Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 16, 2014

Quintin Zhao
Tao Chou
Huawei Technology
Boris Zhang
Telus Communications
Emily Chen

July 15, 2013

P2MP Based mLDP Node Protection Mechanisms for mLDP LSP
draft-zhao-mpls-mldp-protections-05.txt

Abstract

This document specifies the procedures and protocol extensions for the protection of mLDP nodes within Multi-Protocol Label Switching (MPLS) networks using P2MP backup LSPs.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

To meet user demand, operators and service providers continue to deploy multicast applications using the Multicast Label Distribution Protocol (Multicast LDP - mLDP) across MPLS networks. For real-time applications, such as stock trading, on-line games, and multimedia teleconference, traditional node protection mechanisms, such as relying on IGP re-convergence to build the new Label Switched Path (LSP), fail to achieve a protection switching time less than that which is required to minimize the interruption of applications.

Instead of relying on IGP re-convergence to build the new LSP for failure protection, pre-computing and establishing a backup path before the failure delivers a better solution, allowing a more rapid switch over to the backup path when the protected node fails .

Providing a pre-computed backup path requires solutions to two complex problems:

- how to compute a completely disjoint backup path for each node in a multicast tree, and

- how to signal and setup the computed backup path.

For a primary Point-to-Multipoint (P2MP) Label Switched Path (LSP) created by LDP, there are several methods that could be chosen for creating a backup path:

- The use of an RSVP-TE (Resource Reservation Protocol - RSVP - Traffic Engineering) point-to-point (P2P) tunnel as a logical outgoing interface, which consequently utilizes the mature high-availability technology of RSVP-TE.

- The use of an alternative LDP P2P LSP as a packet encapsulation, which avoids the complex configuration of P2P RSVP-TE.

- Creating a P2MP backup LSP using a loop-free alternative route provided by the IGP.

- Using multi-topology technology, wherein the backup topology can be either statically configured or dynamically computed and signaled using IP/LDP Fast-Reroute mechanisms [I-D.ietf-rtgwg-mrt-frr-architecture].

When the backup path is available, there are two methods for packet forwarding and protection switch over:

Method 1 The traffic sender transmits the stream on both the primary and backup path always. Once the local traffic receiver detects a failure, the switch over will be relatively fast. However, the disadvantage of this method is that it consumes bandwidth because duplicate traffic will be sent on the primary and backup paths.

Method 2 The traffic sender transmits only on the primary path before the failure. Traffic is only forwarded to the MP through the backup path when failure is detected. Although bandwidth resource usage is minimized, cooperation is required to provide adequate switching times and to minimize higher-layer and application impact.

Ideally, if the switching time performance can be equal to or better than that of Method 1, it is reasonable to choose Method 2 to avoid bandwidth wastage. This consideration has been taken into account in making the recommendations in this document.

Note that for the computation and configuration of the primary topology, the algorithm used could be the Loop-free Alternate (LFA) based [RFC5286], Maximally Redundant Tree (MRT) based [I-D.ietf-rtgwg-mrt-frr-architecture] , or based on any other algorithms or methods available including static and offline tools; any such method can be used in conjunction with the mechanisms described in this document , which is limited to determining the nodes that should be spanned by the backup paths for the protection of a node in the primary multicast tree. In addition, the mechanism for detecting the node failure that will result in switchover to the backup pathis also outside the scope of this document.

Compared to a P2P LSP based solution, this P2MP LSP based solution not only uses mLDP mechanisms for both the primary path and backup paths, but also avoids unnecessary packet duplication where multiple P2P backup paths for the same node pass through.common nodes.

The remainder of this document specifies the signaling procedures and protocol extensions for the P2MP LSP based mLDP node protection solution which was briefly introduced above.

2. Terminology

This document uses terminology discussed in [RFC6388] and [I-D.ietf-mpls-ldp-multi-topology]. Additional key terms and terminology are listed here:

- Point of Local Repair (PLR): The node upon which traffic is redirected onto the preset backup path.

- Merge Point(MP):The node(s) where the primary path and the backup path rejoin and merge. Since the backup path is a multicast tree there will generally be more than one merge point.

- Primary Transit Node (PTN): The node between the PLR and MP on the primary path.

