Internet Engineering Task Force (IETF)
Request for Comments: 6565
Category: Standards Track
ISSN: 2070-1721

P. Pillay-Esnault
Cisco Systems
P. Moyer
Pollere, Inc.
J. Doyle
Jeff Doyle and Associates
E. Ertekin
M. Lundberg
Booz Allen Hamilton

June 2012

OSPFv3 as a Provider Edge to Customer Edge (PE-CE) Routing Protocol

Abstract

Many Service Providers (SPs) offer Virtual Private Network (VPN) services to their customers using a technique in which Customer Edge (CE) routers are routing peers of Provider Edge (PE) routers. The Border Gateway Protocol (BGP) is used to distribute the customer's routes across the provider's IP backbone network, and Multiprotocol Label Switching (MPLS) is used to tunnel customer packets across the provider's backbone. Support currently exists for both IPv4 and IPv6 VPNs; however, only Open Shortest Path First version 2 (OSPFv2) as PE-CE protocol is specified. This document extends those specifications to support OSPF version 3 (OSPFv3) as a PE-CE routing protocol. The OSPFv3 PE-CE functionality is identical to that of OSPFv2 except for the differences described in this document.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc6565.

Table of Contents

1. Introduction

[RFC4364] offers Service Providers (SPs) a method for providing Layer 3 Virtual Private Network (VPN) services to subtending customer networks. Using the procedures defined in [RFC4364], Provider Edge (PE) routers separate customer VPN routing information into Virtual Routing and Forwarding (VRF) tables. The Border Gateway Protocol (BGP) is used to disseminate customer network VPN routes between PE VRFs configured in the same VPN.

The initial BGP/MPLS IP VPN specification enabled PE routers to learn routes within customer sites through static routing, or through a dynamic routing protocol instantiated on the PE-CE link. Specifically, [RFC4364] (and its predecessor, [RFC2547]) included support for dynamic routing protocols such as BGP, RIP, and OSPFv2. The OSPFv2 as the Provider/Customer Edge Protocol specification [RFC4577] further updates the operation of OSPFv2 as the PE-CE routing protocol by detailing additional extensions to enable intra-domain routing connectivity between OSPFv2-based customer sites.

While [RFC4364] was defined for IPv4-based networks, [RFC4659] extends support to IPv6 VPNs. It is expected that OSPFv3 will be used as the IGP for some IPv6 VPNs just as the OSPFv2 was used for IPv4 VPNs. The advantages of using OSPFv3 as a PE-CE protocol are the same as for the IPv4 VPN deployment.

This document defines the mechanisms required to enable the operation of OSPFv3 as the PE-CE routing protocol. In doing so, it reuses, and extends where necessary, methods defined in [RFC4659] and [RFC4577]. This document also includes the specifications for maintaining intra-domain routing connectivity between OSPFv3-based customer sites across an SP backbone.

We presuppose familiarity with the contents of [RFC4364], [RFC4659], [RFC4577], [RFC4576], [RFC5340], and [RFC2328].

2. Specification of Requirements

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. Requirements

The benefits and considerations associated with deploying OSPFv3 as the PE-CE routing protocol are similar to those described in [RFC4577]. The requirements described in Section 3 of [RFC4577] remain semantically identical for the deployment of OSPFv3.

[RFC5340] describes the modifications required to OSPF to support IPv6. In that specification, many of the fundamental mechanisms associated with OSPFv2 remain unchanged for OSPFv3. Consequently, the operation of OSPFv3 as the PE-CE routing protocol is very similar to OSPFv2 as the PE-CE protocol.

3.1. OSPFv3 Specificities

Section 2 of [RFC5340] describes the differences between OSPFv3 and OSPFv2. Several of these changes will require modifications to the architecture described in [RFC4577]. These differences and their corresponding impact to [RFC4577] are described below:

New LSA types:

For an IPv6 VPN architecture where customers interface with providers through OSPFv3, traditional BGP/OSPF interactions specify that VPN-IPv6 reachability information redistributed into OSPFv3 will be expressed as AS-External OSPFv3 LSAs. Instead, it may be desirable to view these LSAs as inter-area-prefix LSAs. The OSPF Route Type Extended Communities attribute defined in [RFC4577] is extended to include OSPFv3 route types. These new encodings are defined in Section 4.4.

