

Advanced Topics in MPLS-TE Deployment

Virtual path capability and the capacity to engineer precise traffic in a core network have driven Multiprotocol Label Switching (MPLS) towards becoming a standard within service provider core networks.

This paper introduces MPLS and Traffic Engineering, including a summary of achieving Resiliency with the technology. It also addresses the integration of QoS and MPLS. There is a detailed configuration / topology provided in the Annex section, which serves as a reference.

Introduction

Motivation for MPLS

The explosive growth of the Internet presents a serious challenge to service providers and equipment suppliers in terms of the tremendous escalation in traffic. There is also an increased demand to create differentiated IP services and bring these to market quickly is also increasing. Other challenges include the cost of mapping IP over layer 2 networks, as well as difficulties in identifying better network utilization and fault handling.

Service providers already address these issues in several ways: increase bandwidth and/or the number of powerful routers, exploit QoS to better shape and police traffic, utilize available bandwidth more effectively.

Cisco IOS MPLS fuses the intelligence of routing with the performance of switching. It provides significant benefits to networks with a pure IP infrastructure and to those with IP and ATM, or a mix of other layer 2 technologies.

MPLS technology is key to scalable virtual private networks (VPNs) and end-to-end quality of service (QoS), enabling efficient utilization of existing networks to meet future growth and rapid fault correction of link and node failure.

Multi Protocol Label Switching Overview


Unlike legacy routing, MPLS uses labels to forward traffic across the MPLS domains. When packets enter the MPLS domain, labels are imposed on the packets, and the label (not the IP header) determines the next hop. Labels are removed at the egress of the MPLS domain.

When a labeled packet arrives at a Label Switching Router (LSR), the incoming label will determine the path of this packet within the MPLS network. MPLS label forwarding will then swap this label to the appropriate outgoing label and send packets to the next hop.

These labels are assigned to packets based on grouping or forwarding equivalence classes (FECs). Packets belonging to the same FEC receive the same treatment. This MPLS lookup and forwarding system allows explicit control routing, based on destination and source address, allowing easier introduction of new IP services.

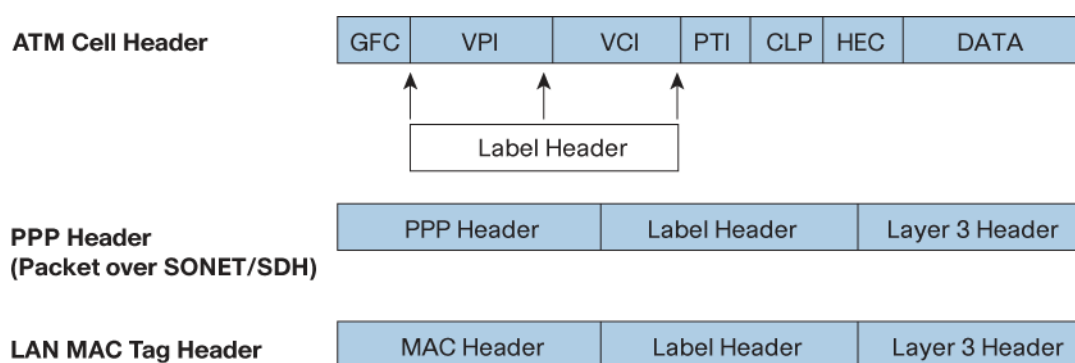
Label switching has traditionally been used as a forwarding scheme. ATM uses the same techniques to forward packets via virtual path identifier/virtual channel identifier (VPI/VCI) regardless of the payload (IP, other). The MPLS standard, published by the Engineering Task Force (IETF), evolved from the Cisco Tag Switching implementation.

The IETF recommendation for label switching is based on 32bit shim headers consisting of: Label (20bits), Exp (3bits), Stack (1bit), TTL (8bits).

Figure 1. MPLS Shim Headers

Label and TTL do not require extensive explanations. The stack bit is used to indicate that the bottom of a stack is reached: this is useful when multi-stacking labels (i.e., MPLS-VPN or link protection). The Exp bits (a.k.a. experimental) are mainly used to carry information relative to Quality of Service.

The label is added between the layer 2 and the layer 3 header (in a packet environment) or in the VPI/VCI field in ATM networks.

Figure 2. Encapsulation of MPLS Labeled Packet

MPLS Traffic Engineering

Although MPLS label switching provides the underlying technologies in forwarding packets through MPLS networks, it does not provide all the components for Traffic Engineering support such as traffic engineering policy.

Traffic Engineering (TE) refers to the process of selecting the paths chosen by data traffic in order to facilitate efficient and reliable network operations while simultaneously optimizing network resource utilization and traffic performance. The goal of TE is to compute path from one given node to another such that the path does not violate any constraints (bandwidth/administrative requirements) and is optimal with respect to some scalar metric. Once the path is computed, TE is responsible for establishing and maintaining forwarding state along such a path.

Traffic Engineering Components

A router capable of supporting MPLS is known as Label Switching Router (LSR). The LSR, found just before the last LSR in the MPLS clouds, is known as the penultimate hop. The end-to-end MPLS path is known as Label Switched Path (LSP). LSP is originated at the head-end router and terminates at the tail-end router.

The existing Interior Gateway Protocols (IGP) are not adequate for traffic engineering. Routing decisions are mostly based on shortest path algorithms that generally use additive metric and do not take into account bandwidth availability or traffic characteristics.

The easiest way to provide such features would be to use an overlay model, which enables virtual topologies on top of the physical networks. The virtual topology is constructed from virtual links that appear as physical links to the routing protocol. Further, the overlay model should be able to provide: (1) constraint based routing, (2) traffic shaping and traffic policing functionality, (3) survivability of the virtual links... These capabilities allow easy movement of traffic from an over subscribed link to an underused one.

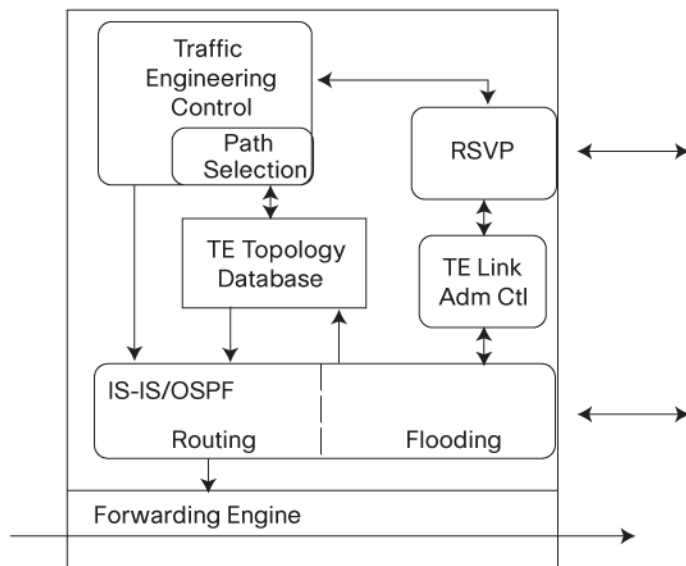
MPLS is the overlay model used by Traffic Engineering. It provides:

1. Explicit label switched paths which are not constrained by the legacy destination based traffic forwarding (as featured in all the existing IGPs)
2. LSPs that can be efficiently maintained
3. Traffic trunks that can be instantiated and mapped into LSPs
4. A set of attributes that can be associated with traffic trunks
5. A set of attributes that can be associated with resources that constrain the placement of LSPs and traffic trunks across them
6. MPLS allows for both traffic aggregation and disaggregation whereas destination based IP forwarding allows only aggregation. "Constraint based routing" and trunk protection can be integrated easily to MPLS.