- Secondary Transit Node (STN): A node on the backup path between the PLR and MP. There might be more than one STN on the backup path between the PLR and MP nodes.

2.1. Requirement 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 RFC 2119 [RFC2119].

3. Requirements

A number of requirements have been identified that will allow the optimal set of mechanisms to be developed. These are:

- Computation of a possible disjoint (link and node) backup path within the multicast tree. In the case that there is no backup path available, there will be no backup path setup using the solution described in this draft will not be applicable.

- Minimize the PLR's switch over time from the primary path to the backup path when failure happens;

- Minimization of operation and maintenance cost;

- The solution should work without other protocol extensions other than the protocol extensions specified in this draft.

- The solution should work for all network topologies deployed in the users' network as long as there is a alternative backup path available.

4. Scope

This document specifies the signaling procedures and protocol extensions for the P2MP LSP based mLDP transit node protection

solution.

The method used for detecting the node failure is out the scope of this document.

Protection of the leaf and root nodes of the multicast tree is also out of scope of this document.

The protection mechanism for the case of multiple failures happening at the same time is out of scope of this document.

In the case that there is no backup path computed from the backup path computation algorithms, then there will be no backup path setup to protect the transit node failures.

5. Example of mLDP Node Protection

By using the procedures introduced in section 5 plus the available backup path computation algorithms, MP can initiate the building of the backup mLDP LSP starting from the PLR, avoiding the PTN and reaching the MPs. In the case of a large number of MPs, the solution introduced in this draft can avoid unnecessary packet duplication between PLR and the merge points. If a backup multicast tree is built rather than individual LSPs from the PLR to each MP then common transit points on the backup tree that would otherwise have multiple unicast LSPs passing through them will be saved some bandwidth on their incoming links.

```
                                     |
                                     V
                               +------------+ Point of Local Repair/
                               |     R1     | Switch over Point
                               +------------+ (Upstream LSR)
                                   /    \
                                10/      \20 (cost)
                                 /        \
                                V          V
             Primary     +----------+  +-----------+ Secondary
           Transit Node  |    R2    |  |     R3    |Transit Node
             (PTN)       +----------+  +-----------+  (STN)
                              |    \     /    |
                            10|   10\   /20   |20
                              |      \ /      |
                              |       \       |
                              |      / \      |
                              V     V   V     V
                         +-----------+ +-----------+ Merge Point
                         |    R4     | |    R5     | (Downstream LSR)
                         +-----------+ +-----------+
                              |               |
                              V               V
```

Figure 1: mLDP Local Protection using mLDP LSP Example

In Figure 1, R2 is on the preferential signalling/data path from R4/5 to R1, and the secondary signalling/data path from R4/R5 to R1 is through R3. In this case , the mLDP LSP will be established according to the IGP preferential path as R1--R2--R4/R5. As an example, this figure takes R2 as the PTN though the PTN can be any Transit Node of the mLDP LSP. (We assume that all the nodes in Figure 1 support this mLDP based node protection method.)

In this scenario, if R2 is protected by two P2P LSPs as R1--R3--R4 and R1--R3--R5 , the traffic will be duplicated on the link between R1 and R3 when the primary traffic is switched into the secondary path, and R3 will receive two copies of the multicast packages. But, if R2 is protected by a mLDP LSP instead, R3 will only receive one copy of the stream, and thus there will be no packet duplication on the link between R1 and R3 when the failure happens.

6. Signaling Procedures for P2MP Based Node Protection

This section introduces the signaling procedures of P2MP LSP's node protection using P2MP-based backup LSP for a Protected Node N.

6.1. The Computation of the Backup Path

Obviously, the backup path can not go through the protected node N. This section discusses how to choose the backup upstream LSR to avoid N.

Firstly, find the candidate upstream LSRs as below:

- MPs should preferentially choose the upstream LSRs on the shortest path as candidates, except node N. If no other upstream LSRs are on the shortest path, MPs should choose the next-hop on N's detour path as a candidate. The detour path can be an IGP-FRR path or other topology-based disjoint paths. The IGP-FRR path can be provided by LFA, U-Turn [[anchor9: EBD: Needs references.]], etc. The disjoint path can be provided by MT, MRT, etc. Choosing the candidates is a local decision and can be determined by configuration.

