Network Working Group
Internet Draft
Intended status: Standards Track
Expires: April 2010

A. Liu
S. Kini
Ericsson

October 13, 2009

# RSVP-TE Graceful Restart under Fast Re-route conditions
# draft-liu-mpls-rsvp-te-gr-frr-00.txt

## Status of this Memo

Abstract

During RSVP-TE graceful restart (GR), the LSR communicates with its directly connected neighbors to synchronize LSP state to recover from control plane failure conditions. However, when the LSP has undergone a Fast Re-route (FRR), the directly connected neighbor of the Point of Local Repair (PLR) or Merge Point (MP) could be down. The FRR condition of the LSP could exist for a substantial period of time. During this period the network is vulnerable to traffic loss if control plane experiences a failure on the PLR or MP. This draft describes a mechanism to extend RSVP-TE GR to work under conditions where FRR has occurred.

Table of Contents



## 1. Introduction

RSVP Graceful Restart ([RFC3473] and [RFC5063]) provides a mechanism to preserve the LSP during control plane failure so that traffic is not impacted. The mechanism uses the RSVP HELLO message defined in [RFC3209] to exchange graceful restart capability information and detect node failure. [RFC4558] introduces node-id based hello message.

RSVP fast reroute (FRR) is specified in [RFC4090] and provides a fast local repair mechanism when link or node failure occurs so that the traffic of protected LSP can be switched on PLR (Point of Local Repair) node to pre-established bypass or detour. The MP (merge point) node merges the traffic back to the protected LSP. When FRR is in effect, the traffic could stay in the bypass or detour for significant period of time. During this period of time, if PLR or MP's control plane restarts, there are certain scenarios where RSVP GR procedures cannot be applied. This document describes those scenarios and an extension to existing RSVP Hello mechanism to allow RSVP GR to operate between non directly connected neighbors. This enables RSVP GR to work when FRR is in effect.

## 2. 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 [RFC2119].

## 3. Terminology

The reader is assumed to be familiar with the terminology defined in [RFC3209], [RFC4090], [RFC3473] and [RFC5063].

## 3. Problem Statement

Per [RFC3209], RSVP Hello message is exchanged only between directly connected neighbors. [RFC3473] extends the RSVP hello mechanism to support RSVP Graceful Restart (GR) functionality. The hello message is used to carry graceful restart capability object and information used to preserve the LSP and recover the LSP state after control plane fails. If the hello session is not established, the graceful restart cannot be achieved.

[RFC4090] specifies a fast local repair mechanism to re-route a protected LSP over a bypass tunnel. When the PLR detects a link/node failure, FRR is triggered and the PLR re-routes the protected LSP over a pre-established bypass tunnel. The protected LSP is merged back at the MP node. When the PLR and MP are not immediate neighbors there is no hello session between them. In such a situation if FRR is in effect and the control plane restarts on PLR or MP node, the protected LSP may be torn down thus leading to traffic loss.

Consider the example in Figure 1. A unidirectional protected LSP is setup as R1-R2-R3-R4-R5. The unidirectional bypass tunnel for node protection is established as R2-R6-R4. R2 is PLR node and R4 is MP node. When link between R2-R3 fails or node R3 fails, traffic flows through the bypass R2-R6-R4. If R2 or R4 restarts, since hello is not running between R2 and R4, GR is not going to take effect and the protected LSP is not preserved. As a result, the traffic will get impacted.

