Network Working Group
Request for Comments: 4798
Category: Standards Track

J. De Clercq
Alcatel-Lucent
D. Ooms
OneSparrow
S. Prevost
BT
F. Le Faucheur
Cisco

February 2007

Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)

Status of This Memo

Copyright Notice

Abstract

This document explains how to interconnect IPv6 islands over a Multiprotocol Label Switching (MPLS)-enabled IPv4 cloud. This approach relies on IPv6 Provider Edge routers (6PE), which are Dual Stack in order to connect to IPv6 islands and to the MPLS core, which is only required to run IPv4 MPLS. The 6PE routers exchange the IPv6 reachability information transparently over the core using the Multiprotocol Border Gateway Protocol (MP-BGP) over IPv4. In doing so, the BGP Next Hop field is used to convey the IPv4 address of the 6PE router so that dynamically established IPv4-signaled MPLS Label Switched Paths (LSPs) can be used without explicit tunnel configuration.

Table of Contents

1. Introduction

There are several approaches for providing IPv6 connectivity over an MPLS core network [RFC4029] including (i) requiring that MPLS networks support setting up IPv6-signaled Label Switched Paths (LSPs) and establish IPv6 connectivity by using those LSPs, (ii) use configured tunneling over IPv4-signaled LSPs, or (iii) use the IPv6 Provider Edge (6PE) approach defined in this document.

The 6PE approach is required as an alternative to the use of standard tunnels. It provides a solution for an MPLS environment where all tunnels are established dynamically, thereby addressing environments where the effort to configure and maintain explicitly configured tunnels is not acceptable.

This document specifies operations of the 6PE approach for interconnection of IPv6 islands over an IPv4 MPLS cloud. The approach requires that the edge routers connected to IPv6 islands be Dual Stack Multiprotocol-BGP-speaking routers [RFC4760], while the core routers are only required to run IPv4 MPLS. The approach uses MP-BGP over IPv4, relies on identification of the 6PE routers by their IPv4 address, and uses IPv4-signaled MPLS LSPs that do not require any explicit tunnel configuration.

Throughout this document, the terminology of [RFC2460] and [RFC4364] is used.

In this document an 'IPv6 island' is a network running native IPv6 as per [RFC2460]. A typical example of an IPv6 island would be a customer's IPv6 site connected via its IPv6 Customer Edge (CE) router to one (or more) Dual Stack Provider Edge router(s) of a Service Provider. These IPv6 Provider Edge routers (6PE) are connected to an IPv4 MPLS core network.

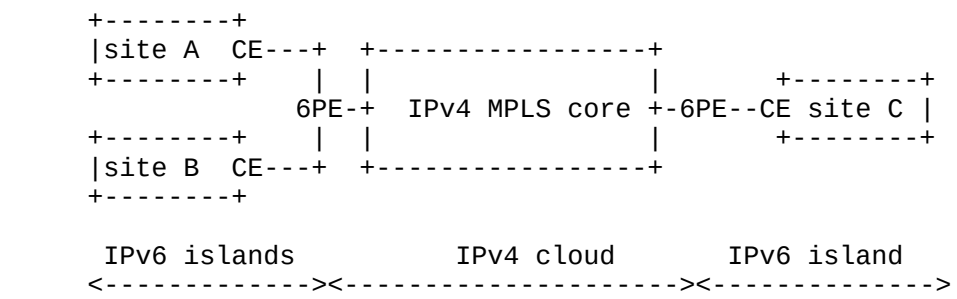

Figure 1

The interconnection method described in this document typically applies to an Internet Service Provider (ISP) that has an IPv4 MPLS network, that is familiar with BGP (possibly already offering BGP/MPLS VPN services), and that wants to offer IPv6 services to some of its customers. However, the ISP may not (yet) want to upgrade its network core to IPv6, nor use only IPv6-over-IPv4 tunneling. With the 6PE approach described here, the provider only has to upgrade some Provider Edge (PE) routers to Dual Stack operations so that they behave as 6PE routers (and route reflectors if those are used for the exchange of IPv6 reachability among 6PE routers) while leaving the IPv4 MPLS core routers untouched. These 6PE routers provide connectivity to IPv6 islands. They may also provide other services simultaneously (IPv4 connectivity, IPv4 L3VPN services, L2VPN services, etc.). Also with the 6PE approach, no tunnels need to be explicitly configured, and no IPv4 headers need to be inserted in front of the IPv6 packets between the customer and provider edge.