These components should be available to support TE:

- Information distribution—sends information about network topology and constraints pertaining to links (i.e., bandwidth)
- Path selection algorithm—computes and selects best paths that obey the constraints
- Route setup—Resource Reservation Protocol TE (RSVP-TE) extension for signaling LSPs setup
- Link Admission Control: decides which tunnel may have resources
- TE control: establishes and maintains trunks
- Forwarding data across the path

Figure 3. MPLS-TE System Block Diagram (Head Router)



Information Distribution

TE relies on the Integrated Gateway Protocol (IGP) protocol to distribute/flood link-related information resource availability, including: bandwidth per priority (0-7) [maximum link bandwidth, maximum reservation bandwidth, reserved bandwidth], link attributes, TE specific link metric, resource class attributes for a link.

IGP has been enhanced to include three new flooded Type Lengths Values (TLV) messages:

- Reservable Bandwidth at each priority (0-7)
- Link Color assignments
- Traffic engineering assigned metrics

A forth TLV message is related to the reservable bandwidth used with by DiffServ aware Traffic Engineering. These new flooded announcements come from traffic engineering provisioning performed on each LSR. Information distribution can occur periodically (timer-based) or it can be event driven (i.e., change in available bandwidth, link configuration, LSP setup failure).

Constrained Based Routing Algorithm (a.k.a. CBR)

LSRs use IGP extensions to create and maintain a TE Link State database (TE-LSDB). This is very similar to the TE-LSDB used by Open Shortest Path First (OSPF)/ Intermediate System to Intermediate System protocol IS-IS: it contains the TE network topology that is updated by IGP flooding whenever a change occurs (establishment of new LSP, change of available bandwidth).

Constraint based algorithm is used to find the best path for an LSP tunnel. It is targeted by the trunk's head-end (i.e. the originator of the tunnel) only when (1) a new tunnel is requested, (2) the current LSP of an existing trunk has failed, (3) to re-optimize an existing trunk.

The following information is considered:

- Attributes of traffic trunks originated at the head-end router (manually configured)
- Attributes associated with resources (IS-IS/OSPF)
- Topology state information (IS-IS/OSPF).

An overview of the algorithm is given hereafter:

- Step 1. Prune off the links that do not have insufficient resources (bandwidth) and violate policy constraints from TE-LSDB
- Step 2. Run Dijkstra on the remaining topology (use IGP metrics or TE metrics if specified)
- Step 3. Select the path with the highest minimum bandwidth, then the ones with the smallest hop-count.

As a result, Constrained Based Routing will give an explicit route known as "Constrained Shortest Path" consisting of a list of {interface/IP address} or {loopback for unnumbered interface}.

How is an LSP Tunnel Set Up? (Signaling the Desired Path)

The head-end router will begin the process of signaling the Constrained Shortest Path. Path (also known as LSP) establishment is based on RSVP-TE messages.

Note: CR-LDP is another feasible protocol, but this paper will focus on RSVP-TE. It is the most widely used protocol, and is available in Cisco IOS Software.

RSVP-TE is an extension of the well-known RSVP protocol (defined in RFC 2205). The protocol defines several messages but this paper will focus on the two that are used for LSP/Path establishment: RSVP-TE PATH and RSVP-TE RESV messages.

RSVP-TE PATH Messages

The Head-end router will transmit RSVP PATH messages. The PATH message will follow the routers listed in the Constrained Shortest Path. The PATH messages contains objects that are used by each LSR along the path:

Label Request Object

- ERO-explicit route object: identifies route from head-end to tail end (in this case it is the Constrained Shortest Path)
- RRO-Record Route Object: keep track of the list of LSR transversed by the PATH message
- Session Attribute: controls LSP setup priority, Holding Priority, preemption, use of local link protection (flag 0x01)...

Session Object: assigns a global label switched path tunnel ID

Sender_: Transmission Specification (Tspec), sent to tail-end to indicate desired reservation characteristics.

The PATH message will follow the routers listed in the Explicit Route Object. At each hop, RRO is updated with the name (or IP address) or the visited LSR.

The Session Attribute contains among other things the setup and holding priorities for the LSP, which is useful when the requested bandwidth is not available at the priority specified in the Setup priority. In this case, if the requested bandwidth is available but is in use by lower priority sessions, then lower priority sessions may be pre-empted to free the necessary bandwidth.

The session Attribute contains also the support of local protection flag.

RSVP-TE RESV Messages

The RESV message is sent back by the tail end upon reception of the PATH message. The tail end must initiate the label distribution process. The following objects are available:

- Label Object: contains the label to be used
- Record Route Object: contains list of LSRs to route RESV message back to head-end
- Style Object: controls label reservation style e.g. Shared explicit...
- Session Object: copied from the PATH message, global label switched path ID

At each LSR, an RESV message identifies and assigns the label value to the incoming interface. Each LSR must allocate a local label (delivered via Label Object) for the next downstream LSR (identified per RRO).

Two control planes are involved in the path setup and any MPLS operation: Trunk Admission Control and Link Admission Control.

Trunk admission control will determine if resources are available along a Label Switched Path. It is also responsible for tearing down existing LSPs with lower priority when necessary. Further, trunk admission control triggers IGP information distribution when there is a change of resource.

Link Admission Control is used within the PATH message (cf. bandwidth reservation). If bandwidth is available, this bandwidth is moved to a waiting pool until a RESV message is received. Otherwise, a path error message will be sent upon reception of RESV [10].

For more information regarding the above messages and a detailed explanation of all the supported objects refer to RSVP-TE specification [10]. Further, a step-by-step LSP setup is shown in Appendix A

Once the Label Switched Path is established, the next step is to forward traffic across this LSP. With Cisco IOS MPLS TE, there are currently two available possibilities:

1. Via Policy Based Routing (including static routes pointing on the tunnel for any destination behind the tail end.)
2. Automatic via the use of Cisco IOS MPLS TE Autoroute Announce.

Autoroute Announce Feature (IGP Shortcut)

Cisco IOS MPLS Autoroute Announce installs the routes announced by the tail-end router and its downstream routers into the routing table (forwarding table) of the head-end router as directly reachable through the tunnel.

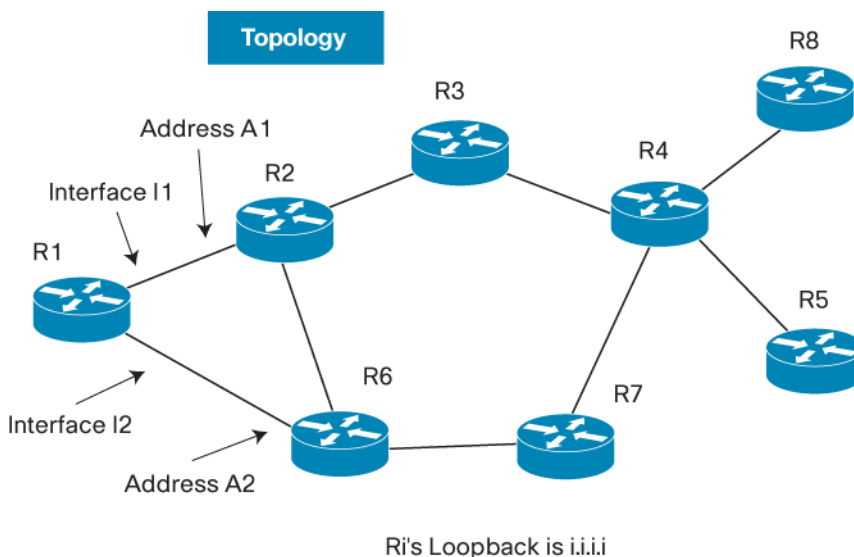
The Constrained Based Routing Algorithm allows MPLS TE to establish a Label Switch Path from the head-end to the tail-end node. By default, those paths will not be announced to the IGP routing protocol. Hence, any prefixes/networks announced by the tail end router and its downstream routers would not be “visible” through those paths.

