

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: August 17, 2013

Zafar Ali
George Swallow
Clarence Filsfils
Ori Gerstel
Matt Hartley
Cisco Systems
February 18, 2013

Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)
Extension for Label Switched Path (LSP) Inquiry

[draft-ali-ccamp-lsp-inquiry-00.txt](#)

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on August 15, 2013.

Copyright Notice

Copyright (c) 2012 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November

10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow

Internet Draft [draft-ali-ccamp-lsp-inquiry-00.txt](#)

modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

RSVP-TE reoptimization procedure for Packet Switch Capable (PSC) tunnels and non-PSC tunnels has some differences. This document highlights these differences, describes how existing procedures can be used for reoptimization of non-PSC tunnels and proposes some enhancements for reoptimization of non-PSC tunnels.

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](#)].

Table of Contents

1.	Introduction	2
1.1.	Inquiry with Resource Locking	4
1.2.	Inquiry without Resource Locking	5
2.	RSVP-TE signaling procedure	6
2.1.	RSVP-TE Signaling for Inquiry with Resource Locking	6
2.2.	RSVP-TE Signaling for Inquiry without Resource Locking	7
3.	Security Considerations	8
4.	IANA Considerations	8
4.1.	RSVP Attribute Bit Flags	8
5.	References	
.....		9
5.1.	Normative References	9
5.2.	Informative References	9

1. Introduction

During reoptimization of PSC tunnel, existing and reoptimization LSPs may use independent labels and may install independent

Internet Draft

[draft-ali-ccamp-lsp-inquiry-00.txt](#)

forwarding states for each LSP. That is, existing and reoptimization LSPs may be simultaneously active in both control and data planes. However, for many non-PSC technologies label itself represents the underlying physical resource and hence cannot be shared between existing and reoptimization LSPs in the data plane. Consequently, for many non-PSC technologies, reoptimization LSPs can only be instantiated in the control plane. Furthermore, reoptimization LSP is not immediately available to carry any traffic and requires additional signaling for its activation in the data plane.

In this document, inquiry refers to a way to signal sharing of resources (e.g., labels, links) between the existing and reoptimization LSPs in the control plane without installing any forwarding states in the data plane (e.g., without installing cross-connections). In most non-PSC networks, inquiry step is required to access feasibility and characteristics of the reoptimization LSP before committing data plane resources and moving traffic to it. This is especially true of scenarios when route computation is not performed by ingress Label Edge Router (LER). These include (but are not limited to):

- . LSPs with loose hops in the Explicit Route Object (ERO), e.g. inter-domain LSPs.
- . Generalized Multi-Protocol Label Switching (GMPLS) User-Network Interface (UNI) where route computation may be performed by the (server layer) core Label Switch Router (LSR) [[RFC4208](#)];

In such cases, ingress LER may like to inquire about feasibility and attributes of a better path. Cost, TE metrics and SRLG values are examples of attributes that ingress LER may like to inquire about (e.g., before making a reoptimization decision). Procedures specified in [[ID.SRLG-RECORDING](#)] and [ID.METRIC-RECORDING] may be used for this purpose.

Reoptimization is an example where inquiry procedure is needed. However, inquiry is also useful for other use cases, e.g., for probing purposes. Hence, for the generalization purposes, LSP

signaled during inquiry is referred as inquiry LSP. Reoptimization LSP is an example of inquiry LSP.

Internet Draft

[draft-ali-ccamp-lsp-inquiry-00.txt](#)

Inquiry LSP may be signaled with or without resource locking in the control plane. This is detailed in the following.

1.1. Inquiry with Resource Locking

Signaling inquiry LSP with resource locking is depicted in Figure 1. Specifically, Path message (Path1) is signaled such that resource activation is Data Plane (DP) is skipped. However, inquiry LSP reserves resources in the Control Plane (CP), i.e., resources/ labels are locked/ committed in the control plane. Nonetheless, resources/ labels are still shared with the existing LSP(s) belonging to the tunnel inquiry LSP belongs to.

