

Network Working Group
Request for Comments: 5150
Category: Standards Track

A. Ayyangar
K. Kompella
Juniper Networks
JP. Vasseur
Cisco Systems, Inc.
A. Farrel
Old Dog Consulting
February 2008

Label Switched Path Stitching with
Generalized Multiprotocol Label Switching Traffic Engineering (GMPLS TE)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

In certain scenarios, there may be a need to combine several Generalized Multiprotocol Label Switching (GMPLS) Label Switched Paths (LSPs) such that a single end-to-end (e2e) LSP is realized and all traffic from one constituent LSP is switched onto the next LSP. We will refer to this as "LSP stitching", the key requirement being that a constituent LSP not be allocated to more than one e2e LSP. The constituent LSPs will be referred to as "LSP segments" (S-LSPs).

This document describes extensions to the existing GMPLS signaling protocol (Resource Reservation Protocol-Traffic Engineering (RSVP-TE)) to establish e2e LSPs created from S-LSPs, and describes how the LSPs can be managed using the GMPLS signaling and routing protocols.

It may be possible to configure a GMPLS node to switch the traffic from an LSP for which it is the egress, to another LSP for which it is the ingress, without requiring any signaling or routing extensions whatsoever and such that the operation is completely transparent to other nodes. This will also result in LSP stitching in the data plane. However, this document does not cover this scenario of LSP stitching.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	3
2. Comparison with LSP Hierarchy	3
3. Usage	4
3.1. Triggers for LSP Segment Setup	4
3.2. Applications	5
4. Routing Aspects	5
5. Signaling Aspects	6
5.1. RSVP-TE Signaling Extensions	7
5.1.1. Creating and Preparing an LSP Segment for Stitching	7
5.1.1.1. Steps to Support Penultimate Hop Popping	8
5.1.2. Stitching the e2e LSP to the LSP Segment	9
5.1.3. RRO Processing for e2e LSPs	10
5.1.4. Teardown of LSP Segments	11
5.1.5. Teardown of e2e LSPs	11
5.2. Summary of LSP Stitching Procedures	12
5.2.1. Example Topology	12
5.2.2. LSP Segment Setup	12
5.2.3. Setup of an e2e LSP	13
5.2.4. Stitching of an e2e LSP into an LSP Segment	13
6. Security Considerations	14
7. IANA Considerations	15
7.1. Attribute Flags for LSP_ATTRIBUTES Object	15
7.2. New Error Codes	15
8. Acknowledgments	16
9. References	16
9.1. Normative References	16
9.2. Informative References	17

1. Introduction

A stitched Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering (TE) Label Switched Path (LSP) is built from a set of different "LSP segments" (S-LSPs) that are connected together in the data plane in such a way that a single end-to-end LSP is realized in the data plane. In this document, we define the concept of LSP stitching and detail the control plane mechanisms and procedures (routing and signaling) to accomplish this. Where applicable, similarities and differences between LSP hierarchy [RFC4206] and LSP stitching are highlighted. Signaling extensions required for LSP stitching are also described here.

It may be possible to configure a GMPLS node to switch the traffic from an LSP for which it is the egress, to another LSP for which it is the ingress, without requiring any signaling or routing extensions whatsoever and such that the operation is completely transparent to other nodes. This results in LSP stitching in the data plane, but requires management intervention at the node where the stitching is performed. With the mechanism described in this document, the node performing the stitching does not require configuration of the pair of S-LSPs to be stitched together. Also, LSP stitching as defined here results in an end-to-end LSP both in the control and data planes.

1.1. 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].

2. Comparison with LSP Hierarchy

LSP hierarchy ([RFC4206]) provides signaling and routing procedures so that:

- a. A Hierarchical LSP (H-LSP) can be created. Such an LSP created in one layer can appear as a data link to LSPs in higher layers. As such, one or more LSPs in a higher layer can traverse this H-LSP as a single hop; we call this "nesting".
- b. An H-LSP may be managed and advertised (although this is not a requirement) as a Traffic Engineering (TE) link. Advertising an H-LSP as a TE link allows other nodes in the TE domain in which it is advertised to use this H-LSP in path computation. If the H-LSP TE link is advertised in the same instance of control plane (TE domain) in which the H-LSP was provisioned, it is then defined as a forwarding adjacency LSP (FA-LSP) and GMPLS nodes can form a forwarding adjacency (FA) over this FA-LSP. There is usually no routing adjacency between end points of an FA. An H-LSP may also be advertised as a TE link in a different TE domain. In this case, the end points of the H-LSP are required to have a routing adjacency between them.
- c. RSVP signaling ([RFC3473], [RFC3209]) for LSP setup can occur between nodes that do not have a routing adjacency.

