

Network Working Group
Request for Comments: 4659
Category: Standards Track

J. De Clercq
Alcatel
D. Ooms
OneSparrow
M. Carugi
Nortel Networks
F. Le Faucheur
Cisco Systems
September 2006

BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document describes a method by which a Service Provider may use its packet-switched backbone to provide Virtual Private Network (VPN) services for its IPv6 customers. This method reuses, and extends where necessary, the "BGP/MPLS IP VPN" method for support of IPv6. In BGP/MPLS IP VPN, "Multiprotocol BGP" is used for distributing IPv4 VPN routes over the service provider backbone, and MPLS is used to forward IPv4 VPN packets over the backbone. This document defines an IPv6 VPN address family and describes the corresponding IPv6 VPN route distribution in "Multiprotocol BGP".

This document defines support of the IPv6 VPN service over both an IPv4 and an IPv6 backbone, and for using various tunneling techniques over the core, including MPLS, IP-in-IP, Generic Routing Encapsulation (GRE) and IPsec protected tunnels. The inter-working between an IPv4 site and an IPv6 site is outside the scope of this document.

Table of Contents

| | |
|--|----|
| 1. Introduction | 2 |
| 2. The VPN-IPv6 Address Family | 4 |
| 3. VPN-IPv6 Route Distribution | 5 |
| 3.1. Route Distribution Among PEs by BGP | 5 |
| 3.2. VPN IPv6 NLRI Encoding | 6 |
| 3.2.1. BGP Next Hop encoding | 6 |
| 3.2.1.1. BGP Speaker Requesting IPv6 Transport | 7 |
| 3.2.1.2. BGP Speaker Requesting IPv4 Transport | 8 |
| 3.3. Route Target | 8 |
| 3.4. BGP Capability Negotiation | 8 |
| 4. Encapsulation | 8 |
| 5. Address Types | 10 |
| 6. Multicast | 11 |
| 7. Carriers' Carriers | 11 |
| 8. Multi-AS Backbones | 11 |
| 9. Accessing the Internet from a VPN | 13 |
| 10. Management VPN | 14 |
| 11. Security Considerations | 14 |
| 12. Quality of Service | 15 |
| 13. Scalability | 15 |
| 14. IANA Considerations | 15 |
| 15. Acknowledgements | 15 |
| 16. References | 16 |
| 16.1. Normative References | 16 |
| 16.2. Informative References | 16 |

1. Introduction

This document describes a method by which a Service Provider may use its packet-switched backbone to provide Virtual Private Network services for its IPv6 customers.

This method reuses, and extends where necessary, the "BGP/MPLS IP VPN" method [BGP/MPLS-VPN] for support of IPv6. In particular, this method uses the same "peer model" as [BGP/MPLS-VPN], in which the customers' edge routers ("CE routers") send their IPv6 routes to the Service Provider's edge routers ("PE routers"). BGP ("Border Gateway Protocol", [BGP, BGP-MP]) is then used by the Service Provider to exchange the routes of a particular IPv6 VPN among the PE routers that are attached to that IPv6 VPN. Eventually, the PE routers distribute, to the CE routers in a particular VPN, the IPv6 routes from other CE routers in that VPN. As with IPv4 VPNs, a key characteristic of this "peer model" is that the (IPv6) CE routers within an (IPv6) VPN do not peer with each other; there is no "overlay" visible to the (IPv6) VPN's routing algorithm.

This document adopts the definitions, acronyms, and mechanisms described in [BGP/MPLS-VPN]. Unless it is stated otherwise, the mechanisms of [BGP/MPLS-VPN] apply and will not be re-described here.

A VPN is said to be an IPv6 VPN when each site of this VPN is IPv6 capable and is natively connected over an IPv6 interface or sub-interface to the Service Provider (SP) backbone via a Provider Edge device (PE).

A site may be both IPv4 capable and IPv6 capable. The logical interface on which packets arrive at the PE may determine the IP version. Alternatively, the same logical interface may be used for both IPv4 and IPv6, in which case a per-packet lookup at the Version field of the IP packet header determines the IP version.

