Network Working Group
Internet Draft
Category: Informational
Created: February 18, 2008
Expires: August 18, 2008

K. Kumaki, Ed.
KDDI R&D Labs
R. Zhang
BT
Y. Kamite
NTT Communications

Requirements for supporting Customer RSVP and RSVP-TE over a BGP/MPLS IP-VPN

draft-kumaki-l3vpn-e2e-rsvp-te-reqts-06.txt

Status of this Memo

Copyright Notice

Abstract

Some service providers want to build a service which guarantees QoS or bandwidth from a local CE to a remote CE through the network. Today, customers expect to run triple play services through BGP/MPLS IP-VPNs. As a result, their requirements for end-to-end QoS of applications are increasing. Depending on the application (e.g., voice, video, bandwidth-guaranteed data pipe, etc.), an end-to-end

native RSVP path and/or an end-to-end MPLS TE LSP are required, and they need to meet some constraints.
This document describes service provider requirements for supporting customer RSVP and RSVP-TE over a BGP/MPLS VPN.

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].

Table of Contents

1. Introduction

Some service providers want to build a service which guarantees QoS or bandwidth from a local CE to a remote CE through the network. A CE could be broadened to include network client equipment owned and operated by the service provider. However, the CE is not part of the MPLS provider network.
Today, customers expect to run triple play services through BGP/MPLS IP-VPNs [RFC4364]. As a result, their requirements for end-to-end QoS of applications are increasing. Depending on the application (e.g., voice, video, bandwidth-guaranteed data pipe, etc.), an end-to-end native RSVP path and/or an end-to-end MPLS TE LSP are required, and they need to meet some constraints. For example, an RSVP path may be used to provide for bandwidth and QoS guarantees. An end-to-end MPLS TE LSP may be used to guarantee bandwidth, and provide for MPLS fast reroute (FRR) [RFC4090] around node and link failure. It should be noted that an RSVP session between two CEs may also be mapped and tunneled into a TE LSP across an MPLS provider network in a most likely scenario.

If service providers offer the above services in BGP/MPLS IP-VPNs, they can have the following two advantages.

The first advantage is for customers to receive these network services while being able to use both private addresses and global addresses as they desire.

The second advantage is for service providers to offer these network services while protecting confidentiality from customers. Customers join a Virtual Routing and Forwarding (VRF) instance and cannot forward packets through the service provider's global forwarding instance, nor can they join the service provider's intra-domain routing.

This document defines a reference model, application scenarios and detailed requirements for supporting customer RSVP and RSVP-TE over a BGP/MPLS IP-VPN.

Also, specification for this solution itself is out of scope in this document.

2. Terminology

LSP: Label Switched Path

TE LSP: Traffic Engineering Label Switched Path

MPLS TE LSP: Multi Protocol Label Switching TE LSP

C-RSVP path: Customer RSVP path: a native RSVP path with bandwidth reservation of X for customers

C-TE LSP: Customer Traffic Engineering Label Switched Path: an end-to-end MPLS TE LSP for customers

P-TE LSP: Provider Traffic Engineering Label Switched Path: a transport TE LSP between two PEs

VPN: Virtual Private Network

CE: Customer Edge Equipment

PE: Provider Edge Equipment that has direct connections to CEs from the Layer3 point of view.

P: Provider Equipment that has backbone trunk connections only.

VRF: Virtual Private Network (VPN) Routing and Forwarding Instance

PCC: Path Computation Client: any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element: an entity (component, application or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

Head-end LSR: ingress LSR

Tail-end LSR: egress LSR

LSR: Label Switched Router

3. Problem Statement

Some service providers think that they can offer advanced services using RSVP or RSVP-TE over BGP/MPLS IP-VPNs. In addition, in many cases, BGP/MPLS IP-VPNs can be used within the service provider network to carry network services. For example, a C-RSVP path with bandwidth reservation of X can be used to transport voice. In order

to achieve sub-50msec recovery during link/node/SRLG failure and to provide strict QoS guarantees, a C-TE LSP with bandwidth X between data centers or customer sites can be used to carry voice and video traffic. Thus, service providers or customers can choose a C-RSVP path or a C-TE LSP to meet their requirements. Please note that there

When service providers offer a C-RSVP path between hosts or CEs over BGP/MPLS IP-VPNs, the CE/host requests an end-to-end C-RSVP path with bandwidth reservation of X to the remote CE/host. However, if a C-RSVP signaling is to send within VPN, the service provider network will face scalability issues. Therefore, in order to solve scalability issues, multiple C-RSVP reservations can be aggregated at PE, where a P-TE LSP head-end can perform admission control using the aggregated C-RSVP reservations. The method that is described in RFC4804 can be considered as a useful approach. In this case, a reservation request from within the context of a VRF can get aggregated onto a P-TE LSP. The P-TE LSP can be pre-established, resized based on the request, or triggered by the request. Service providers, however, can not provide a C-RSVP path over vrf instance as defined in RFC4364. The current BGP/MPLS IP-VPN architecture also does not support an RSVP instance running in the context of a vrf to process RSVP messages and integrated services (int-serv) [RFC1633][RFC2210]. One of solutions is described in [RSVP-L3VPN].

If service providers offer a C-TE LSP from CE to CE over BGP/MPLS IP-VPN, they require that a MPLS TE LSP from a local CE to a remote CE be established. However, if a C-TE LSP signaling is to send within VPN, the service provider network will face scalability issues. Therefore, in order to solve scalability issues, multiple C-TE LSPs can be aggregated at PE, where a P-TE LSP head-end can perform admission control using hierarchical methods. Furthermore, if service providers provide the C-TE LSP over a BGP/MPLS IP-VPN, they can not provide it over vrf instance as defined in RFC4364. The current BGP/MPLS IP-VPN architecture does not support an RSVP-TE instance running in the context of a vrf to process RSVP messages and trigger the establishment of the C-TE LSP over the service provider core network. If every C-TE LSP is to trigger the establishment or resizing of a P-TE LSP, the service provider network will also face scalability issues that arise from maintaining a large number of P-TE LSPs and/or dynamic signaling of these P-TE LSPs.

