

Network Working Group
Internet Draft
Category: Standards Track
Created: October 19, 2009
Expires: April 19, 2010

Seisho Yasukawa (Editor)
NTT

Supporting Multipoint-to-Point Label Switched Paths in
Multiprotocol Label Switching Traffic Engineering

draft-yasukawa-mpls-mp2p-rsvpte-06.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

Traffic engineered Multiprotocol Label Switching (MPLS-TE) uses Resource Reservation Protocol traffic engineering extensions (RSVP-TE) as the signaling protocol to establish Label Switched Paths (LSPs). Although the Resource Reservation Protocol (RSVP) on which RSVP-TE is built supports the convergence of traffic flows toward a common destination, this concept has not been exploited in MPLS-TE which has been limited to point-to-point (P2P) and point-to-multipoint (P2MP) LSPs.

This document describes extensions to RSVP-TE procedures and protocol elements to support multipoint-to-point (MP2P) LSPs.

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 [RFC2119].

Contents

1. Introduction	3
2. Description of MP2P MPLS-TE LSPs	3
2.1 Basic Requirements	5
2.1.1 Non-Requirements	7
2.2 Requirement for new Procedures and Extensions	7
3. Procedures for Signaling MP2P MPLS-TE LSPs	7
3.1 Ingress LSR Initial P2P LSP Establishment	8
3.2 Subsequent LSPs Creation	8
3.3 Path Processing at Merge Points	9
3.3.1 Determining that Merging is Allowed	9
3.3.2 Wildcard Filter Style	10
3.3.3 Fixed Filter Style	11
3.3.4 Shared Explicit Style	11
3.3.5 Record Route Processing on Merged Path Messages	12
3.4 Path Processing Downstream of a Merge Point	13
3.5 Processing at Egress LSRs	13
3.6 Resv Processing at Merge Points	13
3.7 LSP Teardown	13
3.8 Make-Before-Break Support	14
3.9 Fast Re-Route	14
4. Protocol Extensions	16
4.1 Allowing MP2P LSP Merging	16
4.2 Message Formats	16
4.2.1 Path Message	16
4.2.2 PathErr Message	17
5. Backward Compatibility	18
5.1 Legacy LSRs	18
5.2 Support of P2P and P2MP LSPs through an MP2P-enabled network ..	18
6. Path Computation Considerations	18
7. Manageability Considerations	18
7.1 MIB Modules	18
7.2 Operations and Management	19
8. Security Considerations	19
9. IANA Considerations	20
10. Acknowledgments	20
11. References	20
11.1 Normative References	20
11.2 Informative References	21
12. Author's Address	22

1. Introduction

The architecture for Multiprotocol Label Switching (MPLS) is described in [RFC3031]. The signaling for Label Switched Paths (LSPs) in MPLS traffic engineering (MPLS-TE) networks is achieved through specific extensions to the Resource Reservation Protocol (RSVP) defined in [RFC3209]. Although RSVP [RFC2205] includes support for the convergence of traffic flows toward a common destination, the traffic engineering (TE) extensions to RSVP (RSVP-TE) do not leverage this fact, and [RFC3209] only includes support for point-to-point (P2P) LSPs.

Additional extensions to RSVP-TE have been defined to support point-to-multipoint (P2MP) LSPs in [RFC4875].

Multipoint-to-point (MP2P) LSP tunnels offer certain benefits in a core MPLS network. For example, they reduce the amount of LSP state that protocol implementations must maintain within the network, and they may simplify the accounting and prioritization of traffic within individual routers.

This document describes how MP2P LSPs may be established using the protocol elements and procedures inherent in RSVP and RSVP-TE, and describes new procedures and protocol elements where necessary.

2. Description of MP2P MPLS-TE LSPs

MP2P MPLS-TE LSPs are transport tunnels within a core MPLS-TE network. An MP2P MPLS-TE LSP delivers traffic from multiple ingress Label Switching Routers (LSRs) to a single common egress LSR.

Identification of the traffic carried by the tunnel with its originating ingress LSR is obviously not within the scope of the MP2P LSP since the traffic arriving at the single egress of the MP2P LSP comes from more than one ingress. Thus the MP2P LSP must be treated as an edge-to-edge tunnel much in the manner of hierarchical P2P LSPs described in [RFC4206]. For example, directed signaling (MPLS-TE or MPLS using the Label distribution protocol - LDP [RFC5036]) can be used to exchange labels between egress and ingresses so that traffic from each ingress will use a different label on the label stack beneath the label used for the MP2P tunnel.

