MPLS	J.-L. Le Roux, Ed.
Internet-Draft	T. Morin, Ed.
Intended status: Historic	V. Parfaitnotfirstpage
Expires: November 13, 2011	France Telecom - Orange
	L. Fangnotfirstpage
	Cisco
	L. Wangnotfirstpage
	Telenor
	Y. Kamitenotfirstpage
	NTT
	S. Amantenotfirstpage
	Level3
	May 12, 2011

Requirements for Point-To-Multipoint Extensions to the Label
Distribution Protocol
draft-ietf-mpls-mp-ldp-reqs-07

Abstract

This document lists a set of functional requirements that served as
input to the design of Label Distribution Protocol (LDP) extensions for
setting up point-to-multipoint (P2MP) Label Switched Paths (LSP), in
order to deliver point-to-multipoint applications over a Multi Protocol
Label Switching (MPLS) infrastructure.

Requirements Language

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].

Status of this Memo

Table of Contents

[1.](#) Contributing Authors

The following people contributed to the text and content of this document:

    *Shane Amante, Level 3 Communications, LLC

    *Luyuan Fang, Cisco Systems

    *Yuji Kamite, NTT Communications Corporation

    *Jean-Louis Le Roux, France Telecom - Orange

    *Thomas Morin, France Telecom - Orange

    *Vincent Parfait, France Telecom - Orange, Orange Business
     Services

    *Lei Wang, Telenor

[2.](#) Definitions

[2.1.](#) Acronyms

P2P: Point-To-Point

MP2P: Multipoint-to-Point

P2MP: Point-To-MultiPoint

MP2MP: MultiPoint-To-Multipoint

LSP: Label Switched Path

LSR: Label Switching Router

PE: Provider Edge router

P: Provider router

IGP: Interior Gateway Protocol

AS: Autonomous System

2.2. Terminology

The reader is assumed to be familiar with the terminology in [RFC3031], [RFC5036], and [RFC4026].

Ingress LSR:

Router acting as a sender of an LSP

Egress LSR:

Router acting as a receiver of an LSP

P2P LSP:

A LSP that has one unique Ingress LSR and one unique Egress LSR

MP2P LSP:

A LSP that has one or more Ingress LSRs and one unique Egress LSR

P2MP LSP:

A LSP that has one unique Ingress LSR and one or more Egress LSRs

MP2MP LSP:

A LSP that as one or more Leaf LSRs acting indifferently as Ingress or Egress LSR

Leaf LSR:

Egress LSR of a P2MP LSP or Ingress/Egress LSR of a MP2MP LSP

Transit LSR:

A LSR of a P2MP or MP2MP LSP that has one or more Downstream LSRs

Branch LSR:

A LSR of a P2MP or MP2MP LSP that has more than one downstream LSR

Bud LSR:

A LSR of a P2MP or MP2MP LSP that is an egress but also has one or more directly connected downstream LSR(s)

P2MP tree:

The ordered set of LSRs and links that comprise the path of a P2MP LSP from its ingress LSR to all of its egress LSRs.

3. Introduction

Note that this document captures requirements that served as input to the design of [I-D.ietf-mpls-ldp-p2mp] and is published with Historic status with the perspective of documenting the work done.

LDP [RFC5036] has been deployed for setting up point-to-point (P2P) and multipoint-to-point (MP2P) LSPs, in order to offer point-to-point services in MPLS backbones.

There are emerging requirements for supporting delivery of point-to-multipoint applications in MPLS backbones, such as those defined in [RFC4834] and [RFC5501].

For various reasons, including consistency with P2P applications, and taking full advantages of MPLS network infrastructure, it would be highly desirable to use MPLS LSPs for the delivery of multicast traffic. This could be implemented by setting up a group of P2P or MP2P LSPs, but such an approach may be inefficient since it would result in data replication at the ingress LSR and duplicate data traffic within the network. Hence new mechanisms are required that would allow traffic from an Ingress LSR to be efficiently delivered to a number of Egress LSRs in an MPLS backbone on a point-to-multipoint LSP (P2MP LSP), avoiding duplicate copies of a packet on a given link and with MPLS traffic replication at some Branch LSRs.

