Network Working Group
Internet Draft
Category: Informational
Expires: April 2006

J.-L. Le Roux
T. Morin
France Telecom

Vincent Parfait
Equant

Luyuan Fang
AT&T

Lei Wang
Telenor

Yuji Kamite
NTT Communications

Shane Amante
Level 3 Communications

October 2005

Requirements for point-to-multipoint extensions to the Label Distribution Protocol

draft-leroux-mpls-mp-ldp-reqs-02.txt

Status of this Memo

Abstract

This document lists a set of functional requirements for 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. It is intended that solutions that specify LDP procedures for setting up P2MP LSP satisfy these requirements.

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.

Table of Contents

1. Terminology

LSR: Label Switching Router

LSP: MPLS Label Switched Path

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

Leaf LSR: Egress LSR of a P2MP LSP

Transit LSR: A LSR of a P2MP LSP that has one or more downstream LSRs

Branch LSR: A LSR of a P2MP LSP that has more than one downstream LSR

Bud LSR: A LSR of a P2MP LSP that is an egress but also has one or more directly connected downstream LSRs

2. Introduction

Many operators have deployed LDP [LDP] 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 [L3VPN-MCAST-REQ] and [L2VPN-MCAST-REQ].

This requires mechanisms for setting up point-to-multipoint LSPs (P2MP LSP), i.e. LSPs with one Ingress LSR, a set of Egress LSRs, 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 [P2MP-TE-RSVP]. They meet requirements expressed in [P2MP-TE-REQ]. 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 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 applications 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. These requirements should be used as guidelines when specifying LDP extensions. It is intended that solutions that specify LDP procedures for P2MP LSP setup, satisfy these requirements.

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 3 points out the problem statement.

- Section 4 illustrates an application scenario.
- Section 5 addresses detailed requirements.
- Section 6 finally discusses group communication.

3. Problem Statement and Requirements Overview

3.1. Problem Statement

Many operators have deployed LDP [LDP] for setting up P2P and MP2P MPLS LSPs as PE-to-PE tunnels so as to carry point-to-point traffic essentially in Layer 3 and Layer 2 VPN networks.
There are emerging requirements for supporting multicast traffic delivery within these VPN infrastructures ([L3VPN-MCAST-REQ] and [L2VPN-MCAST-REQ]).
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 sub-optimal since it would result in data replication at the ingress LSR, and bandwidth inefficiency (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, avoiding duplicate copies of a packet on a given link.

Such efficient traffic delivery requires setting up P2MP LSPs. A P2MP LSPs is an LSP starting at an Ingress LSR, and ending on a set of one or more Egress LSRs. Traffic sent by the Ingress LSR is replicated on one or more Branch LSRs down to Egress LSRs.

RSVP-TE extensions for setting up P2MP TE LSPs, which meet requirements expressed in [P2MP-TE-REQ], have been defined in [P2MP-TE-RSVP]. This approach is useful, in network environments where Traffic Engineering capabilities are required.
However, for operators that deployed LDP for setting up PE-to-PE unicast MPLS LSPs, and without the need of traffic engineering, an interesting approach would be using LDP extensions for setting up P2MP LSPs.

The following gives a set of guidelines that a specification of LDP extensions for setting up P2MP LSPs should follow.

3.2. 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 arbitrary addition or removal of leaves associated with a P2MP LSP.

The P2MP LDP mechanism MUST interoperate seamlessly with existing P2P and MP2P LDP mechanisms.

It is of paramount importance that the P2MP LDP mechanism MUST NOT impede the operation of existing P2P/MP2P LSPs.

Note that 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 reducing the amount of LDP state to be maintained by a LSR. Detailed requirements for MP2MP LSPs are left for further study.

4. 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 [2547-MCAST] for BGP/MPLS VPNs, or the one defined in [VPLS-MCAST].

MP2P LDP LSPs are setup between PE routers to carry unicast VPN traffic.

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
```

Figure 1: P2MP LSP from PE1 to PE2, PE3, PE4.

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
```

Figure 2: Attachment of PE5 and PE6.