The ISP obtains IPv6 connectivity to its peers and upstreams using means outside of the scope of this document, and its 6PE routers readvertise it over the IPv4 MPLS core with MP-BGP.

The interface between the edge router of the IPv6 island (Customer Edge (CE) router) and the 6PE router is a native IPv6 interface which can be physical or logical. A routing protocol (IGP or EGP) may run between the CE router and the 6PE router for the distribution of IPv6 reachability information. Alternatively, static routes and/or a default route may be used on the 6PE router and the CE router to control reachability. An IPv6 island may connect to the provider network over more than one interface.

The 6PE approach described in this document can be used for customers that already have an IPv4 service from the network provider and additionally require an IPv6 service, as well as for customers that require only IPv6 connectivity.

The scenario is also described in [RFC4029].

Note that the 6PE approach specified in this document provides global IPv6 reachability. Support of IPv6 VPNs is not within the scope of this document and is addressed in [RFC4659].

Deployment of the 6PE approach over an existing IPv4 MPLS cloud does not require an introduction of new mechanisms in the core (other than potentially those described at the end of Section 3 for dealing with dynamic MTU discovery). Configuration and operations of the 6PE approach have a lot of similarities with the configuration and operations of an IPv4 VPN service ([RFC4364]) or IPv6 VPN service ([RFC4659]) over an IPv4 MPLS core because they all use MP-BGP to distribute non-IPv4 reachability information for transport over an IPv4 MPLS Core. However, the configuration and operations of the 6PE approach is somewhat simpler, since it does not involve all the VPN concepts such as Virtual Routing and Forwarding (VRFs) tables.

1.1. 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 RFC 2119 [RFC2119].

2. Protocol Overview

Each IPv6 site is connected to at least one Provider Edge router that is located on the border of the IPv4 MPLS cloud. We call such a router a 6PE router. The 6PE router MUST be dual stack IPv4 and IPv6. The 6PE router MUST be configured with at least one IPv4 address on the IPv4 side and at least one IPv6 address on the IPv6 side. The configured IPv4 address needs to be routable in the IPv4 cloud, and there needs to be a label bound via an IPv4 label distribution protocol to this IPv4 route.

As a result of this, every considered 6PE router knows which MPLS label to use to send packets to any other 6PE router. Note that an MPLS network offering BGP/MPLS IP VPN services already fulfills these requirements.

No extra routes need to be injected in the IPv4 cloud.

We call the 6PE router receiving IPv6 packets from an IPv6 site an ingress 6PE router (relative to these IPv6 packets). We call a 6PE router forwarding IPv6 packets to an IPv6 site an egress 6PE router (relative to these IPv6 packets).

Interconnecting IPv6 islands over an IPv4 MPLS cloud takes place through the following steps:

1. Exchange IPv6 reachability information among 6PE routers with MP-BGP [RFC2545]:

   The 6PE routers MUST exchange the IPv6 prefixes over MP-BGP sessions as per [RFC2545] running over IPv4. The MP-BGP Address Family Identifier (AFI) used MUST be IPv6 (value 2). In doing so, the 6PE routers convey their IPv4 address as the BGP Next Hop for the advertised IPv6 prefixes. The IPv4 address of the egress 6PE router MUST be encoded as an IPv4-mapped IPv6 address in the BGP Next Hop field. This encoding is consistent with the definition of an IPv4-mapped IPv6 address in [RFC4291] as an "address type used to represent the address of IPv4 nodes as IPv6 addresses". In addition, the 6PE MUST bind a label to the IPv6 prefix as per [RFC3107]. The Subsequence Address Family Identifier (SAFI) used in MP-BGP MUST be the "label" SAFI (value 4) as defined in [RFC3107]. Rationale for this and label allocation policies are discussed in Section 3.

2. Transport IPv6 packets from the ingress 6PE router to the egress 6PE router over IPv4-signaled LSPs:

   The ingress 6PE router MUST forward IPv6 data over the IPv4-signaled LSP towards the egress 6PE router identified by the IPv4 address advertised in the IPv4-mapped IPv6 address of the BGP Next Hop for the corresponding IPv6 prefix.

