

NETPLANE

A Comparison of Multiprotocol Label Switching (MPLS) Traffic-Engineering Initiatives

Definition

Multiprotocol label switching (MPLS) is one of several initiatives to enable delivery on the promise of a converged network. By combining the attributes of Layer-2 switching and Layer-3 routing into a single entity, MPLS provides the following benefits:

- enhanced scalability by way of switching technology
- class-of-service (CoS)– and quality-of-service (QoS)–based services (differentiated services)
- no need for an Internet protocol (IP)–over–asynchronous transfer mode (ATM) overlay model and its associated management overhead
- standards-based solution, promoting interoperability
- enhanced traffic-shaping and engineering capabilities

Overview

It is becoming increasingly apparent that the ability to reshape and engineer traffic dynamically to ensure timely prioritization and delivery has become paramount. Currently, there is an initiative to deliver traffic-engineering capabilities to MPLS. There are two distinct approaches that incorporate the obvious benefits traffic engineering provides to the network core. This tutorial examines the fundamental changes being made to signaling protocols used to manage data in a large MPLS network and make it possible to determine the best route, given a set of stated bandwidth constraints. The similarities and differences between constrained-based label distribution protocol (CR–LDP) and traffic engineering (TE)–resource reservation protocol (RSVP) are discussed in detail.

Topics

1. Introduction
 2. MPLS Overview
 3. A Technical Discussion of MPLS Traffic-Engineering Initiatives
 4. Conclusions
- Self-Test
- Correct Answers
- Glossary

1. Introduction

The momentum toward voice and data convergence is driving the Internet to cope with new realities. Historically, the Internet infrastructure and protocols were intended and optimized solely for data. Traditional routing paradigms incorporating Internet gateway protocols (IGPs), such as routing information protocol (RIP) and open shortest path first (OSPF), and exterior gateway protocols (EGPs), such as border gateway protocol 4 (BGP4), no longer represent the optimal solution.

On top of traditional data traffic, the addition of hypertext transfer protocol (HTTP); voice, store, and forward messaging; multimedia traffic; and real-time electronic-commerce applications to the infrastructure are pushing toward ever-higher bandwidth requirements, as well as the ability to guarantee that bandwidth. This new reality is promulgating the development of new models to ensure guaranteed delivery of services such as voice, on par with the public switched telephone network (PSTN), regardless of unexpected interruptions in the network infrastructure. To the network, voice is just considered additional data, with very specific QoS and CoS requirements.

That Was Then, This Is Now

Thirty years ago, the early designers of IP had to address different, but no less difficult, challenges than the designers of MPLS. Back then, the state of the art of computing logic itself dictated how computers could communicate. By today's standards, central processing units (CPUs) were primitive and limited in their capabilities. Dynamic memory was slow and expensive. The facilities used to develop operating systems and communication protocols had yet to completely mature and be fully tested. Hence, the primary focus of early protocol development was to ensure the survivability of a truly decentralized network. Designers concentrated on functionality that supported segmentation, retransmission, and dynamic routing. Successful delivery of the data was the

primary concern. Therefore, at each stop along its journey through the network, the IP datagram was decomposed, verified, analyzed, and reconstructed before it was finally sent on its way. Amazingly, the fundamental infrastructure of the transmission control protocol (TCP)/IP has managed to survive and evolve into the global phenomenon that is the Internet. This is a testament to the tenets of a standardized approach to protocol development.

The Initial Step, RSVP Arrives

In the mid 1990s, traffic levels in large networks and the Internet increased to levels far beyond what conventional routers were able to handle. The network had to support mission-critical applications in which expedited delivery of data over the network was mandatory. RSVP had been designed to offer higher-quality delivery over the existing local-area network (LAN)-based networks. RSVP had been based on the original fundamental requirement of the Internet (i.e., process each IP flow in a hop-by-hop fashion) but now provides limited scheduling and traffic shaping of each forwarded IP datagram at the egress port.

