

MPLS Working Group
Internet-Draft
Intended Status: Experimental RFC
Expires: August 2013

Bhargav Bhikkaji
Balaji Venkat Venkataswami
DELL
Shankar Raman
Gaurav Raina
IIT Madras
February 25, 2013

An Architecture for splicing TE-LSPs in Hierarchical CsC scenarios
draft-balaji-mpls-csc-te-lsp-splice-02

Abstract

Hierarchical Carrier Supporting Carrier deployments involve a Carrier Core which hereinafter is called the Tier-1 provider and two or more VPN sites that are carriers themselves hereinafter called Tier-2 providers that offer MPLS VPN services to their own customers. In such cases normally LDP is used to distribute labels amongst the routers (P and PE devices) in the Tier-2 provider's sites. When RSVP based TE-LSPs are constructed to explicitly route traffic for Tier-2 ISP's customers from the Tier-2 PEs to the CE of the Tier-1 provider and such TE-LSPs exist on multiple sites of the Tier-2 provider, the Tier-2 ISP may require splicing together through an "auto-match-and-splice-together" facility such that traffic flows from the PE of the Tier-2 ISP through the TE-LSP onto the CE of the Tier-1 ISP and then onto the other site and takes a path through a specific TE-LSP from the CE of the other site to the destination Tier-2 PE and then onto the final customer.

This solution offers a lot of advantages such as providing adequate assurance that the bandwidth for the traffic flowing through these spliced TE-LSPs is met. It also provides a explicit routing of the traffic rather than through the regular LDP (which follows IGP) scenarios. Such explicitly routed TE-LSPs would have been constructed taking into account factors such as using under-utilized links for example. Splicing together these TE-LSPs in various sites and doing the splicing on an auto-match based on bandwidth or delay metrics would be a good service to offer to the Tier-2 ISPs customers.

This draft outlines a scheme that offers such a feature and service to the Tier-2 ISPs through the addition of certain additional label exchanges and some additional labels such as the RSVP-stitch label and the RSVP-splicing-LDP label in the label stack which can be used to splice together these tunnels. In case of re-optimization of the LSP end-to-end there is a wide variety of choices for the near-end PE to hook up with a suitable far-end tunnel in the other Tier-2 site.

Explicit tunnel setup can be obviated by merely choosing from a set of already constructed tunnels based on criterion that may involve various parameters. Also fast-reroute in case of remote tunnel failure is taken care of.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/1id-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2013 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.

Table of Contents

1	Introduction	4
1.1	Terminology	8

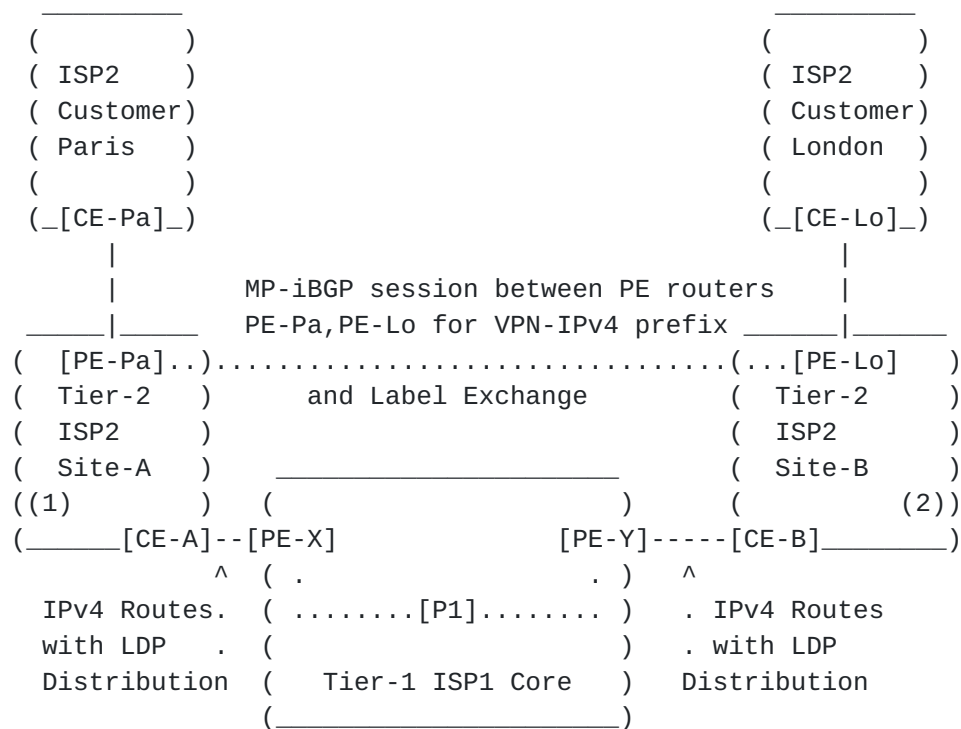
1.2	Problem Statement	8
2.0	Constructing spliced TE-LSPs between the Tier-2 sites . . .	10
2.0.1	RSVP-splicing-LDP label	11
2.0.2	RFC 6511 applicability	11
2.1	Decision at CE-B or the upstream CE in the Tier-2 ISP site. .	11
2.1.1	RSVP-stitch label	12
2.2	Across the Carrier's Core	12
2.3	Decision at PE-Lo	12
2.4	Decision at CE-B	12
2.5	Multiplicity of TE-LSP sections	13
2.6	Illustration	13
2.7	Utility	16
2.8	Tunnel Re-optimization by mere choice and not reconstruction	17
2.9	Fast-Reroute facility	17
3	Security Considerations	19
4	IANA Considerations	19
5	References	19
5.1	Normative References	19
5.2	Informative References	20
	Authors' Addresses	20

