TEAS Working Group
Internet-Draft
Intended Status: Standards Track
Expires: April 4, 2017

M. Taillon
T. Saad, Ed.
R. Gandhi, Ed.
Z. Ali
Cisco Systems

October 1, 2016

Extensions to Resource Reservation Protocol For Fast Reroute of Traffic Engineering GMPLS LSPs
draft-ietf-teas-gmpls-lsp-fastreroute-06

Abstract

This document defines Resource Reservation Protocol - Traffic Engineering (RSVP-TE) signaling extensions to support Fast Reroute (FRR) of Packet Switched Capable (PSC) Generalized Multi-Protocol Label Switching (GMPLS) Label Switched Paths (LSPs). These signaling extensions allow the coordination of a bidirectional bypass tunnel assignment protecting a common facility in both forward and reverse directions of a co-routed bidirectional LSP. In addition, these extensions enable the re-direction of bidirectional traffic and signaling onto bypass tunnels that ensure co-routedness of data and signaling paths in the forward and reverse directions after FRR to avoid RSVP soft-state timeout.

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."

Copyright Notice

Table of Contents

1. Introduction

Packet Switched Capable (PSC) Traffic Engineering (TE) tunnels can be setup using Generalized Multi-Protocol Label Switching (GMPLS) signaling procedures specified in [RFC3473] for both unidirectional and bidirectional LSPs. Fast Reroute (FRR) [RFC4090] has been widely deployed in the packet TE networks today and is desirable for TE GMPLS LSPs. Using FRR methods also allows the leveraging of the existing mechanisms for failure detection and restoration in deployed networks.

The FRR procedures defined in [RFC4090] describe the behavior of the Point of Local Repair (PLR) to reroute traffic and signaling onto the bypass tunnel in the event of a failure for protected LSPs. Those procedures are applicable to the unidirectional protected LSPs signaled using either RSVP-TE [RFC3209] or GMPLS procedures [RFC3473]. When using the FRR procedures defined in [RFC4090] with co-routed bidirectional GMPLS LSPs, it is desired that same PLR and Merge Point (MP) pairs are selected in each direction and both PLR and MP assign the same bidirectional bypass tunnel. This document extends the FRR procedures defined in [RFC4090] to coordinate the bidirectional bypass tunnel assignment and to exchange MP labels between upstream and downstream PLRs of the protected co-routed bidirectional LSP.

When using FRR procedures with co-routed bidirectional GMPLS LSPs, it is possible in some cases for the RSVP signaling refreshes to stop reaching certain nodes along the protected LSP path after the PLRs finish rerouting signaling onto the bypass tunnels. This can occur after a failure event when using node protection bypass tunnels and when RSVP signaling is sent in-band with data. As shown in Figure 2, this is possible even with selecting the same bidirectional bypass tunnels in both directions and the same PLR and MP pairs. This is caused by the asymmetry of paths that may be taken by the bidirectional LSP's signaling in the forward and reverse directions due to upstream and downstream PLRs independently triggering FRR. In

such cases, after FRR, the RSVP soft-state timeout causes the protected bidirectional LSP to be torn down, with subsequent traffic loss.

Protection State Coordination Protocol [RFC6378] is applicable to FRR [RFC4090] for local protection of co-routed bidirectional LSPs in order to minimize traffic disruptions in both directions. However, this does not address the above mentioned problem of RSVP soft-state timeout in control plane.

This document proposes a solution to the RSVP soft-state timeout issue by providing mechanisms in the control plane to complement the

FRR procedures of [RFC4090]. The proposal allows to maintain the RSVP soft-state for co-routed bidirectional protected GMPLS LSPs and achieve co-routedness of the paths followed by the traffic and signaling in the forward and reverse directions after FRR.

Procedures defined in this document apply to GMPLS signaled PSC TE co-routed bidirectional protected LSPs and FRR co-routed bidirectional bypass tunnels. Unless otherwise specified in this document, the FRR procedures defined in [RFC4090] are not modified by this document. FRR mechanism for associated bidirectional GMPLS LSPs where two unidirectional GMPLS LSPs are bound together by using association signaling [RFC7551] is outside the scope of this document.

2. Conventions Used in This Document

2.1. Key Word Definitions

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].

2.2. Terminology

The reader is assumed to be familiar with the terminology in [RFC2205], [RFC3209], [RFC3471], [RFC3473] and [RFC4090].

Downstream PLR: Downstream Point of Local Repair. The PLR that

locally detects a failure in the downstream direction of the traffic flow and reroutes traffic in the same direction of the protected bidirectional LSP RSVP Path signaling. A downstream PLR has a corresponding downstream MP.