For every MPLS TE tunnel configured with Autoroute Announce, the link state IGP will install the routes announced by the tail-end router and its downstream routers into the RIB. Therefore, all the traffic directed to prefixes topologically behind the tunnel head-end is pushed onto the tunnel.

To have a better understanding of this feature, consider an example with and without Autoroute Announce enable. A detailed description of the algorithm is given.

Consider the topology of Figure 4. For the sake of simplicity, assume that R1's loopback address is i.i.i.i.

Figure 4. Topology without Tunnels



The corresponding routing table on Router R1 with normal IGP and no MPLS TE looks like the following.

Figure 5. R1 Routing Table

Routing Table			
Dest	O Intf	Next Hop	Metric
2.2.2.2	I1	A1	1
3.3.3.3	I1	A1	2
4.4.4.4	I1	A1	3
	I2	A2	3
5.5.5.5	I1	A1	4
	I2	A2	4
6.6.6.6	I2	A2	1
7.7.7.7	I2	A2	2
8.8.8.8	I1	A1	4
	I2	A2	4

Considering the same topology as in Figure 4, now let us introduce two MPLS Traffic Engineering tunnels T1 and T2 respectively. Tunnel T1 (resp. T2) will originate in R1 and its tail end is R4 (resp. R5).

MPLS TE Autoroute Announce will be enabled on the two tunnels. Similarly, R1 routing table entries are given in Figure 7.


Figure 6. Topology with Tunnels

Figure 7. R1 Routing Table with Autoroute Announce

Routing Table			
Dest	O Intf	Next Hop	Metric
2.2.2.2	I1	A1	1
3.3.3.3	I1	A1	2
4.4.4.4	T1	R4	3
5.5.5.5	T2	R5	4
6.6.6.6	I2	A2	1
7.7.7.7	I2	A2	2
8.8.8.8	T1	R4	4

The routing tables (Figure 5 and Figure 7) demonstrate that R4 and R5 are directly reachable through tunnel T1 (resp. T2) with MPLS TE Autoroute Announce. Similarly, R8 is now reachable through the tunnel T1 via R4 instead of the “physical” connection.

Without Cisco MPLS TE Autoroute Announce, even though Tunnel T1 is up, route to R8 is done via the “physical” connection (as in Figure 5).

Autoroute Announce Algorithm

Step 1. Run a normal Dijkstra on the topology without tunnel

Step 2. Go through the tree and add a link for each tunnel

a new link's metric = {metric(igp) +/- relative} XOR {absolute} where metric(igp) = metric of the shortest-path as computed by IGP without tunnel

Step 3. Go through the tree and, for each node, prune the link leading to this node and that are not on the best path to that node

If we have 2 links with the same metrics, a link representing a tunnel ending on the node wins

Step 4. Add the IP-prefix leaves to the tree

Step 5. Go through the tree, for each leaf, add an entry in the forwarding table. The metric used for an entry is equal to the sum of the metrics of the links used to reach this leaf from the root of the tree.

Exception: if the path through the tree to a leaf goes via a link that represents a tunnel configured with an absolute metric, then the metric for this leaf (this prefix) is just the absolute metric.

As per the current algorithm, the link state IGP will install the routes announced by the tail-end router and its downstream routers into the RIB. Further, any MPLS traffic engineering tunnel change will be announced into the IGP and will trigger a full SPF calculation in order to make the adjustment to the routes that are associated with the changes. From a layer 3 perspective, most of the MPLS tunnel information change is confined to the tunnel tail end router and its downstream routers. For scalability reason, a new optimized algorithm was introduced: based on the tail-end configuration, either it has children (a.k.a. downstream router), a full or partial SPF calculation is targeted upon changes.

Resiliency, Tunnel Restoration

Network reliability is mandatory in high-speed networks. Disruption can occur due to several reasons: congestion along the LSP, failed link, failed node, administrative change in the LSP. One of the most appealing features of MPLS TE is the possibility to provide non disruptive traffic across the LSP. In case of outage, the upper level application will not notice any service disruption.

Path protection can be achieved at multiple layers of the protocol stack:

- Physical layer (e.g. SONET with Automatic Protection Switch)
- IP (e.g. Routing protocol, IGP, BGP, changes next-hop if there is a change of topology)
- MPLS (Performed by the Head-end upon change of topology)

With MPLS TE, several options exist for path restoration:

- Head-end reroute
- Fast Reroute (link protection)
- Fast Reroute (node protection)

Head-end Reroute

Since the path taken by a trunk is determined by the LSR at the start of the MPLS path (head-end), it seems natural that path restoration is performed by the head-end.

Head-end reroute is mainly targeted by two events: notification by RSVP-TE that the path cannot be maintained (e.g. congestion) or notification by the IGP of a change of topology. (A third one could be a requested action from the CLI to optimize the path).

Upon reception of one of those events, the head-end router will construct a new TE database after pruning the faulty links or area of congestion. The head-end will then re-signal a new path with the “shared-explicit” reservation style. The “shared-explicit” reservation style [10] allows the new LSP to use some of the links used by the previous LSP without double counting for reserved bandwidth.

Once the new LSP is up, the head-end router will change its forwarding table and the original LSP is torn down. Head-end rerouting is best suited for path re-optimization versus fault recovery. Head-end notification is dependent on how fast IGP or RSVP-TE will notice the faulty link or area of congestion and flood this information back to the head-end.

As a guideline, considering IS-IS as the IGP in use, the minimum amount of time in notifying the head end would be: $\text{Time(IGP reaction)} + \text{Time(RSVP-TE signaling)}$ where Time (IGP reaction) can be broken down as: time to detect link failure, time for LSDB change, time to flood new LSDB. An average value would be around 2-3s (under optimum tuning). Computing and signaling time for the new Label Switch Path should also be considered. Therefore, traffic restoration will not happen in less than 2-3s.

Unfortunately, these establishment times for the backup LSP can be too long for some applications of MPLS TE. MPLS Fast Reroute feature provides a mechanism for rapidly repairing (under 50 ms; actual failover time may be greater or less than 50ms, depending on the hardware platform, the number of TE Tunnels and/or Network prefixes) an LSP by routing along a detected failure in order to minimize the length of service interruption experienced while the head-end attempts to establish a replacement LSP.


MPLS TE Fast Reroute substantially reduces the amount of packets lost during a failed link. FastReroute tunnel restoration can be used either in protecting an end-to-end LSP (FRR Path protection) or a local link (FRR Link protection)/node (FRR node Protection) within an LSP path.

This paper will focus on FRR link protection.

Cisco FRR Link Protection

FRR establishes a procedure that allows rerouting around a failed link in the event of a link failure. The LSP is routed to the next-hop using a pre-configured backup tunnel. The backup tunnel must be configured so the LSP can get to the next hop downstream router without traversing the protected link. FRR link protection does not offer any node resiliency support, but it does protect a specific link.


Figure 8. Link Protection to Next Hop, Label/Packet flow without Fast Reroute



A primary LSP (Tunnel 0) is configured from R1 to R6 (Figure 8). The link between R2 and R3 via FRR should be protected. A backup “tunnel,” which will only be used when a failure occurs, is statically configured on R2 going to R3 across R4. When configuring this particular tunnel, the links used by the primary tunnel cannot be used.

The packet flow for LSP Tunnel 0 is shown on Figure 8. R1 will impose the label on any packets destined to R6. R2, R3 perform label swapping, while R5 completes its penultimate hop popping (php).

Upon failure of the link between R2 and R3, R2 will immediately swap the traffic destined to R6 across the backup LSP, hence assuring a substantially reduced amount of packet loss. The all operation will take around 50 ms (actual failover time may be greater or less than 50ms, depending on the hardware platform, the number of TE Tunnels and/or Network prefixes).

