

Internet Engineering Task Force (IETF)
Request for Comments: 9136
Category: Standards Track
ISSN: 2070-1721

J. Rabadan, Ed.
W. Henderickx
Nokia
J. Drake
W. Lin
Juniper
A. Sajassi
Cisco
October 2021

IP Prefix Advertisement in Ethernet VPN (EVPN)

Abstract

The BGP MPLS-based Ethernet VPN (EVPN) (RFC 7432) mechanism provides a flexible control plane that allows intra-subnet connectivity in an MPLS and/or Network Virtualization Overlay (NVO) (RFC 7365) network. In some networks, there is also a need for dynamic and efficient inter-subnet connectivity across Tenant Systems and end devices that can be physical or virtual and do not necessarily participate in dynamic routing protocols. This document defines a new EVPN route type for the advertisement of IP prefixes and explains some use-case examples where this new route type is used.

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 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9136>.

Copyright Notice

Copyright (c) 2021 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 (<https://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
 - 1.1. Terminology
2. Problem Statement
 - 2.1. Inter-Subnet Connectivity Requirements in Data Centers
 - 2.2. The Need for the EVPN IP Prefix Route
3. The BGP EVPN IP Prefix Route
 - 3.1. IP Prefix Route Encoding
 - 3.2. Overlay Indexes and Recursive Lookup Resolution
4. Overlay Index Use Cases
 - 4.1. TS IP Address Overlay Index Use Case
 - 4.2. Floating IP Overlay Index Use Case

- 4.3. Bump-in-the-Wire Use Case
- 4.4. IP-VRF-to-IP-VRF Model
 - 4.4.1. Interface-less IP-VRF-to-IP-VRF Model
 - 4.4.2. Interface-ful IP-VRF-to-IP-VRF with SBD IRB
 - 4.4.3. Interface-ful IP-VRF-to-IP-VRF with Unnumbered SBD IRB
- 5. Security Considerations
- 6. IANA Considerations
- 7. References
 - 7.1. Normative References
 - 7.2. Informative References
- Acknowledgments
- Contributors
- Authors' Addresses

1. Introduction

[RFC7365] provides a framework for Data Center (DC) Network Virtualization over Layer 3 and specifies that the Network Virtualization Edge (NVE) devices must provide Layer 2 and Layer 3 virtualized network services in multi-tenant DCs. [RFC8365] discusses the use of EVPN as the technology of choice to provide Layer 2 or intra-subnet services in these DCs. This document, along with [RFC9135], specifies the use of EVPN for Layer 3 or inter-subnet connectivity services.

[RFC9135] defines some fairly common inter-subnet forwarding scenarios where Tenant Systems (TSs) can exchange packets with TSs located in remote subnets. In order to achieve this, [RFC9135] describes how Media Access Control (MAC) and IPs encoded in TS RT-2 routes are not only used to populate MAC Virtual Routing and Forwarding (MAC-VRF) and overlay Address Resolution Protocol (ARP) tables but also IP-VRF tables with the encoded TS host routes (/32 or /128). In some cases, EVPN may advertise IP prefixes and therefore provide aggregation in the IP-VRF tables, as opposed to propagating individual host routes. This document complements the scenarios described in [RFC9135] and defines how EVPN may be used to advertise IP prefixes. Interoperability between EVPN and Layer 3 Virtual Private Network (VPN) [RFC4364] IP Prefix routes is out of the scope of this document.

Section 2.1 describes the inter-subnet connectivity requirements in DCs. Section 2.2 explains why a new EVPN route type is required for IP prefix advertisements. Sections 3, 4, and 5 will describe this route type and how it is used in some specific use cases.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

AC: Attachment Circuit

ARP: Address Resolution Protocol

BD: Broadcast Domain. As per [RFC7432], an EVI consists of a single BD or multiple BDs. In case of VLAN-bundle and VLAN-based service models (see [RFC7432]), a BD is equivalent to an EVI. In case of a VLAN-aware bundle service model, an EVI contains multiple BDs. Also, in this document, "BD" and "subnet" are equivalent terms.

BD Route Target: Refers to the broadcast-domain-assigned Route Target [RFC4364]. In case of a VLAN-aware bundle service model, all the BD instances in the MAC-VRF share the same Route Target.

BT: Bridge Table. The instantiation of a BD in a MAC-VRF, as per [RFC7432].

CE: Customer Edge

DA: Destination Address

DGW: Data Center Gateway

Ethernet A-D Route: Ethernet Auto-Discovery (A-D) route, as per [RFC7432].

Ethernet NVO Tunnel: Refers to Network Virtualization Overlay tunnels with Ethernet payload. Examples of this type of tunnel are VXLAN or GENEVE.

EVI: EVPN Instance spanning the NVE/PE devices that are participating on that EVPN, as per [RFC7432].

EVPN: Ethernet VPN, as per [RFC7432].

GENEVE: Generic Network Virtualization Encapsulation, as per [RFC8926].

GRE: Generic Routing Encapsulation

GW IP: Gateway IP address

IPL: IP Prefix Length

IP NVO Tunnel: Refers to Network Virtualization Overlay tunnels with IP payload (no MAC header in the payload).

IP-VRF: A Virtual Routing and Forwarding table for IP routes on an NVE/PE. The IP routes could be populated by EVPN and IP-VPN address families. An IP-VRF is also an instantiation of a Layer 3 VPN in an NVE/PE.

IRB: Integrated Routing and Bridging interface. It connects an IP-VRF to a BD (or subnet).

MAC: Media Access Control

MAC-VRF: A Virtual Routing and Forwarding table for MAC addresses on an NVE/PE, as per [RFC7432]. A MAC-VRF is also an instantiation of an EVI in an NVE/PE.

ML: MAC Address Length

ND: Neighbor Discovery

NVE: Network Virtualization Edge

NVO: Network Virtualization Overlay

PE: Provider Edge

RT-2: EVPN Route Type 2, i.e., MAC/IP Advertisement route, as defined in [RFC7432].

RT-5: EVPN Route Type 5, i.e., IP Prefix route, as defined in Section 3.

SBD: Supplementary Broadcast Domain. A BD that does not have any ACs, only IRB interfaces, and is used to provide connectivity among all the IP-VRFs of the tenant. The SBD is only required in IP-VRF-to-IP-VRF use cases (see Section 4.4).

SN: Subnet

TS: Tenant System

VA: Virtual Appliance

VM: Virtual Machine

VNI: Virtual Network Identifier. As in [RFC8365], the term is used as a representation of a 24-bit NVO instance identifier, with the understanding that "VNI" will refer to a VXLAN Network Identifier in VXLAN, or a Virtual Network Identifier in GENEVE, etc., unless it is stated otherwise.