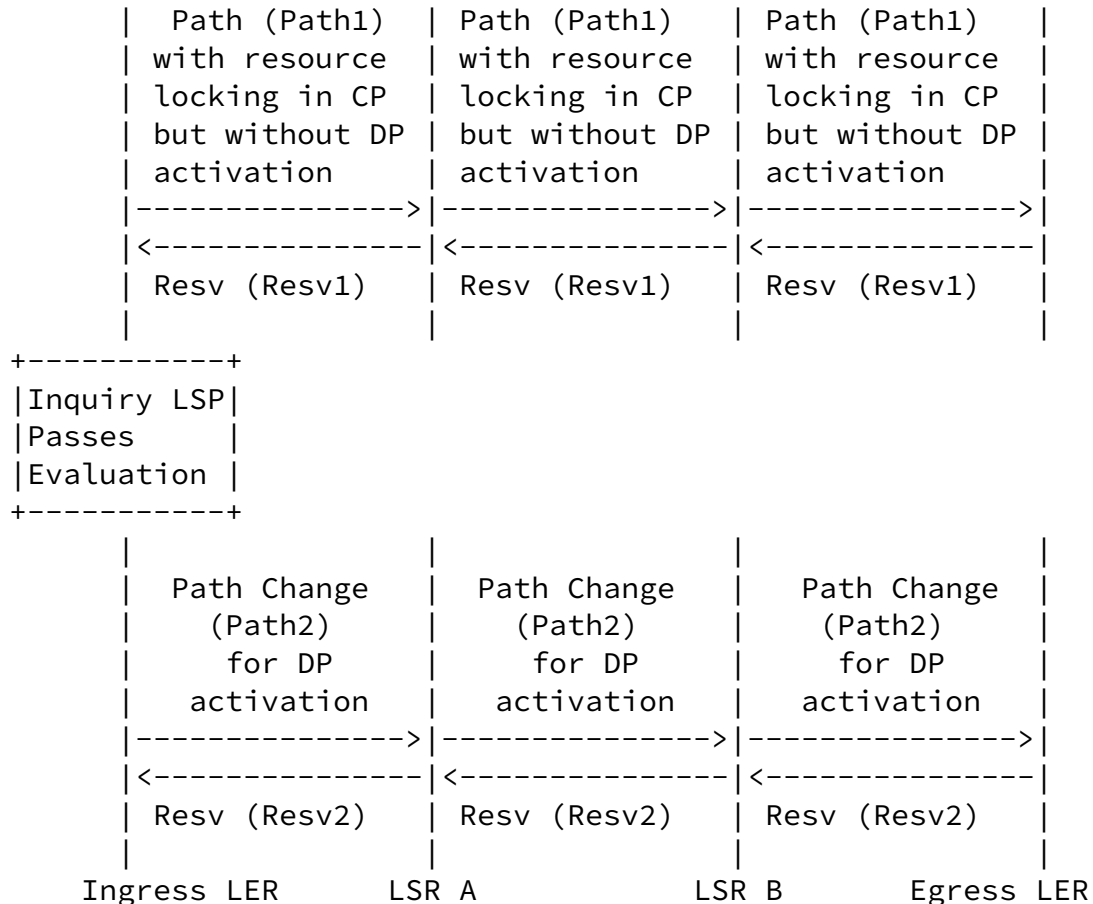


Figure 1: Inquiry LSP Signaling with Resource Locking

By the time Resv1 for the initial Path message (Path1) is received by the ingress LER, the ingress LER knows about feasibility and the requested attributes of the inquiry LSR. Please note that to ensure resource locking at the right priority, inquiry LSP needs to be signaled using the same setup and hold priority as existing LSP.