RSVP-TE extensions for setting up Point-To-Multipoint Traffic Engineered LSPs (P2MP TE LSPs), have been defined in [RFC4875]. They meet requirements expressed in [RFC4461]. This approach is useful, in network environments where P2MP Traffic Engineering capabilities are needed (Optimization, QoS, Fast recovery).

However for operators who want to support point-to-multipoint traffic delivery on an MPLS backbone, without Traffic Engineering needs, and who have already deployed LDP for P2P traffic, an interesting and useful approach would be to rely on LDP extensions in order to setup point-to-multipoint (P2MP) LSPs. This would bring consistency with P2P MPLS applications and would ease the delivery of point-to-multipoint services in an MPLS backbone.

This document focuses on the LDP approach for setting up P2MP LSPs. It lists a detailed set of requirements for P2MP extensions to LDP, so as to deliver P2MP traffic over a LDP-enabled MPLS infrastructure. The original intent was that these requirements should be used as guidelines when specifying LDP extensions.

Note that generic requirements for P2MP extensions to MPLS are out of the scope of this document. Rather this document describes solution specific requirements related to LDP extensions in order to set up P2MP LSPs.

Note also that other mechanisms could be used for setting up P2MP LSPs, such as for instance PIM extensions, but these are out of the scope of this document. The objective is not to compare these mechanisms but rather to focus on the requirements for an LDP extension approach.

The document is structured as follows:

  *Section 4 is an overview of the requirements;

  *Section 5 illustrates an application scenario;

  *Section 6 addresses detailed requirements for P2MP LSPs;

  *Section 7 finally discusses requirements for shared trees and
   MultiPoint-to-MultiPoint (MP2MP) LSPs.

4. Requirements overview

The P2MP LDP mechanism MUST support setting up P2MP LSPs, i.e. LSPs
with one Ingress LSR and one or more Egress LSRs, with traffic
replication at some Branch LSRs.
The P2MP LDP mechanism MUST allow the addition or removal of leaves
associated with a P2MP LSP.
The P2MP LDP mechanism MUST co-exist with current LDP mechanisms and
inherit its capability sets from [RFC5036]. It is of paramount
importance that the P2MP LDP mechanism MUST NOT impede the operation of
existing P2P/MP2P LDP LSPs. Also the P2MP LDP mechanism MUST co-exist
with P2P and P2MP RSVP-TE mechanisms [RFC3209], [RFC4875]. It is of
paramount importance that the P2MP LDP mechanism MUST NOT impede the
operation of existing P2P and P2MP RSVP-TE LSPs.
The P2MP LDP mechanism MAY also allow setting up multipoint-to-
multipoint (MP2MP) LSPs connecting a group of Leaf LSRs acting
indifferently as Ingress LSR or Egress LSR. This may allow a reduction
in the amount of LDP state that needs to be maintained by a LSR.

