

Network Working Group
Internet-Draft
Expires: October 3, 2005

A. Ayyangar, Ed.
Juniper Networks
JP. Vasseur
Cisco Systems, Inc.
April 2005

Label Switched Path Stitching with Generalized MPLS Traffic
Engineering
draft-ietf-ccamp-lsp-stitching-00

Status of this Memo

This document is an Internet-Draft and is subject to all provisions of [Section 3 of RFC 3667](#). By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she become aware will be disclosed, in accordance with [RFC 3668](#).

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>.

This Internet-Draft will expire on October 3, 2005.

Copyright Notice

Abstract

In certain scenarios, there may be a need to combine together two different Generalized Multi-Protocol Label Switching (GMPLS) Label Switched Paths (LSPs) such that in the data plane, a single end-to-end (e2e) LSP is achieved and all traffic from one LSP is

switched onto the other LSP. We will refer to this as "LSP stitching". This document covers cases where: a) the node performing the stitching does not require configuration of every LSP pair to be stitched together b) the node performing the stitching is not the egress of any of the LSPs c) LSP stitching not only results in an end-to-end LSP in the data plane, but there is also a corresponding end-to-end LSP (RSVP session) in the control plane. It might 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, 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.

This document describes the mechanisms to accomplish LSP stitching in the scenarios described above.

Table of Contents

- [1. Introduction](#) [4](#)
- [1.1 Conventions used in this document](#) [5](#)
- [2. Routing aspects](#) [6](#)
- [3. Signaling aspects](#) [7](#)
- [3.1 RSVP-TE signaling extensions](#) [7](#)
- [4. Summary of LSP Stitching procedures](#) [10](#)
- [4.1 LSP segment setup](#) [10](#)
- [4.2 Setup of e2e LSP](#) [10](#)
- [4.3 Stitching of e2e LSP into an LSP segment](#) [10](#)
- [5. Security Considerations](#) [12](#)
- [6. IANA Considerations](#) [13](#)
- [6.1 Attribute Flags for LSP_ATTRIBUTES object](#) [13](#)
- [6.2 New Error Codes](#) [13](#)
- [7. Acknowledgments](#) [14](#)
- [8. References](#) [15](#)
- [8.1 Normative References](#) [15](#)

8.2	Informative References	15
	Authors' Addresses	16
	Intellectual Property and Copyright Statements	17

[1](#). Introduction

LSP hierarchy ([\[2\]](#)) provides signaling and routing procedures so that:

- a. a GMPLS node can form a forwarding adjacency (FA) over the FA LSP
- b. other Label Switching Routers (LSRs) can see this FA LSP as a Traffic Engineering (TE) link
- c. the GMPLS node can nest one or more LSPs over the FA LSP. This

- covers intra-domain LSPs only.
- d. RSVP signaling for LSP setup can occur between nodes that do not have routing adjacency.

LSP stitching is a special case of LSP hierarchy. In case of LSP stitching, instead of an FA LSP, we will create an "LSP segment" between two GMPLS nodes. So an LSP segment for stitching is considered to be the moral equivalent of an FA LSP for nesting. While LSP hierarchy allows more than one LSP to be admitted into the FA-LSP, in case of LSP stitching, the desired switching type of the LSP and the switching capability of the LSP segment are such that at most one LSP may be admitted into an LSP segment. E.g. if LSP-AB is an FA-LSP between nodes A and B, then multiple LSPs, say LSP1, LSP2, LSP3 could 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 permitting LSP1..3 to share the FA-LSP LSP-AB. Each of LSP1..3 may reserve some bandwidth on LSP-AB. On the other hand, if LSP-AB is an LSP segment, then at most one LSP, say LSP1 may be 'stitched to' the LSP segment LSP-AB. No labels are exchanged for LSP1 over the LSP-AB hop (i.e. between A and B directly). Therefore, LSP-AB is dedicated to LSP1 and no other LSPs can be associated with LSP-AB. The LSPs LSP1..3 which are either nested or stitched into another LSP are termed as end-to-end (e2e) LSPs in the rest of this document.

Signaling and routing procedures for LSP stitching are basically similar to that described in [2]). The LSP segment will be seen as a TE link by all nodes in the network. Also, non-adjacent RSVP signaling defined in [2]) will still be required to stitch an LSP to an LSP segment. So, in the control plane, there is one RSVP session corresponding to the e2e LSP as well as one for each LSP segment that the e2e LSP is being stitched to. An LSP segment may be created either via a configuration trigger or dynamically due to an incoming LSP request. In this document we will highlight, where applicable, similarities and differences in the routing and signaling procedures between LSP hierarchy and LSP stitching. Additional signaling extensions required for LSP stitching are also described here.

