

Internet Engineering Task Force (IETF)
Request for Comments: 5824
Category: Informational
ISSN: 2070-1721

K. Kumaki, Ed.
KDDI Corporation
R. Zhang
BT
Y. Kamite
NTT Communications Corporation
April 2010

Requirements for Supporting
Customer Resource ReSerVation Protocol (RSVP)
and RSVP Traffic Engineering (RSVP-TE) over a BGP/MPLS IP-VPN

Abstract

Today, customers expect to run triple-play services through BGP/MPLS IP-VPNs. Some service providers will deploy services that request Quality of Service (QoS) guarantees from a local Customer Edge (CE) to a remote CE across the network. As a result, the application (e.g., voice, video, bandwidth-guaranteed data pipe, etc.) requirements for an end-to-end QoS and reserving an adequate bandwidth continue to increase.

Service providers can use both an MPLS and an MPLS Traffic Engineering (MPLS-TE) Label Switched Path (LSP) to meet their service objectives. This document describes service-provider requirements for supporting a customer Resource ReSerVation Protocol (RSVP) and RSVP-TE over a BGP/MPLS IP-VPN.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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/rfc5824>.

Copyright Notice

Copyright (c) 2010 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 (<http://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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
2. Requirements Language	4
3. Terminology	5
4. Problem Statement	5
5. Application Scenarios	7
5.1. Scenario I: Fast Recovery over BGP/MPLS IP-VPNs	8
5.2. Scenario II: Strict C-TE LSP QoS Guarantees	8
5.3. Scenario III: Load Balance of CE-to-CE Traffic	9
5.4. Scenario IV: RSVP Aggregation over MPLS-TE Tunnels	11
5.5. Scenario V: RSVP over Non-TE LSPs	12
5.6. Scenario VI: RSVP-TE over Non-TE LSPs	13
6. Detailed Requirements for C-TE LSP Model	14
6.1. Selective P-TE LSPs	14
6.2. Graceful Restart Support for C-TE LSPs	14
6.3. Rerouting Support for C-TE LSPs	15
6.4. FRR Support for C-TE LSPs	15
6.5. Admission Control Support on P-TE LSP Head-Ends	15
6.6. Admission Control Support for C-TE LSPs in LDP-Based Core Networks	16
6.7. Policy Control Support for C-TE LSPs	16
6.8. PCE Features Support for C-TE LSPs	16
6.9. Diversely Routed C-TE LSP Support	16
6.10. Optimal Path Support for C-TE LSPs	17
6.11. Reoptimization Support for C-TE LSPs	17
6.12. DS-TE Support for C-TE LSPs	17
7. Detailed Requirements for C-RSVP Path Model	18
7.1. Admission Control between PE-CE for C-RSVP Paths	18
7.2. Aggregation of C-RSVP Paths by P-TE LSPs	18
7.3. Non-TE LSP Support for C-RSVP Paths	18
7.4. Transparency of C-RSVP Paths	18
8. Commonly Detailed Requirements for Two Models	18
8.1. CE-PE Routing	18
8.2. Complexity and Risks	19
8.3. Backward Compatibility	19
8.4. Scalability Considerations	19
8.5. Performance Considerations	19
8.6. Management Considerations	20
9. Security Considerations	20
10. References	21
10.1. Normative References	21
10.2. Informative References	22
Acknowledgments.....	23
Appendix A. Reference Model.....	24
A.1 End-to-End C-RSVP Path Model.....	24
A.2 End-to-End C-TE LSP Model.....	25

1. Introduction

Some service providers want to build a service that guarantees Quality of Service (QoS) and a bandwidth from a local Customer Edge (CE) to a remote CE through the network. A CE includes the network client equipment owned and operated by the service provider. However, the CE may not be part of the MPLS provider network.

Today, customers expect to run triple-play services such as Internet access, telephone, and television through BGP/MPLS IP-VPNs [RFC4364].

As these services evolve, the requirements for an end-to-end QoS to meet the application requirements also continue to grow. Depending on the application (e.g., voice, video, bandwidth-guaranteed data pipe, etc.), a native IP using an RSVP and/or an end-to-end constrained MPLS Traffic Engineering (MPLS-TE) Label Switched Path (LSP) may be required. The RSVP path may be used to provide QoS guarantees and reserve an adequate bandwidth for the data. An end-to-end MPLS-TE LSP may also be used to guarantee a bandwidth, and provide extended functionality like MPLS fast reroute (FRR) [RFC4090] for maintaining the service continuity around node and link, including the CE-PE link, failures. It should be noted that an RSVP session between two CEs may also be mapped and tunneled into an MPLS-TE LSP across an MPLS provider network.

A number of advantages exist for deploying the model previously mentioned. The first is that customers can use these network services while being able to use both private addresses and global addresses. The second advantage is that the traffic is tunneled through the service-provider backbone so that customer traffic and route confidentiality are maintained.

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

A specification for a solution is out of scope in this document.

2. Requirements Language

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].

3. Terminology

This document uses the BGP/MPLS IP-VPN terminology defined in [RFC4364] and also uses Path Computation Element (PCE) terms defined in [RFC4655].

