

Network Working Group  
Request for Comments: 4105  
Category: Informational

J.-L. Le Roux, Ed.  
France Telecom  
J.-P. Vasseur, Ed.  
Cisco Systems, Inc.  
J. Boyle, Ed.  
PDNETs  
June 2005

## Requirements for Inter-Area MPLS Traffic Engineering

### Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The Internet Society (2005).

### Abstract

This document lists a detailed set of functional requirements for the support of inter-area MPLS Traffic Engineering (inter-area MPLS TE). It is intended that solutions that specify procedures and protocol extensions for inter-area MPLS TE satisfy these requirements.

### Table of Contents

1. Introduction .....	2
2. Conventions Used in This Document .....	3
3. Terminology .....	3
4. Current Intra-Area Uses of MPLS Traffic Engineering .....	4
4.1. Intra-Area MPLS Traffic Engineering Architecture .....	4
4.2. Intra-Area MPLS Traffic Engineering Applications .....	4
4.2.1. Intra-Area Resource Optimization .....	4
4.2.2. Intra-Area QoS Guarantees .....	5
4.2.3. Fast Recovery within an IGP Area .....	5
4.3. Intra-Area MPLS TE and Routing .....	6
5. Problem Statement, Requirements, and Objectives of Inter-Area ...	6
5.1. Inter-Area Traffic Engineering Problem Statement .....	6
5.2. Overview of Requirements for Inter-Area MPLS TE .....	7
5.3. Key Objectives for an Inter-Area MPLS-TE Solution .....	8
5.3.1. Preserving the IGP Hierarchy Concept .....	8
5.3.2. Preserving Scalability .....	8
6. Application Scenario.....	9

- 7. Detailed Requirements for Inter-Area MPLS TE .....10
  - 7.1. Inter-Area MPLS TE Operations and Interoperability .....10
  - 7.2. Inter-Area TE-LSP Signaling .....10
  - 7.3. Path Optimality .....11
  - 7.4. Inter-Area MPLS-TE Routing .....11
  - 7.5. Inter-Area MPLS-TE Path Computation .....12
  - 7.6. Inter-Area Crankback Routing .....12
  - 7.7. Support of Diversely-Routed Inter-Area TE LSPs .....13
  - 7.8. Intra/Inter-Area Path Selection Policy .....13
  - 7.9. Reoptimization of Inter-Area TE LSP .....13
  - 7.10. Inter-Area LSP Recovery .....14
    - 7.10.1. Rerouting of Inter-Area TE LSPs .....14
    - 7.10.2. Fast Recovery of Inter-Area TE LSP .....14
  - 7.11. DS-TE support .....15
  - 7.12. Hierarchical LSP Support .....15
  - 7.13. Hard/Soft Preemption .....15
  - 7.14. Auto-Discovery of TE Meshes .....16
  - 7.15. Inter-Area MPLS TE Fault Management Requirements .....16
  - 7.16. Inter-Area MPLS TE and Routing .....16
- 8. Evaluation criteria .....17
  - 8.1. Performances .....17
  - 8.2. Complexity and Risks .....17
  - 8.3. Backward Compatibility .....17
- 9. Security Considerations .....17
- 10. Acknowledgements .....17
- 11. Contributing Authors .....18
- 12. Normative References .....19
- 13. Informative References .....19

1. Introduction

The set of MPLS Traffic Engineering components, defined in [RSVP-TE], [OSPF-TE], and [ISIS-TE], which supports the requirements defined in [TE-REQ], is used today by many network operators to achieve major Traffic Engineering objectives defined in [TE-OVW]. These objectives include:

- Aggregated Traffic measurement
- Optimization of network resources utilization
- Support for services requiring end-to-end QoS guarantees
- Fast recovery against link/node/Shared Risk Link Group (SRLG) failures

Furthermore, the applicability of MPLS to traffic engineering in IP networks is discussed in [TE-APP].

The set of MPLS Traffic Engineering mechanisms, to date, has been limited to use within a single Interior Gateway Protocol (IGP) area.

This document discusses the requirements for an inter-area MPLS Traffic Engineering mechanism that may be used to achieve the same set of objectives across multiple IGP areas.

Basically, it would be useful to extend MPLS TE capabilities across IGP areas to support inter-area resources optimization, to provide strict QoS guarantees between two edge routers located within distinct areas, and to protect inter-area traffic against Area Border Router (ABR) failures.

