Internet Engineering Task Force (IETF)
Request for Comments: 8271
Updates: 4090
Category: Standards Track
ISSN: 2070-1721

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

October 2017

Updates to the Resource Reservation Protocol for Fast Reroute of Traffic Engineering GMPLS Label Switched Paths (LSPs)

Abstract

This document updates the Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Fast Reroute (FRR) procedures defined in RFC 4090 to support Packet Switch Capable (PSC) Generalized Multiprotocol Label Switching (GMPLS) Label Switched Paths (LSPs). These updates 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 updates enable the redirection of bidirectional traffic onto bypass tunnels that ensure the co-routing of data paths in the forward and reverse directions after FRR and avoid RSVP soft-state timeout in the control plane.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc8271.

Table of Contents

1. Introduction

Packet Switch Capable (PSC) Traffic Engineering (TE) Label Switched Paths (LSPs) can be set up using Generalized Multiprotocol Label Switching (GMPLS) signaling procedures specified in [RFC3473] for both unidirectional and bidirectional tunnels. The GMPLS signaling allows sending and receiving the RSVP messages in-band with the data traffic or out-of-band over a separate control channel. 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 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 that both PLR and MP assign the same bidirectional bypass tunnel. This document updates 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 of the signaling messages. This can occur after a failure event when using node protection bypass tunnels. 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 that can occur in the control plane.

This document defines a solution to the RSVP soft-state timeout issue by providing mechanisms in the control plane to complement the FRR procedures of [RFC4090]. This solution allows the RSVP soft state for co-routed, protected bidirectional GMPLS LSPs to be maintained in the control plane and enables co-routing of the traffic paths in the forward and reverse directions after FRR.

The procedures defined in this document apply to PSC TE co-routed, protected bidirectional LSPs and co-routed bidirectional FRR bypass tunnels both signaled by GMPLS. Unless otherwise specified in this document, the FRR procedures defined in [RFC4090] are not modified by this document. The 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", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 can 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 can 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 the protected LSP's rerouted Path message and triggers reroute of traffic and signaling in the upstream direction of the traffic flow using the procedures described in this document.

2.3. Abbreviations

GMPLS: Generalized Multiprotocol Label Switching

LSP: Label Switched Path

LSR: Label Switching Router

MP: Merge Point

MPLS: Multiprotocol Label Switching

PLR: Point of Local Repair

PSC: Packet Switch Capable

RSVP: Resource Reservation Protocol

TE: Traffic Engineering

3. Fast Reroute for Unidirectional GMPLS LSPs

The FRR procedures defined in [RFC4090] for RSVP-TE signaling [RFC3209] are equally applicable to the unidirectional protected LSPs signaled using GMPLS [RFC3473] and are not modified by the updates defined in this document except for the following:

When using the GMPLS out-of-band signaling [RFC3473], after a link failure event, the RSVP messages are not rerouted over the bypass tunnel by the downstream PLR but instead are rerouted over a control channel to the downstream MP.

4. Bypass Tunnel Assignment for Bidirectional GMPLS LSPs

This section describes signaling procedures for FRR bidirectional bypass tunnel assignment for GMPLS signaled PSC co-routed bidirectional TE LSPs for both in-band and out-of-band signaling.

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 PLRs and terminate on the corresponding downstream MPs. 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 PLRs and terminating on the corresponding upstream MPs 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 the protected 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 Coordination

This document defines signaling procedures and a new BYPASS_ASSIGNMENT subobject in the RSVP RECORD_ROUTE Object (RRO) used to coordinate 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 the rerouted traffic flows 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 as long as it has a topological view of the downstream MP and the traversed path information in the Explicit Route Object (ERO). For the protected LSP where the downstream MP cannot be determined from the first Path message (e.g., when using loose hops in the ERO), the downstream PLR needs to wait for the Resv message with RRO in order to assign a bypass tunnel. However, in both cases, the downstream PLR cannot

update the data plane until it receives Resv messages containing the MP labels.