TE LSP: Traffic Engineering Label Switched Path

MPLS-TE LSP: Multiprotocol Label Switching TE LSP

C-RSVP path: Customer RSVP path: a native RSVP path with the 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 Provider Edges (PEs)

LSR: a Label Switched Router

Head-end LSR: an ingress LSR

Tail-end LSR: an egress LSR

4. Problem Statement

Service providers want to deliver triple-play services with QoS guarantees to their customers. Various techniques are available to achieve this. Some service providers will wish to offer advanced services using an RSVP signaling for native IP flows (C-RSVP) or an RSVP-TE signaling for Customer TE LSPs (C-TE LSPs) over BGP/MPLS IP-VPNs.

The following examples outline each method:

A C-RSVP path with the bandwidth reservation of X can be used to transport voice traffic. In order to achieve recovery in under 50 ms during link, node, and Shared Risk Link Group (SRLG) failures, 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.

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 the bandwidth reservation of X to the remote CE/host. However, if a C-RSVP signaling is to send within a VPN, the service-provider

network will face scalability issues because routers need to retain the RSVP state per a customer. Therefore, in order to solve scalability issues, multiple C-RSVP reservations can be aggregated at a 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 Virtual Routing and Forwarding (VRF) instance 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, cannot provide a C-RSVP path over the 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 solution is described in [RSVP-L3VPN].

If service providers offer a C-TE LSP from a CE to a CE over the BGP/MPLS IP-VPN, they require that an MPLS-TE LSP from a local CE to a remote CE be established. However, if a C-TE LSP signaling is to send within the VPN, the service-provider network may face the following scalability issues:

- A C-TE LSP can be aggregated by a P-TE LSP at a PE (i.e., hierarchical LSPs). In this case, only a PE maintains the state of customer RSVP sessions.
- A C-TE LSP cannot be aggregated by a P-TE LSP at a PE, depending on some policies (i.e., continuous LSPs). In this case, both Ps and PEs maintain the state of customer RSVP sessions.
- A C-TE LSP can be aggregated by the non-TE LSP (i.e., LDP). In this case, only a PE maintains the state of customer RSVP-TE sessions. Note that it is assumed that there is always enough bandwidth available in the service-provider core network.

Furthermore, if service providers provide the C-TE LSP over the BGP/MPLS IP-VPN, they currently cannot provide it over the VRF instance as defined in [RFC4364]. Specifically, the current BGP/MPLS IP-VPN architecture does not support the 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

the dynamic signaling of these P-TE LSPs. Section 8.4 of this document, "Scalability Considerations", provides detailed scalability requirements.

Two different models have been 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".

Depending on the service level and the need to meet specific requirements, service providers should be able to choose P-TE LSPs or non-TE LSPs in the backbone network. The selection may be dependent on the service provider's policy and the node's capability to support the mechanisms described.

The items listed below are selectively required to support C-RSVP paths and C-TE LSPs over BGP/MPLS IP-VPNs based on the service level. For example, some service providers need all of the following items to provide a service, and some service providers need only some of them to provide the service. It depends on the service level and policy of service providers. Detailed requirements are described in Sections 6, 7, and 8.

- C-RSVP path QoS guarantees.
- Fast recovery over the BGP/MPLS IP-VPN to protect traffic for the 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.

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. Appendix A, "Reference Model", describes a C-RSVP path, a C-TE LSP, and a P-TE LSP.

In all scenarios, it is the responsibility of the service provider to ensure that enough bandwidth is available to meet the customers' application requirements.

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

In this scenario, as shown in Figure 1, a customer uses a VoIP application between its sites (i.e., between CE1 and CE2). H0 and H1 represent 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 of PE1 fails, C-TE LSP1 needs C-TE LSP2 as a path protection.