Thus, in the models of C-RSVP paths and C-TE LSPs both, the solution must address these scalability concerns.

Two different models are described above. The differences between C-RSVP paths and C-TE LSPs are as follows:

- C-RSVP path model: data packets among CEs are forwarded by "native

IP packets" (i.e. not labeled packets).

- C-TE LSP model: data packets among CEs are forwarded by "labeled IP packets".

The following items are mainly required to support C-RSVP paths and C-TE LSPs over BGP/MPLS IP-VPNs. Detailed requirements are described in sections 6, 7 and 8.

- C-RSVP path QoS guarantees.
- Fast recovery over BGP/MPLS IP-VPN to protect traffic for C-TE LSP against CE-PE link failure and PE node failure.
- Strict C-TE LSP bandwidth and QoS guarantees.
- Resource optimization for C-RSVP paths and C-TE LSPs.
- Scalability for C-RSVP paths and C-TE LSPs.

4. Reference Model

In this section, a C-RSVP path, a C-TE LSP and a P-TE LSP are explained.

4.1 End-to-End C-RSVP Path Model

A C-RSVP path and a P-TE LSP are shown in figure 1 in the context of a BGP/MPLS IP-VPN. A P-TE LSP may be a non-TE LSP (i.e., LDP) in some cases. In some cases, however, it may be difficult to guarantee end-to-end QoS. (e.g. If a P-TE LSP has enough bandwidth in service provider backbone, a C-RSVP path can reserve a bandwidth.)

CE0/CE1 requests an e2e C-RSVP path to CE3/CE2 with bandwidth reservation of X. At PE1, this reservation request received in the context of a VRF will get aggregated onto a pre-established P-TE LSP, or trigger the establishment of a new P-TE LSP. It should be noted that C-RSVP sessions across different BGP/MPLS IP-VPNs can be aggregated onto the same P-TE LSP between the same PE pair, achieving further scalability.

The RSVP control messages (e.g. an RSVP PATH message and an RSVP RESV message) exchanged among CEs are forwarded by IP packets through BGP/MPLS IP-VPN. After CE0 and/or CE1 receive a reservation message from CE2 and/or CE3, CE0/CE1 establishes a C-RSVP path through the BGP/MPLS IP-VPN.

A P-TE LSP is established between PE1 and PE2. This LSP is used by the vrf instance to forward customer packets within BGP/MPLS IP-VPN.

Generally speaking, C-RSVP paths are used by customers and P-TE LSPs are used by service providers.

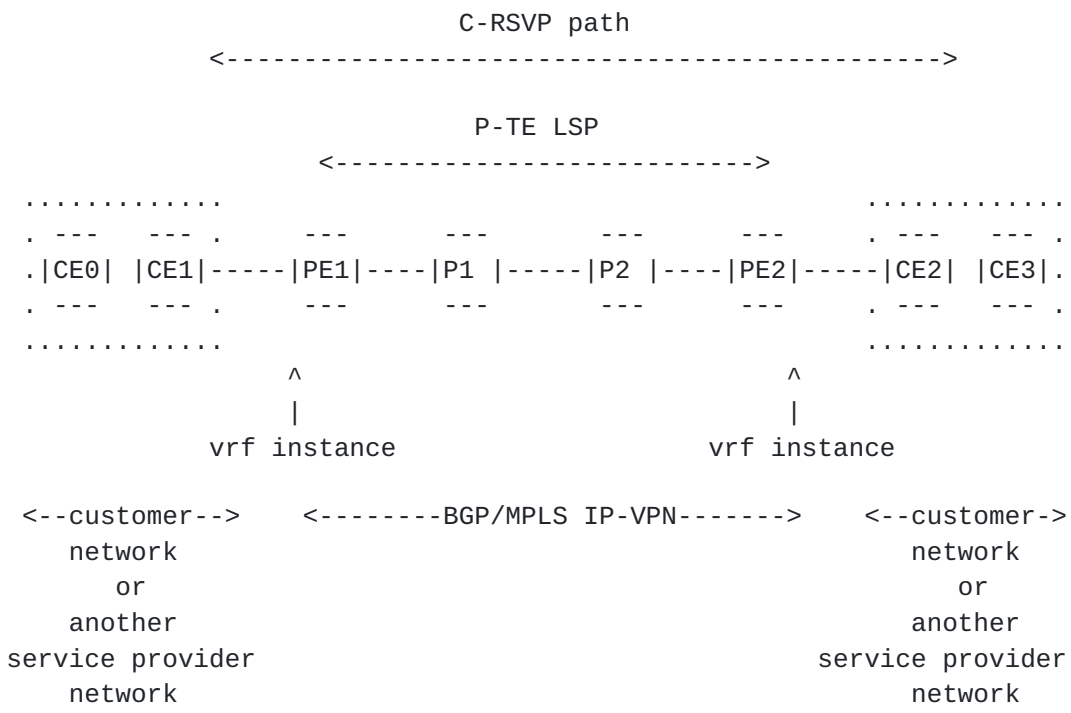

Figure 1 e2e C-RSVP path model

4.2 End-to-End C-TE LSP Model

A C-TE LSP and a P-TE LSP are shown in figure 2 in the context of a BGP/MPLS IP-VPN. A P-TE LSP may be a non-TE LSP (i.e., LDP) in some cases. As described in previous sub-section, it may be difficult to guarantee end-to-end QoS in some cases.

