

## Requirements for support of Inter-Area MPLS Traffic Engineering

### Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 10 of RFC2026](#). Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

### Abstract:

Traffic engineering using MPLS within a single IGP area has become quite prevalent in networks today. This gives a network operator several advantages, including better traffic measurement capabilities, faster restoration times, and of course the ability to better manage congestion points within their network. However, the current uses have been limited to one IGP area. OSPF based networks usually do traffic engineering only in their backbone area, and many ISIS networks still operate at a single level. It would be useful to be able to extend these capabilities across hierarchical areas within a single IGP, as well as to be able to extend across AS boundaries if it is found to be useful and possible. This document first considers current applications of MPLS traffic engineering (TE), and then it addresses how these might be used (or precluded) across IGP areas. For the purposes of this document, IGP areas refers both to OSPF areas and ISIS multi-level hierarchies. Requirements for Inter-AS TE can be found in [INTER-AS-TE].

Please direct comments on this document to [te-wg@ops.ietf.org](mailto:te-wg@ops.ietf.org)

## Table of Contents

- 1.... Current uses of MPLS Traffic Engineering
  - 1.1.. Fundamental functionalities
    - 1.1.1 MPLS LSPs used for Tunneling
  - 1.2.. MPLS TE and routing
  - 1.3.. Intra-area Protection
  - 1.4.. Diffserv and TE
- 2.... General Requirements
- 3.... Inter-area requirements
  - 3.1.. Inter-Area Path Selection
  - 3.2.. Inter-area MPLS TE and routing
  - 3.3.. Inter-area MPLS TE and Protection
  - 3.4.. Diffserv and Inter-area MPLS TE
- 4.... Security Considerations
- 5.... Acknowledgments
- 6.... References
- 7.... Author's Address

## **1. Current uses of MPLS Traffic Engineering**

### **1.1 Fundamental functionalities**

There are a number of primitives that are specified [RSVP-LSP] [RSVP-OSPF] and useful in MPLS TE. Some are itemized here and revisited when discussing how they might be needed or useful for inter-area MPLS TE.

- a) bandwidth specification - allows the sender to specify how much bandwidth is required for an MPLS LSP. This is done in the Tspec.
- b) priorities - allows prioritization of access to link bandwidth. Priorities for setup and hold may be specified, though it is a common practice to set these to be equal, when specified.
- c) ERO - This allows the source to specify the route an LSP should take. The specification is often a strict route determined by the source during path selection, however the protocol allows for loose hops that an LSP should be routed toward hop-by-hop.
- d) RRO - A sender can request that the route an LSP takes is confirmed during the reservation process by inserting a record route object. The semantics of this are the same as the ERO.
- e) protection requested - Different protection techniques are

possible, and the ability to use these is controlled by the source of the LSP [PROTECTION].

f) resource affinities - Links can be assigned to administrative groups and these can be taken into account when selecting a path. The administrative constraints of an LSP are also signaled in the session object of the LSP.

Also, different reservation styles can be used, including shared-explicit which allows for make-before-break functionality during bandwidth and path changes on an LSP. The source is able to specify that shared-explicit is requested in the session object of the LSP.


The fact that all traffic for the LSP is consolidated into one tunnel makes measurement of macroscopic flows much easier than with traditional IP routing. This is also an advantage network operators find in using MPLS TE.

**1.1.1 MPLS LSPs used for Tunneling**

In addition to the carriage of plain, globally routable IPv4 traffic, MPLS is often used to carry VPN traffic between ports of a network. This may be a customer's private IPv4, ATM, Frame Relay or Ethernet traffic. The two endpoints agree on what is in the tunnel via configuration or additional signaling [[RFC2547](#)] [MARTINI]. This is one of the strongest current drivers for new deployments of MPLS. It is also one of the strongest drivers for the need to push MPLS TE across area boundaries, as connectivity is required between network edges.

**1.2 MPLS TE and routing**

At its simplest, MPLS TE allows for the connection between two routers and the tunneling of traffic between them. This might be done for the purposes of measurement, transport of VPN traffic, or to provide a quicker restoration of traffic flow during different failure scenarios. In this mode, the only traffic introduced onto the tunnel is traffic for the destination. For instance, in the following topology if there is an LSP from M1-to-N2, only traffic going through M1 with N2 as its final destination would end up on the LSP.





Boyle

3

Requirements for Inter-Area TE


December 2003

Alternate approaches modify the IGP's shortest path selection algorithm. In this case, if the IGP's shortest paths from M1 are as follows:




Then an LSP from M1-to-N2 might create a shortcut such that the path M1-to-N2 plus the cost of the link from N2 to N3 might be better than M1-N1-N3 path, so traffic hitting M1 and going to N3 would end up the LSP M1-to-N2. Another implementation might only use the LSP for IGP destinations that naturally fall beyond the LSP destination on the shortest path tree. In the above example, M1-to-N2 would not be used for N3, however an LSP from M1-to-N1 would include traffic to N3.

