Internet Engineering Task Force (IETF)
Request for Comments: 6882
Category: Experimental
ISSN: 2070-1721

K. Kumaki, Ed.
KDDI Corporation
T. Murai
Furukawa Network Solution Corp.
D. Cheng
Huawei Technologies
S. Matsushima
Softbank Telecom
P. Jiang
KDDI Corporation

March 2013

Support for Resource Reservation Protocol Traffic Engineering (RSVP-TE) in Layer 3 Virtual Private Networks (L3VPNs)

Abstract

IP Virtual Private Networks (VPNs) provide connectivity between sites across an IP/MPLS backbone. These VPNs can be operated using BGP/MPLS, and a single Provider Edge (PE) node may provide access to multiple customer sites belonging to different VPNs.

The VPNs may support a number of customer services, including RSVP and Resource Reservation Protocol Traffic Engineering (RSVP-TE) traffic. This document describes how to support RSVP-TE between customer sites when a single PE supports multiple VPNs and labels are not used to identify VPNs between PEs.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for examination, experimental implementation, and evaluation.

This document defines an Experimental Protocol for the Internet community. 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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc6882.

Table of Contents

1. Introduction

Service Providers would like to use BGP/MPLS IP VPNs [RFC4364] to support connections between Customer Edge (CE) sites. As described in [RFC5824], these connections can be MPLS Traffic Engineered (TE) Label Switched Paths (LSPs) established using extensions to RSVP [RFC3209] for a number of different deployment scenarios. The requirements for supporting MPLS-TE LSP connections across BGP/MPLS IP VPNs are documented in [RFC5824].

In order to establish a customer MPLS-TE LSP over a BGP/MPLS IP VPN, it is necessary for the RSVP-TE control messages, including the Path and Resv messages described in [RFC3209], to be handled appropriately by the Provider Edge (PE) routers. [RFC4364] allows RSVP messages sent within a VPN's context to be handled just like any other VPN data. In such a solution, the RSVP-TE component at a PE that sends messages toward a remote PE must process the messages in the context of the VPN and must ensure that the messages are correctly labeled. Similarly, when a message sent across the core is received by a PE, both labels must indicate the correct VPN context.

Implementation of the standards-based solution described in the previous paragraph is possible, but requires proper support on the PE. In particular, a PE must be able to process RSVP messages within the context of the appropriate VPN Routing and Forwarding (VRF). This may be easy to achieve in some implementations, but in others, it is not so easy.

This document defines experimental formats and mechanisms that follow a different approach. The documented approach enables the VPN identifier to be carried in the RSVP-TE protocol message so that there is no requirement for label-based VRF identification on the PE.

The experiment proposed by this document does not negate the label-based approach supported by [RFC4364]. The experiment is intended to enable research into alternate methods of supporting RSVP-TE within VPNs.

1.1. Conventions

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. Motivation

If multiple BGP/MPLS IP VPNs are supported at the same PE, new RSVP-TE extensions are required so that RSVP-TE control messages from the CEs can be handled appropriately by the PE.

2.1. Network Example

Figure 1 ("Customer MPLS TE LSPs in the context of BGP/MPLS IP VPNs") shows two VPNs supported by a core IP/MPLS network. Both VPNs have customer sites on the two PEs shown in the figure. The customer sites operate MPLS-TE LSPs.

Here, we make the following set of assumptions:

- VPN1 and VPN2 are for different customers.
- CE1 and CE3 are head-end routers.
- CE2 and CE4 are tail-end routers.
- The same address (e.g., 192.0.2.1) is assigned at CE2 and CE4.