Acceptance of RSVP was not universal for several reasons. To realize actual guaranteed performance, RSVP required each router along the routed path to support RSVP signaling and some level of priority queuing or traffic shaping. It also required devices at the edge of the network to initiate and respond to RSVP requests. The demand for RSVP capabilities was not great enough to require a wholesale upgrade of network hardware.

2. MPLS Overview

Goal of MPLS

Simply put, MPLS provides the ability to support any type of traffic on a large IP network without having to subordinate the design to the limitations of different routing protocols, transport layers, and addressing schemes. The design objective of the Internet Engineering Task Force's (IETF's) MPLS effort was to increase efficiency of data throughput by optimizing packet-processing overhead in the IP network.

In addition to the developments in routing, significant strides are being made in optimizing hardware as well. Increased processing capabilities, lower production costs, and more sophisticated, application-specific integrated circuits (ASICs) make it possible to mass-produce hardware that is capable of forwarding datagrams at wireline speed. Until recently, routing protocols and noncompatible physical layers have been a limitation.

This increase in processing capability and decreased cost, in combination with routing improvements, has made it possible to create a very large and reliable Internet infrastructure and switched paths through this faster and more robust network.

Service Level Agreements (SLAs) and MPLS

In less than 10 years, the public Internet has evolved from a government-funded experiment to a commercial juggernaut. There is a fundamental need to support mission-critical networks using IP routing and enhanced signaling protocols. QoS and CoS have become the watchwords of a converged network.

QoS is defined as those mechanisms that give network administrators the ability to manage traffic's bandwidth, delay, and congestion throughout the network. To realize true QoS, its architecture must be applied end to end, not just at the edge or at select network devices. The solution must provide a wide variety of technologies that can interoperate in such a way as to deliver scalable, feature-rich services throughout the network. The services must provide efficient use of resources by providing the aggregation of large numbers of IP flows where needed while providing simultaneous, fine-tuned granularity to those premium services defined by SLAs. The architecture must provide the devices and capabilities to monitor, analyze, and report detailed network status. Armed with this knowledge, network administrators or network-monitoring software can react quickly to changing conditions, ensuring the enforcement of QoS guarantees. Finally, the architecture must also provide mechanisms to defend against the possibility of theft, prevent denial of service, and anticipate equipment failure.

As an example, virtual private networks (VPNs) are fueling availability and reliability demands with their requirement for dedicated tunnels across the Internet. At present, VPNs are mainly implemented in a site-to-site scenario, requiring a dedicated connection. Service providers, such as GTE Internetworking and UUNET, offer outsourced VPN services and must therefore have a way to deliver predictable and reliable service to their customers. The ability of a VPN to establish and maintain a tunnel will be greatly enhanced by MPLS's ability to establish and guarantee CoS and QoS for a label-switched path (LSP). This will benefit VPN service offerings by allowing predictable connections and the ability to quantify this reliability in an SLA.

Why MPLS?

Two fundamental features make MPLS possible across very large routed IP networks. First, MPLS makes it possible to switch traffic through IP routers that, historically, had to interrogate each IP header before forwarding to the next hop. This is accomplished by applying a Layer-2 label to the IP frame as it enters the

edge of the MPLS-aware network. This label corresponds to an established (configured/signaled) path through the network, also known as an LSP.

Second, at the time a label is applied to the flow, predefined traffic-engineering parameters can be programmed into the forwarding hardware to guarantee levels of traffic bandwidth, delay variation, and congestion control. Once the data begins to flow, the network device must be able to monitor and report the actual level of resources being consumed at each interface.

