Network Working Group
Internet Draft
Intended Status: Experimental
Expires: June 20, 2013

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

December 21, 2012

Support for RSVP-TE in L3VPNs
draft-kumaki-murai-l3vpn-rsvp-te-09.txt

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 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

Table of Contents

Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

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 Path messages and Resv messages described in [RFC3209], to be appropriately handled 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 labelled. Similarly, when a message is received by a PE having been sent across the core, both labels to 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 VRF. This may be achieved easily in some implementations, in others it is not so easy to achieved.

This document defines experimental formats and mechanisms that follows 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.

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 appropriately handled 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 supported by the two PEs shown in the figure. The customer sites operate MPLS-TE LSPs.

Here, we make the following set of assumptions.

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

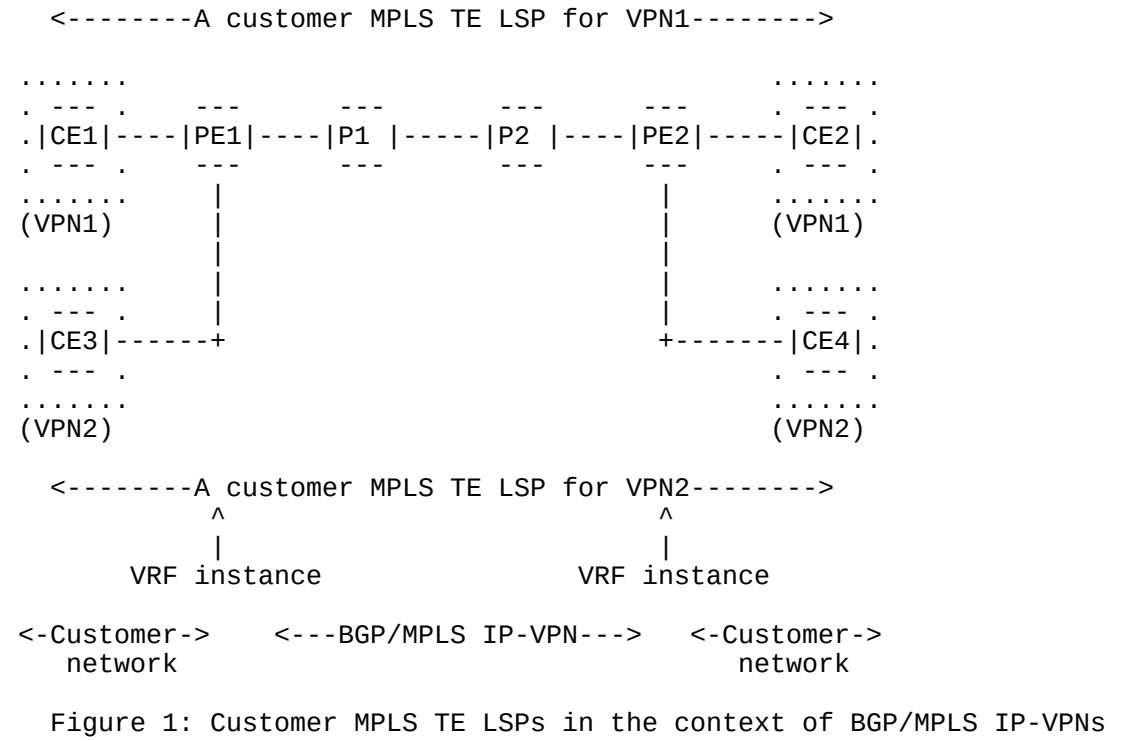

Figure 1: Customer MPLS TE LSPs in the context of BGP/MPLS IP-VPNs

Consider that customers in VPN1 and VPN2 would like to establish a 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 messages, 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 will arrive at PE2 with an MPLS label that identifies the VPN, the messages will be delivered to the RSVP-TE component on PE2 and the context of the core VPN LSP (i.e., the label) will be 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 mechanisms to assign them to the correct VPN.