- The STN node MUST be chosen from the IGP next-hops on the shortest path toward the PLR within the topology specified in the FRR mLDP FEC element by the MT-ID (Multi-Topology Identifier) [[anchor10: EBD: Multi-Topology IDs need some explanation - it appears here without any introduction and I (for one) have no idea why there might be several and what they might cover.]] field. The candidate upstream LSRs MUST NOT include the PTN.

Thus, each node can choose one upstream node from the candidate upstream LSRs as its backup upstream LSR via the algorithm described in Section 2.4.1.1 of [RFC6388].

6.2. The Procedures for MP, PTN, STN and PLR

The procedures for P2MP Based Node Protection described in this document assumes the PLR is capable of node failure detection. In procedures described here covers the scenario where there is only one PTN between the PLR and MP. The cases where there are more than one PTNs between PLR and MP is out scope of this document.

We assume all the involved nodes have advertised their corresponding protection capabilities. And the following section specifies the signaling procedures of P2MP Based Node Protection.

6.2.1. MP's Procedures

Each non-Ingress LSR determines its own upstream LSR and sends out a label mapping message, in accordance with the procedures documented in [RFC6388] without any additional extension. And its upstream LSR will propagate a new label mapping message to its upstream LSR. In

such cases, the non-Ingress LSR is the MP node (as R4, R5 in Figure 1), MP's upstream LSR is the protected node (as R2 in Figure 1), and the protected node's upstream node is PLR (as R1 in Figure 1).

When one MP (as R4/R5 nodes in Figure 1) receives the Notification from the PTN after the MP has sent the label mapping message to the PTN, based on the PLR and PTN info, the MP individually determines its secondary path toward the PLR according to the IGP routes. The algorithms for choosing/computing the backup path can be LFA, MRT or others. After the backup upstream LSR is chosen, MP will send out a Label mapping message with the new FRR FEC (see section 7.3 for details), which includes the mLDP backup tree's key <PLR, protected-node, original-mLDP-FEC> and the MT-ID if the backup path is not in the default topology. Note that the label assigned for the primary path and the secondary path MUST be different to avoid having the MP feed primary traffic to its secondary path's downstream LSRs. In addition, the original-mLDP-FEC of the backup tree key is encoded in a special opaque value as introduced in section 7.3

6.2.2. PTN's Procedures

After the Protected Node (as R2 in Figure 1) determines its upstream LSR (as R1 in Figure 1), it will send the information (PLR's identify, mLDP FEC) via Notification messages to all its downstream nodes(MPs) immediately. If other LSRs become its downstream nodes later, it will send such announcements to its new MP(s).

6.2.3. STN's Procedures

When one node receives such aforementioned label mapping messages which inlcudes the mLDP FRR type of FECs, if it is not the PLR, it can consider itself a STN and will choose its backup upstream node toward PLR on the corresponding topology's shortest IGP path. To avoid the backup LSP going through the PTN, additional path selection rule(s) should be applied. A simple method is for the transit nodes to not choose the specified PTN as its upstream LSR for the backup LSP. Other methods, such as the not-via policy, are under study and will be added in the future. To make the primary and backup topologies rooted from PLR satisfy the 'maximum disjointed' requirement, they can either be configured through static configurations or be signaled dynamically through other mechanisms such as MRT.

When a STN on the backup mLSP fails before the backup LSP is put into use, this will trigger a recalculation of the backup LSP(s).

6.2.4. PLR's Procedures

When PLR(R1) receives the FRR label mapping message, it can identify that it is the PLR by the mLDP backup FEC elements. Thus, it will decode the special opaque value (which contains the primary mLDP FEC element, introduced in section 7.3) and generate the backup forwarding entry for the specific LSP, which is identified by the root address and opaque value in the special opaque value. It will also bind the backup forwarding state to the specific primary entry, which is indicated by the Protected Node address in the message. Note that there might be more than one backup forwarding entry for a specific protected node.

When failure is detected by PLR, it will switch the traffic to the backup paths. MP will also locally choose which traffic to receive and merge this traffic back to the primary LSP. The switch over manner on PLR is specified in detail in the later section of this document.

