



IP Networks for Broadcaster Applications

Yves Hertoghs (yhertogh@cisco.com), Distinguished Systems Engineer, Cisco
Thomas Kernen (thkernen@cisco.com), Consulting Engineer, Cisco
Steve Simlo (ssimlo@cisco.com), Consulting Engineer, Cisco

Table of Contents

Introduction.....	3
Overview of IP Architectures in Broadcast Environments.....	4
An Introduction to Internet Protocol.....	6
The Internet Protocol (IP).....	6
The Role of Multi-Protocol Label Switching (MPLS).....	7
The Role of Ethernet.....	8
Unicast and Multicast IP Forwarding	8
Unicast Routing	8
Multicast.....	9
Achieving Quality of Service and Resilience in IP and MPLS Networks.....	11
Comparing QoS and Resiliency in Packet-Based and Circuit-Switched Networks	11
Achieving Quality of Service Through the IP Differentiated Services Model.....	12
Connection Admission Control.....	13
Comparing IP and MPLS	14
Transporting Contribution and Distribution Video Services over IP.....	14
Video Compression.....	15
Transport and Compression Schemes in IP Video Networks	15
Uncompressed Video Services.....	16
Frame-by-Frame Compression	16
Group-of-Pictures Compression	16
IP Video Adaptation Requirements.....	17
Delay	17
Jitter and Wander.....	17
Clock Synchronization.....	17
Impact of Loss on Different Video Types.....	18
Scheduling Applications.....	19
Convergence Mechanisms for Transporting Video over IP.....	19
IP Convergence in WDM Networks.....	20
Bidirectional Forwarding Detection	20
Routing Protocol Enhancements.....	21
Traffic Engineering.....	21
Multicast-only Fast Re-Route (MoFRR).....	21
Choosing the Right Convergence Technique.....	22
Anycast Source Redundancy.....	23
Packet Retransmission	24
Conclusion.....	24

Table of Figures

Figure 1. Macro view of Broadcaster's Production and Delivery Process.....	4
Figure 2. Video Services Lifecycle.....	5
Figure 3. Adapting Digital Video onto IP	5
Figure 4. Unicast Routing.....	8
Figure 5. Multicast Routing.....	9
Figure 6. Any-Source Multicast	10
Figure 7. Source Specific Multicast.....	10
Figure 8. End-to-end delay.....	17
Figure 9. MPEG-2 video GOP based compression with a slice error due to packet loss "Source material copyright SMPTE, used with permission"	18
Figure 10. IPoDWDM.....	20
Figure 11. Multicast-Only Fast Re-Route	22
Figure 12. Spatial Redundancy.....	23

Introduction

The future of communications is here, and its name is Internet Protocol (IP). Originally regarded as an IT-only transport technology suitable for data and email traffic, IP has quickly become the dominant standard for all types of communications. This change is largely due to the inherent flexibility of IP transport, its cost efficiencies, and the ubiquitous availability of IP networks. Despite these advantages, however, until recently broadcasters have not considered IP ready to support “mission critical” real-time video services. While IP networks have played a role in contribution and production processes, they typically were reserved for non-real-time applications. Today, IP network technology has evolved, and concerns about its ability to support the stringent quality and resiliency demands of real-time video have been addressed. As a result, IP is emerging as an increasingly important technology for broadcasters and service providers, and IP-based transport networks and medianets are now used by broadcasters around the globe.

The advantages of IP extend beyond operational expense (OPEX) and capital expense (CAPEX) cost reductions. Once broadcast services can be managed within the IP domain, broadcasters have the opportunity to transform production, post-production, contribution, and distribution of core video and audio assets. The ability to share video assets quickly and efficiently on a shared IP network infrastructure can unleash unprecedented collaboration, efficiency, and agility throughout the entire broadcast value chain. (Figure 1.) This includes:

- **Production:** Many broadcasters still rely on production systems that are managed as independent applications, supported by dedicated infrastructures and physical tapes. The result is a production workflow that is fragmented and fraught with delays and duplicated efforts. An IP environment supports an end-to-end digital workflow that dynamically moves media through the production process, breaks down operational silos, and supports company-wide collaboration. As a result, digital workflows can reduce OPEX, allow editing functions to be easily shared among different teams, and significantly reduce “time to air” – especially important for news applications.
- **Contribution:** The same innovative approaches that are transforming media production can also be applied to the delivery of video between studio locations and among broadcast partners. Highly flexible and cost-effective IP networks let broadcasters reduce OPEX and rapidly introduce new services, such as high-definition (HD) video.
- **Distribution:** Distribution of national and local Digital Video Broadcast – Terrestrial (DVB-T) services to transmitter sites can also benefit from the CAPEX and OPEX saving of an IP-based network. And, once DVB-T services are managed within the IP domain, they can easily be delivered over fiber, copper, or microwave networks, and within systems encompassing all three.
- **Consumption:** Broadcasters need to deliver TV services to consumers over multiple platforms and multiple screens (TV, PC, and mobile device), both in the home and on the go. They need solutions that can accommodate diverse video formats, quality levels, and compression standards, and deliver the highest quality for the lowest cost. IP provides a common framework for easily adapting and distributing TV services for any platform or device.

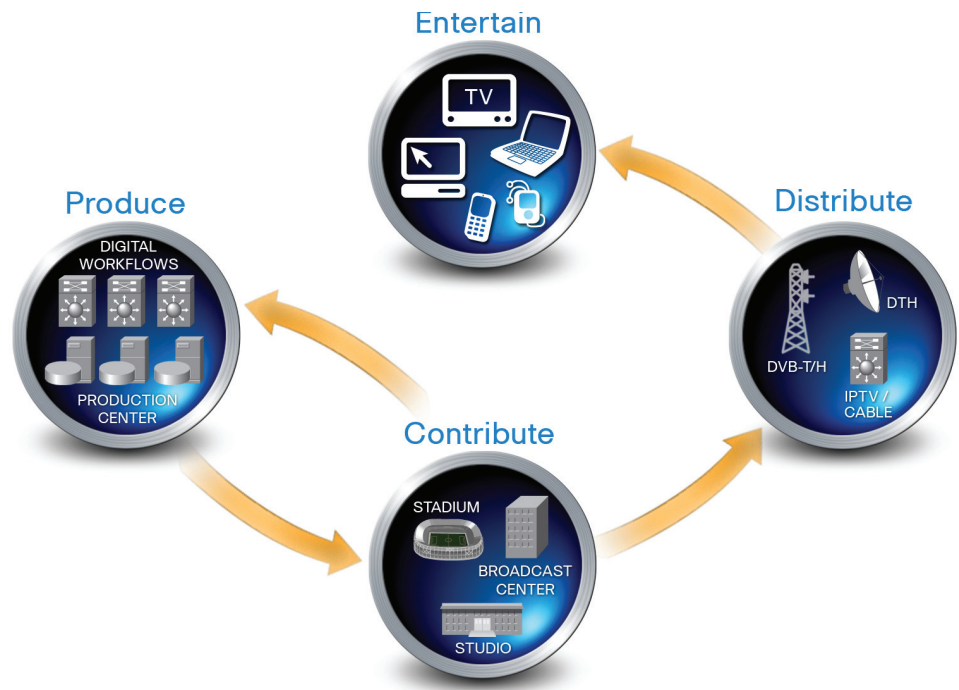


Figure 1. Macro View of Broadcast Production and Delivery Processes

All of these extraordinary capabilities are supported by the unique advantages of IP-based networks in broadcast environments. Unlike any other network type available to broadcasters today, IP networks provide:

- An open, standards-based, widely adopted transport solution, providing reassurance for future longevity as well as competitive pricing
- Exceptional flexibility, with near-infinite bit-rate granularity and easily adaptable routing capabilities
- Substantial OPEX savings through the convergence of multiple services onto a common infrastructure, with these benefits multiplying as more services are migrated to IP

This paper outlines how IP Networks can provide a viable transport solution for broadcasters. It provides an in-depth discussion of IP transport technologies, including the role of IP, Ethernet, Multi Protocol Label Switching (MPLS) and how they compare to legacy transport protocols such as Synchronous Digital Hierarchy (SDH) and Asynchronous Transfer Mode (ATM) for video transport. The paper describes the quality and resiliency techniques that allow modern IP networks to support demanding real-time video services and discusses IP video compression and adaptation mechanisms. Finally, it provides an overview of techniques broadcasters can employ to ensure maximum availability and reliability in IP-enabled broadcast networks.