VSID: Virtual Subnet Identifier

VTEP: VXLAN Termination End Point, as per [RFC7348].

VXLAN: Virtual eXtensible Local Area Network, as per [RFC7348].

This document also assumes familiarity with the terminology of [RFC7365], [RFC7432], and [RFC8365].

2. Problem Statement

This section describes the inter-subnet connectivity requirements in DCs and why a specific route type to advertise IP prefixes is needed.

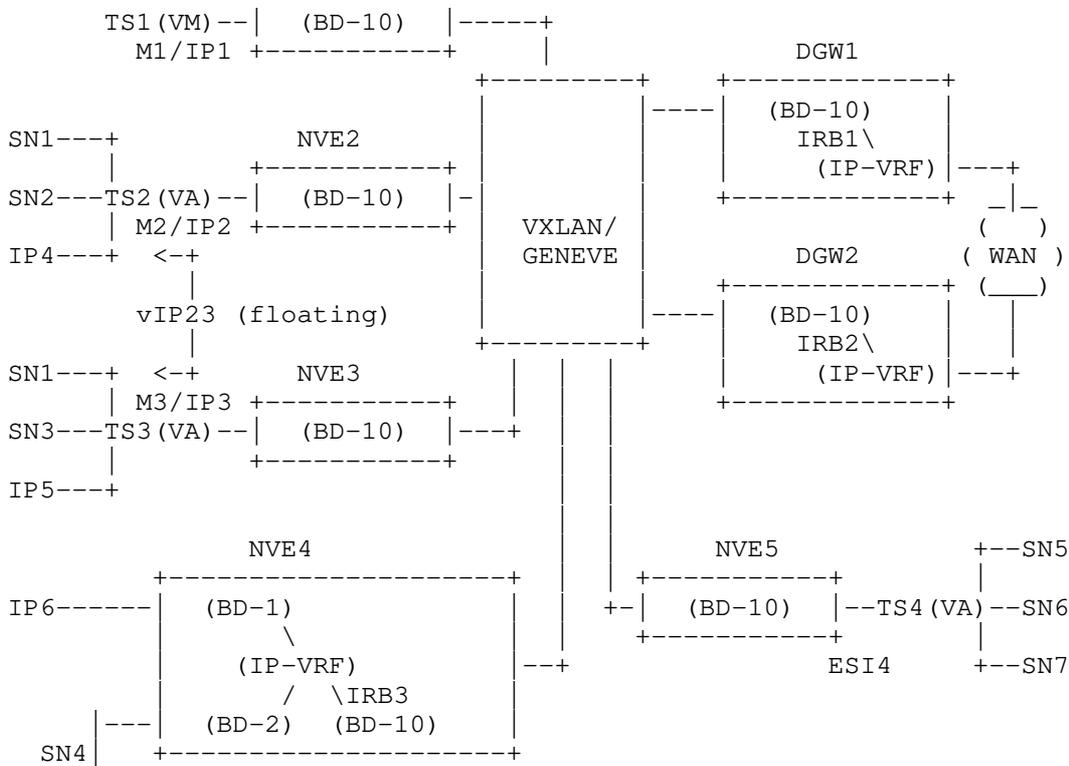
2.1. Inter-Subnet Connectivity Requirements in Data Centers

[RFC7432] is used as the control plane for an NVO solution in DCs, where NVE devices can be located in hypervisors or Top-of-Rack (ToR) switches, as described in [RFC8365].

The following considerations apply to TSs that are physical or virtual systems identified by MAC (and possibly IP addresses) and are connected to BDs by Attachment Circuits:

- * The Tenant Systems may be VMs that generate traffic from their own MAC and IP.
- * The Tenant Systems may be VA entities that forward traffic to/from IP addresses of different end devices sitting behind them.
 - These VAs can be firewalls, load balancers, NAT devices, other appliances, or virtual gateways with virtual routing instances.
 - These VAs do not necessarily participate in dynamic routing protocols and hence rely on the EVPN NVEs to advertise the routes on their behalf.
 - In all these cases, the VA will forward traffic to other TSs using its own source MAC, but the source IP will be the one associated with the end device sitting behind the VA or a translated IP address (part of a public NAT pool) if the VA is performing NAT.
 - Note that the same IP address and endpoint could exist behind two of these TSs. One example of this would be certain appliance resiliency mechanisms, where a virtual IP or floating IP can be owned by one of the two VAs running the resiliency protocol (the Master VA). The Virtual Router Redundancy Protocol (VRRP) [RFC5798] is one particular example of this. Another example is multihomed subnets, i.e., the same subnet is connected to two VAs.
 - Although these VAs provide IP connectivity to VMs and the subnets behind them, they do not always have their own IP interface connected to the EVPN NVE; Layer 2 firewalls are examples of VAs not supporting IP interfaces.

Figure 1 illustrates some of the examples described above.



Note:

ESI4 = Ethernet Segment Identifier 4

Figure 1: DC Inter-subnet Use Cases

Where:

NVE1, NVE2, NVE3, NVE4, NVE5, DGW1, and DGW2 share the same BD for a particular tenant. BD-10 is comprised of the collection of BD instances defined in all the NVEs. All the hosts connected to BD-10 belong to the same IP subnet. The hosts connected to BD-10 are listed below:

- * TS1 is a VM that generates/receives traffic to/from IP1, where IP1 belongs to the BD-10 subnet.
- * TS2 and TS3 are VAs that send/receive traffic to/from the subnets and hosts sitting behind them (SN1, SN2, SN3, IP4, and IP5). Their IP addresses (IP2 and IP3) belong to the BD-10 subnet, and they can also generate/receive traffic. When these VAs receive packets destined to their own MAC addresses (M2 and M3), they will route the packets to the proper subnet or host. These VAs do not support routing protocols to advertise the subnets connected to them and can move to a different server and NVE when the cloud management system decides to do so. These VAs may also support redundancy mechanisms for some subnets, similar to VRRP, where a floating IP is owned by the Master VA and only the Master VA forwards traffic to a given subnet. For example, vIP23 in Figure 1 is a floating IP that can be owned by TS2 or TS3 depending on which system is the Master. Only the Master will forward traffic to SN1.
- * Integrated Routing and Bridging interfaces IRB1, IRB2, and IRB3 have their own IP addresses that belong to the BD-10 subnet too. These IRB interfaces connect the BD-10 subnet to Virtual Routing and Forwarding (IP-VRF) instances that can route the traffic to other subnets for the same tenant (within the DC or at the other end of the WAN).
- * TS4 is a Layer 2 VA that provides connectivity to subnets SN5, SN6, and SN7 but does not have an IP address itself in the BD-10. TS4 is connected to a port on NVE5 that is assigned to Ethernet Segment Identifier 4 (ESI4).