Figure 9. Label/Packet Flow when using Fast Reroute

Upon link failure (Figure 9), R2 rewrites all the incoming labels to go over the backup link. As a result, R2 will reroute traffic around the failure by sending packets into a different downstream interface.

A packet destined for R6 uses a two-level label stack as it goes through the backup link: a backup label, then a label that is used on the link (R2, R3), if it does go through this link (i.e. no link failure). Upon reception of the packet, R4 will pop the “backup label” and send the packet to R3. When the packet is received at R3 through the interface (R4, R3), it has the same label as if it was received from the failing link (R2, R3), or primary tunnel.

This procedure is only possible if the downstream router supplies a label from the “Global Allocation Pool” and it is rerouted at R2 to an interface that is also using Global Label. MPLS Global Label allocation means that there is one label space for all interfaces in the router (i.e.: label 15 from one interface is treated the same as label 15 coming on another interface).

Fast Reroute is initiated for an LSP when the feature is enabled for the associated LSP tunnel as a result of a configuration command on the head-end. The head-end router is responsible for informing all routers along the LSP’s path that the LSP is requesting protection, by setting the session attribute flag bit (to 0x01) in the RSVP PATH message.


The LSP tunnel head-end control module will keep RSVP informed of the status of the Fast Reroute attribute for all active LSPs. When the RSVP module in a Label Switch Router (LSR) [other than tail end] along the LSP’s path learns that the LSP should be protected, it will initiate local Fast Reroute protection procedure to protect the LSP against possible failure of the immediate downstream link.

Upon link failure, all protected LSPs switch to the backup path. FRR performs the operations to prevent the downstream routers (still along the path in use by the LSP) from tearing down the LSP, if the failure is also detected downstream.

A new LSP will be re-signaled by the head-end, either by the reoptimization process, which is automatically running at specific time interval, or upon reception of an RSVP PathErr message at the head-end.

Backup to the Next Next Hop (a.k.a. Node Protection)

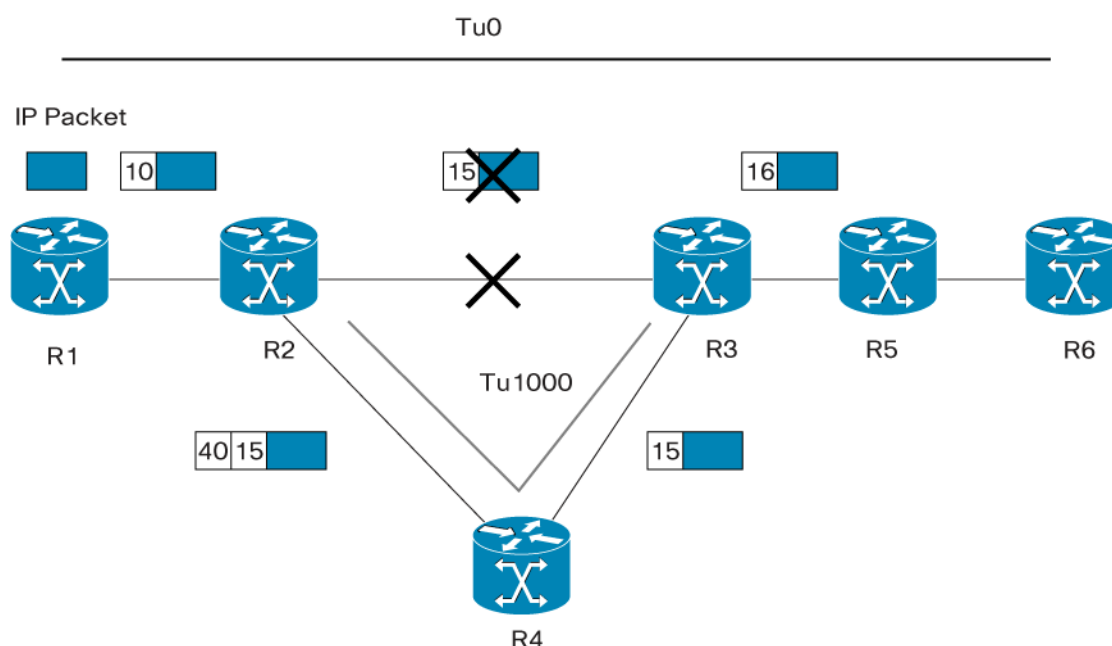
Figure 10. Link Protection to the Next Next Hop



In this scenario (Figure 10), consider a primary LSP: R1-R2-R3-R5-R6. Protecting the node R3 is achieved either through a) R2-R4-R5-R3 or b) R2-R4-R5. In the case of a), the backup path being R1-R2-R4-R5-R3-R5-R6, this is the same as the backup tunnel to the next-hop. Notice though, that the (R3, R5) link will carry the traffic to R6 twice (i.e. double reservation). This solution is not optimal.

Consider Case b) with a backup path of R1-R2-R4-R5-R6 known as “backup to the next next hop” a.k.a. Node Protection.

Node protection is actually more complex than Link Protection as R2 would need to be aware of the label used on the link R3-R5, as R5 expects to receive the correct label through the backup links (Figure 11). The use of an extended Route Record object will allow R2 to learn this label.

Figure 11. Node Protection, Label Flow

[Node Protection in Cisco IOS Software is not currently supported but will be available very soon.]

When to Use the Backup Schema?

FRR Link Protection and Node Protection assume the use of Global Label, as the next hop or next next hop (for Node protection) needs to know the previously used label. Link and Node Protection are mainly designed to backup sensitive link/device servicing several LSPs.

Node Protection also offers the capability to create multiple backup tunnels to the same destination, so traffic can be load balanced across those tunnels.

The objective of Path Protection is to achieve end-to-end protection, so there is no need for Global Label.

Implementing Traffic Engineering

Before MPLS TE is deployed, the traffic load pattern of the network must be evident.

This optimizes the network because MPLS TE selects paths that may have better latency and more available bandwidth. However, the Integrated Gateway Protocol (IGP) may have not necessarily chosen these paths. Typically, this is done through a modeling tool in which a service provider enters the traffic load parameters and network topology, and the tool suggests alternative “best” paths.

The Cisco IOS MPLS AutoBandwidth Allocator speeds installation of MPLS TE tunnels by allowing service providers to set up tunnels with arbitrary bandwidth values, and then dynamically and automatically adjust the bandwidth based on traffic patterns without traffic disruption. Cisco IOS MPLS AutoBandwidth makes the initial set up and re-provisioning of tunnels less complicated, allowing service providers to easily adopt MPLS TE as a traffic management solution.

Cisco AutoBandwidth Allocator

Cisco IOS MPLS AutoBandwidth allocator automatically adjusts the bandwidth size of an MPLS Traffic Engineering (TE) tunnel, based on how much traffic is flowing through the tunnel. This automates the tasks of monitoring and re-configuring bandwidth for an MPLS TE tunnel.

For every MPLS TE tunnel configured for Cisco MPLS AutoBandwidth, the average output rate is sampled based on various configurable parameters. The tunnel bandwidth is then re-adjusted automatically, based upon the largest average output rate noticed during a certain interval or a configured maximum bandwidth value. Practically, Cisco MPLS AutoBandwidth allocator monitors the X minutes (default X = 5 min) average counter, keeping track of the largest X average over some configurable interval Y (default = 24 hours), and then re-adjusting a tunnel bandwidth based upon the largest X average for that interval.

After the Y interval has expired, the initial largest X average is cleared in order to record the new largest X average over the next Y interval.


The X interval is user configurable and is a global parameter that applies to all tunnels configured for AutoBandwidth. The interval Y can be set on a per tunnel basis, and a larger Y interval gives an X average that more closely resembles the actual peak in the traffic bandwidth.