1 Introduction

Some ISP customers of the MPLS/VPN backbone may want to provide MPLS/VPN services for their own customers. These Tier-2 ISPs that we will call them henceforth obtain the services for MPLS/VPN from a Top level Carrier which we will henceforth call as a Tier-1 ISP for their connectivity when these Tier-2 ISPs do not have networks that span across the region, between geographical regions or across the globe.

This type of connectivity is known as hierarchical VPN sometimes referred to as recursive VPN. Its deployment is similar to the other Carrier's Carrier VPNs except Multi-protocol iBGP is introduced for the distribution of the prefixes and label information between ISP sites.

An example topology is provided below...



Legend :

- (1) IGP routes with LDP distribution
- (2) IGP routes with LDP distribution

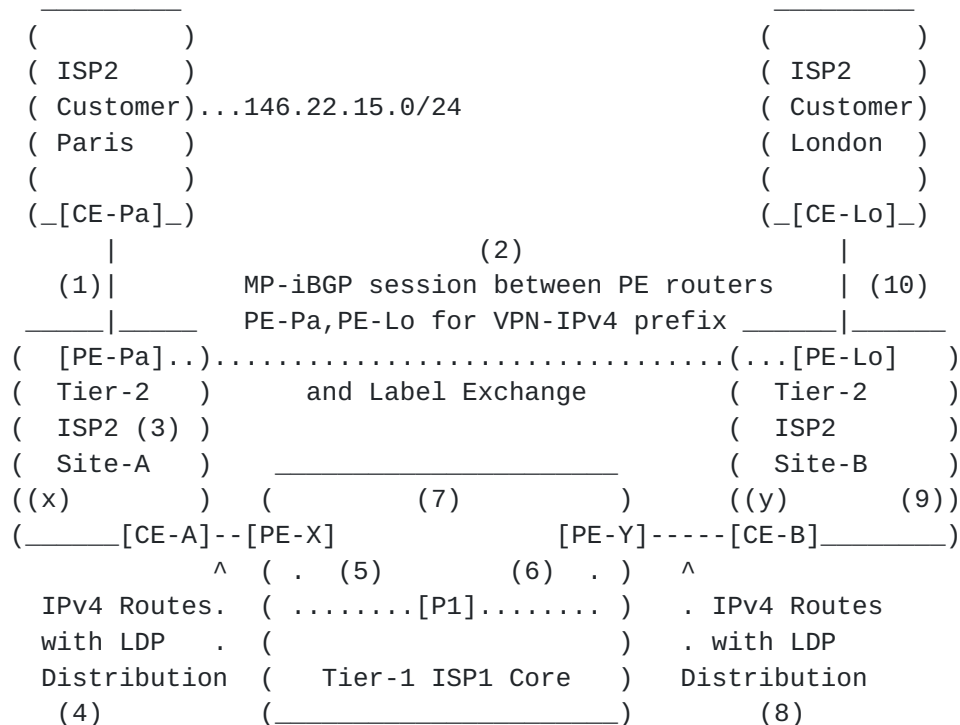
Figure 1: Reference Architecture Diagram

Within each Tier-2 ISP site in Figure 1 the PE-Routers PE-Pa and PE-Lo hold VRF routing information for any VPN customers that are attached to the POP at Paris and London. This is no different than

the standard MPLS/VPN architecture so VPN-IPv4 prefixes are assigned to customer VPN routes and are distributed between Tier-2 ISP sites using MP-iBGP with BGP extended community attributes (Router Target and Site of Origin).

Because each POP site for the Tier-2 ISP at London and Paris may hold several PE routers a full mesh of MP-iBGP is necessary among all PE-routers within the ISP MPLS/VPN topology. However again route reflectors can be deployed to cut down on the number of required peering sessions. In the example shown in Figure 1 it would be possible for example for the Tier-2 ISP2 London and Paris PE-routers to be route reflectors for their own Tier-2 ISP site. One could even deploy totally separate devices and make each PE-router a route reflector client so that MP-iBGP updates can be successfully reflected to all PE-routers that need the VPN information contained within the updates.

In the following figure 2 we provide an example of the relevant label assignment for the 146.22.15.0/24 prefix which is learned from a VPN customer of the Tier-2 ISP in the Paris area.