At each ingress, the MPLS-TE flow from ingress to egress is perceived as a single P2P LSP in the forwarding plane. But where two LSPs from different sources to the same destination pass through a common transit LSR they may be merged in the forwarding plane so that there is just a single downstream LSP. Thus, at the merge-point LSR there are two or more upstream LSP segments, and just one downstream LSP

segment.

LSP merging within the network can be carried out in an ad hoc way on any LSPs that are allowed to be merged. The judgment of whether two LSPs can be merged is dependent on:

- Are the LSPs to the same destination?
- Do the ingresses allow these LSPs to be merged?
- Are the explicit routes downstream of the merge point consistent with merging?
- Is the merge point capable of performing merging?
- Is the merge point willing to perform merging?
- Are the bandwidth implications of merging supported by the downstream network?

The ad hoc way of MP2P LSP establishment means that upstream segments of the MP2P tree are added when new ingress-to-egress LSPs are signaled and those ingress-to-egress LSPs coincide with (discover) other existing LSPs to the same egress that can be merged according to the above criteria.

Two models to manage the bandwidth of merged LSPs may be considered.

- The new upstream segment may be spliced into the MP2P LSP and the additional bandwidth required for the traffic from the new ingress can be added to the bandwidth request for the downstream segment. This is most appropriate for MP2P tunnels carrying unassociated traffic from the ingress LSRs.
- The new upstream segment may be spliced into the MP2P LSP and use (share) the existing bandwidth on the downstream segment. This may be appropriate for certain specific applications where individual ingresses send data one at a time (for example, voice conferencing).

Selection between these two models is critical to gain the benefit of avoiding excess bandwidth allocation where appropriate (the second case), but to avoid oversubscription of the downstream segments in other uses (the first case). LSP signaling that allows LSP merging clearly needs to indicate the type of bandwidth model appropriate for the LSPs that are merged.

Additionally, two models may be considered for how the egress is notified that merging has taken place at an upstream LSR.

- When a new upstream segment is spliced into an MP2P LSP the new ingress LSR may be reported to the egress so that it knows all of the ingress LSRs on the MP2P tree. This mechanism is suitable for some uses of the MP2P LSP such as a tunnel where forwarding adjacencies (FAs, see [RFC4206]) are managed between the end points.
- When a new upstream segment is spliced into an MP2P LSP no further downstream signaling is performed, the egress is not informed of the addition of a new ingress, and the ingress is told that it is connected to the egress. This may be particularly suited to bandwidth sharing applications.

The combination of these two sets of models gives four options for signaling MP2P LSPs downstream of a merge point.

2.1 Basic Requirements

This section lists the basic functional requirements that the signaling of MP2P MPLS-TE LSPs must support.

a. Control of merging function

The ingress LSR MUST be able to specify through signaling whether an LSP may be merged into an MP2P LSP or not.

b. Identification of LSPs that can be merged.

Transit LSRs that are capable of merging LSPs into the MP2P tree MUST be able to distinguish which LSPs can be merged and which are to remain distinct. Further, the transit LSRs MUST be able to determine sufficient information from the signaled LSPs to apply policy to determine whether merging is allowed.

c. Resource sharing downstream of merge points

It is a requirement that the single LSP segment downstream of a merge point MUST have sufficient bandwidth allocated to support all traffic from the LSP segments upstream of the merge point, and that that traffic MUST be able to share the bandwidth on the downstream segment.

d. Explicit route control

It is a requirement that the ingress LSR MUST be able to control

the route of the LSP from source to destination regardless of whether merging is performed at a downstream LSR.

e. Label allocation

Label allocation MUST be supported with a single label on a merged downstream segment and individual labels on upstream segments.

f. LSP identification

Each LSP MUST be uniquely identifiable within the network. Further, in order to support diagnostics and OAM, a single identifier MUST be used on the LSP for the whole of its path between any ingress and the egress. Additionally, a single identifier SHOULD be used for the MP2P LSP across the whole of the LSP tree.

g. Coordination between ingress LSRs

MP2P operation MUST be possible with minimal or zero coordination between ingress LSRs. It is not acceptable to require configuration at each ingress LSR in order to specify which other ingress LSRs may join the same MP2P LSP (for example, through identifying a group). It is acceptable to require:

- Specific configuration for an individual LSP that it will support MP2P merging (see requirement a.)
- Network-wide configuration of identifiers of "classes" of LSP that may be merged.