LSP stitching SHOULD be used when the switching types (capabilities) of the LSP and the LSP segment are such that LSP hierarchy as

described in [2]) is not possible. E.g. if the e2e LSP is a lambda LSP and the LSP segment is also a lambda LSP, then in this case LSP hierarchy is not possible. LSP stitching could also be useful in networks to bypass legacy nodes which may not have certain new capabilities in the control plane and/or data plane. E.g. one suggested usage in case of P2MP RSVP LSPs ([7]) is the use of LSP stitching to stitch a P2MP RSVP LSP to an LSP segment between P2MP capable 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. LSP stitching procedures could also be used for inter-domain TE LSP signaling to stitch an inter-domain LSP to a local intra-domain TE LSP segment ([8]).

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](#) [1].

[2.](#) Routing aspects

An LSP segment is similar to an FA-LSP in that, an LSP segment like an FA-LSP is created and treated like a TE link between two GMPLS nodes whose path transits zero or more GMPLS nodes in the same instance of the GMPLS control plane. These TE links may be numbered or unnumbered. For an unnumbered LSP segment, the assignment and handling of the local and remote link identifiers is specified in [\[9\]](#). Unlike an FA-LSP, a GMPLS node does not have a data plane adjacency with the end point of the LSP segment. This implies that the traffic that arrives at the GMPLS node will be switched into the LSP segment contiguously with a label swap and no label is exchanged directly between the end nodes of the LSP segment itself. Also, a routing adjacency will not be established over an LSP segment. ISIS/OSPF may, however, flood the TE information associated with an LSP segment, which will exist in the TE database (TED) and can then be used for path computation by other GMPLS nodes in the network. The TE parameters defined for an FA in [\[2\]](#)) are also applicable to an LSP segment TE link.

Note that, while an FA-LSP TE link can admit zero or more LSPs over it, an LSP segment can admit at most one LSP over it. So, once an LSP is stitched into an LSP segment, the unreserved bandwidth on the LSP segment is set to zero. This prevents any more LSPs from being computed and admitted over the LSP segment TE link. Multiple LSP segments between the same pair of nodes may be bundled using the concept of Link Bundling ([\[10\]](#)) into a single TE link. When any component LSP segment is allocated for an LSP, the component's unreserved bandwidth MUST be set to zero and the Minimum and Maximum LSP bandwidth of the TE link SHOULD be recalculated.

[3.](#) Signaling aspects

In general, the trigger for the creation or termination of an LSP segment may be either mechanisms which are outside of GMPLS (configuration of LSP segment on the stitching node) or mechanisms within GMPLS (arrival of an LSP setup request with appropriate switching type on the stitching node).

Although not adjacent in routing, the end nodes of the LSP segment will have a signaling adjacency and will exchange RSVP messages directly between each other. When an RSVP-TE LSP is signaled over an LSP segment, the Path message MUST contain an IF_ID RSVP_HOP object [\[4\]](#) and the data interface identification MUST identify the LSP segment. For the purpose of ERO and RRO as well, an LSP segment is treated exactly like an FA.

The main difference between signaling an LSP over an LSP segment instead of over an FA-LSP is that no Labels are allocated and exchanged for the e2e LSP over the LSP segment hop. So, at most one e2e LSP is associated with one LSP segment. If a node at the head-end of an LSP segment receives a Path Msg for an LSP that

identifies the LSP segment in the ERO and the LSP segment bandwidth has already been allocated to some other LSP, then regular rules of RSVP-TE pre-emption apply. If the LSP request over the LSP segment cannot be satisfied, then the node SHOULD send back a PathErr with the error codes as described in [5].

Additional signaling extensions for stitching are described in the next section.

[3.1](#) RSVP-TE signaling extensions

The signaling extensions described here MUST be used if the LSP segment is a packet LSP and an e2e packet LSP may be stitched to it. These extensions are optional for non-packet LSPs and SHOULD be used if no other local mechanisms exist to automatically detect a requirement for stitching at both the ingress and egress nodes of the LSP segment.

If a GMPLS node desires to perform LSP stitching, then it MUST indicate this in the Path message for the LSP segment that it plans to use for stitching. This signaling explicitly informs the egress node for the LSP segment that the ingress node is planning to perform stitching over the LSP segment. This will allow the egress of the LSP segment to allocate the correct label(s) as explained below. Also, so that the head-end node can ensure that correct stitching actions were carried out at the egress node, a new flag is defined below in the RRO subobject to indicate that the LSP segment can be

used for stitching.