5. Application scenario

Figure 1 below illustrates an LDP enabled MPLS provider network, used
to carry both unicast and multicast traffic of VPN customers following
for instance the architecture defined in [I-D.ietf-l3vpn-2547bis-mcast]
for BGP/MPLS VPNs, or the one defined in [I-D.ietf-l2vpn-vpls-mcast].
In this example, a set of MP2P LDP LSPs are setup between Provider Edge
(PE) routers to carry unicast VPN traffic within the MPLS backbone.
And in this example a set of P2MP LDP LSPs are setup between PE routers
acting as Ingress LSRs and PE routers acting as Egress LSRs, so as to
support multicast VPN traffic delivery within the MPLS backbone.
For instance, a P2MP LDP LSP is setup between Ingress LSR PE1 and
Egress LSRs PE2, PE3, and PE4. It is used to transport multicast
traffic from PE1 to PE2, PE3 and PE4. P1 is a Branch LSR, it replicates
MPLS traffic sent by PE1 to P2, P3 and PE2. P2 and P3 are non-branch
transit LSRs, they forward MPLS traffic sent by P1 to PE3 and PE4
respectively.

```
        PE1
        *|                *** P2MP LDP LSP
        *|*****
        P1-----PE2
       */ \*
      */    \*
  *****/     \******
PE3----P2      P3----PE4
       |        |
       |        |
       |        |
      PE5      PE6
```
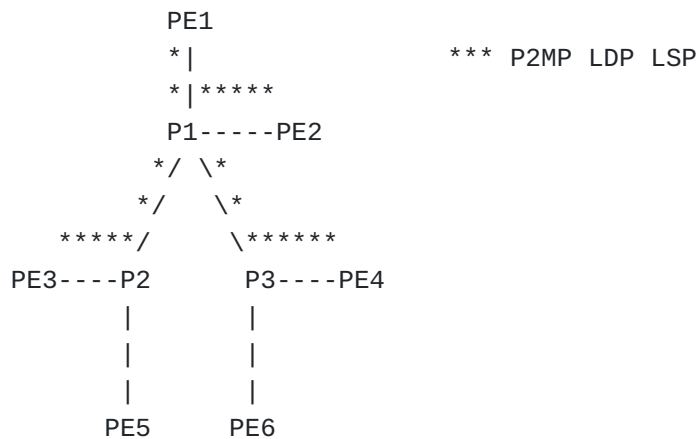
If later there are new receivers attached to PE5 and PE6, then PE5 and
PE6 join the P2MP LDP LSP. P2 and P3 become Branch LSRs and replicate
traffic received from P1, to PE3 and PE5, and to PE4 and PE6
respectively (see Figure 2 below).

```
        PE1
        *|                *** P2MP LDP LSP
        *|*****
        P1-----PE2
       */ \*
      */    \*
  *****/     \******
PE3----P2      P3----PE4
     *|        |*
     *|        |*
     *|        |*
      PE5      PE6
```
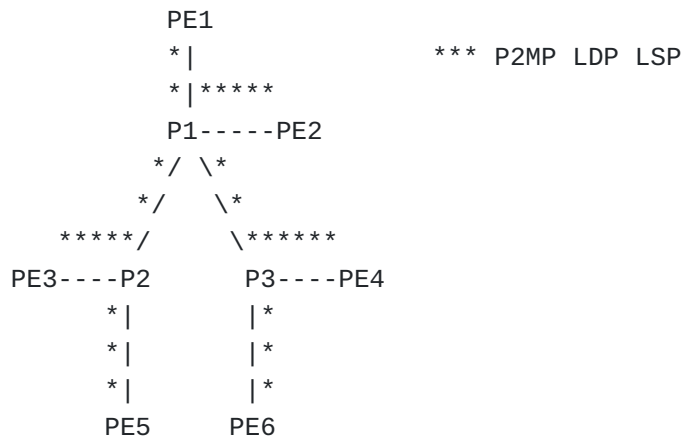
The above example is provided for the sake of illustration. Note that
P2MP LSPs ingress and egress LSRs may not necessarily be PE routers.
Also branch LSRs may not necessarily be P routers.

6. Detailed Requirements

6.1. P2MP LSPs

The P2MP LDP mechanism MUST support setting up P2MP LSPs. Data plane
aspects related to P2MP LSPs are those already defined in [RFC4461].
That is, a P2MP LSP has one Ingress LSR and one or more Egress LSRs.
Traffic sent by the Ingress LSR is received by all Egress LSRs. The
specific aspects related to P2MP LSPs is the action required at a
Branch LSR, where data replication occurs. Incoming labelled data is
appropriately replicated to several outgoing interfaces which may use
different labels.

An LSR SHOULD NOT send more than one copy of a packet on any given link of a P2MP LSP. Exceptions to this are mentioned in Section 6.9 and Section 6.18.
A P2MP LSP MUST be identified by a constant and unique identifier within the whole LDP domain, whatever the number of leaves, which may vary dynamically. This identifier will be used so as to add/remove leaves to/from the P2MP tree.

6.2. P2MP LSP FEC

As with P2P MPLS technology [RFC5036], traffic MUST be classified into a FEC in this P2MP extension. All packets which belong to a particular P2MP FEC and which travel from a particular node MUST use the same P2MP LSP.
If existing FECs cannot be used for this purpose, a new LDP FEC that is suitable for P2MP forwarding MUST be specified.