This document only concerns itself with handling of IPv6 communication between IPv6 hosts located on IPv6-capable sites. Handling of IPv4 communication between IPv4 hosts located on IPv4-capable sites is outside the scope of this document and is covered in [BGP/MPLS-VPN]. Communication between an IPv4 host located in an IPv4-capable site and an IPv6 host located in an IPv6-capable site is outside the scope of this document.

In a similar manner to how IPv4 VPN routes are distributed in [BGP/MPLS-VPN], BGP and its extensions are used to distribute routes from an IPv6 VPN site to all the other PE routers connected to a site of the same IPv6 VPN. PEs use "VPN Routing and Forwarding tables" (VRFs) to maintain the reachability information and forwarding information of each IPv6 VPN separately.

As is done for IPv4 VPNs [BGP/MPLS-VPN], we allow each IPv6 VPN to have its own IPv6 address space, which means that a given address may denote different systems in different VPNs. This is achieved via a new address family, the VPN-IPv6 Address Family, in a fashion similar to that of the VPN-IPv4 address family, defined in [BGP/MPLS-VPN], which prepends a Route Distinguisher to the IP address.

In addition to its operation over MPLS Label Switched Paths (LSPs), the IPv4 BGP/MPLS VPN solution has been extended to allow operation over other tunneling techniques, including GRE tunnels, IP-in-IP tunnels [2547-GRE/IP], L2TPv3 tunnels [MPLS-in-L2TPv3], and IPsec protected tunnels [2547-IPsec]. In a similar manner, this document allows support of an IPv6 VPN service over MPLS LSPs, as well as over other tunneling techniques.

This document allows support for an IPv6 VPN service over an IPv4 backbone, as well as over an IPv6 backbone. The IPv6 VPN service supported is identical in both cases.

The IPv6 VPN solution defined in this document offers the following benefits:

- o From both the Service Provider perspective and the customer perspective, the VPN service that can be supported for IPv6 sites is identical to the one that can be supported for IPv4 sites.
- o From the Service Provider perspective, operations of the IPv6 VPN service require the exact same skills, procedures, and mechanisms as those for the IPv4 VPN service.
- o Where both IPv4 VPNs and IPv6 VPN services are supported over an IPv4 core, the same single set of MP-BGP peering relationships and the same single PE-PE tunnel mesh MAY be used for both.
- o The IPv6 VPN service is independent of whether the core runs IPv4 or IPv6. This is so that the IPv6 VPN service supported before and after a migration of the core from IPv4 to IPv6 is undistinguishable to the VPN customer.

Note that supporting IPv4 VPN services over an IPv6 core is not covered by this document.

2. The VPN-IPv6 Address Family

The BGP Multiprotocol Extensions [BGP-MP] allow BGP to carry routes from multiple "address families". We introduce the notion of the "VPN-IPv6 address family", which is similar to the VPN-IPv4 address family introduced in [BGP/MPLS-VPN].

A VPN-IPv6 address is a 24-octet quantity, beginning with an 8-octet "Route Distinguisher" (RD) and ending with a 16-octet IPv6 address.

The purpose of the RD is solely to allow one to create distinct routes to a common IPv6 address prefix, which is similar to the purpose of the RD defined in [BGP/MPLS-VPN]. In the same way as it is possible per [BGP/MPLS-VPN], the RD can be used to create multiple different routes to the very same system. This can be achieved by creating two different VPN-IPv6 routes that have the same IPv6 part but different RDs. This allows the provider's BGP to install multiple different routes to the same system and allows policy to be used to decide which packets use which route.

Also, if two VPNs were to use the same IPv6 address prefix (effectively denoting different physical systems), the PEs would translate these into unique VPN-IPv6 address prefixes using different RDs. This ensures that if the same address is ever used in two different VPNs, it is possible to install two completely different routes to that address, one for each VPN.

Since VPN-IPv6 addresses and IPv6 addresses belong to different address families, BGP never treats them as comparable addresses.