MPLS Traffic Engineering

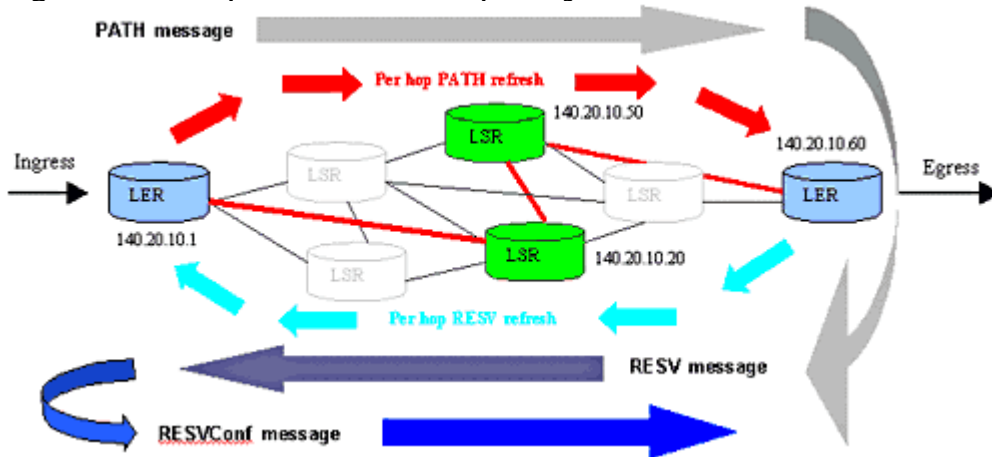
Two different approaches, TE-RSVP and CR-LDP are currently under development by the IETF MPLS Working Group. They characterize the signaling piece of traffic engineering within MPLS. There are two ways to implement an LSP within an MPLS network: control-driven (hop by hop) using LDP and explicitly routed LSP (ER-LSP).

Both TE-RSVP and CR-LDP represent the latter approach. What this implies is that, by having the ability to engineer the route using predetermined CoS and QoS parameters, the optimal LSP for a specific traffic type can be ensured. Further flexibility allows for the definition of loose and strict ER-LSPs. The strict ER-LSP follows a list of nodes using the actual addresses of each node to traverse, while the loose ER-LSP is more adaptive and allows groups of nodes, specified as an autonomous system number, to act as one of the abstract nodes to traverse.

TE-RSVP

MPLS traffic engineering by means of TE-RSVP proposes using extensions to the existing RSVP protocol, request for comment (RFC) 2205. Using TE-RSVP does not mean that a full implementation of RSVP is required to be run on each label edge router (LER) or label switch router (LSR) within an MPLS-aware network. An LER or LSR only requires that the extensions be able to support MPLS-explicit routing. TE-RSVP is a soft-state protocol and uses user datagram protocol (UDP) or IP datagrams as the signaling mechanism for LSP setup communications, including peer discovery, label requests, and mapping and management (see *Figure 1*).

Figure 1. Example of a Loose Explicitly Routed TE-RSVP LSP



In this example, having used BGP to discover the appropriate egress LER to route the traffic to another autonomous system (AS), the ingress LER initiates a PATH message to egress LER through each downstream LSR along the path. Each node receives a PATH message to remember this flow is passing and thus a path state or session is created. The egress LER uses the RESV message to reserve resources with traffic and QoS parameters on each upstream LSR along the path session. Upon receipt at the ingress LER, a RESVConf message is returned to the egress LER confirming the LSP setup. After the loose ER-LSP has been established, refresh messaging is passed between LERs and LSR to maintain path and reservation states. It should be noted that none of the downstream, upstream, or refresh messaging between LER and LSRs is considered to be reliable, because UDP or raw IP datagrams are used as the communication mechanism.

TE-RSVP feature set is robust and provides significant capabilities to provide traffic-engineering services to MPLS.