When re-adjusting the LSP with the new bandwidth constraint, a new Resource Reservation Protocol for Traffic Engineering (RSVP-TE) PATH request is generated, and if the new bandwidth is not available, the last good LSP will continue to be used. The network experiences no traffic interruptions.

Consider the following traffic pattern (bandwidth usage versus time) for a given traffic engineering tunnel (Figure 12). Assume that the bandwidth value used to signal the LSP is far above the highest bandwidth peak observed during the period (over provisioning).

The MPLS TE AutoBandwidth allocator allows the network operator to automatically adjust bandwidth need based on the observation of the highest average bandwidth used during Y interval. At period N, a red dot represents the highest bandwidth for this period. This value will then be used at period N+1 to signal the adjusted LSP; the network provider is able to optimize traffic and bandwidth management

Figure 12. Illustration of Cisco MPLS Bandwidth Allocator Adjusting Bandwidth Over Time



Integrating MPLS and QoS

There are two architectures for adding QoS capabilities to today's network: Integrated Services (IntServ) and Differentiated Services (DiffServ). Integrated Services maintains an end-to-end QoS for an individual or group of flows with the help of a resource reservation protocol (RSVP). [IntServ architecture described in RFC 1633]

DiffServ is one of the two QoS architectures for IP networks defined by the IETF. In this model, packets entering a DiffServ enabled network are grouped into a small number of classes. Each class has a color or mark associated with it (use of the DiffServ Code Point DSCP bits). This makes packet classification extremely scalable and assures appropriate bandwidth and delay guarantees in the network core. Each node within the core network is applied to different queuing and dropping policies on every packet, based on the marking that packet carries (Per Hop Behavior).

MPLS+DiffServ

In a MPLS+DiffServ architecture, packets marked with DiffServ Code Point will enter the MPLS network and per hop behavior is enforced by every LSR along the path. As LSRs do not have any knowledge of the IP header, per hop behavior needs to be achieved by looking at different information.

Two general approaches are used to mark MPLS traffic for QoS handling within an MPLS network.

In the first method, the DiffServ coloring information is mapped in the EXP (experimental) field of the MPLS shim header. This field allows up to eight different QoS marking versus 64 for DSCP. The packet scheduling (PHB) at each hop (in the MPLS clouds) is done based on the EXP. Label Switched Paths that use this approach are called E-LSPs, where QoS information is inferred from the EXP bits.

Alternatively, the label associated with each MPLS packets carries the portion of the DiffServ marking that specifies how a packet should be queued. The dropping precedence portion of the DiffServ marking is carried in the EXP bits (if an MPLS shim header is being used) or on fields available for this purpose on underlying technologies (CLP bit in ATM, DE bit for Frame Relay).

The ingress LSR examines the DSCP in the IP header (resp. CLP/DE for ATM/Frame Relay) and selects an LSP that has been provisioned for that QoS level. At the egress as the label is removed, the packet is sent to the next IP hop with its original DSCP. LSPs using this approach are called L-LSP where QoS information is inferred in part from the MPLS label.

For more information regarding the best choice of mapping (E-LSP vs. L-LSP) and a detailed description of each technique, please refer to [7] and [4].

Table 1. Comparison of E-LSPs and L-LSPs

E-LSPs	L-LSPs
PHB determined from Exp bits	PHB determined from label or {label, EXP/CLP} bits
No additional signaling is required	PHB is signaled at LSP setup (LDP, RSVP-TE, and so on)
EXP->PHB mapping is configured	Label->PHB mapping is signaled EXP/CLP ->PHB mapping is well known (used only for AF)
Shim header is required; E-LSP not possible on ATM links	Shim or link layer header may be used; suitable for ATM links
Up to 8 PHBs per LSP	One PHB per LSP except for AF

Traffic engineering does not differentiate among traffic types. To carry voice and data traffic on the same network, it may be necessary to account separately for the amount of voice traffic being transferred over the network, to provide the necessarily stricter QoS guarantees.

DiffServ Aware Traffic Engineering (DS-TE)

Cisco DiffServ aware Traffic Engineering (DS-TE) not only allows the configuration of a global pool for bandwidth accounting, it also provides a restrictive sub-pool configuration that may be used for high-priority network traffic such as voice or other applications.

Available bandwidth both on the global pool and in the sub-pool are advertised by IGP LSA or TLVs, ensuring each router keeps track of the available bandwidth when admitting new LSPs for voice or high-priority traffic. In this manner, service providers, depending on their service level agreement requirement, can choose to overbook lower-priority classes or even underbook higher-priority traffic to meet tight QoS requirements.

Diff-Serv Aware TE extends MPLS TE to perform constraint based routing (path computation) on a specific (restrictive) set of sub-pools where? bandwidth is dedicated to the high-priority traffics.

This ability to satisfy a more restrictive bandwidth constraint translates into the capability to achieve higher QoS (in terms of delay, jitter or loss) for the traffic using the sub-pool.

DS-TE involves extending OSPF and IS-IS so that the available sub-pool bandwidth at each preemption level is advertised in addition to the available global pool bandwidth at each preemption level. Further, DS-TE modifies constrained based routing to take this more complex advertised information into account, during path computation.

A typical use for DS-TE would be for toll bypass/voice trunking or leased line services emulation, where a point-to-point guarantee is needed in term of bandwidth/delay and jitter bounds.

For more information on Cisco DS TE please look up Cisco MPLS web site at:

<http://www.cisco.com/warp/public/732/Tech/mpls/>

ANNEX A: Step by Step LSP Set Up (RSVP-TE Signaling)


This section is an attempt to illustrate LSP setup when using RSVP-TE as described in [10]. For more information, please refer to RSVP-TE Internet draft [10].

Robert Raszuk's MPLS TE Networkers 2000 presentation [11] contains a very good introduction/overview on LSP path set up.

To illustrate LSP step-by-step establishment, consider the following topology with three routers: R1 would like to establish an LSP towards R3 with a certain number of constraints.


In this case, R1->Head router, R2->penultimate router, R3->Tail router.

Step 1. Tunnel Establishment PATH request originated from R1: note the may optimize flags (0x04) in the SESSION_ATTRIBUTE, and request for 2Mbps bandwidth. Further ERO is present; this path message will be forwarded towards its destination along the path specified in ERO.




Path:
 Common_Header
 Session(**R3-100,0,R1-100**)
 PHOP(**R1-2**)
 Label_Request(IP)
 ERO (**R2-1,R3-1**)
 Session_Attribute (**S(3),H(3),0x04**)
 Sender_Template(**R1-100,00**)
 Sender_Tspec(**2Mbps**)
 Record_Route(**R1-2**)

Step 2. Request received by R2, no change in the packet




Path State:
 Session(R3-100,0,R1-100)
 PHOP(R1-2)
 Label_Request(IP)
 ERO (R2-1,R3-1)
 Session_Attribute (S(3),H(3),0x04)
 Sender_Template(R1-100,00)
 Sender_Tspec(2Mbps)
 Record_Route(R1-2)