A VRF may have multiple equal-cost VPN-IPv6 routes for a single IPv6 address prefix. When a packet's destination address is matched in a VRF against a VPN-IPv6 route, only the IPv6 part is actually matched.

The Route Distinguisher format and encoding is as specified in [BGP/MPLS-VPN].

When a site is IPv4 capable and IPv6 capable, the same RD MAY be used for the advertisement of IPv6 addresses and IPv4 addresses. Alternatively, a different RD MAY be used for the advertisement of the IPv4 addresses and of the IPv6 addresses. Note, however, that in the scope of this specification, IPv4 addresses and IPv6 addresses will always be handled in separate contexts, and that no IPv4-IPv6 interworking issues and techniques will be discussed.

3. VPN-IPv6 Route Distribution

3.1. Route Distribution Among PEs by BGP

As described in [BGP/MPLS-VPN], if two sites of a VPN attach to PEs that are in the same Autonomous System, the PEs can distribute VPN routes to each other by means of an (IPv4) internal Border Gateway Protocol (iBGP) connection between them. Alternatively, each PE can have iBGP connections to route reflectors. Similarly, for IPv6 VPN route distribution, PEs can use iBGP connections between them or use iBGP connections to route reflectors. For IPv6 VPN, the iBGP connections MAY be over IPv4 or over IPv6.

The PE routers exchange, via MP-BGP [BGP-MP], reachability information for the IPv6 prefixes in the IPv6 VPNs and thereby announce themselves as the BGP Next Hop.

The rules for encoding the reachability information and the BGP Next Hop address are specified in the following sections.

3.2. VPN IPv6 NLRI Encoding

When distributing IPv6 VPN routes, the advertising PE router MUST assign and distribute MPLS labels with the IPv6 VPN routes. Essentially, PE routers do not distribute IPv6 VPN routes, but Labeled IPv6 VPN routes [MPLS-BGP]. When the advertising PE then receives a packet that has this particular advertised label, the PE will pop this label from the MPLS stack and process the packet appropriately (i.e., forward it directly according to the label or perform a lookup in the corresponding IPv6-VPN context).

The BGP Multiprotocol Extensions [BGP-MP] are used to advertise the IPv6 VPN routes in the MP_REACH Network Layer Reachability Information (NLRI). The Address Family Identifier (AFI) and Subsequent Address Family Identifier (SAFI) fields MUST be set as follows:

- AFI: 2; for IPv6
- SAFI: 128; for MPLS labeled VPN-IPv6

The NLRI field itself is encoded as specified in [MPLS-BGP]. In the context of this extension, the prefix belongs to the VPN-IPv6 Address Family and thus consists of an 8-octet Route Distinguisher followed by an IPv6 prefix as specified in Section 2, above.

3.2.1. BGP Next Hop encoding

The encoding of the BGP Next Hop depends on whether the policy of the BGP speaker is to request that IPv6 VPN traffic be transported to that BGP Next Hop using IPv6 tunneling ("BGP speaker requesting IPv6 transport") or using IPv4 tunneling ("BGP speaker requesting IPv4 transport").

Definition of this policy (to request transport over IPv4 tunneling or IPv6 tunneling) is the responsibility of the network operator and is beyond the scope of this document. Note that it is possible for that policy to request transport over IPv4 (resp. IPv6) tunneling while the BGP speakers exchange IPv6 VPN reachability information over IPv6 (resp. IPv4). However, in that case, a number of operational implications are worth considering. In particular, an undetected fault affecting the IPv4 (resp. IPv6) tunneling data path and not affecting the IPv6 (resp. IPv4) data path could remain undetected by BGP, which in turn may result in black-holing of traffic.

Control of this policy is beyond the scope of this document and may be based on user configuration.

3.2.1.1. BGP Speaker Requesting IPv6 Transport

When the IPv6 VPN traffic is to be transported to the BGP speaker using IPv6 tunneling (e.g., IPv6 MPLS LSPs, IPsec-protected IPv6 tunnels), the BGP speaker SHALL advertise a Next Hop Network Address field containing a VPN-IPv6 address

- whose 8-octet RD is set to zero, and
- whose 16-octet IPv6 address is set to the global IPv6 address of the advertising BGP speaker.

