White Paper

Best Practices for Video Transit on an MPLS Backbone

Using Point-to-multipoint Label Switched Paths to Simplify and Optimize IP Video Distribution

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408.745.2000
1.888 JUNIPER
www.juniper.net

Part Number: 200106-002 Sept 2007

Table of Contents

Copyright ©2007, Juniper Networks, Inc

Executive Summary

This paper describes how an innovative new technology called point-to-multipoint label switched paths (LSPs) enables service providers to efficiently carry broadcast video (IPTV) traffic across an IP backbone. This technology can be used for distributing IPTV traffic to residential subscribers, across a private network, or within a virtual private network (VPN).

Combined with other recent enhancements, MPLS is now the technology of choice for converging services. With the development of point-to-multipoint LSPs and related multiplay technologies, Juniper Networks has taken a leading role in providing the tools to support the transport of all types of video—from selected content such as video on demand (VoD) to interactive multimedia and broadcast television—across an IP/MPLS backbone.

The Juniper Networks vision for the evolution of networks is an open, next-generation network core[1] that will securely and reliably support all communications applications, from any origin to any destination, at any time. This paper explains this vision and shows how it is supported by point-to-multipoint LSPs.

Introduction

Today, service providers must distribute large volumes of video across backbone networks to an increasing number of distribution sites (often referred to as video hub offices or video serving offices) and subscribers. New deployments rely on IP routing to offer a wide variety of services across a common backbone, including VPNs, voice, Internet access, and video services. As service providers integrate IP-based video, ATM and SONET/SDH become inefficient for delivering this traffic to a wide number of endpoints. ATM and SONET/SDH were not designed to support IP multicast, which is the underlying technology used to deliver broadcast television channels. They also preclude the network from supporting a larger set of other services as part of a multiplay offering.

Modern routing platforms scale to handle the large volume of video traffic to be delivered to various endpoints across the network, along with data flows from other service offerings. Key drawbacks to native IP forwarding are the lack of path control and deterministic resiliency. This has led to the rapid adoption of MPLS, which provides traffic engineering and rapid failure recovery across IP routing devices on par with SONET/SDH.

Although IP/MPLS traffic engineering and instant failover capabilities meet the requirements for video distribution, the lack of replication capabilities with LSPs has hindered the use of IP/MPLS for video transport. Without replication, MPLS LSPs cannot use network resources efficiently. The Juniper Networks point-to-multipoint LSP solution retains the traffic engineering and availability benefits of MPLS while providing replication without the overhead of multicast routing. Juniper Networks delivers this solution as part of a larger suite of technologies to handle video transport across the backbone. The Juniper Networks point-to-multipoint solution efficiently replicates IPTV and other broadcast video traffic from one main content distributor to multiple local studios or head-ends. point-to-multipoint LSPs are supported on all Juniper Networks T-series and M-series routers, including the new T1600 routing node.[2]

[1] For more information, see the *Juniper Networks Next Generation Core Network Solution for Service Providers* at www.juniper.net/solutions/literature/brochures/160013.pdf.

[2] For configuration and usage guidelines related to point-to-multipoint LSPs, see the *MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Copyright ©2007, Juniper Networks, Inc

Overview: Video Distribution and Point-to-multipoint LSPs

Historically, video distribution has been handled by terrestrial ATM or SONET/SDH networks. Today's content and video providers are typically opting for more modern options such as IP/MPLS. However, until recently, IP/MPLS deployments have been limited to point-to-point connections, which are not efficient for video broadcast distribution to multiple destinations. Since a single uncompressed stream may be up to 260 Mbps, sending a separate copy of each stream to each destination can quickly exhaust network bandwidth.

A point-to-multipoint LSP is an MPLS LSP with multiple destinations. By taking advantage of the packet replication capability of the network, point-to-multipoint LSPs avoid unnecessary replication at the ingress router. This solution was first developed in the Internet Engineering Task Force (IETF) and is now deployed in many production networks. Figure 1 illustrates broadcast content being sent over a point-to-multipoint LSP to a number of broadband access networks.