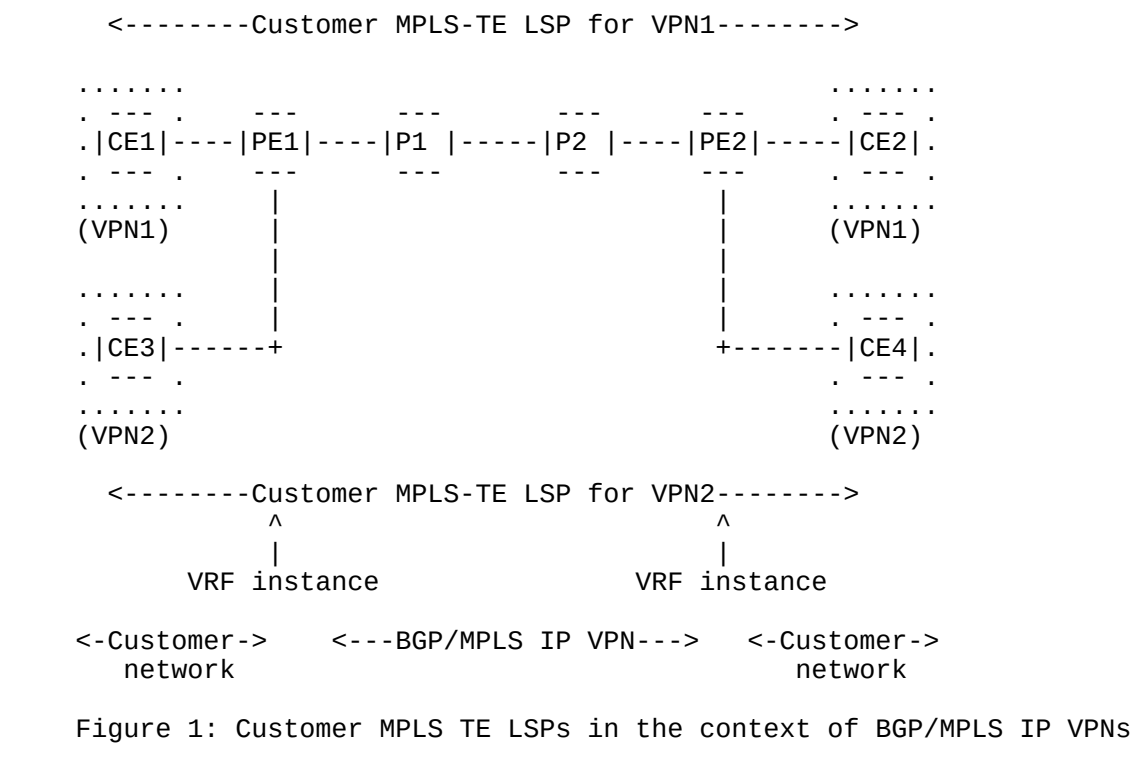

Figure 1: Customer MPLS TE LSPs in the context of BGP/MPLS IP VPNs

Consider that customers in VPN1 and VPN2 would like to establish customer MPLS-TE LSPs between their sites (i.e., between CE1 and CE2, and between CE3 and CE4). In this situation, the following RSVP-TE Path messages would be sent:

1. CE1 would send a Path message to PE1 to establish the MPLS-TE LSP (VPN1) between CE1 and CE2.

2. CE3 would also send a Path message to PE1 to establish the MPLS-TE LSP (VPN2) between CE1 and CE2.

After receiving each Path message, PE1 can identify the customer context for each Path message from the incoming interface over which the message was received. PE1 forwards the messages to PE2 using the routing mechanisms described in [RFC4364] and [RFC4659].

When the Path messages are received at PE2, that node needs to distinguish the messages and determine which applies to VPN1 and which to VPN2 so that the right forwarding state can be established and so that the messages can be passed on to the correct CE. Although the messages arrive at PE2 with an MPLS label that identifies the VPN, the messages are delivered to the RSVP-TE component on PE2, and the context of the core VPN LSP (i.e., the label) is lost. Some RSVP-TE protocol mechanism is therefore needed to embed the VPN identifier within the RSVP-TE message.

Similarly, Resv messages sent from PE2 to PE1 need an RSVP-TE mechanism to assign them to the correct VPN.

3. Protocol Extensions and Procedures

This section defines the additional RSVP-TE objects to meet the requirements described in Section 2. These objects are new variants of the SESSION, SENDER_TEMPLATE, and FILTERSPEC objects. They act as identifiers and allow PEs to distinguish Path/Resv messages per VPN in the context of BGP/MPLS IP VPNs. Section 3.1 defines the new object types, and Section 3.2 defines the specific procedures for handling RSVP messages.