This is potentially followed by another VPN-IPv6 address

- whose 8-octet RD is set to zero, and
- whose 16-octet IPv6 address is set to the link-local IPv6 address of the advertising BGP speaker.

The value of the Length of the Next Hop Network Address field in the MP_REACH_NLRI attribute shall be set to 24 when only a global address is present, and to 48 if a link-local address is also included in the Next Hop field.

If the BGP speakers peer using only their link-local IPv6 address (for example, in the case where an IPv6 CE peers with an IPv6 PE, where the CE does not have any IPv6 global address, and where eBGP peering is achieved over the link-local addresses), the "unspecified address" ([V6ADDR]) is used by the advertising BGP speaker to indicate the absence of the global IPv6 address in the Next Hop Network Address field.

The link-local address shall be included in the Next Hop field if and only if the advertising BGP speaker shares a common subnet with the peer the route is being advertised to [BGP-IPv6].

In all other cases, a BGP speaker shall advertise to its peer in the Next Hop Network Address field only the global IPv6 address of the next hop.

As a consequence, a BGP speaker that advertises a route to an internal peer may modify the Network Address of Next Hop field by removing the link-local IPv6 address of the next hop.

An example scenario where both the global IPv6 address and the link-local IPv6 address shall be included in the BGP Next Hop address field is that where the IPv6 VPN service is supported over a multi-Autonomous System (AS) backbone with redistribution of labeled VPN-

IPv6 routes between Autonomous System Border Routers (ASBR) of different ASes sharing a common IPv6 subnet: in that case, both the global IPv6 address and the link-local IPv6 address shall be advertised by the ASBRs.

3.2.1.2. BGP Speaker Requesting IPv4 Transport

When the IPv6 VPN traffic is to be transported to the BGP speaker using IPv4 tunneling (e.g., IPv4 MPLS LSPs, IPsec-protected IPv4 tunnels), the BGP speaker SHALL advertise to its peer a Next Hop Network Address field containing a VPN-IPv6 address:

- whose 8-octet RD is set to zero, and
- whose 16-octet IPv6 address is encoded as an IPv4-mapped IPv6 address [V6ADDR] containing the IPv4 address of the advertising BGP speaker. This IPv4 address must be routable by the other BGP Speaker.

3.3. Route Target

The use of route target is specified in [BGP/MPLS-VPN] and applies to IPv6 VPNs. Encoding of the extended community attribute is defined in [BGP-EXTCOM].

3.4. BGP Capability Negotiation

In order for two PEs to exchange labeled IPv6 VPN NLRIs, they MUST use BGP Capabilities Negotiation to ensure that they both are capable of properly processing such NLRIs. This is done as specified in [BGP-MP] and [BGP-CAP], by using capability code 1 (multiprotocol BGP), with AFI and SAFI values as specified above, in Section 3.2.

4. Encapsulation

The ingress PE Router MUST tunnel IPv6 VPN data over the backbone towards the Egress PE router identified as the BGP Next Hop for the corresponding destination IPv6 VPN prefix.

When the 16-octet IPv6 address contained in the BGP Next Hop field is encoded as an IPv4-mapped IPv6 address (see Section 3.2.1.2), the ingress PE MUST use IPv4 tunneling unless explicitly configured to do otherwise. The ingress PE MAY optionally allow, through explicit configuration, the use of IPv6 tunneling when the 16-octet IPv6 address contained in the BGP Next Hop field is encoded as an IPv4-mapped IPv6 address. This would allow support of particular deployment environments where IPv6 tunneling is desired but where IPv4-mapped IPv6 addresses happen to be used for IPv6 reachability of the PEs instead of Global IPv6 addresses.

When the 16-octet IPv6 address contained in the BGP Next Hop field is not encoded as an IPv4-mapped address (see Section 3.2.1.1), the ingress PE MUST use IPv6 tunneling.