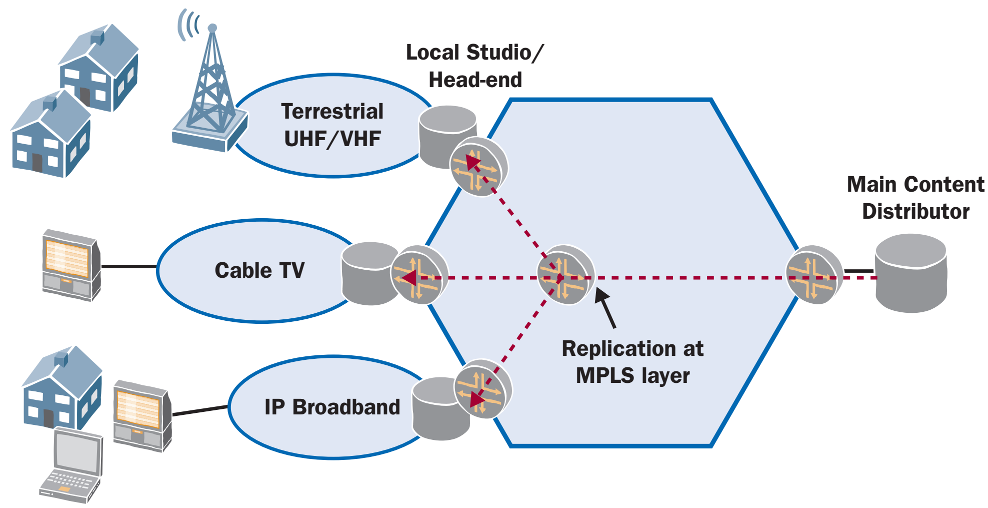

Figure 1: Broadcast Content over Point-to-multipoint LSP

By adding point-to-multipoint support, MPLS networks are able to efficiently deliver both unicast and multicast content over a common network. IP/MPLS provides all of the features offered in existing legacy networks without affecting the subscriber viewing experience. Since IP/MPLS includes traffic engineering for performance and high availability, quality of service (QoS), resource optimization, and security, it serves as an ideal converged backbone—one that enables a wider variety of service offerings for service providers.

Baseline Requirements for Broadcast Content Distribution

The broadcast industry presents many tough challenges. The demands of broadcast television networks are far more stringent than those of corporate Webcasts or distance-learning courseware. In addition to high definition television (HDTV), broadcast content distribution has comparatively strict quality and resiliency needs. For instance, viewers of a corporate presentation on a PC may tolerate a little jitter, but consumers of a live sporting event or a suspenseful movie on the living room TV will find the slightest loss of a frame unacceptable. The requirements of broadcast content distribution are interrelated, but they can be characterized as follows:

Copyright ©2007, Juniper Networks, Inc

Performance

Broadcast content distribution on a large scale means massive quantities of digital content at high speed (up to 260 Mbps) for network TV broadcast. Because of the strict QoS and resiliency requirements, this performance must be maintained under high load.

MPLS provides the ability to rapidly and easily grow the network as site count and content capacity increase. Juniper Networks core routing platforms scale into the multi-terabit-per-second range with wire speed, low-latency forwarding, and the ability to support a high density of 10-Gbps Ethernet and SONET/SDH ports.

Quality of Service

Packet loss is an especially important consideration. For instance, a one-second loss will be experienced as frame freezes, asynchronous dialogue, or slight omissions from the program, and may be displayed for several seconds. The loss of even a handful of packets can result in a noticeable—and unacceptable—blip on the TV screen. Jitter can also be an issue if network delays extend beyond the ability of the set-top box to compensate.

Advanced queuing mechanisms, Resource Reservation Protocol with traffic engineering extensions (RSVP-TE), and Differentiated Services (DiffServ) help resolve these issues. Juniper Networks core routers can maintain jitter performance for high-priority traffic under heavily oversubscribed conditions. The routed network also supports the ability to recover from hot spots between specific locations. These hot spots may be the result of shared facilities, expanded adoption of HDTV, or increased broadcast channels. MPLS traffic engineering allows available bandwidth to be reserved over a selected path. It also plays a role in resiliency, ensuring delivery through a separate backup path by coloring links or nodes and specifying the colors that an LSP connection should follow (or avoid). This ensures that connections that are part of a protected circuit will never travel through a common point. In addition, Juniper Networks routers guarantee that high-priority traffic gets preference over low-priority traffic, even in the event of link node failures.

Resiliency

The network must support rapid recovery from a failure, since video is extremely loss-sensitive. Even with forward-error correction capabilities in the video layer, the reduction of outage times to sub-second intervals is critical to sustain high levels of video quality and minimize error recovery dependencies. Expected failover for video is on the order of that provided by SONET/SDH (50 ms). Traditional, last-generation routers were typically unable to deliver the resiliency requirements for broadcast content distribution.