Legend :

- (x) IGP routes with LDP distribution
- (y) IGP routes with LDP distribution

Figure 2: Reference Diagram for normal control plane exchange for HCSC

Legend for the control plane exchange3 is as follows :

- (1) CE-Pa sends update for 146.22.15.0/24 to PE-Pa
- (2) An MP-iBGP update for Net=146.22.15.0/24 with next hop as PE-Pa and label assignment as 99 is sent to PE-Lo from PE-Pa
- (3) An IGP + LDP update for Net=PE-Pa with label as pop action is sent to CE-A from PE-Pa
- (4) The CE-A device sends an update with Net=PE-Pa with NH=CE-A and a label assignment of 1 to PE-X.
- (7) An MP-iBGP update is sent from PE-X to PE-Y with Net=PE-Pa NH=PE-X and Label as 4.
- (5) An LDP update goes from PE-X with Net as PE-X and Label as pop action to P1.

- (6) An LDP update goes from P1 with Net=PE-X and label as 2 to PE-Y
- (8) An LDP update with Net=PE-Pa and NH=PE-Y with label as 3 from PE-Y to CE-B.
- (9) An IGP update goes from CE-B with Net=PE-Pa with NH as PE-Y to PE-Lo
- (10) An LDP Update goes from CE-B to PE-Lo with Net=PE-Pa and label as 5.

The figure shows again that labels are advertised both within the MPLS/VPN backbone and within each ISP site, for each of the Tier-2 ISP (Tier-1's customer) internal routes. ISP-customer (Tier-2 customer) external routes are distributed between the sites as VPN-IPv4 routes. This mean that the iBGP session between sites becomes an MP-iBGP session so that the VPN-IPv4 routes and associated labels can be successfully advertised.

The actual traffic flow for a packet destined for a host on the 146.22.15.0/24 subnet and arriving at the Tier-2 ISP's PE-Lo router is illustrated in figure 2.

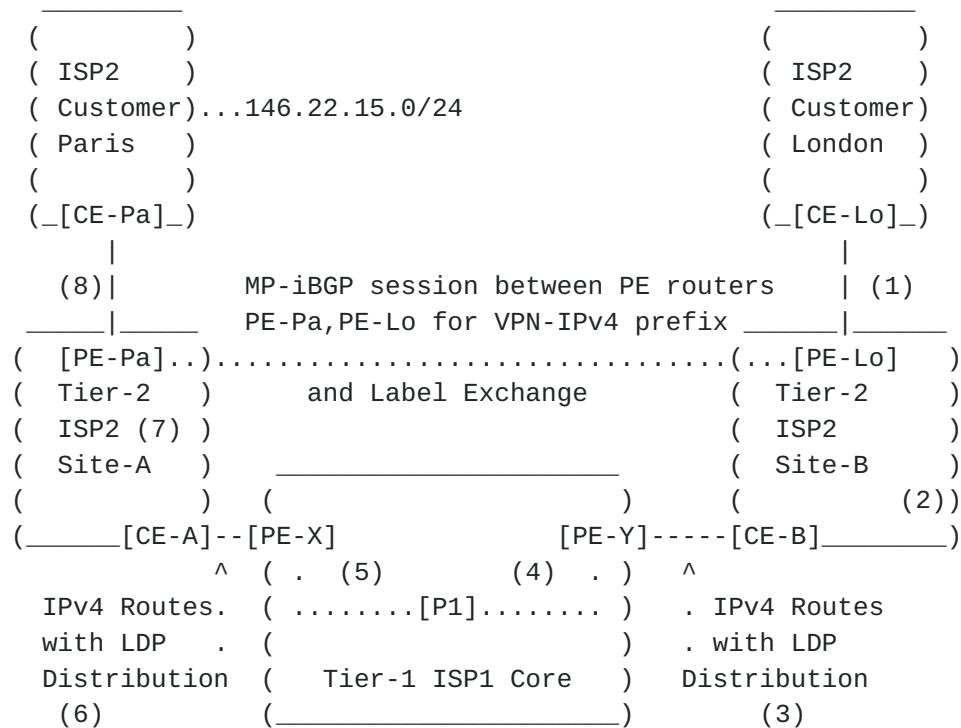


Figure 2: Reference Diagram for normal Data Plane exchange for HCSc

- (1) IP packet with IP-DA as 146.22.15.1
- (2) Label packet with MPLS labels 5,99, IP-DA 146.22.15.1
- (3) Label packet with MPLS labels 3,99, IP-DA 146.22.15.1
- (4) Label packet with MPLS labels 2,4,99, IP-DA 146.22.15.1
- (5) Label packet with MPLS labels 4,99, IP-DA 146.22.15.1
- (6) Label packet with MPLS labels 1,99, IP-DA 146.22.15.1
- (7) Label packet with MPLS labels 99, IP-DA 146.22.15.1
- (8) IP packet with IP-DA as 146.22.15.1

1.1 Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

1.2 Problem Statement

Disparate ISPs or the same ISP could have multiple Tier-2 sites

spread across geographies. This will entail that the these disparate sites be interconnected through a Tier-1 ISP who has presence in these disparate geographies and is willing to provide the interconnect service for these Tier-2 sites.