In order to request LSP stitching on the LSP segment, we define a new flag bit in the Attributes Flags TLV of the LSP_ATTRIBUTES object defined in [3]:

0x02 (TBD): LSP stitching desired bit - This flag will be set in the Attributes Flags TLV of the LSP_ATTRIBUTES object in the Path message for the LSP segment by the head-end of the LSP segment, that desires LSP stitching. This flag MUST not be modified by any other nodes in the network.

An LSP segment can only be used for stitching if appropriate label actions were carried out at the egress node of the LSP segment. In order to indicate this to the head-end node of the LSP segment, the following new flag bit is defined in the RRO Attributes sub-object:
0x02 (TBD): LSP segment stitching ready.

If an egress node receiving a Path message, supports the LSP_ATTRIBUTES object and the Attributes Flags TLV, and also recognizes the "LSP stitching desired" flag 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 sub-code=16 (TBD) "Stitching unsupported" to the head-end node of the LSP segment.

If an egress node receiving a Path message with the "LSP stitching desired" flag set, recognizes the object, the TLV and the flag and also supports the desired stitching behavior, then it MUST allocate a non-NULL label for that LSP segment in the corresponding Resv message. Now, so that the head-end node can ensure that the correct label actions will be carried out by the egress node and that the LSP segment can be used for stitching, the egress node MUST set the "LSP segment stitching ready" bit defined in the RRO Attribute sub-object. Also, when the egress node for the LSP segment receives a Path message for an e2e LSP using this LSP segment, it SHOULD first check if it is also the egress for the e2e LSP. If the egress node is the egress for both the LSP segment as well as the e2e TE LSP, and this is a packet LSP which requires Penultimate Hop Popping (PHP), then the node MUST send back a Resv refresh for the LSP segment with a new label corresponding to the NULL label.

Finally, if the egress node for the LSP segment 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 flag bit, then it SHOULD simply ignore the above request.

An ingress node requesting LSP stitching MUST examine the RRO

Attributes sub-object flag corresponding to the egress node for the LSP segment, to make sure that stitching actions were carried out at the egress node. It MUST NOT use the LSP segment for stitching if the "LSP segment stitching ready" flag is cleared.

The egress node MUST not allocate any Label in the Resv message for the e2e TE LSP. Similarly, in case of bidirectional e2e TE LSP, no Upstream Label is allocated over the LSP segment in the corresponding Path message. An ingress node stitching an e2e TE LSP to an LSP segment MUST ignore any Label object received in the Resv for the e2e TE LSP.

[4.](#) Summary of LSP Stitching procedures

[4.1](#) LSP segment setup

A GMPLS node that originates an LSP segment may decide to use this LSP segment for stitching. The creation of this LSP segment and its use for stitching may be dictated either by configuration or dynamically on arrival of an LSP setup request at the GMPLS node. Successful completion of signaling procedures for the LSP segment as described in [Section 3.1](#) will allow the GMPLS node to : a) advertise this LSP segment as a TE link with the bandwidth of the LSP as the unreserved bandwidth in the IGP and b) carry out stitching procedures to actually stitch an e2e LSP to the LSP segment. Similar to setup, tearing down the LSP segment may also be decided either via local configuration or due to the fact that there is no longer an e2e LSP stitched to the LSP segment. E.g. Let us consider an LSP segment LSP-AB being setup between two nodes A and B. A sends a Path message for the LSP-AB with "LSP stitching desired". If on the egress node B, stitching procedures are successfully carried out, then B will set the "LSP segment stitching ready" in the RRO sent in the Resv. Once A receives the Resv for LSP-AB and sees this bit set in the RRO, it can then use LSP-AB for stitching.

[4.2](#) Setup of e2e LSP

Other nodes in the network (in the same domain) trying to setup an e2e LSP across the network may see the LSP segment as a TE link in their TE databases and may compute a path over this TE link. In case

of an inter-domain e2e LSP, however, the LSP segment TE link, like any other basic TE link in the domain will probably not be advertised outside the domain. In this case, either per-domain path computation ([11]) or PCE based computation will permit setting up e2e LSPs over LSP segments in other domains. The LSP segment TE link may be identified as an ERO hop in the Path message of the e2e LSP message. E.g. Let us consider an e2e LSP LSP1-2 starting one hop before A on R1 and ending on node R2. R1 may compute a path for LSP1-2 over the LSP segment LSP-AB and identify the LSP-AB hop in the ERO.