6.3. PLR's Switching Over Considerations

This section provides two modes for Switch over when failure occurs using the Protection Method 2 described in section 1 where there is only one copy of the traffic sent out from the PLR both before the PTN fails and after the PTN fails.

Depending on the capability of the MP node, MP node can set Node Failure (detection required) Flag in two modes.

Mode 1: The MP sets the Node Failure Required Flag (in the P2MP Based MP Node Protection Status Element) as 'Y'. This means that the MP requires the PLR to have the capability of detecting the PTN's failure. In this case, if the PLR doesn't have the node failure detection capability, then the backup path will not be setup and no protection is setup for the PTN. If the PLR has the capability of detecting the PTN's failure, the backup path can be setup correctly and only after PLR detects that PTN's failure, there will be backup traffic forwarded through the backup path to the MP(s).

Mode 2: The MP sets the Node Failure Required Flag, in the P2MP Based MP Node Protection Status Element, set as 'N'. This means that the MP has the capability of dropping duplicated multicast packages and doesn't require the PLR to have the capability of detecting the PTN's failure. In this case, PLR switches the traffic to the backup path once it detects the link failure between PLR and PTN no matter it is caused

by the PTN's failure or not. In the case that it is a link failure case, and the link protection is also deployed, then the MP will receive two copies of the traffic, one copy from the normal link protection path, and one copy from the node protection path through STN. MP must take the responsibility to drop one of the two duplicate traffics when link fails between PLR and PTN.

6.4. The Procedures after the Reconvergence

When Merge Point(s) see the next hop to Root changed, it/they will advertise the new mapping message(s), and the traffic will re-converge to the new primary path. MP will then withdraw the backup label after the re-convergence. STN will delete the specified backup LSP just as in the procedure of deleting normal P2MP LSP. And the entire backup P2MP LSP will be deleted when the node MP leaves the backup P2MP LSP.

6.5. Considerations for MP2MP LSP Node Protection

When a MP2MP LSP node needs to be protected, it can be treated with the same P2MP LSP node protection procedures for each forwarding direction.

```
                                   |
                                   V
                            +------------+ Point of Local Repair/
                            |     R1     | Switch over Point
                            +------------+ (Upstream LSR for Downstream flow)
                                /    \
                             10/      \20 (cost)
                              /        \
                             V          V
          Primary     +----------+  +-----------+ Secondary
        Transit Node  |    R2    |  |     R3    |Transit Node
          (PTN)       +----------+  +-----------+  (STN)
                           |    \     /    |
                         10|   10\   /20   |20
                           |      \ /      |
                           |       \       |
                           |      / \      |
                           V     V   V     V
                      +-----------+ +-----------+ Merge Point
                      |    R4     | |    R5     | (Downstream LSR
                      +-----------+ +-----------+  for Downstream flow)
                           |               |
                           V               V
```

Figure 2: MP2MP Example (R1 is the PLR)

```
                                   ^
                                   |
                            +------------+ Merge Point
                            |     R1     | (the downstream LSR for the upstream
flow)
                            +------------+
                                ^    ^
                             10/      \20 (cost)
                              /        \
          Primary     +----------+  +-----------+ Secondary
        Transit Node  |    R2    |  |     R3    |Transit Node
          (PTN)       +----------+  +-----------+  (STN)
                           ^   ^       ^   ^
                           |    \     /    |
                         10|   10\   /20   |20
                           |      \ /      |
                           |       \       |
                           |      / \      |
                      +-----------+ +-----------+
           PLR for    |    R4     | |    R5     |
    the upstream flow +-----------+ +-----------+
                           ^               ^
                           |               |
```

Figure 3: MP2MP Example (R4 is the PLR)

For each direction of MP2MP traffic flows (downstream in Figure 2 or upstream in Figure 3, the MP, PLR and MP nodes use the P2MP node protection procedures with the following additional considerations:

6.5.1. MP's Procedure

MP sends a backup label mapping message containing MP2MP downstream FRR FEC elements. When PLR receives a backup label mapping message with a MP2MP downstream flag, it sends the backup label mapping message with mp2mp upstream FRR FEC elements to Pn and then finally to MPs. This procedure just follows the normal MP2MP LSP procedure.