Downstream MP: Downstream Merge Point. The LSR where one or more backup tunnels rejoin the path of the protected LSP in the downstream direction of the traffic flow. The same LSR may be both a downstream MP and an upstream PLR simultaneously.

Upstream PLR: Upstream Point of Local Repair. The PLR that locally detects a failure in the upstream direction of the traffic flow and reroutes traffic in the opposite direction of the protected bidirectional LSP RSVP Path signaling. An upstream PLR has a corresponding upstream MP.

Upstream MP: Upstream Merge Point. The LSR where one or more backup tunnels rejoin the path of the protected LSP in the upstream direction of the traffic flow. The same LSR may be both an upstream

MP and a downstream PLR simultaneously.

Point of Remote Repair (PRR): A downstream MP that assumes the role of upstream PLR upon receiving protected LSP's Path message over the bypass tunnel and triggers reroute of traffic and signaling in the upstream direction of the traffic flow using the procedures described in this document.

2.3. Acronyms and Abbreviations

GMPLS: Generalized Multi-Protocol Label Switching

LSP: An MPLS Label Switched Path

LSR: An MPLS Label Switching Router

MP: Merge Point

MPLS: Multi-Protocol Label Switching

PLR: Point of Local Repair

PSC: Packet Switched Capable

RSVP: Resource ReSerVation Protocol

TE: Traffic Engineering

3. Fast Reroute For Unidirectional GMPLS LSPs

The FRR procedures defined in [RFC4090] are applicable to unidirectional protected LSPs signaled using either RSVP-TE or GMPLS procedures and are not modified by the extensions defined in this document.

4. Fast Reroute for Bidirectional GMPLS LSPs

This section describes signaling procedures for FRR bidirectional bypass tunnel assignment for GMPLS signaled PSC co-routed bidirectional TE LSPs.

4.1. Bidirectional GMPLS Bypass Tunnel Direction

This document defines procedures where bidirectional GMPLS bypass tunnels are signaled in the same direction as the protected GMPLS LSPs. In other words, the bidirectional GMPLS bypass tunnels

originate on the downstream PLR and terminate on the downstream MP. As the originating downstream PLR has the policy information about the locally provisioned bypass tunnels, it always initiates the bypass tunnel assignment. The bidirectional GMPLS bypass tunnels originating from the upstream PLR and terminating on the upstream MP are outside the scope of this document.

4.2. Merge Point Labels

To correctly reroute data traffic over a node protection bypass tunnel, the downstream and upstream PLRs have to know, in advance, the downstream and upstream MP labels of the protected LSP so that data in the forward and reverse directions can be redirected through the bypass tunnel after FRR respectively.

[RFC4090] defines procedures for the downstream PLR to obtain the protected LSP's downstream MP label from recorded labels in the RECORD_ROUTE Object (RRO) of the RSVP Resv message received at the downstream PLR.

To obtain the upstream MP label, the procedures specified in [RFC4090] are used to record the upstream MP label in the RRO of the RSVP Path message of the protected LSP. The upstream PLR obtains the upstream MP label from the recorded labels in the RRO of the received RSVP Path message.

4.3. Merge Point Addresses

To correctly assign a bidirectional bypass tunnel, the downstream and upstream PLRs have to know, in advance, the downstream and upstream MP addresses.

[RFC4561] defines procedures for the downstream PLR to obtain the protected LSP's downstream MP address from the recorded Node-IDs in the RRO of the RSVP Resv message received at the downstream PLR.

To obtain the upstream MP address, the procedures specified in [RFC4561] are used to record upstream MP Node-ID in the RRO of the RSVP Path message of the protected LSP. The upstream PLR obtains the upstream MP address from the recorded Node-IDs in the RRO of the received RSVP Path message.

4.4. RRO IPv4/IPv6 Subobject Flags

RRO IPv4/IPv6 subobject flags are defined in [RFC4090], Section 4.4 and are equally applicable to the FRR procedure for bidirectional GMPLS LSPs.

The procedures defined in [RFC4090] are used by the downstream PLR to signal the IPv4/IPv6 subobject flags upstream in the RRO of the RSVP Resv message of the protected LSP. Similarly, those procedures are used by the downstream PLR to signal the IPv4/IPv6 subobject flags downstream in the RRO of the RSVP Path message of the protected LSP.