This requirement is hard to reconcile with previous requirements and with the assumed use of [RFC2205] and [RFC3209] procedures. Further discussion is provided in Section 3.

h. Identification of senders at egress.

It MUST be possible for the egress LSR to determine all of the ingress LSRs attached to the MP2P LSP. This requires that signaling information from each ingress is passed to the egress.

It SHOULD be possible to enable or disable this function using a signaling control determined by the ingress LSRs. Thus, if the LSP type is such that the egress does not need to know the identification of the ingresses (senders) then the ingresses MUST be able to signal this fact and the merge-point LSRs SHOULD act accordingly. The default position, however, is that the merge-point LSRs MUST inform the egress of all senders in the MP2P

tree.

i. Route recording

It MUST be possible to record the path followed by the MP2P LSP and pass this information to the end points. The ingress only requires to know the path from the ingress to the egress. The egress requires to know the path from each ingress to the egress.

The route reported to the ingress SHOULD indicate which LSRs on the route are currently acting as MP2P merge points.

Some form of information aggregation technique SHOULD be used in reporting the recorded route to the egress so that hops that are common in the tree from multiple ingress LSRs (that is, hops downstream of the merge point) SHOULD NOT be repeated in the information present in the signaling message (Path message).

2.1.1 Non-Requirements

No ingress is required to know of the existence of other ingresses in the MP2P LSP.

2.2 Requirement for new Procedures and Extensions

Section 2.1 summarizes the signaling requirements. Some of the requirements can be met using the merging techniques of [RFC2205] and the LSP identification techniques of [RFC3209], but the other requirements (specifically g, h and i) cannot be met using these techniques. Thus new procedures and extensions are needed.

The remainder of this document introduces new procedures and extensions to support signaling MP2P LSPs, and describes how the requirements are met using existing and new signaling mechanisms.

3. Procedures for Signaling MP2P MPLS-TE LSPs

The procedures described here are modeled on the reservation styles described in [RFC2205]. Further, many ideas are taken from the way that FlowSpecs and FilterSpecs may be combined on Resv messages in [RFC2205] in recognition that the Resv message in [RFC2205] is creating an MP2P flow tree.

Additions are made to allow control of this MP2P behavior on the Path message. The advantage of this over the [RFC2205] techniques is that merging is explicitly announced and is the responsibility of the merge point not of the egress.

3.1 Ingress LSR Initial P2P LSP Establishment

The initial LSP is signaled as a normal point-to-point LSP would be. The ingress LSR cannot know whether future upstream segments will be merged into this LSP or not.

However, the LSP is flagged by the ingress to say whether MP2P merging is allowed using a new flag in the LSP_ATTRIBUTES object [RFC5420] in the Path message. If this flag is not set or is absent, the LSP MUST be treated as a normal P2P LSP and merging MUST NOT be performed.

If MP2P merging is allowed by the ingress, the Path message MUST also include the Style object. The Style object was previously only allowed on a Resv message, [RFC2205]. The Style object directs the merging procedure that is performed at any merge point as described in Section 3.3. Note that the presence of the Style object in the Path message could be used to infer that merging is allowed, but it is better to use an explicit control (as described in the previous paragraph) to allow for flexible future use of the Style object. Therefore, an LSR MUST NOT assume that merging is allowed from the presence of the Style object on a Path message.

The ingress LSR sets all of the fields of the Session and Sender-Template objects as normal [RFC3209] except that the Extended Tunnel ID in the Session object MUST be set to a well-known value that is common across all LSPs that may be merged with this LSP. If any arbitrary merging is allowed then a common, network-wide value SHOULD be configured. If different classes of merging are defined, then multiple network-wide values SHOULD be defined and configured. The value of zero is RECOMMENDED to be used to allow general merging with other LSPs. Other non-zero values MAY be defined according to administrative policy for the network to restrict the LSPs that may be merged, perhaps according to service or application classification. It is RECOMMENDED that only values that do not match valid IP addresses are used for this function since [RFC3209] suggests the use of the sender's IP address in this field for P2P LSPs.