Juniper Networks routers use MPLS Fast Reroute (FRR) to meet the resiliency requirements for broadband content distribution. MPLS can also recover from an outage using local repair techniques. Additionally, MPLS can reallocate lower traffic class bandwidth resources to video in a converged core so that an entire duplication of capacity is not held in reserve as it is with SONET/SDH rings.[3] This reduces the number of idle links in the network while maximizing revenue potential.

The use of a bypass LSP for link protection can also be configured into point-to-multipoint LSPs. The bypass LSP uses a different interface and path to reach the same destination. Similarly, Graceful Restart (GR) can be configured on point-to-multipoint LSPs. This allows a router undergoing a restart to inform its neighbors of its condition, and thus receive a grace period on control plane communications. The restarting router can also still forward MPLS traffic during the restart period.

[3] SONET/SDH rings typically reserve 50 percent of the allocated bandwidth for backup.

Copyright ©2007, Juniper Networks, Inc

Finally, Juniper Networks core routers also provide non-stop forwarding in an MPLS environment. A key feature enabling this support is the clean separation of the forwarding and control plane in the router, along with improved lookup speed and lower-latency buffering. This enables the router to continue forwarding packets without a hitch during software upgrades or when the router is undergoing craft maintenance.

Efficient Transport of Broadcast Video

For bandwidth efficiency, it is necessary to perform an efficient replication within the network to eliminate duplicated traffic over the same link. It is this capability that has historically been missing from MPLS. Standard point-to-point LSPs do not provide efficient distribution. For example, in the point-to-point LSP illustrated in Figure 2, content is sent four times from PE1, even though it is only destined for two routers on the next hop.

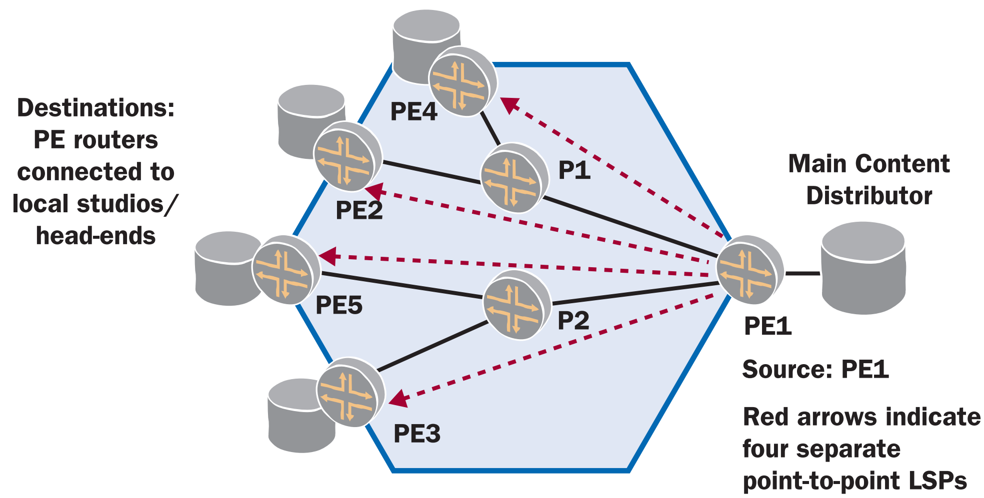

Figure 2: Point-to-point LSP

A possible way to solve this is to run a multicast routing protocol in the core, such as Protocol Independent Multicast (PIM). However, many service providers have good reasons to resist running PIM in the core.

Juniper Networks introduced the point-to-multipoint LSP in order to support the efficient transport of broadcast video. With point-to-multipoint MPLS, a new site becomes a new branch where replication is required. This takes advantage of the replication benefits of multicast in a large-scale network by sending only a single packet from the source location and replicating frames closer to the network edge when needed. There is no need to insert the site into a large multicast SONET ring, which can be a long and expensive process. The size of the link into the new location is based on resource requirements to that specific site, independent of other locations.

Copyright ©2007, Juniper Networks, Inc

Table 1 summarizes the advantages of using point-to-multipoint LSPs instead of an IP multicast protocol, such as PIM, for video transport.

	Point-to-multipoint LSPs	Multicast Routing Protocol (e.g., PIM-SSM) over IP
Protection	Strict service level agreement (SLA) Link protection Graceful Restart (GR) Fast reroute (FRR) “Make before break” capabilities	No such capabilities (PIM is not a traffic engineering protocol)
Resource Reservation	Strict SLA Resource reservation once for each point-to-multipoint tree	No resource reservation mechanisms
Explicit Path Routing	Supports explicit routing along paths, different from hop-by-hop IP	No equivalent support

Table 1: Advantages of Using Point-to-multipoint LSPs

Because the receivers are typically static hubs, the only facet of a multicast routing protocol required for this particular application is replication, which is provided by the point-to-multipoint LSP. High Availability (HA), resiliency, convergence, and traffic engineering are all best handled by MPLS.