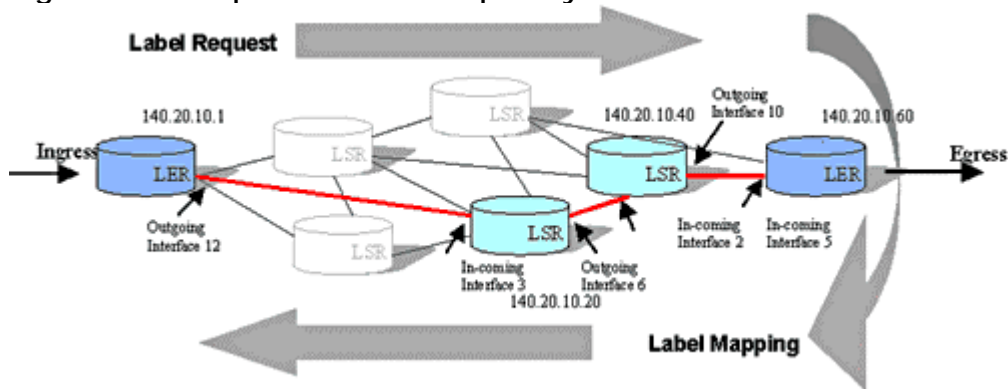
- **QoS and traffic parameters**—These are passed as opaque data to traffic management.
- **failure notification**—Upon failure to establish an LSP, an existing LSP will send failure message, but relies on timers for refresh messages.
- **failure recovery**—This is the “make before break” when rerouting.
- **loop detection**—This is required for loosely routed LSPs only, also supported for repathing.
- **multiprotocol support**—This supports any type of protocol.
- **management**—LSP ID identifies each LSP, thereby allowing ease of management to discrete LSPs.

- **record route objects**—These provide the ability to describe the actual setup path to interested parties.
- **path preemption**—This is the ability to bump or discontinue an existing path so that a higher priority tunnel may be established.

CR-LDP

CR-LDP builds upon LDP, which is already part of MPLS. Although it is not as mature as RSVP, it does not require the implementation of an additional protocol. It uses existing message structures and only extends as necessary to implement traffic engineering. As with TE-RSVP, CR-LDP supports strict and loose explicitly routed LSPs. UDP is used for discovering MPLS peers and TCP is used for control, management, label requests, and mapping.

Figure 2. Example of a Strict Explicitly Routed CR-LDP LSP



In *Figure 2*, a strict CR-LDP LSP has been established between ingress and egress LERs in a service provider network. The path has been predetermined for both ingress and egress and is limited to two specific LSRs. The label requests have been passed down to each downstream device in a hop-by-hop fashion to the egress LER, and mapping has been passed upstream in similar fashion to the ingress LER. As shown in *Figure 2*, the explicit routed path can be so precise as to stipulate the specific LER and LSR IP addresses to be used. This is advantageous, as a particular traffic type (such as voice or VPN) can be matched to the optimal path to leverage bandwidth and prioritization.

CR-LDP traffic-engineering extensions to the LDP feature set are comprehensive and fairly well defined.

- **QoS and traffic parameters**—These offer the ability to define edge rules and per-hop behaviors based on data rates, link bandwidth, and weighting given to those parameters.

- **path preemption**—This is the ability to set prioritization to allow or not allow preemption by another LSP.
- **path reoptimization**—This offers the capability to repath loosely routed LSPs based on traffic pattern changes and includes the option to use route pinning.
- **failure notification**—Upon failure to establish an LSP, this is the notification provided on TCP with supporting failure codes.
- **failure recovery**—These are mapping policies for automatic failure recovery at each device supporting an LSP.
- **loop detection**—This is required for loosely routed LSPs only; LDP already supports loop detection.
- **multiprotocol support**—This supports any type of protocol.
- **management**—LSP ID identifies each LSP, thereby allowing ease of management to discrete LSPs.

3. A Technical Discussion of MPLS Traffic-Engineering Initiatives

Label Switching

A session or flow is defined as a series of IP datagrams that share a common destination IP address and traffic-engineering characteristics. Both CR–LDP and TE–RSVP provide a mechanism to map these sessions to LSPs traveling through the network. At the ingress point, the LSR assigns a label, corresponding to an LSP, to each IP datagram as it is transmitted toward the destination. Thereafter, at each corresponding hop, the label is used to forward the packet to its next hop. Both CR–LDP and RSVP create LSPs by first sending label requests through the network hop-by-hop to the egress point. At each hop, the MPLS–enabled router uses the label and its corresponding IP header information to program the hardware (or firmware) to switch the frame to its next hop. While the actual algorithms may be different, the end result of both signaling protocols is to establish an internal cross-connect from the ingress interface to the egress interface inside the LSR.