As required by the BGP specification [RFC4271], PE routers form a full peering mesh unless Route Reflectors are used.

3. Transport over IPv4-signaled LSPs and IPv6 Label Binding

In this approach, the IPv4-mapped IPv6 addresses allow a 6PE router that has to forward an IPv6 packet to automatically determine the IPv4-signaled LSP to use for a particular IPv6 destination by looking at the MP-BGP routing information.

The IPv4-signaled LSPs can be established using any existing technique for label setup [RFC3031] (LDP, RSVP-TE, etc.).

To ensure interoperability among systems that implement the 6PE approach described in this document, all such systems MUST support tunneling using IPv4-signaled MPLS LSPs established by LDP [RFC3036].

When tunneling IPv6 packets over the IPv4 MPLS backbone, rather than successively prepend an IPv4 header and then perform label imposition

based on the IPv4 header, the ingress 6PE Router MUST directly perform label imposition of the IPv6 header without prepending any IPv4 header. The (outer) label imposed MUST correspond to the IPv4-signaled LSP starting on the ingress 6PE Router and ending on the egress 6PE Router.

While this approach could theoretically operate in some situations using a single level of labels, there are significant advantages in using a second level of labels that are bound to IPv6 prefixes via MP-BGP advertisements in accordance with [RFC3107].

For instance, the use of a second level label allows Penultimate Hop Popping (PHP) on the IPv4 Label Switch Router (LSR) upstream of the egress 6PE router, without any IPv6 capabilities/upgrades on the penultimate router; this is because it still transmits MPLS packets even after the PHP (instead of having to transmit IPv6 packets and encapsulate them appropriately).

Also, an existing IPv4-signaled LSP that is using "IPv4 Explicit NULL label" over the last hop (e.g., because that LSP is already being used to transport IPv4 traffic with the Pipe Diff-Serv Tunneling Model as defined in [RFC3270]) could not be used to carry IPv6 with a single label since the "IPv4 Explicit NULL label" cannot be used to carry native IPv6 traffic (see [RFC3032]), while it could be used to carry labeled IPv6 traffic (see [RFC4182]).

This is why a second label MUST be used with the 6PE approach.

The label bound by MP-BGP to the IPv6 prefix indicates to the egress 6PE Router that the packet is an IPv6 packet. This label advertised by the egress 6PE Router with MP-BGP MAY be an arbitrary label value, which identifies an IPv6 routing context or outgoing interface to send the packet to, or MAY be the IPv6 Explicit Null Label. An ingress 6PE Router MUST be able to accept any such advertised label.

[RFC2460] requires that every link in the IPv6 Internet have an MTU of 1280 octets or larger. Therefore, on MPLS links that are used for transport of IPv6, as per the 6PE approach, and that do not support link-specific fragmentation and reassembly, the MTU must be configured to at least 1280 octets plus the encapsulation overhead.

Some IPv6 hosts might be sending packets larger than the MTU available in the IPv4 MPLS core and rely on Path MTU discovery to learn about those links. To simplify MTU discovery operations, one option is for the network administrator to engineer the MTU on the core facing interfaces of the ingress 6PE consistent with the core MTU. ICMP 'Packet Too Big' messages can then be sent back by the ingress 6PE without the corresponding packets ever entering the MPLS

core. Otherwise, routers in the IPv4 MPLS network have the option to generate an ICMP "Packet Too Big" message using mechanisms as described in Section 2.3.2, "Tunneling Private Addresses through a Public Backbone" of [RFC3032].

Note that in the above case, should a core router with an outgoing link with an MTU smaller than 1280 receive an encapsulated IPv6 packet larger than 1280, then the mechanisms of [RFC3032] may result in the "Packet Too Big" message never reaching the sender. This is because, according to [RFC4443], the core router will build an ICMP "Packet Too Big" message filled with the invoking packet up to 1280 bytes, and when forwarding downstream towards the egress PE as per [RFC3032], the MTU of the outgoing link will cause the packet to be dropped. This may cause significant operational problems; the originator of the packets will notice that his data is not getting through, without knowing why and where they are discarded. This issue would only occur if the above recommendation (to configure MTU on MPLS links of at least 1280 octets plus encapsulation overhead) is not adhered to (perhaps by misconfiguration).