In case of LSP stitching, instead of an H-LSP, an LSP segment (S-LSP) is created between two GMPLS nodes. An S-LSP for stitching is considered to be the moral equivalent of an H-LSP for nesting. An S-LSP created in one layer, unlike an H-LSP, provides a data link to

other LSPs in the same layer. Similar to an H-LSP, an S-LSP could be managed and advertised, although it is not required, as a TE link, either in the same TE domain as it was provisioned or a different one. If so advertised, other GMPLS nodes can use the corresponding S-LSP TE link in path computation. While there is a forwarding adjacency between end points of an H-LSP TE link, there is no forwarding adjacency between end points of an S-LSP TE link. In this aspect, an H-LSP TE link more closely resembles a 'basic' TE link as compared to an S-LSP TE link.

While LSP hierarchy allows more than one LSP to be mapped to an H-LSP, in case of LSP stitching, at most one LSP may be associated with an S-LSP. Thus, if LSP-AB is an H-LSP between nodes A and B, then multiple LSPs, say LSP1, LSP2, and LSP3, can potentially be 'nested into' LSP-AB. This is achieved by exchanging a unique label for each of LSP1..3 over the LSP-AB hop, thereby separating the data corresponding to each of LSP1..3 while traversing the H-LSP LSP-AB. Each of LSP1..3 may reserve some bandwidth on LSP-AB. On the other hand, if LSP-AB is an S-LSP, then at most one LSP, say LSP1, may be stitched to the S-LSP LSP-AB. LSP-AB is then dedicated to LSP1, and no other LSPs can be associated with LSP-AB. The entire bandwidth on S-LSP LSP-AB is allocated to LSP1. However, similar to H-LSPs, several S-LSPs may be bundled into a TE link ([RFC4201]).

The LSPs LSP1..3 that are either nested or stitched into another LSP are termed as e2e LSPs in the rest of this document. Routing procedures specific to LSP stitching are detailed in Section 4.

Targeted (non-adjacent) RSVP signaling defined in [RFC4206] is required for LSP stitching of an e2e LSP to an S-LSP. Specific extensions for LSP stitching are described in Section 5.1. Therefore, in the control plane, there is one RSVP session corresponding to the e2e LSP as well as one for each S-LSP. The creation and termination of an S-LSP may be dictated by administrative control (statically provisioned) or due to another incoming LSP request (dynamic). Triggers for dynamic creation of an S-LSP may be different from that of an H-LSP and will be described in detail in Section 3.1.

3. Usage

3.1. Triggers for LSP Segment Setup

An S-LSP may be created either by administrative control (configuration trigger) or dynamically due to an incoming LSP request. LSP hierarchy ([RFC4206]) defines one possible trigger for dynamic creation of an FA-LSP by introducing the notion of LSP regions based on Interface Switching Capabilities. As per [RFC4206],

dynamic FA-LSP creation may be triggered on a node when an incoming LSP request crosses region boundaries. However, this trigger MUST NOT be used for creation of an S-LSP for LSP stitching as described in this document. In case of LSP stitching, the switching capabilities of the previous hop and the next hop TE links MUST be the same. Therefore, local policies configured on the node SHOULD be used for dynamic creation of LSP segments.

Other possible triggers for dynamic creation of both H-LSPs and S-LSPs include cases where an e2e LSP may cross domain boundaries or satisfy locally configured policies on the node as described in [RFC5151].

3.2. Applications

LSP stitching procedures described in this document are applicable to GMPLS nodes that need to associate an e2e LSP with another S-LSP of the same switching type and LSP hierarchy procedures do not apply. For example, if an e2e lambda LSP traverses an LSP segment TE link that is also lambda-switch capable, then LSP hierarchy is not possible; in this case, LSP switching may be an option.

LSP stitching procedures can be used for inter-domain TE LSP signaling to stitch an inter-domain e2e LSP to a local intra-domain TE S-LSP ([RFC4726] and [RFC5151]).