CE0/CE1 requests an e2e TE LSP path to CE3/CE2 with bandwidth reservation of X. At PE1, this reservation request received in the context of a VRF will get aggregated onto a pre-established P-TE LSP, or trigger the establishment of a new P-TE LSP. It should be noted that C-TE LSPs across different BGP/MPLS IP-VPNs can be aggregated onto the same P-TE LSP between the same PE pair, achieving further scalability.

The RSVP-TE control messages (e.g. a RSVP PATH message and a RSVP RESV message) exchanged among CEs are forwarded by labeled packet through BGP/MPLS IP-VPN. After CE0 and/or CE1 receive a reservation message from CE2 and/or CE3, CE0/CE1 establishes a C-TE LSP through the BGP/MPLS IP-VPN.

A P-TE LSP is established between PE1 and PE2. This LSP is used by the vrf instance to forward customer packets within BGP/MPLS IP-VPN.

Generally speaking, C-TE LSPs are used by customers and P-TE LSPs are used by service providers.

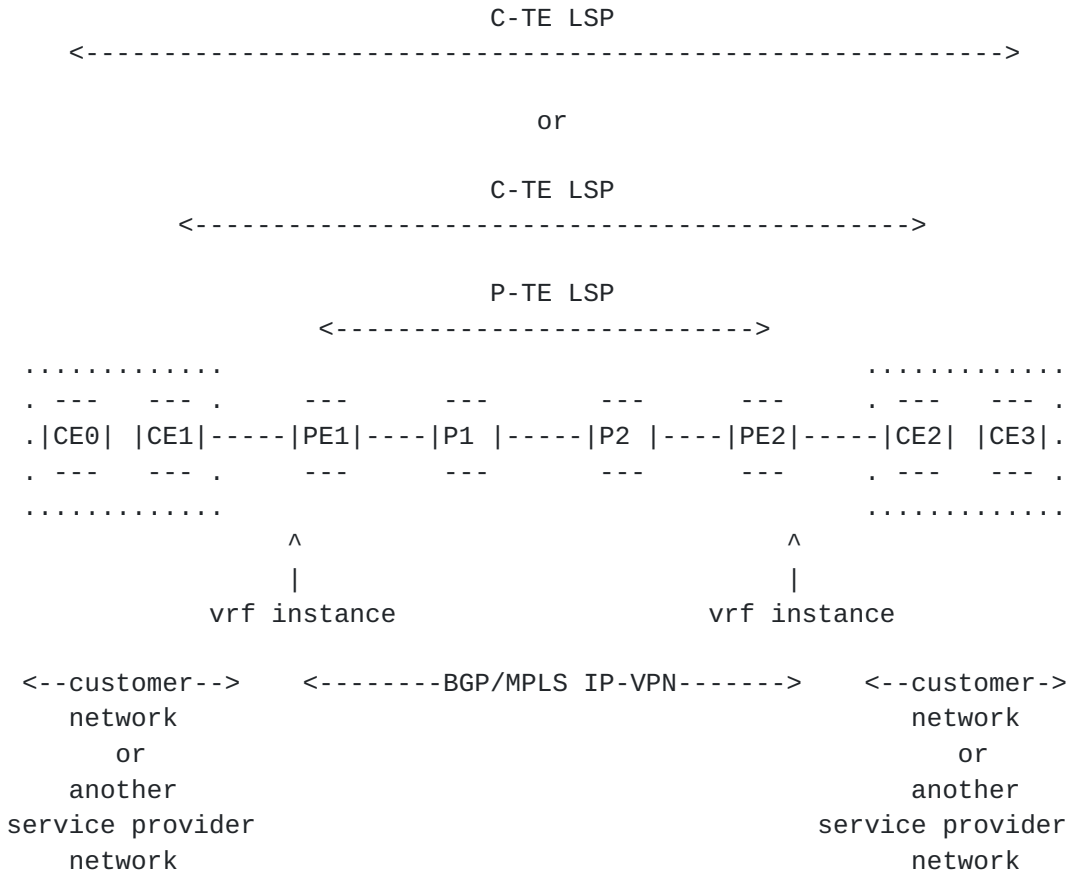

Figure 2 e2e C-TE LSP model

5. Application Scenarios

The following sections present a few application scenarios for C-RSVP paths and C-TE LSPs in BGP/MPLS IP-VPN environments.

5.1 Scenario I: Fast Recovery over BGP/MPLS IP-VPN

In this scenario, as shown in figure 3, a customer uses a VoIP application between its sites (i.e., between CE1 and CE2). H0 and H1 are voice equipment.
In this case, the customer establishes C-TE LSP1 as a primary path and C-TE LSP2 as a backup path. If the link between PE1 and CE1 or the node PE1 fails, C-TE LSP1 needs C-TE LSP2 as a path protection.