First, this document addresses current uses of MPLS Traffic Engineering within a single IGP area. Then, it discusses a set of functional requirements that a solution must or should satisfy in order to support inter-area MPLS Traffic Engineering. Because the scope of requirements will vary between operators, some requirements will be mandatory (MUST), whereas others will be optional (SHOULD). Finally, a set of evaluation criteria for any solution meeting these requirements is given.

## 2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

## 3. Terminology

LSR:	Label Switching Router
LSP:	Label Switched Path
TE LSP:	Traffic Engineering Label Switched Path
Inter-area TE LSP:	TE LSP whose head-end LSR and tail-end LSR do not reside within the same IGP area or whose head-end LSR and tail-end LSR are both in the same IGP area although the TE-LSP transiting path is across different IGP areas.
IGP area:	OSPF area or IS-IS level.
ABR:	Area Border Router, a router used to connect two IGP areas (ABR in OSPF, or L1/L2 router in IS-IS).
CSPF:	Constraint-based Shortest Path First.
SRLG:	Shared Risk Link Group.

#### 4. Current Intra-Area Uses of MPLS Traffic Engineering

This section addresses architecture, capabilities, and uses of MPLS TE within a single IGP area. It first summarizes the current MPLS-TE architecture, then addresses various MPLS-TE capabilities, and finally lists various approaches to integrate MPLS TE into routing. This section is intended to help define the requirements for MPLS-TE extensions across multiple IGP areas.

##### 4.1. Intra-Area MPLS Traffic Engineering Architecture

The MPLS-TE control plane allows establishing explicitly routed MPLS LSPs whose paths follow a set of TE constraints. It is used to achieve major TE objectives such as resource usage optimization, QoS guarantee and fast failure recovery. It consists of three main components:

- The routing component, responsible for the discovery of the TE topology. This is ensured thanks to extensions of link state IGP: [ISIS-TE], [OSPF-TE].
- The path computation component, responsible for the placement of the LSP. It is performed on the head-end LSR thanks to a CSPF algorithm, which takes TE topology and LSP constraints as input.
- The signaling component, responsible for the establishment of the LSP (explicit routing, label distribution, and resources reservation) along the computed path. This is ensured thanks to RSVP-TE [RSVP-TE].

##### 4.2. Intra-Area MPLS Traffic Engineering Applications

###### 4.2.1. Intra-Area Resource Optimization

MPLS TE can be used within an area to redirect paths of aggregated flows away from over-utilized resources within a network. In a small scale, this may be done by explicitly configuring a path to be used between two routers. On a grander scale, a mesh of LSPs can be established between central points in a network. LSPs paths can be defined statically in configuration or arrived at by an algorithm that determines the shortest path given administrative constraints such as bandwidth. In this way, MPLS TE allows for greater control over how traffic demands are routed over a network topology and utilize a network's resources.

Note also that TE LSPs allow measuring traffic matrix in a simple and scalable manner. The aggregated traffic rate between two LSRs is easily measured by accounting of traffic sent onto a TE LSP provisioned between the two LSRs in question.

#### 4.2.2. Intra-Area QoS Guarantees

The DiffServ IETF working group has defined a set of mechanisms described in [DIFF-ARCH], [DIFF-AF], and [DIFF-EF] or [MPLS-DIFF], that can be activated at the edge of or over a DiffServ domain to contribute to the enforcement of a QoS policy (or set of policies), which can be expressed in terms of maximum one-way transit delay, inter-packet delay variation, loss rate, etc. Many Operators have some or full deployment of DiffServ implementations in their networks today, either across the entire network or at least at its edge.

In situations where strict QoS bounds are required, admission control inside the backbone of a network is in some cases required in addition to current DiffServ mechanisms. When the propagation delay can be bounded, the performance targets, such as maximum one-way transit delay, may be guaranteed by providing bandwidth guarantees along the DiffServ-enabled path.

MPLS TE can be simply used with DiffServ: in that case, it only ensures aggregate QoS guarantees for the whole traffic. It can also be more intimately combined with DiffServ to perform per-class of service admission control and resource reservation. This requires extensions to MPLS TE called DiffServ-Aware TE, which are defined in [DSTE-PROTO]. DS-TE allows ensuring strict end-to-end QoS guarantees. For instance, an EF DS-TE LSP may be provisioned between voice gateways within the same area to ensure strict QoS to VoIP traffic.

MPLS TE allows computing intra-area shortest paths, which satisfy various constraints, including bandwidth. For the sake of illustration, if the IGP metrics reflects the propagation delay, it allows finding a minimum propagation delay path, which satisfies various constraints, such as bandwidth.