Multiple instances over a link:

OSPFv3 operates on a per-link basis as opposed to OSPFv2, which operates on a per-IP-subnet basis. The support of multiple OSPFv3 protocol instances on a link changes the architecture described in [RFC4577]. [RFC4577] specifies that each interface belongs to no more than one OSPF instance. For OSPFv3, multiple instances can be established over a single interface and associated with the same VRF.

In addition to establishing multiple OSPFv3 instances over a single PE-CE link, multiple OSPFv3 instances can also be established across a sham link. This enables multiple OSPFv3 instances associated with a VRF to independently establish intra-area connectivity to other OSPFv3 instances attached to a remote PE VRF. Support for multiple OSPFv3 instances across the sham link is described in Section 5.

4. BGP/OSPFv3 Interaction Procedures for PE Routers

4.1. VRFs and OSPFv3 Instances

The relationship between VRFs, interfaces, and OSPFv3 instances on a PE router is described in the following section.

As defined in [RFC4364], a PE router can be configured with one or more VRFs. Each VRF configured on the PE corresponds to a customer VPN and retains the destinations that are reachable within that VPN. Each VRF may be associated with one or more interfaces, which allows multiple sites to participate in the same VPN. If OSPFv3 is instantiated on an interface associated with a VRF, the VRF will be populated with OSPFv3 routing information.

As OSPFv3 supports multiple instances on a single interface, it is therefore possible that multiple customer sites can connect to the same interface of a PE router (e.g., through a Layer 2 switch) using distinct OSPFv3 instances. A PE interface can be associated with only one VRF, and all OSPFv3 instances running on the same interface MUST be associated with the same VRF. Configurations where a PE interface is associated with multiple VRFs are out of scope for this document.

4.1.1. Independent OSPFv3 Instances in PEs

Similar to [RFC4577], the PE must associate at least one OSPFv3 instance for each OSPFv3 domain to which it attaches, and each instance of OSPFv3 MUST be associated with a single VRF.

The support of multiple PE-CE OSPFv3 instances per PE interface does not change the paradigm that an OSPF instance can be associated with only a single VRF. Furthermore, for each instance instantiated on the interface, the PE establishes adjacencies with corresponding CEs associated with the instance. Note that although multiple instances may populate a common VRF, they do not leak routes to one another, unless configured to do so.

4.1.2. OSPFv3 Domain Identifier

The OSPFv3 Domain ID describes the administrative domain of the OSPF instance that originated the route. It has an AS-wide significance and is one of the parameters used to determine whether a VPN-IPv6 route should be translated as an Inter-area-prefix LSA or External LSA. Each OSPFv3 instance MUST have a primary Domain ID that is transported along with the VPN-IPv6 route in a BGP attribute over the VPN backbone. Each OSPFv3 instance may have a set of secondary Domain IDs that applies to other OSPFv3 instances within its administrative domain.

The primary Domain ID may either be configured or be set to a value of NULL. The secondary Domain IDs are only allowed if a non-NULL primary Domain ID is configured. The Domain ID MUST be configured on a per-OSPFv3 instance basis.

The Domain ID is used to determine whether an incoming VPN-IPv6 route belongs to the same domain as the receiving OSPFv3 instance. An incoming VPN-IPv6 route is said to belong to the same domain if a non-NULL incoming Domain ID matches either the local primary or one of the secondary Domain IDs. If the local Domain ID and incoming Domain ID are NULL, it is considered a match.

4.2. OSPFv3 Areas

Sections 4.1.4 and 4.2.3 of [RFC4577] describe the characteristics of a PE router within an OSPFv2 domain. The mechanisms and expected behavior described in [RFC4577] are applicable to an OSPFv3 domain.