The upstream PLR (downstream MP) simply reflects the bypass tunnel assignment in the reverse direction. The absence of the BYPASS_ASSIGNMENT subobject in Path 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 Path RRO:

- The IPv4 or IPv6 subobject containing the 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.

- 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 the 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 to see if the destination address in the BYPASS_ASSIGNMENT matches the address of the upstream PLR. For each BYPASS_ASSIGNMENT subobject that matches, the upstream PLR looks for a tunnel that has a source address matching the downstream PLR that inserted the BYPASS_ASSIGNMENT, as indicated by the Node-ID address and the same Tunnel ID as indicated in the BYPASS_ASSIGNMENT. The RRO can 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 44) and Sub-code "Bypass Tunnel Not Found" (value 1) to the downstream PLR. Upon receiving this error, the downstream PLR SHOULD remove the bypass tunnel assignment and select an alternate bypass tunnel if one available. The RRO containing BYPASS_ASSIGNMENT subobject(s) is then simply forwarded downstream in the RSVP Path message.

A downstream PLR may add, remove, or change the bypass tunnel assignment for a protected LSP resulting in the addition, removal, or modification of the BYPASS_ASSIGNMENT subobject in the Path RRO, respectively. In this case, the downstream PLR SHOULD generate a modified Path message and forward it downstream. The downstream MP SHOULD check the RRO in the received Path message and update the bypass tunnel assignment in the reverse direction accordingly.

4.5.2. One-to-One Bidirectional Bypass Tunnel Assignment

The bidirectional bypass tunnel assignment coordination procedure defined in this document can be used for both the facility backup described in Section 3.2 of [RFC4090] and the one-to-one backup described in Section 3.1 of [RFC4090]. As specified in Section 4.2 of [RFC4090], the DETOUR object can be used in the one-to-one backup method to identify the detour LSPs. In the 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 44) and Sub-code "One-to-One Bypass Already in Use" (value 2) to the downstream PLR. Upon receiving this error, the downstream PLR SHOULD remove the bypass tunnel assignment and select an alternate bypass tunnel if one is available.

4.5.3. Multiple Bidirectional Bypass Tunnel Assignments

The upstream PLR may receive multiple bypass tunnel assignments for a protected LSP from different downstream PLRs, leading to an asymmetric bypass tunnel assignment as shown in the following two examples.

As shown in Examples 1 and 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.

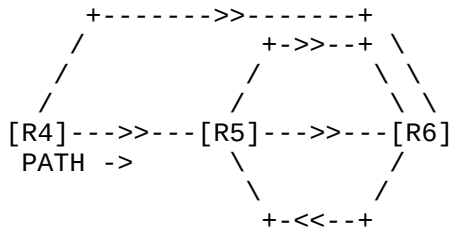

Example 1: Link Protection Is Preferred on Downstream MP

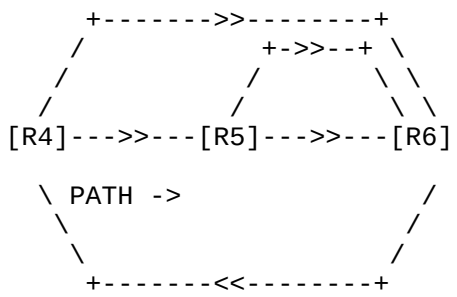

Example 2: Node Protection Is Preferred on Downstream MP

The asymmetry of bypass tunnel assignments can be avoided by using the flags in the SESSION_ATTRIBUTE object defined in Section 4.3 of [RFC4090]. In particular, the "node protection desired" flag is signaled by the head-end node to request node protection bypass tunnels. When this flag is set, both downstream PLR and upstream PLR nodes assign node protection bypass tunnels as shown in Example 2. When the "node protection desired" flag is not set, the downstream PLR nodes may only signal the link protection bypass tunnels avoiding the asymmetry of bypass tunnel assignments shown in Example 1.