After finding feasibility and the requested attributes of the inquiry LSP, the ingress LER evaluates inquiry LSP. E.g., if the inquiry LSP is signaled for reoptimization purposes, the ingress LER

Ali, Swallow, Filsfils, et al Expires August 2013

[Page 4]

Internet Draft

[draft-ali-ccamp-lsp-inquiry-00.txt](#)

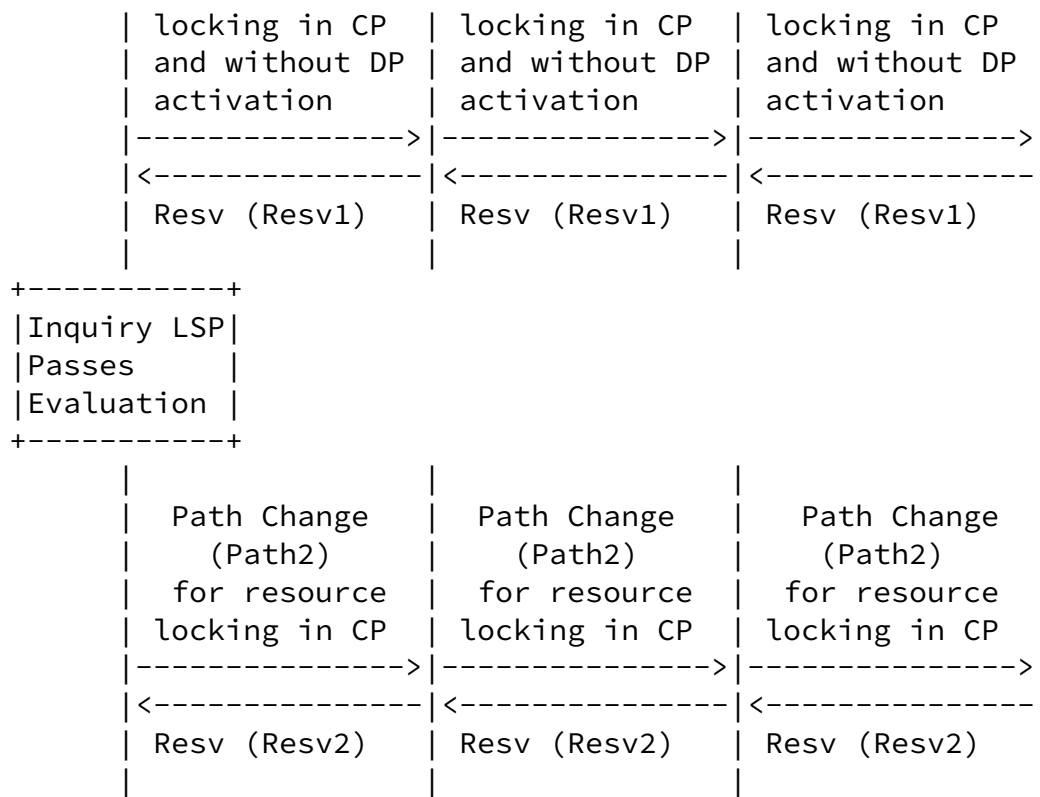
determines if it would like to reoptimize the existing LSP. If ingress LER decides not to reoptimize existing LSP, it deletes the inquiry LSP (by sending Path tear message -

- this option is not shown in Figure 1). However, if the ingress LER decides to reoptimize the tunnel to use the inquire LSP, it initiates activation of the inquiry LSP in the data plane (as shown by Path2 and Resv2 signaling loop in Figure 1). How ingress LER moves traffic from existing LSP to inquire LSP for reoptimization purpose is beyond the scope of this document.

1.2. Inquiry without Resource Locking

When inquiry is performed with resource locking, the resources used by the inquiry LSP are locked in control plane and cannot be used by any LSP other than the LSP(s) belonging to the same tunnel (of course resource preemption based on setup and hold priority of the inquiry LSP is still possible). This limits network availability in the event of a failure, especially when inquiry is used frequently, e.g., for probing purposes. This issue can be addressed by signaling inquiry LSP without resource locking. In this case, the resources used by the inquiry LSP are not locked in control plane and can be used by any LSP, including the existing LSP(s) belonging to the same tunnel. Therefore, if inquiry LSP is signaled without resource locking, additional signaling is required to first lock the resources in the control plane, before the LSP can be activated in the data plane. This is depicted in the following Figure.