#### 4.2.3. Fast Recovery within an IGP Area

As quality-sensitive applications are deployed, one of the key requirements is to provide fast recovery mechanisms, allowing traffic recovery to be guaranteed on the order of tens of msecs, in case of network element failure. Note that this cannot be achieved by relying only on classical IGP rerouting.

Various recovery mechanisms can be used to protect traffic carried onto TE LSPs. They are defined in [MPLS-RECOV]. Protection mechanisms are based on the provisioning of backup LSPs that are used to recover traffic in case of failure of protected LSPs. Among those protection mechanisms, local protection (also called Fast Reroute) is intended to achieve sub-50ms recovery in case of link/node/SRLG

failure along the LSP path [FAST-REROUTE]. Fast Reroute is currently used by many operators to protect sensitive traffic inside an IGP area.

[FAST-REROUTE] defines two modes for backup LSPs. The first, called one-to-one backup, consists of setting up one detour LSP per protected LSP and per element to protect. The second, called facility backup, consists of setting up one or several bypass LSPs to protect a given facility (link or node). In case of failure, all protected LSPs are nested into the bypass LSPs (benefiting from the MPLS label stacking property).

#### 4.3. Intra-Area MPLS TE and Routing

There are several possibilities for directing traffic into intra-area TE LSPs:

- 1) Static routing to the LSP destination address or any other addresses.
- 2) IGP routes beyond the LSP destination, from an IGP SPF perspective (IGP shortcuts).
- 3) BGP routes announced by a BGP peer (or an MP-BGP peer) that is reachable through the TE LSP by means of a single static route to the corresponding BGP next-hop address (option 1) or by means of IGP shortcuts (option 2). This is often called BGP recursive routing.
- 4) The LSP can be advertised as a link into the IGP to become part of IGP database for all nodes, and thus can be taken into account during SPF for all nodes. Note that, even if similar in concept, this is different from the notion of Forwarding-Adjacency, as defined in [LSP-HIER]. Forwarding-Adjacency is when the LSP is advertised as a TE-link into the IGP-TE to become part of the TE database and taken into account in CSPF.

### 5. Problem Statement, Requirements, and Objectives of Inter-Area MPLS TE

#### 5.1. Inter-Area Traffic Engineering Problem Statement

As described in Section 4, MPLS TE is deployed today by many operators to optimize network bandwidth usage, to provide strict QoS guarantees, and to ensure sub-50ms recovery in case of link/node/SRLG failure.

However, MPLS-TE mechanisms are currently limited to a single IGP area. The limitation comes more from the Routing and Path computation components than from the signaling component. This is basically because the hierarchy limits topology visibility of head-

end LSRs to their IGP area, and consequently head-end LSRs can no longer run a CSPF algorithm to compute the shortest constrained path to the tail-end, as CSPF requires the whole topology to compute an end-to-end shortest constrained path.

Several operators have multi-area networks, and many operators that are still using a single IGP area may have to migrate to a multi-area environment, as their network grows and single area scalability limits are approached.

Thus, those operators may require inter-area traffic engineering to:

- Perform inter-area resource optimization.
- Provide inter-area QoS guarantees for traffic between edge nodes located in different areas.
- Provide fast recovery across areas, to protect inter-area traffic in case of link or node failure, including ABR node failures.

For instance, an operator running a multi-area IGP may have voice gateways located in different areas. Such VoIP transport requires inter-area QoS guarantees and inter-area fast protection.

One possible approach for inter-area traffic engineering could consist of deploying MPLS TE on a per-area basis, but such an approach has several limitations:

- Traffic aggregation at the ABR levels implies some constraints that do not lead to efficient traffic engineering. Actually, this per-area TE approach might lead to sub-optimal resource utilization, by optimizing resources independently in each area. What many operators want is to optimize their resources as a whole; in other words, as if there was only one area (flat network).
- This does not allow computing an inter-area constrained shortest path and thus does not ensure end-to-end QoS guarantees across areas.
- Inter-area traffic cannot be protected with local protection mechanisms such as [FAST-REROUTE] in case of ABR failure.

Therefore, existing MPLS TE mechanisms have to be enhanced to support inter-area TE LSPs.

## 5.2. Overview of Requirements for Inter-Area MPLS TE

For the reasons mentioned above, it is highly desired to extend the current set of MPLS-TE mechanisms across multiple IGP areas in order to support the intra-area applications described in Section 4 across areas.