Another common approach is to treat these LSPs as forwarding adjacencies which are included in the LSP source's IGP advertisement [FA-LSP]. In this case, traffic from other nodes can be drawn toward a node with an LSP. For example, if the IGP shortest path tree from M4's perspective is the following:



If the LSP M1-to-N2 is advertised back into the IGP with the right cost, this might modify M4's tree to be as follows:



These are all useful ways to incorporate MPLS into routing with intra-area LSPs.

In summary, MPLS LSPs are used in intra-area TE in the following ways.

- 1) When traffic goes through the source and is to the destination

of the LSP

- 2) When traffic goes through the source and is to the destination of the LSP or somewhere beyond this destination from an IGP SPF perspective.
- 3) The LSP is introduced into the IGP to become part of the path

Boyle

4

Requirements for Inter-Area TE

December 2003

consideration for all nodes with visibility into the advertisement.

### **1.3 Intra-area Protection**

MPLS allows for a variety of approaches to speed up traffic restoration when there is a failure within a network (especially link-based failures). Common approaches include having two LSPs (a working and protect) on separate paths, as well as using MPLS signaling to setup more advanced protection schemes. [PROTECTION] outlines that LSPs may be protected in several approaches. For each LSP, a detour LSP may be routed toward the destination but avoiding the next node (and link). These detours may merge back into the main LSP. Alternatively, an LSP can be established for a given link and if there is a failure of that link, all LSPs that were active on the failed link are switched into the backup LSP. These are both approaches that have found use in today's MPLS networks.

### **1.4 Diffserv and TE**

The MPLS label allows for a CoS value to be used within the "Exp" bits, and this can be used to affect the proper queuing for the packet. This is termed a label-inferred diffserv LSP (L-LSP) in [DIFFSERV-LSP]. An explicit object can also be used to specify the diffserv class of traffic that the LSP is limited to, and that can be used to treat the LSP accordingly. Other approaches are being developed to allow more intimacy with diffserv and route selection, this is termed Diffserv TE, and is specified in [DIFF-TE]. One thing to note about Diffserv TE, besides its pre-emergent state at this time, is that many of the semantics are abstract and inferred by configuration. For instance, Class-Type 0, 1 and 2 are often correlated to Best Effort, Assured Forwarding and Expedited Forwarding, however that is only from convention and the protocol allows much more flexibility to achieve the requirements of the network operator.

## **2. General Requirements**

As outlined in [section 1](#), there are many uses of MPLS TE today, though it is mostly constrained to a single IGP area. Some of the information available in the routing domain of a single area is so condensed when an area boundary is crossed that some of the functionality could be sharply curtailed. For example, the absence of edge to edge link state topology information makes complete explicit route discovery impossible. However, it might still be useful, especially for transporting of VPN traffic from edge to edge of a network (for inter-area). Other current TE uses may be more difficult to achieve inter-area due to the lack of ability to integrate an MPLS LSP into an inter-area routing scheme. For

Boyle

5

Requirements for Inter-Area TE

December 2003

instance, interconnection of a hosting center to a significant peering router in another area may be useful for traffic quantification, explicit route usage, or interaction with EBGp route selection. However this may have to be done from edge router to edge router. It might be impossible to introduce an LSP that originates in one area into the routing such that traffic to another area is drawn towards it. Currently, this sort of topology manipulation is possible within a single area.

It might be possible to achieve limited results and this document attempts to specify some constraints that might be useful guidelines.

There are several reasons that network operators choose to break up their network into different areas. These often include scalability and containment of routing information. The latter can help isolate most of a network from receiving and processing updates that are of no consequence to its routing decisions. In fact, whole routers can be taken down in an area without any router outside of that area being aware. Scalability and containment of routing information should not be compromised to allow inter-area traffic engineering. In fact, to the extent possible, information propagation for path-selection should continue to be localized.

Additionally, The base specification of MPLS TE is architecturally structured and relatively devoid of excessive state propagation in terms of routing or signaling. Its strength in extensibility can also be seen as an Achilles heel, as there is really no limit to what is possible with extensions. It is paramount to maintain architectural vision and discretion when adapting it for use for inter-area and inter-as TE. Additional information carried within an area, or propagated outside of an area (via routing or

signaling) should neither be excessive, patchwork, nor non-relevant.

For these reasons, the solution should stress scalability and generality over piece-meal point application extensions or overly comprehensive approaches that deliver more than may be necessary at this time.

### **3. Inter-area requirements**

All of the features outlined in [section 1](#) can be assumed to have the same meaning and semantics in different areas of a multi-area IGP.

For instance, an affinity of "international" (0x1) in one area can be assumed to be "international" in a different area. Most of the primitives used for intra-area path selection are propagated in the session object of an RSVP LSP.

Boyle

6

Requirements for Inter-Area TE

December 2003

Therefore, all of the functionalities outlined in [section 1.1](#) are possible and necessary for support in inter-area MPLS TE.

