean Carolan
AT&T Labs

Abstract

Advances in IP technology are providing Cable Operators with the opportunity to offer innovative interactive services while sharing transmission capacity between IP and traditional video services. This paper will examine how QoS, MPLS, VPN, and Multicast standards for IP networks fit against the demands for quality video distribution, and will explore how these technologies can be used to move nationwide video content distribution to a national IP-based network.

INTRODUCTION

Cable operators are connecting their local networks with IP services to facilitate Internet Access and Voice over IP services. Within their local networks, operators are using Ethernet to distribute VoD streams. The Next Generation Network Architecture (NGNA) is introducing new applications for IP technologies in the delivery of services. NGNA proposes the delivery of a greater volume of video services, and the routing of video service control signaling, over IP based connections.

Advances in IP technology are enabling these new services, and we envision more innovative services and service integration using IP capabilities for the future. Implementations of QoS, MPLS, VPNs and new codec schemes are enabling video services over IP. This paper will examine how QoS, MPLS, VPN, and Multicast standards for IP networks fit against the demands for quality video distribution, and will explore how these technologies can be used to move nationwide video content distribution to a national IP-based network.

WHY DISTRIBUTE VIDEO OVER IP NETWORKS?

When cable operators offered only traditional video service, satellite access linked their local distribution network access to nationally distributed video content. New Internet Access and VoIP services are fundamentally different than traditional video service; they are more effectively delivered using terrestrially based IP networks, rather than latency-encumbered satellite delivery. Because Internet Access and Voice services send traffic throughout the world, operators need to connect their IP networks to the national and worldwide networks that form the greater Internet. Cable operators are servicing this need by connecting their headends to national terrestrially based IP networks to deliver these services.

For IP networks, scale offers significant cost advantages. A fiber connection into a national network can expand from 100Mbps to 10 Gbps just by upgrading the equipment used at the connection's endpoints. While the bandwidth can increase 100 fold, the cost difference between the 100 Mbps electronics and the 10Gbps electronics is relatively small. A company could effectively reduce the unit cost of their Internet Access and Voice services for a given headend by expanding their existing IP network capacity, and by using the expanded capacity to receive video services as well.

The fiber medium provides much easier capacity expansion than a satellite-based system. Fiber path construction typically installs a bundle of fibers. As the capacity of an individual fiber fills, lighting another fiber in the bundle can easily activate new capacity. Dense Wave Division Multiplexing (DWDM) also enables each single fiber to multiply its capacity. This easy access to capacity creates the opportunity to apply less compression of the digital video during transmission, thus receiving higher quality video .

Additional efficiencies are gained by using a common technology across all services. When Voice, Video and Internet Access are combined into the IP network, a common set of IP network equipment (routers, switches, etc) can be used across all the services. Technical engineering and operations staffs then have fewer technologies to master. An Operations center can use common IP based network monitoring and management tools. The common use of IP technology across all the services creates efficiencies throughout the business operation.

Video content suppliers are now more likely to have access to IP networks. A well-established national IP network likely passes as close to content producers as it does to content distributors. The content suppliers can input their content to the IP network at any point along the path. Just as the cable operator's IP connection links their Internet Access and VoIP customers to the world, their IP network connection can link them to the world for access to new, interesting content.

REQUIREMENTS FOR A QUALITY IP DISTRIBUTION OF VIDEO CONTENT

Video Delivery Requirements

The requirements for distribution of Video-over-IP are guided by two realities. The first of these are the technical capabilities that must be achieved to present an acceptable video experience. The second are the business drivers that then impose additional technical requirements.

At the present time, formal technical standards that support requirements of real time Video-over-IP, as promoted through a recognized standards organization such as IETF and compliant with MPEG standards for video encoding, are being developed based on needs that are still emerging from the industry. We recognize the IP transport must facilitate transmission of the MPEG stream such that the technical requirements that apply at the endpoint are met. The goal is to guarantee a delivery across the national IP network that will support MPEG requirements for video performance at the endpoint, and will also offer a guarantee of reliability and availability as it relates to uninterrupted service.