The solution **MUST** allow setting up inter-area TE LSPs; i.e., LSPs whose path crosses at least two IGP areas.

Inter-area MPLS-TE extensions are highly desired in order to provide:

- Inter-area resources optimization.
- Strict inter-area QoS guarantees.
- Fast recovery across areas, particularly to protect inter-area traffic against ABR failures.

It may be desired to compute inter-area shortest paths that satisfy some bandwidth constraints or any other constraints, as is currently possible within a single IGP area. For the sake of illustration, if the IGP metrics reflects the propagation delay, it may be necessary to be able to find the optimal (shortest) path satisfying some constraints (e.g., bandwidth) across multiple IGP areas. Such a path would be the inter-area path offering the minimal propagation delay.

Thus, the solution **SHOULD** provide the ability to compute inter-area shortest paths satisfying a set of constraints (i.e., bandwidth).

### 5.3. Key Objectives for an Inter-Area MPLS-TE Solution

Any solution for inter-area MPLS TE should be designed with preserving IGP hierarchy concept, and preserving routing and signaling scalability as key objectives.

#### 5.3.1. Preserving the IGP Hierarchy Concept

The absence of a full link-state topology database makes the computation of an end-to-end optimal path by the head-end LSR not possible without further signaling and routing extensions. 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. Containment of routing information **MUST** not be compromised to allow inter-area traffic engineering. Information propagation for path-selection **MUST** continue to be localized. In other words, the solution **MUST** entirely preserve the concept of IGP hierarchy.

#### 5.3.2. Preserving Scalability

Achieving the requirements listed in this document **MUST** be performed while preserving the IGP scalability, which is of the utmost importance. The hierarchy preservation objective addressed in the above section is actually an element to preserve IGP scalability.




The solution also MUST not increase IGP load unreasonably, which could compromise IGP scalability. In particular, a solution satisfying those requirements MUST not require the IGP to carry some unreasonable amount of extra information and MUST not unreasonably increase the IGP flooding frequency.

Likewise, the solution MUST also preserve scalability of RSVP-TE ([RSVP-TE]).

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 no real limit to what is possible with extensions. It is paramount to maintain architectural vision and discretion when adapting it for use for inter-area MPLS TE. Additional information carried within an area or propagated outside of an area (via routing or signaling) should be neither excessive, patchwork, nor non-relevant.

Particularly, as mentioned in Section 5.2, it may be desired for some inter-area TE LSP carrying highly sensitive traffic to compute a shortest inter-area path, satisfying a set of constraints such as bandwidth. This may require an additional routing mechanism, as base CSPF at head-end can no longer be used due to the lack of topology and resource information. Such a routing mechanism MUST not compromise the scalability of the overall system.

6. Application Scenario



- ABR1, ABR2: Area0-Area1 ABRs
- ABR3, ABR4: Area0-Area2 ABRs
- R0, R1, R5: LSRs in area 1
- R2: an LSR in area 0
- R4: an LSR in area 2

Although the terminology and examples provided in this document make use of the OSPF terminology, this document equally applies to IS-IS.

Typically, an inter-area TE LSP will be set up between R0 and R4, where both LSRs belong to different IGP areas. Note that the solution MUST support the capability to protect such an inter-area TE LSP from the failure on any Link/SRLG/Node within any area and the failure of any traversed ABR. For instance, if the TE LSP R0->R4 goes through R1->ABR1->R2, then it can be protected against ABR1 failure, thanks to a backup LSP (detour or bypass) that may follow the alternate path R1->ABR2->R2.

For instance, R0 and R4 may be two voice gateways located in distinct areas. An inter-area DS-TE LSP with class-type EF is set up from R1 to R4 to route VoIP traffic classified as EF. Per-class inter-area constraint-based routing allows the DS-TE LSP to be routed over a path that will ensure strict QoS guarantees for VoIP traffic.

In another application, R0 and R4 may be two pseudo wire gateways residing in different areas. An inter-area LSP may be set up to carry pseudo wires.

In some cases, it might also be possible to have an inter-area TE LSP from R0 to R5 transiting via the backbone area (or any other levels with IS-IS). There may be cases where there are no longer enough resources on any intra area path R0-to-R5, and where there is a feasible inter-area path through the backbone area.

## 7. Detailed Requirements for Inter-Area MPLS TE

### 7.1. Inter-Area MPLS TE Operations and Interoperability

The inter-area MPLS TE solution MUST be consistent with requirements discussed in [TE-REQ], and the derived solution MUST interoperate seamlessly with current intra-area MPLS TE mechanisms and inherit its capability sets from [RSVP-TE].