Scalability (Hard versus Soft State)

Today's routers must receive and process greater and greater numbers of frames. To accomplish this and not cause data bottlenecks, less time must be spent processing each IP frame. As mentioned above, label switching decreases the time required to analyze each IP datagram as it forwards it through the router. Some additional overhead is incurred while creating, maintaining, and destroying the information needed to establish the switched paths. However, this is minimal compared to traditional IP header processing. In networking jargon, the CR-LDP LSP setup is referred to as hard state. This means that all the information is exchanged at the initial setup time, and no additional information is exchanged between routers until the LSP is torn down. When the network management system or other entity determines that the LSP is no longer needed, messages must be exchanged notifying all routers that the resources should be reclaimed. This reclamation process is infrequent and consumes minimal bandwidth and CPU resources.

Conversely, TE-RSVP is referred to as a soft-state protocol. After an initial LSP setup process that is similar to CR-LDP, refresh messages must be periodically exchanged between peers to notify the peers that the connection is still desired. If the refresh messages are not exchanged, a maintenance timer senses the connection is dormant and deletes the state information, returns the label, and reserved bandwidth to the resource pool and notifies the affected peers. The soft-state approach can be viewed as a self-cleaning protocol because eventually all dormant or expired resources are reclaimed. As with CR-LDP, RSVP usually explicitly destroys or tears down LSPs when the network management system decides that the switched path is no longer desired.

Opponents of RSVP point to the soft state refresh overhead as a fundamental weakness in the protocol and therefore not scalable. RFC2208 states that "the resource requirements for running RSVP on a router increase proportionally with the number of RSVP sessions. Supporting numerous small reservations on high-bandwidth links can easily overtax the router and is inadvisable." Contributors to the IETF have taken steps to address this by adding extensions to the RSVP protocol. These extensions combine summarization information with newly defined RSVP objects that are sent inside standard RSVP messages.

To reduce the volume of chatter between two nodes, an RSVP node can group a number of RSVP refresh messages into a single message. This message is sent to the peer router where it is disassembled and each refresh message is processed. This process is referred to as bundling. In addition to bundling messages, the MESSAGE_ID and MESSAGE_ID_ACK objects have also been added to the protocol. These objects are used to hold sequence numbers corresponding to previously sent refresh messages. While the peer router receives a refresh message with a nonchanging Message ID, it assumes that the refresh state is identical to the previous message. Only when the Message ID value changes must

the peer router interrogate the actual information inside the message and act accordingly. To further enhance the summarization process, sets of Message IDs can be sent as a group to the peer router in the form of summary messages.

While this strategy will substantially decrease the time spent exchanging information between peer routers, it does not eliminate the computing time required to generate and process the refresh messages themselves. Time must still be spent checking timers and querying the state of each RSVP session. In short, the criticism surrounding the scalability of RSVP has been addressed, although not completely solved.

While CR-LDP appears to have the advantage of hard state for scalability, it is not without its own set of unique challenges. CR-LDP is based on the concept of an LDP entity. Once two LDP peers or entities have discovered each other, a TCP/IP session is established between them. From that point on, all control plane messages used to establish and maintain LSPs must travel through this reliable transport. As currently written, the MPLS specification requires that all LSPs associated with a particular session must be destroyed if the TCP session is torn down or fails. If hundreds or perhaps thousands of LSPs have been previously established between two LSRs, the impact to the network can be substantial. Members of the IETF are working on proposing solutions to this issue. Conversely, because an RSVP tunnel is a separate entity unto itself, any catastrophic change in its session state is local to itself.