Other ingress LSR procedures and encodings are unchanged from [RFC3209].

3.2 Subsequent LSPs Creation

Subsequent LSPs are identical to initial LSPs described in the previous section. The ingress LSR is not aware whether other LSPs already exist and MP2P merging will happen. Therefore the ingress LSR builds its Path message as already described.

3.3 Path Processing at Merge Points

The main difference for MP2P LSP support takes place at the merge points. Much of the processing depends on the setting of the Style object carried in the received Path messages.

3.3.1 Determining that Merging is Allowed

When an LSR receives a Path message with the MP2P flag set in the LSP_ATTRIBUTES object indicating that MP2P merging is allowed for the LSP, and if the LSR is capable of supporting MP2P merging it performs the following checks to determine an existing LSP with which to merge. The order of the checks is implementation-specific.

- Search all LSP state for LSPs that are allowed to be merged (that is, that were signaled with the MP2P flag set in the LSP_ATTRIBUTES object).
- Find any LSPs that have a common destination set in the Session object and that have the same Extended Tunnel ID in the Session object. Note that the Tunnel ID in the Session object, and the Sender and LSP ID in the Senter-Template object do not form part of this check.
- Select those LSPs that have compatible downstream explicit routes. In this context, an explicit route is compatible if the path of the existing LSP satisfies the Explicit Route object (ERO) carried in the new Path message.
- Select those LSPs that were signaled with the same value of the Style object. See Section 3.3.2, 3.3.3, and 3.3.4 for further discussion of the processing of the Style object.
- Select those LSPs for which MP2P merging is allowed according to policy (possibly using the contents of the Policy object signaled in the received Path messages). This choice MAY include consideration of the hardware capabilities of the merging LSR.
- If more than one LSP is still a candidate for merging select the best LSP to merge with according to policy.

If no merge candidate is found the LSR processes the Path message as normal according to [RFC3209].

If an LSP is found, merging proceeds according to the reservation style carried in the Style object as follows.

3.3.2 Wildcard Filter Style

In [RFC2205] the Wildcard Filter (WF) reservation style allows sharing between all traffic in the same session, that is, to the same destination.

Just as in [RFC2205], the WF M2MP LSP may be thought of as a shared "pipe", whose "size" is the largest of the resource requests, independent of the number of senders using it. In other words, a WF style indicates that resource sharing occurs downstream of the merge point and this style is applicable to resource sharing applications such as voice-conferencing.

A merge-point LSR that receives a Path message for a second LSP that uses the WF style and that identifies a merge candidate LSP also using the WF style carries out the following processing.

- If the new LSP requests bandwidth less than or equal to the bandwidth already allocated downstream for the existing LSP, the merge-point immediately generates a Resv for the new LSP, allocates a label on the new upstream segment, and sends the Resv upstream following normal rules from [RFC3209].

The merge-point LSR installs an entry in the LFIB mapping the new label on the new upstream interface, to the existing label on the downstream interface for the MP2P LSP so that the new LSP is spliced into the MP2P LSP. That is, there are now multiple entries in the LFIB mapping the to the same label on the same downstream interface.

No further downstream signaling is performed.

- If the new LSP requests bandwidth greater than the bandwidth already allocated downstream for the existing LSP, the merge-point LSR sends a trigger Path message downstream for the existing LSP with a TSpec that requests increased bandwidth. If a Resv message is later received indicating that the required bandwidth has been allocated, the merge-point LSR builds and sends a Resv upstream for the new LSP segment only (no new Resv is sent upstream for the previously existing LSP segments). If the requested bandwidth is not made available, the merge-point LSR sends a PathErr message upstream on the new LSP segment just as it would if normal LSP setup had failed owing to lack of resources [RFC3209].

In this mode of operation, the egress does not become aware when new ingress LSRs are spliced into the MP2P tree.

3.3.3 Fixed Filter Style

In [RFC2205] the Fixed Filter (FF) style implies distinct reservations for each flow in the session.

For MP2P LSPs this means that common labels are used on the shared downstream segments, but a separate pool of resource is allocated for each ingress LSR (sender) that shares the MP2P tree.

The consequence is that each LSP (that is, each LSP from each ingress) must be signaled all the way to the egress carrying a distinct LSP identifier and TSpec. Although a possible approach is to send a separate Path message for each such LSP (which would be the technique used in [RFC2205]) this cannot be done because it would leave the choice of MP2P merging to the egress LSR. This situation is OK in [RFC2205] where all routers are capable of merging flows, but is not suitable in MP2P MPLS-TE because the merge node needs control of whether merging is performed.