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

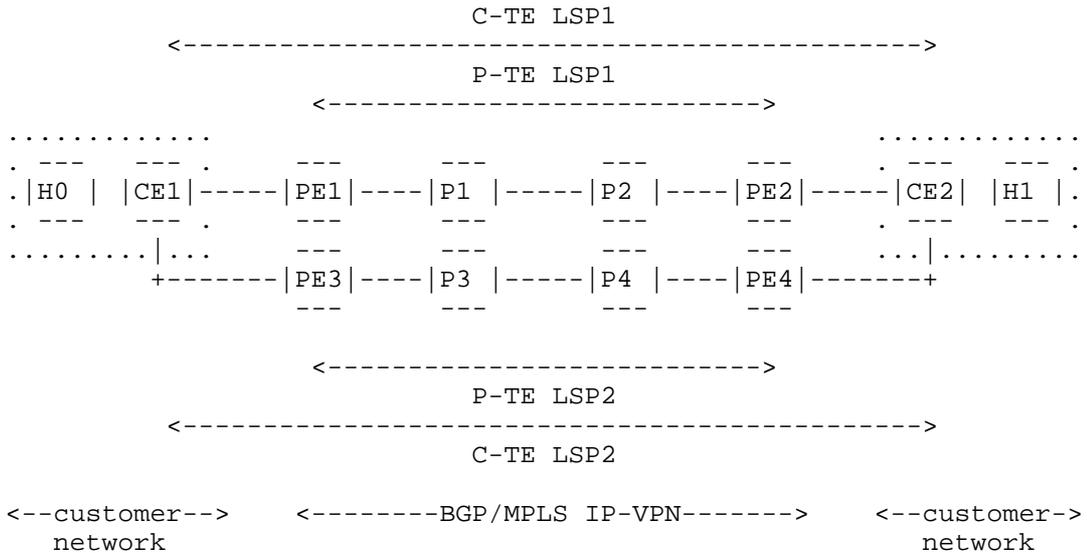


Figure 1. Scenario I

5.2. Scenario II: Strict C-TE LSP QoS Guarantees

In this scenario, as shown in Figure 2, service provider B (SP 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 100-Mbps bandwidth for the voice traffic, and C-TE LSP2 with preemption priority 1 and 200-Mbps bandwidth for the unicast video traffic. On the other hand, service provider A (SP A) also pre-establishes P-TE LSP1 with preemption priority 0 and 1-Gbps bandwidth for the voice traffic, and P-TE LSP2 with preemption priority 1 and 2-Gbps bandwidth for the video

traffic. In this scenario, P-TE LSP1 and P-TE LSP2 should support Diffserv-aware MPLS Traffic Engineering (DS-TE) [RFC4124].

PE1 and PE3 should choose an appropriate P-TE LSP based on the 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 maximum available bandwidth for each P-TE LSP, respectively.

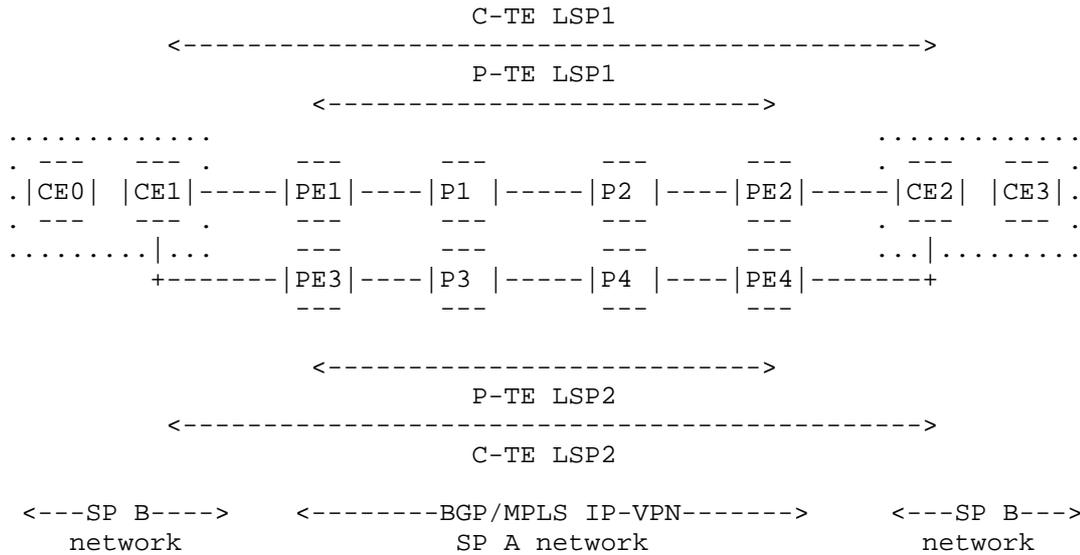


Figure 2. Scenario II

It's possible that the customer and the service provider have differing preemption priorities. In this case, the PE policy will override the customers. In the case where the service provider does not support preemption priorities, then such priorities should be ignored.

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

In this scenario, as shown in Figure 3, service provider C (SP 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 represent 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 100-Mbps bandwidth for the voice traffic, and establishes C-TE LSP2, C-TE LSP4, C-TE LSP6, and C-TE LSP8 with preemption priority 1 and 200-Mbps bandwidth for the video traffic. On the other hand, service provider A also pre-establishes P-TE LSP1 and P-TE LSP3 with preemption priority 0 and 1-Gbps bandwidth for the voice traffic, and P-TE LSP2 and P-TE LSP4 with preemption priority 1 and 2-Gbps bandwidth for the video traffic. In this scenario, 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 the preemption priority. To minimize the traffic disruption due to a single network failure, diversely routed C-TE LSPs are established. In this case, the FRR [RFC4090] is not necessarily required.

Also, unconstrained TE LSPs (i.e., C-TE LSPs/P-TE LSPs with 0 bandwidth) [RFC5330] are applicable to this scenario.