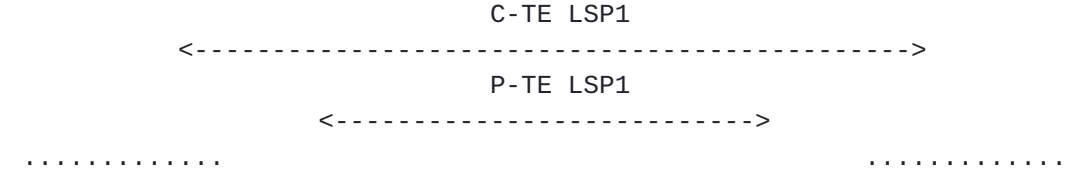

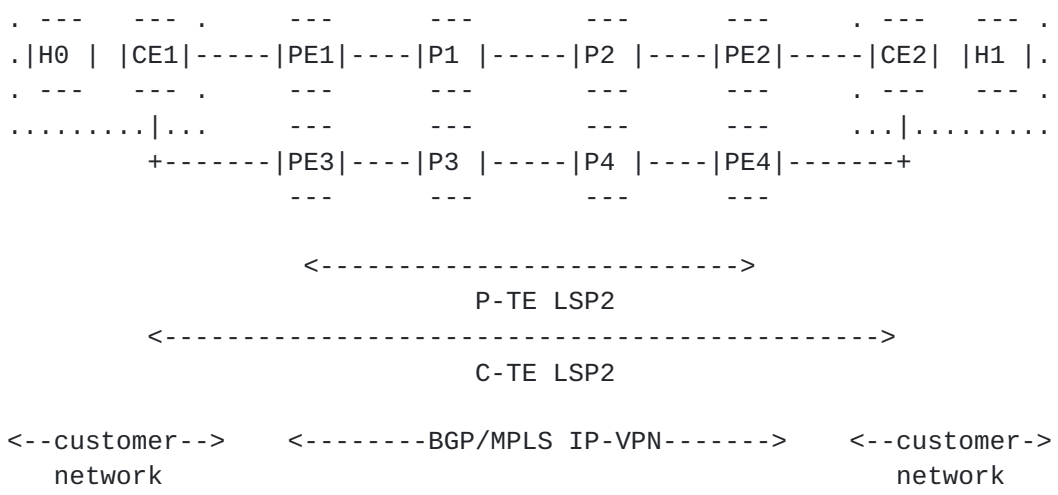

Figure 3 Scenario I

5.2 Scenario II: Strict C-TE LSP QoS Guarantees

In this scenario, as shown in figure 4, a service provider B transports voice and video traffic between its sites (i.e., between CE1 and CE2).
In this case, service provider B establishes C-TE LSP1 with preemption priority 0 and bandwidth 100Mbps for voice traffic, and C-TE LSP2 with preemption priority 1 and bandwidth 200Mbps for unicast video traffic. On the other hand, service provider A also pre-establishes P-TE LSP1 with preemption priority 0 and bandwidth 1Gbps for voice traffic, and P-TE LSP2 with preemption priority 1 and bandwidth 2Gbps for video traffic. These P-TE LSP1 and P-TE LSP2 should support DS-TE. [RFC4124]

PE1 and PE3 should choose an appropriate P-TE LSP based on preemption priority. In this case, C-TE LSP1 must be associated with P-TE LSP1 at PE1 and C-TE LSP2 must be associated with P-TE LSP2 at PE3.

Furthermore, PE1 and PE3 head-ends should control the bandwidth of C-TE LSPs. In this case, PE1 and PE3 can choose C-TE LSPs by the amount of max available bandwidth for each P-TE LSP, respectively.

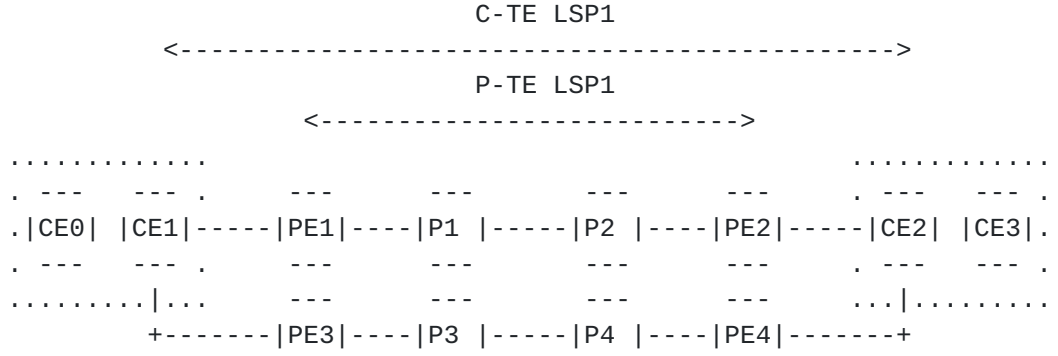

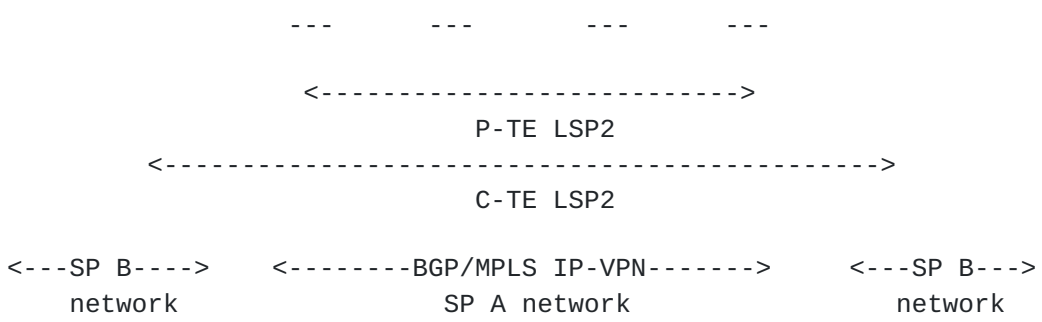

Figure 4 Scenario II

5.3 Scenario III: Load Balance of CE-to-CE Traffic

In this scenario, as shown in figure 5, service provider C uses voice and video traffic between its sites (i.e., between CE0 and CE5/CE7, between CE2 and CE5/CE7, between CE5 and CE0/CE2, and between CE7 and CE0/CE2). H0 and H1 are voice and video equipment.

In this case, service provider C establishes C-TE LSP1, C-TE LSP3, C-TE LSP5 and C-TE LSP7 with preemption priority 0 and bandwidth 100Mbps for voice traffic, and establishes C-TE LSP2, C-TE LSP4, C-TE LSP6 and C-TE LSP8 with preemption priority 1 and bandwidth 200Mbps for video traffic. On the other hand, service provider A also pre-establishes P-TE LSP1 and P-TE LSP3 with preemption priority 0 and bandwidth 1Gbps for voice traffic, and P-TE LSP2 and P-TE LSP4 with preemption priority 1 and bandwidth 2Gbps for video traffic. These P-TE LSP1, P-TE LSP2, P-TE LSP3 and P-TE LSP4 should support DS-TE. [RFC4124]

All PEs should choose an appropriate P-TE LSP based on preemption priority. To minimize the traffic disruption due to a single network failure, diversely routed C-TE LSPs are established. In this case, FRR [RFC4090] is not necessarily required.