Therefore, the solution is to merge the Path messages. The Path message for the first LSP is processed as normal (see Section 3.1), but once merging has been chosen, the incoming Path message for the second LSP is merged, and the trigger Path message sent downstream contains information about both senders and both TSpecs.

The mechanism used to achieve this mirrors the information that will be carried on the Resv. An <FF sender descriptor list> is used to provide a list of Sender Template objects each of which is accompanied by a TSpec. AdSpec and Record Route objects may also be present. The <FF sender descriptor list> MUST NOT be present unless the Style object is present and contains the FF setting. However, the observant reader will note that an <FF sender descriptor list> with only one element is identical to the <sender descriptor> defined in [RFC3209] and used for P2P signaling.

Each <sender descriptor> in the <FF sender descriptor> is copied from the <sender descriptor> on a received Path message. Thus all information from received Path messages is conveyed to the egress LSR.

3.3.4 Shared Explicit Style

[RFC2205] describes the Shared Explicit (SE) style as creating a single reservation that is shared by selected upstream senders. Unlike the WF style, the SE style requires that the senders that share the style are explicitly identified, and so each LSP that is joined to the MP2P LSP needs to be signaled to the egress.

The solution is also modeled on the SE signaling used in the Resv message in [RFC2205]. An <SE sender descriptor list> is used to provide a list of Sender Template objects that are associated with a single TSpec. AdSpec and Record Route objects may also be present. The <SE sender descriptor list> MUST NOT be present unless the Style object is present and contains the SE setting. However, the observant reader will note that an <SE sender descriptor list> with only one element is identical to the <sender descriptor> defined in [RFC3209] and used for P2P signaling.

Each Sender Template and any Record Route or AdSpec object in the <SE sender descriptor> is copied from the a received Path message. Thus all information from received Path messages is conveyed to the egress LSR. But the traffic parameters in the Sender TSpec object are summed to ensure that sufficient bandwidth is allocated on the downstream leg.

3.3.5 Record Route Processing on Merged Path Messages

The route of an LSP can be recorded in a Record Route object (RRO) included in a Path message [RFC3209]. There is no change to this procedure upstream of the merge point for two LSPs.

Downstream of the merge point, the Path message can include multiple RROs: one for each Sender Template object. The presence rules for these RROs is that they exist in the <sender descriptor> for a merged LSP in a Path message downstream of the merge point if and only if the RRO existed in the Path message for the LSP upstream of the merge point.

Note that continuing to record the route of an LSP in all RROs in a Path message would result in duplicate information that might make the Path message too large for convenient processing. To avoid such duplicate information the following three rules MUST be applied:

- When creating a <sender descriptor list> for transmission in a Path message to a downstream LSR the merge-point LSR MUST include an RRO in a <sender descriptor> if there is one in the corresponding <sender descriptor> in the Path message received from upstream.
- When creating a <sender descriptor list> for transmission in a Path message to a downstream LSR, if any <sender descriptor> includes an RRO the merge-point LSR MUST ensure that the first <sender descriptor> contains an RRO. This MUST be achieved by appropriate ordering of <sender descriptor> items and MUST NOT be achieved by inserting an RRO into a <sender descriptor> that would not otherwise include an RRO.

- When adding route record information to a Path message, an LSR MUST add the information only to the first RRO in the Path message.

Thus RROs in Path messages start at ingress LSRs and end at merge-point LSRs, except for the first RRO in the Path message that ends at the egress LSR.

3.4 Path Processing Downstream of a Merge Point

No special processing for Path messages is necessary at an LSR downstream of a merge-point LSR except for the following points.

- RROs MUST be handled as described in the previous section.
- Resources SHOULD be checked based on the contents of the Style object.

3.5 Processing at Egress LSRs

Egress LSRs are responsible for converting TSpec objects on Path messages into FlowSpec objects on Resv messages. The rules for doing this are inherited from [RFC2205] and [RFC3209] applying the information received in Path the messages including the value set in the Style object.

3.6 Resv Processing at Merge Points

Resv messages need careful handling at merge points. The contents of the Resv needs to be split and sent upstream in separate Resv messages according to the Path messages that were received.

