

MPLS Working Group
Internet-Draft
Updates: [4090](#) (if approved)
Intended status: Standards Track
Expires: August 29, 2020

M. Taillon
Cisco Systems, Inc.
T. Saad, Ed.
Juniper Networks
R. Gandhi
Cisco Systems, Inc.
A. Deshmukh
Juniper Networks
M. Jork
128 Technology
V. Beeram
Juniper Networks
February 26, 2020

RSVP-TE Summary Fast Reroute Extensions for LSP Tunnels
draft-ietf-mpls-summary-frr-rsvpte-09

Abstract

This document updates [RFC 4090](#) for the Resource Reservation Protocol (RSVP) Traffic-Engineering (TE) procedures defined for facility backup protection. The updates include extensions that reduce the amount of signaling and processing that occurs during Fast Reroute (FRR), and subsequently, improves scalability when undergoing FRR convergence after a link or node failure. These extensions allow the RSVP message exchange between the Point of Local Repair (PLR) and the Merge Point (MP) nodes to be independent of the number of protected Label Switched Paths (LSPs) traversing between them when facility bypass FRR protection is used. The signaling extensions are fully backwards compatible with nodes that do not support them.

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 <https://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 29, 2020.

Copyright Notice

Copyright (c) 2020 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 (<https://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.

Table of Contents

1.	Introduction	3
2.	Conventions Used in This Document	4
2.1.	Terminology	4
2.2.	Acronyms and Abbreviations	4
3.	Extensions for Summary FRR Signaling	5
3.1.	B-SFRR-Ready Extended ASSOCIATION Object	6
3.1.1.	IPv4 B-SFRR-Ready Extended ASSOCIATION ID	7
3.1.2.	IPv6 B-SFRR-Ready Extended ASSOCIATION ID	8
3.1.3.	Processing Rules for B-SFRR-Ready Extended ASSOCIATION Object	9
3.2.	B-SFRR-Active Extended ASSOCIATION Object	10
3.2.1.	IPv4 B-SFRR-Active Extended ASSOCIATION ID	11
3.2.2.	IPv6 B-SFRR-Active Extended ASSOCIATION ID	12
3.3.	Signaling Procedures Prior to Failure	14
3.3.1.	PLR Signaling Procedure	15
3.3.2.	MP Signaling Procedure	15
3.4.	Signaling Procedures Post Failure	16
3.4.1.	PLR Signaling Procedure	16
3.4.2.	MP Signaling Procedure	17
3.5.	Refreshing Summary FRR Active LSPs	18
4.	Backwards Compatibility	18
5.	Security Considerations	18
6.	IANA Considerations	18
7.	Acknowledgments	19
8.	Contributors	19
9.	References	19
9.1.	Normative References	19
9.2.	Informative References	20
	Authors' Addresses	20

1. Introduction

The Fast Reroute (FRR) procedures defined in [[RFC4090](#)] describe the mechanisms for the Point of Local Repair (PLR) to reroute traffic and signaling of a protected RSVP-TE Label Switched Path (LSP) onto the bypass tunnel in the event of a TE link or node failure. Such signaling procedures are performed individually for each affected protected LSP. This may eventually lead to control plane scalability and latency issues on the PLR and/or the Merge Point (MP) nodes due to limited memory and CPU processing resources. This condition is exacerbated when the failure affects a large number of protected LSPs that traverse the same PLR and MP nodes.

For example, in a large-scale RSVP-TE LSPs deployment, a single Label Switched Router (LSR) acting as a PLR node may host tens of thousands of protected RSVP-TE LSPs egressing the same protected link, and also act as an MP node for a similar number of LSPs that ingress on the same link. In the event of the failure of the link or neighbor node, the RSVP-TE control plane of the node (when acting as a PLR node) becomes busy rerouting protected LSPs over the bypass tunnel(s) in one direction, and (when acting as an MP node) becomes busy merging RSVP states from signaling received over bypass tunnels for LSP(s) in the reverse direction. Subsequently, the head-end Label Edge Routers (LERs) that are notified of the local repair at downstream LSR will attempt to (re)converge the affected RSVP-TE LSPs onto newly computed paths - possibly traversing the same previously affected LSR(s). As a result, the RSVP-TE control plane becomes overwhelmed by the amount of FRR RSVP-TE processing overhead following the link or node failure, and due to other control plane protocol(s) (e.g. the IGP) that undergo convergence on the same node at the same time too.

Today, each protected RSVP-TE LSP is signaled individually over the bypass tunnel after FRR. The changes introduced in this document allow the PLR node to assign multiple protected LSPs to a bypass tunnel group and to communicate this assignment to the MP, such that upon failure, the signaling over the bypass tunnel happens on bypass tunnel group(s). New extensions are defined in this document to update the procedures defined in [[RFC4090](#)] for facility backup protection to enable the MP node to become aware of the PLR node's bypass tunnel assignment group(s) and to allow FRR procedures between the PLR and the MP nodes to be signaled and processed on per bypass tunnel group(s).