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

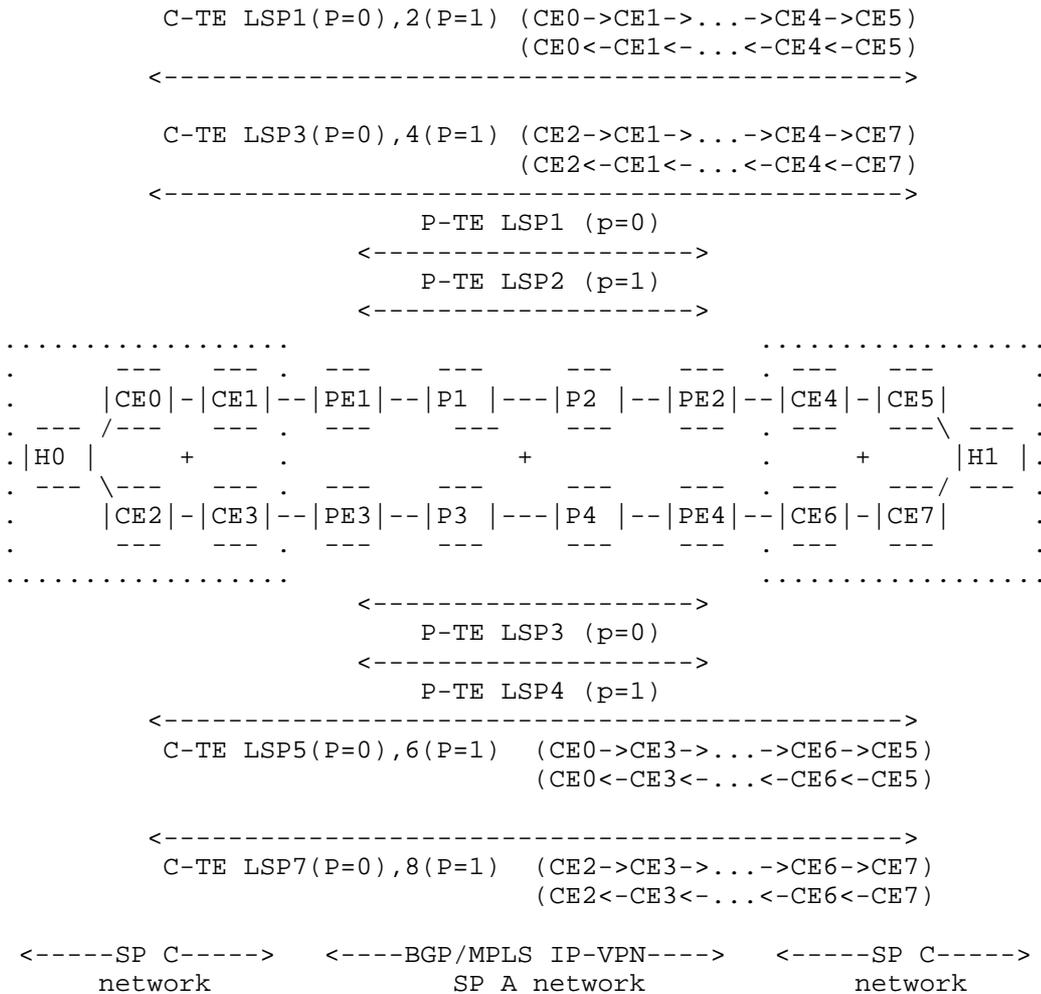


Figure 3. Scenario III

5.4. Scenario IV: RSVP Aggregation over MPLS-TE Tunnels

In this scenario, as shown in Figure 4, the customer has two hosts connecting to 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 from H2 an RSVP path with the 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 end-to-end (e2e) RSVP path request may also be initiated by the CE devices themselves.

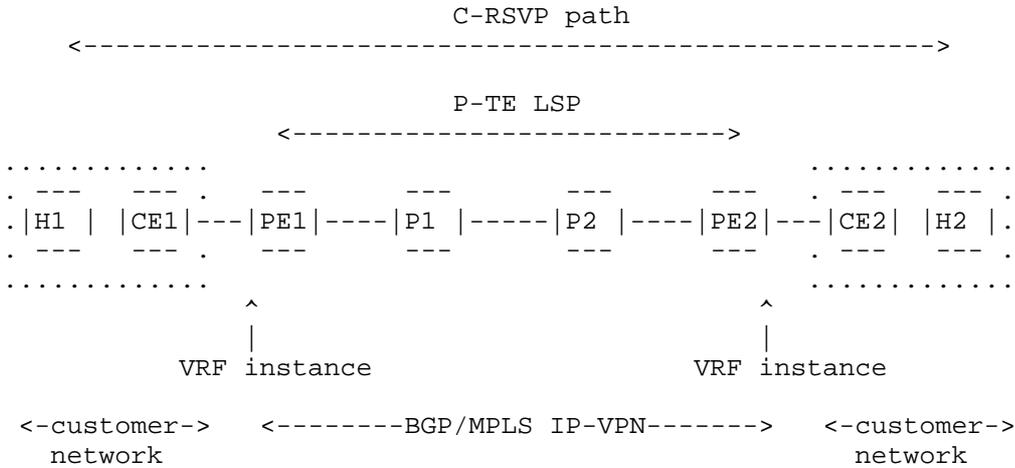


Figure 4. Scenario IV

5.5. Scenario V: RSVP over Non-TE LSPs

In this scenario, as shown in Figure 5, a customer has two hosts connecting to 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 from H2 an RSVP path with the bandwidth reservation of X. In this case, a non-TE LSP (i.e., LDP, etc.) is provided between PEs and has LDP, which supports MPLS Diffserv [RFC3270].

Note that this only provides Diffserv, and not the bandwidth reservation as is done with RSVP-TE.