When a PE receives a packet from an attached CE, it looks up the packet's IPv6 destination address in the VRF corresponding to that CE. This enables it to find a VPN-IPv6 route. The VPN-IPv6 route will have an associated MPLS label and an associated BGP Next Hop. First, this MPLS label is pushed on the packet as the bottom label. Then, this labeled packet is encapsulated into the tunnel for transport to the egress PE identified by the BGP Next Hop. Details of this encapsulation depend on the actual tunneling technique, as follows:

As with MPLS/BGP for IPv4 VPNs [2547-GRE/IP], when tunneling is done using IPv4 tunnels or IPv6 tunnels (resp. IPv4 GRE tunnels or IPv6 GRE tunnels), encapsulation of the labeled IPv6 VPN packet results in an MPLS-in-IP (resp. MPLS-in-GRE) encapsulated packet as specified in [MPLS-in-IP/GRE]. When tunneling is done using L2TPv3, encapsulation of the labeled IPv6 VPN packet results in an MPLS-in-L2TPv3-encapsulated packet, as specified in [MPLS-in-L2TPv3].

As with MPLS/BGP for IPv4 VPNs, when tunneling is done using an IPsec secured tunnel [2547-IPsec], encapsulation of the labeled IPv6 VPN packet results in an MPLS-in-IP- or MPLS-in-GRE-encapsulated packet [MPLS-in-IP/GRE]. The IPsec Transport Mode is used to secure this IPv4 or GRE tunnel from ingress PE to egress PE.

When tunneling is done using IPv4 tunnels (whether IPsec secured or not), the Ingress PE Router MUST use the IPv4 address that is encoded in the IPv4-mapped IPv6 address field of the BGP next hop field as the destination address of the prepended IPv4 tunneling header. It uses one of its IPv4 addresses as the source address of the prepended IPv4 tunneling header.

When tunneling is done using IPv6 tunnels (whether IPsec secured or not), the Ingress PE Router MUST use the IPv6 address that is contained in the IPv6 address field of the BGP next hop field as the destination address of the prepended IPv6 tunneling header. It uses one of its IPv6 addresses as the source address of the prepended IPv6 tunneling header.

When tunneling is done using MPLS LSPs, the LSPs can be established using any label distribution technique (LDP [LDP], RSVP-TE [RSVP-TE], etc.).

When tunneling is done using MPLS LSPs, the ingress PE Router MUST directly push the LSP tunnel label on the label stack of the labeled IPv6 VPN packet (i.e., without prepending any IPv4 or IPv6 header). This pushed label corresponds to the LSP starting on the ingress PE Router and ending on the egress PE Router. The BGP Next Hop field is used to identify the egress PE router and in turn the label to be pushed on the stack. When the IPv6 address in the BGP Next Hop field is an IPv4-mapped IPv6 address, the embedded IPv4 address will determine the tunnel label to push on the label stack. In any other case, the IPv6 address in the BGP Next Hop field will determine the tunnel label to push on the label stack.

To ensure interoperability among systems that implement this VPN architecture, all such systems MUST support tunneling using MPLS LSPs established by LDP [LDP].

5. Address Types

Since Link-local unicast addresses are defined for use on a single link only, those may be used on the PE-CE link, but they are not supported for reachability across IPv6 VPN Sites and are never advertised via MultiProtocol-Border Gateway Protocol (MP-BGP) to remote PEs.

Global unicast addresses are defined as uniquely identifying interfaces anywhere in the IPv6 Internet. Global addresses are expected to be commonly used within and across IPv6 VPN Sites. They are obviously supported by this IPv6 VPN solution for reachability across IPv6 VPN Sites and advertised via MP-BGP to remote PEs and are processed without any specific considerations to their global scope.

Quoting from [UNIQUE-LOCAL]: "This document defines an IPv6 unicast address format that is globally unique and is intended for local communications [IPv6]. These addresses are called Unique Local IPv6 Unicast Addresses and are abbreviated in this document as Local IPv6 addresses. They are not expected to be routable on the global Internet. They are routable inside of a more limited area such as a site. They may also be routed between a limited set of sites."