The <flow descriptor list> on the Resv is simply split according to the <sender descriptors> received on the Path messages. Where bandwidth is shared on downstream links, the merge-point LSR makes appropriate decisions for how much bandwidth to allocate on the upstream links according to the TSpecs received in the Path messages. Where bandwidth is allocated per sender, this is simply copied from the received Resv to the outgoing Resvs.

Note that an RRO MUST NOT be included in a Resv sent upstream from a merge-point LSR unless an RRO was received in the corresponding Path message.

3.7 LSP Teardown

When an ingress LSR decides to tear down an LSP it sends a PathTear message. This message is propagated as per [RFC3209] until it reaches a merge-point LSR.

- If the PathTear is removing the last upstream segment then this is not really a merge-point LSR, and the PathTear MUST be forwarded toward the egress as normal.
- If the LSP uses the WF style, the PathTear MUST be terminated at the merge-point LSR. If the result of removing the associated sender is a reduction in the amount of required resource on the downstream LSP, the merge-point LSR SHOULD send a trigger Path message with a new TSpec to request a reduction in the allocated bandwidth at downstream LSRs, but MAY retain the downstream bandwidth according to local policy.
- If the LSP uses the FF or SE style, the PathTear MUST be terminated at the merge-point LSR and a trigger Path message MUST be sent downstream with the associated <sender descriptor> removed.

3.8 Make-Before-Break Support

Make-before-break is an important aspect of MPLS-TE that allows LSPs to be re-routed in a "hitless" way.

Make-before-break integrates with MP2P MPLS-TE in a very simple way, because of the nature of LSP merging in MP2P LSP establishment. Nevertheless, there are several cases to be considered in order to fully describe the operation of make-before-break. The variables are whether resource sharing or incremental resource allocation are in use for the MP2P LSP, and where the new LSP diverges from the old one in relation to the first merge point for the LSP.

A future version of this document will analyze each of the scenarios.

3.9 Fast Re-Route

Fast Re-Route (FRR) [RFC4090] provides rapid local protection of LSPs in the event of a link or node failure. FRR techniques have been defined for P2P and P2MP LSP.

FRR protection of MP2P LSPs will also be beneficial and suitable techniques need to be defined.

Consider the network fragment shown in Figure 1. Suppose there is an MP2P LSP {ABCGHIJ, DEFGHIJ} with merge point at node G. Failures remote from the merge point are handled as normal for FRR since they impact only a single component LSP. For example, the failure of the link DE can be protected by an FRR tunnel DWE.




Figure 1 : Network fragment to illustrate FRR

But, suppose there is a failure close to the merge point such that the FRR process impacts the MP2P tree? Consider the following separate cases:

- Link BC fails
The protection tunnel is BZG which terminates at the MP2P merge point. When the protected LSP emerges from the protection tunnel it is merged into the MP2P LSP as normal.
- Link BC and link EF fail
The dual failure case causes two protection tunnels BZG and EZG. These tunnels could utilize MP2P techniques and merge at node Z, but there is no significant benefit to this small reduction in LSP state. In any case, whether MP2P or P2P is used for the protection tunnels, at node G the protected LSPs will be merged back into the MP2P LSP.
- Link CG fails
When the link CG fails the protection tunnel CPH is used. This means that original MP2P merge point (G) no longer sees traffic source A, and the end of the protection tunnel (the FRR merge point node H) becomes a merge point for the MP2P LSP.
- Node G fails
The failure of the MP2P merge point (G) obviously impacts traffic from both sources A and D. Protection tunnels CPH and FQH can be used to protect the traffic, and when the end-to-end LSPs emerge from their protection tunnels node H becomes an MP2P merge point.
- Link GH fails
Similarly, if the link GH fails (in this topology) the protection tunnels CPH and FQH will be used even though the failure is downstream of the MP2P merge point. If, however, there was the possibility of building a protection tunnel through another node (for example, GUH - node U is not shown on the figure) the whole MP2P LSP could be protected just as it would be if the link HI failed (see below).

- Link HI fails

In this case the whole MP2P LSP can be protected by single FRR protection tunnel HSI. Note that load balancing on across multiple protection tunnels (HSI and HTI) cannot be achieved based on the source of the traffic since all traffic at node H uses the same labels and node H cannot distinguish the source of the traffic without deeper packet inspection. However, if H was an MP2P merge point, it would be possible to perform this type of load balancing.