For a BD to which an ingress NVE is attached, "Overlay Index" is defined as an identifier that the ingress EVPN NVE requires in order to forward packets to a subnet or host in a remote subnet. As an example, vIP23 (Figure 1) is an Overlay Index that any NVE attached to BD-10 needs to know in order to forward packets to SN1. The IRB3 IP address is an Overlay Index required to get to SN4, and ESI4 is an Overlay Index needed to forward traffic to SN5. In other words, the Overlay Index is a next hop in the overlay address space that can be an IP address, a MAC address, or an ESI. When advertised along with an IP prefix, the Overlay Index requires a recursive resolution to find out the egress NVE to which the EVPN packets need to be sent.

All the DC use cases in Figure 1 require inter-subnet forwarding; therefore, the individual host routes and subnets:

- a) must be advertised from the NVEs (since VAs and VMs do not participate in dynamic routing protocols) and
- b) may be associated with an Overlay Index that can be a VA IP address, a floating IP address, a MAC address, or an ESI. The Overlay Index is further discussed in Section 3.2.

2.2. The Need for the EVPN IP Prefix Route

[RFC7432] defines a MAC/IP Advertisement route (also referred to as "RT-2") where a MAC address can be advertised together with an IP address length and IP address (IP). While a variable IP address length might have been used to indicate the presence of an IP prefix in a route type 2, there are several specific use cases in which using this route type to deliver IP prefixes is not suitable.

One example of such use cases is the "floating IP" example described in Section 2.1. In this example, it is necessary to decouple the advertisement of the prefixes from the advertisement of a MAC address of either M2 or M3; otherwise, the solution gets highly inefficient and does not scale.

For example, if 1,000 prefixes are advertised from M2 (using RT-2) and the floating IP owner changes from M2 to M3, 1,000 routes would be withdrawn by M2 and readvertised by M3. However, if a separate route type is used, 1,000 routes can be advertised as associated with the floating IP address (vIP23), and only one RT-2 can be used for advertising the ownership of the floating IP, i.e., vIP23 and M2 in the route type 2. When the floating IP owner changes from M2 to M3, a single RT-2 withdrawal/update is required to indicate the change. The remote DGW will not change any of the 1,000 prefixes associated with vIP23 but will only update the ARP resolution entry for vIP23 (now pointing at M3).

An EVPN route (type 5) for the advertisement of IP prefixes is described in this document. This new route type has a differentiated role from the RT-2 route and addresses the inter-subnet connectivity scenarios for DCs (or NVO-based networks in general) described in this document. Using this new RT-5, an IP prefix may be advertised along with an Overlay Index, which can be a GW IP address, a MAC, or an ESI. The IP prefix may also be advertised without an Overlay Index, in which case the BGP next hop will point at the egress NVE, Area Border Router (ABR), or ASBR, and the MAC in the EVPN Router's MAC Extended Community will provide the inner MAC destination address to be used. As discussed throughout the document, the EVPN RT-2 does not meet the requirements for all the DC use cases; therefore, this EVPN route type 5 is required.

The EVPN route type 5 decouples the IP prefix advertisements from the MAC/IP Advertisement routes in EVPN. Hence:

- a) The clean and clear advertisements of IPv4 or IPv6 prefixes in a Network Layer Reachability Information (NLRI) message without MAC addresses are allowed.

- b) Since the route type is different from the MAC/IP Advertisement route, the current procedures described in [RFC7432] do not need to be modified.
- c) A flexible implementation is allowed where the prefix can be linked to different types of Overlay/Underlay Indexes: overlay IP addresses, overlay MAC addresses, overlay ESIs, underlay BGP next hops, etc.
- d) An EVPN implementation not requiring IP prefixes can simply discard them by looking at the route type value.

The following sections describe how EVPN is extended with a route type for the advertisement of IP prefixes and how this route is used to address the inter-subnet connectivity requirements existing in the DC.

3. The BGP EVPN IP Prefix Route

The BGP EVPN NLRI as defined in [RFC7432] is shown below:

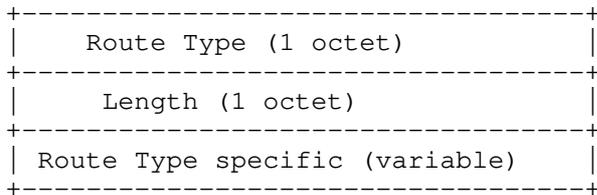


Figure 2: BGP EVPN NLRI

This document defines an additional route type (RT-5) in the IANA "EVPN Route Types" registry [EVPNRouteTypes] to be used for the advertisement of EVPN routes using IP prefixes:

Value: 5
Description: IP Prefix

According to Section 5.4 of [RFC7606], a node that doesn't recognize the route type 5 (RT-5) will ignore it. Therefore, an NVE following this document can still be attached to a BD where an NVE ignoring RT-5s is attached. Regular procedures described in [RFC7432] would apply in that case for both NVEs. In case two or more NVEs are attached to different BDs of the same tenant, they MUST support the RT-5 for the proper inter-subnet forwarding operation of the tenant.

The detailed encoding of this route and associated procedures are described in the following sections.

3.1. IP Prefix Route Encoding

An IP Prefix route type for IPv4 has the Length field set to 34 and consists of the following fields:

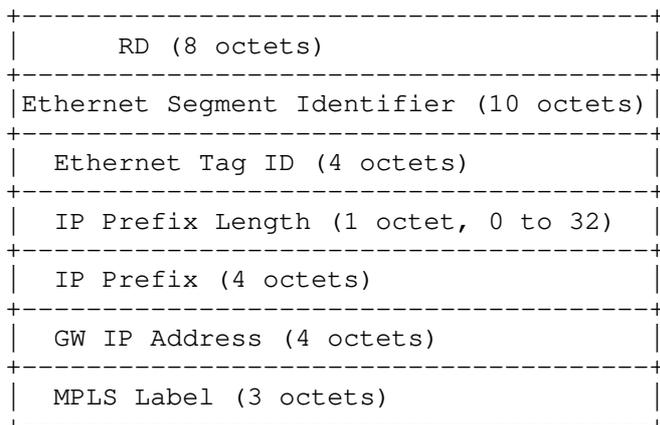


Figure 3: EVPN IP Prefix Route NLRI for IPv4

An IP Prefix route type for IPv6 has the Length field set to 58 and consists of the following fields:

RD (8 octets)
Ethernet Segment Identifier (10 octets)
Ethernet Tag ID (4 octets)
IP Prefix Length (1 octet, 0 to 128)
IP Prefix (16 octets)
GW IP Address (16 octets)
MPLS Label (3 octets)

Figure 4: EVPN IP Prefix Route NLRI for IPv6

Where:

- * The Length field of the BGP EVPN NLRI for an EVPN IP Prefix route MUST be either 34 (if IPv4 addresses are carried) or 58 (if IPv6 addresses are carried). The IP prefix and gateway IP address MUST be from the same IP address family.
- * The Route Distinguisher (RD) and Ethernet Tag ID MUST be used as defined in [RFC7432] and [RFC8365]. In particular, the RD is unique per MAC-VRF (or IP-VRF). The MPLS Label field is set to either an MPLS label or a VNI, as described in [RFC8365] for other EVPN route types.
- * The Ethernet Segment Identifier MUST be a non-zero 10-octet identifier if the ESI is used as an Overlay Index (see the definition of "Overlay Index" in Section 3.2). It MUST be all bytes zero otherwise. The ESI format is described in [RFC7432].
- * The IP prefix length can be set to a value between 0 and 32 (bits) for IPv4 and between 0 and 128 for IPv6, and it specifies the number of bits in the prefix. The value MUST NOT be greater than 128.
- * The IP prefix is a 4- or 16-octet field (IPv4 or IPv6).
- * The GW IP Address field is a 4- or 16-octet field (IPv4 or IPv6) and will encode a valid IP address as an Overlay Index for the IP prefixes. The GW IP field MUST be all bytes zero if it is not used as an Overlay Index. Refer to Section 3.2 for the definition and use of the Overlay Index.
- * The MPLS Label field is encoded as 3 octets, where the high-order 20 bits contain the label value, as per [RFC7432]. When sending, the label value SHOULD be zero if a recursive resolution based on an Overlay Index is used. If the received MPLS label value is zero, the route MUST contain an Overlay Index, and the ingress NVE/PE MUST perform a recursive resolution to find the egress NVE/PE. If the received label is zero and the route does not contain an Overlay Index, it MUST be "treat as withdraw" [RFC7606].

The RD, Ethernet Tag ID, IP prefix length, and IP prefix are part of the route key used by BGP to compare routes. The rest of the fields are not part of the route key.

An IP Prefix route MAY be sent along with an EVPN Router's MAC Extended Community (defined in [RFC9135]) to carry the MAC address that is used as the Overlay Index. Note that the MAC address may be that of a TS.

As described in Section 3.2, certain data combinations in a received route would imply a treat-as-withdraw handling of the route [RFC7606].

3.2. Overlay Indexes and Recursive Lookup Resolution

RT-5 routes support recursive lookup resolution through the use of Overlay Indexes as follows:

- * An Overlay Index can be an ESI or IP address in the address space of the tenant or MAC address, and it is used by an NVE as the next hop for a given IP prefix. An Overlay Index always needs a recursive route resolution on the NVE/PE that installs the RT-5 into one of its IP-VRFs so that the NVE knows to which egress NVE/PE it needs to forward the packets. It is important to note that recursive resolution of the Overlay Index applies upon installation into an IP-VRF and not upon BGP propagation (for instance, on an ASBR). Also, as a result of the recursive resolution, the egress NVE/PE is not necessarily the same NVE that originated the RT-5.
- * The Overlay Index is indicated along with the RT-5 in the ESI field, GW IP field, or EVPN Router's MAC Extended Community, depending on whether the IP prefix next hop is an ESI, an IP address, or a MAC address in the tenant space. The Overlay Index for a given IP prefix is set by local policy at the NVE that originates an RT-5 for that IP prefix (typically managed by the cloud management system).
- * In order to enable the recursive lookup resolution at the ingress NVE, an NVE that is a possible egress NVE for a given Overlay Index must originate a route advertising itself as the BGP next hop on the path to the system denoted by the Overlay Index. For instance:
 - If an NVE receives an RT-5 that specifies an Overlay Index, the NVE cannot use the RT-5 in its IP-VRF unless (or until) it can recursively resolve the Overlay Index.
 - If the RT-5 specifies an ESI as the Overlay Index, a recursive resolution can only be done if the NVE has received and installed an RT-1 (auto-discovery per EVI) route specifying that ESI.
 - If the RT-5 specifies a GW IP address as the Overlay Index, a recursive resolution can only be done if the NVE has received and installed an RT-2 (MAC/IP Advertisement route) specifying that IP address in the IP Address field of its NLRI.
 - If the RT-5 specifies a MAC address as the Overlay Index, a recursive resolution can only be done if the NVE has received and installed an RT-2 (MAC/IP Advertisement route) specifying that MAC address in the MAC Address field of its NLRI.

Note that the RT-1 or RT-2 routes needed for the recursive resolution may arrive before or after the given RT-5 route.

- * Irrespective of the recursive resolution, if there is no IGP or BGP route to the BGP next hop of an RT-5, BGP MUST NOT install the RT-5 even if the Overlay Index can be resolved.
- * The ESI and GW IP fields may both be zero at the same time. However, they MUST NOT both be non-zero at the same time. A route containing a non-zero GW IP and a non-zero ESI (at the same time) SHOULD be treat as withdraw [RFC7606].
- * If either the ESI or the GW IP are non-zero, then the non-zero one is the Overlay Index, regardless of whether the EVPN Router's MAC Extended Community is present or the value of the label. In case the GW IP is the Overlay Index (hence, ESI is zero), the EVPN Router's MAC Extended Community is ignored if present.

* A route where ESI, GW IP, MAC, and Label are all zero at the same time SHOULD be treat as withdraw.

The indirection provided by the Overlay Index and its recursive lookup resolution is required to achieve fast convergence in case of a failure of the object represented by the Overlay Index (see the example described in Section 2.2).

Table 1 shows the different RT-5 field combinations allowed by this specification and what Overlay Index must be used by the receiving NVE/PE in each case. Cases where there is no Overlay Index are indicated as "None" in Table 1. If there is no Overlay Index, the receiving NVE/PE will not perform any recursive resolution, and the actual next hop is given by the RT-5's BGP next hop.