4.3. VRFs and Routes

From the perspective of the CE, the PE appears as any other OSPFv3 neighbor. There is no requirement for the CE to support any mechanisms of IPv6 BGP/MPLS VPNs or for the CE to have any awareness of the VPNs, thereby enabling any OSPFv3 implementation to be used on a CE.

Because the export and import policies might cause different routes to be installed in different VRFs of the same OSPFv3 domain, the VPN backbone cannot be considered as a single router from the perspective of the domain's CEs. Rather, each CE should view its connected PE as a separate router.

The PE uses OSPFv3 to distribute routes to CEs, and MP-BGP [RFC4760] to distribute VPN-IPv6 routes to other (remote) PE routers as defined in [RFC4659]. An IPv6 prefix installed in the VRF by OSPFv3 is changed to a VPN-IPv6 prefix by the addition of an 8-octet Route Distinguisher (RD) as discussed in Section 2 of [RFC4659]. This VPN-IPv6 route can then be redistributed into MP-BGP according to an export policy that adds a Route Target (RT) Extended Communities attribute to the Network Layer Reachability Information (NLRI) [RFC4360].

Domain IDs are used to distinguish between OSPFv3 instances. When an OSPFv3 distributed route is redistributed into MP-BGP, the Domain ID, OSPFv3 Router ID, Area, OSPFv3 Route Type, and Options fields (External Route Type) are also carried in Extended Community Attributes of the MP-BGP route.

A PE receiving a VPN-IPv6 NLRI from MP-BGP uses an import policy to determine, based on the RT, whether the route is eligible to be installed in one of its local VRFs. The BGP decision process selects

which of the eligible routes are to be installed in the associated VRF, and the selected set of VPN-IPv6 routes are converted into IPv6 routes by removing the RD before installation.

An IPv6 route learned from MP-BGP and installed in a VRF might or might not be redistributed into OSPFv3, depending on the local configuration. For example, the PE might be configured to advertise only a default route to CEs of a particular OSPFv3 instance. Further, if the route is to be redistributed into multiple OSPFv3 instances, the route might be advertised using different LSA types in different instances.

If an IPv6 route learned from MP-BGP is to be redistributed into a particular OSPFv3 instance, the OSPF Domain Identifier Extended Communities attribute of the VPN-IPv6 route is used to determine whether the OSPFv3 instance from which the route was learned is the same as the OSPFv3 instance into which the route is to be redistributed.

4.3.1. OSPFv3 Routes on PEs

VRFs may be populated by both OSPFv3 routes from a CE or VPN-IPv6 routes from other PEs via MP-BGP. OSPFv3 routes are installed in a VRF using the OSPFv3 decision process. They may be redistributed into BGP and disseminated to other PEs participating in the VPN. At these remote PEs, the VPN-IPv6 routes may be imported into a VRF and redistributed into the OSPFv3 instance(s) associated with that VRF.

As specified in [RFC4659], routes imported and exported into a VRF are controlled by the Route Target (RT) Extended Communities attribute. OSPFv3 routes that are redistributed into BGP are given an RT that corresponds to the VRF. This RT is examined at remote PEs. In order to import a route, a VRF must have an import RT that is identical to the route’s RT. For routes that are eligible to be imported into the VRF, the standard BGP decision process is used to choose the "best" route(s).

When a route is advertised from a CE to a PE via OSPFv3 and that route is installed in the VRF associated with the CE, the route is advertised to other locally attached CEs under normal OSPFv3 procedures.

The route is also redistributed into MP-BGP to be advertised to remote PEs. The information necessary for accurate redistribution back into OSPFv3 by the remote PEs is carried in the OSPF Route Type, OSPF Domain ID, and OSPF Router ID Extended Communities attributes (Section 4.4). The relevant local OSPFv3 information encoded into these attributes are as follows:

The Area ID of the PE-CE link.

The Route Type, as determined by the LSA type from which the route was learned.

The Options fields (External metric-type).

The Domain ID of the OSPFv3 process. If no Domain ID is configured, the NULL identifier is used.

The PE's Router ID associated with the OSPFv3 instance.