[UNIQUE-LOCAL] also says in its Section 4.7: "Local IPv6 addresses can be used for inter-site Virtual Private Networks (VPN) if appropriate routes are set up. Because the addresses are unique these VPNs will work reliably and without the need for translation. They have the additional property that they will continue to work if the individual sites are renumbered or merged."

In accordance with this, Unique Local IPv6 Unicast Addresses are supported by the IPv6 VPN solution specified in this document for reachability across IPv6 VPN Sites. Hence, reachability to such Unique Local IPv6 Addresses may be advertised via MP-BGP to remote PEs and processed by PEs in the same way as Global Unicast addresses.

Recommendations and considerations for which of these supported address types should be used in given IPv6 VPN environments are beyond the scope of this document.

6. Multicast

Multicast operations are outside the scope of this document.

7. Carriers' Carriers

Sometimes, an IPv6 VPN may actually be the network of an IPv6 ISP, with its own peering and routing policies. Sometimes, an IPv6 VPN may be the network of an SP that is offering VPN services in turn to its own customers. IPv6 VPNs like these can also obtain backbone service from another SP, the "Carrier's Carrier", using the Carriers' Carrier method described in Section 9 of [BGP/MPLS-VPN] but applied to IPv6 traffic. All the considerations discussed in [BGP/MPLS-VPN] for IPv4 VPN Carriers' Carrier apply for IPv6 VPN, with the exception that the use of MPLS (including label distribution) between the PE and the CE pertains to IPv6 routes instead of IPv4 routes.

8. Multi-AS Backbones

The same procedures described in Section 10 of [BGP/MPLS-VPN] can be used (and have the same scalability properties) to address the situation where two sites of an IPv6 VPN are connected to different Autonomous Systems. However, some additional points should be noted

when applying these procedures for IPv6 VPNs; these are further described in the remainder of this section.

Approach (a): VRF-to-VRF connections at the AS (Autonomous System) border routers.

This approach is the equivalent for IPv6 VPNs to procedure (a) in Section 10 of [BGP/MPLS-VPN]. In the case of IPv6 VPNs, IPv6 needs to be activated on the inter-ASBR VRF-to-VRF (sub)interfaces. In this approach, the ASBRs exchange IPv6 routes (as opposed to VPN-IPv6 routes) and may peer over IPv6 or over IPv4. The exchange of IPv6 routes MUST be carried out as per [BGP-IPv6]. This method does not use inter-AS LSPs.

Finally, note that with this procedure, since every AS independently implements the intra-AS procedures for IPv6 VPNs described in this document, the participating ASes may all internally use IPv4 tunneling, or IPv6 tunneling; or alternatively, some participating ASes may internally use IPv4 tunneling while others use IPv6 tunneling.

Approach (b): EBGp redistribution of labeled VPN-IPv6 routes from AS to neighboring AS.

This approach is the equivalent for IPv6 VPNs to procedure (b) in Section 10 of [BGP/MPLS-VPN]. With this approach, the ASBRs use EBGp to redistribute labeled VPN-IPv4 routes to ASBRs in other ASes.

In this approach, IPv6 may or may not be activated on the inter-ASBR links since the ASBRs exchanging VPN-IPv6 routes may peer over IPv4 or IPv6 (in which case, IPv6 obviously needs to be activated on the inter-ASBR link). The exchange of labeled VPN-IPv6 routes MUST be carried out as per [BGP-IPv6] and [MPLS-BGP]. When the VPN-IPv6 traffic is to be transported using IPv6 tunneling, the BGP Next Hop Field SHALL contain an IPv6 address. When the VPN-IPv6 traffic is to be transported using IPv4 tunneling, the BGP Next Hop Field SHALL contain an IPv4 address encoded as an IPv4-mapped IPv6 address.