4.5. Bidirectional Bypass Tunnel Assignment Co-ordination

This document defines signaling procedures and a new BYPASS_ASSIGNMENT subobject in the RSVP RECORD_ROUTE Object (RRO) used to co-ordinate the bidirectional bypass tunnel assignment between the downstream and upstream PLRs.

4.5.1. Bidirectional Bypass Tunnel Assignment Signaling Procedure

It is desirable to coordinate the bidirectional bypass tunnel selected at the downstream and upstream PLRs so that rerouted traffic and signaling flow on co-routed paths after FRR. To achieve this, a new RSVP subobject is defined for RRO that identifies a bidirectional bypass tunnel that is assigned at a downstream PLR to protect a bidirectional LSP.

When the procedures defined in this document are in use, the BYPASS_ASSIGNMENT subobject MUST be added by each downstream PLR in the RSVP Path RRO message of the GMPLS signaled bidirectional protected LSP to record the downstream bidirectional bypass tunnel assignment. This subobject is sent in the RSVP Path RRO message every time the downstream PLR assigns or updates the bypass tunnel assignment. The downstream PLR can assign a bypass tunnel when processing the first Path message of the protected LSP, however, it can not update the forwarding plane until it receives the Resv message containing the downstream MP label.

The upstream PLR (downstream MP) simply reflects the bypass tunnel assignment in the reverse direction. The absence of BYPASS_ASSIGNMENT subobject in RRO means that the relevant node or interface is not protected by a bidirectional bypass tunnel. Hence, the upstream PLR need not assign a bypass tunnel in the reverse direction.

When the BYPASS_ASSIGNMENT subobject is added in the RRO:

- The IPv4 or IPv6 subobject containing Node-ID address MUST also be added [RFC4561]. The Node-ID address must match the source address of the bypass tunnel selected for this protected LSP.

- The BYPASS_ASSIGNMENT subobject MUST be added immediately after

the Node-ID address.

o The Label subobject MUST also be added [RFC3209].

The rules for adding an IPv4 or IPv6 Interface address subobject and Unnumbered Interface ID subobject as specified in [RFC3209] and [RFC4090] are not modified by the above procedure. The options specified in Section 6.1.3 in [RFC4990] are also applicable as long as above mentioned rules are followed when using the FRR procedures defined in this document.

An upstream PLR (downstream MP) SHOULD check all BYPASS_ASSIGNMENT subobjects in the Path RRO in order to assign a reverse bypass tunnel. The upstream PLR that detects a BYPASS_ASSIGNMENT subobject, selects a reverse bypass tunnel that terminates locally with the destination address and tunnel-ID from the subobject, and has a source address matching the Node-ID address. The RRO may contain multiple addresses to identify a node, however, the upstream PLR relies on the Node-ID address preceding the BYPASS_ASSIGNMENT subobject for identifying the bypass tunnel. If the bypass tunnel is not found, the upstream PLR SHOULD send a Notify message [RFC3473] with Error-code - FRR Bypass Assignment Error (value: TBA1) and Sub-code - Bypass Tunnel Not Found (value: TBA3) to the downstream PLR. The RRO containing BYPASS_ASSIGNMENT subobject(s) is then simply forwarded downstream in the RSVP Path message.

The bypass assignment co-ordination procedure described in this Section can be used for both facility backup described in Section 3.2 of [RFC4090] and one-to-one backup described in Section 3.1 of [RFC4090]. As specified in [RFC4090], Section 4.2, the DETOUR_OBJECT can be used in one-to-one backup method to identify detour LSPs. In one-to-one backup method, if the bypass tunnel is already in-use at the upstream PLR, it SHOULD send a Notify message [RFC3473] with Error-code - FRR Bypass Assignment Error (value: TBA1) and Sub-code - One-to-one Bypass Already In-use (value: TBA4) to the downstream PLR.

4.5.2. Multiple Bidirectional Bypass Tunnel Assignments

The upstream PLR may receive multiple bypass tunnel assignments for a protected LSP from different downstream PLRs. The choice of a reverse bypass tunnel is based on the local policy on the upstream PLR. Examples of such a policy could be to prefer link protection over node protection, or to prefer the bypass tunnel to the furthest upstream node.

As shown in Example 1 and Example 2, for the protected bidirectional

GMPLS LSP R4-R5-R6, the upstream PLR R6 receives multiple bypass tunnel assignments, one from downstream PLR R4 for node protection and one from downstream PLR R5 for link protection. In Example 1, R6 prefers the link protection bypass tunnel from downstream PLR R5 whereas in Example 2, R6 prefers the node protection bypass tunnel from downstream PLR R4.