3.1. Object Definitions

This experiment will be carried out using the following private Class Types. This document identifies these Class Types as "C-Type = EXPn".

```
Class = SESSION, LSP_TUNNEL_VPN-IPv4 C-Type = EXP1
Class = SESSION, LSP_TUNNEL_VPN-IPv6 C-Type = EXP2
Class = SENDER_TEMPLATE, LSP_TUNNEL_VPN-IPv4 C-Type = EXP3
Class = SENDER_TEMPLATE, LSP_TUNNEL_VPN-IPv6 C-Type = EXP4
Class = FILTER SPECIFICATION, LSP_TUNNEL_VPN-IPv4 C-Type = EXP5
Class = FILTER SPECIFICATION, LSP_TUNNEL_VPN-IPv6 C-Type = EXP6
```

3.1.1. LSP_TUNNEL_VPN-IPv4 and LSP_TUNNEL_VPN-IPv6 SESSION Object

The LSP_TUNNEL_VPN-IPv4 (or LSP_TUNNEL_VPN-IPv6) SESSION object appears in RSVP-TE messages that ordinarily contain a SESSION object and that are sent between the ingress PE and egress PE in either direction. This object MUST NOT be included in any RSVP-TE message that is sent outside of the provider's backbone.

The LSP_TUNNEL_VPN-IPv6 SESSION object is analogous to the LSP_TUNNEL_VPN-IPv4 SESSION object, using a VPN-IPv6 address ([RFC4659]) instead of a VPN-IPv4 address ([RFC4364]).

Experimenters MUST ensure that there is no conflict between the private Class Types used for this experiment and other Class Types used by the PEs.

The formats of the SESSION objects are as follows:

Class = SESSION, LSP_TUNNEL_VPN-IPv4 C-Type = EXP1

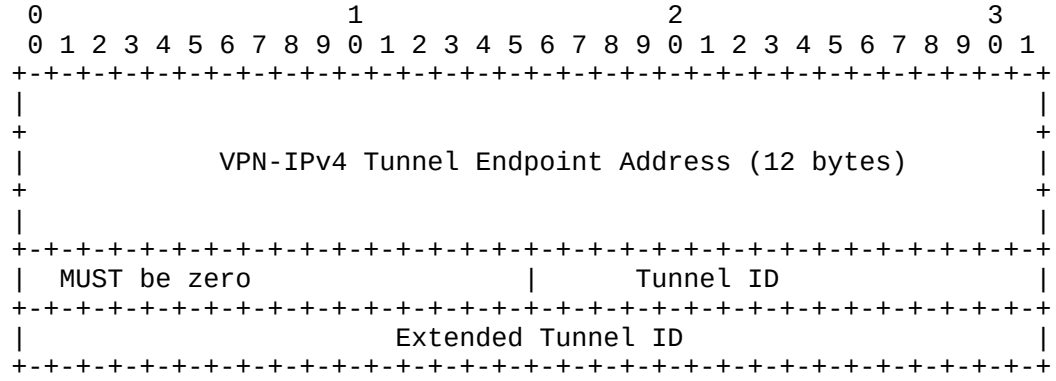

Class = SESSION, LSP_TUNNEL_VPN-IPv6 C-Type = EXP2

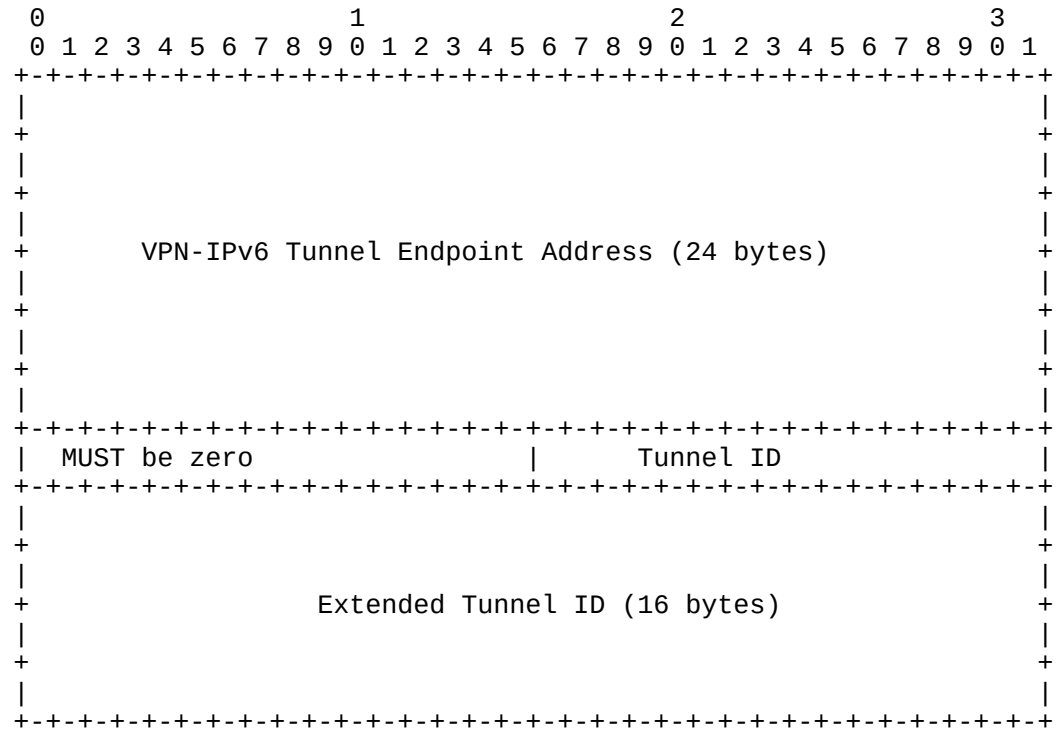

The VPN-IPv4 or VPN-IPv6 tunnel endpoint address field contains an address of the VPN-IPv4 or VPN-IPv6 address family encoded as specified in [RFC4364] or [RFC4659], respectively.

The Tunnel ID and Extended Tunnel ID are identical to the same fields in the LSP_TUNNEL_IPv4 and LSP_TUNNEL_IPv6 SESSION objects as per [RFC3209].

3.1.2. LSP_TUNNEL_VPN-IPv4 and LSP_TUNNEL_VPN-IPv6 SENDER_TEMPLATE Objects

The LSP_TUNNEL_VPN-IPv4 (or LSP_TUNNEL_VPN-IPv6) SENDER_TEMPLATE object appears in RSVP-TE messages that ordinarily contain a SENDER_TEMPLATE object and that are sent between ingress PE and egress PE in either direction, such as Path, PathError, and PathTear messages. The object MUST NOT be included in any RSVP-TE messages that are sent outside of the provider’s backbone.

The format of the object is as follows:

Class = SENDER_TEMPLATE, LSP_TUNNEL_VPN-IPv4 C-Type = EXP3

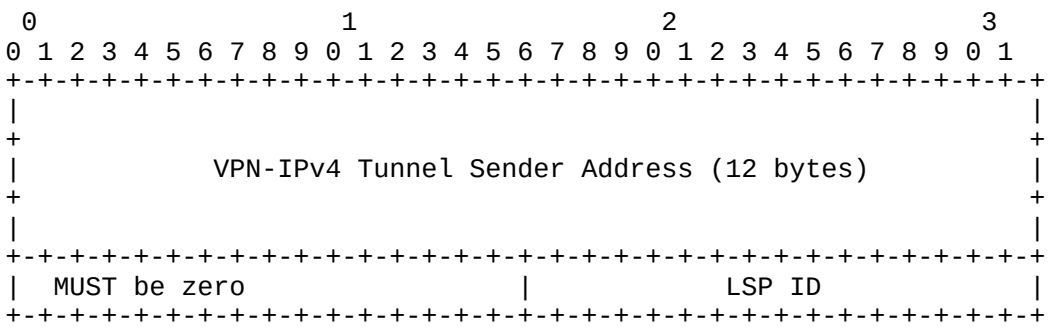

Class = SENDER_TEMPLATE, LSP_TUNNEL_VPN-IPv6 C-Type = EXP4

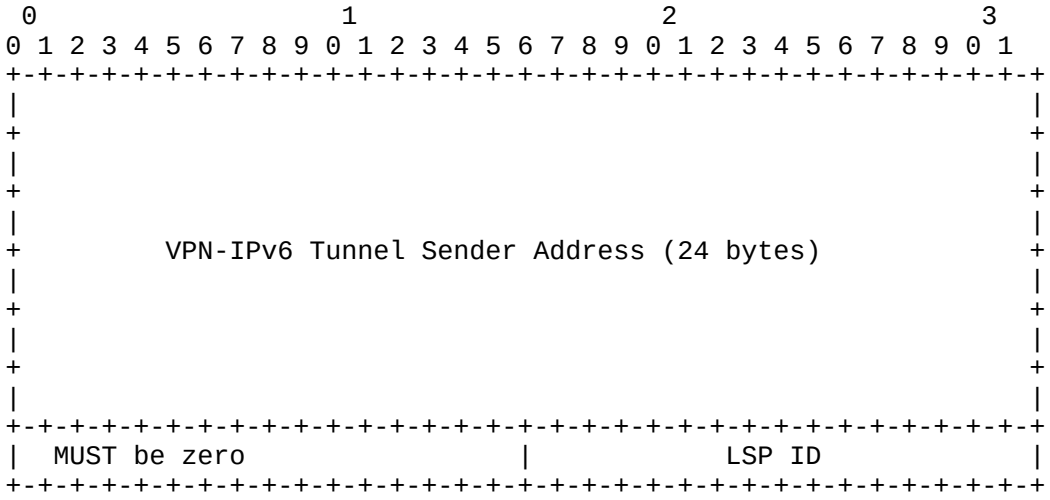

The VPN-IPv4 or VPN-IPv6 tunnel sender address field contains an address of the VPN-IPv4 or VPN-IPv6 address family encoded as specified in [RFC4364] or [RFC4659], respectively.

The LSP ID is identical to the LSP ID field in the LSP_TUNNEL_IPv4 and LSP_TUNNEL_IPv6 SENDER_TEMPLATE objects as per [RFC3209].

3.1.3. LSP_TUNNEL_VPN-IPv4 and LSP_TUNNEL_VPN-IPv6 FILTER_SPEC Objects

The LSP_TUNNEL_VPN-IPv4 (or LSP_TUNNEL_VPN-IPv6) FILTER_SPEC object appears in RSVP-TE messages that ordinarily contain a FILTER_SPEC object and that are sent between ingress PE and egress PE in either direction, such as Resv, ResvError, and ResvTear messages. The object MUST NOT be included in any RSVP-TE messages that are sent outside of the provider's backbone.

Class = FILTER SPECIFICATION, LSP_TUNNEL_VPN-IPv4 C-Type = EXP5

The format of the LSP_TUNNEL_VPN-IPv4 FILTER_SPEC object is identical to the LSP_TUNNEL_VPN-IPv4 SENDER_TEMPLATE object.

Class = FILTER SPECIFICATION, LSP_TUNNEL_VPN-IPv6 C-Type = EXP6

The format of the LSP_TUNNEL_VPN-IPv6 FILTER_SPEC object is identical to the LSP_TUNNEL_VPN-IPv6 SENDER_TEMPLATE object.

3.1.4. VPN-IPv4 and VPN-IPv6 RSVP_HOP Objects

The formats of the VPN-IPv4 and VPN-IPv6 RSVP_HOP objects are identical to the RSVP_HOP objects described in [RFC6016].

3.2. Handling the Messages

This section describes how the RSVP-TE messages are handled. Handling of these messages assumes that, in the context of BGP/MPLS IP VPNs, the ingress and egress PEs have RSVP-TE capabilities.

3.2.1. Path Message Processing at the Ingress PE

When a Path message arrives at the ingress PE (PE1 in Figure 1), the PE needs to establish suitable Path state and forward the Path message on to the egress PE (PE2 in Figure 1). Below, we describe the message handling process at the ingress PE.

1. CE1 sends a Path message to PE1 to establish the MPLS-TE LSP (VPN1) between CE1 and CE2. The Path message is addressed to the eventual destination (the receiver at the remote customer site) and carries the IP Router Alert option, in accordance with [RFC2205]. The ingress PE must recognize the router alert, intercept these messages, and process them as RSVP-TE signaling messages.

2. When the ingress PE receives a Path message from a CE that is addressed to the receiver, the VRF that is associated with the incoming interface can be identified. (This step does not deviate from current behavior.)