#### **3.1 Inter-Area Path Selection**

Obviously some functions such as a full route selection for an ERO will not be possible across area boundaries. For this reason, and since the most necessary primitives for route selection are propagated within an LSPs Path message, an approach of ERO route expansion is suggested as the best approach for inter-area LSP establishment. An alternate approach of propagation of sufficient topology information to make a complete path determination is not recommended due to the additional information required as well as the incongruence of this approach with the overall containment provided by the network hierarchy. The alignment of these is mentioned as a goal in [section 2](#).

In overview, the requirements below call for what may be called a crankback method which employs head end direction for regional optimization.

A source should specify a strict ERO to an egress point of its area with the next hop being at a minimum the destination for the LSP. The area border router will expand this ERO to the destination (or its border router if the destination is more than one area away from the source). When insufficient topology exist at a route

expansion at a border router, a Path Error message will be sent to the source, who may attempt to establish the LSP through another border router. Similarly, when the source is more than one area away from the destination, the destination's border router may send back a Path Error. Ideally, the source's border router would try another border router into the destination's area, however with current protocol, the Path Error will propagate to the source. The source now needs to know if it should try again at the border router of the source's area. This is the challenge of inter-area traffic engineered LSP establishment

A solution is needed which elegantly and efficiently allows for the LSP to be routed through the best pair of border routers. As these may change over time, there is a need to show how border-router optimization is handled.

Additionally, more optimal routes may become available within an area. To the extent possible, optimization of the LSP onto these routes should be possible, potentially without impacting state in other areas. The head-end of the LSP should be able to have control on the candidacy of an LSP for optimization in a remote area.

### **3.2 Inter-area MPLS TE and routing**

Clearly an inter-area LSP can be used to carry traffic that is destined for the destination of the LSP. (item (1) at the end of [section 1.2](#)).

Much information may also be present which allows other traffic to be included in the LSP. Area border routers propagate information for more than their directly attached interfaces. In OSPF, network summary LSAs are sent throughout the network by ABRs (and external networks are propagated by ASBRs). This information allows the LSP also to be used to reach networks beyond an LSP to a border router or an ASBR (item (2) at the end of [section 1.2](#)).

However, it may be difficult to allow the existence of an LSP (or two from source to destination and vice-versa) to attract traffic to an LSP's source. The only way that might work would be to mirror existing network summary and external advertisements. However since these can flood throughout a network, they could have a severe impact on IGP database size. Incorporation of LSPs back into routing across areas is not currently required and is not recommended outside of experimental investigation. (item 3 at the



end of [section 1.2](#))

### **[3.3](#) Inter-area MPLS TE and Protection**

Facility protection of a single link is necessarily intra-area and not applicable to inter-area discussions.

Per-path protection is necessary and can hopefully use the same path establishment mechanisms (along with necessary REROUTE and DETOUR objects) to achieve the same restoration levels as are achieved within an area.

Since the complete ERO may not be present, there may not be sufficient information to initially calculate a valid detour for all cases. However, a minimization of protocol modifications and state propagation may still be desired, even if it came at the expense of delayed detour establishment. For instance, one approach might be to wait for an LSP's RRO to help craft a detour which explicitly merges back into the main LSP quickly after the area boundary. Or an approach might be to simply try a detour through the next most reasonable ABR. The LSP will likely merge back in naturally, however there can be some topologies where the backup ABR can not expand the path without avoiding one of the nodes listed in the DETOUR object.

### **[3.4](#) Diffserv and Inter-area MPLS TE**

Just as session specific parameters such as affinities are assumed

Boyle

8

Requirements for Inter-Area TE

December 2003

to be consistent within one IGP, the same can be said for diffserv support on an LSP. COS mappings on L-LSPs are expected to be consistent, and E-LSPs are explicitly interpreted by definition.

Diffserv TE should be directly translatable at border-routers, as the class-type of an LSP is explicit for class-types greater than 0 and not-existent in the path message for class-type 0.

## **[4.](#) Security Considerations**

None.

## **[5.](#) Acknowledgments**

Thanks to Craig Pierantozzi, Laura Cunningham and Jerry Ash, whose feedback and editorial review have helped refine this document.

## **6. References**

[DIFF-TE]	<a href="#">draft-ietf-tewg-diff-te-proto-05.txt</a>
[DIFFSERV-LSP]	<a href="#">RFC3270</a>
[FA-LSP]	<a href="#">draft-ietf-mpls-lsp-hierarchy-08.txt</a>
[INTER-AS-TE]	<a href="#">draft-ietf-tewg-interas-mpls-te-req-02.txt</a>
[MARTINI]	<a href="#">draft-martini-l2circuit-trans-mpls-11.txt</a>
[PROTECTION]	<a href="#">draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt</a>
[RFC2547]	<a href="#">RFC2547</a>
[RSVP-LSP]	<a href="#">RFC3209</a>
[RSVP-OSPF]	<a href="#">RFC3630</a>

## **7. Author's Address**

Jim Boyle  
Email: [jboyle@pdnets.com](mailto:jboyle@pdnets.com)