IP for Video Heads

By:

Dave Brown
Cable Solutions Manager
Cisco Systems

davebr@cisco.com

June 24 — 27, 2008

SCTE Cable-Tec Expo® 2008

Table of Contents

Introduction

Internet Protocol (IP) is a relative newcomer to cable video delivery which traditionally has been sent directly from a video headend (where the video was acquired and processed) to the Hybrid Fiber Coaxial (HFC) network for delivery to subscribers. Over the last five years, cable operators have found that IP networks allow them to achieve a number of key benefits. Cable operators can now:

- Consolidate their headends to achieve CapEx and OpEx savings
- Send video over virtually unlimited distances without any quality degradation as occurs in RF networks
- Dynamically adapt to network conditions and automatically re-route around failures
- Deliver video traffic over the same network as their high-speed data and digital voice services

Before IP, cable operators ran multiple parallel networks to deliver different services. In particular, video was distributed over a completely different network than the IP network used for voice or data traffic.

Moving to a “converged” network—capable of transporting multiple services simultaneously—is tremendously desirable. It reduces capital and operational expenditures since there is only one network to build instead of multiple networks. It also improves efficiency through the statistical multiplexing of multiple services.

This paper provides an overview of IP and related networking technologies. The paper also illustrates how these technologies are used in modern cable networks.

Ethernet and IP Networking Technologies

Before getting into IP technologies, it is helpful to first consider the "Layer 2" technology—Ethernet—on which most IP networks are built. Ethernet originally worked in a shared mode using an algorithm called Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Using CSMA/CD, Ethernet-connected devices listened to the Ethernet segment and determined if any other device was transmitting. If no, the device could start sending its Ethernet frames. If two devices happened to start transmitting at the same time, there would be a collision, and each station would back off for a random amount of time and try again.

The simplest networking device, which operates in the CSMA/CD mode, is called a *hub*. A *hub* has no intelligence built in, and simply broadcasts Ethernet frames received on one port to all other ports. Hubs do not scale well and can result in "broadcast storms" if there is any loop in the network. They also result in too many collisions when the number of devices on a segment grows too large, and too many devices contend for bandwidth. As a result, cable networks are not built with hub technology, and instead, use a combination of switches and routers (described below).

Today's networks no longer use CSMA/CD; switches and routers each operate in a mode where they buffer the full contents of an Ethernet frame or an IP packet and then send them out only on the appropriate ports. To determine the "appropriate" ports on which to send out the data, there are two different methods of delivering video traffic over the Ethernet network: bridging (called switching when it is done in hardware) and routing. Layer 2 switches perform bridging, while routers perform routing. Layer 3 switching lies somewhere in between, able to perform either Layer 2 switching or basic IP routing.

Ethernet Bridging occurs at the Ethernet layer, and the bridging device does not make use of the IP address *at all* as it delivers the traffic. Instead, bridges and switches "learn" by inspecting the Media Access Control (MAC) addresses of incoming traffic, and building a learning table that tells them which port to use to forward the packets. (MAC addresses are usually given in hex format, of the form <00-FF>:<00-FF>:<00-FF>:<00-FF>:<00-FF>:<00-FF>, for example C8:F9:32:FC:BB:99.) Figure 1 shows an example network, where the MAC address of station A has been learned by the switch.

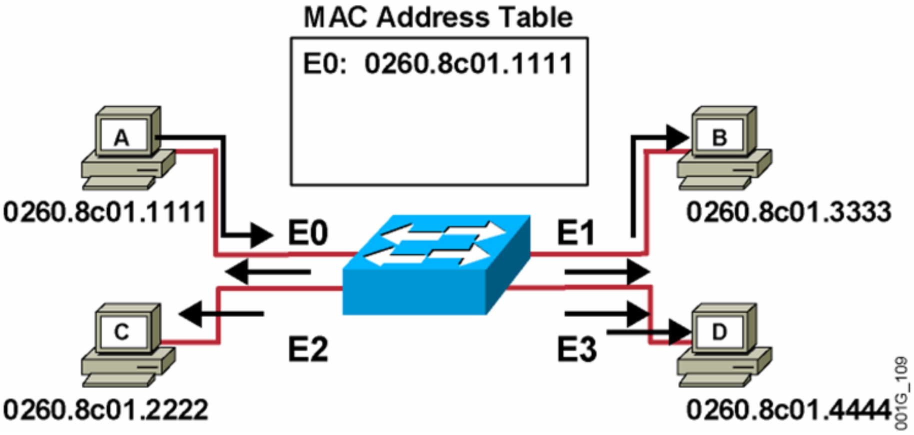

Figure 1

Bridging (or Layer 2 switching) can be used in small networks, but is not very scalable. Furthermore, it requires the use of a protocol to avoid the endless forwarding of unknown Ethernet packets around "loops" in the network. The Spanning Tree protocol is the best known loop prevention protocol, although many others have also been devised.