Each Tier-2 site could have pre-built TE-LSPs that possess different characteristic with respect to bandwidth, resource color, delay etc. In order that these disparate sites be interconnected by using such pre-built TE-LSPs through matching the characteristics of these TE-LSPs appropriately between a pair of Tier-2 sites there exists a need for a solution that exports RSVP-TE built TE-LSPs in the MP-iBGP / MP-eBGP protocol between such sites through prior arrangement. The solution also requires that such exported TE-LSPs are spliced together with a Tier-1 ISP's Grand Forwarding Adjacency LSP which spans across multiple areas of a Tier-1 ISP/ ISPs (in the case of inter-AS).

Thus each Tier-2 site could have pre-built TE-LSP having different characteristics such as

- available bandwidth characteristics
- Number of hops as a characteristic
- Delay bound characteristics.

Tier-2 sites do not have capability to choose TE-tunnel in respective remote site based on such characteristics etc.. Similar need exists for Inter-AS scenarios as well with the multiple Tier-1 ISPs serving disparate Tier-2 sites.

RSVP-TE tunnels mechanisms are available at Inter-AS level. But there are certain issues with them.

- These tunnels are setup end-to-end
- Requires administrative permissions across ASes

Security issues with respect to requesting LSPs to be setup in one shot from PE in Tier-2 to PE in another Tier-2 is thus manifested. Providers may have policies that prevent LSPs being setup right through to their end customer facing PEs. In case of re-optimization of the LSP end-to-end there is a wide variety of choices for the near-end PE to hookup with a suitable far-end tunnel in the other Tier-2 site. Explicit tunnel setup can be obviated by merely choosing from a set of already constructed tunnels based on criterion that may involve various parameters.

To overcome these issues the following draft provides a solution in that problem statement's space.

2.0 Constructing spliced TE-LSPs between the Tier-2 sites

It is possible that there be requirements to establish TE LSPs for a variety of reasons between the PE routers in the Tier-2 sites for example between PE-Lo and PE-Pa. These variety of reasons could include Traffic engineering necessities that may arise for the sake of allocating a certain amount of bandwidth for sets of traffic from the Tier-2 customer sites or for guaranteeing a certain delay budget or merely for utilizing under-utilized links (which otherwise would not have been used if left to IGP paths).

Consider a situation where there exists requirements to establish TE-LSPs between PE-Lo and PE-Pa for traffic direction from PE-Lo towards PE-Pa.

These TE-LSP needs to traverse the Tier-1 ISP core. To traverse the core it is possible to envisage that there exists sufficient bandwidth between PE-X and PE-Y. One possibility is to establish a TE-LSP between PE-X and PE-Y and use that LSP as a forwarding adjacency LSP for carrying the traffic over the core. The other possibility is to use normal LDP to carry the traffic over the core.

In both cases the TE-LSP established at the Tier-2 Customer sites need to be spliced together. And that splicing should be done automatically with reference to the characteristics that the TE-LSP advertises. In case of using both schemes for traversing the core, the TE-LSP in Site-B needs to be spliced to the section in Site-A.

Again it is possible that the section of the TE-LSP in Site-B was constructed independently of Site-A TE-LSP section. The TE-LSP section in Site-B being between PE-Lo and CE-B and the Site-A TE-LSP section between CE-A and PE-Pa.

Now there has to be a mechanism of conveying the section information between Site-A and Site-B. For traffic direction from Site-B to Site-A the draft solution that we propose intends to convey this TE-LSP information with TE-LSP characteristics such as bandwidth, delay, cost et... through a MP-iBGP update from CE-A to CE-B. This mechanism conveys the existence of a TE-LSP between PE-Pa and CE-A.

For the reverse direction of traffic the MP-iBGP update for the vice-versa occurs.

In our case in the diagram Figure 3 the PE-Pa-to-CE-A TE-LSP is advertised to CE-B. This information is thus sent from CE-A to CE-B. This is sent thus as a MP-iBGP update. This MP-iBGP update is sent to the PE-Pa and the PE-Lo Provider Edge routers as well. This is required since the PE-Lo in our example can take a decision to use

the TE-LSP or the normal LDP path at its end based on knobs configured or based on certain policy decisions at that time of sending the traffic where such policies could be configured. These policy decisions could be built in as an algorithm with a set of rules. This mechanism is outside the scope of the current document.

2.0.1 RSVP-splicing-LDP label

CE-B then generates two different labels towards PE-Y one for LDP and another for RSVP-splicing-LDP. The LDP label is used when the packet reaches CE-B towards PE-Y when the labels in the packet are LDP based labels. If the packet arrives with a RSVP Label for the TE-LSP between PE-Lo and CE-B at the head of the stack of labels at CE-B then the RSVP-splicing-LDP label is used.

This also means that the MP-iBGP exchange between PE-X and PE-Y has to have two classes of labels one for LDP and another for RSVP-splicing-LDP.

Additionally PE-X to CE-A labels have to be of two kinds. One for LDP and another for RSVP-splicing-LDP.