The use of PIM and point-to-multipoint LSP is not always an “either/or” situation. For example, you can use point-to-multipoint LSPs to distribute multicast traffic to PIM islands situated downstream from egress routers. This is enabled by the ability to control whether a reverse path forwarding (RPF) check is performed for a source and group entry before the route is installed in the multicast forwarding cache.[4]

Replication in Point-to-multipoint LSPs

A traditional point-to-point LSP has one ingress point and one egress point, but a point-to-multipoint LSP has a single ingress node with multiple egress nodes. This replication process is illustrated in Figure 3. Router PE1 is configured with a point-to-multipoint LSP to routers PE2, PE3, PE4, and PE5. When router PE1 sends a packet on the point-to-multipoint LSP to routers P1 and P2, router P1 replicates the packet and forwards it to routers PE2 and PE4. Router P2 sends the packet to router PE3 and PE5.

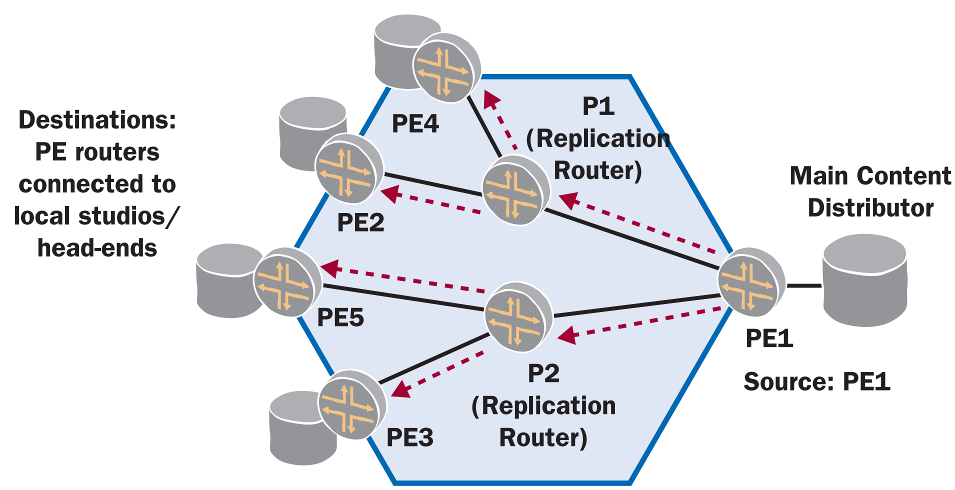

Figure 3: Replication in a Point-to-multipoint LSP

[4] This is necessary because the interface receiving the multicast traffic on a point-to-multipoint LSP egress router might not always be the RPF interface.

Copyright ©2007, Juniper Networks, Inc

This maps a single ingress to multiple egress points and is unidirectional; traffic flows from ingress to egress points, but not vice versa. Packet replication occurs in the MPLS layer on the P routers.

In developing this technology, Juniper Networks has extended the functionality of RSVP-TE for both MPLS and Generalized MPLS (GMPLS). A point-to-multipoint LSP is set up with RSVP-TE constraints, allowing efficient packet replication at various branching points in the network. IP traffic (unicast or multicast) can be transported along this LSP, and Layer 2 traffic can be transported using Circuit Cross-connect (CCC) encapsulation. Traffic engineered paths may be defined statically or through the use of Constrained Shortest Path First (CSPF), which allows the network to dynamically determine the best path based upon criteria such as bandwidth requirements, hop limitations, and network priority.

For particularly important traffic, distributors might want to use application-level protection (i.e., two copies of the video stream are sent along diverse paths and the video receiver can seamlessly switch between streams). The local distributor can choose between the primary path and the backup path, as illustrated in Figure 4.