Also, unconstrained TE LSPs (i.e., C-TE LSPs/P-TE LSPs with 0 bandwidth) [ZERO-BANDWIDTH] are applicable to this scenario.

Furthermore, load balancing for a communication between H0 and H1 can be done by setting up full mesh C-TE LSPs between CE0/CE2 and CE5/CE7.

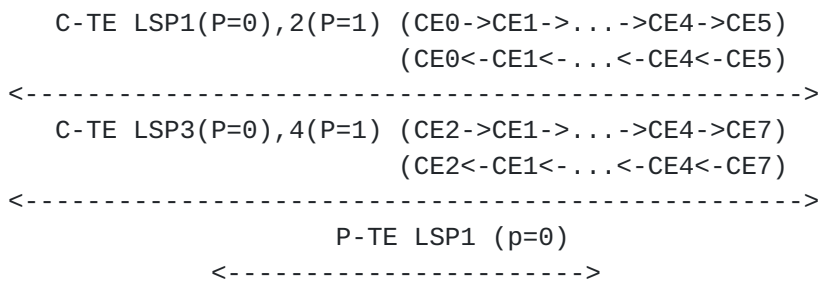

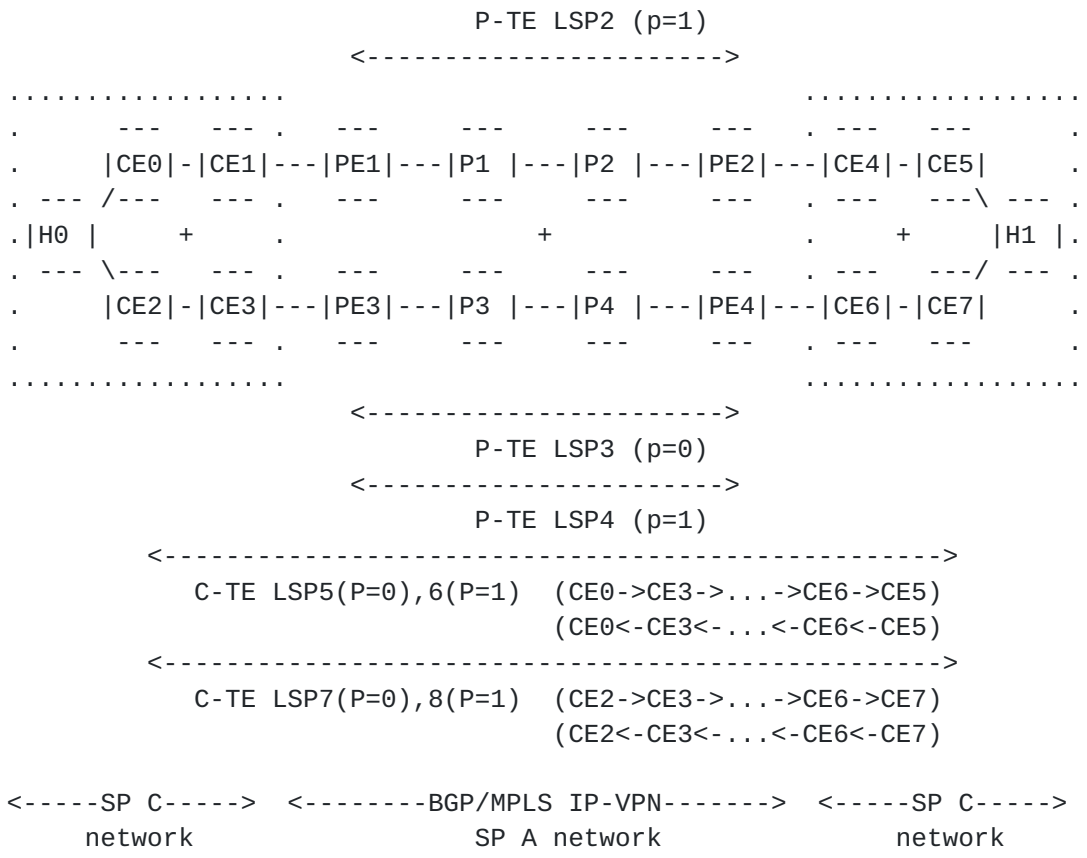

Figure 5 Scenario III

5.4 Scenario IV: RSVP Aggregation over MPLS TE Tunnels

In this scenario, as shown in figure 6, the customer has two hosts connecting off CE1 and CE2 respectively. CE1 and CE2 are connected to PE1 and PE2, respectively, within a VRF instance belonging to the same VPN. The requesting host (H1) may request to H2 an RSVP path with bandwidth reservation of X. This reservation request from within the context of VRF will get aggregated onto a pre-established P-TE/DS-TE LSP based upon procedures similar to [RFC4804]. As in the case of [RFC4804], there may be multiple P-TE LSPs belonging to different DS-TE class-types. Local policies can be implemented to map the incoming RSVP path request from H1 to the P-TE LSP with the appropriate class-type. Please note that the e2e RSVP path request may also be initiated by the CE devices themselves.