2.0.2 RFC 6511 applicability

For RSVP tunnels proposed in this mechanism it is important that non-PHP behaviour be observed by the egress LSRs in the Carrier's core and in the TE-LSP sections in the Tier-2 ISP as per [[RFC6511](#)].

2.1 Decision at CE-B or the upstream CE in the Tier-2 ISP site.

If the CE-B in our example has received MP-iBGP updates about the TE-LSP at the remote site (CE-A to PE-Lo) it can take a decision as to whether it has to use that TE-LSP or use an ordinary LDP based LSP by choosing either the LDP label or the RSVP-splicing-LDP label. MP-iBGP updates can be expected to keep the information of the TE-LSPs at the remote Tier-2 site current by periodically but not too often updating the information through a MP-iBGP update. This MP-iBGP update should have characteristics of the TE-LSP at the remote end. The characteristics relate to bandwidth and/or delay and/or MTU etc. The exact set of characteristics would also include available bandwidth on that TE-LSP. The end-point PE on the remote side connecting to the Tier-2 ISP's customer is also part of the update. In our case the PE-Pa and CE-B will know that to reach the 144.22.15.0/24 prefix there exists a TE-LSP from CE-A to PE-Pa. The MP-iBGP update from CE-A about the TE-LSP section from CE-A to PE-Pa also contains a label called the RSVP-stitch label. This RSVP-stitch label will have to be

imposed by the upstream CE-B at the Tier-2 ISP Site-B.

2.1.1 RSVP-stitch label

When the packet to 144.22.15.1 heads from PE-Lo towards CE-B, the RSVP label for the TE-LSP to be used for that traffic is the topmost label in the packet while the VPN instance label is the second label. Assume that the PE-Lo has chosen to use the TE-LSP with the stitch option in the remote Tier-2 site, then the packets are forwarded on the TE-LSP as specified. At CE-B two things happen. The RSVP label at the head of the stack of labels is swapped with the the RSVP-stitch label.

An outer label is added which is the RSVP-splicing-LDP label propositioned by PE-Y to CE-B instead of the regular LDP label. The forwarding tables are spliced with the RSVP-splicing-LDP label to be used if the incoming labeled traffic is exiting out of the RSVP tunnel at CE-B with the RSVP-stitch label.

2.2 Across the Carrier's Core

This then carries the packet to PE-Y where the outer label which is either a LDP label or a forwarding adjacency TE-LSP RSVP label is added. This makes it four labels in the Carrier's core. Once the packet reaches PE-X the RSVP-stitch label is exposed and the packet is sent to the CE-A with the RSVP-splicing-LDP label at the top of the stack. CE-A on receiving this has already stitched the forwarding action for such a label in its forwarding table to pop the RSVP-splicing-LDP label and swap the RSVP-stitch label so that the TE-LSP section from CE-A to PE-Pa is used. The packet is then sent through the TE-LSP section thereof. This action is programmed whenever a RSVP-splicing-LDP label is the incoming label into the CE-A.

The packet then reaches the end of that TE-LSP section on to the Tier-2 Provider's Site-A customer site.

2.3 Decision at PE-Lo

The decision to send the packets for a prefix through the TE-LSP at Tier-2 Site-B is first made by PE-Lo since it has information that TE-LSP between itself and CE-B exists and that there also exists a TE-LSP section at Site-A between PE-Pa and CE-A.

2.4 Decision at CE-B

On arriving at CE-B the traffic is also subject to another decision at the CE-B as to whether to use the LDP label or the RSVP-splicing-LDP label. The latter would take the traffic through the TE-LSP

section in Site-A of the Tier-2 ISP.

Thus policies can be enforced as per [section 2.3](#) and/or 2.4. The flexibility is left to the Tier-2 ISP administrator. The choice is two-fold.

2.5 Multiplicity of TE-LSP sections

There could be multiple TE-LSP section between CE-A and PE-Pa. These are conveyed through the MP-iBGP updates from CE-A to CE-B. For the reverse direction of traffic the vice-versa applies. So CE-B in our example could choose which one of these TE-LSP sections in Site-A could be the most appropriate. A suitable decision algorithm may be arrived at given the choices to be made at CE-B in our example.

2.6 Illustration

In other words to diagrammatically illustrate the methodology described above we provide the control and data plane exchanges for the same...