Overview of IP Architectures in Broadcast Environments

The creation of content and its distribution is a multi-stage process that involves a broad range of stakeholders, skill sets, and technologies. Video services follow a lifecycle from initial acquisition, through the production and packaging of the content, to final playout to the distribution network that delivers the content to viewers. (Figure 2.) Each stage in this lifecycle has its own requirements and challenges. This paper focuses on the contribution and primary distribution stages.

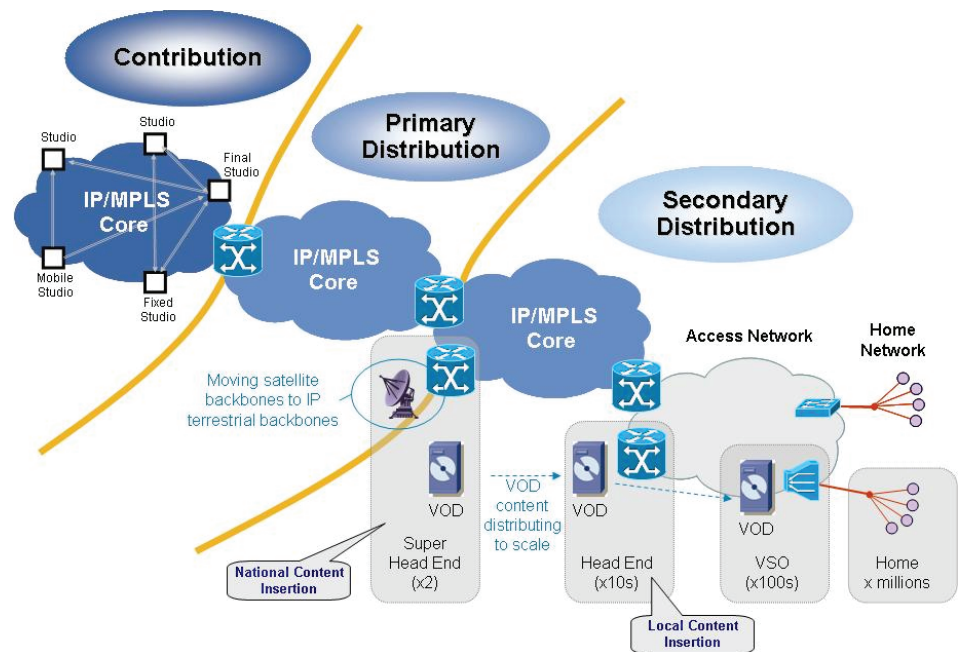


Figure 2. Video Services Lifecycle

The first stage in the lifecycle is the acquisition of the video content into the IP domain. Adapting digital video onto an IP network is achieved using either cameras with a built in Ethernet/IP network interface card, or via a standalone IP video adaptor (sometimes referred to as a "IP video gateway" or "IP video encoder") as shown in Figure 3.

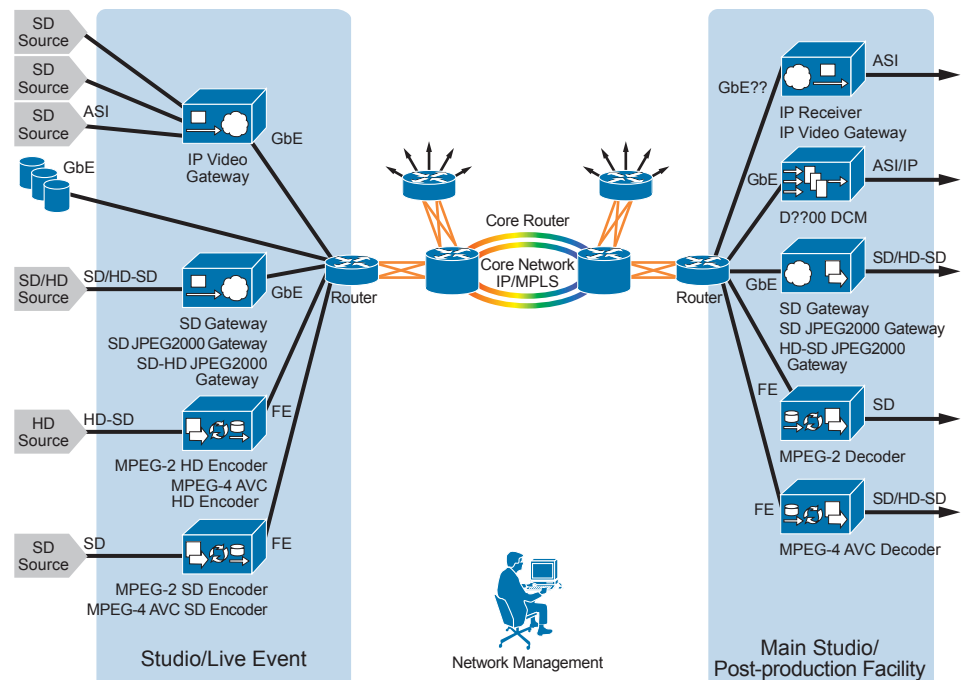


Figure 3. Adapting Digital Video onto IP

An Introduction to Internet Protocol

This section outlines the Internet Protocol suite. It discusses unicast and multicast packet forwarding, as well as techniques for achieving Quality of Service (QoS) and high availability in an IP Network. This section also explains the role of Ethernet and of MPLS in IP networks.

The Internet Protocol (IP)

Originally, IP was designed for communication across the Internet. In recent years, however, it has become the de facto communication protocol for all types of traffic in private and public networks. In today's enterprises, nearly all communication is IP based, allowing enterprise networks to support data, voice, video, storage, and other services on a common, standards-based infrastructure. Service Providers have also adopted the Internet Protocol suite almost universally, allowing them to converge their various services across a common IP-based backbone. Services such as Internet access, voice (both private branch exchange [PBX] interconnects and Public Switched Telephone Network [PSTN] services), business interconnect services (typically via virtual private networks [VPNs]) and increasingly, video, are now delivered over IP networks. For all organizations relying on IP, the common driver is the flexibility and cost savings afforded by converging services across a common, cost-efficient, standards-based infrastructure.

IP is also becoming the preferred protocol for delivering broadcast video services. Broadcasters are using IP transport not only in secondary distribution networks (i.e. IP television [IPTV] over residential broadband systems), but also increasingly for Primary Distribution and Contribution networks. While some broadcasters previously questioned whether IP could support video services, the latest achievements in quality of service, resilience, fast repair, switching speeds, and scalability have made IP networks reliable enough to become a viable option for video contribution networks. Consequently, broadcasters can now converge services and technologies over a common IP infrastructure, and enjoy the same OPEX and CAPEX advantages that enterprises and service providers have enjoyed for many years.

