TEWG Working Group
Internet Draft

JL Le Roux,(Ed.)
France Telecom

JP Vasseur, (Ed.)
Cisco System Inc.

Jim Boyle, (Ed.)
PDNETs

Category: Informational
Expires: August 2004

February 2004

Requirements for Inter-area MPLS Traffic Engineering

draft-leroux-tewg-interarea-mpls-te-req-00.txt

Status of this Memo

Abstract

This document lists a detailed set of functional requirements for the support of inter-area MPLS Traffic Engineering (inter-area MPLS TE) which could serve as a guideline to develop the required set of protocol extensions.

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 RFC-2119.

0. Summary

(This section to be removed before publication.)

0.1. Related I-d's

This draft is actually a merge of two previous requirements IDs related to inter-area MPLS-TE requirements:

draft-boyle-tewg-interarea-reqts-01.txt
draft-leroux-tewg-inter-area-te-rqs-00.txt

Table of Contents

1. Introduction

The set of MPLS Traffic Engineering tools, defined in [RSVP-TE], [OSPF-TE] and [ISIS-TE], that supports the requirements defined in [TE-REQ], is used today by many network operators to achieve major Traffic Engineering objectives defined in [TE-OVW] and summarized below:

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

However, the current set of MPLS Traffic Engineering mechanisms have to date been limited to use within a single 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 ABR failures.

This document first addresses current uses of MPLS Traffic Engineering within a single IGP area. This helps, then, in discussing a set of functional requirements a solution must or should satisfy in order to support inter-area MPLS Traffic Engineering. Since 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. Contributing Authors

This document was the collective work of several. 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 10, and is not repeated below.):

Ting-Wo Chung
Bell Canada
181 Bay Street, Suite 350,
Toronto,
Ontario, Canada, M5J 2T3
Email: ting_wo.chung@bell.ca

Yuichi Ikejiri
NTT Communications Corporation
1-1-6, Uchisaiwai-cho,
Chiyoda-ku, Tokyo 100-8019
JAPAN
Email: y.ikejiri@ntt.com

Raymond Zhang
Infonet Services Corporation
2160 E. Grand Ave.
El Segundo, CA 90025
USA
Email: raymond_zhang@infonet.com

Parantap Lahiri
MCI
22001 loudoun Cty Pky
Ashburn, VA 20147
USA
E-mail: parantap.lahiri@mci.com

Kenji Kumaki
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460,
JAPAN
E-mail : ke-kumaki@kddi.com

3. Terminology

LSR: Label Switch Router

TE LSP: MPLS 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 both head-end LSR and tail-end LSR are in the same IGP area but the TE LSP transiting path may be across different IGP areas.

IGP area: OSPF area or IS-IS level.

ABR: Area Border Router, router used to connect two IGP areas (ABR in OSPF or L1/L2 router in IS-IS).

CSPF: Constraint-based Shortest Path First.

4. Current intra-area uses of MPLS Traffic Engineering

This section addresses capabilities and uses of MPLS-TE within a single IGP area. It first addresses various capabilities offered by these mechanisms and then lists various approaches to integrate MPLS-TE into routing. This section is intended to help defining the requirements for MPLS-TE extensions across multiple IGP areas.

4.1. Intra-area MPLS Traffic Engineering Applications

4.1.1. Intra-area resources optimization

MPLS-TE can be used within an area to redirect paths of aggregated flows away from over-utilized resources within a network topology. In a small scale, this may be done by explicitly configuring a path to be used between two routers. In a grander scale a mesh of LSPs between central points in a network can be established with the LSPs configured with paths defined statically in configuration or arrived at by an algorithm that determines the shortest path given constraints such as bandwidth or other administrative constraints. In this way, MPLS-TE allows for greater control of how traffic demands utilize a network topology. As mentioned in Section 1, uses to date have been limited to within a single IGP area.

Note also that TE LSPs allow to measure traffic matrix in a simple and scalable manner. Basically, 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.

4.1.2. 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 or over a DiffServ domain to contribute to the enforcement of a (set of) QoS policy(ies), 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 the edge of the network.

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 and defined in [DS-TE-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 satisfying various constraint including bandwidth. For the sake of illustration, if the IGP metrics reflects the propagation delay, it allows finding a minimum propagation delay path satisfying various constraints like bandwidth.