As defined in [[RFC2961](#)], Summary Refresh procedures use MESSAGE_ID to refresh the RSVP Path and Resv states to help with scaling. The Summary FRR procedures introduced in this document build on those concepts to allow the MESSAGE_ID(s) to be exchanged on per bypass tunnel assignment group, and continue use Summary Refresh procedures

while reducing the amount of messaging that occurs after rerouting signaling over the bypass tunnel post FRR.

2. Conventions Used in This Document

2.1. Terminology

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. Acronyms and Abbreviations

The reader is assumed to be familiar with terms and abbreviations used in [[RFC3209](#)] and [[RFC4090](#)].

The following abbreviations are also used in this document:

LSR: Label Switching Router

LER: Label Edge Router

MPLS: Multiprotocol Label Switching

LSP: Label Switched Path

MP: Merge Point node as defined in [[RFC4090](#)]

PLR: Point of Local Repair node as defined in [[RFC4090](#)]

FRR: Fast Reroute as defined in [[RFC4090](#)]

B-SFRR-Ready: Bypass Summary FRR Ready Extended ASSOCIATION object. Added by the PLR node for each LSP protected by the bypass tunnel.

B-SFRR-Active: Bypass Summary FRR Active Extended ASSOCIATION object. Used to notify the MP node that one or more groups of protected LSP(s) have been rerouted over the associated bypass tunnel.

MTU: Maximum transmission unit.

3. Extensions for Summary FRR Signaling

The RSVP ASSOCIATION object is defined in [[RFC4872](#)] as a means to associate LSPs with each other. For example, in the context of GMPLS-controlled LSP(s), the ASSOCIATION object is used to associate a recovery LSP with the LSP(s) it is protecting. The Extended ASSOCIATION object is introduced in [[RFC6780](#)] to expand on the possible usage of the ASSOCIATION object and generalize the definition of the Extended Association ID field.

This document defines the use of the Extended ASSOCIATION object to carry the Summary FRR information and associate the protected LSP(s) with the bypass tunnel that protects them. Two new Association Types for the Extended ASSOCIATION object, and new Extended Association IDs are proposed in this document to describe the Bypass Summary FRR Ready (B-SFRR-Ready) and the Bypass Summary FRR Active (B-SFRR-Active) associations.

The PLR node creates and manages the Summary FRR LSP groups (identified by Bypass_Group_Identifiers) and shares the group identifier(s) with the MP via signaling.

A PLR node SHOULD assign the same Bypass_Group_Identifier to all protected LSPs provided that the protected LSPs:

- o share the same outgoing protected interface,
- o are protected by the same bypass tunnel, and
- o are assigned the same tunnel sender address that is used for backup path identification after FRR as described in [[RFC4090](#)].

This minimizes the number of bypass tunnel SFRR groups, and optimizes the amount of signaling that occurs between the PLR and the MP nodes after FRR.

A PLR node that supports Summary FRR procedures adds an Extended ASSOCIATION object with B-SFRR-Ready Extended Association ID in the RSVP Path message of the protected LSP. The PLR node adds the protected LSP Bypass_Group_Identifier, information from the assigned bypass tunnel, and MESSAGE_ID object into the B-SFRR-Ready Extended Association ID. The MP uses the information contained in the received B-SFRR-Ready Extended Association ID to refresh and merge the protected LSP Path state after FRR occurs.

An MP node that supports Summary FRR procedures adds the B-SFRR-Ready Extended ASSOCIATION object and respective Extended Association ID in the RSVP Resv message of the protected LSP to acknowledge the PLR's

bypass tunnel assignment, and provide the MESSAGE_ID object that the MP node will use to refresh the protected LSP Resv state after FRR occurs.

The MP maintains the PLR node group assignments learned from signaling, and acknowledges the group assignments to the PLR node via signaling. Once the PLR node receives the group assignment acknowledgment from the MP, the FRR signaling can proceed based on Summary FRR procedures as described in this document.

The B-SFRR-Active Extended ASSOCIATION object with Extended Association ID is sent by the PLR node after activating the Summary FRR procedures. The B-SFRR-Active Extended ASSOCIATION object with Extended Association ID is sent within the RSVP Path message of the bypass tunnel to inform the MP node that one or more groups of protected LSPs protected by the bypass tunnel are now being rerouted over the bypass tunnel.

3.1. B-SFRR-Ready Extended ASSOCIATION Object

The Extended ASSOCIATION object is populated using the rules defined below to associate a protected LSP with the bypass tunnel that is protecting it when Summary FRR procedures are enabled.

The Association Type, Association ID, and Association Source MUST be set as defined in [[RFC4872](#)] for the ASSOCIATION Object. More specifically:

Association Source:

The Association Source is set to an address of the PLR node.

Association Type:

A new Association Type is defined for B-SFRR-Ready as follows:

Value	Type
-----	-----
(TBD-1)	Bypass Summary FRR Ready Association (B-SFRR-Ready)

The Extended ASSOCIATION object's Global Association Source MUST be set according to the rules defined in [[RFC6780](#)].

The B-SFRR-Ready Extended ASSOCIATION ID is populated by the PLR node when performing Bypass Summary FRR Ready association for a protected LSP. The rules governing its population are described in the subsequent sections.