Path (Path1)	Path (Path1)	Path (Path1)
without resource	without resource	without resource



Internet Draft [draft-ali-ccamp-lsp-inquiry-00.txt](#)

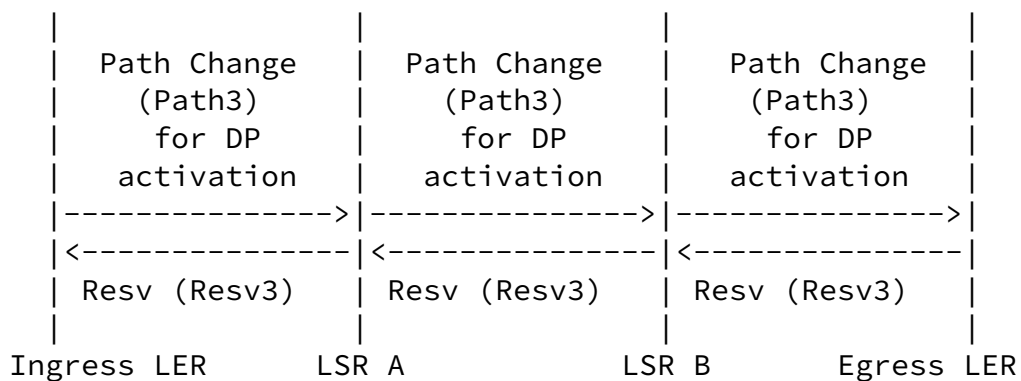


Figure 2: Inquiry LSP Signaling without Resource Locking

Initial Path message (Path1) is signaled such that resource locking in control plane and resource activation in data plane is skipped. By the time Resv1 for the initial Path message (Path1) is received by the ingress LER, the ingress LER knows about feasibility and the requested attributes of the inquiry LSP. The inquiry LSP evaluation process works in the same fashion as discussed in [Section 1.1](#).

As Path1 is signaled without resource locking, there is no guarantee

that the resources for the inquiry LSP will be available when ingress LER decides to activate the inquiry LSP. Therefore, the ingress LER first signals a path change (Path2) to lock the resource in the control plane before inquiry LSP can be activated in the data plane.

2. RSVP-TE signaling procedure

This section describes the signaling procedure for LSP inquiry with and without resource locking.

2.1. RSVP-TE Signaling for Inquiry with Resource Locking

[RFC6001] specifies procedure for signaling Soft Forwarding Adjacency (Soft FA). It defines the Pre-Planned LSP flag in the Attribute Flags TLV, which can be carried in an LSP_ATTRIBUTES object defined in [RFC5420]. The Pre-Planned LSP flag can also be used to signal inquiry LSP with resource locking.

The processing rules for the Pre-Planned LSP flag are unchanged from [RFC6001]. The procedures for the processing the Attribute Flags TLV are also unchanged from [RFC5420]. The following description is provided to help describe usage of the Pre-Planned LSP flag in the context of signaling the inquiry LSP. Example of Figure 1 is used to describe the usage.

- . Path1 message in Figure 1 is sent with the Pre-Planned LSP flag set to 1. As per [RFC6001], this enables provisioning of inquiry LSP in control plane only. In order to enable sharing if resources within the same tunnel, Path1 message is sent with shared explicit reservation style [RFC3209].
- . In order to activate inquiry LSP in the data plane, Path2 message (please see Figure 1) is sent with the Pre-Planned LSP flag set to 0.