The chief characteristic of IP that distinguishes it from traditional technologies such as ATM and SDH is that it is packet-based. With traditional "connection-oriented" technologies, a path must be set up across the network from origin to destination before any traffic can be sent. IP offers a fundamentally different paradigm, in which the network itself determines the optimal path for transmitting traffic to its destination at any given moment, and routes traffic dynamically. In the IP model, no transmission path is set up to the destination in advance. Instead, an end station wraps data inside a packet "container," stamps a destination (and origin) address on it, and sends it into the network. The network then uses the IP addresses to transport the packet to its destination through "connection-less" packet forwarding or "IP routing." The nodes forwarding these IP packets (routers) constantly update each other about the reachability of IP addresses and/or networks through the use of IP routing protocols. Today's IP routing protocols allow every router in the network to individually build a full topology view of the IP network.

The connection-less approach of IP networks offers several advantages. First, since no paths must be established in advance, provisioning is easier and more cost-efficient. IP networks are also inherently resilient: since no paths are pre-established, an IP network will always reroute around any link or router failure (assuming the network has been designed with resilient nodes and links). This allows IP networks to survive multiple link and node failures – something not always possible with path-protected networking technologies such as SDH.

The Role of Multi-Protocol Label Switching (MPLS)

Multi-Protocol Label Switching is a technology that builds on “Layer 3” or routing-layer IP capabilities to simplify and improve the exchange of IP packets. In MPLS networks, MPLS-enabled routers use IP routing protocols to exchange information with each other. However, the information exchanged goes beyond the reachability of IP routes to include “Layer 2” information about network links, such as bandwidth, latency, and utilization. Routers at the edge of an MPLS network encapsulate packets with MPLS headers containing one or more “label stack” entries. These label stack entries contain a 20-bit value (a label), that can be used to forward packets. (Functionally, this label replaces the IP address, which is now “hidden” within the MPLS packet.) The label points to the next hop MPLS router. By stacking MPLS labels, network engineers can create hierarchies inside the network, since intermediate MPLS routers will only act upon the top or outer label. Labels further down the stack provide information for “applications” at the edge of the network, such as an IP VPN identifier, Layer 2 tunnel ID, and more. Note that the outer MPLS label is only specific to the link. The MPLS network swaps this outer label on a node-by-node basis (analogous to Data Link Connection Identifiers [DLCIs] in Frame Relay networks or Virtual Path or Virtual Circuit Identifiers [VPIs/VCI] in ATM networks).

Traditional IP routers examine the IP headers and make individual forwarding decisions on a hop-by-hop basis. This is essentially the way connection-less networks work. MPLS routers perform the IP lookup only once when the packet enters the network. At that point, the MPLS router replaces the routing information with a label, and downstream MPLS nodes make forwarding decisions based only on this label, effectively creating a more “connection-oriented” approach. This approach offers some advantages over traditional IP routing.

MPLS allows the router performing the MPLS encapsulation to assign a label based on more than the destination IP address of the packet (e.g. traffic class, ingress interface). This allows for the creation of different paths across the MPLS network, even if the ultimate IP destination is the same. The router performing the MPLS encapsulation can assign a label based on its own identity, so the receiving router can then infer from which router this packet came. This is impossible with traditional IP routing.

MPLS also allows engineers to force a packet to follow a given route across the network without having to encode the desired path inside the packet. The MPLS nodes merely forward based on the labels, but the labels can be installed for a pre-computed explicit path. This path can also be installed with a certain amount of bandwidth assigned. Using this technique, traffic engineering capabilities can be applied to networks running IP protocols, making them more familiar to network administrators used to path-based, connection-oriented networks. For example, the IP protocol Resource Reservation Protocol – Traffic Engineering (RSVP-TE) allows bandwidth reservations to be made across an MPLS path.

MPLS networks also allow for extra labels to be pre-established at every MPLS node to provide a pre-established backup path for switching packets in the event of a local link failure. This backup path is not end-to-end, but merges with the primary path at downstream nodes. Since the trigger to switch to the backup path is a local link failure (and does not rely on end-to-end signaling), MPLS networks can achieve switching times of 50 milliseconds. This is often referred to as MPLS Fast Re-Route (MPLS-FRR). MPLS-FRR can be applied to point-to-point label-switched paths or point-to-multipoint label-switched paths (referred to as P2MP MPLS-TE).

Note that MPLS can use exactly the same per-hop QoS model as IP networks, as explained below. However, MPLS allows network engineers to employ per-path bandwidth reservations for certain applications, if desired.

The Role of Ethernet

Service providers worldwide are increasingly using Ethernet (often referred to as Carrier Ethernet) in Wide Area Networks (WANs) and IP/MPLS backbones to improve cost-effectiveness. In fact, Ethernet is now often used to interconnect the IP routers that make up the backbones of the largest global networks, allowing speeds of tens of Gigabits per second (Gbps). By using the same common Ethernet technology in network backbones that is used inside enterprise networks, service providers have dramatically cut the cost of delivering LAN-to-WAN interconnects and are benefiting from the economies of scale of Ethernet technologies. Broadcasters can take advantage of Ethernet in IP-based contribution and distribution networks to realize the same advantages.

Using Ethernet technology to interconnect IP/MPLS routers is also a relatively simple proposition. Nothing must be provisioned to make that interconnect, as Ethernet has its own addressing scheme using supplier-provided Ethernet addresses, and the IP protocol automatically discovers these addresses.

Unicast and Multicast IP Forwarding

An IP router has two fundamental models for forwarding packets, unicast and multicast.

Unicast Routing

In the unicast model (Figure 4), the router looks at the destination IP address of each packet and uses this as an index into the unicast routing table. This will point to the outgoing interface and/or next-hop IP router to which the router must send the packet.




Figure 4. Unicast Routing

Multicast

In the multicast model (Figure 5), the router forwards IP packets to multiple different destinations simultaneously. In this model, the destination address is a multicast destination group address, or a special set of defined addresses. The network understands which multicast group addresses to forward on specific interfaces, depending on either static configuration or on end-stations signaling their interest in receiving traffic. Effectively, the multicast model builds a tree-like topology across the routers from the multicast sources to requesting receivers (referred to as a multicast distribution tree). In order to avoid forwarding loops in multicast topologies with redundant links, every IP router does a route lookup. This lookup references the source IP address of the packet. If a packet arrives on an interface pointing towards the source address, the router accepts and forwards the packet. If a packet arrives on an interface that does not point towards the source IP address, the packet will be dropped. This mechanism is referred to as Reverse Path Forwarding (RPF) check.