3.1.1. IPv4 B-SFRR-Ready Extended ASSOCIATION ID

The IPv4 Extended ASSOCIATION ID for the B-SFRR-Ready association type is carried inside the IPv4 Extended ASSOCIATION object and has the following format:

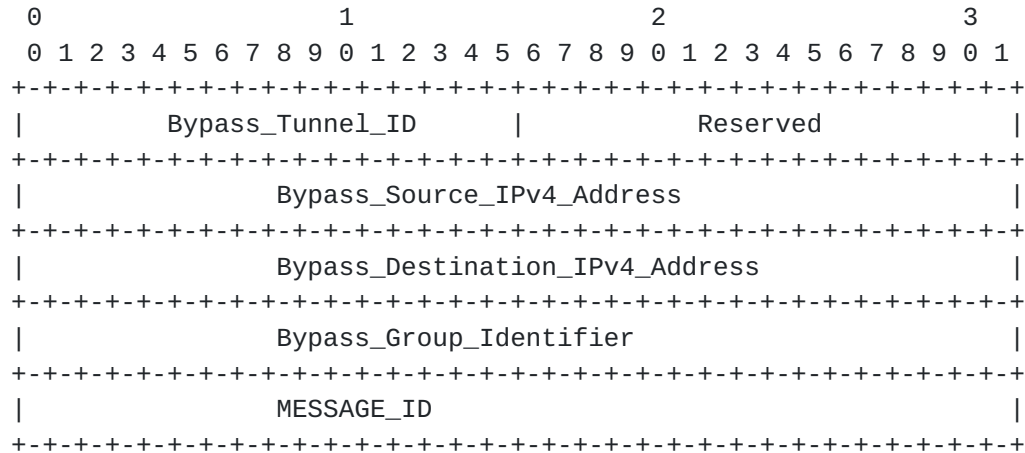


Figure 1: The IPv4 Extended ASSOCIATION ID for B-SFRR-Ready

Bypass_Tunnel_ID: 16 bits

The bypass tunnel identifier.

Reserved: 16 bits

Reserved for future use. MUST be set to zero when sending and ignored on receipt.

Bypass_Source_IPv4_Address: 32 bits

The bypass tunnel source IPV4 address.

Bypass_Destination_IPv4_Address: 32 bits

The bypass tunnel destination IPV4 address.

Bypass_Group_Identifier: 32 bits

The bypass tunnel group identifier that is assigned to the LSP.

MESSAGE_ID

A MESSAGE_ID object as defined by [[RFC2961](#)].

3.1.2. IPv6 B-SFRR-Ready Extended ASSOCIATION ID

The IPv6 Extended ASSOCIATION ID for the B-SFRR-Ready association type is carried inside the IPv6 Extended ASSOCIATION object and has the following format:

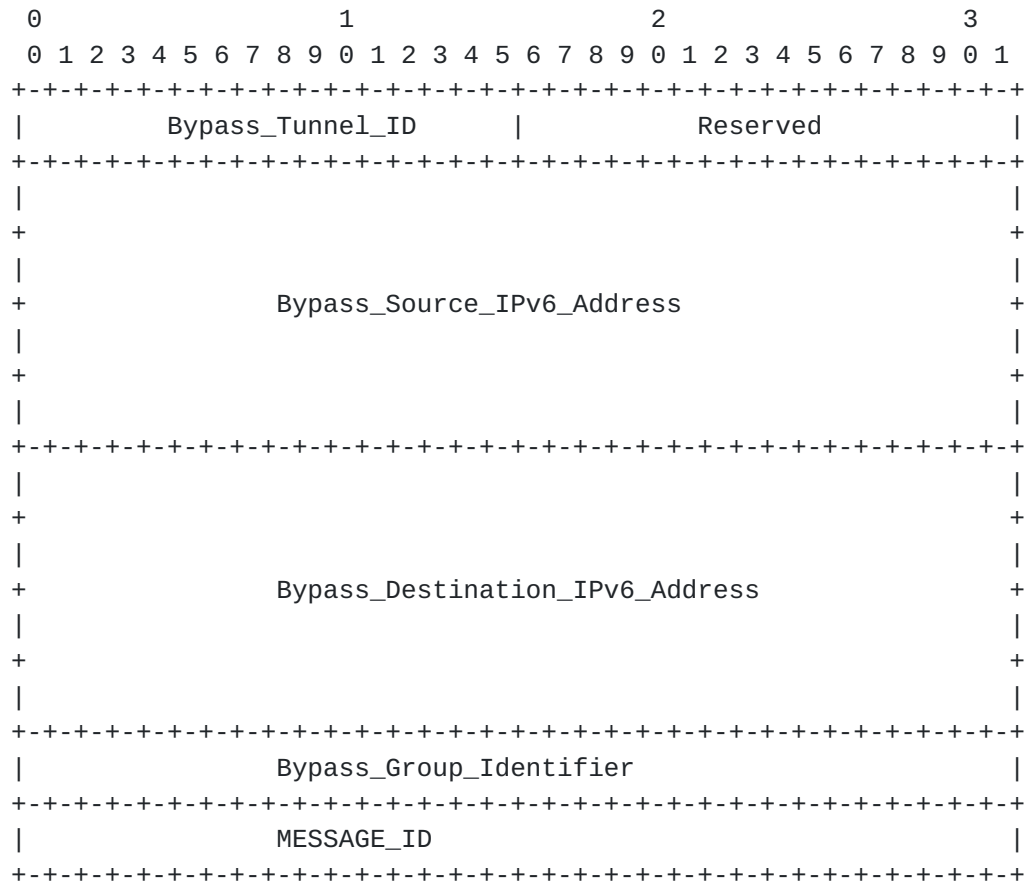


Figure 2: The IPv6 Extended ASSOCIATION ID for B-SFRR-Ready

Bypass_Tunnel_ID: 16 bits

The bypass tunnel identifier.

Reserved: 16 bits

Reserved for future use. MUST be set to zero when sending and ignored on receipt.

Bypass_Source_IPv6_Address: 128 bits

The bypass tunnel source IPV6 address.

Bypass_Destination_IPv6_Address: 128 bits

The bypass tunnel destination IPV6 address.

Bypass_Group_Identifier: 32 bits

The bypass tunnel group identifier that is assigned to the LSP.

MESSAGE_ID

A MESSAGE_ID object as defined by [[RFC2961](#)].

3.1.3. Processing Rules for B-SFRR-Ready Extended ASSOCIATION Object

A PLR node assigns a bypass tunnel and Bypass_Group_Identifier for each protected LSP. The same Bypass_Group_Identifier is used for the set of protected LSPs that share the same bypass tunnel, traverse the same egress link and are not already rerouted. The PLR node MUST generate a MESSAGE_ID object with Epoch and Message_Identifier set according to [[RFC2961](#)]. The MESSAGE_ID object flags MUST be cleared when transmitted by the PLR node and ignored when received at the MP node.

A PLR node MUST generate a new Message_Identifier each time the contents of the B-SFRR-Ready Extended ASSOCIATION ID changes (e.g. when the PLR node changes the bypass tunnel assignment).

A PLR node notifies the MP node of the bypass tunnel assignment via adding a B-SFRR-Ready Extended ASSOCIATION object and Extended Association ID in the RSVP Path message for the protected LSP using procedures described in [Section 3.3](#).

An MP node acknowledges the assignment to the PLR node by signaling the B-SFRR-Ready Extended ASSOCIATION object and Extended Association

ID within the RSVP Resv message of the protected LSP. With the exception of the MESSAGE_ID objects, all other fields of the received in the B-SFRR-Ready Extended ASSOCIATION ID in the RSVP Path message are copied into the B-SFRR-Ready Extended ASSOCIATION ID to be added in the Resv message. The MESSAGE_ID object is set according to [\[RFC2961\]](#). The MESSAGE_ID object flags MUST be cleared when transmitted by the MP node and ignored when received at the PLR node. A new Message_Identifier MUST be used to acknowledge an updated PLR node's assignment.

A PLR node considers the protected LSP as Summary FRR capable only if all the fields in the B-SFRR-Ready Extended ASSOCIATION ID that are sent in the RSVP Path message match the fields received in the RSVP Resv message (with exception of the MESSAGE_ID). If the fields do not match, or if B-SFRR-Ready Extended ASSOCIATION object is absent in a subsequent refresh, the PLR node MUST consider the protected LSP as not Summary FRR capable.

A race condition may arise for a previously Summary FRR capable protected LSP when the MP node triggers a refresh that does not contain the B-SFRR-Ready Extended ASSOCIATION object, while at the same time, the PLR triggers Summary FRR procedures due to a fault occurring concurrently. In this case, it is possible that the PLR triggers Summary FRR procedures on the protected LSP before it can receive and process the refresh from the MP node. As a result, the MP will receive a Srefresh with a Message_Identifier that is not associated with any state. As per [\[RFC2961\]](#), this results in the MP generating an Srefresh NACK for this Message_Identifier and sending it back to the PLR. The PLR processes the Srefresh NACK and replays the full Path state associated with the Message_Identifier, and subsequently recovering from this condition.

[3.2.](#) B-SFRR-Active Extended ASSOCIATION Object

The Extended ASSOCIATION object for B-SFRR-Active association type is populated by a PLR node to indicate to the MP node (bypass tunnel destination) that one or more groups of Summary FRR protected LSPs that are being protected by the bypass tunnel are being rerouted over the bypass tunnel.

The B-SFRR-Active Extended ASSOCIATION object is carried in the RSVP Path message of the bypass tunnel and signaled downstream towards the MP (bypass tunnel destination).

The Association Type, Association ID, and Association Source MUST be set as defined in [\[RFC4872\]](#) for the ASSOCIATION Object. More specifically:

Association Source:

The Association Source is set to an address of the PLR node.