6.3. P2MP LDP routing

As with P2P and MP2P LDP LSPs, the P2MP LDP mechanism MUST support hop-by-hop LSP routing. P2MP LDP-based routing SHOULD rely upon the information maintained in LSR Routing Information Bases (RIB).
It is RECOMMENDED that the P2MP LSP routing rely upon the unicast route to the Ingress LSR to build a reverse path tree.

6.4. Setting up, tearing down and modifying P2MP LSPs

The P2MP LDP mechanism MUST support the establishment, maintenance and teardown of P2MP LSPs in a scalable manner. This MUST include both the existence of a large amount of P2MP LSPs within a single network and a large amount of leaf LSRs for a single P2MP LSP (see also Section 6.17 for scalability considerations and figures).
In order to scale well with a large number of leaves it is RECOMMENDED to follow a leaf-initiated P2MP LSP setup approach. For that purpose, leaves will have to be aware of the P2MP LSP identifier. The ways a Leaf LSR discovers P2MP LSPs identifiers rely on the applications that will use P2MP LSPs, and are out of the scope of this document.
The P2MP LDP mechanism MUST allow the dynamic addition and removal of leaves to and from a P2MP LSP, without any restriction (provided there is network connectivity). It is RECOMMENDED that these operations be leaf-initiated. These operations MUST NOT impact the data transfer (packet loss, duplication, delay) towards other leaves. It is RECOMMENDED that these operations do not cause any additional processing except on the path from the added/removed Leaf LSR to the Branch LSR.

6.5. Label Advertisement

The P2MP LDP mechanism MUST support downstream unsolicited label advertisement mode. This is well suited to a leaf-initiated approach and is consistent with P2P/MP2P LDP operations.
Other advertisement modes MAY also be supported.

6.6. Data Duplication

Data duplication refers to the receipt of multiple copies of a packet by any leaf. Although this may be a marginal situation, it may also be detrimental for certain applications. Hence, data duplication SHOULD as much as possible be avoided, and limited to (hopefully rare) transitory conditions.
Note, in particular, that data duplication might occur if P2MP LSP rerouting is being performed (See also Section 6.8).

6.7. Detecting and Avoiding Loops

The P2MP LDP extension MUST have a mechanism to detect routing loops. This MAY rely on extensions to the LDP Loop detection mechanism defined in [RFC5036]. A loop detection mechanism MAY require recording the set of LSRs traversed on the P2MP Tree. The P2MP loop avoidance mechanism MUST NOT impact the scalability of the P2MP LDP solution.
The P2MP LDP mechanism SHOULD have a mechanism to avoid routing loops in the data plane even during transient events.
Furthermore, the P2MP LDP mechanism MUST avoid routing loops in the data plane, that may trigger unexpected non-localized exponential growth of traffic.

6.8. P2MP LSP Re-routing

The P2MP LDP mechanism MUST support the rerouting of a P2MP LSP in the following cases:

    *Network failure (link or node);

    *A better path exists (e.g. new link, metric change);

    *Planned maintenance.

Given that P2MP LDP routing should rely on the RIB, the achievement of the following requirements relies on the underlying routing protocols (IGP, etc.).

6.8.1. Rerouting upon Network Failure

The P2MP LDP mechanism MUST allow for rerouting of a P2MP LSP in case of link or node failure(s), by relying upon update of the routes in the

RIB. The rerouting time SHOULD be minimized as much as possible so as to reduce traffic disruption.
A mechanism MUST be defined to prevent constant P2MP LSP teardown and rebuild which may be caused by the instability of a specific link/node in the network. This can rely on IGP dampening but may be completed by specific dampening at the LDP level.

6.8.2. Rerouting on a Better Path

The P2MP LDP mechanism MUST allow for rerouting of a P2MP LSP in case a better path is created in the network, for instance as a result of a metric change, a link repair, or the addition of links or nodes. This will rely on update of the routes in the RIB.

6.8.3. Rerouting upon Planned Maintenance

The P2MP LDP mechanism MUST support planned maintenance operations. It MUST be possible to reroute a P2MP LSP before a link/node is deactivated for maintenance purposes. Traffic disruption and data duplication SHOULD be minimized as much as possible during such planned maintenance. P2MP LSP rerouting upon planned maintenance MAY rely on a make before break procedure.