Most cable networks, instead, use IP routing, which uses the IP addressing information at "Layer 3" to forward traffic to its destination. The advantage of using IP addresses over using Ethernet MAC addresses (as is done in the bridging case) is that MAC addresses are not location-specific and are generally assigned by the hardware manufacturer. In contrast, IP addresses are assigned by the network operator specifically to indicate where in the network the IP device is located.

IP Addressing

IP addresses are 32 bits in length, and are broken up into 4 different octets (each separated by a ".") IP addresses are typically represented in decimal format, in the format:

<0-255>.<0-255>.<0-255>.<0-255> (for example 244.32.17.221)

IP addresses can be subdivided into two components: a "network" address and a "user" address. The network address indicates a single logical grouping of devices, whereas the user address indicates a specific device within the logical grouping, or network. A "netmask" is configured in the router so that it knows which part of the IP address represents the network and which part represents the user. The netmask can also be represented in the form:

<0-255>.<0.255>.<0-255>.<0-255>

The netmask can be better understood by converting to binary, with each octet being represented by 8 bits. In binary, a netmask will be represented by a series of "1s" followed by a series of "0s". The "1s" indicate the "network" portion of the IP address, while the "0s" indicate the "user" portion of the IP address. For example, the netmask might be as follows:

Netmask: 11111111.11111111.11110000.00000000

In decimal, this netmask would be represented as 255.255.240.0. Using another popular notation, it could be represented as "/20", where 20 represents the number of 1s in the binary representation of the netmask.

This netmask indicates that the first 20 bits of an IP address represent the network, and the last 12 bits the user device. To determine the "network" of the example IP address above of 244.32.17.221, convert to binary and perform an "AND" operation with the netmask:

11110100.00100000.00010001.11011101 AND
<u>11111111.11111111.11110000.00000000</u>
11110100.00100000.00010000.00000000

Converting back to decimal, the "network" for this address is 244.32.16.0. Therefore the router will advertise network 244.32.16.0, and all clients connected to that network can be assumed to be in the same logical location. When an IP packet is addressed to 244.32.17.221, it will first be sent to the router advertising network 244.32.16.0, and the router will then locally ensure delivery to the correct IP destination.

Address Resolution Protocol

Once a packet has been routed to the router advertising a given network, that router must then send it to the final destination device. How it does this depends on the Layer 2 protocol being used.

In an Ethernet network, any given IP address maps to a single Ethernet MAC address—the hardware address assigned to the Ethernet interface on the device.

When a router receives a packet for a destination IP address, it may not know the MAC address of the device. As such, the router sends out a broadcast to all users on the network, asking, "Who is 244.32.17.221?" The device with that address then responds with its MAC address: "C8:F9:32:FC:BB:99 is 244.32.17.221." The router stores this address in its "ARP table" and subsequent packets can be sent to that IP address without first issuing an ARP request.

IP Transport Protocols for Video Networks

Three commonly used IP transport protocols exist:

- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Real-time Transport Protocol (RTP).

TCP is known as a *reliable* protocol, since the protocol defines sequence numbers and automatically requests re-transmission of lost packets. TCP also incorporates a windowing mechanism that automatically reduces the transmission rate in response to packet loss. TCP is normally not used for real-time content distribution since any replaced packets are likely to be "too late,"and the windowing mechanisms inappropriately can reduce the transmission rate for video traffic. As such, TCP is not used as a transport protocol for cable video traffic.

User Datagram Protocol (UDP) is the protocol most commonly used today for video distribution over a cable network. It is the de facto standard for sending transport streams over IP. With UDP, there are no sequence numbers and no re-transmission requests. If a UDP packet is lost, video artifacts may appear on the end-user's viewing device. The severity of the video impairment depends on the type of packet lost. A lost "I" frame will generally result in the most severe impairment.

Either Single Program Transport Streams (SPTSs) or Multi Program Transport Streams (MPTSs) can be transmitted inside a (UDP-encapsulated) IP packet. Most cable networks limit packet size to 1500 bytes or fewer—the limit for Ethernet frames unless Ethernet "jumbo frame" extensions are enabled. Since an MPEG packet is 188 bytes, that means that seven MPEG packets can be transmitted in a single IP packet. (This fact also highlights the need for very low packet loss rates for video distribution networks since a single lost packet means seven MPEG packets will be dropped.)

As resiliency requirements have grown and IP network latencies have improved, Real-time Transport Protocol (RTP) has gained steady momentum as an alternative to UDP. RTP is like TCP in that it maintains sequence numbers. RTP can be used in a number of ways to increase resiliency. One method is to send an identical multicast stream over diverse paths in the IP network. The receiving device can detect when packets have been lost over one path and fill the buffer with the correct packet (based on sequence number) from the alternate path. Another way of using the sequence numbers is to actually make re-transmission requests in real-time when a lost packet is detected, and have a backup device send a replacement packet before the buffer empties. Such real-time packet replacement can occur in less than 100 ms in a well-designed network.