A Multi-Exit-Discriminator (MED) attribute SHOULD also be set to the value of the OSPFv3 metric associated with the route plus 1, when the OSPFv3 route is redistributed into the MP-BGP.

4.3.2. VPN-IPv6 Routes Received from MP-BGP

When a PE receives a valid VPN-IPv6 route from MP-BGP and has identified an association with a local VRF, it must determine:

Whether a route to the corresponding IPv6 prefix is to be installed in the VRF;

Whether the installed IPv6 route is to be redistributed to one or more local OSPFv3 instances; and

What OSPFv3 LSA type is to be used when advertising the route into each OSPFv3 instance.

An IPv6 route derived from a received VPN-IPv6 route is not installed in the associated local VRF if:

The BGP decision process identifies a better route to the destination NLRI; or

A configured import policy prohibits the installation of the route.

The PE advertises the IPv6 route learned from MP-BGP to attached CEs via OSPFv3 if:

No configured filtering prohibits redistributing the route to OSPFv3;

No configured policy blocks the route in favor of a less-specific summary route; and

Redistribution of a BGP learned IPv6 route into OSPF is based on local policy.

The subsequent sections discuss the advertisement of routes learned from MP-BGP and the rules for determining to which LSA types and to which CEs to advertise the routes.

When the PE sends an LSA to a CE, it sets the DN-bit in the LSA to prevent looping. The DN-bit is discussed in Section 4.5.1.

4.3.2.1. OSPF Inter-Area Routes

A PE advertises an IPv6 route using an Inter-Area-Prefix (type 0x2003) LSA under the following circumstances:

The OSPFv3 domain from which the IPv6 route was learned is the same (as determined by the Domain ID) as the domain of the OSPFv3 instance into which it is to be redistributed; and

The IPv6 route was advertised to a remote PE in an Intra-Area-Prefix (type 0x2009) OR an Inter-Area-Prefix (type 0x2003) LSA.

Note that under these rules, the PE represents itself as an Area Border Router (ABR) regardless of whether or not the route is being advertised into the same area number from which the remote PE learned it (that is, whether the VPN-IPv6 route carries the same or different area numbers).

4.3.2.2. OSPF Intra-Area Route

A route is advertised as an intra-area route using an Intra-Area-Prefix (type 0x2009) LSA only when sham links are used, as described in Section 5. Otherwise, routes are advertised as either inter-area (Section 4.3.2.1) or external / Not-So-Stubby Area (NSSA) (Section 4.3.2.3) routes.

4.3.2.3. OSPF External Routes and NSSA Routes

A PE considers an IPv6 route to be external under the following circumstances:

The OSPFv3 domain from which the route was learned is different (as determined by the Domain ID) from the domain of the OSPFv3 instance into which it is redistributed; or

The OSPFv3 domain from which the route was learned is the same as the domain of the OSPFv3 instance into which it is redistributed, AND it was advertised to the remote PE in an AS-External-LSA (type 0x4005) or an NSSA-LSA (type 0x2007); or

The route was not learned from an OSPFv3 instance.

To determine if the learned route is from a different domain, the Domain ID associated with the VPN-IPv6 route (in the OSPF Domain ID Extended Communities attribute or attributes) is compared with the local OSPFv3 Domain ID, if configured. Compared Domain IDs are considered identical if:

1. All 8 bytes are identical; or

2. Both Domain IDs are NULL (all zeroes).

Note that if the VPN-IPv6 route does not have a Domain ID in its attributes, or if the local OSPFv3 instance does not have a configured Domain ID (i.e., in either case), the route is considered to have a NULL Domain ID.

An IPv6 route that is determined to be external might or might not be advertised to a connected CE, depending on the type of area to which the PE-CE link belongs and whether there is a configured policy restricting its advertisement.

If there are multiple external routes to the same prefix, the standard OSPFv3 decision process is used to select the "best" route.

If the external route is to be advertised and the area type of the PE-CE link is NSSA, the PE advertises the route in an NSSA-LSA (type 0x2007); otherwise, the external route is advertised in an AS-External-LSA (type 0x4005).