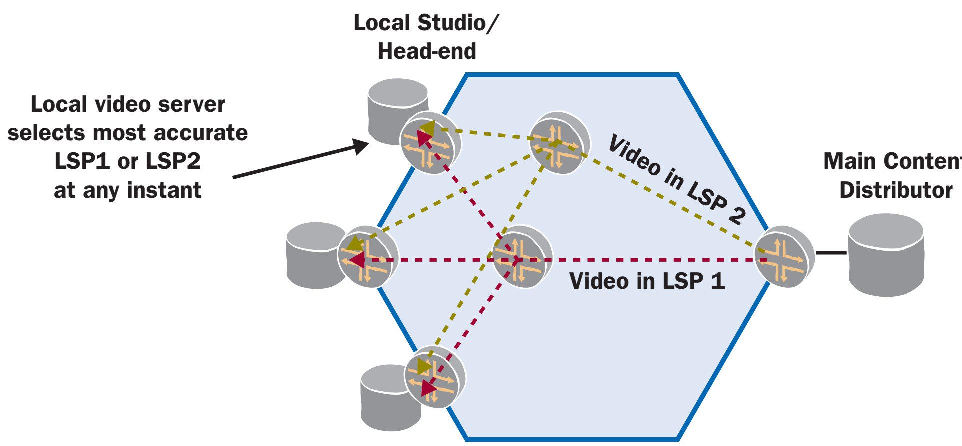

Figure 4: Backup Path Selection

Copyright ©2007, Juniper Networks, Inc

Typical Point-to-multipoint Application

Figure 5 shows a typical broadcast application using Layer 2 multicast with point-to-multipoint LSPs. Primary and secondary point-to-multipoint LSPs are set up between the content source and the local distributors.

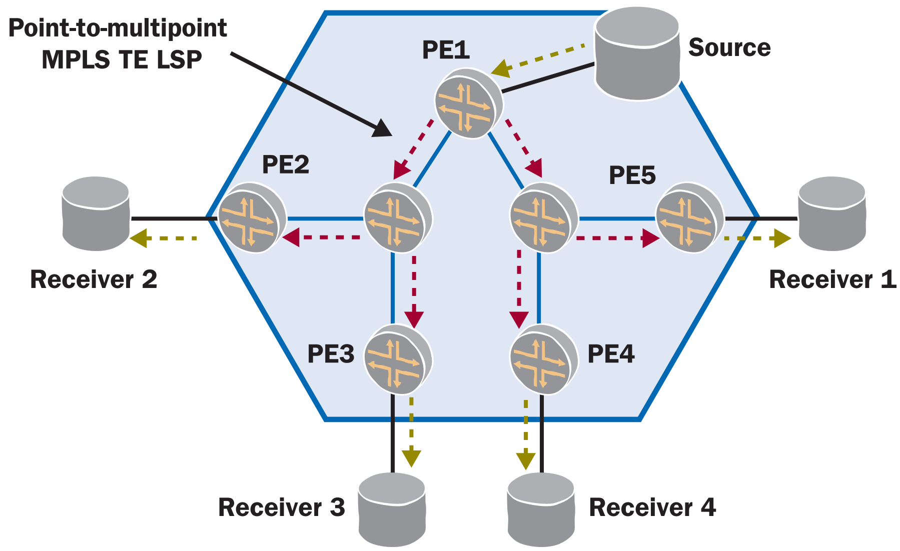

Figure 5: Broadcast Content Distribution

Note that there might be a different source at different times of day, and the broadcast would occur over different point-to-multipoint LSPs. For example, at a certain time of day, PE1 might be the head-end of a point-to-multipoint LSP (as shown in Figure 5), but at another time of day one of the other PEs might be the source.

In this application, the broadcast TV would typically be transformed into IP packets and encoded into Ethernet frames or ATM cells. The PE router at the ingress (PE1) receives the stream of encoded IP packets from the video encoder, and maps this traffic to the point-to-multipoint LSP. Thus, the ATM cells have MPLS labels appended and they travel down the point-to-multipoint tree so that each receiver gets a copy. This saves bandwidth utilization on the network by avoiding the need to establish a point-to-point LSP for each stream to a receiver.

RSVP's call admission control is used for the video traffic. Each point-to-multipoint LSP has a unique bandwidth reservation, and traffic is policed to the size of the bandwidth reservation. Doing so ensures that, once traffic enters the LSP, resources are available to carry the traffic to its destination.

This approach retains all the functionality and QoS guarantees available to Layer 2 services as they migrate to IP/MPLS. As the virtual circuit is mapped to an LSP, it also retains the FRR capabilities of MPLS.

IP unicast, IP multicast, and Layer 2 traffic can all be mapped to point-to-multipoint LSPs. This opens up the possibility for a phased migration, which may be necessary when non-IP video is transported in ATM virtual circuits (VCs). Here, you would take the ATM VC solution (as noted on the left in Figure 6) and, in the first step, map the ATM circuits to the LSP (as shown in Figure 6 on the right). Using this approach, it is possible to eventually migrate over to an all-IP solution.