When multiple bypass tunnel assignments are received, the upstream PLR SHOULD send a Notify message [RFC3473] with Error Code "FRR Bypass Assignment Error" (value 44) and Sub-code "Bypass Assignment Cannot Be Used" (value 0) to the downstream PLR to indicate that it cannot use the bypass tunnel assignment in the reverse direction. Upon receiving this error, the downstream PLR MAY remove the bypass tunnel assignment and select an alternate bypass tunnel if one is available.

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 procedure for restoring co-routing after FRR as specified in Section 5.2.2.

5. Fast Reroute for Bidirectional GMPLS LSPs with In-Band Signaling

When a bidirectional bypass tunnel is used after a link failure, the following procedure is followed when using the in-band signaling:

- The downstream PLR reroutes protected LSP traffic and RSVP Path signaling over the bidirectional bypass tunnel using the procedures defined in [RFC4090]. The RSVP Path messages are modified as described in Section 6.4.3 of [RFC4090].

- The upstream PLR reroutes protected LSP traffic upon detecting the link failure or upon receiving an RSVP Path message over the bidirectional bypass tunnel.

- The upstream PLR also reroutes protected LSP RSVP Resv signaling after receiving the modified RSVP Path message over the bidirectional bypass tunnel. The upstream PLR uses the procedure defined in Section 7 of [RFC4090] to detect that RSVP Path messages have been rerouted over the bypass tunnel by the downstream PLR. The upstream PLR does not modify the RSVP Resv message before sending it over the bypass tunnel.

The above procedure allows both traffic and RSVP signaling to flow on symmetric paths in the forward and reverse directions of a protected bidirectional GMPLS LSP. The following sections describe the handling for link protection and node protection bypass tunnels.

5.1. Link Protection for Bidirectional GMPLS LSPs

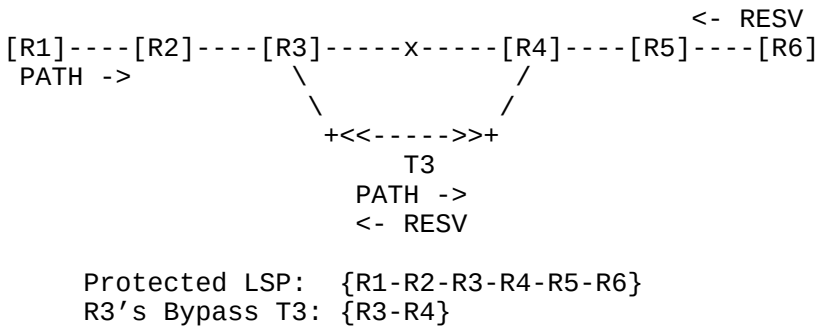

Figure 1: Flow of RSVP Signaling after Link Failure and FRR

Consider the TE network shown in Figure 1. Assume that 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.

5.1.1. Behavior after Link Failure

Consider the link R3-R4 on the protected LSP path failing. The downstream PLR R3 and upstream PLR R4 independently trigger fast reroute to redirect traffic onto bypass tunnel 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 an RSVP Path message over bypass tunnel T3.

5.1.2. Revertive Behavior after Fast Reroute

The revertive behavior 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 (in Figure 1) 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 traffic flow of the protected LSP over the restored link and stops sending it 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.

5.2. Node Protection for Bidirectional GMPLS LSPs

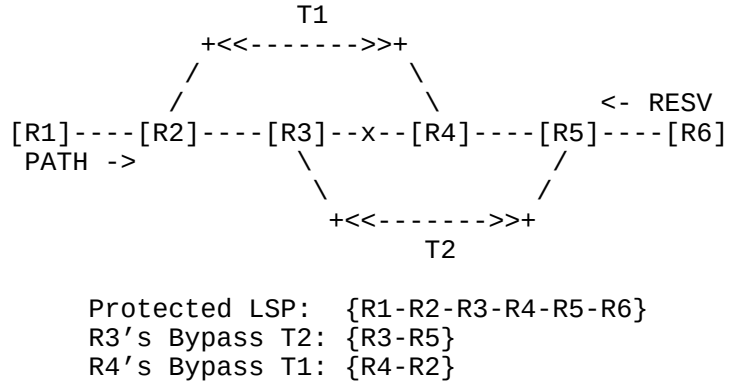

Figure 2: Flow of RSVP Signaling after Link Failure and FRR

Consider the TE network shown in Figure 2. Assume that 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 solution introduces two phases for 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 restores the co-routing of signaling and data traffic in the forward and reverse directions after the first phase.