For the forwarding entries, MP node binds its primary MP2MP downstream NHLFE entry to backup MP2MP downstream ILM entry and binds its backup MP2MP upstream NHLFE entry to primary MP2MP upstream ILM entry.

For the forwarding entries, MP node binds its primary MP2MP downstream NHLFE entry to backup MP2MP downstream ILM entry and binds its backup MP2MP upstream NHLFE entry to primary MP2MP upstream ILM entry.

6.5.2. PLR's Procedure

PLR node binds its backup MP2MP downstream NHLFE entry to primary MP2MP downstream ILM entry, also binds its primary MP2MP upstream NHLFE entry to backup MP2MP upstream ILM entry.

6.5.3. Switch over Procedure

When the protected node fails, both the affected downstream and upstream nodes function as PLR and switch the downstream flow and upstream flow to their respective backup paths.

7. Protocol Extensions for mLDP Based Node Protection

Numerical fields in the formats defined in this section are encoded as unsigned integers in network octet and bit order.

7.1. mLDP Based MP Protection Capability Parameter TLV

A new Capability Parameter TLV is defined as mLDP Based MP Protection Capability for node protection. The format is illustrated as the following:

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |1|0| mLDP Based MP Prot.(IANA) |          Length (= 2)         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |S| Reserved    |
   +-+-+-+-+-+-+-+-+
```

Figure 4: mLDP Based MP Protection Capability

mLDP Based MP Prot.: TBA1 (to be assigned by IANA)

S: As specified in [RFC5561]

This is an unidirectional capability announcement.

An LSR, which supports the mLDP based protection procedures, should advertise mLDP Based MP Protection Capability TLV to its LDP

speakers. Without receiving this capability announcement, an LSR MUST NOT send any messages including the mLDP Based MP Node Protection Status Element and mLDP Backup FEC Element to its peer.

7.2. mLDP Based MP Node Protection Status Elements

A new type of LDP MP Status Value Element is introduced for notifying upstream LSR information. It is encoded as follows:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |mLDP FRR Type  |      Length                   |   Reserved    |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    ~                    PLR Node Address                           ~
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: FRR LDP MP Status Value Element

```
mLDP FRR Type:    Type TBA2 (to be assigned by IANA)

Length:           If the Address Family is IPv4, the Length MUST be
                  5;
                  if the Address Family is IPv6, the Length MUST be
                  17.