The standard measures for video quality include the following:

- Inter-packet jitter
- Packet loss
- Packet arrival order
- Availability

Latency, while critical for real-time interactive applications like telephony, is not as critical for video distribution. Latency of a well run terrestrial IP network will be substantially less than that of satellite television. Inter-packet jitter - the variance in latency from packet to packet – can, for any reasonable jitter that is expected in an IP network, be accommodated by buffering of IP packets at the receiver. Packet arrival order can also be overcome through buffering and packet reordering in the IP endpoints.

Packet loss is a critical measure of performance for IP video distribution. While transport protocols, like TCP, can request retransmission of lost packets, this is not practical in a video distribution application where a single source is sending a multicast stream of packets to a large number of endpoints. The stream must support transmission, and associated IP overhead, of up to 3.75Mb/s of MPEG2 data for a Standard Definition ATSC transmission; 19.3Mb/s for a High Definition ATSC transmission. A single lost IP packet can translate into a loss of 7 MPEG2 packets.

Network Reliability

Service availability is another key technical capability required to offer broadcast quality content delivery to a cable headend. Since satellite is typically the delivery method for all channels, satellite reliability can be considered a benchmark for video service delivery. Operators don't think about availability on an individual channel basis when all channels arrive via the same transmission path.

However, when various delivery options exist and slight changes in availability performance can be traded against economic benefits, a range of acceptable availability might become part of the delivery system decision. For example, IP networks can be designed for 99.99% availability and 99.999% availability. These two metrics represent a difference of 48 minutes of availability in a year. Yet they represent a significant difference in cost of network to design, build, and operate. Given the option to make the availability/distribution-price decision on a per channel basis, a cable operator might elect to receive some channels on the lower cost connection and others on the higher cost connection. If satellite is a third, more reliable service delivery, the cable operator might elect to move some lower value channels from satellite deliver to IP network delivery in order to create satellite bandwidth for more high value content.

While no ubiquitous industry standard currently exists for this, it is expected that a video delivery system should employ an architecture that would meet very high level of availability. While a 99.999% availability metric represents 5.3 minutes of outage in a year, cable operators expect interruptions to be few, of short duration and restoration must be seamlessly engineered so that continuity of the video is preserved for the user. Based on these criteria, a successful re-convergence of an interrupted video stream should occur within 1 second or less of detection.

Performance benchmarks and seamless end-user reliability guarantees are not all that drive the IP video distribution technical requirements. In fact there are several other important business criteria that drive addition technical requirements that must be met in order to offer a viable service. These criteria are:

- Costs
- Competition
- Functionality
- Efficiency

The costs of converting to and using Video-over-IP distribution must meet certain thresholds for initial investment (e.g.

CapEx) and the ongoing cost of running the business. The first threshold is met when the reliability and performance of the IP distribution reduces the investment in expensive terminal equipment to control and groom the received video product. The ongoing costs are even more important to the MSO and are also tightly coupled with the three other criteria: competition, functionality and efficiency. Satellite distribution is not the only competition for terrestrially-based Video-over-IP; other competition comes from satellite broadcast providers (e.g. Dish Network, DirectTV). The level of service offered by these competitors is what actually establishes the market's benchmarks for the Video-over-IP distribution service. The bottom line is that IP video distribution must at least marginally beat the cost of alternate providers for equivalent services or provide significantly better service(s) at only incremental cost increases

One key advantage that Video-over-IP must provide over these competitors is the ability to deliver a significantly larger (almost limitless) amount of content. A starting point for the service should begin where the competition leaves off. For example:

- 200 – 400 channels of SD programming
- 25 channels of HD programming
- 100 channels of CD quality music

Functionality is another important consideration for the MSO. Frequently, a Conditional Access connection is required to support the video service delivery, and the CA authority is located away from the headend. These implementations today need to coordinate access to the satellite video services through a separate terrestrial CA network connection. As previously mentioned, the IP connectivity is two-way. It can already support VoIP and Internet Access service. The same IP connection can be used for the CA connection. The use of IP for video service delivery and CAS may create new opportunities for CA mechanisms.

The final criterion, efficiency, is very important to the MSO as video content options expand, consumers become more sensitive to picture quality, and new innovative services create interaction between video, voice, data, and wireless services. The video service delivery must be capable of expanding overall capacity rapidly, support varying levels of video program compression to ensure high quality content, and provide delivery protocols that easily inter-work with the other services. Finally, it should be easy to add new endpoints to the video service distribution to facilitate easily adding content suppliers and local content distribution networks.