The DN-bit of the LSA advertising the external route MUST be set, as described in Section 4.5.1.

If the VPN-IPv6 route indicates a route Type-1 metric, the PE should advertise the external route with that metric-type; otherwise, the metric-type of the external IPv6 route is set to Type-2 by default. Note that, by default, a PE should advertise an external route with a Type-2 metric if the IPv6 route’s Domain ID is different than the local OSPFv3 instance, unless specified otherwise by local policy.

4.4. BGP Extended Communities Attributes

OSPFv3 routes from one site are translated and delivered transparently to the remote site as BGP VPN-IPv6 routes. The original OSPFv3 routes carry OSPFv3-specific information that needs to be communicated to the remote PE to ensure transparency. BGP Extended Communities are used to carry the needed information to enable the receiving side to reconstruct a database just as in the OSPFv2 case.

All OSPFv3 routes added to the VRF routing table on a PE router are examined to create a corresponding VPN-IPv6 route in BGP. Each of the OSPFv3 routes MUST have the corresponding BGP Extended Communities Attributes that contain and preserve the OSPFv3 information of the original OSPFv3 route. The BGP Extended Communities attributes defined in [RFC4577] are reused for convenience.

OSPF Domain Identifier Extended Communities Attribute

Each OSPFv3 Instance within a VRF MUST have a Domain ID. The Domain ID is configured per OSPFv3 Instance. The OSPFv3 Domain ID is a 6-byte number, and its default value is 0. This attribute has a 2-byte type field, encoded with a value of 0x0005, 0x0105, or 0x0205.

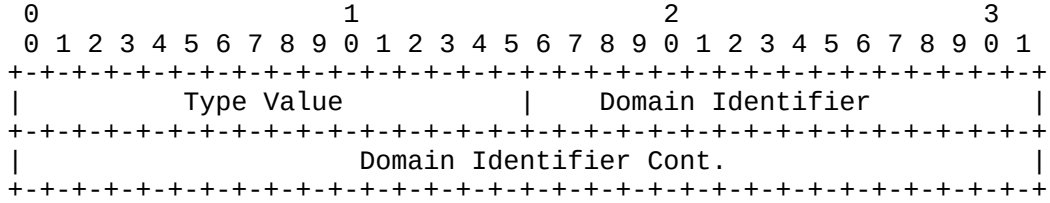

The OSPF Domain Identifier Extended Communities Attribute

OSPFv3 Domain IDs field : 6 bytes

Each OSPFv3 Instance within a VRF MUST have a Domain ID and its default value (if none is configured) is 0. The Domain ID is configured per OSPFv3 Instance.

OSPF Router ID Extended Communities Attribute

The OSPFv3 Router ID is a 32-bit number as in OSPFv2. This attribute has a 2-byte type field, encoded with a value of 0x0107.

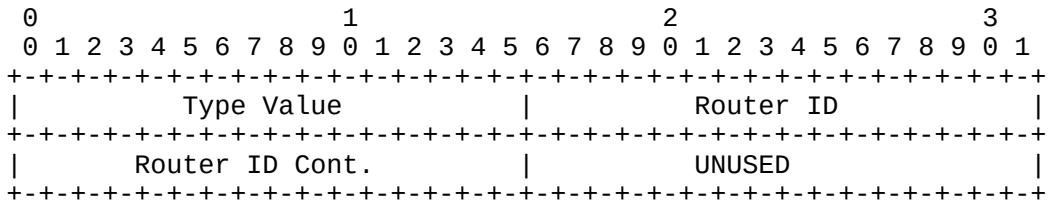

The OSPF Router ID Extended Communities Attribute

OSPFv3 Router ID field : 4 bytes

The OSPFv3 Router ID is a 32-bit number as in OSPFv2. Setting this field is OPTIONAL, and its default value is 0.

OSPF Route Type Extended Communities Attribute

The OSPF Route Type Extended Communities Attribute MUST be present. It contains a 2-byte type field, encoded with a value of 0x0306. The remaining 6 bytes are divided into 3 fields, an Area Number, a Route Type, and an Options field.