Routing in the IP Video Network

Routing in the IP video network is generally accomplished using the Open Shortest Path First (OSPF) protocol. Border Gateway Protocol 4 (BGP-4) is also used for some applications, but will not be covered in detail here.

Each router in an OSPF network maintains a "link state" table. Through this table, the router knows which link it must use to reach any other router in the network. If any link fails, the router can very quickly re-calculate and determine which secondary link can be used as an alternate route to any destination formerly reached on that failed link. Routers periodically send full "link state advertisements" to other routers informing them of what routers are reachable through which links. "Hello" messages are also sent by default every 10 seconds, informing routers of other routers in the network with links that are "up."

IP Multicast

Until now, the discussion has focused on basic unicast IP, for one-to-one communications such as is used for Video on Demand (VoD). But a range of IP addresses is also set aside for use as IP multicast addresses, for one-to-many applications such as Switched Digital Video (SDV). Specifically, all IP addresses beginning with 1110 (meaning they are in the range **224.x.y.z** through **239.x.y.z)** are reserved for multicast. Each address in this range represents not a single receiving end station, but *every* station that has requested access to that multicast group.

Some of the addresses in the multicast range are reserved for specific purposes: **239.x.y.z** addresses are “administratively scoped” addresses, or private addresses. **232.x.y.z** addresses are reserved for source specific multicast (described in more detail in a section below).

Netmasks are assigned for IP multicast addresses in much the same way as they are assigned for IP unicast.

Benefits of IP Multicast

As discussed above, the Internet Protocol efficiently handles both one-to-one (unicast) and one-to-many (multicast) modes of operation. This is a particularly important characteristic since video services use a combination of unicast and multicast. IP multicast is much more efficient for one-to-many services since a single copy of each packet can be replicated by routers throughout the network, thereby saving bandwidth. See Figure 2.

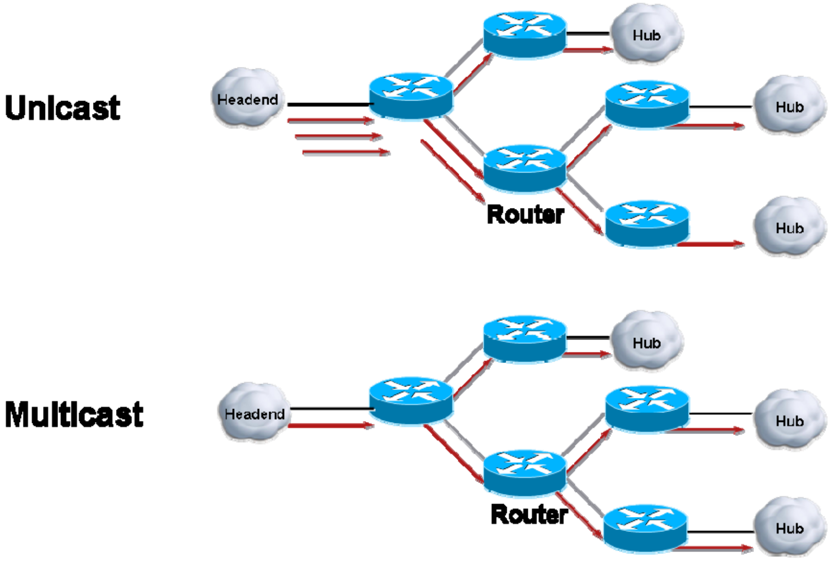

Figure 2

Broadcast video and switched digital video (SDV) are one-to-many in nature and use multicast. Since IP natively supports multicast, it can send a *single copy* of a broadcast channel (e.g., CNN) around a typical metro ring network, and feed that same channel to each hub site on the ring.

How Does IP Multicast Work?

The basic idea of IP multicast is to create a middle ground between unicast (where only a single device receives each packet) and broadcast (where every device receives each packet). With IP multicast, devices request which “multicast group” they are interested in being a part of. Routers then ensure delivery of the traffic to each receiving station.

In a cable network, multicast groups are heavily used for a variety of purposes. For example, a multicast group might be used to represent a specific multi-program transport stream bound for a specific quadrature amplitude modulation (QAM) modulator being used across multiple hub sites. In this case, a QAM in each hub site needs to send a request to join the multicast group. Alternatively, it could be a distributed ad splicer that sends the request to join the multicast group, which then performs ad insertion on the channels in that transport stream and constructs a new MPTS which is sent out to a *different* IP multicast (or unicast) address for delivery to a QAM modulator.

IP multicast uses two different protocols for building multicast "trees" for distributing multicast traffic: Protocol Independent Multicast – Dense Mode (PIM-DM) and Protocol Independent Multicast – Sparse Mode (PIM-SM). Dense mode is best for local area networks where receivers are concentrated in one area, whereas sparse mode is best for distribution of multicast traffic across a large network with a relatively small number of devices receiving the multicast traffic over a large geographic area. Cable networks use "sparse mode" multicast distribution.