Local policies can be implemented to map the customer's reserved flow to the LSP with the appropriate Traffic Class [RFC5462] at PE1.

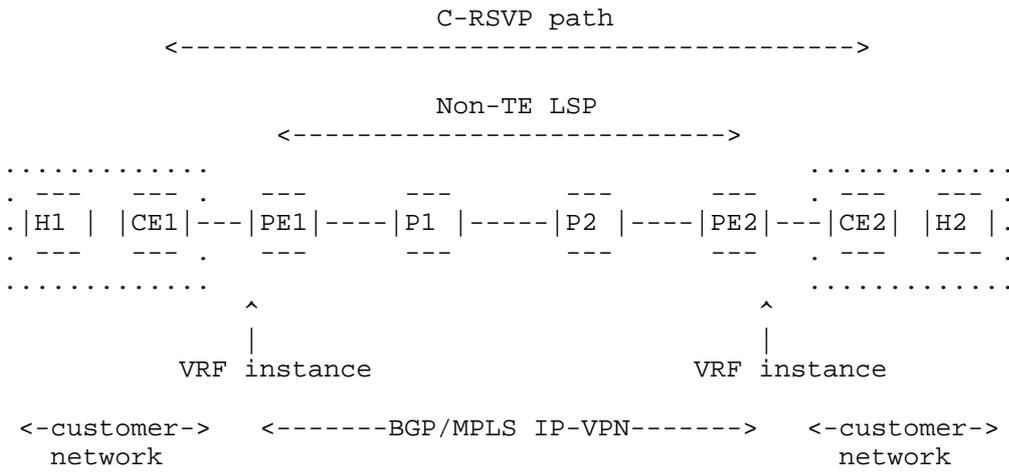


Figure 5. Scenario V

5.6. Scenario VI: RSVP-TE over Non-TE LSPs

In this scenario, as shown in Figure 6, a customer uses a VoIP application between its sites (i.e., between CE1 and CE2). H0 and H1 represent voice equipment. In this case, a non-TE LSP means LDP, and 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 of PE1 fails, C-TE LSP1 needs C-TE LSP2 as a path protection.

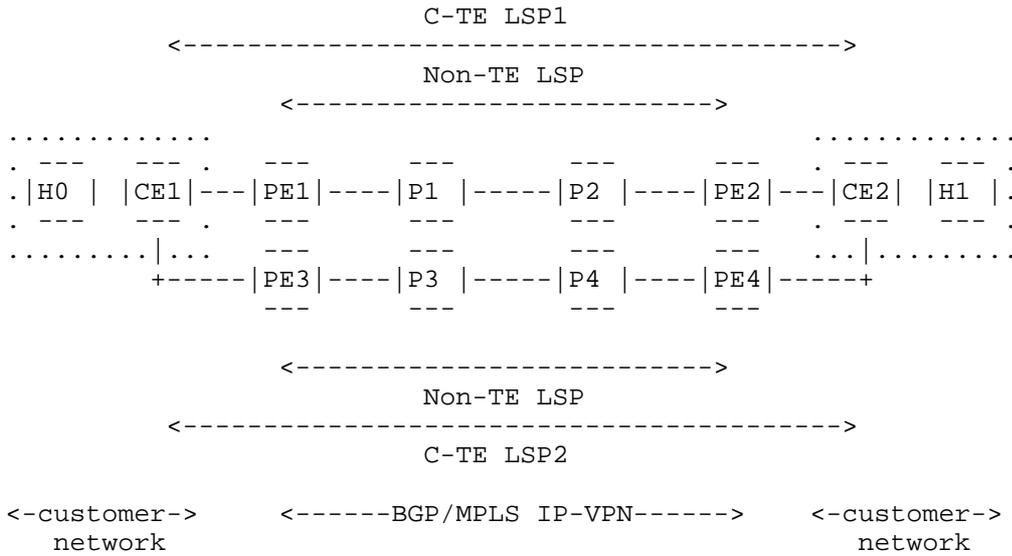


Figure 6. Scenario VI

6. Detailed Requirements for the C-TE LSP 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 LSPs a PE uses for a C-RSVP path and a C-TE LSP. When a PE receives a native RSVP and/or a path message from a CE, it MUST be able to decide which P-TE LSPs it uses. In this case, various kinds of P-TE LSPs exist in the 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: (Differentiated Services Code Point (DSCP) or Traffic Class bits)

6.2. Graceful Restart Support for C-TE LSPs

The solution SHOULD support the graceful restart capability, where the C-TE LSP traffic continues to be forwarded during a PE graceful restart. Graceful restart mechanisms related to this architecture are described in [RFC3473], [RFC3623], and [RFC4781].

6.3. Rerouting Support for C-TE LSPs

The solution MUST provide the rerouting of a C-TE LSP in case of link, node, and SRLG failures, or in case of preemption. Such rerouting may be controlled by a CE or by a PE, depending on the failure. In a dual-homed environment, the ability to perform rerouting MUST be provided against a CE-PE link failure or a PE failure, if another CE-PE link or PE 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 a 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 link protection or node protection.