6.9. Support for multi-access networks

The P2MP LDP mechanism SHOULD provide a way for a Branch LSR to send a single copy of the data onto an interface to a multi-access network (e.g. an Ethernet LAN) and reach multiple adjacent downstream nodes. This requires that the same label be negotiated with all downstream LSRs for the LSP.
When there are several candidate upstream LSRs on an interface to a multi-access LAN, the P2MP LDP mechanism SHOULD provide a way for all downstream LSRs of a given P2MP LSP to select the same upstream LSR, so as to avoid traffic replication. In addition, the P2MP LDP mechanism SHOULD allow for an efficient balancing of a set of P2MP LSPs among a set of candidate upstream LSRs on a LAN interface.

6.10. Support for encapsulation in P2P and P2MP TE tunnels

The P2MP LDP mechanism MUST support nesting P2MP LSPs into P2P and P2MP TE tunnels.
The P2MP LDP mechanism MUST provide a way for a Branch LSR of a P2MP LSP, which is also a Head End LSR of a P2MP TE tunnel, to send a single copy of the data onto the tunnel and reach all downstream LSRs on the P2MP LSP, which are also Egress LSRs of the tunnel. As with LAN interfaces, this requires that the same label be negotiated with all downstream LSRs of the P2MP LDP LSP.

6.11. Label spaces

Labels for P2MP LSPs and P2P/MP2P LSPs MAY be assigned from shared or dedicated label spaces.
Note that dedicated label spaces will require the establishment of separate P2P and P2MP LDP sessions.

6.12. IPv4/IPv6 support

The P2MP LDP mechanism MUST support the establishment of LDP sessions over both IPv4 and IPv6 control planes.

6.13. Multi-Area/AS LSPs

The P2MP LDP mechanism MUST support the establishment of multi-area P2MP LSPs, i.e. LSPs whose leaves do not all reside in the same IGP area as the Ingress LSR. This SHOULD be possible without requiring the advertisement of Ingress LSRs' addresses across IGP areas.
The P2MP LDP mechanism MUST also support the establishment of inter-AS P2MP LSPs, i.e. LSPs whose leaves do not all reside in the same AS as the Ingress LSR. This SHOULD be possible without requiring the advertisement of Ingress LSRs' addresses across ASes.

6.14. OAM

LDP management tools ([RFC3815], etc.) will have to be enhanced to support P2MP LDP extensions. This may yield a new MIB module, which may possibly be inherited from the LDP MIB.
Built-in diagnostic tools MUST be defined to check the connectivity, trace the path and ensure fast detection of data plane failures on P2MP LDP LSPs.
Further and precise requirements and mechanisms for P2MP MPLS OAM purpose are out of the scope of this document and are addressed in [RFC4687].

6.15. Graceful Restart and Fault Recovery

LDP Graceful Restart mechanisms [RFC3478] and Fault Recovery mechanisms [RFC3479] SHOULD be enhanced to support P2MP LDP LSPs.

6.16. Robustness

A solution MUST be designed to re-establish connectivity for P2MP and MP2P LSPs in the event of failures, provided there exists network connectivity between ingress and egress nodes (i.e. designed without introducing single points of failure).

6.17. Scalability

Scalability is a key requirement for the P2MP LDP mechanism. It MUST be designed to scale well with an increase in the number of any of the following:

    *number of Leaf LSRs per P2MP LSP;

    *number of Downstream LSRs per Branch LSR;

    *number of P2MP LSPs per LSR.

In order to scale well with an increase in the number of leaves, it is RECOMMENDED that the size of a P2MP LSP state on a LSR, for one particular LSP, depend only on the number of adjacent LSRs on the LSP.

6.17.1. Orders of magnitude expected in operational networks

Typical orders of magnitude that we expect should be supported are:

    *tens of thousands of P2MP trees spread out across core network
     routers;

    *hundreds, or a few thousands, of leaves per tree;

See also section 4.2 of [RFC4834].