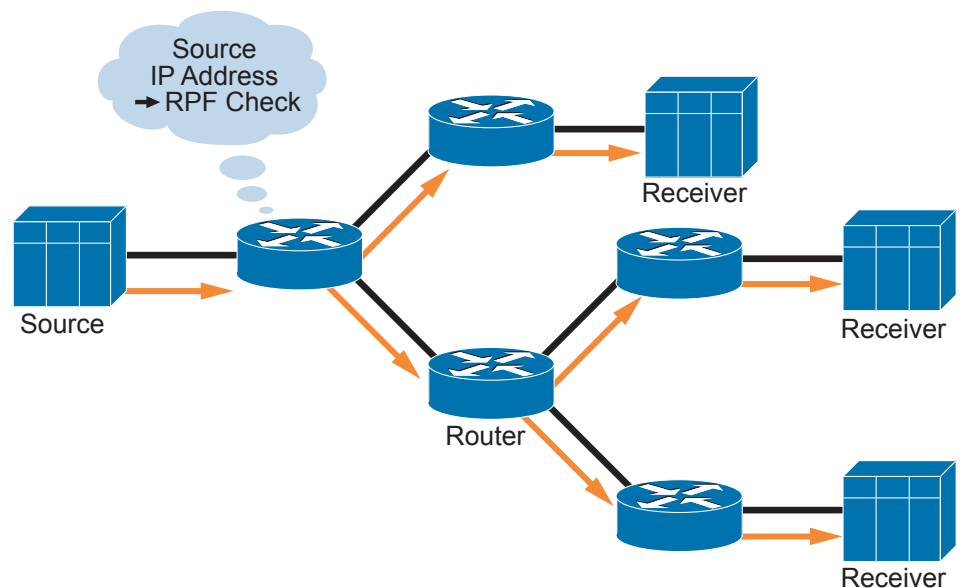


Figure 5. Multicast Routing

The most popular protocol used to build multicast distribution trees is called Protocol Independent Multicast or PIM. PIM uses the unicast routing table independent of how that table was built (hence, the reference to protocol independence). Network engineers can employ two types of PIM: Any-Source Multicast (ASM) and Source Specific Multicast (SSM).

In ASM (Figure 6), the routers establish multicast trees according to destination, independent of the source(s) of the multicast flows. ASM uses the concept of a “shared tree,” i.e. a multicast tree that has a known root (known as the rendezvous point) in order to forward multicast streams without regard for the source address. Each router in the network that wants to receive multicast traffic for a certain group becomes part of the shared tree rooted at the rendezvous point. In this model, the rendezvous point is (initially) the only router with knowledge of individual sources and will also build trees towards these sources when required.




Figure 6. Any-Source Multicast

In SSM (Figure 7), the routers build multicast trees and forward packets based on both the unicast source and the multicast destination. SSM has the advantage of better access control, since it does not forward two separate source multicast streams via a common shared tree, preventing traffic collisions and providing better security. This model also simplifies multicast operations, since SSM does not need a shared tree and a rendezvous point. SSM is very well suited to secondary distribution video services, since this application entails the distribution of video from a few sources to many receivers.




Figure 7. Source-Specific Multicast

Achieving Quality of Service and Resilience in IP and MPLS Networks

Real-time audio and video services are extremely sensitive to packet loss and delay. As a result, any IP infrastructure operating in a broadcast environment must meet stringent performance and availability requirements. It must provide:

- Extremely low jitter, or variation in the timing between the arrival of packets or signal pulses (stipulated by the European Broadcasting Union, for example, as less than 10 milliseconds)
- Very low delay (typically less than 80 milliseconds)
- Extremely low (ideally zero) packet loss, since even a single dropped packet can have a major effect on video quality

As discussed, IP and MPLS networks are inherently resilient, and the connection-less nature of IP means that traffic will continue flowing in the event of a link or node failure. However, IP and MPLS networks do not by default retransmit packets that may have been lost during network reconvergence. To accomplish this, network engineers can use higher-layer protocols to signal applications to retransmit certain lost packets, if desired. However, retransmission often has the disadvantage of delaying or slowing down the application, rendering it unacceptable for real-time video delivery. Fortunately, there are techniques that network engineers can employ in IP and MPLS networks to address packet loss during reconvergence more effectively.

One approach is to configure the application to add packets to the stream so that it contains enough information to reassemble the stream even if some packets are lost. This is known as Application Layer Forward Error Correction (AL-FEC¹). Another approach is sending the stream twice (either across different links and nodes or at different timeslots). This is referred to as “Live-Live” delivery. Note that AL-FEC and Live-Live techniques are not exclusive to IP and MPLS networks. These techniques can also apply to any transport technology, since all technologies take time to reconverge after a link and/or node failure. IP does offer the advantage of rerouting around individual link and/or node failures dynamically, however, whereas SDH networks require an extra end-to-end protection path for every configured path in the system.

Comparing QoS and Resiliency in Packet-Based and Circuit-Switched Networks

Modern IP networks are extremely responsive to link or node failure. IP networks react to a failure by sending out updates in all directions, causing each router to recalculate its own view of the new topology. Several years ago, failure detection was often slow (of the order of seconds or even minutes), as it relied on a router noticing that neighbors had “gone away.” Today, most router topologies are based on point-to-point links (often using Ethernet), so there is no longer a need to rely on a router detecting the loss of a neighbor. Instead, IP routers usually notice local link failures almost immediately. If a link spans an optical Wavelength Division Multiplexing (WDM) infrastructure, modern routers have integrated lower-level WDM signaling (known as ITU-T G.709). This allows a router to recognize degraded links as well as totally failed links and adjust its IP forwarding accordingly. In addition, numerous other improvements have been made at the IP routing protocol level in recent years which make today’s networks converge within a few hundred milliseconds for both unicast and multicast services. These improvements are known as IP Fast Convergence. Together, these mechanisms allow modern IP networks to meet the same stringent resiliency requirements as circuit-switched SDH and ATM systems.

¹ Application Layer FEC does not replace physical layer FEC schemes, such as those used in Digital Subscriber Line (DSL) or optical fiber transmission systems.

Where IP transport differs from SDH or ATM is in the way it handles reliable quality of service. With SDH or ATM, circuits are set up with a specific “end-to-end” bandwidth. If there is no traffic on these circuits, the associated bandwidth on those links is unused. Therefore, these networks are restricted to signaling (or provisioning) only as many circuits as the links can forward, and they eliminate packet-level congestion. However, in systems that require redundancy (such as real-time video networks), the number of circuits that must be pre-configured can double, as that bandwidth must be reserved and cannot be used for other applications. As a result, this model is extremely inefficient in terms of utilization of available bandwidth.

IP networks operate within a very different paradigm. With IP, there are no circuits. IP routers statistically multiplex different traffic flows onto links without first checking whether this will congest the interface. For traditional IP applications, this congestion is not an issue. Applications like web browsing, for example, handle momentary congestion quite well by making use of the Transport Control Protocol (TCP). TCP “slows down” traffic flows in reaction to congestion and signals the application to resend any lost segments using a system of segment numbers and “windowing.” (The TCP “window size” is a value indicating how much data can be sent without requiring an acknowledgment that the data has been successfully received.) If the network has more capacity at a given moment in time (a common occurrence given the very “bursty” nature of Internet data traffic), TCP senses this and speeds up transmission by increasing the window size. If the amount of TCP flows on a single link would lead to congestion due to packet buffer overruns, packet drops will alert TCP to shrink its TCP window size, automatically lowering the rates of the individual flows on that link. In most cases, buffer overruns should be avoided, as too many TCP packets get dropped. Network engineers typically employ Congestion Avoidance mechanisms such as Random Early Detection (RED) to accomplish this. RED randomly drops single packets from TCP flows, with the probability of dropping increasing depending on buffer utilization and (if desired) the individual flow rate. This avoids scenarios in which IP packet discards create simultaneous congestion conditions on multiple parallel TCP flows.

Applications like voice and video that are highly sensitive to network congestion do not use TCP, instead employing the much simpler User Datagram Protocol (UDP) to carry packets. UDP has no segment numbers or windowing mechanism, so it cannot react to packet loss. However, in these types of applications, it is better to drop packets than introduce delay by waiting for a retransmission. In cases where retransmission is possible (e.g. if the receiving end can buffer packets for a couple seconds), mechanisms like Real-time Transport Protocol (RTP) offer sequencing and retransmission capabilities for UDP-based transport. RTP is often used across “lossy” media (such as DSL access networks). RTP can also be used to synchronize two redundant video streams and monitor the IP transport without having to look into the IP payload where the video signal is located.