Any Source Multicast (ASM) vs. Source Specific Multicast (SSM)

The original technique for IP multicast was called Any Source Multicast (ASM). With ASM, a multicast receiver can specify the multicast group it would like to join, but not the sender (or source) of the traffic. There are two primary protocols used for ASM: Internet Group Management Protocol (IGMP) version 2, and PIM-SM.

IGMPv2 is a client-server protocol where the edge router acts as the server and an edge QAM or other edge device acts as the client. In this example, the QAM devices use IGMPv2 to tell the router when they want to "join" or "leave" a particular multicast group.

In a SDV network, there are two different schools of thought on when the QAM should perform the "join" of a multicast stream. One method is to join all SDV channels as soon as the QAM is connected, thereby minimizing latency for channel changes since all channels are received by the QAM all the time. The downside to this approach is that it only works when there is sufficient network bandwidth available to deliver all SDV channels – which becomes less likely as the number of SDV channels increases.

The alternative is to perform a "leave" when no subscribers are viewing a channel (e.g., when the last subscriber in a service group viewing a channel switches to a different channel) and send a "join" when a subscriber chooses a new channel that was previously unwatched by others in the service group. Network bandwidth usage is less with this approach, although channel change latency can be slightly increased.

ASM, as described above, does not allow control of the multicast source. However, in a cable network, one normally wants to control the source of the multicast traffic.

As an example of why it may be useful to select the source of a video stream, suppose two encoders in the headend (one primary and one secondary) are each encoding the same video stream. A client running ASM would have no control over the source, and upon requesting the multicast group may receive the feed from either the primary or the secondary. For better control, the cable operator might prefer to instead receive the feed only from the primary source and request the feed from the secondary source only in the event of a failure. Such a behavior requires Source Specific Multicast (SSM).

Source Specific Multicast (SSM) requires protocol changes; in particular, the use of IGMP version 3 and an updated version of PIM called PIM-SSM (Protocol Independent Multicast - Source Specific Multicast). Using SSM, the IGMPv3 request from a client device can specify both a multicast group address, as well as a source IP address in the join request.

The challenge with migrating to SSM is that many older devices do not support the newer protocols (IGMPv3 and PIM-SSM) required to support it. To support these older clients, but allow deployment of an

SSM network, routers can be configured to map IGMPv2 requests into an SSM request for a specific source.

Multi Protocol Label Switched (MPLS)

Multi Protocol Label Switching (MPLS) is a set of IP protocols that allows a fundamentally “connectionless” IP network to deliver connection-oriented services. MPLS was designed as a means of providing the basic benefits of an Asynchronous Transfer Mode (ATM) network – namely a deterministic path across a complex network, along with Quality of Service (QoS) – over an IP network. MPLS also provides the ability to provide network-based Virtual Private Network (VPN) services. For example, Virtual Private LAN Services (VPLS) is an MPLS-based service that provides Layer 2 VPNs, allowing businesses to connect multiple offices at different physical locations as if they were all part of a single network.

In MPLS networks, an IP packet is classified at ingress and assigned to a specific flow. A 20-bit “label” is then assigned to identify that flow, along with an 8-bit “Time to Live” or TTL. Three bits are assigned as “Experimental” bits that can be used for Quality of Service. One bit is used to indicate “bottom of stack,” enabling the routers in the network to know which label is the last label in networks where multiple MPLS labels are stacked on each other.

“Label switching” is performed through the network along the MPLS label switched path to the MPLS destination, at which point the MPLS labels are removed and the IP packet is forwarded. See Figure 3.

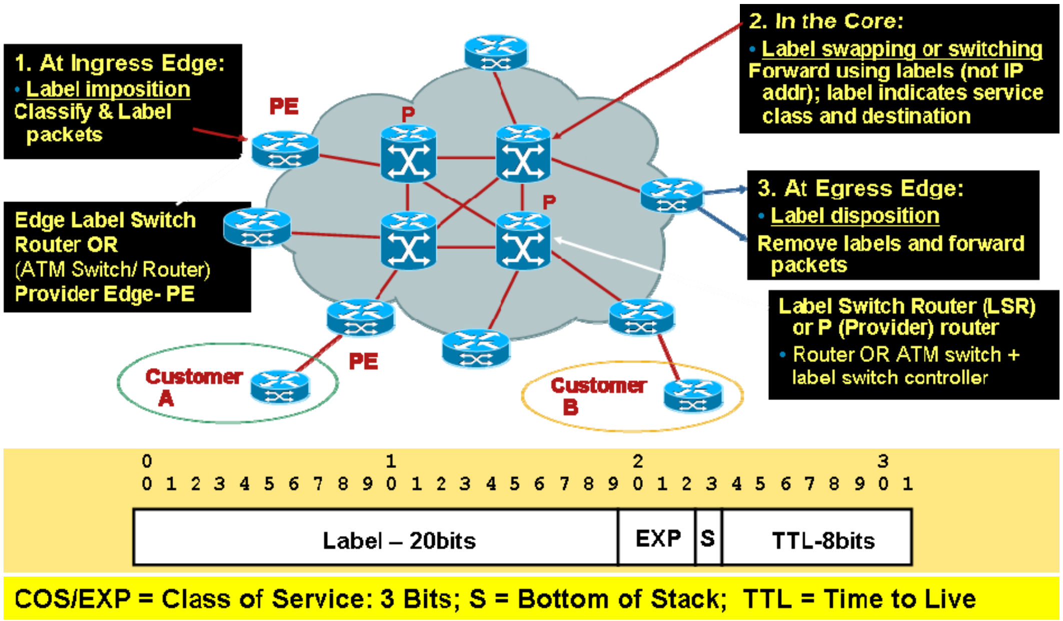
1. At Ingress Edge:
• Label imposition
Classify & Label packets
2. In the Core:
• Label swapping or switching
Forward using labels (not IP addr); label indicates service class and destination
3. At Egress Edge:
• Label disposition
Remove labels and forward packets
Edge Label Switch Router OR (ATM Switch/ Router) Provider Edge- PE
Label Switch Router (LSR) or P (Provider) router
• Router OR ATM switch + label switch controller
PE
P
P
P
PE
Customer A
Customer B
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Label – 20bits
EXP
S
TTL-8bits
COS/EXP = Class of Service: 3 Bits; S = Bottom of Stack; TTL = Time to Live