The following protection MUST be supported:

1. CE link protection
2. PE node protection
3. CE node protection

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 for C-TE LSPs. C-TE LSPs may potentially try to reserve the bandwidth that exceeds the bandwidth of the 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 SHOULD 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 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. Admission Control Support for C-TE LSPs in LDP-Based Core Networks

The solution MUST support admission control for a C-TE LSP at a PE in the LDP-based core network. Specifically, PEs MUST have a configurable limit on the maximum amount of bandwidth that can be reserved by C-TE LSPs for a given VRF instance (i.e., for a given customer). Also, a PE SHOULD have a configurable limit on the total amount of bandwidth that can be reserved by C-TE LSPs between PEs.

6.7. Policy Control Support for C-TE LSPs

The solution MUST support the 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 and map, or reject, requests for a given affinity.
3. Accept and map, or reject, requests with a specified setup and/or preemption priorities.
4. Accept or reject requests for fast reroutes.
5. Ignore the requested setup and/or preemption priorities, and select a P-TE LSP based on a local policy that applies to the CE-PE link or the VRF.
6. Ignore the requested affinity, and select a P-TE LSP based on a local policy that applies to the CE-PE link or the VRF.
7. Perform mapping in the data plane between customer Traffic Class bits and transport P-TE LSP Traffic Class bits, as signaled per [RFC3270].

6.8. PCE Features Support for C-TE LSPs

The solution SHOULD support the 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 features ([RFC4655], [RFC5440]).

In this case, CE routers or PE routers may be Path Computation Clients (PCCs), and PE routers and/or P routers may be PCEs. Furthermore, the solution SHOULD support a mechanism for dynamic PCE discovery. Specifically, all PCEs are not necessarily discovered automatically, and only specific PCEs that know VPN routes should be discovered automatically.

6.9. Diversely Routed C-TE LSP Support

The solution MUST provide for setting up diversely routed C-TE LSPs over the 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 that connect to different PEs, it is desirable that multiple C-TE LSPs over the VRF instance be established between a pair of LSRs. When two CEs have multiple uplinks that connect to different PEs, it is desirable that multiple C-TE LSPs over the VRF instance be established between two different pairs of LSRs. In these cases, for example, the following points will be beneficial to customers.

1. load balance of the 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)

6.10. Optimal Path Support for C-TE LSPs

The solution MUST support the optimal path for a C-TE LSP over the VRF instance. Depending on an application (e.g., voice and video), an optimal path is needed for a C-TE LSP over the VRF instance. In the case of a TE LSP, an optimal route may be the shortest path based on the TE metric applied. For a non-TE LSP using LDP, the IGP metric may be used to compute optimal paths.

6.11. Reoptimization Support for C-TE LSPs

The solution MUST support the reoptimization of a C-TE LSP over the VRF instance. These LSPs MUST be reoptimized using "make-before-break" [RFC3209].

In this case, it is desirable for a CE to be configured with regard to the timer-based or event-driven reoptimization. Furthermore, customers SHOULD be able to reoptimize a C-TE LSP manually. To provide for delay-sensitive or jitter-sensitive traffic (i.e., voice traffic), C-TE LSP path computation and route selection are expected to be optimal for the specific application.

6.12. DS-TE Support for C-TE LSPs

The solution MUST support DS-TE [RFC4124] for a C-TE LSP over the VRF instance. In the event that the service provider and the customer have differing bandwidth constraint models, then only the service-provider bandwidth model should be supported.

Applications, which have different traffic characteristics, are used in BGP/MPLS IP-VPN environments. Service providers try to achieve the 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 Diffserv class of service on a separate C-TE LSP, DS-TE allows this traffic to utilize resources available to the given class on both shortest paths and non-shortest paths, and also to follow paths that meet TE constraints that are specific to the given class.

7. Detailed Requirements for the C-RSVP Path Model

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

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

The solution MUST support admission control at the ingress PE. PEs MUST control RSVP messages per a VRF instance.

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 manually or dynamically by operators and MAY be established if triggered by C-RSVP messages. Also, the P-TE LSP SHOULD support DS-TE.

7.3. Non-TE LSP Support for C-RSVP Paths

The solution SHOULD support non-TE LSPs (i.e., LDP-based LSP, etc.). Non-TE LSPs are established by LDP [RFC5036] between PEs and support MPLS Diffserv [RFC3270]. The solution MAY support local policies to map the customer's reserved flow to the LSP with the appropriate Traffic Class at the PE.

7.4. Transparency of C-RSVP Paths

The solution SHOULD NOT change RSVP messages from the local CE to the remote CE (Path, Resv, Path Error, Resv Error, etc.). The solution SHOULD allow customers to receive RSVP messages transparently between CE sites.

8. Commonly Detailed Requirements for Two Models

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

8.1. CE-PE Routing

The solution SHOULD 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 avoid introducing 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 avoid impacting existing RSVP and MPLS-TE mechanisms, respectively, but should allow for a smooth migration or co-existence.

8.4. Scalability Considerations

The solution SHOULD minimize the impact on network scalability from a C-RSVP path and a C-TE LSP over the VRF instance. As identified in earlier sections, PCE provides a method for offloading computation of C-TE LSPs and helps with the solution scalability.