AN OVERVIEW OF IP TECHNOLOGIES

Listed below are the various IP protocols that can be applied to allow video streams on an IP network meet the requirements stated above. These brief descriptions are here just to provide a high level review of the terms, as they will be used throughout the paper. More detailed explanations can be obtained from the IETF and vendor web sites.

IP Encapsulation of Video Frames

MPEG2 frames can be encapsulated in an IP packet. Seven MPEG2 frames are typically combined into one IP packet since this creates a packet size within the 1500 byte limit of Ethernet and enables the packet to be easily moved between layer two transmission protocols. While MPEG2 is currently the most common video stream protocol, other encoding protocols are also easily placed into IP packets. These IP packets are typically sent as UDP frames. The UDP protocol does not include the ability for the receiving end point to request the retransmission of a lost packet. A national network with 50msec latency might actually allow for 100msec round trip to be used to retransmit a packet. However, Forward Error Correction (FEC) and interleaving packets are more common methods for correcting for lost packets. The Real-Time Protocol (RTP) is also utilized to help sequence packets on the receiving end. If two packets arrive out of order, the sequence numbers in the RTP protocol will allow the receiving end to assemble the video stream in the correct sequence.

Quality of Service (QoS)

QoS standards allow traffic to be marked for specific handling when the network is congested. The most basic handling of IP traffic is called "best effort". There is no special handling of this class of traffic. The network will try its best to get the packet through as fast as the network will allow. Packets are processed in the order they arrives, first in/first out. The highest quality for traffic handling is called Real-time class. This class gets the top priority from the network equipment. This class of traffic will only be dropped if all the capacity allocated to this class is consumed. Network operators typically allocate enough capacity to this level to support all Real-time traffic they've agreed to accept to avoid any dropped packets once it has entered their network. Other classes between Best effort and Real-time define specific behavior during network congestion periods to prioritize and drop packets based on the needs of that traffic class.

Virtual Private Network (VPN)

This segregates traffic on the network such that the network operator can keep some traffic flows separate from each other and from the public Internet. Public Internet traffic just reacts to the IP destination, and will pass any packet to any requested destination. VPN allows the operator to establish additional rules for traffic flow. These rules can include restricting participation in the traffic flows, encryption of the packets in a particular flow, and packet routing based on VPN identification instead of IP address,

Multicast

Most traffic flows in an IP network are point-to-point transmissions. A single source wants to deliver a packet to a single destination, a Unicast flow. Multicast is used when one source wants to send the same information simultaneously to multiple recipients. The multicast routing protocol builds a tree distribution map for all the recipients on the network. A single copy of the packet moves through the network until it arrives at a branch in the distribution tree. At the branch, the network duplicates the packet and sends one copy of the packet down each branch of the tree.

MPLS Traffic Engineering (MPLS-TE)

Traffic Engineering is an extension of the MPLS standard that provides the network operator with more control over the path packets take through the network. This control serves to aid capacity management as well as fast failover recovery. The operator can define specific paths for defined MPLS flows. TE also allows the operator to define a specific failover path for an MPLS flow. This pre-defined failover has become known as Fast ReRoute (FRR).

FRR enables the network to recover a data stream in < 100msec because the alternate path is already known. Since it was predefined, the network doesn't need to take time to discover alternate routes.

POTENTIAL IP IMPLEMENTATIONS FOR IP VIDEO CONTENT DISTRIBUTION

The most basic implementation would be to Unicast an IP video stream across a wide area IP network. This implementation would require the video stream to be replicated at the source for each headend destination. It provides little if any protection of the content, implies a best effort delivery and does not scale very well when there are multiple destinations. It's easy to see that this doesn't fit very well against requirements for quality video programming distribution. The following is a discussion of how other protocols can be applied to improve upon this basic implementation.

Model A: A Multicast solution.

We can make the basic implementation more efficient by implementing multicast. With video services, we expect a single source and many recipients. Multicast provides the ability for the network to take in one video stream and distribute it to multiple recipients by replicating the stream only when necessary. This greatly reduces the capacity demands from the basic implementation and enables the operator to transmit many channels to many end-points within a reasonable network capacity allocation. For example, one allocation of 2 Gigabytes could distribute 300 channels for SDTV and 40 channels of HDTV to any recipient connected to the network.