Copyright ©2007, Juniper Networks, Inc

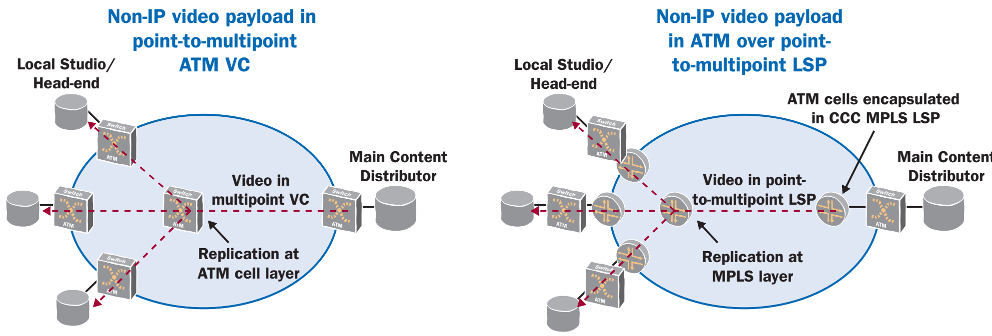

Figure 6: Migration of Video Distribution Solution

Other Video Transport Scenarios

Juniper Networks also delivers tools for supporting other video applications, such as unicast streams for VoD and multicast VPN (MVPN). VoD provides broadband subscribers with the ability to view programs at any time, thus eliminating the finite choice of broadcast channels. VoD is typically a unicast application, but has high bandwidth—though not as high as content distribution, which might transmit tens or hundreds of channels at a time. Typically, VoD also has high QoS requirements, although business-oriented or distance-learning applications may have more relaxed resolution and jitter requirements than entertainment applications.

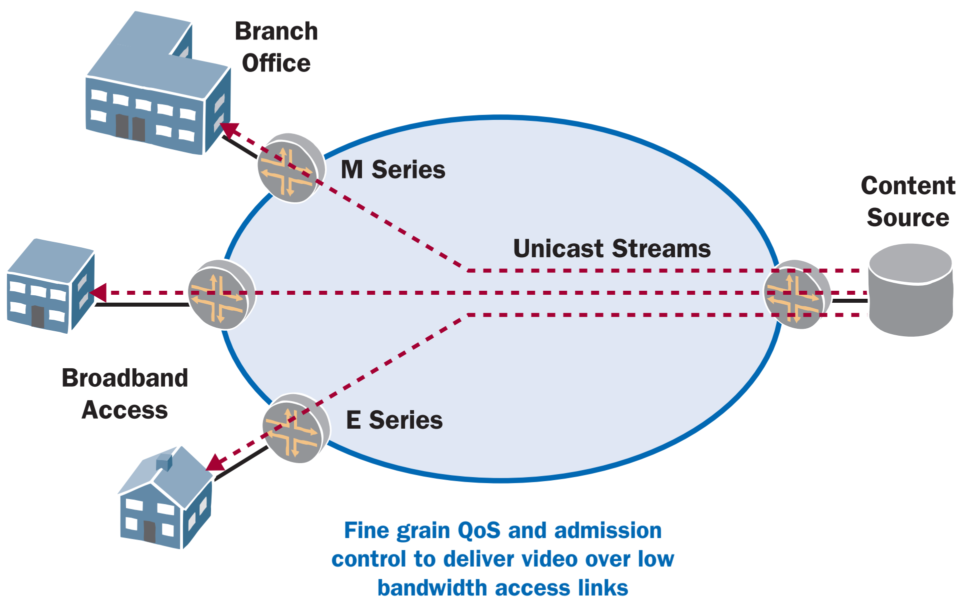

Figure 7: Video on Demand

Copyright ©2007, Juniper Networks, Inc

Video can also be distributed over MVPNs, which would typically be used for corporate Webcasts. Moving forward, Juniper Networks intends to make this more efficient by incorporating point-to-multipoint LSPs into this tool.