5. Detailed Requirements

5.1. P2MP LSPs

The P2MP LDP mechanism MUST support setting up P2MP LSPs.
Data plane aspects related to P2MP LSPs are those already defined in [P2MP-TE-REQ]. 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. Only one copy of a packet MUST be sent on a given link of a P2MP LSP.

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.

5.2. P2MP LSP FEC

As with P2P MPLS technology [LDP], 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.

As such, a solution MUST specify a FEC that is suitable for P2MP forwarding. Such P2MP FEC MUST be distinguished clearly from the exiting P2P FEC.

5.3. 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).

The P2MP LDP mechanism MUST NOT require the use of a multicast routing protocol for setting up P2MP LSPs.

It is RECOMMENDED that the P2MP LDP mechanism rely on unicast routing for setting up P2MP LSPs. In particular, the P2MP LSP routing SHOULD rely upon an IP shortest path to the Ingress LSR.

5.4. 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.

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. It is RECOMMENDED that these operations be leaf-initiated.
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.

5.5. Label Advertisement

The P2MP LDP mechanism SHOULD support downstream unsolicited label advertisement mode. This is well suited to a leaf-initiated approach and is consistent with P2P/MP2P LDP operations.

In order to follow a leaf initiated LSP setup approach, the P2MP LDP mechanism SHOULD support the Ordered label distribution control mode. Note that the Independent control mode is not relevant in a P2MP context, because the upstream LSRs cannot distribute labels independently like P2P/MP2P LDP, they must wait for label distribution from downstream LSRs.

Upstream label allocation ([MPLS-UPSTREAM]) may be particularly useful to avoid packet replication on LAN interfaces of a P2MP LSP, or when encapsulating the P2MP LSP into a P2MP TE tunnel.

Hence the P2MP LDP mechanism SHOULD also support upstream solicited label advertisement mode, where the solicitation is made by the downstream LSR, but the label is assigned by the upstream LSR. Note that the existing base LDP specification [RFC3036] does not specify upstream solicited label advertisement. Hence specific extensions SHOULD be defined.

5.6. 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 be avoided as much as possible, 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 5.6).

5.7. Avoiding loops

The P2MP LDP mechanism SHOULD have a mechanism to avoid routing loops even during transient events.

Furthermore, the P2MP LDP mechanism MUST avoid routing loops that may trigger unexpected non-localized exponential growth of traffic. Note that any loop-avoidance mechanism MUST respect scalability requirements.

5.8. 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 must rely on the RIB, the achievement of the following requirements implies also the underlying routing protocols (IGP).

5.8.1. Rerouting upon Network Failure

The P2MP LDP mechanism MUST allow for rerouting of a P2MP LSP in case of link or node failure(s). The rerouting time SHOULD be as much as possible minimized 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.

5.8.2. 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, or the addition of links or nodes.
Traffic disruption SHOULD be as much as possible minimized during such rerouting. It SHOULD be feasible to avoid packet loss during such rerouting.
Unnecessary data duplication during such rerouting SHOULD also be as much as possible minimized.

Note that there is likely to be a tension between packet loss minimization and packet duplication minimization objectives.

5.8.3. 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 SHOULD be as much as possible minimized during such rerouting. It SHOULD be feasible to avoid packet loss during such rerouting.
Unnecessary traffic duplication during such rerouting SHOULD also be as much as possible minimized.

Note that there is likely to be a tension between packet loss minimization and packet duplication minimization objectives.

5.9. Support for LAN interfaces

The P2MP LDP mechanism MUST provide a way for a Branch LSR to send a single copy of the data onto an Ethernet LAN interface and reach multiple adjacent downstream nodes. This requires that the same label be negotiated will all downstream LSRs for the LSP. In order to ease such negotiation an upstream label assignment approach may be used.

5.10. 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 LPS, 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 LDP label be negotiated with all downstream LSRs for the P2MP LDP LSP. In order to ease such negotiation, an upstream label assignment approach may be used.