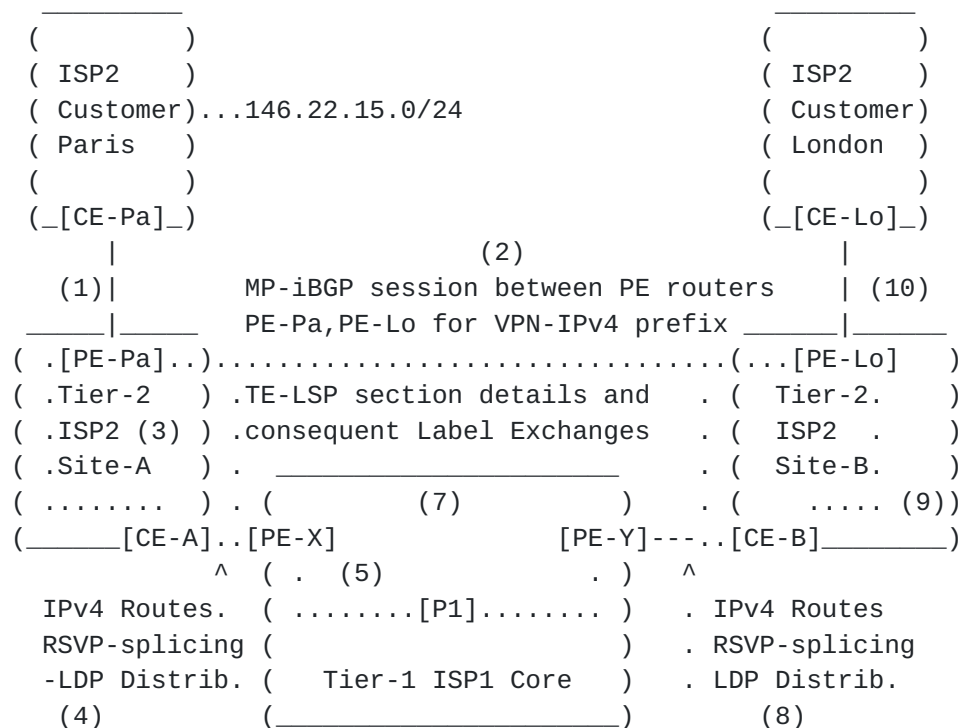


Figure 3: Reference Diagram for proposed control plane exchange for HCSc with stitch and splice TE-LSP method

Assumption is that there exist TE-LSP sections in Site-A and Site-B

between CE-A and PE-Pa in the direction specified and between PE-Lo and CE-B in that specified direction. Methods to adopt Non-PHP behaviour at CE-B is to be implemented as per [[RFC6511](#)].

We demonstrate the use of the Tier-1 ISP core RSVP TE-LSP that ties together the two TE-LSP sections in Site-A and Site-B in the data plane example. This RSVP TE-LSP too has non-PHP behaviour for its egress LSR PE-X for traffic in the London to Paris direction. The same non-PHP behaviour for the RSVP TE-LSP between CE-A and PE-Pa is also in vogue.

Legend for the control plane exchange is as follows :

- (1) CE-Pa sends update for 146.22.15.0/24 to PE-Pa
- (2) An MP-iBGP update for Net=146.22.15.0/24 with next hop as PE-Pa and label assignment as 99 is sent to PE-Lo from PE-Pa
- (2.1) An MP-iBGP update for TE-LSP section between CE-A to PE-Pa describing a RSVP-stitch label 1000 with characteristics of Site-A TE-LSP is sent to CE-B and PE-Lo.
- (3) An RSVP TE-LSP between CE-A and PE-Pa with label as 500 with Non-PHP behaviour is assumed to exist
- (4) The CE-A device sends an LDP update with Net=PE-Pa with NH=CE-A and a label assignment of 12 to PE-X where the label 12 is a RSVP-splicing-LDP label. It also sends a normal LDP label for the same FEC.
- (7) An MP-iBGP update is sent from PE-X to PE-Y with Net=PE-Pa NH=PE-X and Label as 4 where label 4 is identified as a RSVP-splicing-LDP label.
- (5) An RSVP forwarding Adjacency TE-LSP is assumed to exist between PE-X and PE-Y from latter to former with non-PHP behaviour as per [[RFC6511](#)]. The labels between PE-X and P1 is 2001 and PE-Y and P1 is 2000.
- (8) An LDP update with Net=PE-Pa and NH=PE-Y with label as 11 from PE-Y to CE-B where 11 is a RSVP-splicing-LDP label. There is also an LDP update sent for normal LDP for the same FEC.
- (9) An RSVP TE-LSP between PE-Lo and CE-B with label as 400 with non-PHP behaviour is assumed to exist.
- (10) null