LSP stitching may also be useful in networks to bypass legacy nodes that may not have certain new capabilities in the control plane and/or data plane. For example, one suggested usage in the case of point-to-multipoint (P2MP) RSVP LSPs ([RFC4875]) is the use of LSP stitching to stitch a P2MP RSVP LSP to an LSP segment between P2MP-capable Label Switching Routers (LSRs) in the network. The LSP segment would traverse legacy LSRs that may be incapable of acting as P2MP branch points, thereby shielding them from the P2MP control and data path. Note, however, that such configuration may limit the attractiveness of RSVP P2MP and should carefully be examined before deployment.

4. Routing Aspects

An S-LSP is created between two GMPLS nodes, and it may traverse zero or more intermediate GMPLS nodes. There is no forwarding adjacency between the end points of an S-LSP TE link. So although in the TE topology, the end points of an S-LSP TE link are adjacent, in the data plane, these nodes do not have an adjacency. Hence, any data plane resource identifier between these nodes is also meaningless.

The traffic that arrives at the head end of the S-LSP is switched into the S-LSP contiguously with a label swap, and no label is associated directly between the end nodes of the S-LSP itself.

An S-LSP MAY be treated and managed as a TE link. This TE link MAY be numbered or unnumbered. For an unnumbered S-LSP TE link, the schemes for assignment and handling of the local and remote link identifiers as specified in [RFC3477] SHOULD be used. When appropriate, the TE information associated with an S-LSP TE link MAY be flooded via ISIS-TE [RFC4205] or OSPF-TE [RFC4203]. Mechanisms similar to that for regular (basic) TE links SHOULD be used to flood S-LSP TE links. Advertising or flooding the S-LSP TE link is not a requirement for LSP stitching. If advertised, this TE information will exist in the TE database (TED) and can then be used for path computation by other GMPLS nodes in the TE domain in which it is advertised. When so advertising S-LSPs, one should keep in mind that these add to the size and complexity of the link-state database.

If an S-LSP is advertised as a TE link in the same TE domain in which it was provisioned, there is no need for a routing adjacency between end points of this S-LSP TE link. If an S-LSP TE link is advertised in a different TE domain, the end points of that TE link SHOULD have a routing adjacency between them.

The TE parameters defined for an FA in [RFC4206] SHOULD be used for an S-LSP TE link as well. The switching capability of an S-LSP TE link MUST be equal to the switching type of the underlying S-LSP; i.e., an S-LSP TE link provides a data link to other LSPs in the same layer, so no hierarchy is possible.

An S-LSP MUST NOT admit more than one e2e LSP into it. If an S-LSP is allocated to an e2e LSP, the unreserved bandwidth SHOULD be set to zero to prevent further e2e LSPs from being admitted into the S-LSP.

Multiple S-LSPs between the same pair of nodes MAY be bundled using the concept of Link Bundling ([RFC4201]) into a single TE link. In this case, each component S-LSP may be allocated to at most one e2e LSP. When any component S-LSP is allocated for an e2e LSP, the component's unreserved bandwidth SHOULD be set to zero and the Minimum and Maximum LSP bandwidth of the TE link SHOULD be recalculated. This will prevent more than one LSP from being computed and admitted over an S-LSP.

5. Signaling Aspects

The end nodes of an S-LSP may or may not have a routing adjacency. However, they SHOULD have a signaling adjacency (RSVP neighbor relationship) and will exchange RSVP messages with each other. It

may, in fact, be desirable to exchange RSVP Hellos directly between the LSP segment end points to allow support for state recovery during Graceful Restart procedures as described in [RFC3473].

In order to signal an e2e LSP over an LSP segment, signaling procedures described in Section 8.1.1 of [RFC4206] MUST be used. Additional signaling extensions for stitching are described in the next section.

5.1. RSVP-TE Signaling Extensions

The signaling extensions described here MUST be used for stitching an e2e packet or non-packet GMPLS LSP ([RFC3473]) to an S-LSP.

Stitching an e2e LSP to an LSP segment involves the following two-step process:

1. Creating and preparing the S-LSP for stitching by signaling the desire to stitch between end points of the S-LSP; and
2. Stitching the e2e LSP to the S-LSP.