Achieving Quality of Service Through the IP Differentiated Services Model

An IP/MPLS network supports the concept of “Per-Hop Behaviors” (PHBs), which allow network engineers to classify incoming traffic into traffic classes. IP networks can schedule packets out of an outgoing interface in accordance with the PHB indicated by the traffic class. This behavior is referred to as the Differentiated Services Model (or DiffServ for short). The advantage of using PHBs is that, in the absence of a certain high-priority traffic class, other traffic classes can re-use the configured bandwidth. This is fundamentally different (and inherently more efficient) than circuit-switched architectures, in which bandwidth must be “nailed up” across both the active and backup paths. Note that DiffServ scheduling of IP packets through routers and on network links does not introduce a significant amount of delay. Today’s implementations achieve end-to-end jitter (or variations in delay) of less than 1 millisecond.

One commonly employed PHB is “Expedited Forwarding” (EF), which schedules traffic to be forwarded out of an interface as soon as it arrives at the packet scheduler for that interface. This is a good PHB for traffic that is delay-sensitive, such as voice or video. It also prevents congestion for that specific traffic class from occurring, as packets will always be scheduled first. Typically, traffic matching classes conforming to this PHB must be controlled to ensure that this PHB does not “starve out” other traffic classes.

Another common PHB is “Assured Forwarding” (AF), which defines a guaranteed minimum bandwidth (often expressed as a percentage of the total link bandwidth) for traffic assigned to this PHB. If a router forwards traffic conforming to the traffic class associated with this PHB, that traffic can use at least the configured bandwidth value (and may also burst up to line-rate if extra capacity is available). If network engineers can control traffic in the AF PHB such that it never exceeds the configured minimum bandwidth across the network, congestion for that class will never occur, and no packet loss in that traffic class will occur even in the case of interface congestion.

To implement “Best Effort” services, network engineers can configure an AF PHB with no minimum bandwidth guarantee. (The traffic class can still burst up to the available bandwidth of the link, minus any concurrent AF and EF traffic.) Typically, no traffic control is employed for this traffic class, as it will use whatever bandwidth is available. Therefore, applications that can handle packet loss quite well are normally assigned to this class.

The advantage of using PHBs in an IP network (as opposed to using a circuit-switched architecture) is that in the absence of certain traffic, other traffic classes can re-use the configured bandwidth. Naturally, this allows for much more efficient utilization of available bandwidth. This becomes particularly important when the network is used for a mix of different services, such as concurrent voice, video, and data.

The following section explains how network engineers can control traffic classes associated with EF and AF PHBs so that they are never congested. Using these techniques, broadcasters can ensure that video networks never experience traffic loss even under conditions of heavy link utilization.

Connection Admission Control

To protect against delay and packet loss, broadcasters must eliminate network congestion and tightly control the amount of traffic traversing all links in the network. Controlling how much traffic a network is forwarding at any given moment can be accomplished through a simple policing function, in which all packets that exceed a given rate are discarded at the ingress points of the IP network. For simple topologies and Voice-over-IP (VoIP) applications using the EF PHB, this mechanism works quite well. In these scenarios, the bandwidth is low, and the amount of concurrent voice calls and total voice traffic is quite predictable.

For more demanding applications such as video, network engineers can configure the application to “check” the network for the number of existing traffic flows that are sharing the same traffic class conforming to a certain PHB before setting up a connection. This technique is referred to as “Connection Admission Control” or CAC.

CAC can be performed at the application level when the application needs to model the resources used inside the network in real time. If it seems that no more resources can be used (i.e., the network has reached the maximum capacity set aside for this traffic class), the application does not even attempt to set up the connection. Alternatively, a simple scheduling application can be used to control the number of concurrent connections occurring on a link-by-link basis. This technique is often referred to as “Off-Path CAC,” as the application does not query the network upfront.

“On-Path CAC,” which queries the network before setting up a connection, offers a more accurate admission control mechanism based on the actual capacity of the link at a given time. However, On-Path CAC requires more intelligence within the network, as the application uses an IP-based protocol to query the network in real time, namely the Resource Reservation Protocol (RSVP). In the On-Path CAC model, the application routes RSVP packets across the network and checks how much bandwidth is used in a specific traffic class at a given moment in the current topology. The network replies to the application with a simple yes or no answer: “yes” if resources are available to support the connection without packet loss due to congestion, or “no” if congestion could occur.

The major advantage of On-Path CAC is that it dynamically adapts to changes in the topology. Even in the event of a link failure, the application maintains awareness of available capacity, and admits or denies connections accordingly. Note, however, that in this use case, RSVP does not really “reserve” hardware resources across a given path to the destination. Rather, it merely queries the router to check for the current utilization of the traffic classes. The full RSVP protocol suite does allow for hardware resource reservation hop by hop, however, this technique has proven to be unscalable in today’s IP routing environments. As a result, IP network engineers commonly use only the CAC capabilities of RSVP.

Comparing IP and MPLS

It can be argued that for applications that demand fast convergence, MPLS has an advantage. MPLS allows for bandwidth reservation and traffic engineering, making it potentially more attractive to network administrators used to this paradigm. However, this advantage comes at a cost: introducing traffic engineering leads to more operational overhead. The extra overhead is amplified when deploying point-to-multipoint traffic engineered connections, as each endpoint (or “leaf”) router must be provisioned individually. Today, improvements in IP convergence mechanisms have made it possible to offer connectionless operations with very fast convergence times for both unicast and multicast services. As described above, modern IP networks can meet even the most stringent quality and resiliency demands using the DiffServ model combined with ingress policing and/or CAC.

MPLS (as the name implies) is protocol independent, as it can tunnel IP, but also Ethernet, ATM, Frame Relay and many other technologies. IP tunnel technologies providing similar services have also been developed, but are not as widespread in use.

IP routing can be combined with MPLS to deliver a virtualized infrastructure, in which multiple “IP customers” share a common MPLS backbone. This is known as IP Virtual Private Networks (IP-VPNs). Today’s MPLS technologies allow IP-VPN services for both unicast and multicast applications.

MPLS allows the mixing of traffic engineered and non-traffic engineered paths, which may then be selected on a per-application or per-service basis. The following sections elaborate on the various options and protocol choices that can be used for IP-based “Broadcaster Services” and consider the pros and cons of each approach.

Transporting Contribution and Distribution Video Services over IP

Delivering video services over an IP network involves more than encapsulating uncompressed or compressed video into IP packets and transporting it. It requires an understanding of the interaction between the type of video compression in use, the transport of the packetized video services, and the IP adapter requirements.

The following sections detail the different scenarios and their respective requirements that must be considered for the successful delivery of video services. The way in which video is compressed (or not) has a direct impact on various system attributes, such as link bandwidth, end-to-end delay, jitter and wander. As a result, the type of compression employed directly affects the requirements of the video adaptors that adapt the ASI² and SDI³ video streams onto the IP infrastructure.

Video Compression

The main purpose of video compression is to overcome the bandwidth constraints of the network transport infrastructure. Compression typically involves a tradeoff between bandwidth availability, cost of transmission, and the level of quality required for the video services at each of the different stages between capturing the content and delivering it to the end user. The appropriate video compression (and the requirements of the underlying network) depend on the specific application. For example, a video feed for a live news program may demand the lowest possible latency. Typically, this means a video feed with minimal compression and an extremely high bit rate. When broadcasting sporting events (or when transporting video feeds among teams in a production facility), broadcasters typically prioritize video quality above all else, requiring very high bit rates.