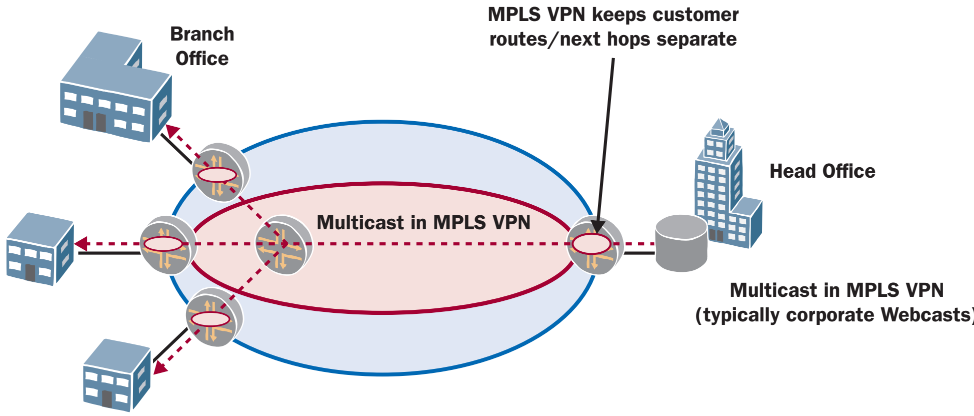

Figure 8: Multicast VPN

Additional Services for Content Distributors

In the use of MPLS for this application, content providers are already enjoying some of the benefits of other IP/MPLS services. For example, MPLS VPNs are currently being used for transmitting non-live content (such as taped programs) to distributor head-ends for further processing. This type of activity has traditionally been handled by copying the content to tape and transmitting it by postal system or courier. However, new file formats have allowed sufficient compression so that the content can be transmitted in an MPLS VPN. It is easier for the service provider to have one network carrying traffic from multiple broadcast companies. It is anticipated that eventually a lot more wide-area traffic will not be real time and will be transported in this way. When a common network infrastructure is used for both real-time video and file transfers, it is much simpler to adjust the differences in loads between these two scenarios_especially compared to migrating capacity between two distinct networks.

Conclusion

Juniper Networks has taken a leading role in developing the tools to support the distribution of video across an IP/MPLS backbone. Along with today's leading solutions for video distribution, Juniper Networks is providing the vision to meet tomorrow's backbone video needs, and is also taking a leading role in the standards process to ensure interoperability in all facets of future video distribution. The point-to-multipoint LSP, along with the high-performance and HA features of Juniper Networks core routers, is a key feature that enables service providers to migrate high-value broadcast video distribution from legacy transport to an IP/MPLS core.

Further information on these MPLS enhancements please refer to these documents:

> RFC 4875: Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)

> RFC 4461: Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)

Copyright ©2007, Juniper Networks, Inc

Appendix A: Technical Details of the Point-to-multipoint Solution

This section takes a closer look at the mechanisms behind point-to-multipoint LSPs. The labels in Figure 9 introduce the concepts of a Source PE (Spe), Receiver PE (Rpe), and point-to-multipoint-ID (Pid) where:

> Spe: PE connected to the traffic source

> Rpe: PE connected to one or more receivers

> Pid: maps a set of Rpes to a point-to-multipoint tree

Using the following diagram as an example, a point-to-multipoint LSP is first set up from Spe1 to Rpe1, Rpe2, and Rpe3.

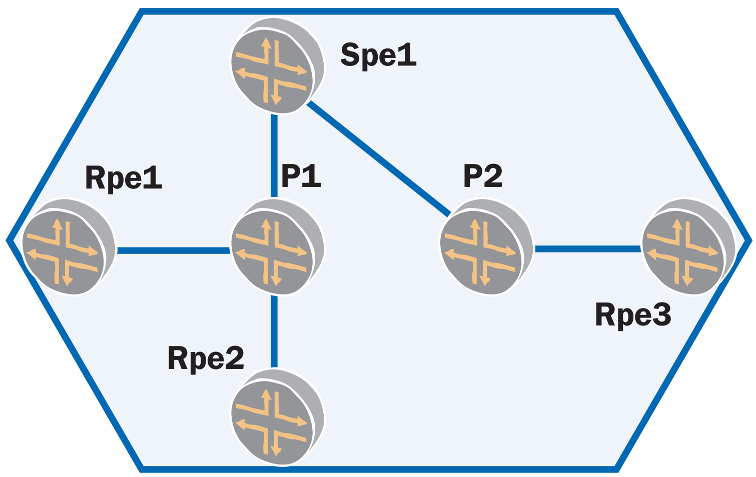

Figure 9: Components of the Solution

As far as the control plane is concerned, the source (Spe1) initiates individual point-to-point LSPs to each receiver for a given point-to-multipoint LSP. The point-to-point LSPs use a common point-to-multipoint Session Object and an identifier to distinguish a constituent branch LSP. Thus, a point-to-multipoint LSP is actually created from a set of point-to-point LSPs. A P2P LSP is a sub-LSP of the point-to-multipoint LSP. There is one point-to-point LSP from Spe to each Rpe. The RSVP protocol ties the point-to-point LSPs together into a point-to-multipoint LSP.