This approach requires that there be inter-AS LSPs. As such, the corresponding (security) considerations described for procedure (b) in Section 10 of [BGP/MPLS-VPN] apply equally to this approach for IPv6.

Finally, note that with this procedure, as with procedure (a), since every AS independently implements the intra-AS procedures for IPv6 VPNs described in this document, the participating ASes may all internally use IPv4 tunneling or IPv6 tunneling; alternatively, some participating ASes may internally use IPv4 tunneling while others use IPv6 tunneling.

Approach (c): Multihop EBGp redistribution of labeled VPN-IPv6 routes between source and destination ASes, with EBGp redistribution of labeled IPv4 or IPv6 routes from AS to neighboring AS.

This approach is equivalent for exchange of VPN-IPv6 routes to procedure (c) in Section 10 of [BGP/MPLS-VPN] for exchange of VPN-IPv4 routes.

This approach requires that the participating ASes either all use IPv4 tunneling or all use IPv6 tunneling.

In this approach, VPN-IPv6 routes are neither maintained nor distributed by the ASBR routers. The ASBR routers need not be dual stack. An ASBR needs to maintain labeled IPv4 (or IPv6) routes to the PE 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 (or IPv6) routes. This results in the creation of an IPv4 (or IPv6) label switch path from ingress PE router to egress PE router. Now, PE routers in different ASes can establish multi-hop EBGp connections to each other over IPv4 or IPv6 and can exchange labeled VPN-IPv6 routes over those EBGp connections. Note that the BGP Next Hop field of these distributed VPN-IPv6 routes will contain an IPv6 address when IPv6 tunneling is used or an IPv4-mapped IPv6 address when IPv4 tunneling is used.

The considerations described for procedure (c) in Section 10 of [BGP/MPLS-VPN] with respect to possible use of route-reflectors, with respect to possible use of a third label, and with respect to LSPs spanning multiple ASes apply equally to this IPv6 VPN approach.

9. Accessing the Internet from a VPN

The methods proposed by [BGP/MPLS-VPN] to access the global IPv4 Internet from an IPv4 VPN can be used in the context of IPv6 VPNs and the global IPv6 Internet. Note, however, that if the IPv6 packets from IPv6 VPN sites and destined for the global IPv6 Internet need to traverse the SP backbone, and that if this is an IPv4 only backbone, these packets must be tunneled through that IPv4 backbone.

Clearly, as is the case outside the VPN context, access to the IPv6 Internet from an IPv6 VPN requires the use of global IPv6 addresses.

In particular, Unique Local IPv6 addresses cannot be used for IPv6 Internet access.

10. Management VPN

The management considerations discussed in Section 12 of [BGP/MPLS-VPN] apply to the management of IPv6 VPNs.

Where the Service Provider manages the CE of the IPv6 VPN site, the Service Provider may elect to use IPv4 for communication between the management tool and the CE for such management purposes. In that case, regardless of whether a customer IPv4 site is actually connected to the CE (in addition to the IPv6 site), the CE is effectively part of an IPv4 VPN in addition to belonging to an IPv6 VPN (i.e., the CE is attached to a VRF that supports IPv4 in addition to IPv6). Considerations presented in [BGP/MPLS-VPN], on how to ensure that the management tool can communicate with such managed CEs from multiple VPNs without allowing undesired reachability across CEs of different VPNs, are applicable to the IPv4 reachability of the VRF to which the CE attaches.

Where the Service Provider manages the CE of the IPv6 VPN site, the Service Provider may elect to use IPv6 for communication between the management tool and the CE for such management purposes. Considerations presented in [BGP/MPLS-VPN], on how to ensure that the management tool can communicate with such managed CEs from multiple VPNs without allowing undesired reachability across CEs of different VPNs, are then applicable to the IPv6 reachability of the VRF to which the CE attaches.

11. Security Considerations

The extensions defined in this document allow MP-BGP to propagate reachability information about IPv6 VPN routes.

Security considerations for the transport of IPv6 reachability information using BGP are discussed in RFC2545, Section 5, and are equally applicable for the extensions described in this document.