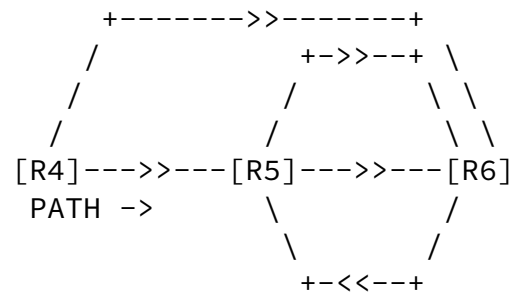

Example 1: Link protection is preferred on downstream MP

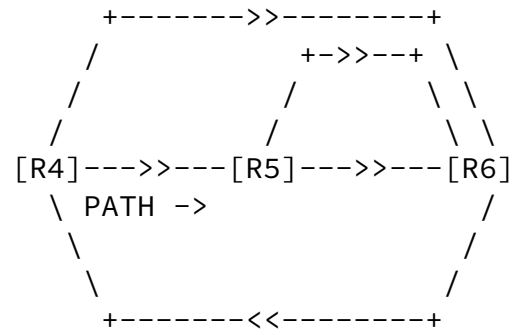

Example 2: Node protection is preferred on downstream MP

In both examples above, the upstream PLR SHOULD send a Notify message [RFC3473] with Error-code - FRR Bypass Assignment Error (value: TBA1) and Sub-code - Bypass Assignment Cannot Be Used (value: TBA2) to the downstream PLR to indicate that it cannot use the bypass tunnel assignment. The upstream PLR can then remove the bypass assignment and select an alternate bypass tunnel.

If multiple bypass tunnel assignments are present on the upstream PLR R6 at the time of a failure, any resulted asymmetry gets corrected using the re-coroute procedure after FRR as specified in Section 6.2 of this document.

When a bidirectional bypass tunnel is used, after a link failure, following procedure is followed:

- The downstream PLR reroutes traffic and RSVP Path signaling over the bidirectional bypass tunnel using the procedures defined in [RFC4090].

- Upstream PLR reroutes traffic upon detecting the link failure or upon receiving RSVP Path message over the bidirectional bypass tunnel.

- Upstream PLR also reroutes RSVP Resv signaling after receiving RSVP Path message over the bidirectional bypass tunnel.

This procedure allows both traffic and RSVP signaling to flow on symmetric paths in the forward and reverse directions of a protected bidirectional GMPLS LSP. This is described in Section 5 for link protection bypass tunnels and Section 6 for node protection bypass tunnels.

5. Link Protection for Bidirectional GMPLS LSPs

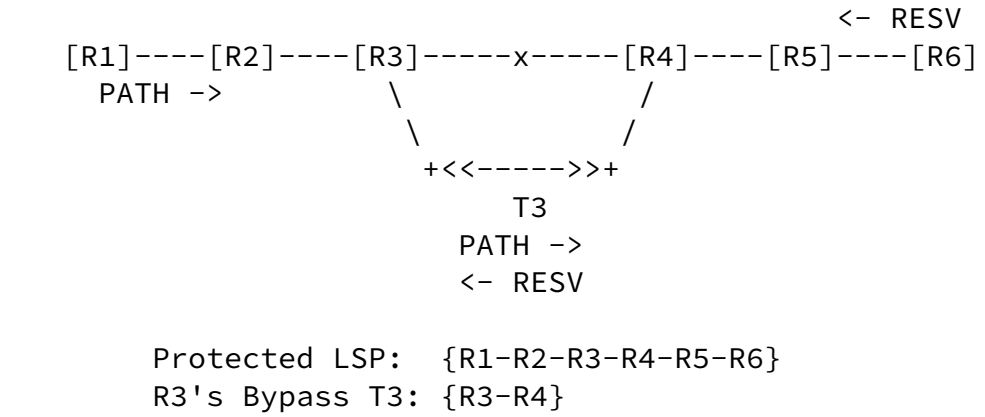

Figure 1: Flow of RSVP signaling after link failure and FRR

Consider the TE network shown in Figure 1. Assume every link in the network is protected with a link protection bypass tunnel (e.g.

bypass tunnel T3). For the protected co-routed bidirectional LSP whose head-end is on node R1 and tail-end is on node R6, each traversed node (a potential PLR) assigns a link protection co-routed bidirectional bypass tunnel.