5.1.1. Creating and Preparing an LSP Segment for Stitching

If a GMPLS node desires to create an S-LSP, i.e., one to be used for stitching, then it MUST indicate this in the Path message for the S-LSP. This signaling explicitly informs the S-LSP egress node that the ingress node is planning to perform stitching over the S-LSP. Since an S-LSP is not conceptually different from any other LSP, explicitly signaling 'LSP stitching desired' helps clarify the data plane actions to be carried out when the S-LSP is used by some other e2e LSP. Also, in the case of packet LSPs, this is what allows the egress of the S-LSP to carry out label allocation as explained below. Also, so that the head-end node can ensure that correct stitching actions will be carried out at the egress node, the egress node MUST signal this information back to the head-end node in the Resv, as explained below.

In order to request LSP stitching on the S-LSP, we define a new bit in the Attributes Flags TLV of the LSP_ATTRIBUTES object defined in [RFC4420]:

LSP stitching desired bit - This bit SHOULD be set in the Attributes Flags TLV of the LSP_ATTRIBUTES object in the Path message for the S-LSP by the head end of the S-LSP that desires LSP stitching. This bit MUST NOT be modified by any other nodes in the network. Nodes other than the egress of the S-LSP SHOULD ignore this bit. The bit number for this flag is defined in Section 7.1.

An LSP segment can be used for stitching only if the egress node of the S-LSP is also ready to participate in stitching. In order to indicate this to the head-end node of the S-LSP, the following new bit is defined in the Flags field of the Record Route object (RRO) Attributes subobject: "LSP segment stitching ready". The bit number for this flag is defined in Section 7.1.

If an egress node of the S-LSP receiving the Path message supports the LSP_ATTRIBUTES object and the Attributes Flags TLV, and also recognizes the "LSP stitching desired" bit, but cannot support the requested stitching behavior, then it MUST send back a PathErr message with an error code of "Routing Problem" and an error value of "Stitching unsupported" to the head-end node of the S-LSP. The new error value is defined in Section 7.2.

If an egress node receiving a Path message with the "LSP stitching desired" bit set in the Flags field of received LSP_ATTRIBUTES object recognizes the object, the TLV TLV, and the bit and also supports the desired stitching behavior, then it MUST allocate a non-NULL label for that S-LSP in the corresponding Resv message. Also, so that the head-end node can ensure that the correct label (forwarding) actions will be carried out by the egress node and that the S-LSP can be used for stitching, the egress node MUST set the "LSP segment stitching ready" bit defined in the Flags field of the RRO Attribute subobject.

Finally, if the egress node for the S-LSP supports the LSP_ATTRIBUTES object but does not recognize the Attributes Flags TLV, or supports the TLV as well but does not recognize this particular bit, then it SHOULD simply ignore the above request.

An ingress node requesting LSP stitching MUST examine the RRO Attributes subobject Flags corresponding to the egress node for the S-LSP, to make sure that stitching actions are carried out at the egress node. It MUST NOT use the S-LSP for stitching if the "LSP segment stitching ready" bit is cleared.

5.1.1.1. Steps to Support Penultimate Hop Popping

Note that this section is only applicable to packet LSPs that use Penultimate Hop Popping (PHP) at the last hop, where the egress node distributes the Implicit NULL Label ([RFC3032]) in the Resv Label. These steps MUST NOT be used for a non-packet LSP and for packet LSPs where PHP is not desired.

When the egress node of a packet S-LSP receives a Path message for an e2e LSP that uses the S-LSP, the egress of the S-LSP SHOULD first check to see if it is also the egress of the e2e LSP. If the egress node is the egress for both the S-LSP and the e2e TE LSP, and this is

a packet LSP that requires PHP, then the node MUST send back a Resv trigger message for the S-LSP with a new label corresponding to the Implicit or Explicit NULL Label. Note that this operation does not cause any traffic disruption because the S-LSP is not carrying any traffic at this time, since the e2e LSP has not yet been established.

If the e2e LSP and the S-LSP are bidirectional, the ingress of the e2e LSP SHOULD first check whether it is also the ingress of the S-LSP. If so, it SHOULD re-issue the Path message for the S-LSP with an Implicit or Explicit NULL Upstream Label, and only then proceed with the signaling of the e2e LSP.