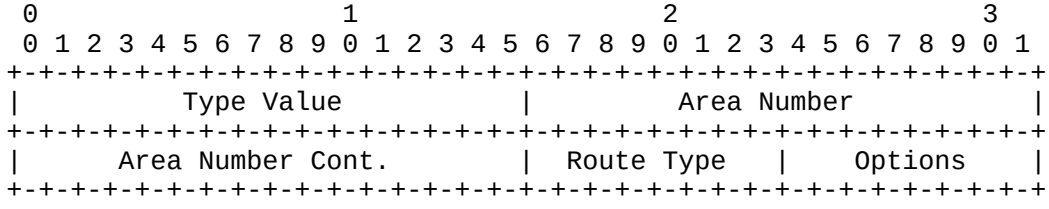

The OSPF Route Type Extended Communities Attribute

Area Number : 4 bytes

The area number indicates the 32-bit Area ID to which the route belongs.

Route Types : 1 byte

To accommodate OSPFv3 LSA types (as registered by [RFC5340]), the Route Type field is encoded as follows:

Route Type Code	Route Type	LSA Type	Description
3	Inter-area-prefix	0x2003	Inter-Area-Prefix-LSA
5	External	0x4005	AS-External-LSA
7	NSSA	0x2007	NSSA-LSA
1 or 2	Intra-area-prefix	0x2009	Intra-Area-Prefix-LSA

Route Type Field Encoding

Options : 1 byte

The Options field indicates the options that are associated with the OSPFv3 route.

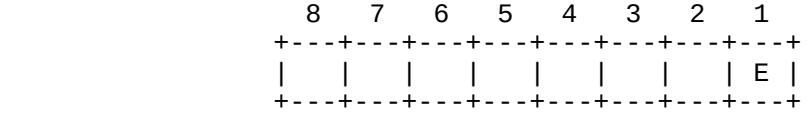

The OSPFv3 Route Options Field

The least significant bit (i.e., bit E) in this field designates the external metric-type. If the bit is clear, the route carries a Type-1 external metric; if the bit is set, the route carries a Type-2 external metric.

4.5. Loop Prevention Techniques

In some topologies, it is possible for routing loops to occur due to the nature and manner of route reachability propagation. One such example is the case of a dual-homed CE router connected to two PEs; those PE routers would receive reachability information both through their CE and their peer PE. As there is transparent transport of OSPFv3 routes over the VPN backbone, it is not possible for the PE routers to determine whether they are within a loop.

The loop scenarios in OSPFv3 topologies are identical to those in the OSPFv2 topologies described in Sections 4.2.5.1 and 4.2.5.2 of [RFC4577]. Of the two loop prevention mechanisms described in the aforementioned sections, only the DN-bit option will be supported in the OSPFv3 implementation.

4.5.1. OSPFv3 Down Bit

[RFC4576] describes the usage of the DN-bit for OSPFv2 and is applicable for OSPFv3 for Inter-area-prefix LSAs, NSSA LSAs, and External LSAs. Similarly, the DN-bit MUST be set in Inter-area-prefix LSAs, NSSA LSAs, and AS-External LSAs, when these are originated from a PE to a CE, to prevent those prefixes from being re-advertised into BGP. As in [RFC4577], any LSA with the DN-bit set must not be used for route calculations on PE routers.

The DN-bit MUST be clear in all other LSA types. The OSPFv3 DN-bit format is described in Appendix A.4.1.1 of [RFC5340].

4.5.2. Other Possible Loops

The mechanism described in Section 4.5.1 of this document is sufficient to prevent looping if the DN-bit information attached to a prefix is preserved in the OSPF domain. As described in Section 4.2.5.3 of [RFC4577], caution must be exercised if mutual redistribution that is performed on a PE causes loss of loop prevention information.

5. OSPFv3 Sham Links

This section modifies the specification of OSPFv2 sham links (defined in Section 4.2.7 of [RFC4577]) to support OSPFv3. Support for OSPFv3 sham links is an OPTIONAL feature of this specification.