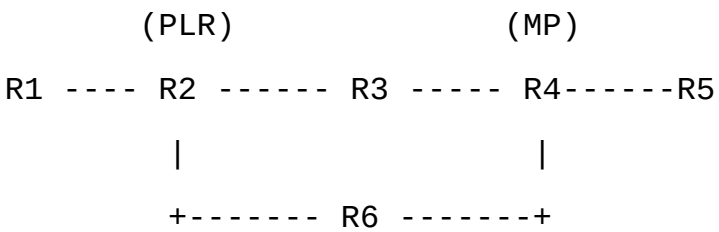


Figure 1

## 4. Extensions to RSVP Hello Message Handling

A targeted hello session must be established through the exchange of hello messages between nodes that are not immediate neighbors but have a bypass tunnel between them. These Hellos may be sent via IP forwarding or via the bypass tunnel. If a bypass tunnel does not have an independent data plane failure detection mechanism (e.g. BFD) then a targeted Hello session sending Hellos via the bypass can act as one. If the bypass has independent failure detection mechanisms, the Hello should be sent via IP forwarding. For Hellos sent using IP forwarding an IP TTL value of 255 is recommended whereas for Hellos sent via bypass tunnel the TTL of 1 should be used. The remote address must be the node id (IPv4 or IPv6) of the targeted neighbor and must also be used in the destination fields of the hello packet. The corresponding hello message handling procedures described in [RFC3209]. [RFC3473] and [RFC4558] still apply. The RSVP graceful restart procedures described in [RFC3473] and [RFC5063] also applies to these non directly connected neighbors.

At least one targeted hello session per pair of non directly connected PLR node and MP node must be established. If there are

multiple bypasses between the nodes, only one targeted hello session should be established, unless a targeted hello session is used to indicate data plane liveness. If there are multiple bypasses to a node, the selection of which one to use for sending the hello message is local policy. A hello session that uses IP forwarding should be used to avoid any forwarding problems specific to a LSP.

A non directly connected neighbor may be configured to setup the targeted hello session. This configuration may be derived from the detour/bypass configuration.

A targeted hello session may be automatically initiated by the node that is initiating the bypass. The peer node may create the hello session on receiving a hello. E.g. in Figure 1, R2 may create the Hello session to R4 when the bypass LSP R2-R6-R4 is initiated or a protected LSP R1-R2-R3-R4-R5 is associated with the bypass LSP. Also R4 may create a Hello session to R2 on seeing the Hello from R2.

A targeted hello session may also be created automatically on receiving a PATH message from a neighbor that is not directly connected but has a LSP (bypass) to it. E.g. in Figure 1, if R2 does not initiate Hello session creation, R4 may initiate creation of a Hello session with R2 on receiving a PATH message of the protected LSP from R2.

If RSVP graceful restarted is enabled, the Restart_Cap Object should be included in the hello message following procedures described in the [RFC3473].

Once the targeted hello session is established between the non-direct neighbors, the RSVP graceful restart procedures described in [RFC3473] and [RFC5063] should be followed if either node restarts.

The support of the targeted hello can be enabled or disabled by configuration which is beyond the scope of this document.

The hello session between direct neighbors should be able to co-exist with the targeted hello session.

## 5. Security Considerations

This document extends the RSVP hello message exchange to non-direct neighbors. The security considerations pertaining to the original [RFC3209] remain relevant. RSVP message security is described in [RFC2747] and provides integrity and authentication of the hello message.

## 6. IANA Considerations

This document makes no request of IANA.

## 7. References

### 7.1. Normative References

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

[RFC5063] Satyanarayana, A., Rahman, R., "Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart", RFC 5063, September 2007.

[RFC4090] Pan, P., Swallow, G., Atlas, A., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005

[RFC4558] Ali, Z., Rahman, R., Prairie, D., Papadimitriou, D., "Node-ID Based Resource Reservation Protocol (RSVP) Hello: A Clarification Statement", RFC 4558, June 2006

[RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.

## 8. Acknowledgments

The authors would like to thank Venkatesan Pradeep and Loa Andersson for their comments.

## Authors' Addresses

Autumn Liu
Ericsson
300 Holger Way, San Jose, CA 95134

Email: autumn.liu@ericsson.com

Sriganesh Kini
Ericsson
300 Holger Way, San Jose, CA 95134

Email: sriganesh.kini@ericsson.com