The proposed solution MUST allow provisioning at the head-end with end-to-end RSVP signaling (potentially with loose paths) traversing across the interconnected ABRs, without further provisioning required along the transit path.

### 7.2. Inter-Area TE-LSP Signaling

The solution MUST allow for the signaling of inter-area TE LSPs, using RSVP-TE.

In addition to the signaling of classical TE constraints (bandwidth, admin-groups), the proposed solution MUST allow the head-end LSR to specify a set of LSRs explicitly, including ABRs, by means of strict or loose hops for the inter-area TE LSP.

In addition, the proposed solution SHOULD also provide the ability to specify and signal certain resources to be explicitly excluded in the inter-area TE-LSP path establishment.

### 7.3. Path Optimality

In the context of this requirement document, an optimal path is defined as the shortest path across multiple areas, taking into account either the IGP or TE metric [METRIC]. In other words, such a path is the path that would have been computed by making use of some CSPF algorithm in the absence of multiple IGP areas.

As mentioned in Section 5.2, the solution SHOULD provide the capability to compute an optimal path dynamically, satisfying a set of specified constraints (defined in [TE-REQ]) across multiple IGP areas. Note that this requirement document does not mandate that all inter-area TE LSPs require the computation of an optimal (shortest) inter-area path. Some inter-area TE-LSP paths may be computed via some mechanisms that do not guarantee an optimal end-to-end path, whereas some other inter-area TE-LSP paths carrying sensitive traffic could be computed by making use of mechanisms allowing an optimal end-to-end path to be computed dynamically. Note that regular constraints such as bandwidth, affinities, IGP/TE metric optimization, path diversity, etc., MUST be taken into account in the computation of an optimal end-to-end path.

### 7.4. Inter-Area MPLS-TE Routing

As mentioned in Section 5.3, IGP hierarchy does not allow the head-end LSR to compute an end-to-end optimal path. Additional mechanisms are required to compute an optimal path. These mechanisms MUST not alter the IGP hierarchy principles. Particularly, in order to maintain containment of routing information and to preserve the overall IGP scalability, the solution SHOULD avoid any dynamic-TE-topology-related information from leaking across areas, even in a summarized form.

Conversely, this does not preclude the leaking of non-topology-related information that is not taken into account during path selection, such as static TE Node information (TE router ids or TE node capabilities).

### 7.5. Inter-Area MPLS-TE Path Computation

Several methods may be used for path computation, including the following:

- Per-area path computation based on ERO expansion on the head-end LSR and on ABRs, with two options for ABR selection:
  - 1) Static configuration of ABRs as loose hops at the head-end LSR.
  - 2) Dynamic ABR selection.
- Inter-area end-to-end path computation, which may be based on (for instance) a recursive constraint-based searching thanks to collaboration between ABRs.

Note that any path computation method may be used provided that it respect key objectives pointed out in Section 5.3.

If a solution supports more than one method, it should allow the operator to select by configuration, and on a per-LSP basis, the desired option.

### 7.6. Inter-Area Crankback Routing

Crankback routing, as defined in [CRANKBACK], may be used for inter-area TE LSPs. For paths computed thanks to ERO expansions with a dynamic selection of downstream ABRs, crankback routing can be used when there is no feasible path from a selected downstream ABR to the destination. The upstream ABR or head-end LSR selects another downstream ABR and performs ERO expansion.

Note that this method does not allow computing an optimal path but just a feasible path. Note also that there can be  $O(N^2)$  LSP setup failures before finding a feasible path, where  $N$  is the average number of ABR between two areas. This may have a non-negligible impact on the LSP setup delay.

Crankback may also be used for inter-area LSP recovery. If a link/node/SRLG failure occurs in the backbone or tail-end area, the ABR upstream to the failure computes an alternate path and reroutes the LSP locally.

An inter-area MPLS-TE solution MAY support [CRANKBACK]. A solution that does, MUST allow [CRANKBACK] to be activated/deactivated via signaling, on a per-LSP basis.