Video compression works by reducing the amount of data used to describe the video frames. This can result in a reduction in visible quality compared to uncompressed streams, but the quality can be maintained at a level that is still deemed adequate for the specific application. Compression technologies can be classified within two main categories:

- Acquisition and production technologies: In these scenarios, an end user can deal with some amount of video information loss. The computer-based systems in use at earlier stages in the production and distribution chain, however, require substantially more information to allow for the creation, editing, and production of video content. Often, video streams in this part of the broadcast process (i.e., contribution) are uncompressed or lightly compressed.
- Transmission technologies: Networks that deliver video for viewing by an end user, such as secondary distribution networks, are often highly compressed. Typically, they depend on the capabilities of the human vision system (HVS) to recover from this level of compression (and the associated quality loss) and compensate for it.

Transport and Compression Schemes in IP Video Networks

There are three main schemes for transporting video services across an IP network:

- Uncompressed
- Frame-by-frame compressed
- Group-of-Pictures compressed

The following sections describe each scheme in detail.

² ASI: Asynchronous Serial Interface

³ SDI: Serial Digital Interface

Uncompressed Video Services

The ideal scenario for any broadcaster is to be able to transport all streams uncompressed whenever possible. This is because uncompressed transmission eliminates any video quality degradation or delays introduced by cascading and concatenation of compression and decompression cycles along the transmission path. Delivering uncompressed video is not always possible, however, due to the extreme bandwidth demands of video services – especially HD video.

Standard-Definition (SD) sources operate at a raw bit rate of 270 Megabits per second (Mbps)⁴, and so easily fit on a Gigabit Ethernet transport. Such a service usually contains a video source, one or multiple (up to 16) embedded audio channels, and additional ancillary data.

In an IP video network, adaptors simply encapsulate the uncompressed structured video data into IP packets. Such networks may also use error correction mechanisms at the IP layer (such as AL-FEC). Uncompressed video is often used in high-end contribution services, such as sports contribution, if the bandwidth is available.

Frame-by-Frame Compression

In today's modern network infrastructures, in which Gigabit Ethernet (GE) prevails, HD sources typically must be compressed to overcome the fact that they operate natively at 1.485 Gbps, dual 1.485 Gbps, or 2.970 Gbps⁵. In frame-by-frame compression schemes, each video frame (or field) is individually compressed and self-contained. As a result, decompressing the video stream does not require any information from previous or subsequent frames. In this compression scheme, data streams are typically encapsulated within a wrapper (such as Material Exchange Format, or MXF) in order to include all of the required information, including metadata from the ancillary data that is part of the service. The streams are transmitted over RTP⁶.

Frame-by-frame compression is commonly used in contribution networks for applications that require low delay for interactivity. For these applications, the level of compression and encoding/decoding cycles (known as "generations") must be kept to a minimum in order to reduce artifacts. Examples of frame-by-frame compression codecs include JPEG2000 and MPEG-4 AVC (H.264) when operating in I-frame-only mode.

Group-of-Pictures Compression

Group-of-Pictures (GOP) compression is based on the concept of encoding a key or "anchor" frame at the beginning of a group of pictures. All subsequent frames that are part of the GOP are then derived from that key frame (or from other frames that are part of that group).

A service containing the video, audio, and ancillary data streams is typically multiplexed within an MPEG-2 Transport Stream. The service is built up either as a single-program or multiple-program stream. With GOP compression, the quality degradation that results from any data loss depends on which set of data within the stream was lost during an outage. The data loss may have an impact for the length of the GOP (or even longer), depending on the video codec, GOP length, and other encoding settings in use.

GOP compression is commonly used in broadcast distribution networks, which typically cannot accommodate either uncompressed or frame-by-frame compressed services due to bandwidth constraints. Examples of GOP compression codecs are MPEG-2 and MPEG-4 AVC (H.264).

⁴ As defined in SMPTE 259M/ITU BT.656

⁵ As per SMPTE 292M, SMPTE 372M, and SMPTE 424M respectively

⁶ The most common codec in use today is JPEG2000, but others such as H.264/AVC-I, Dirac Pro or DNxHD (SMPTE VC-3) may also be used.

IP Video Adaptation Requirements

IP video adaptors take in ASI or SDI signals and adapt them to IP. Adaptors act as either transmitters (also known as encoders) or receivers (also known as decoders). To deliver high-quality video services, receivers must be able to compensate for any variations introduced by the network or inherited from the payload (i.e., as a result of compression or the encapsulating and de-encapsulating of the video information inside the IP packets). IP infrastructures must account for:

- Delay
- Jitter and wander
- Clock synchronization

The following sections describe each factor in detail.

Delay

The end-to-end delay in an IP video system is the sum of multiple individual elements (Figure 8):

- **The encoding delay** depends on the compression settings and the generation of AL-FEC (if applied).
- **The queuing delay** is introduced by the network components buffering the Ethernet frames to avoid packet loss and through prioritization. (As explained previously, a properly engineered and DiffServ-compliant IP network can ensure minimal delays.)
- **The serialization delay** is caused by any packet network component that is storing a frame and sending it to line. Serialization delays in modern high-speed networks are very low (i.e. on a 10-GE link, a 1500-byte frame serialization delay would be around 1-2µs per hop).
- **The transmission delay** is caused by link distance and, in the case of optical transmission, introduces approximately 1 millisecond delay per 100-150 kilometers.




Figure 8. End-to-End Delay

Note that the overall delay introduced by the IP/MPLS network, is typically very low. In short, with a properly engineered IP network, the delay budget is influenced mainly by the use of encoding/decoding, compression, and AL-FEC.

Jitter and Wander

Network-introduced jitter and wander have no direct impact on the video services transported but do need to be compensated for in the receiver buffer at the IP layer. However, properly engineered DiffServ IP networks are known to deliver jitter of less than 1 millisecond.

Clock Synchronization

In order to compensate for video jitter and wander, the receiver clock must be synchronized with the source clock. This can be accomplished using an external reference clock (or “master clock”) or by deriving the clock from the received signal. IP networks therefore require a buffer to compensate for the jitter and the use of a framestore for accurate video frame/field signal phasing.

Impact of Loss on Different Video Types

The primary concern for video services is packet loss. Loss can be attributed to four primary causes:

- Excess delay
- Congestion
- Physical errors
- Network convergence events

Loss due to excess delays introduced by the network can be prevented by a properly designed and capacity-planned DiffServ IP network. Congestion in video services can also be avoided through the use of IP DiffServ-based QoS, together with Off-Path or On-Path CAC. As today's IP transport networks typically make use of high-quality cabling (fiber) for backbone connections, physical errors are usually not a problem. This means that controlling network convergence is the chief mechanism for reducing loss in video networks.

In the event of loss due to network convergence, the impact depends on the compression scheme employed, as follows:

- **Uncompressed video:** In the case of a data loss from which the service cannot recover, the receiving IP video adapter will drop the corrupted video line and insert the missing line from the previous field or frame for the time of the network convergence event. In most cases this is imperceptible to the receiver/viewer.
- **Frame-by-frame compressed video:** In the case of data loss that cannot be recovered, the IP video adapter will discard the corrupted frame and reinsert the previous one to compensate for the loss. Loss may be perceivable to the receiver for the duration of the event.
- **GOP-based compression:** When using GOP-based compression, a network convergence event that lasts only tens of microseconds can affect video quality up to the GOP size, and possibly beyond. This effect can be on the order of seconds with some encoding profiles, as highlighted in Figure 9.