Figure 3

While MPLS can be used for video services, in cable networks, it is more often used for business services – including the aforementioned VPLS services, as well as Layer 3 VPN services. MPLS can also be used to provide traffic-engineered links over a core network. Video may be one of the services riding over MPLS links.

Ensuring High Resiliency in the IP Network

IP networks can make use of an array of different mechanisms for ensuring a high level of reliability in the network. This section considers four different mechanisms: Quality of Service (QoS), source diversity, path diversity, and path resiliency.

Quality of Service

IP routers gradually evolved from products with limited performance and software-based processing into extremely high performance devices that can do wire-speed, low-jitter, and low-latency traffic forwarding. Furthermore, the DiffServ working group in the IETF developed new recommendations for ensuring Quality of Service (QoS) in IP networks. Leading router vendors incorporated the recommendations and developed QoS functionality in their products. These QoS features can be used to ensure that video receives highest priority in a converged network.

The DiffServ operational model requires four basic functions to be provided by routers in the network. They first must perform packet classification when a packet enters the network, and mark the packet with the appropriate DiffServ codepoint. The packet then enters an input Q, and is subsequently policed based on the committed access rate for the service. Finally, those packets that have not been dropped by the policing engine are sent to an output queue, with one queue for each class of traffic. Video traffic will generally be given the highest priority, thus minimizing the likelihood that the packet will be dropped.

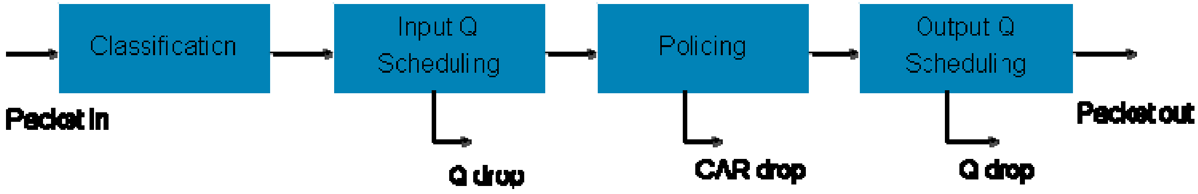

Figure 4

Source Diversity

An IP network can easily be engineered for source diversity. This can be done in a number of different ways. First, with multicast-based services that use source specific multicast, the receiving devices can be configured to request an alternate source in the event of a failure.

Path Diversity

Engineering an IP network to provide path diversity means configuring it to assure that primary and backup video streams are sent over different routes—thereby traversing different links and different nodes. As a result, if any link or router fails, the secondary video will not be impacted. Many networks are designed as rings, and can easily be engineered for path diversity by setting OSPF weights to ensure that the primary and the secondary video traffic take paths in opposite directions around the ring.

Path Resiliency

Path resiliency makes use of a variety of mechanisms to ensure that a path recovers quickly following any type of failure or disruption. Path resiliency can be accomplished in a variety of ways:

- All well-designed routers have redundant control modules, so that a secondary control module can take over in the event of any failure. They can also be configured with high availability features, whereby the routing tables and learning tables of redundant control modules are kept synchronized so that failover is near-hitless if a failover does occur.

- “Fast convergence” of routing protocols can be built into the network through a number of features and configuration options, including:

    - Bidirectional Forwarding Detection (BFD) which detects partial link failures where a device can transmit, but not receive on a given port.

    - “Fast Hellos” which can be configured in OSPF. By default, Hellos are sent only once every 10 seconds, meaning it may take 10 seconds for a failure to be detected. With “Fast Hellos,” failures can be detected in less than one second.