4.1.3. Fast recovery within an area

As traffic sensitive applications are deployed, one of the key requirements is to provide fast recovery mechanisms(on the order of tens of msecs) in case of network element failure, which cannot be achieved with current IGP.

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 one, called one-to-one backup, consists in setting up a detour LSP per protected LSP and per element to protect. The second one called facility-backup consists in 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).

4.2. Intra-area MPLS-TE and routing

There are several possibilities to direct traffic into intra-area TE LSPs:

1) Static routing to the LSP destination address or any other addresses,
2) Traffic to the destination of the TE LSP or somewhere beyond this destination from an IGP SPF perspective.
3) The LSP can be advertised as a link into the IGP to become part of IGP database for all nodes, and thus 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].
4) Traffic sent to a set of routes announced by a (MP-)BGP peer that is reachable through the TE-LSP by means of a single static route to the corresponding BGP next-hop address (2) or by means of IGP SPF (3). This is often called BGP recursive routing.

5. Problem Statement, Requirements and Objectives of inter-area MPLS-TE

5.1. Inter-Area Traffic Engineering Problem Statement

As described in section 1, 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. This is basically due to the fact that 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.

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.

Hence, 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 in 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 no lead to efficient traffic engineering. Actually such per-area TE approach might lead to sub-optimal resource utilization, by optimizing resources independently in each area. And 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.

5.2. Requirements for inter-area MPLS-TE

For the reasons mentioned above, it is highly desired to extend the current set of MPLS-TE mechanism across multiple IGP areas in order to support the intra area applications described in section 1 across areas.

Basically, the solution MUST allow setting up inter-area TE LSPs, ie LSPs whose path crosses at least two IGP areas.

Inter-area MPLS-TE extensions are highly desired to provide inter-area resources optimization, to provide strict inter-area QoS guarantees, and to provide fast recovery across areas, particularly in order to protect inter-area traffic against ABR failures.

It may be desired to compute inter-area shortest path that satisfy some bandwidth constraints or any other constraints, as currently possible within a single IGP area. For the sake of illustration, if the IGP metrics reflects the propagation delay, it may be needed to be able to find the optimal (shortest) path satisfying some constraints (i.e 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 path 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 having as key objectives to preserve IGP hierarchy concept, and to preserve routing and signaling scalability.

5.3.1. Preserve the IGP hierarchy concept

The absence of a full link state topology database makes the computation of an end-to-end 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 should not be compromised to allow inter-area traffic engineering. Information propagation for path-selection should continue to be localized. These requirements are summarized as follows:
The solution MUST entirely preserve the concept of IGP hierarchy. In other words, flooding of TE link information across areas MUST be avoided.

5.3.2. Preserve Scalability

Being able to achieve 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 MUST also not increase IGP load which could compromise IGP scalability. In particular, a solution satisfying those requirements MUST not require for 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 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 MPLS-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.

Particularly, as mentioned in 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 like bandwidth. This may require an additional routing mechanism, as base CSPF at head-end can not longer be used due to the lack of topology and resources information. Such routing mechanism MUST not compromise the scalability of the overall system.

6. Application Scenario

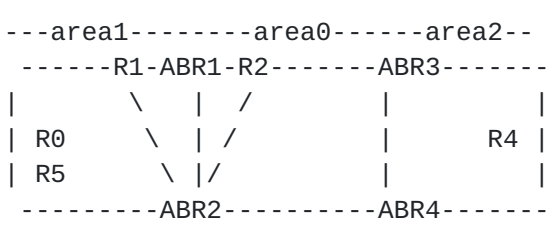

- ABR1, ABR2: Area0-Area1 ABRs
- ABR3, ABR4: Area0-Area2 ABRs

- 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 setup from R1 to R4 to route VoIP traffic classified as EF. Per-class inter-area constraint based routing allows to route the DS-TE LSP 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 setup to carry pseudo wire connections.

In some cases, it might also be possible to have an inter-area TE LSP from R0 to R5 transiting via the back-bone area (or any other levels with IS-IS). Basically, there may be cases where there is no longer enough resources on intra area path R0-to-R5, while 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 be such that it will 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 signalling (eventually with loose paths) traversing across the interconnected ABRs, without further provisioning required along the transit path.