A sham link enables a VPN backbone to act as an intra-area link. It is needed when two sites are connected by an intra-area "backdoor" link and the inter-area VPN backbone route would be less preferable due to OSPF route preference rules. The figure below shows the instantiation of a sham link between two VPN sites.

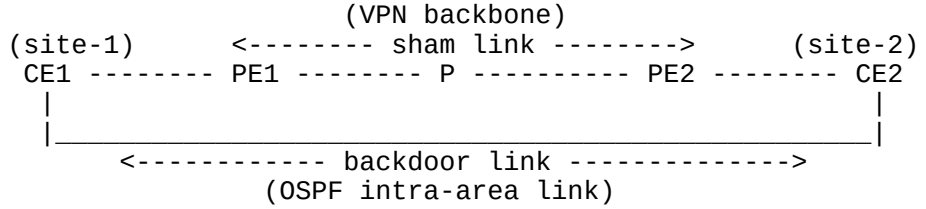

Sham Link

Much of the operation of sham links remains semantically identical to what was previously specified. There are, however, several differences that need to be defined to ensure the proper operation of OSPFv3 sham links.

One of the primary differences between sham links for OSPFv3 and sham links as specified in [RFC4577] is for configurations where multiple OSPFv3 instances populate a VRF. It may be desirable to provide separate intra-area links between these instances over the same sham link. To achieve this, multiple OSPFv3 instances may be established across the PE-PE sham link to provide intra-area connectivity between PE-CE OSPFv3 instances.

Note that even though multiple OSPFv3 instances may be associated with a VRF, a sham link is still thought of as a relation between two VRFs.

Another modification to OSPFv2 sham links is that OSPFv3 sham links are now identified by 128-bit endpoint addresses. Since sham link endpoint addresses are now 128 bits, they can no longer default to the RouterID, which is a 32-bit number. Sham link endpoint addresses MUST be configured.

Sham link endpoint addresses MUST be distributed by BGP as routeable VPN IPv6 addresses, each with an IPv6 address prefix that is 128 bits long. As specified in Section 4.2.7.1 of [RFC4577], these endpoint addresses MUST NOT be advertised by OSPFv3; if there is no BGP route to the sham link endpoint address, that address is to appear unreachable, so that the sham link appears to be down.

If there is a BGP route to the remote sham link endpoint address, the sham link appears to be up. Conversely, if there is no BGP route to the sham link endpoint address, the sham link appears to be down.

5.1. Creating a Sham Link

The procedures for creating an OSPFv3 sham link are identical to those specified in Section 4.2.7.2 of [RFC4577]. Note that the creation of OSPFv3 sham links requires the configuration of both local and remote 128-bit sham link endpoint addresses. The local sham link endpoint address associated with a VRF MAY be used by all OSPFv3 instances that are attached to that VRF. The OSPFv3 PE-PE "link" Instance ID in the protocol packet header is used to demultiplex multiple OSPFv3 instance protocol packets exchanged over the sham link.

5.2. OSPF Protocol on Sham Link

Much of the operation of OSPFv3 over a sham link is semantically the same as the operation of OSPFv2 over a sham link, as described in Section 4.2.7.3 of [RFC4577]. This includes the methodology for sending and receiving OSPFv3 packets over sham links, as well as

Hello/Router Dead Intervals. Furthermore, the procedures associated with the assignment of sham link metrics adhere to those set forth for OSPFv2. OSPFv3 sham links are treated as on-demand circuits.

Although the operation of the OSPFv3 protocol over the sham link is the same as OSPFv2, multiple OSPFv3 instances may be instantiated across this link. By instantiating multiple instances across the sham link, distinct intra-area connections can be established between PE-PE OSPFv3 instances associated with the endpoint addresses.

For example, if two OSPFv3 instances (O1, O2) attach to a VRF V1, and on a remote PE, two other OSPFv3 instances (O3, O4) attach to a VRF V2, it may be desirable to connect O1 and O3 with an intra-area link, and O2 and O4 with an intra-area link. This can be accomplished by instantiating two OSPFv3 instances across the sham link, which connects V1 and V2. O1 and O3 can be mapped to one of the sham link OSPFv3 instances; O2 and O4 can be mapped to the other sham link OSPFv3 instance.