4. Protocol Extensions

This section describes the protocol extensions and changes to message encodings required to support the function described in Section 3.

4.1 Allowing MP2P LSP Merging

An ingress LSR may allow M2MP LSP merging by setting the "MP2P Merge Allowed" flag in the LSP Attributes TLV carried in the LSP_ATTRIBUTES object [RFC5420] in the Path message that it sends for this LSP. The bit number for the "MP2P Merge Allowed" flag is defined in Section 9.

The LSP_ATTRIBUTES object is used rather than the LSP_REQUIRED_ATTRIBUTES object because transit LSRs that do not support MP2P LSP merging are not required to take any action.

An LSR that recognizes the LSP_ATTRIBUTES object, but does not recognize the "MP2P Merge Allowed" flag will ignore the flag according to the processing rules in [RFC5420], and MUST forward the flag unmodified in any Path message that it sends for this LSP. An LSR that recognizes the LSP_ATTRIBUTES object and the "MP2P Merge Allowed" flag, but which does not support MP2P LSP merging or does not wish to support MP2P merging for the signaled LSP MUST forward the flag unmodified in any Path message that it sends for this LSP.

4.2 Message Formats

As described in Section 3, some changes are required to the existing RSVP-TE message formats. These are shown in the sections that follow. The messages are described using the Backus-Naur Format notation from [RFC3209] as defined in [RFC5511].

The following messages are not changed by this document: PathTear, Resv, ResvTear, ResvErr, ResvConf.

4.2.1 Path Message

The Path message is extended by the optional inclusion of the Style object. If this object is present the <sender descriptor> is replaced

by a <sender descriptor list>.

```
<Path Message> ::=  <Common Header> [ <INTEGRITY> ]
                    <SESSION> <RSVP_HOP>
                    <TIME_VALUES>
                    [ <EXPLICIT_ROUTE> ]
                    <LABEL_REQUEST>
                    [ <SESSION_ATTRIBUTE> ]
                    [ <LSP_ATTRIBUTES> ]
                    [ <POLICY_DATA> ... ]
                    <sender descriptor>
                    | <STYLE> <sender descriptor list>
```

The new <sender descriptor list> is defined to mirror the Resv <flow descriptor list> as follows.

```
<sender descriptor list> ::=  <WF sender descriptor list>
                             | <FF sender descriptor list>
                             | <SE sender descriptor>

<WF sender descriptor list> ::= <sender descriptor>

<FF sender descriptor list> ::= <sender descriptor>
                               | <FF sender descriptor list> <sender descriptor>

<SE sender descriptor list> ::= <sender descriptor>
                               <SE sender list>

<SE sender list> ::=  <SENDER_TEMPLATE>
                    [ <ADSPEC> ]
                    [ <RECORD_ROUTE> ]
```

The definition of <sender descriptor> is found in [RFC3209].

4.2.2 PathErr Message

The PathErr message is modified as follows.

```
<PathErr message> ::=  <Common Header> [ <INTEGRITY> ]
                    <SESSION> <ERROR_SPEC>
                    [ <POLICY_DATA> ...]
                    [ <sender descriptor>
                    | <STYLE> <sender descriptor list>]
```

<sender descriptor> is as defined in [RFC3209] and <sender descriptor list> is as defined in the previous section.

5. Backward Compatibility

This section examines how MP2P support may co-exist with previous MPLS-TE function.

5.1 Legacy LSRs

The presence of the Style object on a Path message will cause a legacy LSR to reject the message with a PathErr message containing an Error Code that indicates that an unsupported object was found. Therefore, in order to support MP2P LSPs it is necessary that all LSRs that may lie on the route of an MP2P LSP must be upgraded to support the procedures and protocol extensions described in this document.

Where this is not practical, it is possible to consider establish FA LSPs [RFC4206] to tunnel past the legacy LSPs. This would probably require management coordination.

5.2 Support of P2P and P2MP LSPs through an MP2P-enabled network

It is clearly a requirement that an MPLS-TE network that has been enabled for MP2P MUST continue to support P2P and P2MP TE LSPs. There is nothing in the procedures and extensions defined in this document that prevents this.