What Does a Typical Cable IP Network Look Like?

Before IP, video streams were generally processed in the headend, sent over Asynchronous Serial Interface (ASI) or Digital Head End interface (DHEI) ports in native MPEG format, and then sent either directly to QAMs over Asynchronous Transfer Mode (ATM), Synchronous Optical Networking (SONET), or proprietary transport network to their destinations. Encoders, groomers, and receivers, VoD pumps, and other video devices used these same standard ASI and DHEI interfaces. These various transports were all point-to-point technologies and point-to-multipoint traffic (e.g., broadcast). They inefficiently used network capacity since multiple copies of the video traffic were transmitted.

Within the last five years, the network evolved to IP-based technology. Devices in the video headend typically now come with one or more Ethernet interfaces—be they Fast Ethernet, Gigabit Ethernet, or 10 Gigabit Ethernet. MPEG packets are encapsulated in IP and sent (over Ethernet) to a configurable IP destination address. These IP-encapsulated video packets are then sent across the IP network and ultimately terminate on a QAM modulator that de-encapsulates the video, modulates the video stream, upconverts it to the appropriate RF frequency, and sends it out into the HFC network.

The flexibility and scalability of IP means that there are many different ways to build the network. One common trend that can be observed, however, is that IP technology is being used to consolidate headends in order to reduce costs. Rather than having many headends with associated satellite receivers, a much smaller number of headends can be used to distribute content terrestrially (over fiber connections) using IP over dark fiber of Dense Wavelength Division Multiplexing (DWDM).

Figure 5 shows an example of how such a network might be set up. Here there are two "national" headends distributing common video feeds, while a regional IP network distributes video between systems in a region. Within a regional network, headends can back each other up (while taking advantage of the relatively inexpensive bandwidth in a region – as compared to more expensive long-haul capacity). Video feeds are distributed over the IP network. Capacity of links in such a system can vary significantly depending on the size of the cable operator, but most use either Gigabit Ethernet or 10 Gigabit Ethernet links. The largest operators might use as much as 40 Gigabits of capacity between locations.

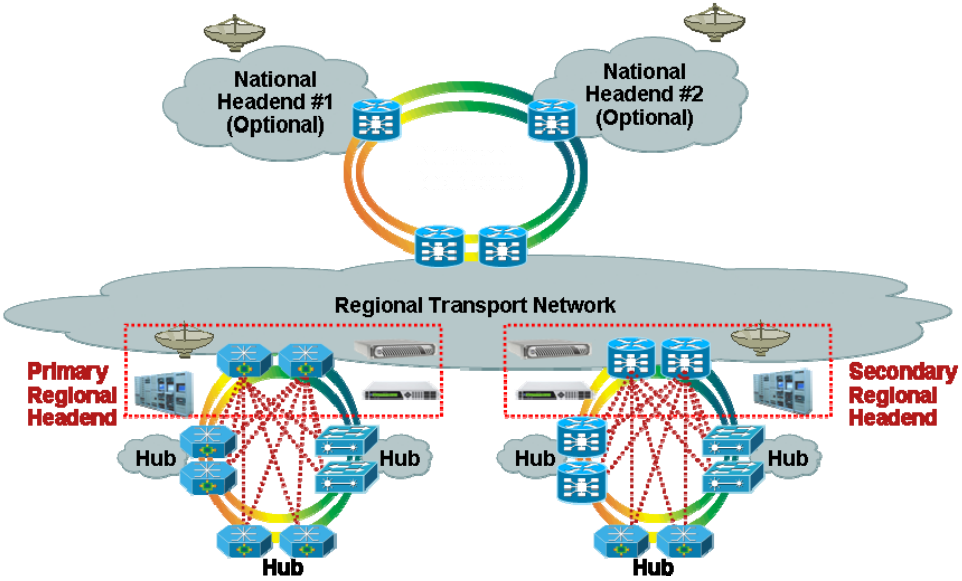
National Headend #1 (Optional)
National Headend #2 (Optional)
Regional Transport Network
Primary Regional Headend
Secondary Regional Headend
Hub
Hub
Hub
Hub
Hub
Hub

Figure 5

Systems are often set up using a ring-based topology, though this varies depending on fiber availability and other factors. Ring topologies provide inherent path diversity since video traffic from the headend can be sent on both sides of the ring.