3. Protocol Extensions and Procedures

This section provides the additional RSVP-TE objects to meet the requirements described in Section 2. These are new variants of the SESSION, SENDER_TEMPLATE and FILTERSPEC objects. These new objects will act as identifiers and allow PEs to The new object types are defined in Section 3.1, and the specific procedure is described in Section 3.2.

3.1 Object Definitions

3.1.1 LSP_TUNNEL_VPN-IPv4 and LSP_TUNNEL_VPN-IPv6 SESSION Object

The LSP_TUNNEL_VPN-IPv4 (or VPN-IPv6) SESSION object appears in RSVP-TE messages that ordinarily contain a SESSION object and are sent between ingress PE and egress PE in either direction. The 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]).

This experimentation will be carried out using private Class Types. These can be identified in this document 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

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 objects are as follows:

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

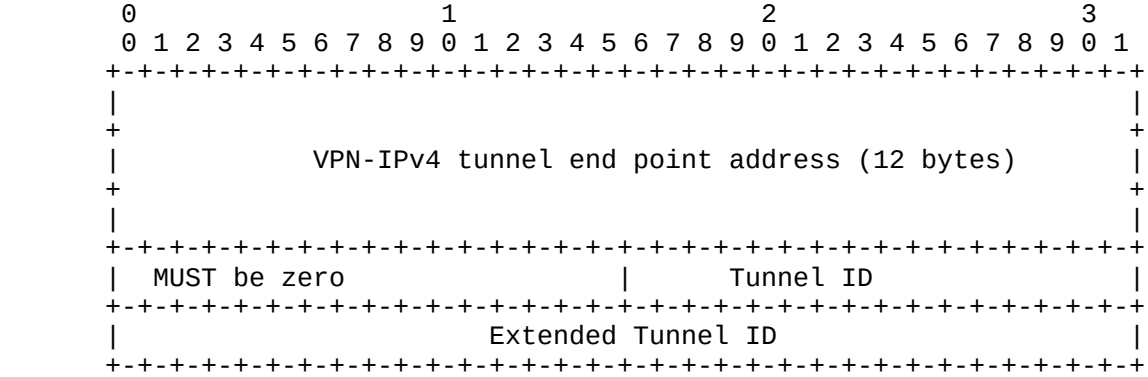

Class = SESSION, LSP_TUNNEL_VPN-IPv6 C-Type = EXP2

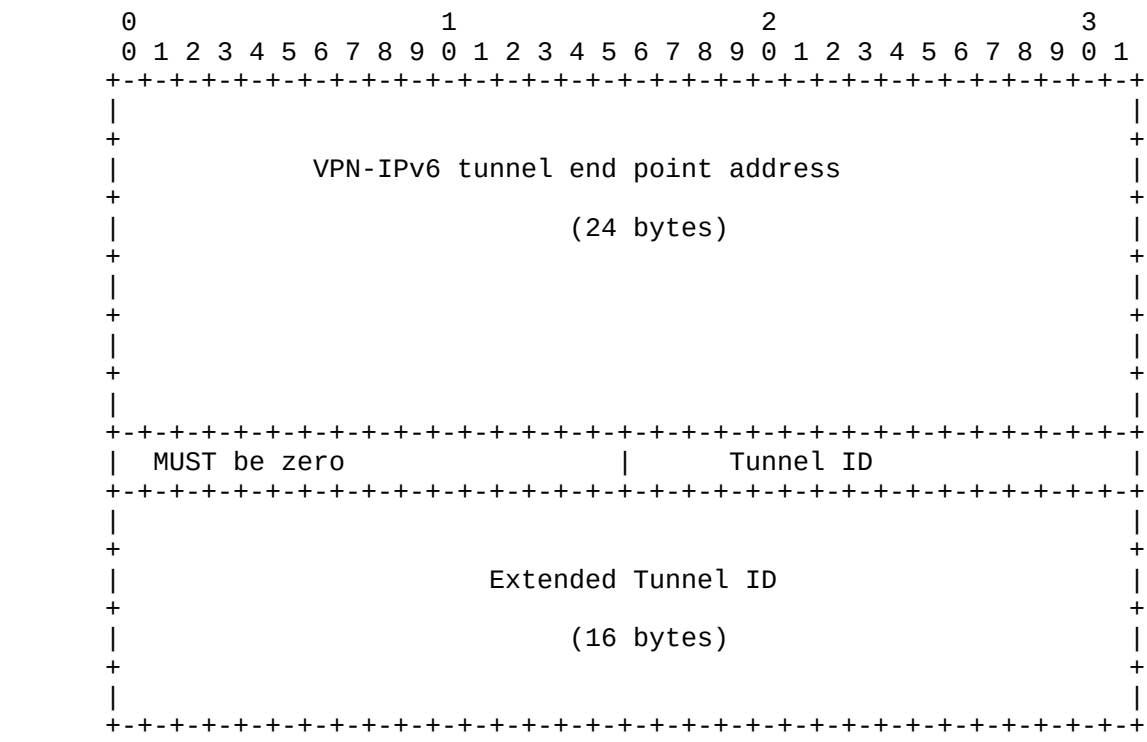

The VPN-IPv4 tunnel end point address (respectively, VPN-IPv6 tunnel end point address) field contains an address of the VPN-IPv4 (respectively, VPN-IPv6) address family encoded as specified in [RFC4364](respectively, [RFC4659]).

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 VPN-IPv6) SENDER_TEMPLATE object appears in RSVP-TE messages that ordinarily contain a SENDER_TEMPLATE object and are sent between ingress PE and egress PE in either direction (such as Path, PathError, and PathTear). 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