[4.3](#) Stitching of e2e LSP into an LSP segment

When the Path message for an e2e LSP arrives at the GMPLS stitching node, the LSP segment to stitch the e2e LSP to is determined. The Path message for the e2e LSP is then sent directly to the LSP segment end node with the destination IP address of the Path message set to the address of the LSP segment end node. The Router Alert option MUST not be set in this case. Furthermore, when the message arrives at the end node, RSVP TTL checks MUST be disabled. The LSP segment MUST be identified in the IF_ID RSVP_HOP (PHOP) object of the Path

message. It is assumed that the receiver of this Path message can identify the LSP segment based on the data interface identification in the IF_ID RSVP_HOP. When the Resv is sent back for the e2e LSP, no Label is allocated on the LSP segment hop. E.g. When the Path message for the e2e LSP LSP1-2 arrives at node A, and the LSP segment LSP-AB to stitch LSP1-2 to has been identified (either based on explicit hop in ERO or due to local decision), then Path message for LSP1-2 is sent directly to node B with the IF_ID RSVP_HOP identifying the LSP segment 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. If B is not the egress, then Path message for LSP1-2 is propagated to R2. The Resv for LSP1-2 is sent back from B directly to A with no Label in it. Node A then propagates the Resv to R1. This stitches an e2e LSP LSP1-2 to an LSP segment LSP-AB between nodes A and B. In the data plane, this yields a series of label swaps from R1 to R2 along LSP LSP1-2.

[5.](#) Security Considerations

Similar to [\[2\]](#)), this document permits that the control interface over which RSVP messages are sent or received need not be the same as the data interface which the message identifies for switching traffic. Also, the 'sending interface' and 'receiving interface' may change as routing changes. So, these cannot be used to establish

security association between neighbors. Mechanisms described in [6] should be re-examined and may need to be altered to define new security associations based on receiver's IP address instead of the sending and receiving interfaces. Also, this document allows the IP destination address of Path and PathTear messages to be the IP address of a nexthop node (receiver's address) instead of the RSVP session destination address. So, [6] should be revisited to check if IPSec AH is now a viable means of securing RSVP-TE messages.

6. IANA Considerations

The following values have to be defined by IANA for this document. The registry is <http://www.iana.org/assignments/rsvp-parameters>.

6.1 Attribute Flags for LSP_ATTRIBUTES object

The following new flag bit is being defined for the Attributes Flags TLV in the LSP_ATTRIBUTES object. The numeric value should be assigned by IANA.

LSP stitching desired bit - 0x02 (Suggested value)

This flag 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 to be used in the RRO Attributes sub-object.

6.2 New Error Codes

The following new error sub-code is being defined under the RSVP error-code "Routing Problem" (24). The numeric error sub-code value should be assigned by IANA.

Stitching unsupported - sub-code 16 (Suggested value)

This error code is to be used only in a RSVP PathErr.

[7.](#) Acknowledgments

The authors would like to thank Adrian Farrel and Kireeti Kompella for their comments and suggestions.

8. References

8.1 Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [2] Kompella, K. and Y. Rekhter, "LSP Hierarchy with Generalized MPLS TE", Internet-Draft [draft-ietf-mpls-lsp-hierarchy-08](#), September 2002.

- [3] Farrel, A., Papadimitriou, D. and J. Vasseur, "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using RSVP-TE", Internet-Draft [draft-ietf-mpls-rsvpte-attributes-04](#), July 2004.

- [4] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [5] 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.
- [6] Baker, F., Lindell, B. and M. Talwar, "RSVP Cryptographic Authentication", [RFC 2747](#), January 2000.

8.2 Informative References

- [7] Aggarwal, R., "Extensions to RSVP-TE for Point to Multipoint TE LSPs", Internet-Draft [draft-ietf-mpls-rsvp-te-p2mp-01](#), January 2005.
- [8] Ayyangar, A. and J. Vasseur, "Inter domain GMPLS Traffic Engineering - RSVP-TE extensions", Internet-Draft [draft-ietf-ccamp-inter-domain-rsvp-te-00](#), February 2005.
- [9] Kompella, K. and Y. Rekhter, "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", [RFC 3477](#), January 2003.
- [10] Kompella, K., Rekhter, Y. and L. Berger, "Link Bundling in MPLS Traffic Engineering", Internet-Draft [draft-ietf-mpls-bundle-06](#), December 2004.
- [11] Vasseur, J., "A Per-domain path computation method for computing Inter-domain Traffic Engineering (TE) Label Switched

Path (LSP)",
Internet-Draft [draft-vasseur-ccamp-inter-domain-pd-path-comp-00](#), Febru
.

Authors' Addresses

Arthi Ayyangar (editor)
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: arthi@juniper.net

Jean Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, MA 01719
US

Email: jpv@cisco.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

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 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.

Copyright Statement

Copyright (C) The Internet Society (2005). 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.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.