3. The tunnel endpoint address of the receiver is looked up in the appropriate VRF, and the BGP next hop for that tunnel endpoint address is identified. The next hop is the egress PE.

4. A new LSP_TUNNEL_VPN-IPv4/VPN-IPv6 SESSION object is constructed, containing the Route Distinguisher (RD) that is part of the VPN-IPv4/VPN-IPv6 route prefix for this tunnel endpoint address, and the IPv4/IPv6 tunnel endpoint address from the original SESSION object.

5. A new LSP_TUNNEL_VPN-IPv4/IPv6 SENDER_TEMPLATE object is constructed, with the original IPv4/IPv6 tunnel sender address from the incoming SENDER_TEMPLATE plus the RD that is used by the PE to advertise the prefix for the customers VPN.

6. A new Path message is sent containing all the objects from the original Path message, replacing the original SESSION and SENDER_TEMPLATE objects with the new LSP_TUNNEL_VPN-IPv4/VPN-IPv6 type objects. This Path message is sent directly to the egress PE (the next hop that was determined in Step 3) without the IP Router Alert option.

3.2.2. Path Message Processing at the Egress PE

Below, we describe the message handling process at the egress PE.

1. When a Path message arrives at the egress PE (PE2 in Figure 1), it is addressed to the PE itself and is handed to RSVP for processing.

2. The router extracts the RD and IPv4/IPv6 address from the LSP_TUNNEL_VPN-IPv4/VPN-IPv6 SESSION object and determines the local VRF context by finding a matching VPN-IPv4 prefix with the specified RD that has been advertised by this router into BGP.

3. The entire incoming RSVP message, including the VRF information, is stored as part of the Path state.

4. The egress PE can now construct a Path message that differs from the Path message it received in the following ways:

   a. Its tunnel endpoint address is the IP address extracted from the SESSION object.

   b. The SESSION and SENDER_TEMPLATE objects have been converted back to IPv4-type/IPv6-type by discarding the attached RD.

   c. The RSVP_HOP object contains the IP address of the outgoing interface of the egress PE and a Logical Interface Handle (LIH), as per normal RSVP processing.

5. The egress PE then sends the Path message towards its tunnel endpoint address over the interface identified in Step 4c. This Path message carries the IP Router Alert option, as required by [RFC2205].

3.2.3. Resv Processing at the Egress PE

When a receiver at the customer site originates a Resv message for the session, normal RSVP procedures apply until the Resv, making its way back towards the sender, arrives at the "egress" PE (it is the egress with respect to the direction of data flow, i.e., PE2 in Figure 1). Upon arriving at PE2, the SESSION and FILTER_SPEC objects in the Resv message, and the VRF in which the Resv was received, are used to find the matching Path state that was stored previously.

The PE constructs a Resv message to send to the RSVP HOP stored in the Path state, i.e., the ingress PE (PE1 in Figure 1). The LSP TUNNEL IPv4/IPv6 SESSION object is replaced with the same LSP_TUNNEL_VPN-IPv4/VPN-IPv6 SESSION object received in the Path message. The LSP TUNNEL IPv4/IPv6 FILTER_SPEC object is replaced with a LSP_TUNNEL_VPN-IPv4/VPN-IPv6 FILTER_SPEC object, which copies the VPN-IPv4/VPN-IPv6 address from the LSP TUNNEL SENDER_TEMPLATE received in the matching Path message.

The Resv message MUST be addressed to the IP address contained within the RSVP_HOP object in the Path message.

3.2.4. Resv Processing at the Ingress PE

When the ingress PE receives a Resv message (the ingress with respect to data flow, i.e., PE1 in Figure 1), the PE determines the local VRF context and associated Path state for this Resv message by decoding the received SESSION and FILTER_SPEC objects. It is now possible to generate a Resv message to send to the appropriate CE. The Resv

message sent to the ingress CE contains the LSP TUNNEL IPv4/IPv6 SESSION and LSP TUNNEL FILTER_SPEC objects, which are derived from the appropriate Path state.