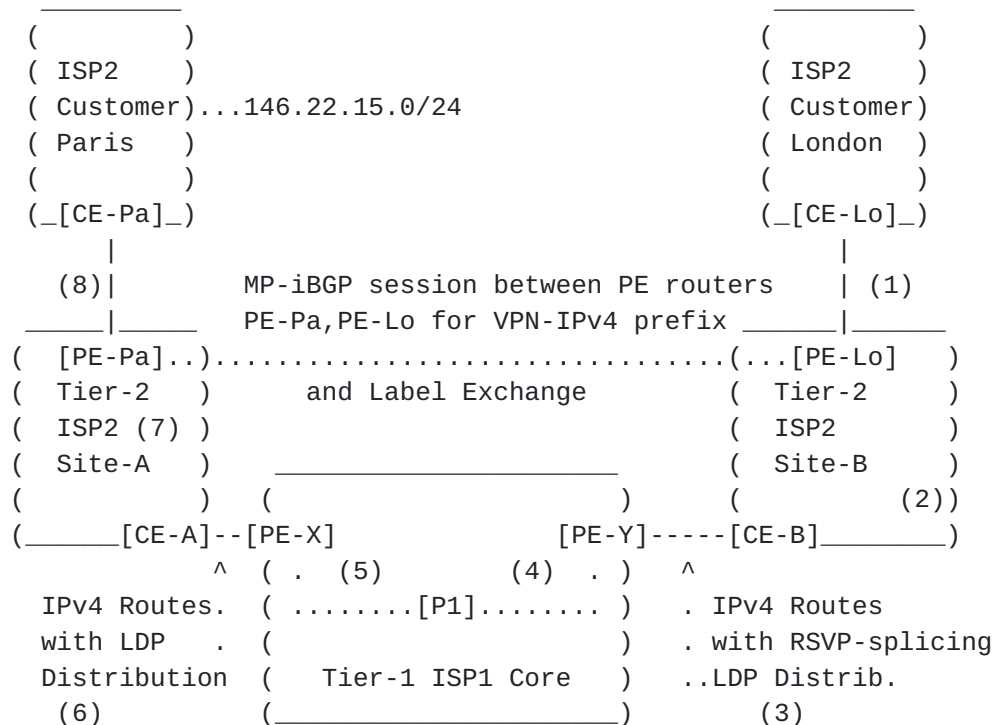


Figure 4: Reference Diagram for proposed Data Plane exchange for HCSc splicing method

Assume TE-LSPs exist between PE-Lo and CE-B in Site-B and CE-A and PE-Pa in Site-A. Also assume a forwarding adjacency LSP constructed using RSVP exists between PE-Y and PE-X in the said direction from Y to X.

- (1) IP packet with IP-DA as 146.22.15.1
- (2) Label packet with RSVP label 400,99, IP-DA 146.22.15.1
- (3) Label packet with MPLS labels 11,1000,99, IP-DA 146.22.15.1
- (4) Label packet with MPLS labels 2000,4,1000,99, IP-DA 146.22.15.1
- (5) Label packet with MPLS labels 2001,4,1000,99, IP-DA 146.22.15.1
- (6) Label packet with MPLS labels 12,1000,99, IP-DA 146.22.15.1
- (7) Label packet with MPLS labels 500,99, IP-DA 146.22.15.1
- (8) IP packet with IP-DA as 146.22.15.1

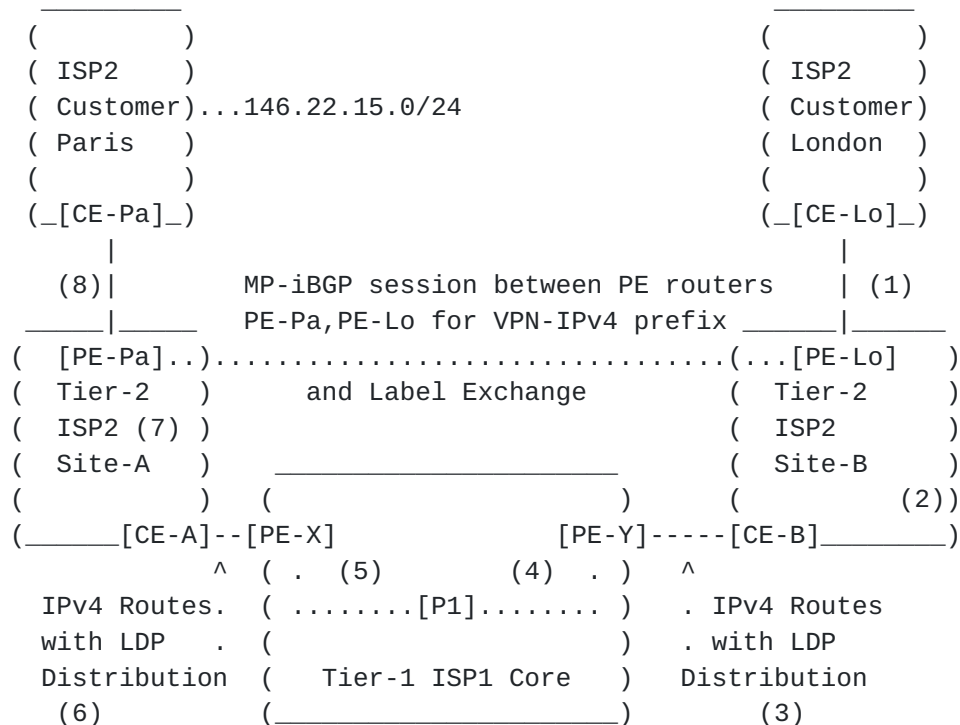


Figure 5: Reference Diagram for Data Plane exchange for proposed scheme with HCSC with LDP labels in the Carrier's Core.

Note : Control plane has not been shown for sake of brevity.

Assume TE-LSPs exist between PE-Lo and CE-B in Site-B and CE-A and PE-Pa in Site-A. Also assume there is no forwarding adjacency LSP constructed using RSVP exists between PE-Y and PE-X in the said direction from Y to X. There are only LDP labels assigned in that direction.