5.11. Label spaces

Labels for P2MP LSPs and P2P/MP2P LSPs MAY be assigned from shared or dedicated label spaces.
MPLS Context Specific Label Spaces ([UPSTREAM-LABEL]) and particularly Upstream Node label spaces and Tunnel label spaces may be required to support upstream label assignment and avoid packet replication on LAN or P2MP TE Tunnel interfaces.

Note that dedicated label spaces will require the establishment of separate P2MP LDP sessions.

5.12. IPv4/IPv6 support

The P2MP LDP mechanism MUST be equally applicable to IPv4 and IPv6 traffic. Likewise, it SHOULD be possible to convey both kinds of traffic in a given P2MP LSP facility.

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

5.13. Multi-Area 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. This SHOULD be possible without requiring the advertisement of Leaf LSRs' addresses across IGP areas.

5.14. OAM

LDP management tools ([LDP-MIB]) MUST be enhanced to support P2MP LDP extensions. This may yield a new MIB module, which may possibly be inherited from the LDP MIB.

In order to facilitate correct management, P2MP LDP LSPs MUST have unique identifiers, otherwise it is impossible to determine which LSP is being managed.

There is a need for a mechanism to check the connectivity and trace the path of P2MP LDP LSPs. Note that the P2MP LSP Ping and Traceroute mechanisms, defined in [P2MP-LSP-PING], and allowing to check the connectivity and trace the path of P2MP RSVP-TE LSPs, could be enhanced to support P2MP LDP LSPs.

There is also a need for a mechanism allowing for 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

[P2MP-OAM-REQ].

5.15. Graceful Restart and Fault Recovery

LDP Graceful Restart mechanisms [LDP-GR] and Fault Recovery [LDP-FT] mechanisms SHOULD be enhanced to support P2MP LDP LSPs.

Particularly [LDP-GR] applies only to downstream unsolicited label distribution. Hence new mechanisms are required to account for upstream label assignment, particularly on LANs.

5.16. Robustness

A solution SHOULD avoid whatever single points of failures or propose some technical solutions for a failover mechanism.

5.17. 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.

5.17.1. Orders of magnitude of the expected numbers of P2MP LSPs 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 [L3VPN-MCAST-REQ].

5.18. 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 interoperate seamlessly with current LDP mechanisms and inherit its capability sets from [LDP]. 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.

6. Shared Trees

For traffic delivery between a group of N Leaf LSRs which are acting indifferently as Ingress or Egress LSRs, 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 all Leaf LSRs.
  For instance with Multicast L3 VPN applications, it would be possible to build a shared tree by combining:
    - 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 the root PE to all PEs of the group.

  (see also section 6.6 of [2547-MCAST]).

- 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.

6.1. MP2MP LSPs

Editorial note: There is currently no clear analysis of the gain of the MP2MP MPLS approach (with the potential impact on LDP), versus using application-level shared trees. This is why the requirement for MP2MP LSPs is currently optional

The P2MP LDP mechanism MAY allow setting up MP2MP LSP. A MP2MP LSP is a LSP connecting a group of Leaf LSRs acting indifferently as Ingress or Egress LSRs. Traffic sent by any Leaf LSRs is received by all other Leaf LSRs of the group. A sender LSR does not receive back the traffic sent.

All detailed requirements for P2MP LSPs listed in section 5, apply equally to MP2MP LSPs. Particularly it is RECOMMENDED that the size of a MP2MP state on a LSR, for one particular MP2MP LSP, depend only on the number of adjacent LSRs on the LSP, and not on the number of Leaf LSRs.

Additional detailed requirements specific to MP2MP LSPs are left for further study.

7. Evaluation criteria

7.1. Performances

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

(1) Time (in msec) to add or remove a Leaf LSR;
(2) Time (in msec) to repair a P2MP LSP in case of link or node failure;
(3) Scalability (state size, number of messages, message size).

Particularly, 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.

7.2. 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 P2MP LDP solution.

8. Security Considerations

This document does not introduce any new security issue beyond those inherent to LDP, and a P2MP LDP solution may rely on the security mechanisms defined in [LDP] (e.g. TCP MD-5).