The Spe initiates individual point-to-point LSPs to each Rpe for a given point-to-multipoint LSP. With RSVP signaling, each Spe originates a PATH message and each Rpe responds with RESV messages. The forwarding operations on the replication routers are handled with a single-labeled packet going into the core and multilabeled packets going out, each on a different output interface. Figure 10 highlights the solution.

Copyright ©2007, Juniper Networks, Inc

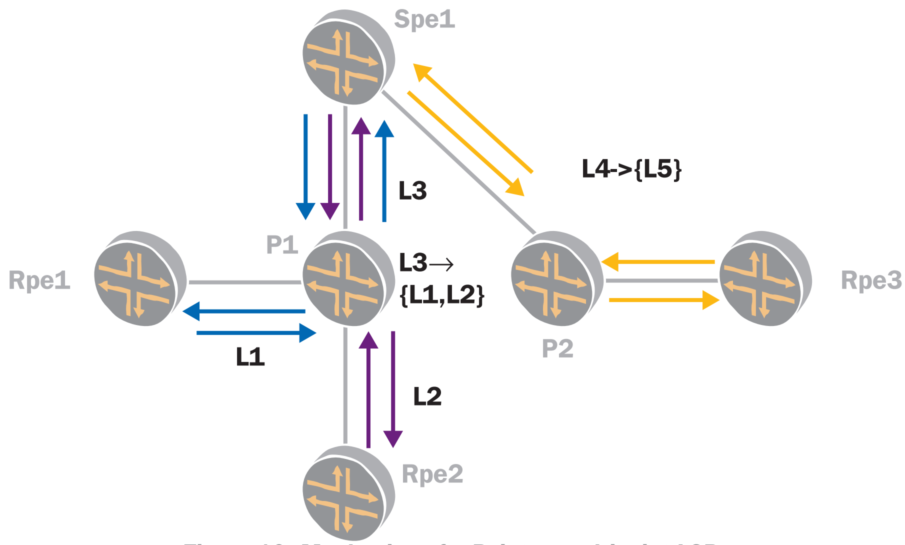

Figure 10: Mechanics of a Point-to-multipoint LSP

In this example, the blue arrows are PATH messages associated with a sub LSP, and the return arrows (back to the Spe) are RESV messages. Router P1 realizes that the two RESV messages (i.e., the one received from Rpe1 and the one received from Rpe2) correspond to the same point-to-multipoint LSP, and thus uses the same label for the two corresponding reservation messages that it sends to Spe1.

Copyright ©2007, Juniper Networks, Inc

Appendix B: References

The following references provide additional information on Juniper M- and T-series routers, advanced MPLS support, and advanced applications solutions.

Documentation

JUNOS Software Release 8.4: MPLS Applications Configuration Guide
http://www.juniper.net/techpubs/software/junos/junos84/index.html

White Papers

Architectural Issues in Carrier Class Operating Systems

A summary of the software engineering practices and architectural issues that lead to carrier class operating systems. Highlights advantages of the JUNOS operating system.

www.juniper.net/solutions/literature/white_papers/200209.pdf

Carrier Class Routing with the JUNOS Advantage

A discussion on the control plane in Juniper Networks M- and T-series routers.

www.juniper.net/solutions/literature/solutionbriefs/351081.pdf

Efficient Scaling for MultiService Networks

Details the packet processing advantages, including scaling and service flexibility, of the ASIC-driven SRAM search technology used in Juniper Networks M- and T-series routers.

www.juniper.net/solutions/literature/white_papers/200207.pdf

Juniper Networks Next Generation Core Solution for Service Providers

Details the agility and service awareness of the T1600 router in conjunction with the Session Resource Controller (SRC) portfolio and JUNOS software for delivering advanced services such as VOD and broadcast video across an , next-generation network core.

www.juniper.net/solutions/literature/brochures/160013.pdf

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

Copyright 2007, Juniper Networks, Inc. All rights reserved. Juniper Networks and the Juniper Networks logo are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered trademarks, or registered service marks in this document are the property of Juniper Networks or their respective owners. All specifications are subject to change without notice. Juniper Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Copyright ©2007, Juniper Networks, Inc