7.2. Protocol signalling and path computation

The proposed solution MUST allow the head-end LSR to explicitly specify a set of LSRs, 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.

If multiple signalling methods are proposed in the solution (e.g. contiguous LSP, stitched or nested LSP), the head-end LSR MUST have the ability to signal the required signalling method on a per-LSP basis.

Several options may be used for path computations among those

- Per-area path computation based on ERO expansion with two options for ABR selection:
  - -Static loose hop ABR configuration at the head-end LSR.
  - -Dynamic loose hop ABR determination.
- Inter-area end-to-end path computation, that may be based for instance on 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 5.3

7.3. Path optimality

As already mentioned in 5.2, the solution SHOULD provide the capability for the head-end LSR to dynamically compute an optimal path 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 LSP require the computation of an optimal (shortest) inter-area path: some inter-area TE LSP path may be computed via some mechanisms not guaranteeing an optimal end to end path whereas some other inter-area TE LSP paths carrying sensitive traffic could be computed making use of some mechanisms allowing to dynamically compute an optimal end-to-end path. Note that regular constraints like bandwidth, affinities, IGP/TE metric optimization, path diversity, etc MUST also be taken into account in the computation of an optimal end-to-end path.

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. In other words, such a path is the path that would have been computed making use of some CSPF algorithm in the absence of multiple IGP areas.

Note that mechanism allowing to compute an optimal path are likely to consume more CPU resources than mechanisms computing only sub-optimal paths. So a solution should support both mechanism, and SHOULD allow the operator to select by configuration, and on a per-LSP basis, the required level of optimality.

7.4. Support of diversely routed inter-areas TE LSPs

There are several cases where the ability to compute diversely routed TE LSP paths may be desirable. For instance, in case of LSP protection, primary and backup LSPs should be diversely routed. Another example is the requirement to set up multiple TE LSPs between a pair of LSRs residing in different IGP areas in case a single TE LSP satisfying the set of requirements could not be found.

Hence, the solution SHOULD be able to provide the ability to compute diversely routed inter-area TE LSP paths. In particular, if such paths obeying the set of constraints exist, the solution SHOULD be able to compute them. For the sake of illustration, there are some algorithms that may not always allow to find diversely routed TE LSPs because they make use of a two steps approach that cannot guarantee to compute two diversely routed TE LSP paths even if such a solution exist. This is in contrast with other methods that simultaneously compute the set of diversely routed paths and that can always find such paths if they exist. Moreover, the solution SHOULD not require

extra-load in signalling and routing in order to reach that objective.

7.5. 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. In case the shortest path is an inter-area path, an operator may either want to avoid, as far as possible, crossing area and thus 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 MUST allow to enable or disable IGP area crossing, for TE LSPs whose head-end and tail-end reside in the same IGP area.

7.6. Reoptimization of inter-area TE LSP

The solution MUST provide the ability to reoptimize in a non 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 parameter such a reoptimization on a timer or event-driven basis. It should also be possible to trigger such a reoptimization manually.

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

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

In 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 on the reoptimization process. Hence, the inter-area head-end LSR, once informed of a more optimal path in some downstream IGP areas, could decide (or not) to gracefully perform a reoptimization, according to the inter-area TE LSP characteristics.

7.7. Failure handling and rerouting of an inter-area LSP.

In case of inter-area TE LSP failure in the backbone or tail-end area, it may be interesting to allow the ABR upstream to the failure to try to recover the LSP using a procedure such as defined in [CRANKBACK]. This may reduce the recovery delay, but with the risk of multiple crankbacks, and sub-optimality.
The solution SHOULD provide the ability to allow/disallow crankback via signalling on a per-LSP basis.

7.8. 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] in both the case of an intra-area network element failure (link/SRLG/Node) and 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 in particular in the case of an ABR node failure that 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 (like bounded propagation delay increase along the backup path).

7.9. 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].

7.10. Hierarchical LSP support

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

7.11. Soft preemption

The solution SHOULD support the ability to make use of the soft

preemption mechanisms for inter-area TE LSP such as one defined in [SOFT-PREEMPT].