Figure 9. MPEG-2 Video GOP-Based Compression with a Slice Error Due to Packet Loss*

*Source material copyright SMPTE, used with permission

Scheduling Applications

Broadcast services are either permanent (24/7) or Occasional Use (OU), and are therefore setup for a given period of time. In order to prevent resource shortages while running these real-time services, broadcaster networks historically run on dedicated infrastructures that are set up so that all services have to be accounted for by 'booking' the required capacity and endpoints. Booking operators can use tools ranging from spreadsheets to graphical applications that incorporate all the endpoints and network nodes may be used.

Graphical-based scheduling applications allow booking operators to provision services without having to understand the details of IP technology, node, and endpoint settings. They also ensure that other bookings do not interfere with new requests or solicit resources already in use. Once the booking operator submits the configurations, the scheduling application pushes them toward all required devices (encoder/decoder endpoints, network nodes, QoS settings), via an automated process. This prevents configuration errors and allows for the use of pre-validated service templates.

These graphical applications are considered "off-path," since they interact with the network elements from a remote location that is not in line with the transmission path. "On-path" scheduling occurs when the source signals to the network its intent to transmit data, and requires a bandwidth reservation between the endpoints. (On-Path CAC is discussed in more detail in the "Achieving Quality of Service and Resilience in IP Networks" section of this document.) Note that scheduling applications can still employ On-Path CAC protocols on the network nodes.

Convergence Mechanisms for Transporting Video over IP

IP networks are fundamentally "connectionless" in nature, as described previously. In simple terms, packets are delivered into "the cloud" at one point in the network. Using the destination address of the packet, the network then makes a series of "hop-by-hop" forwarding decisions regarding where to send the packet. When the packet arrives at a router directly attached to the device referenced by the destination address, the packet is delivered. In many cases, an IP application does not care which specific path the packet follows through the network.

The emergence of MPLS has ushered in the advent of Traffic Engineering, allowing network engineers to define a specific path so that the network always forwards certain flows a certain way. This mechanism is employed mainly in environments in which the traffic rate of specific flows is relatively high compared to the bandwidth available. In such environments, bandwidth "hotspots" can occur. Traffic Engineering allows these hotspots to be avoided by steering high-rate flows around them.

In IP networks, "convergence" refers to the process whereby all routers agree on optimal routes through a network. When a network event (such as a link or node failure) changes the status quo, the routers send update messages, which in turn cause the routing algorithms to recalculate a new topology. When all routers agree on a new topology, the network is said to have converged. Minimizing network convergence times is the most important component in controlling loss in an IP Video network.

Modern IP routers use a separate control-plane and forwarding-plane. As a result, the router supports tasks associated with network convergence separately from packet-forwarding, and handles each more efficiently. Today, IP networks can converge within a few hundred milliseconds. The routers detect link failures almost immediately and feed this information back into the routing protocol subsystem. As a result, modern IP networks can maximize network availability and minimize loss to support real-time video services.

The following sections discuss the various mechanisms available to ensure that an IP-enabled broadcaster network can achieve maximum availability and reliability. They examine methods for ensuring that networks maintain service in the event of various outages, including various link recovery mechanisms, and source redundancy and stream redundancy schemes.

IP Convergence in WDM Networks

As discussed previously, modern IP routers can incorporate WDM technologies that allow for convergence of a degraded network interface even before the interface has failed completely. This technology is referred to as IP-over-Dense Wavelength Division Multiplexing, or IPoDWDM. (Figure 10.)

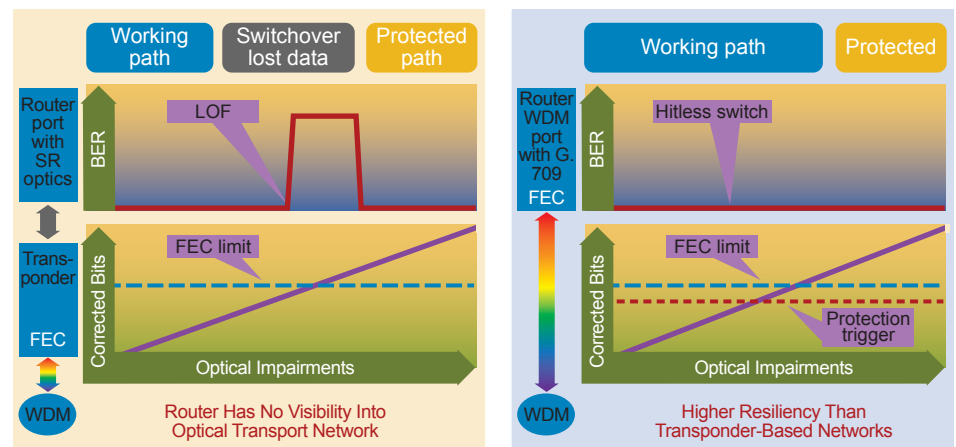


Figure 10. IPoDWDM

IPoDWDM is based on the integration of DWDM transponder capabilities into a port of an IP router. With this integration, the router can monitor for errors at the optical layer and trigger a switch to a protected path before any data loss is incurred. (See the right side of Figure 10.) Contrast this model with the non-integrated approach of conventional platforms shown on the left. In this model, the router initiates a switch to a protected path only when it detects a “Loss of Framing” (LOF) at the optical layer. Naturally, this advanced convergence capability can greatly enhance service availability.

Bidirectional Forwarding Detection

When two routers are not directly connected (e.g. when they are interconnected by an Ethernet switch), network engineers can employ alternate mechanisms to ensure that the network rapidly detects and adapts to topology changes. Technologies such as Bidirectional Forwarding Detection (BFD) can signal routing protocols in response to any discontinuity between two routers.

The advantage of BFD is that it is used only as a means of measuring continuity between two routers across a Layer-2 path. That means that protocols running between the routers do not have to rely on their individual timers, but can reference the BFD state instead.

Routing Protocol Enhancements

Modern IP networks employ routing protocol enhancements that ensure extremely fast convergence after link failures. For example, network engineers can configure specific IP network addresses with higher priority. As a result, the routing tables converge first for these addresses (e.g. prioritizing the convergence of the video source or video adaptors, or converging the “important” IP addresses in a network). All of these enhancements are referred to as IP Fast Convergence (IP FC), and often do not require extra network-wide engineering. On today’s high-end routers, unicast routing can often converge in less than 200 milliseconds, and multicast routing can converge in less than 500 milliseconds for more than 800 concurrent multicast groups.

Traffic Engineering

Another feature of modern IP networks is the ability to create specific paths through the network for specific flows. Technologies such as MPLS allow network engineers to build specific paths through an IP core and to carefully steer specific flows onto those paths. With MPLS TE, it is possible to always ensure that the more bandwidth-intensive (or higher-demanding) application streams receive the best possible service from the IP network.

Using MPLS TE, network engineers can also create a highly available backup scenario by providing a backup tunnel to protect against the failure of a specific network link. This technique is called MPLS TE Fast Re-Route, as described previously. By ensuring that the backup tunnel always follows a different path (excluding the link being protected), MPLS TE FRR ensures that when a link fails, the protected streams are automatically routed via the specified alternate path. Typically, MPLS TE FRR can reroute around a failure in less than 50 milliseconds.

Multicast-only Fast Re-Route (MoFRR).

A further enhancement for providing highly available multicast services is Multicast-only Fast Re-Route (MoFRR). The name of the technique is somewhat misleading, in that it implies some connection with (or dependency on) MPLS TE FRR techniques. In fact, MoFRR does not require an underlying MPLS infrastructure and delivers resilient service in both pure IP and MPLS environments. Figure 11 shows the MoFRR approach.