Step 3. R2 is forwarding the request to R3, PHOP and ERO have been updated



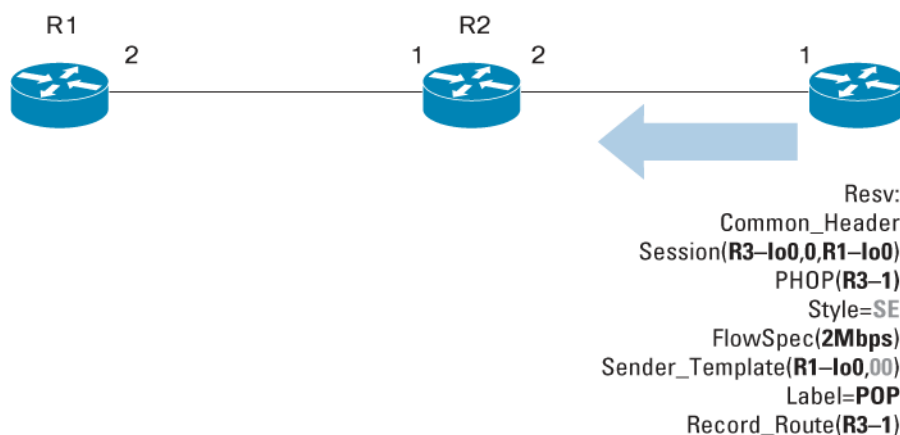
Path:
 Common_Header
 Session(R3-100,0,R1-100)
 PHOP(R2-2)
 Label_Request(IP) ERO (R3-1)
 Session_Attribute (S(3),H(3),0x04)
 Sender_Template(R1-100,00)
 Sender_Tspec(2Mbps)
 Record_Route(R1-2, R2-2)

Step 4. PATH request completed (received by R3)

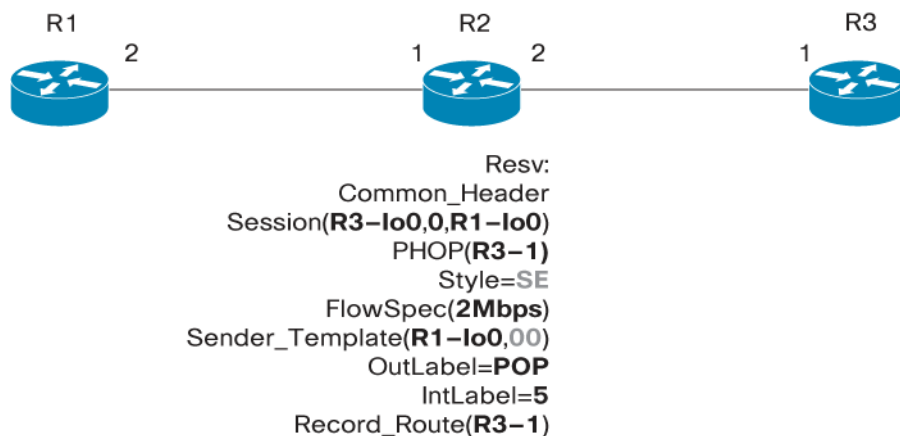


Path State:
 Session(R3-100,0,R1-100)
 PHOP(R2-2)
 Label_Request(IP)
 ERO ()
 Session_Attribute (S(3),H(3),0x04)
 Sender_Template(R1-100,00)
 Sender_Tspec(2Mbps)
 Record_Route(R1-2, R2-2,R3-1)

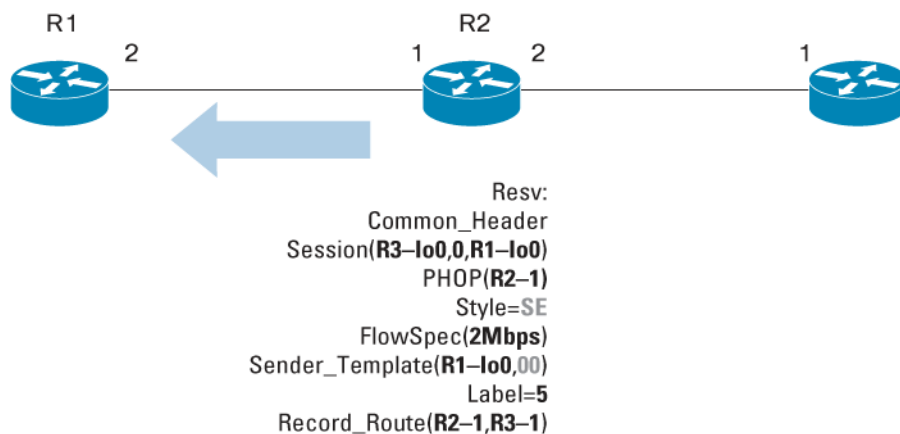
Step 5. RESV reply from R3 to R1, 2Mbps request granted, note also Label=POP i.e. R3 is the pultimate router in this LSP. Further, Reservation Style (STYLE) is set to Shared Explicit (SE) i.e. there is a single reservation on this link for the sender R1.




Step 6. R2 will forward any incoming packet from interface R2-1 to interface R2-2, at the same time R2 will POP the label for the outgoing packets.



Step 7. RESV message R2-R1, R2 requests R1 to use label=5 for any packets on the link R1-2/R2-1 destined to R3



Step 8. LSP is established, any outgoing packets on R1-2 will be tagged with Label=5. Next step is to send the traffic down this LSP either by static route, Autoroute announce or Policy based Routing.




Resv state:
 Session(R3-100,0,R1-100)
 PHOP(R2-1)
 Style=SE
 FlowSpec (2Mbps)
 Sender_Template(R1-100,00)
 Label=5
 Record_Route(R1-2,R2-1,R3-1)

ANNEX B: Configuration/Topology Example

Basic Traffic Engineering

Figure B1 MPLS Topology



Before configuring MPLS TE tunnels, all routers within the domain need to support MPLS, RSVP-TE and either IS-IS or OSPF with TE extension.

In this configuration, it is assumed that all routers participating in MPLS TE are within a single area or domain as far as IGP is concerned. Some work is underway to support Traffic Engineering in a multi-area environment.

Basic Traffic Engineering Requirements:

- CEF is necessary for all MPLS features.
 ("ip cef" or "ip cef distributed" distributed should be used for 7500)
- Loopback must be in IGP.
- Tunnel is always unidirectional.

- Cisco IOS will not route IP across an interface with no IP address.
- With IS-IS must transition to “wide metric”

We assume that the IGP in use is IS-IS. A configuration using OSPF is very similar and requires very few changes.

For OSPF, TE is typically in area 0

```
router ospf <pid>
  mpls traffic-eng area <area>
  mpls traffic-eng router-id <rtr>
```

We would like to establish a TE tunnel from R1 to R5

Head-end configuration (R1)

```
ip cef
mpls traffic-eng tunnels
interface loopback0
  ip address 192.168.1.1 255.255.255.255
  ip router isis

interface R1-R2
  ip address 10.10.10.1 255.255.255.252
  ip router isis
  mpls traffic-eng tunnels <-- enable TE
  ip rsvp bandwidth 100000 100000 <--enable RSVP, needed on both ends of any link an
LSP could pass over

router isis
  metric style wide
mpls traffic-eng level-2
mpls traffic-eng router-id loopback0

.....
! Setting up the tunnel from R1 to R5
! We will use a dynamic path i.e. best path is found by Constrain Based Routing
Algorithm
! It is also possible to set up an explicit path, in this case the explicit path
consisting of
! {router Id, interface} should be specified.
interface tunnell
  ip unnumbered loopback0
  no ip direct-broadcast
```

```
tunnel destination 192.168.1.5
tunnel mode mpls traffic-eng
tunnel mpls traffic-eng autoroute announce <-- announce tunnel tail reachability to
RIB, cause IGP to use the tunnel in its enhanced SPF calculation
tunnel mpls traffic-eng bandwidth 100
tunnel mpls traffic-eng priority 1 1
tunnel mpls traffic-eng path-option 1 dynamic <-- define how path will be calculated
tunnel mpls traffic-eng record-route
```