Security and Reliability

As mentioned above, one of the benefits of the MPLS architecture is its well-defined separation of the routing decisions and forwarding of the data. Once the path has been established and the data is being forwarded (or switched) in the device's hardware, the frame is no longer promoted up to the upper layers and visible to the software. There is minimal chance that unauthorized individuals will be able to sniff the data or redirect the flow from its intended destination. Data is only allowed to enter and exit the LSP at locations authorized and configured by the MPLS control software (control plane). This minimizes the possibility of certain types of spoofing and flooding attacks possible in non-MPLS-controlled IP networks. Furthermore, CR-LDP uses a TCP/IP connection, thus offering a reliable and more secure connection between peers. The LDP and TE-RSVP specification also supports the use of MD5 signature password support to further ensure that the TCP/IP session is secure.

The TCP/IP connection capabilities also offer timely error notification if there is a communication failure between peers. This notification can then be quickly reported to the local network management system so that appropriate actions can be initiated. The sensing MPLS router can then initiate LDP withdraw and release messages to peers so that recovery actions can begin in earnest.

RSVP, on the other hand, uses UDP and raw IP datagrams to communicate between peers. This raises two reliability concerns: vulnerability to security attacks and fast recovery. While IPSec and other encryption or authentication schemes can be used to guarantee valid RSVP PATH and RESV messages, spoofing attacks could impair performance of TE–RSVP. Furthermore, connection failure will only be detected after a TE–RSVP neighbor fails to receive a refresh message from one of its peers. Depending on the configured refresh time intervals, the detection could take seconds or possibly minutes before recovery actions can be initiated at the end nodes of the affected LSP.

Data Aggregation—Support for Fine- and Coarse-Grained Flows

From its inception, LDP has been designed to establish switched paths that service a single IP host or an aggregation of thousands. The term used for this capability in the MPLS documents is forward equivalency class (FEC). Each FEC is specified as a set of packets that are mapped to a corresponding LSP. An IP address prefix describing an entire IP subnet can be designated as the destination of the LSP or FEC. As such, all traffic with destination IP host addresses inside that one subnet can travel through a single LSP. CR–LDP combines this addressing flexibility with the concept of differentiated services. At ingress, the LSR can assign a certain set of traffic parameters or constraints to be applied to each packet as it traverses the network. The combined concepts of FECs and differentiated services make it possible to aggregate traffic at the core of carrier backbones.

RSVP on the other hand, was initially designed to offer reserved bandwidth capabilities to a single IP address. RFC2205 describes an RSVP session as defined by its triple: DestAddress, Protocol ID, DestPort. The term *microflow* is used to describe a session containing a single IP host address as its destination. Clearly, modifications were required to address the needs in the core of the network. Recently a new Internet draft has been submitted to the IETF MPLS Working Group to address this issue of supporting differentiated services in networks using TE–RSVP. More work is expected to follow.

Minor Differences of Interest

There are a number of minor differences between CR–LDP and TE–RSVP that are of interest. They are listed here to give the reader a better understanding of how to determine the best MPLS signaling protocol.

Upon discovering an LDP peer, a TCP/IP session is established for the reliable exchange of control information. Each peer submits its type and range of labels to be used to establish LSPs. The LDP session chooses an intersection of the two

ranges. If there is no set of labels that intersect, the session is torn down. With RSVP, there is no negotiation of label space; such space must be configured via network management. If the network is very large or contains a larger number of heterogeneous interfaces, the effort to configure the labels could be considerable.

CR-LDP can specify the source route for an LSP by including an explicit route TLV in the label request message. TE-RSVP also has the ability to supply the same routing information with its explicit route object. Both support the concept of loosely routed paths. These are paths that can pass through a given network where any number of nodes might serve as the transit node. Both TE-RSVP and CR-LDP will respond back to the ingress the success or failure of the setup. Both CR-LDP and TE-RSVP allow route pinning. This term refers to the ability to force an LSP to stay in place after setup and not be rerouted by preemption. In CR-LDP, the act of pinning can only take place at setup time, but RSVP can set up pinning by modifying the PATH messages at any time.