PLR Node Address: The host address of the PLR Node.
```

7.3. mLDP Backup FEC Element Encoding

A new type of mLDP backup FEC Element is introduced, it is used for notifying upstream LSR information. It is encoded as follows:

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |mLDP FEC T-FRR |         Address Family        | Address Length|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   ~                        PLR Node Address                       ~
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |N| Status code | FEC-Type      |        MT-ID                  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                     Protected Node Address                    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |    Opaque Length              |    Opaque Value ...           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               +
   ~                                                               ~
   |                                                               |
   |                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: mLDP Backup FEC Element

mLDP FEC Type-FRR: Type TBA3 (to be assigned by IANA)

Address Length: If the Address Family is IPv4, the Address Length MUST be 9;
if the Address Family is IPv6, the Address Length MUST be 33.

PLR Node Address: The host address of the PLR Node.

Protected Node Address: The host address of the Protected Node.

Status code: 1 = Primary path for traffic forwarding
2 = Secondary path for traffic forwarding

FEC-Type: 6 = P2MP FEC type
7 = MP2MP-up FEC type
8 = MP2MP-down FEC type

MT-ID: Multi-Topology ID. Unsigned 16 bit integer.

Opaque Length: The length of the opaque value, in octets.

Opaque Value: One or more MP opaque value elements, which is the same definition in [RFC6388]. For the FRR mLDP FEC element, the Opaque Value MUST be encoded as the Recursive Opaque Value,

which is defined in [RFC6512]. The value fields of the Recursive Opaque Value contain the original primary path's mLDP FEC element.

The encoding for the Recursive Opaque Value, as defined in [RFC6512], is shown in Figure 7.

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |  Type = 7     |         Length                |               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+               |
   ~                                                               ~
   |                 P2MP or MP2MP FEC Element                     |
   |                               +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7: Recursive Opaque Value

The Opaque Value is encoded by the MP node and decoded by PLR. Any other nodes on the path MUST NOT interpret the opaque value.

8. IANA Considerations

The document introduces following new protocol elements that require IANA consideration and assignments:

- o Code Points for "MP Protection Capability" TLV from the "LDP TLV Type Name Space" registry within the LDP Parameters

Registry:

Range/Value	Description
TBA1	MP Protection Capability (defined in section 7.1)

Figure 8: New Code Points for MP Protection Capability Extensions

- o Code Points for mLDP Based MP Node Protection Status Elements from the "LDP TLV Type Name Space" registry within the LDP Parameters

Registry:

Range/Value	Description
TBA2	Based MP Node Protection Status (defined in section 7.2)

Figure 9: Based MP Node Protectiin on Status

- o Code Points for mLDP FRR Type FEC from the the LDP registry "Forwarding Equivalence Class (FEC) Type Name Space". within the LDP Parameters

Registry:

Range/Value	Description
TBA3	mLDP FRR Type FEC (defined in section 7.3)

Figure 10: mLDP FRR Type FEC

9. Security Considerations

The same security considerations apply as for the base LDP specification, as described in [RFC5036]. The protocol extensions specified in this document do not provide any authorization mechanism for controlling the set of LSRs that may attempt to join a mLDP protection session. If such authorization is desirable, additional mechanisms outside the scope of this document are needed.

Note that authorization policies should be implemented and/or configured at all the nodes involved.

10. Acknowledgements

We would like to thank Nicolai Leymann and Daniel King for their valuable suggestions to this draft. We also would like to thank Robin Li, Lujun Wan, Kenji fujihira, Martin Vigoureux, Yaacov Weingarten, Eric Osborne and Elwyn Davis for their detailed comments and suggestions to the draft.

11. References

11.1. 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.

[RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, July 2009.

[RFC6348] Le Roux, JL. and T. Morin, "Requirements for Point-to-Multipoint Extensions to the Label Distribution Protocol", RFC 6348, September 2011.

[RFC6388] Wijnands, IJ., Minei, I., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, November 2011.

[RFC6512] Wijnands, IJ., Rosen, E., Napierala, M., and N. Leymann, "Using Multipoint LDP When the Backbone Has No Route to the Root", RFC 6512, February 2012.

11.2. Informative References

[I-D.ietf-mpls-ldp-multi-topology]
Zhao, Q., Fang, L., Zhou, C., Li, L., and K. Raza, "LDP Extensions for Multi Topology Routing", draft-ietf-mpls-ldp-multi-topology-08 (work in progress), May 2013.

[I-D.wijnands-mpls-mldp-node-protection]
Wijnands, I., Rosen, E., Raza, K., Tantsura, J., Atlas, A., and Q. Zhao, "mLDP Node Protection", draft-wijnands-mpls-mldp-node-protection-04 (work in progress), June 2013.

[I-D.ietf-rtgwg-mrt-frr-architecture]
Atlas, A., Kebler, R., Envedi, G., Csaszar, A., Tantsura, J., Konstantynowicz, M., and R. White, "An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees", draft-ietf-rtgwg-mrt-frr-architecture-03 (work in progress), July 2013.

[I-D.enyedi-rtgwg-mrt-frr-algorithm]
Atlas, A., Envedi, G., Csaszar, A., and A. Gopalan, "Algorithms for computing Maximally Redundant Trees for IP/LDP Fast- Reroute", draft-enyedi-rtgwg-mrt-frr-algorithm-02 (work in progress), October 2012.

Authors' Addresses

Quintin Zhao
Huawei Technology
125 Nagog Technology Park
Acton, MA 01719
US

Email: quintin.zhao@huawei.com

Tao Chou
Huawei Technology
156 Beiqing Rd
Haidian District, Beijing 100095
China

Email: tao.chou@huawei.com

Boris Zhang
Telus Communications
200 Consilium Pl Floor 15
Toronto, ON M1H 3J3
Canada

Phone:
Email: Boris.Zhang@telus.com

Emily Chen
2717 Seville Blvd, Apt 1205
Clearwater, FL 33764
US

Email: emily.chen220@gmail.com