### 7.7. Support of Diversely-Routed Inter-Area TE LSPs

There are several cases where the ability to compute diversely-routed TE-LSP paths may be desirable. For instance, in the case of LSP protection, primary and backup LSPs should be diversely routed. Another example is the requirement to set up multiple diversely-routed TE LSPs between a pair of LSRs residing in different IGP areas. For instance, when a single TE LSP satisfying the bandwidth constraint cannot be found between two end-points, a solution would consist of setting up multiple TE LSPs so that the sum of their bandwidth satisfy the bandwidth requirement. In this case, it may be desirable to have these TE LSPs diversely routed in order to minimize the impact of a failure, on the traffic between the two end-points.

Thus, the solution MUST be able to establish diversely-routed inter-area TE LSPs when diverse paths exist. It MUST support all kinds of diversity (link, node, SRLG).

The solution SHOULD allow computing an optimal placement of diversely-routed LSPs. There may be various criteria to determine an optimal placement. For instance, the placement of two diversely routed LSPs for load-balancing purposes may consist of minimizing their cumulative cost. The placement of two diversely-routed LSPs for protection purposes may consist of minimizing the cost of the primary LSP while bounding the cost or hop count of the backup LSP.

### 7.8. Intra/Inter-Area Path Selection Policy

For inter-area TE LSPs whose head-end and tail-end LSRs reside in the same IGP area, there may be intra-area and inter-area feasible paths. If the shortest path is an inter-area path, an operator either may want to avoid, as far as possible, crossing area and thus may prefer selecting a sub-optimal intra-area path or, conversely, may prefer to use a shortest path, even if it crosses areas. Thus, the solution should allow IGP area crossing to be enabled/disabled, on a per-LSP basis, for TE LSPs whose head-end and tail-end reside in the same IGP area.

### 7.9. Reoptimization of Inter-Area TE LSP

The solution MUST provide the ability to reoptimize in a minimally disruptive manner (make before break) an inter-area TE LSP, should a more optimal path appear in any traversed IGP area. The operator should be able to parameterize such a reoptimization according to a timer or event-driven basis. It should also be possible to trigger such a reoptimization manually.

The solution SHOULD provide the ability to reoptimize an inter-area TE LSP locally within an area; i.e., while retaining the same set of transit ABRs. The reoptimization process in that case MAY be controlled by the head-end LSR of the inter-area LSP, or by an ABR. The ABR should check for local optimality of the inter-area TE LSPs established through it on a timer or event driven basis. The option of a manual trigger to check for optimality should also be provided.

In some cases it is important to restrict the control of reoptimization to the Head-End LSR only. Thus, the solution MUST allow for activating/deactivating ABR control of reoptimization, via signaling on a per LSP-basis.

The solution SHOULD also provide the ability to perform an end-to-end reoptimization, potentially resulting in a change on the set of transit ABRs. Such reoptimization can only be controlled by the Head-End LSR.

In the case of head-end control of reoptimization, the solution SHOULD provide the ability for the inter-area head-end LSR to be informed of the existence of a more optimal path in a downstream area and keep a strict control over the reoptimization process. Thus, the inter-area head-end LSR, once informed of a more optimal path in some downstream IGP areas, could decide to perform a make-before-break reoptimization gracefully (or not to), according to the inter-area TE-LSP characteristics.

## 7.10. Inter-Area LSP Recovery

### 7.10.1. Rerouting of Inter-Area TE LSPs

The solution MUST support rerouting of an inter-area TE LSP in case of SRLG/link/node failure or preemption. Such rerouting may be controlled by the Head-End LSR or by an ABR (see Section 7.6, on crankback).

### 7.10.2. Fast Recovery of Inter-Area TE LSP

The solution MUST provide the ability to benefit from fast recovery, making use of the local protection techniques specified in [FAST-REROUTE] both in the case of an intra-area network element failure (link/SRLG/node) and in that of an ABR node failure. Note that different protection techniques SHOULD be usable in different parts of the network to protect an inter-area TE LSP. This is of the utmost importance, particularly in the case of an ABR node failure, as this node typically carries a great deal of inter-area traffic. Moreover, the solution SHOULD allow computing and setting up a backup tunnel following an optimal path that offers bandwidth guarantees

during failure, along with other potential constraints (such as bounded propagation delay increase along the backup path).

The solution SHOULD allow ABRs to be protected, while providing the same level of performances (recovery delay, bandwidth consumption) as provided today within an area.

Note that some signaling approaches may have an impact on FRR performances (recovery delay, bandwidth consumption). Typically, when some intra-area LSPs (LSP-Segment, FA-LSPs) are used to support the inter-area TE LSP, the protection of ABR using [FAST-REROUTE] may lead to higher bandwidth consumption and higher recovery delays. The use of [FAST-REROUTE] to protect ABRs, although ensuring the same level of performances, currently requires a single end-to-end RSVP session (contiguous LSP) to be used, without any intra-area LSP. Thus, the solution MUST provide the ability, via signalling on a per-LSP basis, to allow or preclude the use of intra-area LSPs to support the inter-area LSPs.