2.2. RSVP-TE Signaling for Inquiry without Resource Locking

In order to indicate signaling of an LSP without resource provisioning in neither control plane nor data plane, a new flag in the Attribute Flags TLV, which can be carried in an

LSP_ATTRIBUTES Object, is defined as follows:

- . Pre-Planned LSP Without Resource Locking flag (to be assigned by IANA, recommended bit position TBD)

The Pre-Planned LSP Without Resource Locking flag is meaningful on a Path message. The procedure for the processing the Attribute Flags TLV follows [RFC5420]. The flag is used as described in the following.

- . If the Pre-Planned LSP Without Resource Locking flag is set to 1, the transit nodes SHOULD NOT reserve resources required by the LSP in the control plane and MUST NOT install any forwarding states associated with the LSP in the data plane. However, all LERs and LSRs are required to remember the resource (labels, links, etc.) assignments and the RSVP-TE states associated with this LSP. These resources are not locked and hence can be claimed by another LSP. If the Pre-Planned LSP Without Resource Locking flag is set to 1, the Pre-Planned LSP flag MUST be ignored. In our example of Figure 2, Path1 message is sent with the Pre-Planned LSP Without Resource Locking flag set to 1.
- . If the Pre-Planned LSP Without Resource Locking flag is set to 0 and the Pre-Planned LSP flag is set to 1, the transit nodes MUST commit resources associated with the LSP in the control plane. However, if a resource is already claimed by another LSP, the processing LSR/ LER MUST send a Path Error with code: "Admission Control Failure (1)" and subcode "LSP Admission Failure (4) and Path State Removal flag. A

processing LSR/ LER MUST NOT install any forwarding states associated with the LSP in the data plane. Path2 message in Figure 2 is sent with the Pre-Planned LSP Without Resource Locking flag set to 0 and the Pre-Planned LSP flag is set to 1. It is RECOMMENDED that no other modifications be made to other RSVP objects during this operation (signaling Path2).

- . Activation of LSP in data plane follows the procedure specific in [RFC6001]. E.g., Path2 message in Figure 2 is sent with the Pre-Planned LSP flag set to 0.

3. Security Considerations

This document does not introduce any additional security issues above those identified in [[RFC5920](#)], [[RFC2205](#)], [[RFC3209](#)], and [[RFC3473](#)] and [[RFC4874](#)].

4. IANA Considerations

4.1. RSVP Attribute Bit Flags

The IANA has created a registry and manages the space of attributes bit flags of Attribute Flags TLV as described in [section 11.3 of \[RFC5420\]](#). It is requested that the IANA makes assignments from the Attribute Bit Flags defined in this document.

This document introduces the following new Attribute Bit Flag:

- Bit number: TBD
- Defining RFC: this I-D
- Name of bit: Pre-Planned LSP Without Resource Locking Flag

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 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.
- [RFC6001] Papadimitriou, D., Vigoureux, M., Shiomoto, K.,

Brungard, D., and JL. Le Roux, "Generalized MPLS (GMPLS) Protocol Extensions for Multi-Layer and Multi-Region Networks (MLN/MRN)", [RFC 6001](#), October 2010.

[RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", [RFC 5420](#), February 2009.

5.2. Informative References

[ID.SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria, "RSVP-TE Extensions for Collecting SRLG Information", [draft-ietf-ccamp-rsvp-te-srlg-collect.txt](#), work in progress.

[ID.TE-METRIC-RECORDING] Ali, Z., Swallow, G., Filsfils, C., et al "RSVP-TE extension for recording TE Metric of a Label Switched Path", [draft-ali-ccamp-te-metric-recording](#), work in progress.

[RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", [RFC 4208](#), October 2005.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.

[RFC2209] Braden, R. and L. Zhang, "Resource ReSerVation Protocol (RSVP) -- Version 1 Message Processing Rules", [RFC 2209](#), September 1997.

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", [RFC 5920](#), July 2010.

Authors' Addresses

Zafar Ali
Cisco Systems.
Email: zali@cisco.com

George Swallow
Cisco Systems
swallow@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
cfilsfil@cisco.com

Ori Gerstel
Cisco Systems
Email: ogerstel@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com