5.2.1. Behavior after Link Failure

Consider a link R3-R4 (in Figure 2) on the protected LSP path failing. The downstream PLR R3 and upstream PLR R4 independently trigger fast reroute procedures to redirect the protected LSP 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, that 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 bypass tunnel T1, it does not reroute RSVP Resv messages over the reverse bypass tunnel T1.

5.2.2. Behavior after Link Failure to Restore Co-routing

The downstream MP R5 that receives the 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 restore co-routing 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 the 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. This can happen when the downstream PLR

has changed the bypass tunnel assignment but the upstream PLR has not yet processed the updated Path RRO and programmed the data plane when link failure occurs.

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

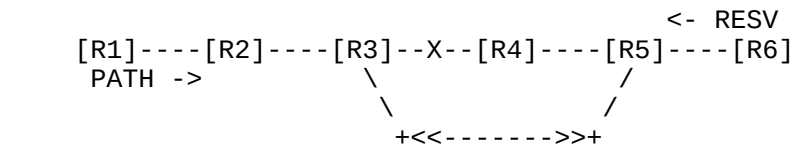

Bypass Tunnel T2

traffic + signaling

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

Figure 3: Flow of RSVP Signaling after FRR and Restoring Co-routing

Figure 3 describes the path taken by the traffic and signaling after restoring co-routing 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 restoring co-routing.

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.

5.2.2.1. Restoring Co-routing in Data Plane after Link Failure

The downstream MP (upstream PLR) MAY optionally support restoring co-routing in the data plane as follows. If the downstream MP has assigned a bidirectional bypass tunnel, as soon as the downstream MP receives the protected LSP packets on the bypass tunnel, it MAY switch the upstream traffic on to the bypass tunnel. In order to identify the protected LSP packets through the bypass tunnel, Penultimate Hop Popping (PHP) of the bypass tunnel MUST be disabled. The downstream MP checks whether the protected LSP signaling is rerouted over the found bypass tunnel, and if not, it performs the signaling procedure described in Section 5.2.2.

5.2.3. Revertive Behavior after Fast Reroute

The revertive behavior 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 (in Figures 2 and 3) is restored, the 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 (when the protected LSP is present) starts sending the 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 the traffic over the bypass tunnel.
- When upstream PLR R4 receives the protected LSP Path messages over the restored link, if not already done, the node R4 (when the protected LSP is present) 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.

5.2.4. Behavior after Node Failure

Consider the node R4 (in Figure 3) on the protected LSP path failing. The downstream PLR R3 and upstream PLR R5 independently trigger fast reroute procedures to redirect the protected LSP traffic onto bypass tunnel T2 in forward and reverse directions. The downstream PLR R3 also reroutes RSVP Path messages over the bypass tunnel T2 using the procedures described in [RFC4090]. The upstream PLR R5 reroutes RSVP Resv signaling after receiving the modified RSVP Path message over the bypass tunnel T2.

5.3. Unidirectional Link Failures

Unidirectional link failures can result in the traffic flowing on asymmetric paths in the forward and reverse directions. In addition, unidirectional link failures can 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 Figures 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 procedure to restore co-routing (defined in Section 5.2.2) after receiving RSVP Path messages of the protected LSP over the bypass tunnel from the downstream PLR (node R3 in Figures 1 and 2).

6. Fast Reroute For Bidirectional GMPLS LSPs with Out-of-Band Signaling

When using the GMPLS out-of-band signaling [RFC3473], after a link failure event, the RSVP messages are not rerouted over the bidirectional bypass tunnel by the downstream and upstream PLRs but are instead rerouted over the control channels to the downstream and upstream MPs, respectively.