Association Type:

A new Association Type is defined for B-SFRR-Active as follows:

Value	Type
(TBD-2)	Bypass Summary FRR Active Association (B-SFRR-Active)

Extended ASSOCIATION ID for B-SFRR-Active:

The B-SFRR-Active Extended ASSOCIATION ID is populated by the PLR node for the Bypass Summary FRR Active association. The rules to populate the Extended ASSOCIATION ID in this case are described below.

3.2.1. IPv4 B-SFRR-Active Extended ASSOCIATION ID

The IPv4 Extended ASSOCIATION ID for the B-SFRR-Active association type is carried inside the IPv4 Extended ASSOCIATION object and has the following format:

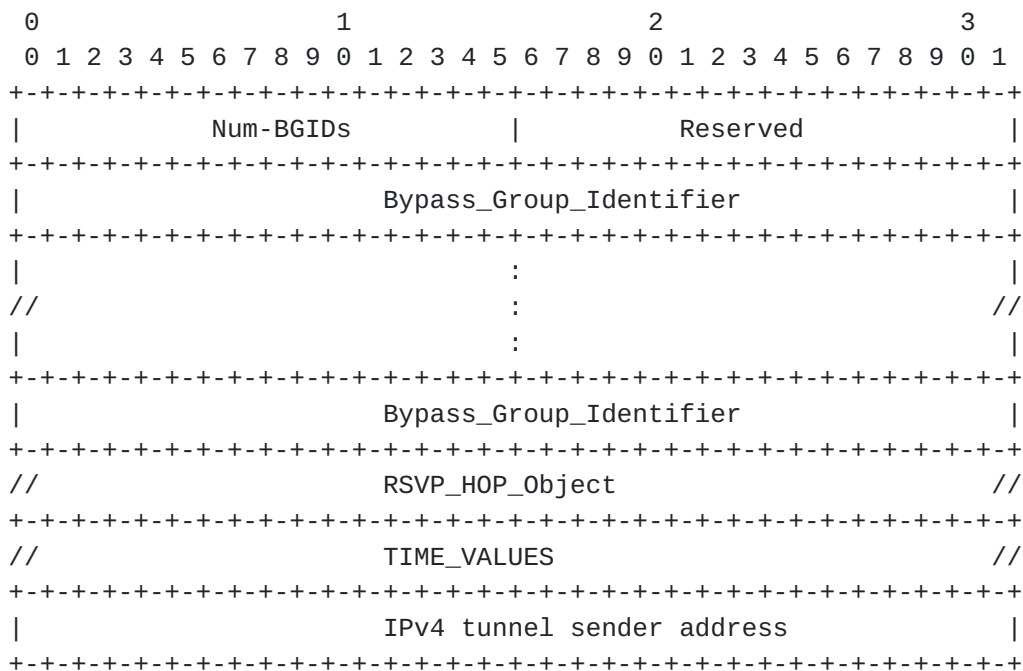


Figure 3: The IPv4 Extended ASSOCIATION ID for B-SFRR-Active

Num-BGIDs: 16 bits

Number of Bypass_Group_Identifier fields.

Reserved: 16 bits

Reserved for future use.

Bypass_Group_Identifier: 32 bits each

A Bypass_Group_Identifier that was previously signaled by the PLR using the Extended ASSOCIATION object in the B-SFRR-Ready Extended Association ID. One or more Bypass_Group_Identifiers MAY be included.

RSVP_HOP_Object: Class 3, as defined by [[RFC2205](#)]

Replacement RSVP HOP object to be applied to all LSPs associated with each of the following Bypass_Group_Identifiers. This corresponds to C-Type = 1 for IPv4 RSVP HOP.

TIME_VALUES object: Class 5, as defined by [[RFC2205](#)]

Replacement TIME_VALUES object to be applied to all LSPs associated with each of the preceding Bypass_Group_Identifiers after receiving the B-SFRR-Active Extended ASSOCIATION Object.

IPv4 tunnel sender address:

The IPv4 address that the PLR node sets to identify backup path(s) as described in [Section 6.1.1 of \[RFC4090\]](#). This address is applicable to all groups identified by Bypass_Group_Identifier(s) carried in the B-SFRR-Active Extended ASSOCIATION ID.

[3.2.2.](#) IPv6 B-SFRR-Active Extended ASSOCIATION ID

The IPv6 Extended ASSOCIATION ID for the B-SFRR-Active association type is carried inside the IPv6 Extended ASSOCIATION object and has the following format:

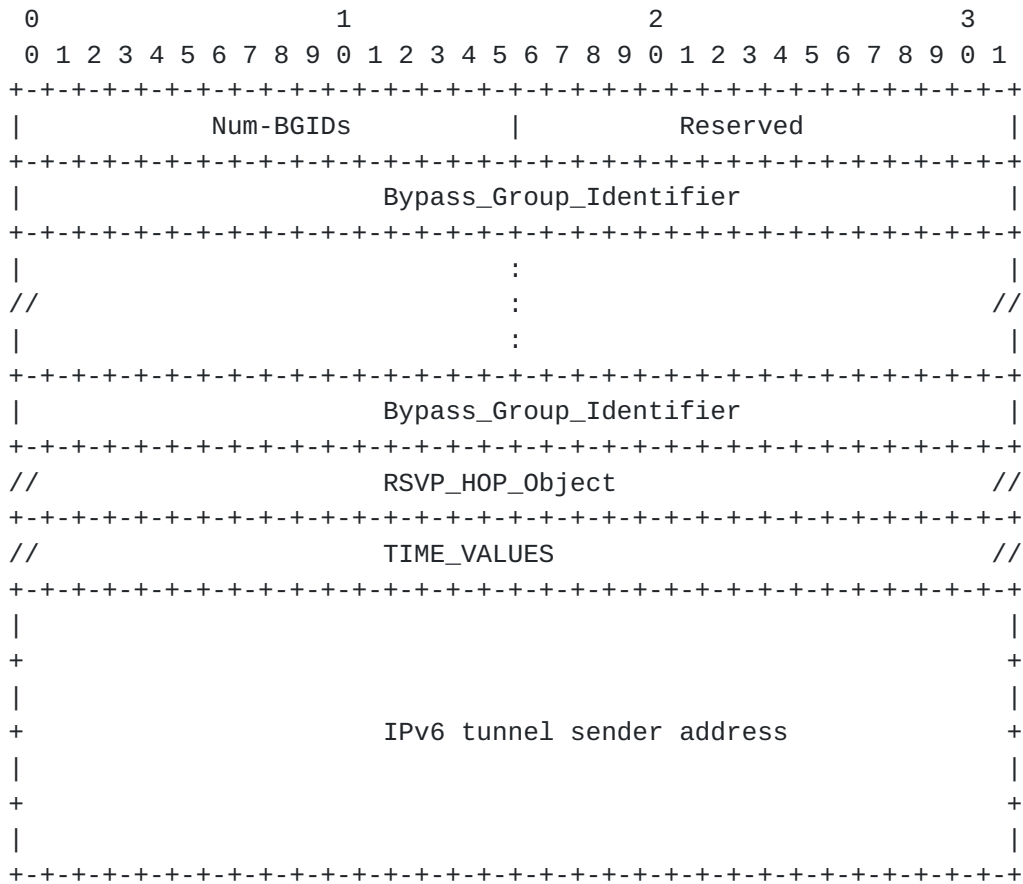


Figure 4: The IPv6 Extended ASSOCIATION ID for B-SFRR-Active

Num-BGIDs: 16 bits

Number of Bypass_Group_Identifier fields.

Reserved: 16 bits

Reserved for future use.

Bypass_Group_Identifier: 32 bits each

A Bypass_Group_Identifier that was previously signaled by the PLR using the Extended ASSOCIATION object in the B-SFRR-Ready Extended Association ID. One or more Bypass_Group_Identifier MAY be included.

RSVP_HOP_Object: Class 3, as defined by [[RFC2205](#)]

Replacement RSVP HOP object to be applied to all LSPs associated with each of the following Bypass_Group_Identifiers. This corresponds to C-Type = 2 for IPv6 RSVP HOP.

TIME_VALUES object: Class 5, as defined by [[RFC2205](#)]

Replacement TIME_VALUES object to be applied to all LSPs associated with each of the following Bypass_Group_Identifiers after receiving the B-SFRR-Active Extended ASSOCIATION Object.

IPv6 tunnel sender address:

The IPv6 address that the PLR node sets to identify backup path(s) as described in [Section 6.1.1 of \[RFC4090\]](#). This address is applicable to all groups identified by Bypass_Group_Identifier(s) carried in the B-SFRR-Active Extended ASSOCIATION ID.

[3.3.](#) Signaling Procedures Prior to Failure

Before Summary FRR procedures can be used, a handshake MUST be completed between the PLR and MP nodes. This handshake is performed using the Extended ASSOCIATION object that carries the B-SFRR-Ready Extended Association ID in both the RSVP Path and Resv messages of the protected LSP.

The facility backup method introduced in [[RFC4090](#)] takes advantage of MPLS label stacking (PLR node imposing additional MPLS label post FRR) to allow rerouting of protected traffic over the backup path. The backup path may have stricter MTU requirement and due to label stacking at PLR node, the protected traffic may exceed the backup path MTU. The operator is assumed to engineer their network to allow rerouting of protected traffic and the additional label stacking at PLR node to not exceed the backup path MTU.

When using procedures defined in this document, the PLR node MUST ensure the bypass tunnel assignment can satisfy the protected LSP MTU requirements post FRR. This avoids any packets from being dropped due to exceeding the MTU size of the backup path after traffic is rerouted on to the bypass tunnel post the failure. [Section 2.6 in \[RFC3209\]](#) describes a mechanism to determine whether a node needs to fragment or drop a packet when it exceeds the Path MTU discovered using RSVP signaling on primary LSP path. A PLR can leverage the RSVP discovered Path MTU on the backup and primary LSP paths to ensure MTU is not exceeded before or after rerouting the protected traffic on to the bypass tunnel.

3.3.1. PLR Signaling Procedure

The B-SFRR-Ready Extended ASSOCIATION object is added by each PLR node in the RSVP Path message of the protected LSP to record the bypass tunnel assignment. This object is updated every time the PLR node updates the bypass tunnel assignment and that triggers an RSVP Path change message.