- (1) IP packet with IP-DA as 146.22.15.1
- (2) Label packet with RSVP label 400,99, IP-DA 146.22.15.1
- (3) Label packet with MPLS labels 11,1000,99, IP-DA 146.22.15.1
- (4) Label packet with MPLS labels 3000,4,1000,99, IP-DA 146.22.15.1
- (5) Label packet with MPLS labels 4,1000,99, IP-DA 146.22.15.1
- (6) Label packet with MPLS labels 12,1000,99, IP-DA 146.22.15.1
- (7) Label packet with MPLS labels 500,99, IP-DA 146.22.15.1
- (8) IP packet with IP-DA as 146.22.15.1

2.7 Utility

It is possible to envisage the following advantages as coming out of this proposed architecture.

- o TE-LSPs in multiple sites may be needed to be spliced together.
- o Such TE-LSPs may have been constructed to give a specific QoS for select FECs / Prefixes.
- o Without this scheme that ties the TE-LSP sections in multiple sites together, traffic may pass through TE-LSP with a given QoS in one site and may not pass through similar TE-LSP sections in other sites.
- o Splicing them together with a TE-LSP in the Tier-1 ISP gives the traffic a complete QoS experience from start to finish.
- o Multiple TE-LSP sections with different characteristics may exist in multiple sites. As per MP-iBGP updates coming in the CE devices in the Tier-2 ISP sites may choose one of them to provide the said QoS to such traffic.
- o The decision / choice to use either LDP in the sites and/or in the Carrier's core may be done by suitable algorithms that sense degradation in delay or bandwidth or a cost metric.

2.8 Tunnel Re-optimization by mere choice and not reconstruction

Through merely choosing from an available choice of multiple TE-tunnel sections in the other Tier-2 site the appropriate far-end tunnel can be hooked up with and re-signalling of the entire tunnel can be obviated. By merely matching the characteristics with the criterion applied a suitable label that will ensure choice of the far-end tunnel can be applied. This does obviate the need for re-construction of the tunnel in the far-end site. Suitable tunnels could have been constructed a-priori for this very purpose before the choice is made.

2.9 Fast-Reroute facility

A quicker fast re-route facility can be ensured if the BGP advertisement of the TE-Tunnel in the far-end site withdraws it from the near-end site. If the BGP advertisement is too late regular RSVP messaging that is targeted at the head end of the tunnel could inform such a head-end of the need to switch over to another tunnel and an appropriate choice can be made of the suitable label to ensure this.

Alternate tunnels with their respective suitable labels to impose could be chosen well in advance and the near-end PE in the Tier-2 site closest to the Tier-1 site could run a BFD session with the far-end PEs in the far-end Tier-2 sites to check the health of the far-end tunnel. When the BFD session fails and reports an error in the

health of the far-end tunnel then the appropriate alternate far-end tunnel could be chosen and a suitable label imposed at the near-end to facilitate choice of the alternate far-end tunnel.

3 Security Considerations

No additional security considerations exist except those that apply already in the current specifications.

4 IANA Considerations

MP-iBGP PDU formats for TE-LSP characteristics and for passing the RSVP-stitch label need to be added to this document.

Changes to LDP updates to indicate plain LDP labels and RSVP-splicing-LDP labels need to be incorporated. A single bit or type / code value needs to indicate whether the LDP label exchanged is either a LDP label or a RSVP-splicing-LDP label.

These will be done in the future versions of the document.

5 References

5.1 Normative References

- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [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.
- [RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", [RFC 3473](#), January 2003.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", [RFC 4875](#), May 2007.

[RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", [RFC 5420](#), February 2009.

5.2 Informative References

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), February 2006.

[RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), January 2007.

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

[RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", [RFC 5921](#), July 2010.

[MPLSVPN-ARCH] Jim Guichard et.al, MPLS and VPN Architectures, ISBN-1-58705-002-1

[RFC6511] Zafar Ali et.al, Non-Penultimate Hop Popping Behavior and Out-of-Band Mapping for RSVP-TE Label Switched Paths

Authors' Addresses

Bhargav Bhikkaji
Dell-Force10,
350 Holger Way,
San Jose, CA
U.S.A

Email: Bhargav_Bhikkaji@dell.com

Balaji Venkat Venkataswami
Dell-Force10,
Olympia Technology Park,
Fortius block, 7th & 8th Floor,

Plot No. 1, SIDCO Industrial Estate,
Guindy, Chennai - 600032.

TamilNadu, India.

Tel: +91 (0) 44 4220 8400

Fax: +91 (0) 44 2836 2446

EMail: BALAJI_VENKAT_VENKAT@dell.com

Shankar Raman

Department of Computer Science and Engineering

IIT Madras,

Chennai - 600036

TamilNadu,

India.

EMail: mjsraman@cse.iitm.ac.in

Prof.Gaurav Raina

Department of Electrical Engineering,

IIT Madras,

Chennai - 600036,

TamilNadu,

India.

EMail: gaurav@ee.iitm.ac.in