5.1.2. Stitching the e2e LSP to the LSP Segment

When a GMPLS node receives an e2e LSP request, depending on the applicable trigger, it may either dynamically create an S-LSP based on procedures described above or map an e2e LSP to an existing S-LSP. The switching type in the Generalized Label Request of the e2e LSP MUST be equal to the switching type of the S-LSP. Other constraints like the explicit path encoded in the Explicit Route object (ERO), bandwidth, and local TE policies MUST also be used for S-LSP selection or signaling. In either case, once an S-LSP has been selected for an e2e LSP, the following procedures MUST be followed in order to stitch an e2e LSP to an S-LSP.

The GMPLS node receiving the e2e LSP setup Path message MUST use the signaling procedures described in [RFC4206] to send the Path message to the end point of the S-LSP. In this Path message, the node MUST identify the S-LSP in the RSVP_HOP. An egress node receiving this RSVP_HOP should also be able to identify the S-LSP TE link based on the information signaled in the RSVP_HOP. If the S-LSP TE link is numbered, then the addressing scheme as proposed in [RFC4206] SHOULD be used to number the S-LSP TE link. If the S-LSP TE link is unnumbered, then any of the schemes proposed in [RFC3477] SHOULD be used to exchange S-LSP TE link identifiers between the S-LSP end points. If the TE link is bundled, the RSVP_HOP SHOULD identify the component link as defined in [RFC4201].

In case of a bidirectional e2e TE LSP, an Upstream Label MUST be signaled in the Path message for the e2e LSP over the S-LSP hop. However, since there is no forwarding adjacency between the S-LSP end points, any label exchanged between them has no significance. So the node MAY choose any label value for the Upstream Label. The label value chosen and signaled by the node in the Upstream Label is out of the scope of this document and is specific to the implementation on that node. The egress node receiving this Path message MUST ignore the Upstream Label in the Path message over the S-LSP hop.

The egress node receiving this Path message MUST signal a Label in the Resv message for the e2e TE LSP over the S-LSP hop. Again, since there is no forwarding adjacency between the egress and ingress S-LSP nodes, any label exchanged between them is meaningless. So the egress node MAY choose any label value for the Label. The label value chosen and signaled by the egress node is out of the scope of this document and is specific to the implementation on the egress node. The egress S-LSP node SHOULD also carry out data plane operations so that traffic coming in on the S-LSP is switched over to the e2e LSP downstream, if the egress of the e2e LSP is some other node downstream. If the e2e LSP is bidirectional, this means setting up label switching in both directions. The Resv message from the egress S-LSP node is IP routed back to the previous hop (ingress of the S-LSP). The ingress node stitching an e2e TE LSP to an S-LSP MUST ignore the Label object received in the Resv for the e2e TE LSP over the S-LSP hop. The S-LSP ingress node SHOULD also carry out data plane operations so that traffic coming in on the e2e LSP is switched into the S-LSP. It should also carry out actions to handle traffic in the opposite direction if the e2e LSP is bidirectional.

Note that the label exchange procedure for LSP stitching on the S-LSP hop is similar to that for LSP hierarchy over the H-LSP hop. The difference is the lack of the significance of this label between the S-LSP end points in case of stitching. Therefore, in case of stitching, the recipients of the Label/Upstream Label MUST NOT process these labels. Also, at most one e2e LSP is associated with one S-LSP. If a node at the head end of an S-LSP receives a Path message for an e2e LSP that identifies the S-LSP in the ERO and the S-LSP bandwidth has already been allocated to some other LSP, then regular rules of RSVP-TE pre-emption apply to resolve contention for S-LSP bandwidth. If the LSP request over the S-LSP cannot be satisfied, then the node SHOULD send back a PathErr with the error codes as described in [RFC3209].

5.1.3. RRO Processing for e2e LSPs

RRO procedures for the S-LSP specific to LSP stitching are already described in Section 5.1.1. In this section, we will look at the RRO processing for the e2e LSP over the S-LSP hop.

An e2e LSP traversing an S-LSP SHOULD record in the RRO for that hop, an identifier corresponding to the S-LSP TE link. This is applicable to both Path and Resv messages over the S-LSP hop. If the S-LSP is numbered, then the IPv4 or IPv6 address subobject ([RFC3209]) SHOULD be used to record the S-LSP TE link address. If the S-LSP is unnumbered, then the Unnumbered Interface ID subobject as described in [RFC3477] SHOULD be used to record the node's Router ID and Interface ID of the S-LSP TE link. In either case, the RRO subobject

SHOULD identify the S-LSP TE link end point. Intermediate links or nodes traversed by the S-LSP itself SHOULD NOT be recorded in the RRO for the e2e LSP over the S-LSP hop.