Mid-point (transit) configuration, in our case R2

```
ip cef
mpls traffic-eng tunnels
interface loopback0
    ip address 192.168.1.2 255.255.255.255
    ip router isis

interface R2-R1
    ip address 10.10.10.2 255.255.255.252
    ip router isis
    mpls traffic-eng tunnels
    ip rsvp bandwidth 1000 1000

interface R2-R4
    ip address 11.11.11.1 255.255.255.252
    ip router isis
    mpls traffic-eng tunnels
    ip rsvp bandwidth 100000 100000
! Make sure to enable traffic engineering on any interface one would
! like to be included in the Traffic engineering database
!

interface R2-R3
    ip address 12.12.12.1 255.255.255.252
    ip router isis
    mpls traffic-eng tunnels
    ip rsvp bandwidth 100000 100000
```

```
router isis
  metric style wide
  mpls traffic-eng level-2
  mpls traffic-eng router-id loopback0
```

Tail-end configuration

Almost the same configuration as the mid point router

```
ip cef
mpls traffic-eng tunnels
interface loopback0
  ip address 192.168.1.5 255.255.255.255
  ip router isis

interface R5-R4
  ip address 14.14.14.2 255.255.255.252
  ip router isis
  mpls traffic-eng tunnels
  ip rsvp bandwidth 100000 100000
```

.....

```
router isis
  metric style wide
  mpls traffic-eng level-2
  mpls traffic-eng router-id loopback0
```

At this stage the tunnel should be operable. To enable traffic forwarding into the tunnel on the head-end R1 or to configure static route for any destination behind the tail-end, use “tunnel mpls traffic-eng autoroute announce”.

The following command can be used to verify that traffic is routed through the tunnel.

```
show mpls traffic-eng tunnel
show ip route 192.168.1.5
show mpls traffic-eng autoroute
ping 192.168.1.5
show interface tunnel1 accounting
show interface r1-r2 accounting
```

To create an explicit path from R1 to R5, the head-end configuration should be slightly changed.

```
;define the explicit path
ip explicit-path identifier-explicit-path-name
```

```
next-address 10.10.10.2
```

```
next-address 11.11.11.2
```

```
next-address 14.14.14.2
```

under “interface tunnel1” add

```
tunnel mpls traffic-eng path-option 1 explicit name identifier-explicit-path-name
```

Note: For every LSP, explicit and dynamic path can be specified simultaneously. Choice of which path will be installed first is based on the “path-option” value.

Tunnel Protection

Assume that the primary LSP is R1-R2-R4-R5. We would like to protect the link between R2 and R4 with a backup R2-R3-R4 (via Tunnel 1000)

For the time being link protection (Fast Reroute) is only supported on POS interface.

First, build a backup tunnel R2-R4 going through R3.

Backup tunnel configuration on R2

```
interface tunnel1000
```

```
ip unnumbered loopback0
```

```
mpls traffic-eng tunnels
```

```
tunnel destination 13.13.13.2
```

```
tunnel mode mpls traffic-eng
```

```
tunnel mpls traffic-eng priority 0 0
```

```
tunnel mpls traffic-eng path-option 1 explicit backup-tunnel1000
```

```
ip rsvp bandwidth 1 1
```

```
ip explicit-path backup-tunnel1000
```

```
next address 12.12.12.2
```

```
next address 13.13.13.2
```

Protected link configuration on R2: link protection is supported on POS only.

```
interface R2-R4
```

```
ip address 11.11.11.1
```

```
mpls traffic-eng tunnels
```

```
mpls traffic-eng backup tunnel1000
```

```
pos ais-shut <-- we assume POS interface
```

```
pos report lrldi <-- we assume POS interface
```

```
ip rsvp bandwidth 2480000 2480000
```

On the head-end router, Tunnel1 should be aware that it has a backup tunnel ready to take over in case of link failure. It will set the 0x01 local link protection flag then the head-end will signal for the LSP for tunnel1.

Under “interface Tunnel1” add “tunnel mpls traffic-eng fast-reroute”


```

interface tunnel1
  ip unnumbered loopback0
  no ip direct-broadcast
  tunnel destination 192.168.1.5
  tunnel mode mpls traffic-eng
  tunnel mpls traffic-eng autoroute announce <-- announce tunnel tail reachability to
RIB, cause IGP use the tunnel in its enhanced SPF calculation
  tunnel mpls traffic-eng bandwidth 100
  tunnel mpls traffic-eng priority 1 1
  tunnel mpls traffic-eng path-option 1 dynamic <-- define how path will be calculated
  tunnel mpls traffic-eng record-route
  tunnel mpls traffic-eng fast-reroute

```

MPLS-TE+ QoS

Figure B2 Achieving End-to-end QoS Support



Consider the architecture defined in Figure B2. We would like to enforce DiffServ from end-to-end. The two IP clouds respectively on the left and right of the MPLS cloud are already DiffServ aware. The goal is to extend DiffServ support into the MPLS cloud.

The easiest and most natural solution (when a limited number of classes is used), is to map the DSCP marking directly to the MPLS Experimental bits at the ingress point of the MPLS cloud. Conversely, another mapping is done at the egress point and ensures that DiffServ is enforced from end-to-end.

Configuration DiffServ mapping to MPLS Exp:

On the ingress MPLS router:

- Mark and police traffic according to contract
- Define IP Precedence/DSCP mapping to EXP


```

class-map match-all PREMIUM-IP
  match ip dscp ef
class-map match-all BUSINESS-IP
  match ip dscp af31 af32 af33
!
policy-map IN-POLICY
  class PREMIUM-IP
    police 1280000 32000 32000
      conform-action set-mpls-exp-transmit 5
      exceed-action drop
  class BUSINESS-IP
    police 22000000 550000 550000
      conform-action set-mpls-exp-transmit 4
      exceed-action set-mpls-exp-transmit 3
  class default
    set mpls experimental 0

```

Note: The police command should be in one line. For readability purpose, it was split over several lines.

The IN-POLICY has to be applied on the inbound interface.

On the “Outbound” interface facing the MPLS clouds:

- Traffic classified by EXP
- LLQ (MDDR) for queuing MPLS packets
- WRED based on EXP to implement dropping precedence
- IP Precedence copied to MPLS EXP if no mapping defined in input policy

Define a policy-map OUT-POLICY, apply this policy on the outbound interface. (interface facing MPLS cloud).

Within the MPLS clouds, for the P routers:

- Service traffic based on EXP marking
- LLQ (MDRR) for MPLS packets
- WRED based on EXP

!Define a Cos Queue Group e.g. OUT-POLICY

```

cos-queue-group OUT-POLICY
  precedence 0 queue 0
  precedence 3 queue 1
  precedence 4 queue 1
  precedence 5 queue low-latency
  precedence 0 random-detect-label 0
  precedence 3 random-detect-label 1

```

```

precedence 4 random-detect-label 2
random-detect-label 0 300 500 1
random-detect-label 1 100 300 1
random-detect-label 2 300 500 1
queue 0 50
queue 1 50
queue low-latency strict-priority
!

```

Apply this policy on the P router outgoing interface

```

!
interface POS2/0
 ip addr X.X.X.X 255.255.255.252
 . . . . .
 tx-cos OUT-POLICY
!

```

Note: The above syntax is GSR specific. The same result can be achieved by using MQC (Modular QoS CLI) on platform such as 72xx/75xx and on the outcoming Engine 3 Linecards for GSR.

On the egress MPLS, DSCP bits will be set back to their original value

Configuration DS-TE

DS-TE provides the possibility of dedicating specific LSPs for high-priority/sensitive traffic where a higher quality of service performance (in terms of delay, jitter or loss) is required.