5.1. Behavior After Link Failure

Consider the link R3-R4 on the protected LSP path fails. The downstream PLR R3 and upstream PLR R4 independently trigger fast

reroute to redirect traffic onto bypass tunnels T3 in the forward and reverse directions. The downstream PLR R3 also reroutes RSVP Path messages onto the bypass tunnel T3 using the procedures described in [RFC4090]. The upstream PLR R4 reroutes RSVP Resv messages onto the reverse bypass tunnel T3 upon receiving RSVP Path message over bypass tunnel T3.

5.2. Revertive Behavior After Fast Reroute

Revertive behavior as defined in [RFC4090], Section 6.5.2, is applicable to the link protection of bidirectional GMPLS LSPs. When using the local revertive mode, after the link R3-R4 is restored, following node behaviors apply:

- The downstream PLR R3 starts sending the Path messages and traffic flow of the protected LSP over the restored link and stops sending them over the bypass tunnel.
- The upstream PLR R4 starts sending the Resv messages and traffic flow of the protected LSP over the restored link and stops sending them over the bypass tunnel.
- When upstream PLR R4 receives the protected LSP Path messages over the restored link, if not already done, it starts sending Resv messages and traffic flow of the protected LSP over the restored link and stops sending them over the bypass tunnel.

6. Node Protection for Bidirectional GMPLS LSPs

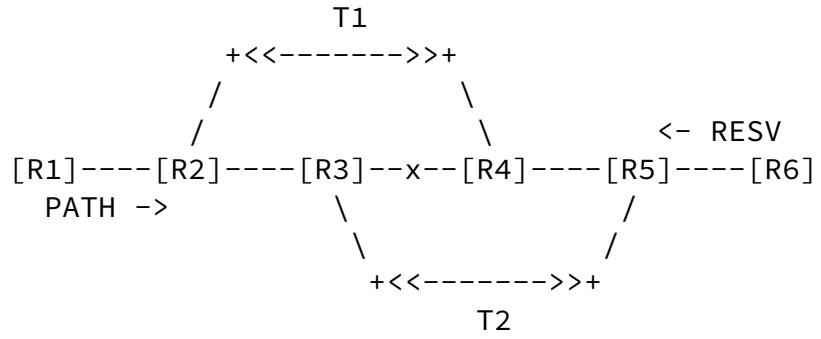

Protected LSP: {R1-R2-R3-R4-R5-R6}
R3's Bypass T2: {R3-R5}
R4's Bypass T1: {R4-R2}

Figure 2: Flow of RSVP signaling after link failure and FRR

Consider the TE network shown in Figure 2. Assume every link in the network is protected with a node protection bypass tunnel. For the protected co-routed bidirectional LSP whose head-end is on node R1 and tail-end is on node R6, each traversed node (a potential PLR) assigns a node protection co-routed bidirectional bypass tunnel.

The proposed solution introduces two phases to invoking FRR procedures by the PLR after the link failure. The first phase comprises of FRR procedures to fast reroute data traffic onto bypass tunnels in the forward and reverse directions. The second phase re-coroutes the data and signaling in the forward and reverse directions after the first phase.

6.1. Behavior After Link Failure

Consider a link R3-R4 on the protected LSP path fails. The downstream PLR R3 and upstream PLR R4 independently trigger fast reroute procedures to redirect traffic onto respective bypass tunnels T2 and T1 in the forward and reverse directions. The downstream PLR R3 also reroutes RSVP Path messages over the bypass tunnel T2 using the procedures described in [RFC4090]. Note, at this point, node R4 stops receiving RSVP Path refreshes for the protected bidirectional LSP while protected traffic continues to flow over bypass tunnels. As node R4 does not receive Path messages over the bypass tunnel, it does not reroute RSVP Resv messages over the reverse bypass tunnel.

6.2. Behavior After Link Failure To Re-coroute

The downstream MP R5 that receives rerouted protected LSP RSVP Path message through the bypass tunnel, in addition to the regular MP processing defined in [RFC4090], gets promoted to a Point of Remote Repair (PRR) role and performs the following actions to re-coroute signaling and data traffic over the same path in the reverse direction:

- Finds the bypass tunnel in the reverse direction that terminates on the downstream PLR R3. Note: the downstream PLR R3's address can be extracted from the "IPV4 tunnel sender address" in the SENDER_TEMPLATE Object of the protected LSP (see [RFC4090], Section 6.1.1).

- If reverse bypass tunnel is found and the protected LSP traffic is not already rerouted over the found bypass tunnel T2, the PRR R5 activates FRR reroute procedures to direct traffic over the found bypass tunnel T2 in the reverse direction. In addition, the PRR R5 also reroutes RSVP Resv over the bypass tunnel T2 in the reverse direction.