5.1.4. Teardown of LSP Segments

S-LSP teardown follows the standard procedures defined in [RFC3209] and [RFC3473]. This includes procedures without and with setting the administrative status. Teardown of S-LSP may be initiated by the ingress, egress, or any other node along the S-LSP path. Deletion/teardown of the S-LSP SHOULD be treated as a failure event for the e2e LSP associated with it, and corresponding teardown or recovery procedures SHOULD be triggered for the e2e LSP. In case of S-LSP teardown for maintenance purpose, the S-LSP ingress node MAY treat this to be equivalent to administratively shutting down a TE link along the e2e LSP path and take corresponding actions to notify the ingress of this event. The actual signaling procedures to handle this event is out of the scope of this document.

5.1.5. Teardown of e2e LSPs

e2e LSP teardown also follows standard procedures defined in [RFC3209] and [RFC3473] either without or with the administrative status. Note, however, that teardown procedures of e2e LSP and of S-LSP are independent of each other. So it is possible that while one LSP follows graceful teardown with administrative status, the other LSP is torn down without administrative status (using PathTear/ResvTear/PathErr with state removal).

When an e2e LSP teardown is initiated from the head end, and a PathTear arrives at the GMPLS stitching node, the PathTear message like the Path message MUST be IP routed to the LSP segment egress node with the destination IP address of the Path message set to the address of the S-LSP end node. Router Alert MUST be off and RSVP Time to Live (TTL) check MUST be disabled on the receiving node. PathTear will result in deletion of RSVP states corresponding to the e2e LSP and freeing of label allocations and bandwidth reservations on the S-LSP. The unreserved bandwidth on the S-LSP TE link SHOULD be readjusted.

Similarly, a teardown of the e2e LSP may be initiated from the tail end either using a ResvTear or a PathErr with state removal. The egress of the S-LSP MUST propagate the ResvTear/PathErr upstream, and MUST use IP addressing to target the ingress of the LSP segment.

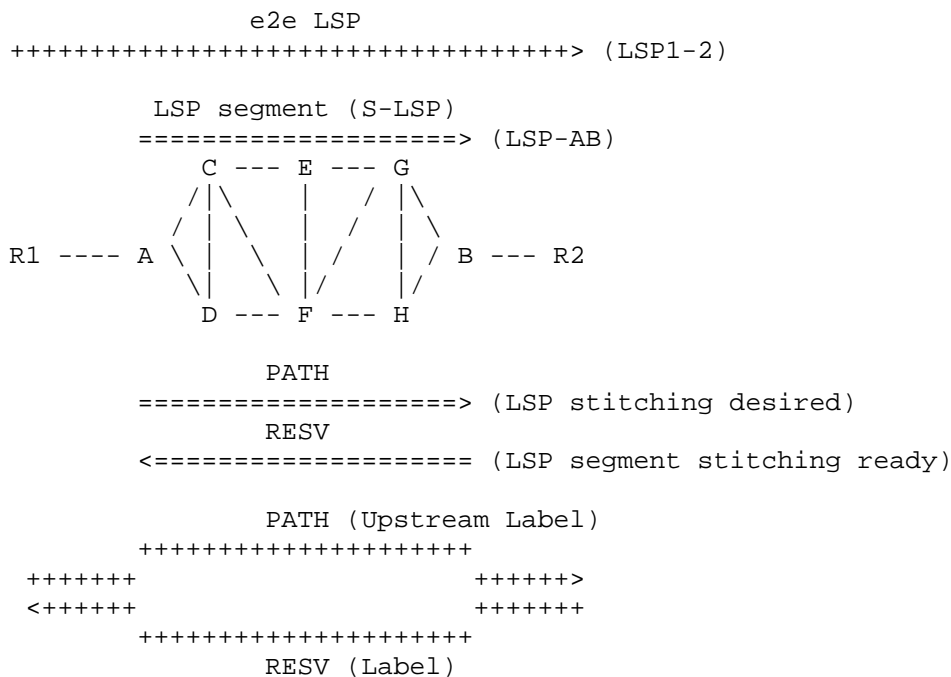
Graceful LSP teardown using ADMIN_STATUS as described in [RFC3473] is also applicable to stitched LSPs.