6.18. Backward Compatibility

In order to allow for a smooth migration, the P2MP LDP mechanism SHOULD offer as much backward compatibility as possible. In particular, the solution SHOULD allow the setup of a P2MP LSP along non-Branch Transit LSRs that do not support P2MP LDP extensions.
Also, the P2MP LDP solution MUST co-exist with current LDP mechanisms and inherit its capability sets from [RFC5036]. The P2MP LDP solution MUST NOT impede the operation of P2P/MP2P LSPs. A P2MP LDP solution MUST be designed in such a way that it allows P2P/MP2P and P2MP LSPs to be signalled on the same interface.

7. Shared Trees

For traffic delivery between a group of N LSRs that act as egress and/ or egress nodes on different P2MP flows, it may be useful to setup a shared tree connecting all these LSRs, instead of having N P2MP LSPs. This would reduce the amount of control and forwarding state that has to be maintained on a given LSR.
There are actually two main options for supporting such shared trees:

    *This could rely on the applications protocols that use LDP LSPs.
     A shared tree could consist of the combination of a MP2P LDP LSP

from Leafs LSRs to a given root node, with a P2MP LSP from this
root to Leaf LSRs. For instance with Multicast L3 VPN
applications, it would be possible to build a shared tree by
combining (see [I-D.ietf-l3vpn-2547bis-mcast]):

   -a MP2P unicast LDP LSP, from each PE of the group to a
    particular root PE acting as tree root,

   -with a P2MP LDP LSP from this root PE to each PE of the group.

*Or this could rely on a specific LDP mechanism allowing to setup
 multipoint-to-multipoint MPLS LSPs (MP2MP LSPs).

The former approach (Combination of MP2P and P2MP LSPs at the
application level) is out of the scope of this document while the
latter (MP2MP LSPs) belong to the scope of this document. Requirements
for the set up of MP2MP LSPs are listed below.

7.1. Requirements for MP2MP LSPs

A Multipoint-to-multipoint (MP2MP) LSP is an LSP connecting a group of
Leaf LSRs acting as Egress and/or Ingress LSRs. Traffic sent by any
Leaf LSR is received by all other Leaf LSRs of the group.
Procedures for setting up MP2MP LSPs with LDP SHOULD be specified. An
implementation that support P2MP LDP LSPs MAY also support MP2MP LDP
LSP.
The MP2MP LDP procedures MUST NOT impede the operations of P2MP LSP.
Requirements for P2MP LSPs, set forth in Section 6, apply equally to
MP2MP LSPs. Particular attention should be given on the below
requirements:

   *The solution MUST support recovery upon link and transit node
    failure and be designed to re-establish connectivity for MP2MP
    LSPs in the event of failures, provided there exists network
    connectivity between ingress and egress nodes (i.e. designed
    without introducing single points of failure);

   *The size of MP2MP state on a LSR, for one particular MP2MP LSP,
    SHOULD only depend on the number of adjacent LSRs on the LSP;

   *Furthermore, the MP2MP LDP mechanism MUST avoid routing loops
    that may trigger exponential growth of traffic. Note that this
    requirement is more challenging with MP2MP LSPs as a LSR may need
    to receive traffic for a given LSP on multiple interfaces.

There are additional requirements specific to MP2MP LSPs:

   *It is RECOMMENDED that a MP2MP MPLS LSP is built based on the
    unicast route to a specific LSR called root LSR;

*It is RECOMMENDED to define several root LSRs (e.g. a primary and
    a backup) to ensure redundancy upon root LSR failure;

   *The receiver SHOULD NOT receive back a packet it has sent on the
    MP2MP LSP;

   *The solution SHOULD avoid that all traffic between any pair of
    leaves is traversing a root LSR (similarly to PIM-Bidir trees)
    and SHOULD provide the operator with means to minimize the delay
    between two leaves;

   *It MUST be possible to check connectivity of a MP2MP LSP in both
    directions.

8. Evaluation criteria

8.1. Performance

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

(1) Efficiency of network resource usage;

(2) Time to add or remove a Leaf LSR;

(3) Time to repair a P2MP LSP in case of link or node failure;

(4) Scalability (state size, number of messages, message size).

Particularly the P2MP LDP mechanism SHOULD be designed with as key
objective to minimize the additional amount of state and additional
processing required in the network.
Also, the P2MP LDP mechanism SHOULD be designed so that convergence
times in case of link or node failure are minimized, in order to limit
traffic disruption.