- If reverse bypass tunnel is not found, the PRR R5 immediately tears down the protected LSP.

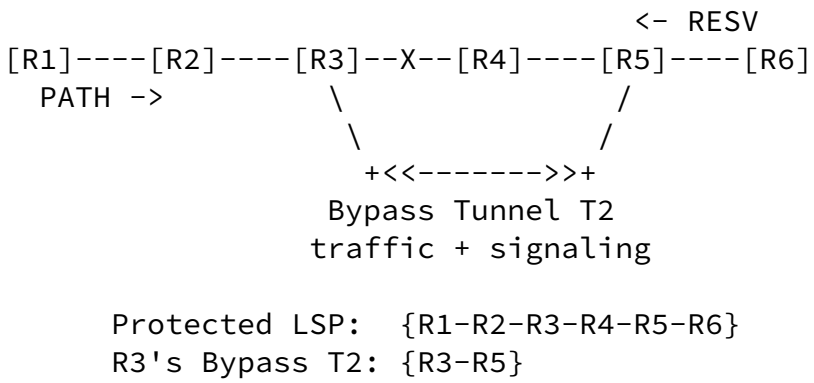

Figure 3: Flow of RSVP signaling after FRR and re-coroute

Figure 3 describes the path taken by the traffic and signaling after completing re-coroute of data and signaling in the forward and reverse paths described above. Node R4 will stop receiving the Path and Resv messages and it will timeout the RSVP soft-state, however, this will not cause the LSP to be torn down. RSVP signaling at node R2 is not affected by the FRR and re-corouting.

If downstream MP R5 receives multiple RSVP Path messages through multiple bypass tunnels (e.g. as a result of multiple failures), the PRR SHOULD identify a bypass tunnel that terminates on the farthest downstream PLR along the protected LSP path (closest to the protected bidirectional LSP head-end) and activate the reroute procedures mentioned above.

The downstream MP (upstream PLR) MAY optionally support re-corouting in data plane as follows. If the downstream MP is pre-configured with bidirectional bypass tunnel, as soon as the downstream MP receives the protected LSP packets on this bypass tunnel, it MAY switch the upstream traffic on to this bypass tunnel. In order to identify the protected LSP packets through this bypass tunnel, Penultimate Hop Popping (PHP) of the bypass tunnel MUST be disabled. The signaling procedure described above in this Section will still apply, and downstream MP checks whether the protected LSP traffic and signaling is already rerouted over the found bypass tunnel, if not, perform the above signaling procedure.

6.3. Revertive Behavior After Fast Reroute

Revertive behavior as defined in [RFC4090], Section 6.5.2, is applicable to the node protection of bidirectional GMPLS LSPs. When

using the local revertive mode, after the link R3-R4 is restored, following node behaviors apply:

- The downstream PLR R3 starts sending the Path messages and traffic flow of the protected LSP over the restored link and stops sending them over the bypass tunnel.

- The upstream PLR R4 starts sending the Resv messages and traffic flow of the protected LSP over the restored link towards downstream PLR R3 and forwarding the Path messages towards PRR R5 and stops sending them over the bypass tunnel.

- When upstream PLR R4 receives the protected LSP Path messages over the restored link, if not already done, it starts sending Resv messages and traffic flow over the restored link towards downstream PLR R3 and forwarding the Path messages towards PRR R5 and stops sending them over the bypass tunnel.

- When PRR R5 receives the protected LSP Path messages over the restored path, it starts sending Resv messages and traffic flow over the restored path and stops sending them over the bypass tunnel.

7. Unidirectional Link Failures

Unidirectional link failures may result in the traffic flowing on asymmetric paths in the forward and reverse directions. In addition, unidirectional link failures may cause RSVP soft-state timeout in the control-plane in some cases. As an example, if the unidirectional link failure is in the upstream direction (from R4 to R3 in Figures 1 and 2), the downstream PLR (node R3) can stop receiving the Resv messages of the protected LSP from the upstream PLR (node R4 in Figures 1 and 2) and this can cause RSVP soft-state timeout to occur on the downstream PLR (node R3).

A unidirectional link failure in the downstream direction (from R3 to R4 in Figure 1 and 2), does not cause RSVP soft-state timeout when using the FRR procedures defined in this document, since the upstream PLR (node R4 in Figure 1 and node R5 in Figure 2) triggers the re-coroute procedure (defined in Section 6.2 of this document) after receiving RSVP Path messages of the protected LSP over the bypass tunnel from the downstream PLR (node R3 in Figures 1 and 2).