See Figure 6 below:

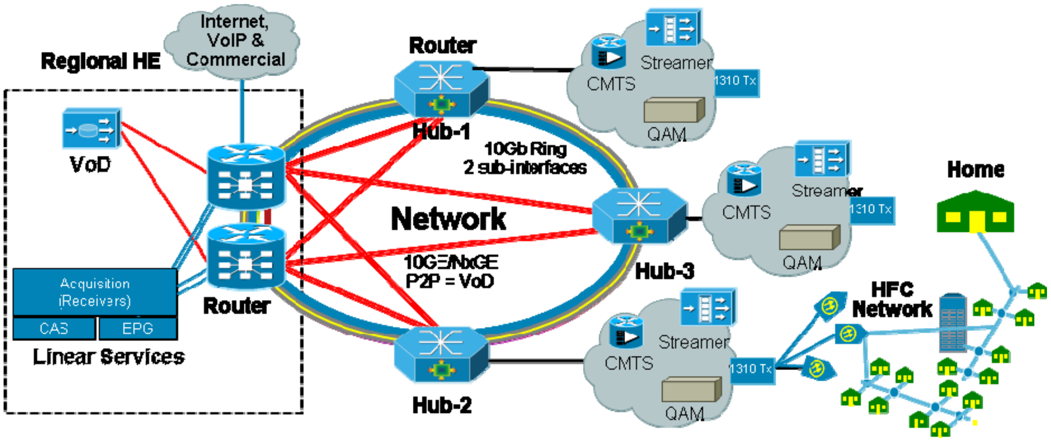

Figure 6

The ring is very efficient for multicast distribution of services such as linear television (both broadcast and switched digital video). Video feeds are dropped off at each hub site and then carried on to other destinations on the ring. The IP destination address for SPTSs and MPTSs originating out of the headend will generally be a multicast group address. Depending on architecture and MPTS/SPTS content, the receiving devices will generally be either QAMs or ad splicers—either of which will use IGMP to “join” the multicast group. PIM-SM configured on the routers will ensure that the video is delivered to appropriate destinations.

In the case where multicast video is delivered to distributed ad splicers, these ad splicers will perform ad insertion and then send on the processed traffic to a *different* multicast group, which the downstream QAM will subsequently join. For example, a headend might deliver an MPTS on IP address 225.47.6.18, an address joined by a hub site ad splicer. After performing ad splicing, the ad splicer will then become the *multicast source* for a new multicast group, such as 226.88.3.12. Downstream QAMs will then join 226.88.3.12.

Video on Demand is relatively simpler from an IP perspective, although it requires significant IP bandwidth capacity. VoD feeds in the network require significant capacity and are less efficiently distributed over a ring-based architecture. As a result, VoD is typically carried on point-to-point connections from the headend to each hub site. These point-to-point links can be either physical links (i.e., an actual fiber direct from headend to hub) or a logical link (i.e., a DWDM wavelength that is physically sent around the ring, but is provisioned as a dedicated link to the specified hub site). The “IP destination” of VoD traffic will be an edge QAM in the hub site to which a subscriber is connected. Based on IP address and the UDP “port number” of the packet, the QAM will map the VoD to an appropriate QAM port (and RF frequency).

Traditional Cable Video vs. IPTV

Many believe that cable will eventually evolve to deliver IP-encapsulated video all the way to the set-top box – as is done today in Telco-based IP television (IPTV) delivery. From a *network* perspective, there is a relatively small difference between traditional MPEG video delivery and cable IPTV. The only difference is that IP-encapsulated video is delivered to a cable modem termination system (CMTS) which delivers DOCSIS video through a QAM port, whereas traditional MPEG video is delivered through the QAM without traversing a CMTS. Otherwise, IP video distribution through the transport network and the headend processing is identical.

The real challenge in delivering IPTV over cable is in the set-top boxes. All cable set-top boxes that have been deployed to date are capable of receiving and decoding only traditional MPEG video. A new generation of "hybrid" set-top boxes now rolling out of factories are capable of receiving and decoding either traditional MPEG video or IP encapsulated video. These same set-top boxes also enable either H.264/MPEG-4 or MPEG-2 video decoding. H.264/MPEG-4 is about twice as efficient from a bandwidth perspective as MPEG-2.

With this new generation of set-top boxes, cable operators can gradually introduce new IP video services without any disruption to their existing IP services. For example, they might enable set-top users to access user generated content made available by PC-based users – without forcing the cable operator to transcode the videos and import them into a Video on Demand (VoD) system prior to use. These set-top boxes also enable the migration of existing video services to IP-based video services. While there is significant debate about what added value there is in delivering today's linear video over IP, these hybrid set-top boxes enable the capability if and when the cable operator is ready to make the switch. Figure 7 shows how hybrid set-top boxes can receive either DOCSIS or traditional MPEG services in a cable network.