3.2.5. Other RSVP Messages

Processing of other RSVP messages (i.e., PathError, PathTear, ResvError, ResvTear, and ResvConf) generally follows the rules defined in [RFC2205]. The following additional rules MUST be observed for messages transmitted within the VPN, i.e., between the PEs:

- The SESSION, SENDER_TEMPLATE, and FILTER_SPEC objects MUST be converted from LSP_TUNNEL_IPv4/LSP_TUNNEL_IPv6 [RFC3209] to LSP_TUNNEL_VPN-IPv4/LSP_TUNNEL_VPN-IPv6 form, respectively, and back again, in the same manner as described above for Path and Resv messages.

- The appropriate type of RSVP_HOP object (VPN-IPv4 or VPN-IPv6) MUST be used, as described in Section 8.4 of [RFC6016].

- Depending on the type of RSVP_HOP object received from the neighbor, the message MUST be MPLS encapsulated or IP encapsulated.

- The matching state and VRF MUST be determined by decoding the corresponding RD and IPv4 or IPv6 address in the SESSION and FILTER_SPEC objects.

- The message MUST be directly addressed to the appropriate PE, without using the Router Alert Option.

4. Management Considerations

MPLS-TE-based BGP/MPLS IP VPNs are based on a peer model. If an operator would like to configure a new site to an existing VPN, configuration of both the CE router and the attached PE router is required. The operator is not required to modify the configuration of PE routers connected to other sites or to modify the configuration of other VPNs.

4.1. Impact on Network Operation

It is expected that the use of the extensions specified in this document will not significantly increase the level of operational traffic.

Furthermore, the additional extensions described in this document will have no impact on the operation of existing resiliency mechanisms available within MPLS-TE.

5. Security Considerations

This document defines RSVP-TE extensions for BGP/MPLS IP VPNs. The general security issues for RSVP-TE are described in [RFC3209], [RFC4364] addresses the specific security considerations of BGP/MPLS VPNs. General security considerations for MPLS are described in [RFC5920].

In order to secure the control plane, techniques such as the TCP Authentication Option (TCP-AO) [RFC5925] MAY be used authenticate BGP messages.

To ensure the integrity of an RSVP request, the RSVP Authentication mechanisms defined in [RFC2747], and updated by [RFC3097], SHOULD be used.

6. References

6.1. Normative References

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

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

6.2. Informative References

[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.

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

[RFC3097] Braden, R. and L. Zhang, "RSVP Cryptographic Authentication -- Updated Message Type Value", RFC 3097, April 2001.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, September 2006.

[RFC5824] Kumaki, K., Ed., Zhang, R., and Y. Kamite, "Requirements for Supporting Customer Resource ReSerVation Protocol (RSVP) and RSVP Traffic Engineering (RSVP-TE) over a BGP/MPLS IP-VPN", RFC 5824, April 2010.

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

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

[RFC6016] Davie, B., Le Faucheur, F., and A. Narayanan, "Support for the Resource Reservation Protocol (RSVP) in Layer 3 VPNs", RFC 6016, October 2010.

7. Acknowledgments

The authors would like to express thanks to Makoto Nakamura and Daniel King for their helpful and useful comments and feedback.

8. Contributors

Chikara Sasaki
KDDI R&D Laboratories, Inc.
2-1-15 Ohara Fujimino
Saitama 356-8502
Japan
EMail: ch-sasaki@kddilabs.jp

Daisuke Tatsumi
KDDI Corporation
2-3-2 Nishishinjuku Shinjuku-ku
Tokyo 163-8003
Japan
EMail: da-tatsumi@kddi.com

Authors' Addresses

Kenji Kumaki
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460
Japan
EMail: ke-kumaki@kddi.com

Tomoki Murai
Furukawa Network Solution Corp.
5-1-9, Higashi-Yawata, Hiratsuka
Kanagawa 254-0016
Japan
EMail: murai@fnsc.co.jp

Dean Cheng
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
USA
EMail: dean.cheng@huawei.com

Satoru Matsushima
Softbank Telecom
1-9-1,Higashi-Shimbashi,Minato-Ku
Tokyo 105-7322
Japan
EMail: satoru.matsushima@g.softbank.co.jp

Peng Jiang
KDDI Corporation
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460
Japan
EMail: pe-jiang@kddi.com