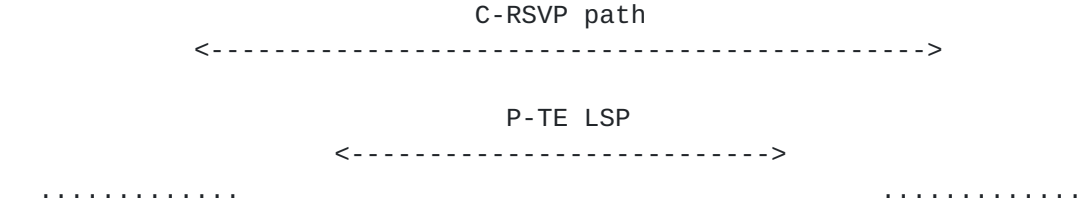

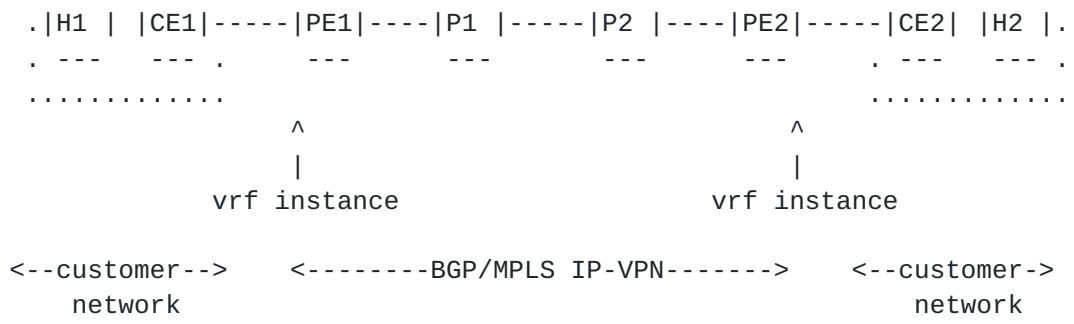

Figure 6 Scenario IV

5.5 Scenario V: RSVP over Non-TE LSP

In this scenario, as shown in figure 7, the customer has two hosts connecting off CE1 and CE2 respectively. CE1 and CE2 are connected to PE1 and PE2, respectively, within a VRF instance belonging to the same VPN. The requesting host (H1) may request to H2 an RSVP path with bandwidth reservation of X. In this case, a non-TE LSP (i.e. LDP etc) is provided between PEs and supports MPLS diffserv [RFC3270]. Local policies can be implemented to map customer's reserved flow to the LSP with the appropriate EXP at PE1. Please note that there is always enough bandwidth available in service provider backbone.

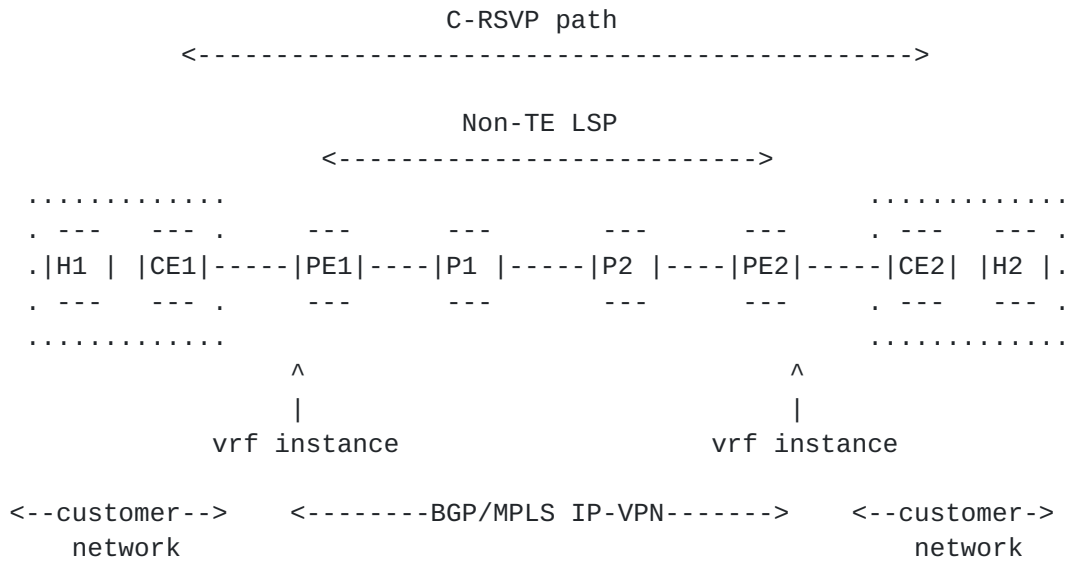

Figure 7 Scenario V

6. Detailed Requirements for C-TE LSPs Model

This section describes detailed requirements for C-TE LSPs in BGP/MPLS IP-VPN environments.

6.1 Selective P-TE LSPs

The solution MUST provide the ability to decide which P-TE LSP a PE uses for a C-RSVP path and a C-TE LSP. When a PE receives a native RSVP and/or a path messages from a CE, it may be able to decide which P-TE LSP it uses. In this case, various kinds of P-TE LSPs exist in service provider network. For example, the PE MUST choose an appropriate P-TE LSP based on local policies such as:

1. preemption priority
2. affinity
3. class-type
4. on the data plane: (DSCP or EXP bits)

6.2 Graceful Restart Support for C-TE LSPs

The solution SHOULD provide graceful restart capability for a C-TE LSP over vrf instance. Graceful restart mechanisms related to this architecture are described in [RFC3473] [RFC3623] [RFC4781].

6.3 Rerouting Support for C-TE LSPs

The solution MUST provide rerouting of a C-TE LSP in case of link/node/SRLG failures or preemption. Such rerouting may be controlled by a CE or by a PE depending on the failure. Rerouting capability MUST be provided against a CE-PE link failure or a PE failure if another is available between the head-end and the tail-end of the C-TE LSP.

6.4 FRR Support for C-TE LSPs

The solution MUST support FRR [RFC4090] features for a C-TE LSP over vrf instance.
In BGP/MPLS IP-VPN environments, a C-TE LSP from a CE traverses multiple PEs and Ps, albeit tunneled over a P-TE LSP. In order to avoid PE-CE link/PE node/SRLG failures, a CE (a customer's head-end router) needs to support a fast local protection or a fast path protection.