9. Acknowledgments

We would like to thank Christian Jacquenet (France Telecom), Hitoshi Fukuda (NTT Communications), Ina Minei (Juniper) and Dean Cheng (Cisco Systems) for their highly useful comments and suggestions.

We would also like to thank authors of [P2MP-TE-REQ] from which some text of this document has been inspired.

10. References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3667] Bradner, S., "IETF Rights in Contributions", BCP 78, RFC 3667, February 2004.

[BCP79] Bradner, S., "Intellectual Property Rights in IETF Technology", RFC 3979, March 2005.

[LDP] L. Andersson, P. Doolan, N. Feldman, A. Fredette, B. Thomas, "LDP Specification", RFC 3036, January 2001

[L3VPN-MCAST-REQ] T. Morin, Ed., "Requirements for Multicast in L3 Provider-Provisioned VPNs", draft-ietf-l3vpn-ppvpn-mcast-reqts, work in progress.

[L2VPN-MCAST-REQ] Y. Kamite et al. " Requirements for Multicast Support in Virtual Private LAN Services", draft-kamite-l2vpn-vpls-mcast-reqts, work in progress

[P2MP-TE-REQ] S. Yasukawa, et. al., "Requirements for Point-to-Multipoint capability extension to MPLS", draft-ietf-mpls-p2mp-sig-requirement, work in progress.

[P2MP-TE-RSVP] R. Aggarwal, D. Papadimitriou, S. Yasukawa, et. al.., "Extensions to RSVP-TE for Point to Multipoint TE LSPs", draft-ietf-mpls-rsvp-te-p2mp, work in progress.

[MPLS-UPSTREAM-LABEL] R. Aggarwal, Y. Rekhter, E. Rosen, "MPLS Upstream Label Assignment and Context Specific Label Space", draft-raggarwa-mpls-upstream-label, work in progress.

[LDP-MIB] J. Cuchiarra et al. " Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)", RFC3815, June 2004.

[LDP-GR] M. Leelanivas, Y. Rekhter, R. Aggarwal, " Graceful Restart Mechanism for Label Distribution Protocol" RFC3478, February 2003.

[LDP-FT] A. Farrel, " Fault Tolerance for the Label Distribution Protocol (LDP)", RFC3479, February 2003.

[2547-MCAST] E. Rosen, R. Aggarwal, et. al., "Multicast in MPLS/BGP IP VPNs", draft-ietf-l3vpn-2547bis-mcast, work in progress.

[P2MP-OAM-REQ] S. Yasukawa, A. Farrel, D. King, T. Nadeau, "OAM Requirements for Point-To-Multipoint MPLS Networks", draft-yasukawa-mpls-p2mp-oam-reqs, work in progress.

[P2MP-LSP-PING] S. Yasukawa, A. Farrel, Z. Ali, B. Fenner, "Detecting data plane failures in P2MP MPLS-TE, Extensions to LSP-Ping", draft-ietf-mpls-p2mp-lsp-ping, work in progress.

[VPLS-MCAST] R.Aggarwal, Y Kamite, L Fang, VPLS Multicast draft-raggarwa-l2vpn-vpls-mcast-01.txt

11. Authors' Addresses:

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: jeanlouis.leroux@francetelecom.com

Thomas Morin
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
Email: thomas.morin@francetelecom.com

Vincent Parfait
EQUANT
1041 Route des Dolines
Sophia Antipolis
06560 Valbonne
FRANCE
Email: vincent.parfait@equant.com

Luyuan Fang
AT&T
200 Laurel Avenue
Middletown, NJ 07748
USA
Email: luyuanfang@att.com

Lei Wang
Telenor
Snaroyveien 30
Fornebu 1331
NORWAY
Email: lei.wang@telenor.com

Yuji Kamite
NTT Communications Corporation
Tokyo Opera City Tower
3-20-2 Nishi Shinjuku, Shinjuku-ku,
Tokyo 163-1421,
JAPAN
Email: y.kamite@ntt.com

Shane Amante
Level 3 Communications, LLC

1025 Eldorado Blvd
Broomfield, CO 80021
USA
Email: shane@level3.net

12. Intellectual Property Statement

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.