#### 7.11. DS-TE support

The proposed inter-area MPLS TE solution SHOULD also satisfy core requirements documented in [DSTE-REQ] and interoperate seamlessly with current intra-area MPLS DS-TE mechanism [DSTE-PROTO].

#### 7.12. Hierarchical LSP Support

In the case of a large inter-area MPLS deployment, potentially involving a large number of LSRs, it may be desirable/necessary to introduce some level of hierarchy in order to reduce the number of states on LSRs (such a solution implies other challenges). Thus, the proposed solution SHOULD allow inter-area TE-LSP aggregation (also referred to as LSP nesting) so that individual TE LSPs can be carried onto one or more aggregating LSPs. One such mechanism, for example, is described in [LSP-HIER].

#### 7.13. Hard/Soft Preemption

As defined in [MPLS-PREEMPT], two preemption models are applicable to MPLS: Soft and Hard Preemption.

An inter-area MPLS-TE solution SHOULD support the two models.

In the case of hard preemption, the preempted inter-area TE LSP should be rerouted, following requirements defined in Section 7.10.1.

In the case of soft preemption, the preempted inter-area TE LSP should be re-optimized, following requirements defined in Section 7.9.

#### 7.14. Auto-Discovery of TE Meshes

A TE mesh is a set of LSRs that are fully interconnected by a full mesh of TE LSPs. Because the number of LSRs participating in some TE mesh might be quite large, it might be desirable to provide some discovery mechanisms allowing an LSR to discover automatically the LSRs members of the TE mesh(es) that it belongs to. The discovery mechanism SHOULD be applicable across multiple IGP areas, and SHOULD not impact the IGP scalability, provided that IGP extensions are used for such a discovery mechanism.

#### 7.15. Inter-Area MPLS TE Fault Management Requirements

The proposed solution SHOULD be able to interoperate with fault detection mechanisms of intra-area MPLS TE.

The solution SHOULD support [LSP-PING] and [MPLS-TTL].

The solution SHOULD also support fault detection on backup LSPs, in case [FAST-REROUTE] is deployed.

#### 7.16. Inter-Area MPLS TE and Routing

In the case of intra-area MPLS TE, there are currently several possibilities for routing traffic into an intra-area TE LSP. They are listed in Section 4.2.

In the case of inter-area MPLS TE, the solution MUST support static routing into the LSP, and also BGP recursive routing with a static route to the BGP next-hop address.

ABRs propagate IP reachability information (summary LSA in OSPF and IP reachability TLV in ISIS), that MAY be used by the head-end LSR to route traffic to a destination beyond the TE-LSP tail-head LSR (e.g., to an ASBR).

The use of IGP shortcuts MUST be precluded when TE-LSP head-end and tail-end LSRs do not reside in the same IGP area. It MAY be used when they reside in the same area.

The advertisement of an inter-area TE LSP as a link into the IGP, in order to attract traffic to an LSP source, MUST be precluded when TE-LSP head-end and tail-end LSRs do not reside in the same IGP area. It MAY be used when they reside in the same area.



## 8. Evaluation criteria

### 8.1. Performances

The solution will be evaluated with respect to the following criteria:

- (1) Optimality of the computed inter-area TE-LSP primary and backup paths, in terms of path cost.
- (2) Capability to share bandwidth among inter-area backup LSPs protecting independent facilities.
- (3) Inter-area TE-LSP setup time (in msec).
- (4) RSVP-TE and IGP scalability (state impact, number of messages, message size).

### 8.2. Complexity and Risks

The proposed solution SHOULD not introduce complexity to the current operating network to such a degree that it would affect the stability and diminish the benefits of deploying such a solution over SP networks.

### 8.3. Backward Compatibility

In order to allow for a smooth migration or co-existence, the deployment of inter-area MPLS TE SHOULD not affect existing MPLS TE mechanisms. In particular, the solution SHOULD allow the setup of an inter-area TE LSP among transit LSRs that do not support inter-area extensions, provided that these LSRs do not participate in the inter-area TE procedure. For illustration purposes, the solution MAY require inter-area extensions only on end-point LSRs, on ABRs, and, potentially, on Points of Local Repair (PLR) protecting an ABR.