The following protection MUST be supported:

1. CE link protection
2. PE node protection
3. CE node protection (supposed that there are one or more C-TE nodes at customer sites)

6.5 Admission Control Support on P-TE LSP Head-Ends

The solution MUST support admission control on a P-TE LSP tunnel head-end. C-TE LSPs may potentially reserve over the bandwidth of a

P-TE LSP. The P-TE LSP tunnel head-end SHOULD control the number of C-TE LSPs and/or the bandwidth of C-TE LSPs.
For example, the transport TE LSP head-end MUST have a configurable limit on the maximum number of C-TE LSPs that it can admit from a CE.
As for the amount of bandwidth that can be reserved by C-TE LSPs: there could be two situations:
1. Let the P-TE LSP do its natural bandwidth admission
2. Set a cap on the amount of bandwidth and have the configuration option to:
a. Reserve the minimum of the cap bandwidth or the C-TE LSP bandwidth on the P-TE LSP if the required bandwidth is available
b. Reject the C-TE LSP if the required bandwidth by the C-TE LSP is not available

6.6 Policy Control Support for C-TE LSPs

The solution MUST support policy control for a C-TE LSP at a PE.
The PE MUST be able to perform the following:
1. Limit the rate of RSVP messages per CE link
2. Accept or reject requests for a given affifinity
3. Accept or reject requests with a specified setup and/or pre-emption priorities.
4. Accept or reject requests for fast reroutes
5. Neglect the requested setup and/or pre-emption priorities and select a P-TE LSP based on a local policy that applies to the CE-PE link or VRF.
6. Neglect the requested affinity and select a P-TE LSP based on a local policy that applies to the CE-PE link or VRF.
7. Perform mapping in data plane between customer exp bits and transport P-TE LSP exp bits.

6.7 PCE Features Support for C-TE LSPs

The solution MAY support PCE architecture for a C-TE LSP establishment in the context of a vrf instance. When a C-TE LSP is provided, CEs, PEs and Ps may support PCE [RFC4655] [PCEP] features. In this case, CE routers or PE routers may be PCCs and PE routers and/or P routers may be PCEs.

6.8 Diversely Routed C-TE LSPs Support

The solution MUST provide for setting up diversely routed C-TE LSPs over vrf instance. These diverse C-TE LSPs MAY be traversing over two different P-TE LSPs that are fully disjoint within a service provider network. When a single CE has multiple uplinks which connect to different PEs, it is desirable that multiple C-TE LSPs over vrf instance are established between a pair of LSRs. When two CEs have multiple uplinks which connect to different PEs, it is desirable that

multiple C-TE LSPs over vrf instance are established between two

different pairs of LSRs. In these cases, for example, the following points will be beneficial to customers.

1. load balance of CE-to-CE traffic across diverse C-TE LSPs so as to minimize the traffic disruption in case of a single network element failure
2. path protection (e.g. 1:1, 1:N)

6.9 Optimal Path Support for C-TE LSPs

The solution MUST support an optimal path for a C-TE LSP over vrf instance.
Depending on an application (e.g. voice and video), an optimal path is needed for a C-TE LSP over vrf instance. An optimal path may be a shortest path based on TE metric or IGP metric.

6.10 Reoptimization Support for C-TE LSPs

The solution MUST support reoptimization of a C-TE LSP over vrf instance. These LSPs MUST be reoptimized using make-before-break.
In this case, it is desirable for a customer's head-end LSR to be configured with regard to timer-based or event-driven reoptimization.
Furthermore, customers SHOULD be able to reoptimize a C-TE LSP manually.
To provide delay- or jitter-sensitive traffic (i.e. voice traffic), a C-TE LSP is expected to be kept optimal.

6.11 DS-TE Support for C-TE LSPs

The solution MUST support DS-TE [RFC4124] for a C-TE LSP over vrf instance.
Applications, which have different traffic characteristics, are used in BGP/MPLS IP-VPN environments. Service providers try to achieve fine-grained optimization of transmission resources, efficiency and further enhanced network performance. It may be desirable to perform TE at a per-class level.
By mapping the traffic from a given diff-serv class of service on a separate C-TE LSP, it allows this traffic to utilize resources available to the given class on both shortest paths and non-shortest paths, and follow paths that meet TE constraints which are specific to the given class.

7. Detailed Requirements for C-RSVP Paths Model

This section describes detailed requirements for C-RSVP paths in BGP/MPLS IP-VPN environments.

7.1 Admission Control between PE-CE for C-RSVP Paths

The solution MUST support admission control at ingress/egress PE. PEs MUST control RSVP messages per a vrf.

7.2 Aggregation of C-RSVP Paths by P-TE LSPs

The solution SHOULD support C-RSVP paths aggregated by P-TE LSPs. P-TE LSPs SHOULD be pre-established by manually or dynamically, MAY be established triggered by C-RSVP message. Also, P-TE LSP SHOULD support DS-TE.

7.3 Non-TE LSPs support for C-RSVP Paths

The solution MUST support non-TE LSPs (i.e. LDP-based LSP, etc). They are provided between PEs and supports MPLS diffserv [RFC3270]. Local policies can be implemented to map customer's reserved flow to the LSP with the appropriate EXP at PE.
Please note that there is always enough bandwidth available in service provider backbone.

7.4 Transparency of C-RSVP Paths

The solution SHOULD NOT change RSVP messages from local CE to remote CE (Path, Resv, Path Error, Resv Error, etc). Customers SHOULD deal RSVP messages transparently between CE sites.

8. Common Detailed Requirements for Two Models

This section describes common detailed requirements for C-TE LSPs and C-RSVP paths in BGP/MPLS IP-VPN environments.