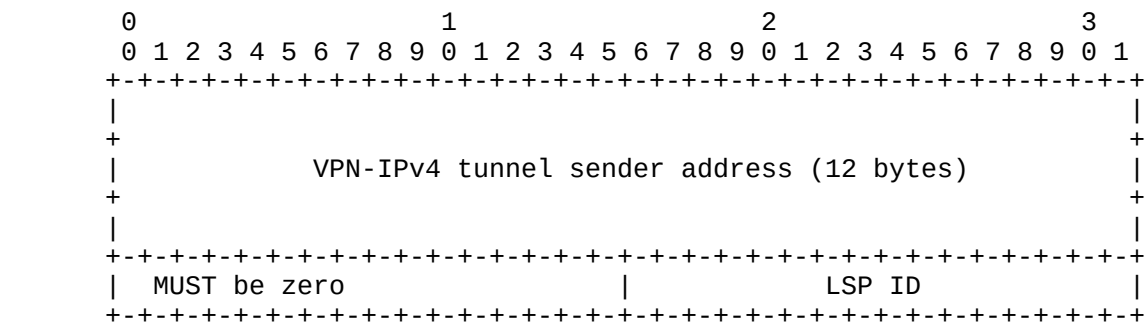

Class = SENDER_TEMPLATE, LSP_TUNNEL_VPN-IPv6 C-Type = EXP4

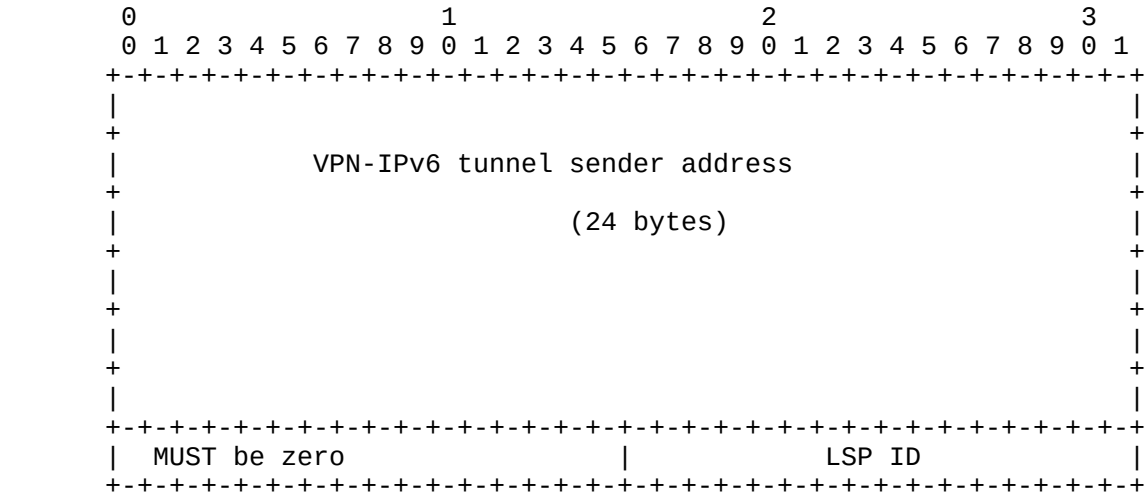

The VPN-IPv4 tunnel sender address (respectively, VPN-IPv6 tunnel sender address) field contains an address of the VPN-IPv4 (respectively, VPN-IPv6) address family encoded as specified in [RFC4364] (respectively, [RFC4659]).

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 VPN-IPv6) FILTER_SPEC object appears in RSVP-TE messages that ordinarily contain a FILTER_SPEC object and are sent between ingress PE and egress PE in either direction (such as Resv, ResvError, and ResvTear). 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 format of the VPN-IPv4 and VPN-IPv6 RSVP_HOP objects are identical to objects described in [RFC6016].

3.2 Handling

It assumes that ingress PEs and egress PEs in the context of BGP/MPLS IP-VPNs have RSVP-TE capabilities.

3.2.1 Path Message Processing at 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). In this section we described the message handling process at the ingress PE.

1. CE1 would send 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 signalling 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 end point address of the receiver is looked up in the appropriate VRF, and the BGP Next-Hop for that tunnel end point 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 end point address, and the IPv4/IPv6 tunnel end point 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 as being looked up in step 3 above) without IP Router Alert.

3.2.2 Path Message Processing at Egress PE

In this section we described 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 which differs from the Path message it received in the following ways:

   a. Its tunnel end point address is the IP address extracted from the SESSION object;

   b. The SESSION and SENDER_TEMPLATE objects are 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 an Logical Interface Handle (LIH), as per normal RSVP processing.

5. The egress PE then sends the Path message on towards its tunnel end point address over the interface identified above. This Path message carries the IP Router-Alert option as required by [RFC2205].