Model B: A VPN Multicast.

An IP network configured with VPN and Multicast makes a significant improvement over the basic implementation. Implementing the VPN protocol helps protect the video content from being intercepted by unauthorized parties during transmission. The VPN protocol restricts the traffic to specific end delivery points. The network operator controls access to this traffic stream to authorized recipients.

The application of QoS markings to Multicast IP packets is currently not available across all vendors and routers, so the reliable delivery of these streams will usually be dependent on the network operator allocating sufficient bandwidth throughout the network to avoid congestion delay. In addition, a route failure in a traditional IP implementation can cause up to a 10 second outage while the IP network recalculates routes around the failed connection. (An advanced network can provide SLAs that are significantly shorter.) Once the route is reestablished, it can take many more seconds for the video stream end point to resynchronize, re-establish buffers, and return the video stream to a stable flow. This amount of video stream loss far exceeds our requirements for availability.

Model C: MPLS-TE for Fast Network Recovery and QoS

It is desirable to have sub-second restoration capability for video distribution, which can pose problems when required of an IP network. For conventional IP networks, it typically takes five to ten seconds to have traffic rerouted around connectivity failures; either failed links or failed nodes. It's possible to tune nodes on the IP network (using "hello interval", "dead timer", and by leveraging a calculation of the hold time between two consecutive SPF calculations) to improve IGP convergence time. However, it's still not easy to shave the recovery time down to a level below one second. FRR (FastReroute) in MPLS TE technology enables fail over time of less than 50 ms; an interval that matches the link restoration capabilities of SONET. Fast Reroute is initiated for a Label-Switched Path (LSP) when the feature is enabled for the associated LSP tunnel as a result of a configuration command on the head-end. The head-end router is responsible for informing all routers along the LSP's path that the LSP is requesting protection. The LSP tunnel head-end control module will keep RSVP informed of the status of the Fast Reroute attribute for all active LSPs. When the RSVP module in a Label Switch Router (LSR) [other than tail end] along the LSP's path learns that the LSP should be protected, it will initiate local Fast Reroute protection procedure to protect the LSP against possible failure of the immediate downstream link. Upon link failure, all protected LSPs switch to the backup path. FRR performs the operations to prevent the downstream routers (still along the path in use by the LSP) from tearing down the LSP, if the failure is also detected downstream.

Content delivery services require more significant guarantees for bandwidth rates and for Quality-of-Service (QoS) from the IP network than conventional IP services do. As currently formulated, the leading IETF-endorsed architecture for QoS maintenance of differentiated services, "Diffserv", is strong on simplicity and weak on bandwidth guarantees. In the case of network congestion events, different services would compete for the available link bandwidth. A strict priority queue that includes bandwidth policing for real-time traffic could be enforced, but packet loss and latency still cannot be guaranteed if the rate of the incoming real time traffic stream is higher than the available bandwidth for the real time traffic. The network operator needs to know how much video traffic will be coming into the network so allocate the necessary bandwidth through the network. DiffServ-aware Traffic Engineering (TE) is a tool for network operators to implement the appropriate bandwidth allocations. DS-TE is meant to enable computing path per class with different bandwidth constraints, and perform admission control over different bandwidth pools. OSPF extensions for DS TE allow advertisement of unreserved TE bandwidth, at each preemption level, for each class type. In DS aware TE tunnels setup time, LSP signaling includes class type as a tunnel parameter, in addition to bandwidth, label, explicit route, affinity, preemption, adaptability and resilience. Class-type aware call admission control will be performed at each LSR during the DS TE tunnel setup. Rate limiting at the head end of the DS TE tunnel can be configured to ensure the traffic into the tunnel does not exceed the provisioned tunnel bandwidth.

Unfortunately, MPLS does not support multicast. MPLS tagging assumes a packet coming into the network can be mapped to one exit point on the network. MPLS can accept a packet coming into the network from one of many possible entry points, assign a tag representing the appropriate exit point, and efficiently direct that packet to the correct single exit point. Multicast wants to do the opposite. Multipoint processing assumes a packet entry at a single exit point should be distributed to many exit points. Relevant IETF working groups are discussing changes to MPLS that could support multicast traffic, but it may be a year or more before those changes begin to appear in network equipment.