6. Path Computation Considerations

An issue arises for path computation for MP2P LSPs. If this computation is performed on ingress LSRs using the Traffic Engineering Database (TED) built from information distributed by the Interior Gateway Protocols (IGPs) then the computation engine will not be aware of bandwidth sharing possibilities and will possibly exclude optimal paths assuming that bandwidth is not available. This problem only occurs where bandwidth sharing downstream of the merge points is desirable.

In the case of this particular application it may be desirable to consider coordinated or centralized path computation such as might be performed by stateful PCE [RFC4655].

7. Manageability Considerations

7.1 MIB Modules

The LSR MIB module [RFC3813] contains full support for many-to-one associations through the cross-connect table. No extensions are required.

The MPLS TE MIB module [RFC3812] needs an applicability statement to show how a Path message that contains multiple senders may be mapped to entries in the MIB tables. However, no extensions to those MIB tables are required.

7.2 Operations and Management

Operations and Management (OAM) tools are essential to the successful operation of MPLS networks. [RFC4378] provides a framework for OAM in MPLS networks, and this is applicable for MP2P MPLS-TE networks. [RFC4377] sets out requirements for OAM in MPLS networks, and those are applicable to MP2P MPLS-TE.

It may be observed that the structure of an LSP in LDP is similar to that in MP2P MPLS-TE and so many of the same issues and solutions will apply.

8. Security Considerations

MP2P LSPs open up new security concerns not raised in [RFC3209]. In particular, the WF style is vulnerable to an ingress being joined to an MP2P LSP without the knowledge of the egress LSR and being able to send data to the egress. The peer-level (neighbor) authentication supported by RSVP [RFC2205] becomes more important in this case and it is RECOMMENDED that the WF style only be used in a network where complete security is assumed, or where RSVP security is enabled between all neighbors.

The policy features available at merge-point LSRs can also be used to enhance the security of MP2P LSPs by applying careful checks, including security checks before allowing an upstream LSP segment to be merged into an MP2P LSP. The hop-by-hop security of RSVP may be assumed to provide guarantees of the identity of each LSP sender, but greater security can be achieved by providing sender authentication within the Policy object, and potentially by establishing a security relationship between each ingress LSR and each merge-point LSR. The latter, however, is a considerable burden in a large and well-meshed network; it might be more appropriate to utilise a common security authentication server such as a policy server.

Otherwise, these extensions raise no security issues that are not covered [RFC3209].

9. IANA Considerations

IANA is requested to assign a bit number for use by the "MP2P Merge Allowed" flag in the Attribute Flags TLV carried in the LSP_ATTRIBUTES object.

Bit number 8 is suggested.

The interpretation of this flag should be tracked by IANA as follows.

- It has meaning in the Attribute Flags TLV on a Path
- It does not have meaning in the Attribute Flags TLV on a Resv
- It does not have meaning in the RRO Attributes Subobject.

See Section 4.1 for a description of the use of this flag.

10. Acknowledgments

The authors thank Dimitri Papadimitriou and Francois Le Faucheur for useful discussions.

11. References

11.1 Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to indicate requirements levels", RFC 2119, March 1997.
- [RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1, Functional Specification", RFC 2205, September 1997.
- [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.
- [RFC5420] Farrel, A., Papadimitriou, D., Vasseur, J.P., and Ayyangar, A., "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 5420, February 2009.
- [RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, April 2009.

11.2 Informative References

- [RFC3031] Rosen, E., Viswanathan, A., and Callon, R., "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3812] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Traffic Engineering (TE) Management Information Base (MIB)", RFC 3812, June 2004.
- [RFC3813] Srinivasan, C., Viswanathan, A., and T. Nadeau, "Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base (MIB)", RFC 3813, June 2004.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.
- [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and Matsushima, S. "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC4377, February 2006.
- [RFC4378] Allan, D., and Nadeau, T., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC4378, February 2006.
- [RFC4655] Farrel, A., Vasseur, J.P., and Ash, G., "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4875] Aggarwal, R., Papadimitriou, D., and Yasukawa, S., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC5036] Andersson, L., Doolan, P., Feldman, N., Fredette, A., and Thomas, B., "LDP Specification", RFC 5036, October 2007.

12. Authors' Addresses

Seisho Yasukawa
NTT Corporation
9-11, Midori-Cho 3-Chome
Musashino-Shi, Tokyo 180-8585,
Japan
Email: yasukawa.seisho@lab.ntt.co.jp

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

Full Copyright Statement

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