If the S-LSP was statically provisioned, tearing down of an e2e LSP MAY not result in tearing down of the S-LSP. If, however, the S-LSP was dynamically set up due to the e2e LSP setup request, then, depending on local policy, the S-LSP MAY be torn down if no e2e LSP is utilizing the S-LSP. Although the S-LSP may be torn down while the e2e LSP is being torn down, it is RECOMMENDED that a delay be introduced in tearing down the S-LSP once the e2e LSP teardown is complete, in order to reduce the simultaneous generation of RSVP errors and teardown messages due to multiple events. The delay interval may be set based on local implementation. The RECOMMENDED interval is 30 seconds.

5.2. Summary of LSP Stitching Procedures

5.2.1. Example Topology

The following topology will be used for the purpose of examples quoted in the following sections.



5.2.2. LSP Segment Setup

Let us consider an S-LSP LSP-AB being set up between two nodes A and B that are more than one hop away. Node A sends a Path message for the LSP-AB with "LSP stitching desired" set in the Flags field of the

LSP_ATTRIBUTES object. If the egress node B is ready to carry out stitching procedures, then B will respond with "LSP segment stitching ready" set in the Flags field of the RRO Attributes subobject, in the RRO sent in the Resv for the S-LSP. Once A receives the Resv for LSP-AB and sees this bit set in the RRO, it can then use LSP-AB for stitching. Node A cannot use LSP-AB for stitching if the bit is cleared in the RRO.

5.2.3. Setup of an e2e LSP

Let us consider an e2e LSP LSP1-2 starting one hop before A on R1 and ending on node R2, as shown above. If the S-LSP has been advertised as a TE link in the TE domain, and R1 and A are in the same domain, then R1 may compute a path for LSP1-2 over the S-LSP LSP-AB and identify the LSP-AB hop in the ERO. If not, R1 may compute hops between A and B and A may use these ERO hops for S-LSP selection or signaling a new S-LSP. If R1 and A are in different domains, then LSP1-2 is an inter-domain LSP. In this case, S-LSP LSP-AB, similar to any other basic TE link in the domain, will not be advertised outside the domain. R1 would use either per-domain path computation ([RFC5152]) or PCE-based computation ([RFC4655]) for LSP1-2.

5.2.4. Stitching of an e2e LSP into an LSP Segment

When the Path message for the e2e LSP LSP1-2 arrives at node A, A matches the switching type of LSP1-2 with the S-LSP LSP-AB. If the switching types are not equal, then LSP-AB cannot be used to stitch LSP1-2. Once the S-LSP LSP-AB to which LSP1-2 will be stitched has been determined, the Path message for LSP1-2 is sent (via IP routing, if needed) to node B with the IF_ID RSVP_HOP identifying the S-LSP LSP-AB. When B receives this Path message for LSP1-2, if B is also the egress for LSP1-2, and if this is a packet LSP requiring PHP, then B will send a Resv refresh for LSP-AB with the NULL Label. In this case, since B is not the egress, the Path message for LSP1-2 is propagated to R2. The Resv for LSP1-2 from B is sent back to A with a Label value chosen by B. B also sets up its data plane to swap the Label sent to either G or H on the S-LSP with the Label received from R2. Node A ignores the Label on receipt of the Resv message and then propagates the Resv to R1. A also sets up its data plane to swap the Label sent to R1 with the Label received on the S-LSP from C or D. This stitches the e2e LSP LSP1-2 to an S-LSP LSP-AB between nodes A and B. In the data plane, this yields a series of label swaps from R1 to R2 along e2e LSP LSP1-2.

6. Security Considerations

From a security point of view, the changes introduced in this document model the changes introduced by [RFC4206]. That is, the control interface over which RSVP messages are sent or received need not be the same as the data interface that the message identifies for switching traffic. But the capability for this function was introduced in [RFC3473] to support the concept of out-of-fiber control channels, so there is nothing new in this concept for signaling or security.

The application of this facility means that the "sending interface" or "receiving interface" may change as routing changes. So these interfaces cannot be used to establish security associations between neighbors, and security associations MUST be bound to the communicating neighbors themselves.