Upon receiving an RSVP Resv message with B-SFRR-Ready Extended ASSOCIATION object, the PLR node checks if the expected sub-objects from the B-SFRR-Ready Extended ASSOCIATION ID are present. If present, the PLR node determines if the MP has acknowledged the current PLR node's assignment.

To be a valid acknowledgement, the received B-SFRR-Ready Extended ASSOCIATION ID contents within the RSVP Resv message of the protected LSP MUST match the latest B-SFRR-Ready Extended ASSOCIATION object and Association ID contents that the PLR node had sent within the RSVP Path message (with exception of the MESSAGE_ID).

Note, when forwarding an RSVP Resv message upstream, the PLR node SHOULD remove any/all B-SFRR-Ready Extended ASSOCIATION objects whose Bypass_Source_IPv4_Address or Bypass_Source_IPv6_Address field matches any of the PLR node addresses.

3.3.2. MP Signaling Procedure

Upon receiving an RSVP Path message with a B-SFRR-Ready Extended ASSOCIATION object, an MP node processes all (there may be multiple PLR nodes for a single MP node) B-SFRR-Ready Extended ASSOCIATION objects that have the MP node address as Bypass Destination address in the Extended Association ID.

The MP node first ensures the existence of the bypass tunnel and that the Bypass_Group_Identifier is not already FRR active. That is, an LSP cannot join a group that is already FRR rerouted.

The MP node builds a mirrored Summary FRR Group database per PLR node by associating the Bypass_Source_IPv4_Address or Bypass_Source_IPv6_Address that is carried in the IPv4 or IPv6 B-SFRR-Ready Extended ASSOCIATION IDs respectively.

The MESSAGE_ID is extracted and recorded for the protected LSP Path state. The MP node signals a B-SFRR-Ready Extended Association object and Extended Association ID in the RSVP Resv message of the protected LSP. With the exception of the MESSAGE_ID objects, all other fields of the received B-SFRR-Ready Extended ASSOCIATION object in the RSVP Path message are copied into the B-SFRR-Ready Extended

ASSOCIATION object to be added in the Resv message. The MESSAGE_ID object is set according to [\[RFC2961\]](#) with the Flags being clear.

Note, an MP may receive more than one RSVP Path message with the B-SFRR-Ready Extended ASSOCIATION object from different upstream PLR node(s). In this case, the MP node is expected to save all the received MESSAGE_IDs received from the different upstream PLR node(s). After a failure, the MP node determines and activates the state(s) associated with the Bypass_Group_Identifier(s) received in the RSVP Path message containing B-SFRR-Active Extended ASSOCIATION object that is signaled over the bypass tunnel from the PLR node, as described [Section 3.4](#)

When forwarding an RSVP Path message downstream, the MP node SHOULD remove any/all B-SFRR-Ready Extended ASSOCIATION object(s) whose Bypass_Destination_IPv4_Address or Bypass_Destination_IPv6_Address field matches any of the MP node addresses.

[3.4.](#) Signaling Procedures Post Failure

Upon detection of a fault (egress link or node failure) the PLR node will first perform the object modification procedures described by [Section 6.4.3 of \[RFC4090\]](#) for all affected protected LSPs. For the Summary FRR capable LSPs that are assigned to the same bypass tunnel a common RSVP_HOP and SENDER_TEMPLATE MUST be used.

The PLR node MUST signal non-Summary FRR capable LSPs over the bypass tunnel before signaling the Summary FRR capable LSPs. This is needed to allow for the case where the PLR node recently changed a bypass assignment and the MP has not processed the change yet.

The B-SFRR-Active Extended ASSOCIATION object is sent within the RSVP Path message of the bypass tunnel to reroute RSVP state of Summary FRR capable LSPs.

[3.4.1.](#) PLR Signaling Procedure

After a failure event, when using the Summary FRR path signaling procedures, an individual RSVP Path message is not signaled for each Summary FRR LSP. Instead, to reroute Summary FRR LSPs via the bypass tunnel, the PLR node adds the B-SFRR-Active Extended Association object in the RSVP Path message of the RSVP session of the bypass tunnel.

The RSVP_HOP_Object field in the B-SFRR-Active Extended ASSOCIATION ID is set to a common object that will be applied to all LSPs associated with the Bypass_Group_Identifiers that are carried in the B-SFRR-Active Extended ASSOCIATION ID.

The PLR node adds the Bypass_Group_Identifier(s) of group(s) that have common group attributes, including the tunnel sender address, to the same B-SFRR-Active Extended ASSOCIATION ID. Note that multiple ASSOCIATION objects, each carrying a B-SFRR-Active Extended ASSOCIATION ID, can be carried within a single RSVP Path message of the bypass tunnel and sent towards the MP as described in [\[RFC6780\]](#).

The previously received MESSAGE_ID(s) from the MP are activated on the PLR. As a result, the PLR starts sending Srefresh messages containing the specific Message_identifier(s) for the states to be refreshed.