The solution MUST address the scalability of C-RSVP paths and C-TE LSPs for the following protocols.

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, overload 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 overhead.
5. Impact and scalability of the data/forwarding plane due to added overhead.

8.6. Management Considerations

The solution MUST address the manageability of C-RSVP paths and C-TE LSPs for the following considerations.

1. Need for a MIB module for the control plane (including mapping of P-TE LSPs and C-TE LSPs) and bandwidth monitoring.
2. Need for diagnostic tools (this includes traceroute and Ping).

The solution MUST allow routers to support the MIB module for C-RSVP paths and C-TE LSPs per a VRF instance. If a CE is managed by service providers, the solution MUST allow service providers to collect MIB information for C-RSVP paths and C-TE LSPs from the CE per a customer.

Diagnostic tools can detect failures of the control plane and data plane for general MPLS-TE LSPs [RFC4379]. The solution MUST allow routers to be able to detect failures of the control plane and the data plane for C-TE LSPs over a VRF instance.

MPLS Operations, Administration, and Maintenance (OAM) for C-TE LSPs MUST be supported within the context of VRF, except for the above.

9. Security Considerations

Any solution should consider the following general security requirements:

1. The solution SHOULD NOT divulge the service-provider topology information to the customer network.
2. The solution SHOULD minimize the service-provider network's vulnerability to Denial of Service (DoS) attacks.
3. The solution SHOULD minimize the misconfiguration of DSCP marking, preemption, and holding priorities of the customer traffic.

The following additional security issues for C-TE LSPs relate to both the control plane and the data plane.

In terms of the control plane, in both the C-RSVP path and C-TE LSP models, a PE receives IPv4 or IPv6 RSVP control packets from a CE. If the CE is a router that is not trusted by service providers, the PE MUST be able to limit the rate and number of IPv4 or IPv6 RSVP control packets.

In terms of the data plane, in the C-TE LSP model, a PE receives labeled IPv4 or IPv6 data packets from a CE. If the CE is a router that is not trusted by service providers, the PE MUST be able to limit the rate of labeled IPv4 or IPv6 data packets. If the CE is a

trusted router for service providers, the PE MAY be able to limit the rate of 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 must be avoided. Therefore, a PE MUST isolate the impact of such customers' RSVP/RSVP-TE packets from other customers.

In the event that C-TE LSPs are diversely routed over VRF instances, the VRF should indicate to the CE how such diversity was provided.

10. References

10.1. Normative References

- [RFC1633] Braden, R., Clark, D., and S. Shenker, "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.
- [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.
- [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, May 2002.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.
- [RFC3623] Moy, J., Pillay-Esnault, P., and A. Lindem, "Graceful OSPF Restart", RFC 3623, November 2003.

- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [RFC4124] Le Faucheur, F., Ed., "Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering", RFC 4124, June 2005.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.
- [RFC4781] Rekhter, Y. and R. Aggarwal, "Graceful Restart Mechanism for BGP with MPLS", RFC 4781, January 2007.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, February 2009.

10.2. Informative References

- [RSVP-L3VPN] Davie, B., Le Faucheur, F., and A. Narayanan, "Support for RSVP in Layer 3 VPNs", Work in Progress, November 2009.
- [RFC4804] Le Faucheur, F., Ed., "Aggregation of Resource ReSerVation Protocol (RSVP) Reservations over MPLS TE/DS-TE Tunnels", RFC 4804, February 2007.
- [RFC5330] Vasseur, JP., Ed., Meyer, M., Kumaki, K., and A. Bonda, "A Link-Type sub-TLV to Convey the Number of Traffic Engineering Label Switched Paths Signalled with Zero Reserved Bandwidth across a Link", RFC 5330, October 2008.

[RFC5440] Vasseur, JP., Ed., and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

11. Acknowledgments

The authors would like to express thanks to Nabil Bitar, David McDysan, and Daniel King for their helpful and useful comments and feedback.

Appendix A. Reference Model

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

All scenarios in this appendix assume the following:

- A P-TE LSP is established between PE1 and PE2. This LSP is used by the VRF instance to forward customer packets within a BGP/MPLS IP-VPN.
- The service provider has ensured that enough bandwidth is available to meet the service requirements.

A.1. End-to-End C-RSVP Path Model

A C-RSVP path and a P-TE LSP are shown in Figure 7, 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 the case of a non-TE mechanism, however, it may be difficult to guarantee an end-to-end bandwidth, as resources are shared.

CE0/CE1 requests an e2e C-RSVP path to CE3/CE2 with the 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. [RFC4804] defines this scenario in more detail.

The RSVP control messages (e.g., an RSVP PATH message and an RSVP RESV message) exchanged among CEs are forwarded by IP packets through the 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.