The extensions described in this document for offering IPv6 VPNs use the exact same approach as the approach described in [BGP/MPLS-VPN]. As such, the same security considerations apply with regards to Data Plane security, Control Plane security, and PE and P device security as described in [BGP/MPLS-VPN], Section 13.

12. Quality of Service

Since all the QoS mechanisms discussed for IPv4 VPNs in Section 14 of [BGP/MPLS-VPN] operate in the same way for IPv4 and IPv6 (Diffserv, Intserv, MPLS Traffic Engineering), the QoS considerations discussed in [BGP/MPLS-VPN] are equally applicable to IPv6 VPNs (and this holds whether IPv4 tunneling or IPv6 tunneling is used in the backbone.)

13. Scalability

Each of the scalability considerations summarized for IPv4 VPNs in Section 15 of [BGP/MPLS-VPN] is equally applicable to IPv6 VPNs.

14. IANA Considerations

This document specifies (see Section 3.2) the use of the BGP AFI (Address Family Identifier) value 2, along with the BGP SAFI (Subsequent Address Family Identifier) value 128, to represent the address family "VPN-IPv6 Labeled Addresses", which is defined in this document.

The use of AFI value 2 for IPv6 is as currently specified in the IANA registry "Address Family Identifier", so IANA need not take any action with respect to it.

The use of SAFI value 128 for "MPLS-labeled VPN address" is as currently specified in the IANA registry "Subsequence Address Family Identifier", so IANA need not take any action with respect to it.

15. Acknowledgements

We would like to thank Gerard Gastaud and Eric Levy-Abegnoli, who contributed to this document.

In Memoriam

The authors would like to acknowledge the valuable contribution to this document from Tri T. Nguyen, who passed away in April 2002 after a sudden illness.

16. References

16.1. Normative References

- [BGP/MPLS-VPN] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [BGP-EXTCOM] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.
- [BGP-MP] Bates, T., Rekhter, Y., Chandra, R., and D. Katz, "Multiprotocol Extensions for BGP-4", RFC 2858, June 2000.
- [IPv6] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [MPLS-BGP] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, May 2001.
- [BGP-CAP] Chandra, R. and J. Scudder, "Capabilities Advertisement with BGP-4", RFC 3392, November 2002.
- [LDP] Andersson, L., Doolan, P., Feldman, N., Fredette, A., and B. Thomas, "LDP Specification", RFC 3036, January 2001.
- [BGP-IPv6] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, March 1999.

16.2. Informative References

- [V6ADDR] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [UNIQUE-LOCAL] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [2547-GRE/IP] Rekhter and Rosen, "Use of PE-PE GRE or IP in RFC2547 VPNs", Work in Progress.
- [2547-IPsec] Rosen, De Clercq, Paridaens, T'Joens, Sargor, "Use of PE-PE IPsec in RFC2547 VPNs", Work in Progress, August 2005.

- [RSVP-TE] 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.
- [MPLS-in-IP/GRE] Worster, T., Rekhter, Y., and E. Rosen, "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, March 2005.
- [MPLS-in-L2TPv3] Townsley, M., et al., "Encapsulation of MPLS over Layer-2 Tunneling Protocol Version 3", Work in Progress, February 2006.
- [BGP] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

Authors' Addresses

Jeremy De Clercq
Alcatel
Copernicuslaan 50, 2018 Antwerpen, Belgium

EMail: jeremy.de_clercq@alcatel.be

Dirk Ooms
OneSparrow
Belegstraat 13, 2018 Antwerpen, Belgium

EMail: dirk@onesparrow.com

Marco Carugi
Nortel Networks S.A.
Parc d'activites de Magny-Les Jeunes Bois CHATEAUFORT
78928 YVELINES Cedex 9 - France

EMail: marco.carugi@nortel.com

Francois Le Faucheur
Cisco Systems, Inc.
Village d'Entreprise Green Side - Batiment T3
400, Avenue de Roumanille
06410 Biot-Sophia Antipolis
France

EMail: flefauch@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