ESI	GW IP	MAC*	Label	Overlay Index
Non-Zero	Zero	Zero	Don't Care	ESI
Non-Zero	Zero	Non-Zero	Don't Care	ESI
Zero	Non-Zero	Zero	Don't Care	GW IP
Zero	Zero	Non-Zero	Zero	MAC
Zero	Zero	Non-Zero	Non-Zero	MAC or None**
Zero	Zero	Zero	Non-Zero	None***

Table 1: RT-5 Fields and Indicated Overlay Index

Table Notes:

- * MAC with "Zero" value means no EVPN Router's MAC Extended Community is present along with the RT-5. "Non-Zero" indicates that the extended community is present and carries a valid MAC address. The encoding of a MAC address MUST be the 6-octet MAC address specified by [IEEE-802.1Q]. Examples of invalid MAC addresses are broadcast or multicast MAC addresses. The route MUST be treat as withdraw in case of an invalid MAC address. The presence of the EVPN Router's MAC Extended Community alone is not enough to indicate the use of the MAC address as the Overlay Index since the extended community can be used for other purposes.
- ** In this case, the Overlay Index may be the RT-5's MAC address or "None", depending on the local policy of the receiving NVE/PE. Note that the advertising NVE/PE that sets the Overlay Index SHOULD advertise an RT-2 for the MAC Overlay Index if there are receiving NVE/PEs configured to use the MAC as the Overlay Index. This case in Table 1 is used in the IP-VRF-to-IP-VRF implementations described in Sections 4.4.1 and 4.4.3. The support of a MAC Overlay Index in this model is OPTIONAL.
- *** The Overlay Index is "None". This is a special case used for IP-VRF-to-IP-VRF where the NVE/PEs are connected by IP NVO tunnels as opposed to Ethernet NVO tunnels.

If the combination of ESI, GW IP, MAC, and Label in the receiving RT-5 is different than the combinations shown in Table 1, the router will process the route as per the rules described at the beginning of this section (Section 3.2).

Table 2 shows the different inter-subnet use cases described in this document and the corresponding coding of the Overlay Index in the route type 5 (RT-5).

=====+

Section	Use Case	Overlay Index in the RT-5
4.1	TS IP address	GW IP
4.2	Floating IP address	GW IP
4.3	"Bump-in-the-wire"	ESI or MAC
4.4	IP-VRF-to-IP-VRF	GW IP, MAC, or None

Table 2: Use Cases and Overlay Indexes for Recursive Resolution

The above use cases are representative of the different Overlay Indexes supported by the RT-5 (GW IP, ESI, MAC, or None).

4. Overlay Index Use Cases

This section describes some use cases for the Overlay Index types used with the IP Prefix route. Although the examples use IPv4 prefixes and subnets, the descriptions of the RT-5 are valid for the same cases with IPv6, except that IP Prefixes, IPL, and GW IP are replaced by the corresponding IPv6 values.

4.1. TS IP Address Overlay Index Use Case

Figure 5 illustrates an example of inter-subnet forwarding for subnets sitting behind VAs (on TS2 and TS3).

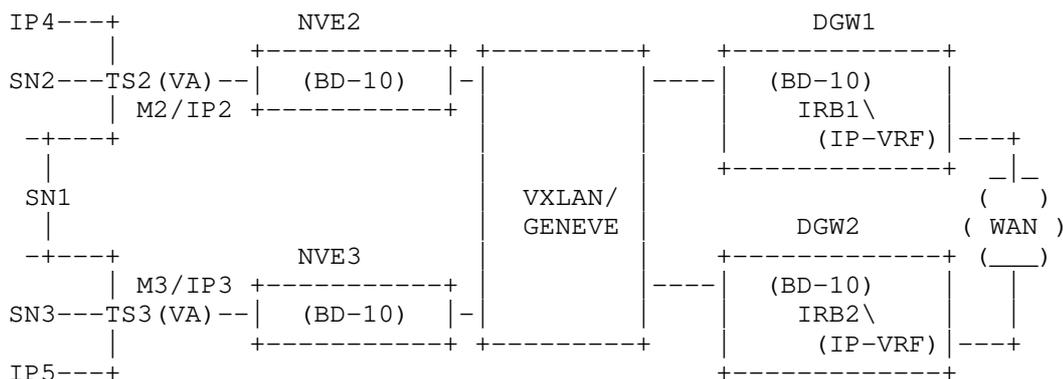


Figure 5: TS IP Address Use Case

An example of inter-subnet forwarding between subnet SN1, which uses a 24-bit IP prefix (written as SN1/24 in the future), and a subnet sitting in the WAN is described below. NVE2, NVE3, DGW1, and DGW2 are running BGP EVPN. TS2 and TS3 do not participate in dynamic routing protocols, and they only have a static route to forward the traffic to the WAN. SN1/24 is dual-homed to NVE2 and NVE3.

In this case, a GW IP is used as an Overlay Index. Although a different Overlay Index type could have been used, this use case assumes that the operator knows the VA's IP addresses beforehand, whereas the VA's MAC address is unknown and the VA's ESI is zero. Because of this, the GW IP is the suitable Overlay Index to be used with the RT-5s. The NVEs know the GW IP to be used for a given prefix by policy.

(1) NVE2 advertises the following BGP routes on behalf of TS2:

- * Route type 2 (MAC/IP Advertisement route) containing: ML = 48 (MAC address length), M = M2 (MAC address), IPL = 32 (IP prefix length), IP = IP2, and BGP Encapsulation Extended Community [RFC9012] with the corresponding tunnel type. The MAC and IP addresses may be learned via ARP snooping.
- * Route type 5 (IP Prefix route) containing: IPL = 24, IP = SN1, ESI = 0, and GW IP address = IP2. The prefix and GW IP

are learned by policy.

(2) Similarly, NVE3 advertises the following BGP routes on behalf of TS3:

- * Route type 2 (MAC/IP Advertisement route) containing: ML = 48, M = M3, IPL = 32, IP = IP3 (and BGP Encapsulation Extended Community).
- * Route type 5 (IP Prefix route) containing: IPL = 24, IP = SN1, ESI = 0, and GW IP address = IP3.

(3) DGW1 and DGW2 import both received routes based on the Route Targets:

- * Based on the BD-10 Route Target in DGW1 and DGW2, the MAC/IP Advertisement route is imported, and M2 is added to the BD-10 along with its corresponding tunnel information. For instance, if VXLAN is used, the VTEP will be derived from the MAC/IP Advertisement route BGP next hop and VNI from the MPLS Label1 field. M2/IP2 is added to the ARP table. Similarly, M3 is added to BD-10, and M3/IP3 is added to the ARP table.
- * Based on the BD-10 Route Target in DGW1 and DGW2, the IP Prefix route is also imported, and SN1/24 is added to the IP-VRF with Overlay Index IP2 pointing at the local BD-10. In this example, it is assumed that the RT-5 from NVE2 is preferred over the RT-5 from NVE3. If both routes were equally preferable and ECMP enabled, SN1/24 would also be added to the routing table with Overlay Index IP3.

(4) When DGW1 receives a packet from the WAN with destination IPx, where IPx belongs to SN1/24:

- * A destination IP lookup is performed on the DGW1 IP-VRF table, and Overlay Index = IP2 is found. Since IP2 is an Overlay Index, a recursive route resolution is required for IP2.
- * IP2 is resolved to M2 in the ARP table, and M2 is resolved to the tunnel information given by the BD FIB (e.g., remote VTEP and VNI for the VXLAN case).
- * The IP packet destined to IPx is encapsulated with:
 - Inner source MAC = IRB1 MAC.
 - Inner destination MAC = M2.
 - Tunnel information provided by the BD (VNI, VTEP IPs, and MACs for the VXLAN case).