7.12. Auto-discovery

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 automatically discover 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.

7.13. 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].

7.14. Inter-area MPLS-TE and routing

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

In case of inter-area MPLS-TE, the solution MUST support static routing into the LSP (1), and also BGP recursive routing with a static route to the BGP next-hop address.

ABRs propagate IP reacheability information (summary LSA in OSPF and IP reacheability 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 advertisement of an inter-area TE LSP as a link into the IGP, 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 SHOULD clearly be evaluated with respects to the following criteria:
(1) Optimality of the computed inter-area TE LSP path,
(2) Optimality of the computed backup tunnel path protecting against the failure of an ABR, capability to share bandwidth among backup tunnels protecting independent facilities.
(3) Inter-area TE LSP set up time,
(4) Scalability and state impact.

Other criteria may be added in further revisions of this document.

8.2. Complexity and risks

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

8.3. Backward Compatibility

The deployment of inter-area MPLS TE SHOULD not have impact on existing MPLS TE mechanisms to allow for a smooth migration or co-existence.

9. Security Considerations

Inter-area MPLS-TE does not raise any new security issue, beyond those of intra-area MPLS-TE.

10. Normative References

[TE-REQ], Awduche et. al., "Requirements for Traffic Engineering over MPLS", RFC2702, September 1999.

[OSPF-TE] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF Version 2", RFC3630, September 2003.

[ISIS-TE] Li, T., Smit, H., "IS-IS extensions for Traffic Engineering", draft-ietf-isis-traffic-04.txt (work in progress)

[RSVP-TE] Awduche, et al, "Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[FAST-REROUTE] Ping Pan, et al, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt, December 2003.

[LSPPING] Kompella, K., Pan, P., Sheth, N., Cooper, D.,Swallow, G., Wadhwa, S., Bonica, R., " Detecting Data Plane Liveliness in MPLS", Internet Draft <draft-ietf-mpls-lsp-ping-02.txt>, October 2002. (Work in Progress)

[MPLS-TTL] Agarwal, R., et al, "Time to Live (TTL) Processing in MPLS Networks", RFC 3443 Updates RFC 3032) ", January 2003.

[LSP-HIER] Kompella K., Rekhter Y., "LSP Hierarchy with Generalized MPLS TE", draft-ietf-mpls-lsp-hierarchy-08.txt, March 2002.

[MPLS-RECOV] V. Sharma, F. Hellstrand, "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery", RFC 3469, February 2003

[CRANKBACK] Farrel, A., Ed., "Crankback Signaling Extensions for MPLS Signalingö, draft-ietf-ccamp-crankback-01.txt, January 2004.

[DSTE-REQ] Le faucheur, F., et al, ô Requirements for Support of Differentiated Services-aware MPLS Traffic Engineeringö, RFC3564.

[DSTE-PROTO]Le faucheur, F., Ed., ôProtocol extensions for support of Differentiated-Service-aware MPLS Traffic Engineeringö, draft-ietf-tewg-diff-te-proto-06.txt, January 2004.

[SOFT-PREEMPT]Meyer, M., et al, ôMPLS Traffic Engineering Soft preemptionö, draft-ietf-mpls-soft-preemption-01.txt, October 2003.

11. Informative References

[MPLS-ARCH], Rosen, et. al., "Multiprotocol Label Switching Architecture", RFC 3031, January 2001

[DIFF-ARCH], Blake, et. al., "An Architecture for Differentiated Services", RFC 2475, December 1998

[DIFF-AF], Heinanen,et. al., "Assured Forwarding PHB Group", RFC 2597, June 1999.

[DIFF-EF], Davie, et. al., "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, March 2002

[MPLS-DIFF], Le Faucheur, et. al., "MPLS Support of Differentiated Services", RFC 3270, May 2002

[TE-OVW], Awduche, et. al., "Overview and Principles of Internet Traffic Engineering", RFC 3272,May 2002

[TE-APP], Boyle, et. al., "Applicability Statement of Traffic Engineering", RFC 3346, August 2002.

12. Editors' Address:

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

Jean-Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough , MA - 01719
USA
Email: jpv@cisco.com

Jim Boyle
Email: jboyle@pdnets.com

Full Copyright Statement