3.2.3 Resv Processing at 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 "egress" with respect to the direction of data flow, i.e. PE2 in figure 1). On arriving at PE2, the SESSION and FILTER_SPEC objects in the Resv, and the VRF in which the Resv was received, are used to find the matching Path state 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. 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 Ingress PE

Upon receiving a Resv message at the ingress PE (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 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 will contain LSP TUNNEL IPv4/IPv6 SESSION and LSP TUNNEL FILTER_SPEC objects, 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 message in general follows the rules defined in [RFC2205], with additional rules that MUST be observed for messages transmitted within the VPN, i.e., between the PEs as follows:

- 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 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 (respectively, 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 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 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], update by [RFC3097], SHOULD be used.

6. IANA Considerations

This document makes no request for IANA actions.

7. References

7.1 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 Swallow, G., "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

7.2 Informative References

[RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and Jamin, S., "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., Zhang, R. and Kamite, Y., "Requirements for supporting Customer RSVP and RSVP-TE over a BGP/MPLS IP-VPN", RFC 5824, April 2010.

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

[RFC5925] J. Touch, et. al., "The TCP Authentication Option", RFC5925, June 2010.

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

8. Acknowledgments

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

9. Author's Addresses

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

Tomoki Murai
Furukawa Network Solutions 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, U.S.A.
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

10. Contributors' Addresses

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
Garden Air Tower
Iidabashi, Chiyoda-ku,
Tokyo 102-8460, JAPAN
Email: da-tatsumi@kddi.com