8. Message and Object Definitions

8.1. BYPASS_ASSIGNMENT Subobject

The BYPASS_ASSIGNMENT subobject is used to inform the downstream MP of the bypass tunnel being assigned by the PLR. This can be used to coordinate the bypass tunnel assignment for the protected LSP by the downstream and upstream PLRs in the forward and reverse directions

respectively prior or after the failure occurrence.

This subobject SHOULD be inserted into the Path RRO by the downstream PLR. It SHOULD NOT be inserted into an RRO by a node which is not a downstream PLR. It MUST NOT be changed by downstream LSRs and MUST NOT be added to a Resv RRO.

The BYPASS_ASSIGNMENT subobject in RRO has the following format:

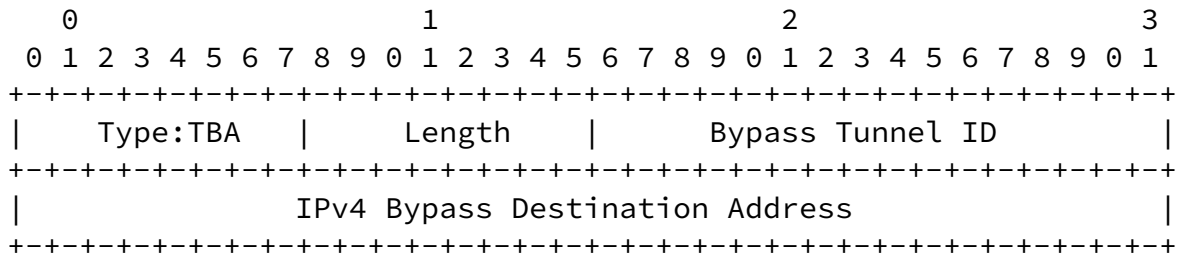

Figure 4: BYPASS ASSIGNMENT IPv4 RRO Subobject

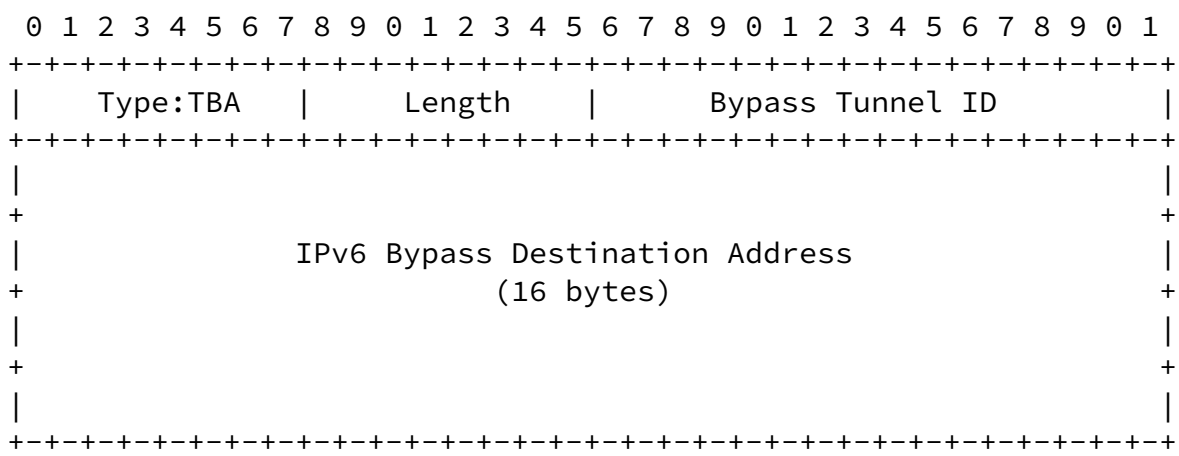

Figure 5: BYPASS_ASSIGNMENT IPv6 RRO Subobject

Type

Downstream Bypass Assignment. Value is TBA by IANA.

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The length is 8 bytes with IPv4 address and 20 bytes with IPv6 object.

Bypass Tunnel ID

The bypass tunnel identifier (16 bits).

Bypass Destination Address

The bypass tunnel IPv4 or IPv6 destination address.

8.2. FRR Bypass Assignment Error Notify Message

New Error-code - FRR Bypass Assignment Error (value: TBA1) and its sub-codes are defined for the ERROR_SPEC Object (C-Type 6) [RFC2205] in this document, that is carried by the Notify message (Type 21) defined in [RFC3473] Section 4.3. This Error is sent by the upstream PLR to the downstream PLR to notify a bypass assignment error. In the Notify message, the IP destination address is set to the node address of the downstream PLR that had initiated the bypass assignment. In the ERROR_SPEC Object, IP address is set to the node address of the upstream PLR that detected the bypass assignment error. This Error MUST NOT be sent in a Path Error message. This Error does not cause protected LSP to be torn down.