## 9. Security Considerations

This document does not introduce new security issues beyond those inherent in MPLS TE [RSVP-TE] and an inter-area MPLS-TE solution may use the same mechanisms proposed for that technology. It is, however, specifically important that manipulation of administratively configurable parameters be executed in a secure manner by authorized entities.

## 10. Acknowledgements

We would like to thank Dimitri Papadimitriou, Adrian Farrel, Vishal Sharma, and Arthi Ayyangar for their useful comments and suggestions.

## 11. Contributing Authors

This document was the collective work of several authors. The text and content of this document was contributed by the editors and the co-authors listed below (the contact information for the editors appears in Section 14 and is not repeated below):

Ting-Wo Chung  
Bell Canada  
181 Bay Street, Suite 350,  
Toronto,  
Ontario, Canada, M5J 2T3

EEmail: ting\_wo.chung@bell.ca

Yuichi Ikejiri  
NTT Communications Corporation  
1-1-6, Uchisaiwai-cho,  
Chiyoda-ku, Tokyo 100-8019  
JAPAN

EEmail: y.ikejiri@ntt.com

Raymond Zhang  
Infonet Services Corporation  
2160 E. Grand Ave.  
El Segundo, CA 90025  
USA

EEmail: raymond\_zhang@infonet.com

Parantap Lahiri  
MCI  
22001 Loudoun Cty Pky  
Ashburn, VA 20147  
USA

EEmail: parantap.lahiri@mci.com

Kenji Kumaki  
KDDI Corporation  
Garden Air Tower  
Iidabashi, Chiyoda-ku,  
Tokyo 102-8460,  
JAPAN

EEmail: ke-kumaki@kddi.com

## 12. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.
- [TE-REQ] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.
- [DSTE-REQ] Le Faucheur, F. and W. Lai, "Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering", RFC 3564, July 2003.

## 13. Informative References

- [TE-OVW] Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., and X. Xiao, "Overview and Principles of Internet Traffic Engineering", RFC 3272, May 2002.
- [RSVP-TE] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [OSPF-TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [ISIS-TE] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.
- [TE-APP] Boyle, J., Gill, V., Hannan, A., Cooper, D., Awduche, D., Christian, B., and W. Lai, "Applicability Statement for Traffic Engineering with MPLS", RFC 3346, August 2002.
- [FAST-REROUTE] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [LSP-PING] Kompella, K., Pan, P., Sheth, N., Cooper, D., Swallow, G., Wadhwa, S., Bonica, R., "Detecting Data Plane Liveliness in MPLS", Work in Progress.
- [MPLS-TTL] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, January 2003.

- [LSP-HIER] Kompella, K., and Y. Rekhter, "LSP Hierarchy with Generalized MPLS TE", Work in Progress.
- [MPLS-RECOV] Sharma, V. and F. Hellstrand, "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery", RFC 3469, February 2003.
- [CRANKBACK] Farrel, A., Ed., "Crankback Signaling Extensions for MPLS Signaling", Work in Progress.
- [MPLS-DIFF] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.
- [DSTE-PROTO] Le Faucheur, F., et al., "Protocol Extensions for Support of Differentiated-Service-aware MPLS Traffic Engineering", Work in Progress.
- [DIFF-ARCH] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [DIFF-AF] Heinanen, J., Baker, F., Weiss, W., and J. Wroclawski, "Assured Forwarding PHB Group", RFC 2597, June 1999.
- [DIFF-EF] Davie, B., Charny, A., Bennet, J.C., Benson, K., Le Boudec, J., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002.
- [MPLS-PREEMPT] Farrel, A., "Interim Report on MPLS Pre-emption", Work in Progress.
- [METRIC] Le Faucheur, F., Uppili, R., Vedrenne, A., Merckx, P., and T. Telkamp, "Use of Interior Gateway Protocol (IGP) Metric as a second MPLS Traffic Engineering (TE) Metric", BCP 87, RFC 3785, May 2004.

## 14. Editors' Addresses

Jean-Louis Le Roux  
France Telecom  
2, avenue Pierre-Marzin  
22307 Lannion Cedex  
France

EEmail: jeanlouis.leroux@francetelecom.com

Jean-Philippe Vasseur  
Cisco Systems, Inc.  
300 Beaver Brook Road  
Boxborough, MA - 01719  
USA

EEmail: jpv@cisco.com

Jim Boyle

EEmail: jboyle@pdnets.com

## Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.