The RSVP record route object can be used to request that the list of nodes actually involved in the path setup be reported back to the ingress. This can assist the network administrator when gathering information on network status and troubleshooting. CR-LDP does not have any way to request the trace route for an established LSP. See *Tables 1* and *2* for an explanation of the similarities and differences between TE-RSVP and CR-LDP.

Table 1. Similarities between TE-RSVP and CR-LDP

Characteristics	CR-LDP	TE-RSVP	Comments
initiate setup	label request message	PATH message containing LABEL_REQUEST object	
setup accomplished	mapping message	RESV message	
differentiated services defined	DIFF-SERV_PSC TLV	DIFFSERV_PSC object	Both contain the DiffServ code point or DSCP information and are included in the setup request message.
support for point-to-multipoint LSPs	no	no	This is yet to be defined by the IETF.
source route capability	This is carried in EXPLICIT_ROUTE list TLV.	This is carried in EXPLICIT_ROUTE object.	Specify route used to set up switched path.

Table 2. Differences between TE-RSVP and CR-LDP

Characteristics	CR-LDP	TE-RSVP	Comments
development stage	new	old with extensions being added, support for legacy networks	RSVP objects being modified to be used in a MPLS environment
signaling transport	UDP for discovery, TCP for sessions	raw IP datagrams or UDP encapsulation for message exchange	nondeterministic failure detection with RSVP; TCP failure can have catastrophic impact on LSPs with CR-LDP
connection state	hard state	soft state	soft state said to be nonscalable; RSVP to support aggregation of refreshes (also known as refresh reduction)
reliability	failure produces proactive signaling action	dependent on soft-state timer response to detect failure	nondeterministic failure detection with RSVP
manageability	LSR, LDP, TE MIBs	modified RSVP and LSR MIBs	
extensibility	vendor-specific, opaque, and experimental TLVs	experimental objects	very similar in function
scalability	hard-state connections reduce session signaling overhead	requires refresh reduction, aggregation to minimize soft state overhead	
interoperability	well-defined support for most transports: ATM, frame relay, Ethernet	tunneling through ATM network must be manually configured	

4. Conclusion

MPLS was designed out of the need to address new connection-oriented needs of the new Internet. It is adapting and evolving to new technologies just as the IP protocol itself has been evolving over the past 30 years. As with all new protocols, there is still a fair amount of work to be accomplished. The need to support traffic-engineered routes in the Internet has required new extensions to

traditional IGP and EGP protocols such as OSPF and BGP. Advances in fiber optics are requiring more modifications to MPLS in the areas of routing as well as signaling.

Both CR–LDP and TE–RSVP provide very similar functionality for establishing traffic-engineered, labeled switched paths. Each has its strengths and weaknesses. While LDP is the younger of the two protocols, RSVP has been previously deployed and has operational experience. It is true that there have been extensive enhancements to RSVP in order to support the needs of MPLS. As both CR–LDP and TE–RSVP evolve they will offer more and more similar functionality.

Eventually, MPLS traffic engineering should evolve into a single entity that combines the best-of-breed attributes from both TE–RSVP and CR–LDP. In the meantime, any MPLS implementation by original equipment manufacturers (OEMs) developing LER or LSR platforms should consider supporting both TE–RSVP and CR–LDP to ensure interoperability.

It should be noted that the intrinsic value of having an entry point to providing policy-based management to the core is extremely compelling. The lure of MPLS and the benefits of IP circuit switching for latency-sensitive traffic is no longer a case of "the emperor's new clothes."