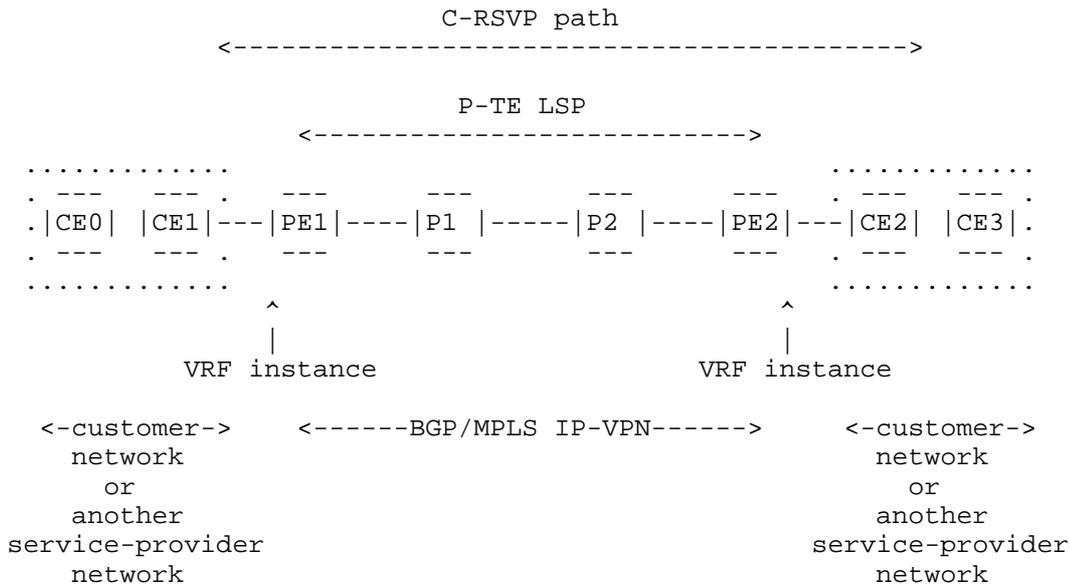


Figure 7. e2e C-RSVP Path Model

A.2. End-to-End C-TE LSP Model

A C-TE LSP and a P-TE LSP are shown in Figure 8, 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 the previous sub-section, it may be difficult to guarantee an end-to-end QoS in some cases.

CE0/CE1 requests an e2e TE LSP path to CE3/CE2 with the 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., an RSVP PATH message and an RSVP RESV message) exchanged among CEs are forwarded by a labeled packet through the 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 the BGP/MPLS IP-VPN.

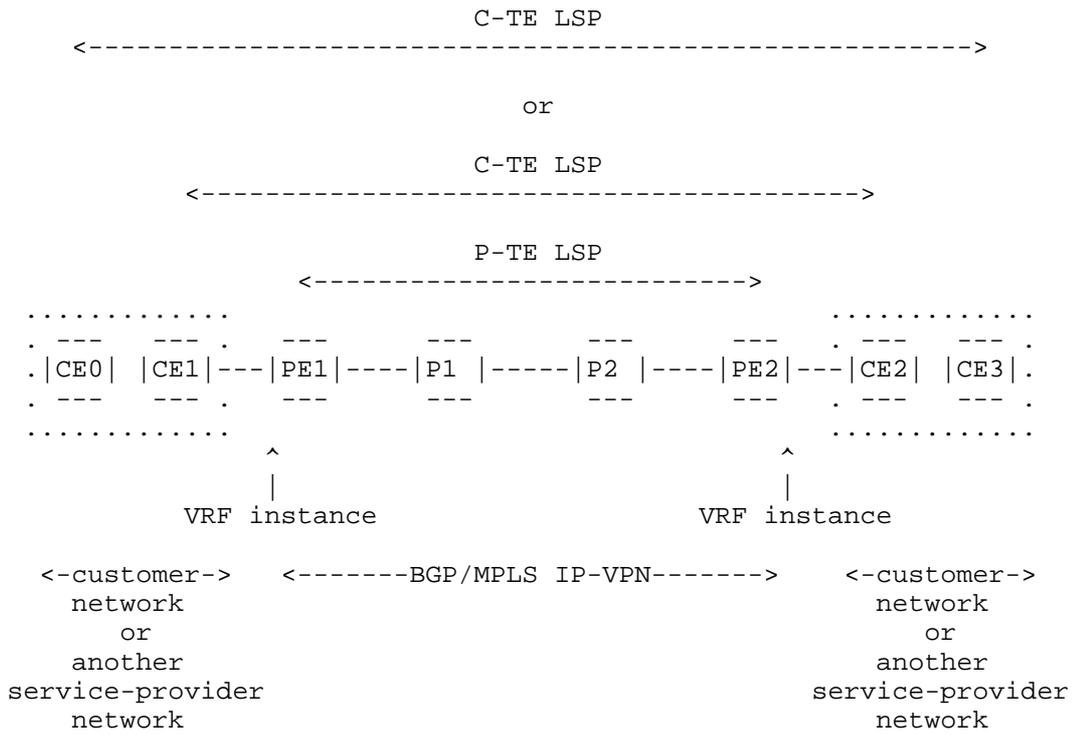


Figure 8. e2e C-TE LSP Model

Authors' Addresses

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

Raymond Zhang
BT
Farady Building, PP1.21
1 Knightrider Street
London EC4V 5BT
UK
EMail: raymond.zhang@bt.com

Yuji Kamite
NTT Communications Corporation
Granpark Tower
3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan
EMail: y.kamite@ntt.com