[RFC2747] provides a solution to this issue: in Section 2.1, under "Key Identifier", an IP address is a valid identifier for the sending (and by analogy, receiving) interface. Since RSVP messages for a given LSP are sent to an IP address that identifies the next/previous hop for the LSP, one can replace all occurrences of 'sending [receiving] interface' with 'receiver's [sender's] IP address' (respectively). For example, in Section 4, third paragraph, instead of:

"Each sender SHOULD have distinct security associations (and keys) per secured sending interface (or LIH). ... At the sender, security association selection is based on the interface through which the message is sent."

it should read:

"Each sender SHOULD have distinct security associations (and keys) per secured receiver's IP address. ... At the sender, security association selection is based on the IP address to which the message is sent."

Thus, the mechanisms of [RFC2747] can be used unchanged to establish security associations between control plane neighbors.

This document allows the IP destination address of Path and PathTear messages to be the IP address of a next hop node (receiver's address) instead of the RSVP session destination address. This means that the use of the IPsec Authentication Header (AH) (ruled out in [RFC2747]

because RSVP messages were encapsulated in IP packets addressed to the ultimate destination of the Path or PathTear messages) is now perfectly applicable, and standard IPsec procedures can be used to secure the message exchanges.

An analysis of GMPLS security issues can be found in [MPLS-SEC].

7. IANA Considerations

IANA has made the following codepoint allocations for this document.

7.1. Attribute Flags for LSP_ATTRIBUTES Object

The "RSVP TE Parameters" registry includes the "Attributes Flags" sub-registry.

IANA has allocated the following new bit (5) defined for the Attributes Flags TLV in the LSP_ATTRIBUTES object.

LSP stitching bit - Bit Number 5

This bit is only to be used in the Attributes Flags TLV on a Path message.

The 'LSP stitching desired' bit has a corresponding 'LSP segment stitching ready' bit (Bit Number 5) to be used in the RRO Attributes subobject.

The following text has been included in the registry:

Bit No	Name	Attribute Flags	Path Path	Path Resv	RRO	Reference
5	LSP stitching desired	Yes		No	Yes	[RFC5150]

7.2. New Error Codes

The "Resource ReSerVation Protocol (RSVP) Parameters" registry includes the "Error Codes and Globally-Defined Error Value Sub-Codes" sub-registry.

IANA has assigned a new error sub-code (30) under the RSVP error-code "Routing Problem" (24).

This error code (30) is to be used only in an RSVP PathErr.

The following text has been included in the registry:

24 Routing Problem [RFC3209]

30 = Stitching unsupported [RFC5150]

8. Acknowledgments

The authors would like to thank Dimitri Papadimitriou and Igor Bryskin for their thorough review of the document and discussions regarding the same.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.
- [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.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [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.
- [RFC4420] Farrel, A., Ed., Papadimitriou, D., Vasseur, J.-P., and A. Ayyangar, "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4420, February 2006.

9.2. Informative References

- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC3477] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.
- [RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.
- [RFC4205] Kompella, K., Ed., and Y. Rekhter, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4205, October 2005.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4726] Farrel, A., Vasseur, J.-P., and A. Ayyangar, "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [RFC5151] Farrel, A., Ed., Ayyangar, A., and JP. Vasseur, "Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 5151, February 2008.
- [RFC5152] Vasseur, JP., Ed., Ayyangar, A., Ed., and R. Zhang, "A Per-Domain Path Computation Method for Establishing Inter-Domain Traffic Engineering (TE) Label Switched Paths (LSPs)", RFC 5152, February 2008.

[MPLS-SEC] Fang, L., Ed., Behringer, M., Callon, R., Le Roux, J.
L., Zhang, R., Knight, P., Stein, Y., Bitar, N., and R.
Graveman., "Security Framework for MPLS and GMPLS
Networks", Work in Progress, July 2007.

Authors' Addresses

Arthi Ayyangar
Juniper Networks
1194 N. Mathilda Avenue
Sunnyvale, CA 94089
EMail: arthi@juniper.net

Kireeti Kompella
Juniper Networks
1194 N. Mathilda Avenue
Sunnyvale, CA 94089
EMail: kireeti@juniper.net

JP Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, MA 01719
EMail: jpv@cisco.com

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

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

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.