Figure 11. Multicast-Only Fast Re-Route

The MoFRR case shown in Figure 11 involves a receiver (R) that is connected to a source (S) via a router (D), which has more than one available path to that source (S). In a standard PIM environment, the router attached to the receiver would choose one of the upstream paths for the stream in question. If a failure occurred on the active path, the router would detect a change and begin sending multicast “join” requests via the alternate path to begin receiving the stream from the new path. In the traditional multicast model, some amount of time would always elapse between the loss of the active stream on the primary path and the recovery via the alternate path.

With MoFRR, the network avoids the delay incurred in waiting for the backup path to be built by always maintaining the backup path and always receiving the alternate stream (via B), alongside the active stream (via C). The router with two paths available therefore always receives two streams, and simply discards one of those streams as long as the primary path is available. Obviously, MoFRR incurs more network bandwidth and requires more processing power on the router to make the discard decisions. However, the model does achieve a hitless switchover to the standby in the event of the loss of the primary path.

Choosing the Right Convergence Technique

The appropriate convergence technique for a given application depends on the amount of loss that application can handle. Note, however, that achieving minimum loss often leads to extra complexity and its associated costs. Some applications, such as those that are compressed using GOP, always have a finite chance of losing important information inside the packet streams, regardless of the underlying convergence technique (i.e. whether an MPEG I-frame is lost during a 50-millisecond or 200-millisecond outage, the visual outcome is the same.) For uncompressed (or frame-by-frame compression), the video loss is proportional to the time it takes the network to converge.

The only way to achieve a lossless experience during a network convergence event is to add redundancy at the application level. This is accomplished either through adding extra information to the IP stream (using AL-FEC), or by sending the stream twice (referred to as Live-Live, or spatial redundancy, as discussed previously). More specifically, the spatial redundancy technique sends the stream twice across a different part of the topology, as illustrated in Figure 12.




Figure 12. Spatial Redundancy

The spatial redundancy approach relies on the availability of diverse paths between a source and receiver. Network engineers can employ various mechanisms for engineering these diverse paths. In the case in which two different physical topologies exist, MoFRR (described above) can be used to engineer diverse paths. Another mechanism, referred to as Multi-Topology Routing, utilizes control plane software within routers to create different logical topologies within a single physical topology. Whichever mechanism network engineers use to provision the diverse paths, the approach requires that there be a point in the network where both streams are received (shown as the Edge Distribution device in Figure 12). That device is then responsible for making an intelligent discard decision of one of the two streams. In the event that the primary stream is lost the receiver device must select the same primary stream via the other path. In the case of MoFRR, the router itself forwards two streams, and its neighbor is responsible for dropping the duplicate packets.

Anycast Source Redundancy

In addition to the approaches described above, network engineers can engineer further high availability with respect to the endpoints themselves. One mechanism is a technique called “Anycast Source Redundancy.” This technique requires that there be two copies of the same content sent into the network with identical appearance from an IP standpoint. In other words, the same content stream is made available from two different locations but with the same source and group address from each location.

Additionally, this technique requires that the source head-ends be capable of signaling to the network the availability of the respective streams. Each content stream emitted by a specific source must be signaled with a different availability message from an IP perspective. In routing terms, this means that each individual stream is sourced with a unique Unicast Source Address (even if each stream is emerging from the same source and the same physical interface).

Most network engineers using this technique configure the source (or some monitoring system adjacent to the source) to generate a “reachability advertisement” on behalf of each stream that it is currently sourcing. For example, a source sending (S1,G1), (S2,G2), etc. would simply announce that S1, S2, etc., were all reachable. In the event that a single stream became unavailable, the announcement would change for this stream to signal that it is no longer reachable on this interface. These announcements are accepted by the first-hop router attached to the source, which triggers a routing update that is dynamically propagated to the entire network. In this way, all routers in the network can rapidly learn if there is a stream that has become unavailable at a particular source/head-end.

With Anycast Source Redundancy, a discrete component failure that affects only one stream does not require a wholesale rebuild of all trees sourced from the same head-end. Every router sees two routes to the same source, and selects the closest one. In the event of link failures, the distance to the sources can change, or one source can even disappear from the routing table. Effectively, the technique always yields the optimal forwarding of multicast traffic from the closest available source. If a loss of any of the source streams occurs, the reachability of that stream (and only that stream) is signaled to the network. This triggers a change in the routing table such that all routers now only see one route to a given source stream.

Packet Retransmission

As discussed previously, network engineers can use RTP retransmission to retransmit video frames if they are lost across the last-mile connection. DSL networks, for example, are highly susceptible to interference from many impairments, and are therefore known to have relatively high bit error rates. In such scenarios, RTP retransmission can help tremendously. In order to avoid application delays, the RTP retransmission can be performed from a network-based appliance that is located close to the receiver. This network-based appliance receives the same multicast streams as the receivers. It monitors for errors signaled from the receiver and fulfills requests for packet retransmission within the time constraints of the receiver’s jitter buffer. A software client on the receiver synchronizes the retransmission from the network-based appliance and the streams coming from the original multicast source.

Conclusion

This paper has described how IP and MPLS technologies have the capability to deliver highly available point-to-point and point-to-multipoint services to meet today's broadcast video requirements. With the inherent capabilities of modern IP routers and high-speed optical and Ethernet transport, broadcasters can take advantage of the substantial OPEX and CAPEX savings of IP networks. With no need to statically provision circuits and bandwidth, they can achieve much more efficient utilization of their networks and resources, while ensuring the bandwidth and QoS to support even the most demanding applications.

IP and MPLS technologies may once have been considered best suited for data and Internet communications, but they have clearly evolved. Today, broadcasters can employ both connection-less (with QoS) and connection-oriented/traffic engineered options, depending on the application requirements and the administrator's choice. There is no longer a need to deploy an SDH-type transport infrastructures to support video services, as IP and MPLS networks can now meet the same stringent service-level agreements (SLA). Additionally, IP and MPLS networks offer the extra advantage of service convergence, allowing broadcasters to use the transport infrastructure for other non-video services, such as Internet access and VoIP.

The dynamic nature of IP networks, employing techniques such as IP Fast Convergence and the one-to-many nature of PIM-SSM, make them an excellent choice for distributing video in secondary distribution networks. Coupled with MoFRR and Live-Live techniques, network engineers can achieve lossless video delivery in these networks – an essential requirement for supporting GOP-based compression techniques. The network can also employ MPLS-based Layer-3 VPNs to offer different instances, often per service/content provider. In contribution environments where traffic flows may be more static, predictable, and scheduled, and where no compression is typically applied, a traffic engineered approach such as MPLS-TE or P2MP MPLS-TE can provide an effective solution.

In order to support these combined requirements and capabilities, Cisco is developing a unique suite of networking technologies. These technologies span Cisco's broad range of video-optimized products, and allow broadcasters to create a next-generation "medianet" – an intelligent network that is optimized end-to-end for the delivery of extraordinary media experiences. Ultimately, they allow broadcasters to create a single, scalable IP architecture that extends from the point of content ingest through every aspect of editing and production, across video contribution and distribution networks, all the way to the customer's screen.

By embracing this medianet approach, broadcasters can:

- Transform the customer experience by delivering more content, mobility, personalization, and control
- Assure a high-quality customer experience end-to-end
- Virtualize content and applications through every phase of the media value chain to drive down CAPEX and OPEX
- Monetize content and advertising in new ways, across more platforms

To find out more, visit:

<http://www.cisco.com/web/solutions/medianet/sp.html>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0903R)