8.2. Complexity and Risks

The proposed solution SHOULD NOT introduce complexity to the current
LDP operations to such a degree that it would affect the stability and
diminish the benefits of deploying such solution.

9. Security Considerations

It is expected that addressing the requirements defined in this
document should not introduce any new security issue beyond those
inherent to LDP, and that a P2MP LDP solution will rely on the security
mechanisms defined in [RFC5036] (e.g. TCP MD5 Signature).
An evaluation of the security features for MPLS networks may be found
in [RFC5920], and where issues or further work is identified by that

document, new security features or procedures for the MPLS protocols
will need to be developed.

10. IANA Considerations

This informational document makes no requests to IANA for action.

11. Acknowledgments

We would like to thank Christian Jacquenet, Hitoshi Fukuda, Ina Minei,
Dean Cheng, and Benjamin Niven-Jenkins, for their highly useful
comments and suggestions. We would also like to thank Adrian Farrel for
reviewing this document before publication.
We would also like to thank authors of [RFC4461] from which some of the
text in this document has been inspired.

12. References

12.1. Normative References

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
[RFC3031]	Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
[RFC5036]	Andersson, L., Minei, I. and B. Thomas, "LDP Specification", RFC 5036, October 2007.
[RFC3815]	Cucchiara, J., Sjostrand, H. and J. Luciani, "Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)", RFC 3815, June 2004.
[RFC3478]	Leelanivas, M., Rekhter, Y. and R. Aggarwal, "Graceful Restart Mechanism for Label Distribution Protocol", RFC 3478, February 2003.
[RFC3479]	Farrel, A., "Fault Tolerance for the Label Distribution Protocol (LDP)", RFC 3479, February 2003.
[RFC4461]	Yasukawa, S., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.

12.2. Informative References

[RFC4834]	Morin, T, "Requirements for Multicast in Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4834, April 2007.
[I-D.ietf-mpls-ldp-p2mp]	Minei, I, Wijnands, I, Kompella, K and B Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint

	Label Switched Paths", Internet-Draft draft-ietf-mpls-ldp-p2mp-13, April 2011.
[I-D.ietf-l3vpn-2547bis-mcast]	Aggarwal, R, Bandi, S, Cai, Y, Morin, T, Rekhter, Y, Rosen, E, Wijnands, I and S Yasukawa, "Multicast in MPLS/BGP IP VPNs", Internet-Draft draft-ietf-l3vpn-2547bis-mcast-10, January 2010.
[I-D.ietf-l2vpn-vpls-mcast]	Aggarwal, R, Kamite, Y, Fang, L and Y Rekhter, "Multicast in VPLS", Internet-Draft draft-ietf-l2vpn-vpls-mcast-08, October 2010.
[RFC4687]	Yasukawa, S., Farrel, A., King, D. and T. Nadeau, "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", RFC 4687, September 2006.
[RFC4875]	Aggarwal, R., Papadimitriou, D. and S. Yasukawa, "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
[RFC4026]	Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, March 2005.
[RFC3209]	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.
[RFC5920]	Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.
[RFC5501]	Kamite, Y., Wada, Y., Serbest, Y., Morin, T. and L. Fang, "Requirements for Multicast Support in Virtual Private LAN Services", RFC 5501, March 2009.

Authors' Addresses

Jean-Louis Le Roux editor Le Roux France Telecom - Orange EMail: jeanlouis.leroux@orange-ftgroup.com

Thomas Morin editor Morin France Telecom - Orange EMail: thomas.morin@orange-ftgroup.com

Vincent Parfait notfirstpage Parfait France Telecom - Orange, Orange Business Services EMail: vincent.parfait@orange-ftgroup.com

Luyuan Fang notfirstpage Fang Cisco Systems, Inc. EMail: lufang@cisco.com

Lei Wang notfirstpage Wang Telenor EMail: lei.wang@telenor.com

Yuji Kamite notfirstpage Kamite NTT Communications Corporation
EMail: y.kamite@ntt.com

Shane Amante notfirstpage Amante Level 3 Communications, LLC EMail:
shane@level3.net