The RSVP soft-state timeout after FRR as described in Section 5.2 is equally applicable to the GMPLS out-of-band signaling as the RSVP signaling refreshes can stop reaching certain nodes along the protected LSP path after the downstream and upstream PLRs finish rerouting of the signaling messages. However, unlike with the in-band signaling, unidirectional link failures as described in Section 5.3 do not result in soft-state timeout with GMPLS out-of-band signaling. Apart from this, the FRR procedure described in Section 5 is equally applicable to the GMPLS out-of-band signaling.

7. Message and Object Definitions

7.1. 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 that 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 IPv4 subobject in RRO has the following format:

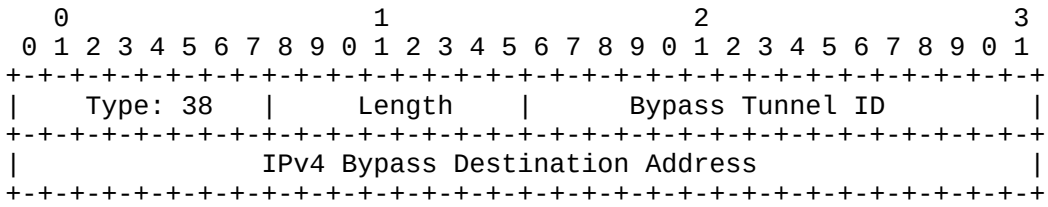

Figure 4: BYPASS ASSIGNMENT IPv4 RRO Subobject

Type

Downstream Bypass Assignment. Value is 38.

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The length is 8 bytes.

Bypass Tunnel ID

The bypass tunnel identifier (16 bits).

Bypass Destination Address

The bypass tunnel IPv4 destination address.

The BYPASS_ASSIGNMENT IPv6 subobject in RRO has the following format:

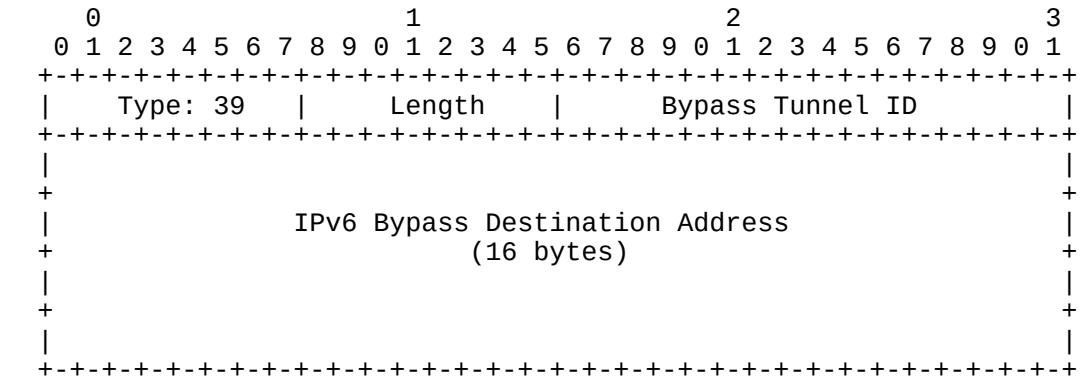

Figure 5: BYPASS_ASSIGNMENT IPv6 RRO Subobject

Type

Downstream Bypass Assignment. Value is 39.

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The length is 20 bytes.

Bypass Tunnel ID

The bypass tunnel identifier (16 bits).

Bypass Destination Address

The bypass tunnel IPv6 destination address.

7.2. FRR Bypass Assignment Error Notify Message

New Error Code "FRR Bypass Assignment Error" (value 44) 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 message 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, the 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 the protected LSP to be torn down.

8. Compatibility

New RSVP subobject BYPASS_ASSIGNMENT is defined for the 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 7, this subobject is not carried in the RSVP Resv message and is ignored by sending the Notify message for "FRR Bypass Assignment Error" (with Sub-code "Bypass Assignment Cannot Be Used") defined in this document. Nodes not supporting the Notify message defined in this document will ignore it but forward it without modification.