(5) When the packet arrives at NVE2:

- * Based on the tunnel information (VNI for the VXLAN case), the BD-10 context is identified for a MAC lookup.
- * Encapsulation is stripped off and, based on a MAC lookup (assuming MAC forwarding on the egress NVE), the packet is forwarded to TS2, where it will be properly routed.

(6) Should TS2 move from NVE2 to NVE3, MAC Mobility procedures will be applied to the MAC route M2/IP2, as defined in [RFC7432]. Route type 5 prefixes are not subject to MAC Mobility procedures; hence, no changes in the DGW IP-VRF table will occur for TS2 mobility -- i.e., all the prefixes will still be pointing at IP2 as the Overlay Index. There is an indirection for, e.g., SN1/24, which still points at Overlay Index IP2 in the routing table, but IP2 will be simply resolved to a different tunnel based on the outcome of the MAC Mobility procedures for the MAC/IP Advertisement route M2/IP2.

Note that in the opposite direction, TS2 will send traffic based on its static-route next-hop information (IRB1 and/or IRB2), and regular EVPN procedures will be applied.

4.2. Floating IP Overlay Index Use Case

Sometimes TSs work in active/standby mode where an upstream floating IP owned by the active TS is used as the Overlay Index to get to some subnets behind the TS. This redundancy mode, already introduced in Sections 2.1 and 2.2, is illustrated in Figure 6.

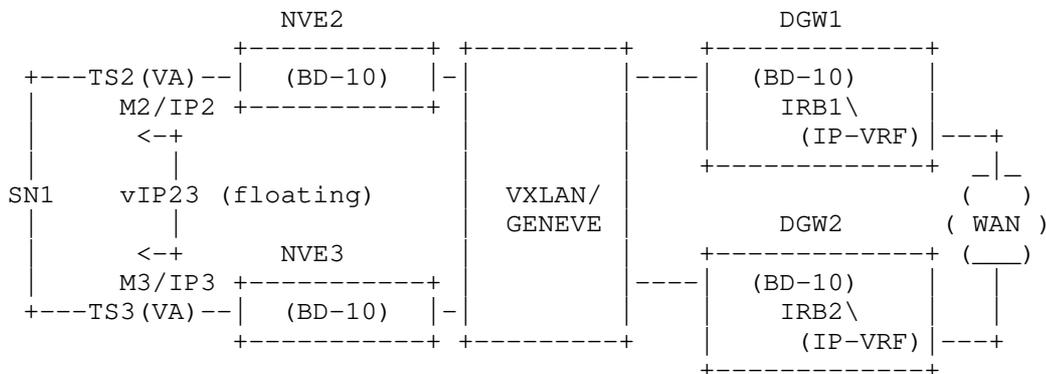


Figure 6: Floating IP Overlay Index for Redundant TS

In this usecase, a GW IP is used as an Overlay Index for the same reasons as in Section 4.1. However, this GW IP is a floating IP that belongs to the active TS. Assuming TS2 is the active TS and owns vIP23:

- (1) NVE2 advertises the following BGP routes for TS2:
 - * Route type 2 (MAC/IP Advertisement route) containing: ML = 48, M = M2, IPL = 32, and IP = vIP23 (as well as BGP Encapsulation Extended Community). The MAC and IP addresses may be learned via ARP snooping.
 - * Route type 5 (IP Prefix route) containing: IPL = 24, IP = SN1, ESI = 0, and GW IP address = vIP23. The prefix and GW IP are learned by policy.
- (2) NVE3 advertises the following BGP route for TS3 (it does not advertise an RT-2 for M3/vIP23):
 - * Route type 5 (IP Prefix route) containing: IPL = 24, IP = SN1, ESI = 0, and GW IP address = vIP23. The prefix and GW IP are learned by policy.
- (3) DGW1 and DGW2 import both received routes based on the Route Target:
 - * M2 is added to the BD-10 FIB along with its corresponding tunnel information. For the VXLAN use case, the VTEP will be derived from the MAC/IP Advertisement route BGP next hop and VNI from the VNI field. M2/vIP23 is added to the ARP table.
 - * SN1/24 is added to the IP-VRF in DGW1 and DGW2 with Overlay Index vIP23 pointing at M2 in the local BD-10.
- (4) When DGW1 receives a packet from the WAN with destination IPx, where IPx belongs to SN1/24:
 - * A destination IP lookup is performed on the DGW1 IP-VRF table, and Overlay Index = vIP23 is found. Since vIP23 is an Overlay Index, a recursive route resolution for vIP23 is required.
 - * vIP23 is resolved to M2 in the ARP table, and M2 is resolved

to the tunnel information given by the BD (remote VTEP and VNI for the VXLAN case).

- * The IP packet destined to IPx is encapsulated with:
 - Inner source MAC = IRB1 MAC.
 - Inner destination MAC = M2.
 - Tunnel information provided by the BD FIB (VNI, VTEP IPs, and MACs for the VXLAN case).

(5) When the packet arrives at NVE2:

- * Based on the tunnel information (VNI for the VXLAN case), the BD-10 context is identified for a MAC lookup.
- * Encapsulation is stripped off and, based on a MAC lookup (assuming MAC forwarding on the egress NVE), the packet is forwarded to TS2, where it will be properly routed.

(6) When the redundancy protocol running between TS2 and TS3 appoints TS3 as the new active TS for SN1, TS3 will now own the floating VIP23 and will signal this new ownership using a gratuitous ARP REPLY message (explained in [RFC5227]) or similar. Upon receiving the new owner's notification, NVE3 will issue a route type 2 for M3/vIP23, and NVE2 will withdraw the RT-2 for M2/vIP23. DGW1 and DGW2 will update their ARP tables with the new MAC resolving the floating IP. No changes are made in the IP-VRF table.

4.3. Bump-in-the-Wire Use Case

Figure 7 illustrates an example of inter-subnet forwarding for an IP Prefix route that carries subnet SN1. In this use case, TS2 and TS3 are Layer 2 VA devices without any IP addresses that can be included as an Overlay Index in the GW IP field of the IP Prefix route. Their MAC addresses are M2 and M3, respectively, and are connected to BD-10. Note that IRB1 and IRB2 (in DGW1 and DGW2, respectively) have IP addresses in a subnet different than SN1.