Self-Test

1. Historically, the Internet infrastructure and protocols were intended and optimized solely for data.
 - a. true
 - b. false
2. Acceptance of RSVP was universal.
 - a. true
 - b. false
3. To realize QoS, the architecture must be applied _____.
 - a. end to end
 - b. at the edge
 - c. at select network devices

4. With MPLS, IP routers must interrogate each IP header before forwarding to the next hop.
 - a. true
 - b. false
5. How many ways are there to implement an LSP within an MPLS network?
 - a. one
 - b. two
 - c. three
 - d. four
6. Which of the following accurately describes strict ER–LSP?
 - a. adaptive
 - b. follows a list of nodes using the actual addresses of each node
 - c. allows groups of nodes to act as one of the abstract nodes to traverse
7. Which of the following is true of using TE–RSVP?
 - a. A full implementation of RSVP must be run on each LER or LSR.
 - b. TE–RSVP is a hard-state protocol.
 - c. An LER or LSR only requires that the extensions be able to support MPLS–explicit routing.
 - d. TE–RSVP does not use UDP.
8. These provide the ability to describe the actual setup path to interested parties.
 - a. record route objects
 - b. loop detectors
 - c. path preemptors
 - d. traffic parameters

9. CR–LDP is referred to as a _____ protocol; TE–RSVP is referred to as a _____ protocol.
- a. soft-state; hard-state
 - b. soft-state; soft-state
 - c. hard-state; hard-state
 - d. hard-state; soft-state
10. CR–LDP and TE–RSVP both support the concept of loosely routed paths.
- a. true
 - b. false

Correct Answers

1. Historically, the Internet infrastructure and protocols were intended and optimized solely for data.
- a. true**
 - b. false
- See Topic 1.
2. Acceptance of RSVP was universal.
- a. true
 - b. false**
- See Topic 1.
3. To realize QoS, the architecture must be applied _____.
- a. end to end
 - b. the edge
 - c. select network devices
- See Topic 2.

4. With MPLS, IP routers must interrogate each IP header before forwarding to the next hop.

a. true

b. false

See Topic 2.

5. How many ways are there to implement an LSP within an MPLS network?

a. one

b. two

c. three

d. four

See Topic 2.

6. Which of the following accurately describes strict ER–LSP?

a. adaptive

b. follows a list of nodes using the actual addresses of each node

c. allows groups of nodes to act as one of the abstract nodes to traverse

See Topic 2.

7. Which of the following is true of using TE–RSVP?

a. A full implementation of RSVP must be run on each LER or LSR.

b. TE–RSVP is a hard-state protocol.

c. An LER or LSR only requires that the extensions be able to support MPLS–explicit routing.

d. TE–RSVP does not use UDP.

See Topic 2.

8. These provide the ability to describe the actual setup path to interested parties.

a. record route objects

- b. loop detectors
- c. path preemptors
- d. traffic parameters

See Topic 2.

9. CR-LDP is referred to as a _____ protocol; TE-RSVP is referred to as a _____ protocol.
- a. soft-state; hard-state
 - b. soft-state; soft-state
 - c. hard-state; hard-state
 - d. hard-state; soft-state

See Topic 3.

10. CR-LDP and TE-RSVP both support the concept of loosely routed paths.

- a. true
- b. false

See Topic 3.

Glossary

AS

autonomous system

ASIC

application-specific integrated circuit

ATM

asynchronous transfer mode

BGP

border gateway protocol

CoS

cost of service

CPU

central processing unit

CR–LDP

constrained-based label distribution protocol

EGP

exterior gateway protocol

ER–LSP

explicitly routed label-switched path

FEC

forward equivalency class

HTTP

hypertext transfer protocol

IETF

Internet Engineering Task Force

IGP

Internet Gateway Protocol

IP

Internet protocol

LAN

local-area network

LER

label edge router

LSP

label-switched path

LSR

label switch router

MPLS

multiprotocol label switching

OSPF

open shortest path first

PSTN

public switched telephone network

QoS

quality of service

RFC

request for comment

RIP

routing information protocol

RSVP

resource reservation protocol

SLA

service-level agreement

TCP

transmission control protocol

TE-RSVP

traffic engineering resource reservation protocol

UDP

user datagram protocol

VPN

virtual private network