9. 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 the 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. If an MP does not find a matching bypass tunnel with given source and destination addresses locally, it ignores the BYPASS_ASSIGNMENT subobject. Due to this, security risks introduced by inserting a random address in this subobject is minimal. The Notify message for the "FRR Bypass Assignment Error" defined in this document does not result in tear-down of the protected LSP and does not affect service.

Security considerations for RSVP-TE and GMPLS signaling extensions are covered in [RFC3209] and [RFC3473]. Further, general considerations for securing RSVP-TE in MPLS-TE and GMPLS networks can be found in [RFC5920]. This document updates the mechanisms defined in [RFC4090], which also discusses related security measures that are also applicable to this document. As specified in [RFC4090], a PLR and its selected merge point trust RSVP messages received from each other. The security considerations pertaining to the original RSVP protocol [RFC2205] also remain relevant to the updates in this document.

10. IANA Considerations

10.1. BYPASS_ASSIGNMENT Subobject

IANA manages the "Resource Reservation Protocol (RSVP) Parameters" registry (see <http://www.iana.org/assignments/rsvp-parameters>). IANA has assigned 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 the RECORD_ROUTE Object:

Type	Description	Carried in Path	Carried in Resv	Reference
38	BYPASS_ASSIGNMENT IPv4 subobject	Yes	No	RFC 8271
39	BYPASS_ASSIGNMENT IPv6 subobject	Yes	No	RFC 8271

10.2. 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 44: FRR Bypass Assignment Error
- Sub-code 0: Bypass Assignment Cannot Be Used
- Sub-code 1: Bypass Tunnel Not Found
- Sub-code 2: One-to-One Bypass Already in Use

11. References

11.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <https://www.rfc-editor.org/info/rfc2205>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <https://www.rfc-editor.org/info/rfc3473>.

[RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <https://www.rfc-editor.org/info/rfc4090>.

[RFC4561] Vasseur, J., Ed., Ali, Z., and S. Sivabalan, "Definition of a Record Route Object (RRO) Node-Id Sub-Object", RFC 4561, DOI 10.17487/RFC4561, June 2006, <https://www.rfc-editor.org/info/rfc4561>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

11.2. Informative References

[RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, DOI 10.17487/RFC3471, January 2003, <https://www.rfc-editor.org/info/rfc3471>.

[RFC4990] Shiomoto, K., Papneja, R., and R. Rabbat, "Use of Addresses in Generalized Multiprotocol Label Switching (GMPLS) Networks", RFC 4990, DOI 10.17487/RFC4990, September 2007, <https://www.rfc-editor.org/info/rfc4990>.

[RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <https://www.rfc-editor.org/info/rfc5920>.

[RFC6378] Weingarten, Y., Ed., Bryant, S., Osborne, E., Sprecher, N., and A. Fulignoli, Ed., "MPLS Transport Profile (MPLS-TP) Linear Protection", RFC 6378, DOI 10.17487/RFC6378, October 2011, <https://www.rfc-editor.org/info/rfc6378>.

[RFC7551] Zhang, F., Ed., Jing, R., and R. Gandhi, Ed., "RSVP-TE Extensions for Associated Bidirectional Label Switched Paths (LSPs)", RFC 7551, DOI 10.17487/RFC7551, May 2015, <https://www.rfc-editor.org/info/rfc7551>.

Acknowledgements

The authors would like to thank George Swallow for many useful comments and suggestions. The authors would like to thank Lou Berger for the guidance on this work and for providing review comments. The authors would also like to thank Nobo Akiya, Loa Andersson, Matt Hartley, Himanshu Shah, Gregory Mirsky, Mach Chen, Vishnu Pavan Beeram, and Alia Atlas for reviewing this document and providing valuable comments. A special thanks to Adrian Farrel for his thorough review of this document.

Contributors

Frederic Jounay
Orange
Switzerland

Email: frederic.jounay@salt.ch

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

Manav Bhatia
Nokia
Bangalore, India

Email: manav.bhatia@nokia.com