4. Crossing Multiple IPv4 Autonomous Systems

This section discusses the case where two IPv6 islands are connected to different Autonomous Systems (ASes).

Like in the case of multi-AS backbone operations for IPv4 VPNs described in Section 10 of [RFC4364], three main approaches can be distinguished:

a. eBGP redistribution of IPv6 routes from AS to neighboring AS

   This approach is the equivalent for exchange of IPv6 routes to procedure (a) described in Section 10 of [RFC4364] for the exchange of VPN-IPv4 routes.

   In this approach, the 6PE routers use IBGP (according to [RFC2545] and [RFC3107] and as described in this document for the single-AS situation) to redistribute labeled IPv6 routes either to an Autonomous System Border Router (ASBR) 6PE router, or to a route reflector of which an ASBR 6PE router is a client. The ASBR then uses eBGP to redistribute the (non-labeled) IPv6 routes to an ASBR in another AS, which in turn distributes them to the 6PE routers in that AS as described earlier in this specification, or perhaps to another ASBR, which in turn distributes them etc.

There may be one, or multiple, ASBR interconnection(s) across any two ASes. IPv6 needs to be activated on the inter-ASBR links and each ASBR 6PE router has at least one IPv6 address on the interface to that link.

No inter-AS LSPs are used. There is effectively a separate mesh of LSPs across the 6PE routers within each AS.

In this approach, the ASBR exchanging IPv6 routes may peer over IPv6 or IPv4. The exchange of IPv6 routes MUST be carried out as per [RFC2545].

Note that the peering ASBR in the neighboring AS to which the IPv6 routes were distributed with eBGP, should in its turn redistribute these routes to the 6PEs in its AS using IBGP and encoding its own IPv4 address as the IPv4-mapped IPv6 BGP Next Hop.

b. eBGP redistribution of labeled IPv6 routes from AS to neighboring AS

This approach is the equivalent for exchange of IPv6 routes to procedure (b) described in Section 10 of [RFC4364] for the exchange of VPN-IPv4 routes.

In this approach, the 6PE routers use IBGP (as described earlier in this document for the single-AS situation) to redistribute labeled IPv6 routes either to an Autonomous System Border Router (ASBR) 6PE router, or to a route reflector of which an ASBR 6PE router is a client. The ASBR then uses eBGP to redistribute the labeled IPv6 routes to an ASBR in another AS, which in turn distributes them to the 6PE routers in that AS as described earlier in this specification, or perhaps to another ASBR, which in turn distributes them, etc.

There may be one, or multiple, ASBR interconnection(s) across any two ASes. IPv6 may or may not be activated on the inter-ASBR links.

This approach requires that there be label switched paths established across ASes. Hence the corresponding considerations described for procedure (b) in Section 10 of [RFC4364] apply equally to this approach for IPv6.

In this approach, the ASBR exchanging IPv6 routes may peer over IPv4 or IPv6 (in which case IPv6 obviously needs to be activated on the inter-ASBR link). When peering over IPv6, the exchange of labeled IPv6 routes MUST be carried out as per [RFC2545] and [RFC3107]. When peering over IPv4, the exchange of labeled IPv6

routes MUST be carried out as per [RFC2545] and [RFC3107] with encoding of the IPv4 address of the ASBR as an IPv4-mapped IPv6 address in the BGP Next Hop field.

c. Multi-hop eBGP redistribution of labeled IPv6 routes between source and destination ASes, with eBGP redistribution of labeled IPv4 routes from AS to neighboring AS.

   This approach is the equivalent for exchange of IPv6 routes to procedure (c) described in Section 10 of [RFC4364] for exchange of VPN-IPv4 routes.

   In this approach, IPv6 routes are neither maintained nor distributed by the ASBR routers. The ASBR routers need not be dual stack, but may be IPv4/MPLS-only routers. An ASBR needs to maintain labeled IPv4 /32 routes to the 6PE routers within its AS. It uses eBGP to distribute these routes to other ASes. ASBRs in any transit ASes will also have to use eBGP to pass along the labeled IPv4 /32 routes. This results in the creation of an IPv4 label switched path from the ingress 6PE router to the egress 6PE router. Now 6PE routers in different ASes can establish multi-hop eBGP connections to each other over IPv4, and can exchange labeled IPv6 routes (with an IPv4-mapped IPv6 BGP Next Hop) over those connections.