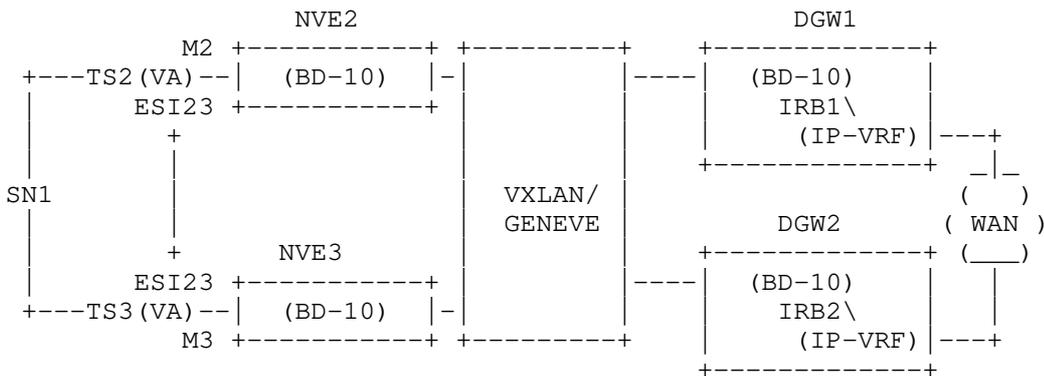


Figure 7: Bump-in-the-Wire Use Case

Since TS2 and TS3 cannot participate in any dynamic routing protocol and neither has an IP address assigned, there are two potential Overlay Index types that can be used when advertising SN1:

- a) an ESI, i.e., ESI23, that can be provisioned on the attachment ports of NVE2 and NVE3, as shown in Figure 7 or
- b) the VA's MAC address, which can be added to NVE2 and NVE3 by policy.

The advantage of using an ESI as the Overlay Index as opposed to the VA's MAC address is that the forwarding to the egress NVE can be done purely based on the state of the AC in the Ethernet segment (notified by the Ethernet A-D per EVI route), and all the EVPN multihoming

redundancy mechanisms can be reused. For instance, the mass withdrawal mechanism described in [RFC7432] for fast failure detection and propagation can be used. It is assumed per this section that an ESI Overlay Index is used in this use case, but this use case does not preclude the use of the VA's MAC address as an Overlay Index. If a MAC is used as the Overlay Index, the control plane must follow the procedures described in Section 4.4.3.

The model supports VA redundancy in a similar way to the one described in Section 4.2 for the floating IP Overlay Index use case, except that it uses the EVPN Ethernet A-D per EVI route instead of the MAC advertisement route to advertise the location of the Overlay Index. The procedure is explained below:

- (1) Assuming TS2 is the active TS in ESI23, NVE2 advertises the following BGP routes:
 - * Route type 1 (Ethernet A-D route for BD-10) containing: ESI = ESI23 and the corresponding tunnel information (VNI field), as well as the BGP Encapsulation Extended Community as per [RFC8365].
 - * Route type 5 (IP Prefix route) containing: IPL = 24, IP = SN1, ESI = ESI23, and GW IP address = 0. The EVPN Router's MAC Extended Community defined in [RFC9135] is added and carries the MAC address (M2) associated with the TS behind which SN1 sits. M2 may be learned by policy; however, the MAC in the Extended Community is preferred if sent with the route.
- (2) NVE3 advertises the following BGP route for TS3 (no AD per EVI route is advertised):
 - * Route type 5 (IP Prefix route) containing: IPL = 24, IP = SN1, ESI = 23, and GW IP address = 0. The EVPN Router's MAC Extended Community is added and carries the MAC address (M3) associated with the TS behind which SN1 sits. M3 may be learned by policy; however, the MAC in the Extended Community is preferred if sent with the route.
- (3) DGW1 and DGW2 import the received routes based on the Route Target:
 - * The tunnel information to get to ESI23 is installed in DGW1 and DGW2. For the VXLAN use case, the VTEP will be derived from the Ethernet A-D route BGP next hop and VNI from the VNI/VSID field (see [RFC8365]).
 - * The RT-5 coming from the NVE that advertised the RT-1 is selected, and SN1/24 is added to the IP-VRF in DGW1 and DGW2 with Overlay Index ESI23 and MAC = M2.
- (4) When DGW1 receives a packet from the WAN with destination IPx, where IPx belongs to SN1/24:
 - * A destination IP lookup is performed on the DGW1 IP-VRF table, and Overlay Index = ESI23 is found. Since ESI23 is an Overlay Index, a recursive route resolution is required to find the egress NVE where ESI23 resides.
 - * The IP packet destined to IPx is encapsulated with:
 - Inner source MAC = IRB1 MAC.
 - Inner destination MAC = M2 (this MAC will be obtained from the EVPN Router's MAC Extended Community received along with the RT-5 for SN1). Note that the EVPN Router's MAC Extended Community is used in this case to carry the TS's MAC address, as opposed to the MAC address of the NVE/PE.
 - Tunnel information for the NVO tunnel is provided by the

Ethernet A-D route per EVI for ESI23 (VNI and VTEP IP for the VXLAN case).

(5) When the packet arrives at NVE2:

- * Based on the tunnel demultiplexer information (VNI for the VXLAN case), the BD-10 context is identified for a MAC lookup (assuming a MAC-based disposition model [RFC7432]), or the VNI may directly identify the egress interface (for an MPLS-based disposition model, which in this context is a VNI-based disposition model).
- * Encapsulation is stripped off and, based on a MAC lookup (assuming MAC forwarding on the egress NVE) or a VNI lookup (in case of VNI forwarding), the packet is forwarded to TS2, where it will be forwarded to SN1.

(6) If the redundancy protocol running between TS2 and TS3 follows an active/standby model and there is a failure, TS3 is appointed as the new active TS for SN1. TS3 will now own the connectivity to SN1 and will signal this new ownership. Upon receiving the new owner's notification, NVE3's AC will become active and issue a route type 1 for ESI23, whereas NVE2 will withdraw its Ethernet A-D route for ESI23. DGW1 and DGW2 will update their tunnel information to resolve ESI23. The inner destination MAC will be changed to M3.

4.4. IP-VRF-to-IP-VRF Model

This use case is similar to the scenario described in Section 9.1 of [RFC9135]; however, the new requirement here is the advertisement of IP prefixes as opposed to only host routes.

In the examples described in Sections 4.1, 4.2, and 4.3, the BD instance can connect IRB interfaces and any other Tenant Systems connected to it. EVPN provides connectivity for:

1. Traffic destined to the IRB or TS IP interfaces, as well as
2. Traffic destined to IP subnets sitting behind the TS, e.g., SN1 or SN2.

In order to provide connectivity for (1), MAC/IP Advertisement routes (RT-2) are needed so that IRB or TS MACs and IPs can be distributed. Connectivity type (2) is accomplished by the exchange of IP Prefix routes (RT-5) for IPs and subnets sitting behind certain Overlay Indexes, e.g., GW IP, ESI, or TS MAC.