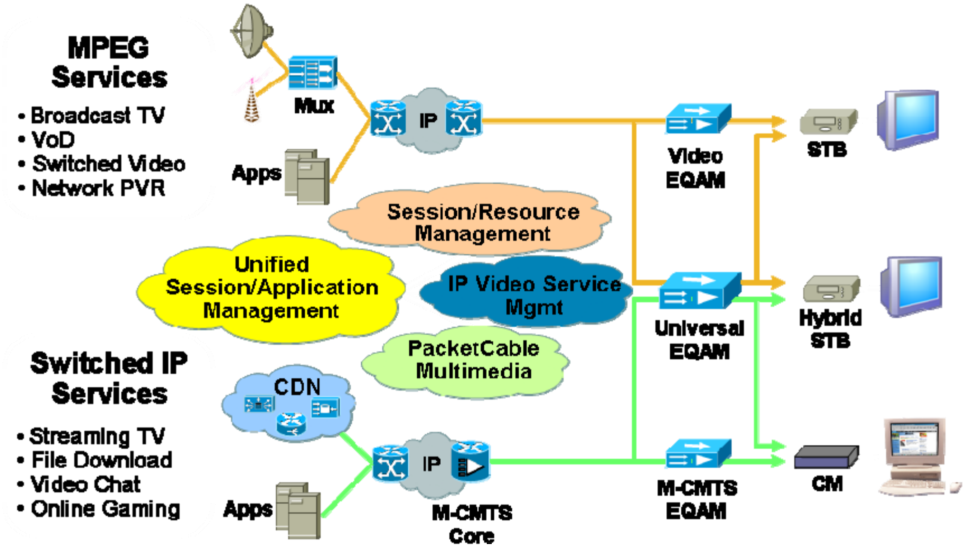

Figure 7

Full Migration to Cable IPTV?

There are varying views within the cable industry on when (or if) there will be a full migration from the traditional video services of today to an IPTV delivery model – where all services are coming through the DOCSIS path.

One perspective is that existing services will continue to be delivered without IP encapsulation, and only new services will be delivered over IP. The belief of those holding this perspective is that there is nothing to be gained by moving these services to IP – so why do it?

Another perspective is that all services will eventually move to IP. Those holding this view believe that the simplicity of delivering all IP services over a “big pipe” will provide operational simplicity and statistical gains (e.g., due to increased stat mux gain and variable bit rate (VBR) encoding), and that set-top boxes can be cost-reduced if they must only support IP video.

Time will tell whether all cable video is eventually delivered over IP through the DOCSIS network. Regardless, nearly all observers agree that cable operators will need to add IP video delivery capability to their networks.

Summary and Conclusion

Cable operators have evolved from relatively simply operations delivering analog video services to sophisticated multimedia and entertainment operations delivering a full suite of voice, video, and data services. IP networks have proven to be key enablers of this transformation, and are uniquely capable of serving as a converged platform that enables all these services to be delivered over a single network – while reducing capital and operational expenditures in the process.

A number of properties of IP make it ideally suited for video delivery:

- IP supports efficient unicast and multicast delivery mechanisms – for efficient support of both VoD and linear video services.
- IP is connectionless, and utilizes dynamic routing protocols to automatically direct traffic on an optimum path.
- IP supports sophisticated QoS and other resiliency mechanisms that can ensure minimal latency, jitter, and packet loss as needed for video service delivery.
- IP is massively scalable. The Internet itself, after all, is based on IP. IP interfaces can run at 10 Gigabit Ethernet, 40 Gigabit Ethernet and even higher speeds.

As a result of these benefits, cable operators around the world have been moving more and more to IP technology. The flexible IP network will serve as the basis for new services and applications that will allow cable operators to differentiate themselves for years to come.

Abbreviations and Acronyms

App	Application
ARP	Address Resolution Protocol
ASI	Asynchronous Serial Interface
ASM	Any Source Multicast
ATM	Asynchronous Transfer Mode
BFD	Bidirectional Forwarding Detection
BGP-4	Border Gateway Protocol 4
CAS	Conditional Access System
CDN	Content Delivery Network
CM	Cable Modem
CMTS	Cable Modem Termination System
CSMA/CD	Carrier Sense Multiple Access with Collision Detection
DHEI	DigiCable Headend Expansion Interface
DOCSIS	Data over Cable Service Interface Specification
DWDM	Dense Wavelength Division Multiplexing
EPG	Electronic Program Guide
EQAM	Edge QAM
HFC	Hybrid Fiber Coaxial
IETF	Internet Engineering Task Force
IGMPv2	Internet Group Management Protocol version 2
IP	Internet Protocol
IPTV	Internet Protocol Television
LSR	Label Switch Router
MAC	Media Access Control
M-CMTS	Modular CMTS
MPEG	Moving Picture Experts Group
MPLS	Multi Protocol Label Switching
MPTS	Multi Program Transport Stream
OSPF	Open Shortest Path First protocol
P2P	Point to Point
PE	Provider Edge
PIM-DM	Protocol Independent Multicast – Dense Mode
PIM-SM	Protocol Independent Multicast – Sparse Mode
PVR	Personal Video Recorder
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RTP	Real-time Transport Protocol
SDV	Switched Digital Video
SONET	Synchronous Optical Networking
SPTS	Single Program Transport Stream
SSM	Source Specific Multicast
STB	Set-Top Box
TCP	Transmission Control Protocol
TTL	Time to Live
UDP	User Datagram Protocol
VBR	Variable Bit Rate
VoD	Video on Demand
VoIP	Voice over Internet Protocol
VPLS	Virtual Private LAN Services
VPN	Virtual Private Network