8.1 CE-PE Routing

The solution MUST support the following routing configuration on the CE-PE links with either RSVP or RSVP-TE on the CE-PE link:

1. static routing
2. BGP routing
3. OSPF
4. OSPF-TE (RSVP-TE case only)

8.2 Complexity and Risks

The solution SHOULD NOT introduce unnecessary complexity to the current operating network to such a degree that it would affect the stability and diminish the benefits of deploying such a solution over SP networks.

8.3 Backward Compatibility

The deployment of C-RSVP paths and C-TE LSPs SHOULD NOT impact existing RSVP and MPLS TE mechanisms respectively, but allow for a smooth migration or co-existence.

8.4 Scalability Considerations

The solution MUST have a minimum impact on network scalability from a C-RSVP path and a C-TE LSP over vrf instance.
Scalability of C-RSVP paths and C-TE LSPs MUST address the following consideration.

1. RSVP (e.g. number of RSVP messages, retained state etc).
2. RSVP-TE (e.g. number of RSVP control messages, retained state, message size etc).
3. BGP (e.g. number of routes, flaps, overloads events etc).

8.5 Performance Considerations

The solution SHOULD be evaluated with regard to the following criteria.

1. Degree of path optimality of the C-TE LSP.
2. TE LSP setup time.
3. Failure and restoration time.
4. Impact and scalability of the control plane due to added overheads and so on.
5. Impact and scalability of the data/forwarding plane due to added overheads and so on.

8.6 Management Considerations

Manageability of C-RSVP paths and C-TE LSPs MUST addresses the following considerations.

1. Need for a MIB module for control plane and monitoring.
2. Need for diagnostic tools.

MIB module for C-RSVP paths and C-TE LSPs MUST collect per a vrf instance.
If a CE is managed by service providers, MIB information for C-RSVP paths and C-TE LSPs from the CE MUST be collected per a customer.

Today, diagnostic tools can detect failures of control plane and data plane for general MPLS TE LSPs [RFC4379].
The diagnostic tools MUST detect failures of control and data plane for C-TE LSPs over a vrf instance.

MPLS OAM for C-TE LSPs MUST be supported within the context of VRF except for the above.

In BGP/MPLS IP-VPN environments, from a CE point of view, IP TTL decreases at a local PE and a remote PE. But from a PE point of view, both IP TTL and MPLS TTL decreases between PEs.

9. Security Considerations

Security issues for C-TE LSPs relate to both control plane and data plane.

In terms of control plane, in the models of C-RSVP paths and C-TE LSPs both, a PE receives IPv4 or IPv6 RSVP control packets from a CE. If the CE is an untrusted router for service providers, the PE MUST be able to limit IPv4 or IPv6 RSVP control packets. If the CE is a trusted router for service providers, the PE MAY be able to limit IPv4 or IPv6 control packets.

In terms of data plane, in the model of C-TE LSPs, a PE receives labeled IPv4 or IPv6 data packets from a CE. If the CE is an untrusted router for service providers, the PE MUST be able to limit labeled IPv4 or IPv6 data packets. If the CE is a trusted router for service providers, the PE MAY be able to limit labeled IPv4 or IPv6 data packets. Specifically, the PE must drop MPLS-labeled packets if the MPLS label was not assigned over the PE-CE link on which the packet was received. The PE must also be able to police traffic to the traffic profile associated with the LSP on which traffic is received on the PE-CE link.

Moreover, flooding RSVP/RSVP-TE control packets from malicious customers enables other customers to impact themselves on their communication. Therefore, a PE MUST isolate the impact of such customer's packets from other customers.

In BGP/MPLS IP-VPN environments, from a CE point of view, IP TTL should decrease at a local PE and a remote PE to hide service provider network topology.

10. IANA Considerations

This requirement document makes no requests for IANA action.

11. References

11.1 Normative References

[RFC1633] Braden, R., et al., "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.

[RFC3270] Le Faucheur, F., "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.

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

[RFC3623] Moy, J., et al., "Graceful OSPF Restart", RFC3623, November 2003.

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

[RFC4124] Le Faucheur, F., "Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering", RFC 4124, June 2005.

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

[RFC4379] Kompella, K. and G. Swallow, "Detecting MPLS Data Plane Failures", RFC 4379, February 2006.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "Path Computation Element (PCE) Architecture", RFC 4655, August 2006.

[RFC4781] Rekhter, Y., and Aggarwal, R., "Graceful Restart Mechanism for BGP with MPLS", RFC 4781, January 2007.

11.2 Informative References

[RSVP-L3VPN] Davie, B., et al., "Support for RSVP in Layer 3 VPNs", Work in Progress, June 2007.

[PCEP] Vasseur, J.-P., et al., "Path Computation Element(PCE) communication Protocol (PCEP) - Version 1", Work in Progress, February 2007.

[RFC4804] Le Faucheur, F., et al., "Aggregation of RSVP Reservations over MPLS TE/DS-TE Tunnels", RFC4804, February 2007.

[ZERO-BANDWIDTH] Vasseur, J.-P., et al., "A Link-Type sub-TLV to convey the number of Traffic Engineering Label Switched Paths signaled with zero reserved bandwidth across a link", Work in Progress, February 2008.

12. Acknowledgments

The author would like to express the thanks to Nabil Bitar for his helpful and useful comments and feedback.

13. Author's Addresses

Kenji Kumaki
KDDI R&D Laboratories, Inc.
2-1-15 Ohara Fujimino
Saitama 356-8502, JAPAN
Email: ke-kumaki@kddi.com

Raymond Zhang
BT Infonet
2160 E. Grand Ave.
El Segundo, CA 90025
Email: raymond.zhang@bt.infonet.com

Yuji Kamite
NTT Communications Corporation
Tokyo Opera City Tower
3-20-2 Nishi Shinjuku, Shinjuku-ku
Tokyo 163-1421, Japan
Email: y.kamite@ntt.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).