   IPv6 need not be activated on the inter-ASBR links.

   The considerations described for procedure (c) in Section 10 of [RFC4364] with respect to possible use of multi-hop eBGP connections via route-reflectors in different ASes, as well as with respect to the use of a third label in case the IPv4 /32 routes for the PE routers are NOT made known to the P routers, apply equally to this approach for IPv6.

   This approach requires that there be IPv4 label switched paths established across the ASes leading from a packet’s ingress 6PE router to its egress 6PE router. Hence the considerations described for procedure (c) in Section 10 of [RFC4364], with respect to LSPs spanning multiple ASes, apply equally to this approach for IPv6.

   Note also that the exchange of IPv6 routes can only start after BGP has created IPv4 connectivity between the ASes.

5. Security Considerations

The extensions defined in this document allow BGP to propagate reachability information about IPv6 routes over an MPLS IPv4 core network. As such, no new security issues are raised beyond those that already exist in BGP-4 and use of MP-BGP for IPv6.

The security features of BGP and corresponding security policy defined in the ISP domain are applicable.

For the inter-AS distribution of IPv6 routes according to case (a) of Section 4 of this document, no new security issues are raised beyond those that already exist in the use of eBGP for IPv6 [RFC2545].

For the inter-AS distribution of IPv6 routes according to case (b) and (c) of Section 4 of this document, the procedures require that there be label switched paths established across the AS boundaries. Hence the appropriate trust relationships must exist between and among the set of ASes along the path. Care must be taken to avoid "label spoofing". To this end an ASBR 6PE SHOULD only accept labeled packets from its peer ASBR 6PE if the topmost label is a label that it has explicitly signaled to that peer ASBR 6PE.

Note that for the inter-AS distribution of IPv6 routes, according to case (c) of Section 4 of this document, label spoofing may be more difficult to prevent. Indeed, the MPLS label distributed with the IPv6 routes via multi-hop eBGP is directly sent from the egress 6PE to ingress 6PEs in another AS (or through route reflectors). This label is advertised transparently through the AS boundaries. When the egress 6PE that sent the labeled IPv6 routes receives a data packet that has this particular label on top of its stack, it may not be able to verify whether the label was pushed on the stack by an ingress 6PE that is allowed to do so. As such, one AS may be vulnerable to label spoofing in a different AS. The same issue equally applies to the option (c) of Section 10 of [RFC4364]. Just as it is the case for [RFC4364], addressing this particular security issue is for further study.

6. Acknowledgements

We wish to thank Gerard Gastaud and Eric Levy-Abegnoli who contributed to this document. We also wish to thank Tri T. Nguyen, who initiated this document, but unfortunately passed away much too soon. We also thank Pekka Savola for his valuable comments and suggestions.

7. References

7.1. Normative References

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

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, March 1999.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.

[RFC3036] Andersson, L., Doolan, P., Feldman, N., Fredette, A., and B. Thomas, "LDP Specification", RFC 3036, January 2001.

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 2001.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

7.2. Informative References

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 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.

[RFC4029] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", RFC 4029, March 2005.

[RFC4182] Rosen, E., "Removing a Restriction on the use of MPLS Explicit NULL", RFC 4182, September 2005.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

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

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

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

Authors' Addresses

Jeremy De Clercq
Alcatel-Lucent
Copernicuslaan 50
Antwerpen 2018
Belgium

EMail: jeremy.de_clercq@alcatel-lucent.be

Dirk Ooms
OneSparrow
Belegstraat 13
Antwerpen 2018
Belgium

EMail: dirk@onesparrow.com

Stuart Prevost
BT
Room 136 Polaris House, Adastral Park, Martlesham Heath
Ipswich Suffolk IP5 3RE
England
EMail: stuart.prevost@bt.com

Francois Le Faucheur
Cisco
Domaine Green Side, 400 Avenue de Roumanille
Biot, Sophia Antipolis 06410
France

EMail: flefauch@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.