[3.4.2.](#) MP Signaling Procedure

Upon receiving an RSVP Path message with a B-SFRR-Active Extended Association object, the MP performs normal merge point processing for each protected LSP associated with each Bypass_Group_Identifier, as if it had received an individual RSVP Path message for that LSP.

For each Summary FRR capable LSP that is being merged, the MP first modifies the Path state as follows:

1. The RSVP_HOP object is copied from the RSVP_HOP_Object field in the B-SFRR-Active Extended ASSOCIATION ID.
2. The TIME_VALUES object is copied from the TIME_VALUES field in the B-SFRR-Active Extended ASSOCIATION ID. The TIME_VALUES object contains the refresh time of the PLR node to generate refreshes and that would have exchanged in a Path message sent to the MP after the failure when no Summary FRR procedures are in effect.
3. The tunnel sender address field in the SENDER_TEMPLATE object is copied from the tunnel sender address field of the B-SFRR-Active Extended ASSOCIATION ID.
4. The ERO object is modified as per [Section 6.4.4 of \[RFC4090\]](#). Once the above modifications are completed, the MP node performs the merge processing as per [\[RFC4090\]](#).
5. The previously received MESSAGE_ID(s) from the PLR node are activated. The MP is allowed to send Srefresh messages containing the specific Message_identifier(s) for the states to be refreshed.

A failure during merge processing of any individual rerouted LSP MUST result in an RSVP Path Error message.

An individual RSVP Resv message for each successfully merged Summary FRR LSP is not signaled. The MP node SHOULD immediately use Summary Refresh procedures to refresh the protected LSP Resv state.

3.5. Refreshing Summary FRR Active LSPs

Refreshing of Summary FRR active LSPs is performed using Summary Refresh as defined by [[RFC2961](#)].

4. Backwards Compatibility

The (Extended) ASSOCIATION object is defined in [[RFC4872](#)] with a class number in the form 11bbbbbb, where b=0 or 1. This ensures compatibility with non-supporting node(s) in accordance with the procedures specified in [[RFC2205](#)], [Section 3.10](#) for unknown-class objects, Such nodes will ignore the object and forward it without any modification.

5. Security Considerations

This document updates an existing RSVP object. Thus, in the event of the interception of a signaling message, slightly more information could be deduced about the state of the network than was previously the case.

When using procedures defined in this document, FRR signaling for rerouting of protected LSP(s) states on to the bypass tunnel can be performed on a group of protected LSP(s) with a single RSVP message. This allows an intruder to potentially impact and manipulate a set of protected LSP that are assigned to the same bypass tunnel group. Note that such attack is even possible without the mechanisms proposed in this document; albeit, at an extra cost resulting from the excessive per LSP signaling that will occur.

Existing mechanisms for maintaining the integrity and authenticity of RSVP protocol messages [[RFC2747](#)] can be applied. Other considerations mentioned in [[RFC4090](#)] and [[RFC5920](#)] also apply.

6. IANA Considerations

IANA maintains the "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" registry. The "Association Type" sub-registry is included in this registry.

This registry has been updated by new Association Type for Extended ASSOCIATION Object defined in this document as follows:

Value	Name	Reference
-----	-----	-----
TBD-1	B-SFRR-Ready Association	Section 3.1
TBD-2	B-SFRR-Active Association	Section 3.2

7. Acknowledgments

The authors would like to thank Alexander Okonnikov, Loa Andersson, Lou Berger, Eric Osborne, Gregory Mirsky, Mach Chen for reviewing and providing valuable comments to this document.

8. Contributors

Nicholas Tan
Arista Networks

Email: ntan@arista.com

9. References

9.1. 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>>.
- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", [RFC 2747](#), DOI 10.17487/RFC2747, January 2000, <<https://www.rfc-editor.org/info/rfc2747>>.
- [RFC2961] Berger, L., Gan, D., Swallow, G., Pan, P., Tommasi, F., and S. Molendini, "RSVP Refresh Overhead Reduction Extensions", [RFC 2961](#), DOI 10.17487/RFC2961, April 2001, <<https://www.rfc-editor.org/info/rfc2961>>.
- [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>>.

- [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>>.
- [RFC4872] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", [RFC 4872](#), DOI 10.17487/RFC4872, May 2007, <<https://www.rfc-editor.org/info/rfc4872>>.
- [RFC6780] Berger, L., Le Faucheur, F., and A. Narayanan, "RSVP ASSOCIATION Object Extensions", [RFC 6780](#), DOI 10.17487/RFC6780, October 2012, <<https://www.rfc-editor.org/info/rfc6780>>.
- [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>>.

9.2. Informative References

- [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>>.

Authors' Addresses

Mike Taillon
Cisco Systems, Inc.

Email: mtaillon@cisco.com

Tarek Saad (editor)
Juniper Networks

Email: tsaad@juniper.net

Rakesh Gandhi
Cisco Systems, Inc.

Email: rgandhi@cisco.com

Abhishek Deshmukh
Juniper Networks

Email: adeshmukh@juniper.net

Markus Jork
128 Technology

Email: mjork@128technology.com

Vishnu Pavan Beeram
Juniper Networks

Email: vbeeram@juniper.net