9. Compatibility

New RSVP subobject BYPASS_ASSIGNMENT is defined for RECORD_ROUTE Object in this document that is carried in the RSVP Path message. Per [RFC3209], nodes not supporting this subobject will ignore the subobject but forward it without modification. As described in Section 8, this subobject is not carried in the RSVP Resv message. A new Notify message for FRR Bypass Assignment Error is defined in this document. Nodes not supporting this message will ignore it but forward it without modification.

10. Security Considerations

This document introduces a new BYPASS_ASSIGNMENT subobject for the RECORD_ROUTE Object that is carried in an RSVP signaling message. Thus in the event of the interception of a signaling message, more information about LSP's fast reroute protection can be deduced than was previously the case. This is judged to be a very minor security

risk as this information is already available by other means.

Otherwise, this document introduces no additional security considerations. For general discussion on MPLS and GMPLS related security issues, see the MPLS/GMPLS security framework [RFC5920].

11. IANA Considerations

11.1. BYPASS_ASSIGNMENT Subobject

IANA manages the "RSVP PARAMETERS" registry located at <http://www.iana.org/assignments/rsvp-parameters>. IANA is requested to assign a value for the new BYPASS_ASSIGNMENT subobject in the "Class Type 21 ROUTE_RECORD - Type 1 Route Record" registry.

This document introduces a new subobject for RECORD_ROUTE Object:

Type	Description	Carried in Path	Carried in Resv	Reference
TBA By IANA	BYPASS_ASSIGNMENT IPv4 subobject	Yes	No	This document
TBA By IANA	BYPASS_ASSIGNMENT IPv6 subobject	Yes	No	This document

11.2. FRR Bypass Assignment Error Notify Message

IANA maintains the "Resource Reservation Protocol (RSVP) Parameters" registry (see <http://www.iana.org/assignments/rsvp-parameters>). The "Error Codes and Globally-Defined Error Value Sub-Codes" subregistry is included in this registry.

This registry has been extended for the new Error-code and Sub-codes defined in this document as follows:

- Error-code TBA1: FRR Bypass Assignment Error

- Sub-code TBA2: Bypass Assignment Cannot Be Used
- Sub-code TBA3: Bypass Tunnel Not Found
- Sub-code TBA4: One-to-one Bypass Already In-use

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2205] Braden, R., Ed., 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.

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

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

[RFC4561] Vasseur, J.P., Ed., Ali, Z., and S. Sivabalan, "Definition of a Record Route Object (RRO) Node-Id Sub-Object", RFC 4561, June 2006.

12.2. Informative References

[RFC3471] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC4990] Shiomoto, K., Papneja, R., and R. Rabbat, "Use of Addresses in Generalized Multiprotocol Label Switching (GMPLS) Networks", RFC 4990, September 2007.

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

[RFC6378] Weingarten, Y., Bryant, S., Osborne, E., Sprecher, N., and A. Fulignoli, "MPLS Transport Profile (MPLS-TP) Linear Protection", RFC 6378, October 2011.

[RFC7551] Zhang, F., Ed., Jing, R., and Gandhi, R., Ed., "RSVP-TE Extensions for Associated Bidirectional LSPs", RFC 7551, May 2015.

Acknowledgements

Authors would like to thank George Swallow for his detailed and useful comments and suggestions. Authors would also like to thank Nobo Akiya, Loa Andersson, Matt Hartley and Gregory Mirsky for reviewing this document. A special thanks to Adrian Farrel for his thorough review of this document.

Contributors

Frederic Jounay
Orange CH

EMail: frederic.jounay@salt.ch

Manav Bhatia
Nokia
Banglore, India

EMail: manav.bhatia@nokia.com

Lizhong Jin
Shanghai, China

EMail: lizho.jin@gmail.com

Authors' Addresses

Mike Taillon
Cisco Systems, Inc.

EMail: mtaillon@cisco.com

Tarek Saad (editor)
Cisco Systems, Inc.

EMail: tsaad@cisco.com

Rakesh Gandhi (editor)
Cisco Systems, Inc.

EMail: rgandhi@cisco.com

Zafar Ali
Cisco Systems, Inc.

EMail: zali@cisco.com