5.3. OSPF Packet Forwarding on Sham Link

The rules associated with route redistribution, stated in Section 4.2.7.4 of [RFC4577], remain unchanged in this specification. Specifically:

> If the next-hop interface for a particular route is a sham link, then the PE SHOULD NOT redistribute that route into BGP as a VPN-IPv6 route.
>
> Any other route advertised in an LSA that is transmitted over a sham link MUST also be redistributed (by the PE flooding the LSA over the sham link) into BGP.

When redistributing these LSAs into BGP, they are encoded with the BGP Extended Communities Attributes, as defined in Section 4.4 of this document.

When forwarding a packet, if the preferred route for that packet has the sham link as its next-hop interface, then the packet MUST be forwarded according to the corresponding BGP route (as defined in [RFC4364] and [RFC4659]).

6. Multiple Address Family Support

The support of multiple address families (AFs) in OSPFv3 is described in [RFC5838]. [RFC5838] differentiates between AFs by using reserved ranges of Instance IDs for each AF.

The architecture described in this document is fully compatible with [RFC5838]. The OSPFv3 PE-CE protocol can support multiple address families across a VPN backbone. All AFs redistributed from OSPFv3 into BGP on a PE MUST contain the BGP Extended Communities Attributes as described in Section 4.4.

7. Security Considerations

The extensions described in this document are specific to the use of OSPFv3 as the PE-CE protocol and do not introduce any new security concerns other than those already defined in Section 6 of [RFC4577].

8. IANA Considerations

An early version of this document resulted in the allocation of OSPFv3 Route Attributes (0x0004) entry in the BGP IPv6 Address Specific Extended Community. This allocation is no longer required. IANA has marked the OSPFv3 Route Attributes (0x0004) entry in the BGP IPv6 Address Specific Extended Community registry as deprecated. The BGP Extended Communities Attributes in this document have already been registered by IANA.

9. Acknowledgments

The authors would like to thank Kelvin Upson, Seiko Okano, Matthew Everett, Dr. Vineet Mehta, Paul Wells, and Marek Karasek for their support of this work. Thanks to Peter Psenak, Abhay Roy, Acee Lindem, Nick Weeds, Robert Hanzl, and Daniel Cohn for their Last Call comments. Special thanks to Stewart Bryant, Stephen Farrel, and Fred Baker for their thorough review.

10. References

10.1. Normative References

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

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, April 1998.

[RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, February 2006.

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

[RFC4576] Rosen, E., Psenak, P., and P. Pillay-Esnault, "Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4576, June 2006.

[RFC4577] Rosen, E., Psenak, P., and P. Pillay-Esnault, "OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4577, June 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.

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

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[RFC5838] Lindem, A., Ed., Mirtorabi, S., Roy, A., Barnes, M., and R. Aggarwal, "Support of Address Families in OSPFv3", RFC 5838, April 2010.

10.2. Informative References

[RFC2547] Rosen, E. and Y. Rekhter, "BGP/MPLS VPNs", RFC 2547, March 1999.

Authors' Addresses

Padma Pillay-Esnault
Cisco Systems
510 McCarty Blvd.
Milpitas, CA 95035
USA

EMail: ppe@cisco.com

Peter Moyer
Pollere, Inc.
325M Sharon Park Drive #214
Menlo Park, CA 94025
USA

EMail: pete@pollere.net

Jeff Doyle
Jeff Doyle and Associates
9878 Teller Ct.
Westminster, CO 80021
USA

EMail: jdoyle@doyleassociates.net

Emre Ertekin
Booz Allen Hamilton
5220 Pacific Concourse Drive
Los Angeles, CA 90045
USA

EMail: ertekin_emre@bah.com

Michael Lundberg
Booz Allen Hamilton
8283 Greensboro Drive
McLean, VA 22102
USA

EMail: lundberg_michael@bah.com