From a configuration standpoint, DS-TE comprises two main components:

- Configure two bandwidth pools in the core (“global pool” and “sub-pool”): use one pool—the sub-pool, for tunnels that carry traffic requiring strict bandwidth guarantees or delay guarantees, use the other pool- the global pool, for tunnels that carry traffic requiring only Differentiated service/Best effort. Within the MPLS core, assure that the traffic sent in the “sub pool” LSP is placed in a “high-priority/low latency” queue at the outbound interface of every LSR across the path. Further, assure also that this queue is never over subscribed.
- On the edge, rate limit the traffic before entering the “sub-pool” LSP tunnel. The aggregate rate of all traffic entering the “sub pool” tunnel should be less than or equal to the bandwidth capacity of the “sub-pool” tunnel. Excess traffic can be dropped or can be marked differently for preferential discard.

DS-TE is enabled within the MPLS core. MPLS DS-TE configuration is slightly different compared to vanilla Traffic Engineering. DS-TE is enabled in the core using extended version of the commands “tunnel mpls traffic-eng bandwidth sub-pool xxxx” and “ip rsvp bandwidth xxxxx yyyyy sub-pool zzzzz”.

Head-end Configuration

```

! we will only indicate the changes
!
interface R1-R2

```

```

ip address 10.10.10.1 255.255.255.0

ip router isis

mpls traffic-eng tunnels <-- enable TE

ip rsvp bandwidth 100000 100000 sub-pool 60000

router isis

metric style wide

mpls traffic-eng level-2

mpls traffic-eng router-id loopback0

.....

! Setting up the tunnel from R1 to R5

! tunnel10 is a DS-TE tunnel user for Traffics which requires tight QoS requirements.

! we will use a dynamic path i.e. best path is found by Constrain Based Routing
Algorithm

!

interface tunnel10

ip unnumbered loopback0

no ip direct-broadcast

tunnel destination 192.168.1.5

tunnel mode mpls traffic-eng

no tunnel mpls traffic-eng autoroute announce <-- Do not Announce Tunnel via IGP

tunnel mpls traffic-eng bandwidth sub-pool 40000

tunnel mpls traffic-eng priority 0 0

tunnel mpls traffic-eng path-option 1 dynamic <-- define how path will be calculated

tunnel mpls traffic-eng record-route

```

As destinations behind the tunnel are not announce by "Autoroute announce," we would need to insert a static route which will use Tunnel 10. Thus we make sure that only the desired traffic will use Tunnel 10.

On the inbound interface:

- Create a class of traffic matching ACL 100 named "ds-te-class" and apply it to all packets destined to "OUR-DESTINATION".
- Create a policy ds-te-input-policy, where
 - Packet in the class "ds-te-class" are rate-limited to:
 - 8 million bits per second
 - normal burst of 1 million bytes
 - maximum burst of 2 millions bytes
 - Packets conforming this rate are marked with an MPLS EXP value of 5 and are forwarded
 - Packets, which exceed this rate, are dropped
 - All other packets are marked with an MPLS EXP of 0 and forwarded.

On the R1-R2 interface:

- All MPLS packets with the EXP bit set to 5 will be placed in high-priority (or low-latency) queue

Note: Based on the hardware in use (either 7xxx or GSR), the command lines used to achieve the above requirements as far as QoS is concerned might be different. QoS commands for the GSR are stated in regular Cisco CLI, whereas the one for the 7xxx are based on the Modular QoS CLI.

Head End is a 7xxx

```
class-map match-all ds-te-class
  match access-group 100
access-list 100 permit ip any host "our-destination"
policy-map ds-te-input-policy
  class ds-te-class
    police 8000000 1000000 2000000 conform-action set-mpls-exp-transmit 5
    exceed-action drop
  class class-default
    set-mpls-exp-transmit 0
```

Apply this policy to the inbound (ingress) interface

```
interface "inbound"
  service-policy input ds-te-input-policy
  .....
```

!on the outbound MPLS interface

!Put all MPLs traffic with EXP bit set to 5 to High Priority (or low latency) queue

```
class-map match-all exp5-traffic
  match mpls experimental 5
policy-map output-interface-policy
  class exp5-traffic
    priority 32
interface R1-R2
  service-policy output output-interface-policy
  ....
```

Head End is a GSR

Same configuration translated to GSR using regular Cisco CLI for QoS

! For the ingress interface

!

```
access-list 100 permit ip any "our-destination" 0.0.0.255
interface "ingress"
  rate-limit input access-group 100 8000000 1000000 2000000 \
    conform-action set-mpls-exp-transmit 5\
```

```

        exceed-action set-mpls-exp-transmit 0
! On the outbound interface (R1-R2)
!
interface R1-R2
....
tx-cos exp-class-ds-te

```

Where "exp-class-ds-te" is defined as:

```

cos-queue-group exp-class-ds-te
  precedence 0 random-detect-label 0
  precedence 5 queue low-latency
  precedence 5 random-detect-label 5
  random-detect-label 0 100 200 1
  random-detect-label 5 2000 3000 1
  queue low-latency strict-priority

```

Midpoint Configuration

Both inbound and outbound interfaces on the Midpoint router are configured identically to the outbound interface of the Head-end router.

Tail-End Configuration

The inbound interfaces (facing the MPLS cloud) of the Tail-End router are configured identically to the inbound interfaces of the midpoint routers.

In the above example, only DS-TE traffic is flowing in the MPLS core. Imagine a configuration, where the traffic destined to "our-destination" use DS-TE and the other traffics take advantage of MPLS TE.

As in the previous configuration, DS-TE LSP will still be using the low latency/high priority queue; the vanilla TE traffic will use "normal" queue. Further, on all the inbound interfaces, where DS-TE LSP and vanilla TE LSP might run concurrently, a rate limit policy needs to be established. This policy will allow us to make sure that DS-TE bandwidth is always available if there is a need and the vanilla TE LSP will not take all the bandwidth.

For More Information

Additional information about the Cisco IOS MPLS technology can be found at <http://www.cisco.com/go/mpls/> or by contacting your local Cisco representative.

References and Recommended Reading

- [1] Callon, R., Doolan, P., Feldman, N., Fredette, A., Swallow, G. and A. Viswanathan, "A Framework for Multiprotocol Label Switching," Work in Progress.
- [2] Awduche D., Malcolm J., Agogbua J., O'Dell M., McManus J. "Requirements for Traffic Engineering over MPLS," RFC 2702, September 1999
- [3] Davie B., Rekhter Y. "MPLS Technology and Application" Morgan Kaufmann Publishers.
- [4] MPLS Traffic Engineering
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120s/120s5/mpls_te.htm

- [5] Braden B., et al., "Resource Reservation Protocol (RSVP)-Version 1 Functional Specification," RFC 2205, September 1997
- [7] Le Faucheur et al., "MPLS Support of Differentiated Services," draft-ietf-mpls-diff-ext-09.txt, April 2001
- [8] Awduche D., et al., "RSVP-TE: Extensions to RSVP for LSP Tunnels," draft-ietf-mpls-rsvp-lsp-tunnel-09.txt, August 2001
- [9] Robert Raszuk "Networkers 2000: MPLS Traffic Engineering"
<http://www.cisco.com/networkers/nw00/pres/pdf2000.htm>
- [10] Le Faucheur et al., "Protocol extensions for support of Diff-Serv-Aware MPLS Traffic Engineering," draft-lefaucheur-diff-te-proto-00.txt, July 2001.
- [11] Osborne E., Aziz Z., "Deploying MPLS TE and Backbone VPN,"
http://www.cisco.com/networkers/nw00/pres/2202_7-6.pdf
- [12] Diff-Serv-Aware Traffic Engineering (DS-TE)
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120st/120st11/ds_te.htm
- [13] Configuring MPLS Basic Traffic Engineering Using IS-IS <http://www.cisco.com/warp/public/105/mplsteisis.html>
- [14] Configuring MPLS Basic Traffic Engineering Using OSPF
http://www.cisco.com/warp/public/105/mpls_te_ospf.html
- [15] MPLS QoS FAQ



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, CCSI, CCENT, 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)