As it happens, MPLS-TE no longer has a monopoly on Fast Reroute; standards bodies and the vendor community are working on Fast Reroute on native IP connections. This approach must also be evaluated and compared to MPLS-TE as standards emerge.

Model D: We need a Solution that incorporates QoS, Multicast, Fast Reroute, and VPN.

Possibly the best solution is a combination of all these technologies. We need multicast to make efficient use of network capacity. We need QoS to ensure consistent, on-time delivery of the packet stream. We need VPN to enable access control. Finally, we need Fast Reroute capability to minimize interruptions to the video stream caused by network failure events. Standards bodies and vendors are working to make the whole combination available.

Point to Multipoint Traffic Engineering Label Switched Path (P2MP TE LSP) is currently being proposed in support of the construction of a Point-to-Multipoint (P2MP) backbone network for multicast services. In such a scheme, a P2MP Label Switched Path (LSP) will be set up between an ingress Provider Edge (PE) and multiple egress PEs; the ingress PE would accommodate a multicast source, and the multiple egress PEs would accommodate multicast receivers. Ingress/egress PEs at the edge of the multicast network will handle subsequent multicast routing. The P2MP LSP will be set up with TE constraints and will allow efficient packet replication at various branching points in the network. The proposed P2MP TE LSP would be established by setting up multiple standard P2P TE LSPs. If each P2P sub-LSP is protected by its backup-tunnel, the multicast video traffic can be protected by the standard FRR TE mechanism, therefore, ensure recovery within 50 ms in case of link/node failure events.

A caveat is: even though there are obvious benefits of deploying TE tunnels in IP network, there are concerns about its scalability and the complexity it adds to network operation. For a facility based ISP that owns the physical links and infrastructure of its IP network, capacity constraint is a relatively minor issue compared to other ISPs which have to purchase or lease capacity from other providers. It is hard to justify sending all IP traffic into fully meshed TE tunnels ubiquitously deployed for a facility based ISP. Instead, only special traffic, such as VoIP, broadcast video, video conferences, or VoD transported in IP network, are candidates to be carried in MPLS TE tunnels. This requires the traffic that enters a configured MPLS TE tunnel get preferential treatment over all other traffic by all routers' queuing and congestion avoiding mechanism along the path. TE queues for configured MPLS TE tunnels in every router the tunnel traverses had been proposed. The proposed TE queues for P2P unicast tunnels can be extended to P2MP multicast tunnels.

IN CONCLUSION

Cable operators have made IP protocols an important part of their network services for Internet Access, VoD, and VoIP. NGNA is creating additional opportunities for IP based services in the network. Advances in IP technologies provide are creating the opportunity to move national broadcast video distribution to IP networks. The application of Multicast and VPN with sufficient bandwidth allocation can provide very reliable video distribution and could be used for some channels today. The application of Fast Reroute capabilities can bring recovery from link outages to around 100msec and make IP video delivery even more reliable. Mixing national video distribution with Internet Access and VoIP traffic creates economies of scale throughout the business. Putting all service delivery on IP enables new opportunities for delivery to end consumers and innovative service integration. Moving national video content distribution to IP networks could be the next big service breakthrough for Cable operators.

REFERENCES:

1. Extended RSVP-TE for Point-to-Multipoint LSP Tunnels, draft-yasukawa mpls-rsvp-p2mp-04.txt by S. Yasukawa, A. Kullberg, and L. Berger
2. Requirements for Point to Multipoint extension to RSVP-TE, draft-yasukawa-mpls-p2mp-requirement-01.txt
3. Performance Analysis of MPLS TE Queues for QoS Routing**,** Yihan Li, Shivendra Panwar, and C.J. (Charlie) Liu, Proceedings of Applied Telecommunication Symposium/ASTC, 2004 (pp.170-174), April 18-22, Arlington, Virginia.
4. MPLS TE Queue Creation as a Mechanism for QoS Routing, C.J. (Charlie) Liu, Proceedings of Applied Telecommunicatio Symposium/ASTC March 30-April 3, 2003, Orlando, Florida