In some cases, IP Prefix routes may be advertised for subnets and IPs sitting behind an IRB. This use case is referred to as the "IP-VRF-to-IP-VRF" model.

[RFC9135] defines an asymmetric IRB model and a symmetric IRB model based on the required lookups at the ingress and egress NVE. The asymmetric model requires an IP lookup and a MAC lookup at the ingress NVE, whereas only a MAC lookup is needed at the egress NVE; the symmetric model requires IP and MAC lookups at both the ingress and egress NVE. From that perspective, the IP-VRF-to-IP-VRF use case described in this section is a symmetric IRB model.

Note that in an IP-VRF-to-IP-VRF scenario, out of the many subnets that a tenant may have, it may be the case that only a few are attached to a given IP-VRF of the NVE/PE. In order to provide inter-subnet connectivity among the set of NVE/PEs where the tenant is connected, a new SBD is created on all of them if a recursive resolution is needed. This SBD is instantiated as a regular BD (with no ACs) in each NVE/PE and has an IRB interface that connects the SBD to the IP-VRF. The IRB interface's IP or MAC address is used as the Overlay Index for a recursive resolution.

Depending on the existence and characteristics of the SBD and IRB

interfaces for the IP-VRFs, there are three different IP-VRF-to-IP-VRF scenarios identified and described in this document:

1. Interface-less model: no SBD and no Overlay Indexes required.
2. Interface-ful with an SBD IRB model: requires SBD as well as GW IP addresses as Overlay Indexes.
3. Interface-ful with an unnumbered SBD IRB model: requires SBD as well as MAC addresses as Overlay Indexes.

Inter-subnet IP multicast is outside the scope of this document.

4.4.1. Interface-less IP-VRF-to-IP-VRF Model

Figure 8 depicts the Interface-less IP-VRF-to-IP-VRF model.

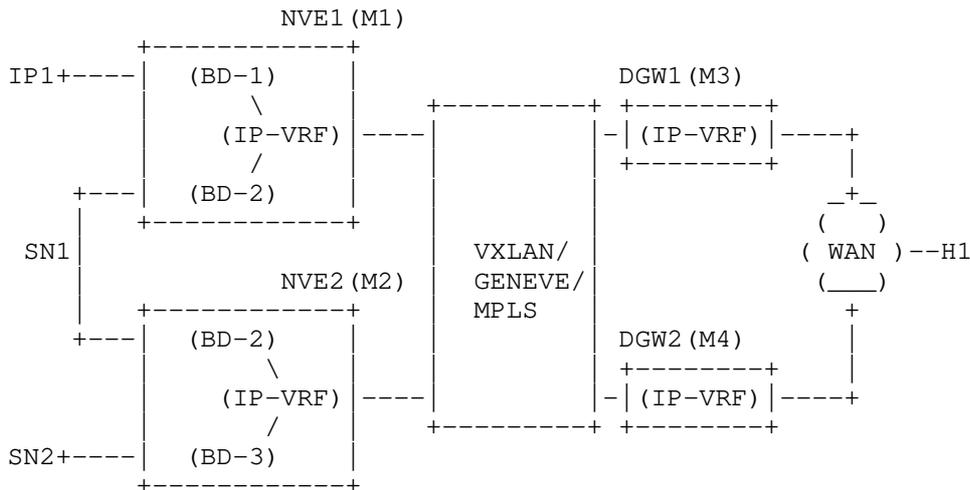


Figure 8: Interface-less IP-VRF-to-IP-VRF Model

In this case:

- a) The NVEs and DGWs must provide connectivity between hosts in SN1, SN2, and IP1 and hosts sitting at the other end of the WAN -- for example, H1. It is assumed that the DGWs import/export IP and/or VPN-IP routes to/from the WAN.
- b) The IP-VRF instances in the NVE/DGWs are directly connected through NVO tunnels, and no IRBs and/or BD instances are instantiated to connect the IP-VRFs.
- c) The solution must provide Layer 3 connectivity among the IP-VRFs for Ethernet NVO tunnels -- for instance, VXLAN or GENEVE.
- d) The solution may provide Layer 3 connectivity among the IP-VRFs for IP NVO tunnels -- for example, GENEVE (with IP payload).

In order to meet the above requirements, the EVPN route type 5 will be used to advertise the IP prefixes, along with the EVPN Router's MAC Extended Community as defined in [RFC9135] if the advertising NVE/DGW uses Ethernet NVO tunnels. Each NVE/DGW will advertise an RT-5 for each of its prefixes with the following fields:

- * RD as per [RFC7432].
- * Ethernet Tag ID = 0.
- * IP prefix length and IP address, as explained in the previous sections.
- * GW IP address = 0.
- * ESI = 0.

* MPLS label or VNI corresponding to the IP-VRF.

Each RT-5 will be sent with a Route Target identifying the tenant (IP-VRF) and may be sent with two BGP extended communities:

- * The first one is the BGP Encapsulation Extended Community, as per [RFC9012], identifying the tunnel type.
- * The second one is the EVPN Router's MAC Extended Community, as per [RFC9135], containing the MAC address associated with the NVE advertising the route. This MAC address identifies the NVE/DGW and MAY be reused for all the IP-VRFs in the NVE. The EVPN Router's MAC Extended Community must be sent if the route is associated with an Ethernet NVO tunnel -- for instance, VXLAN. If the route is associated with an IP NVO tunnel -- for instance, GENEVE with an IP payload -- the EVPN Router's MAC Extended Community should not be sent.

The following example illustrates the procedure to advertise and forward packets to SN1/24 (IPv4 prefix advertised from NVE1):

(1) NVE1 advertises the following BGP route:

- * Route type 5 (IP Prefix route) containing:
 - IPL = 24, IP = SN1, Label = 10.
 - GW IP = set to 0.
 - BGP Encapsulation Extended Community [RFC9012].
 - EVPN Router's MAC Extended Community that contains M1.
 - Route Target identifying the tenant (IP-VRF).

(2) DGW1 imports the received routes from NVE1:

- * DGW1 installs SN1/24 in the IP-VRF identified by the RT-5 Route Target.
- * Since GW IP = ESI = 0, the label is a non-zero value, and the local policy indicates this interface-less model, DGW1, will use the label and next hop of the RT-5, as well as the MAC address conveyed in the EVPN Router's MAC Extended Community (as the inner destination MAC address) to set up the forwarding state and later encapsulate the routed IP packets.

(3) When DGW1 receives a packet from the WAN with destination IPx, where IPx belongs to SN1/24:

- * A destination IP lookup is performed on the DGW1 IP-VRF table. The lookup yields SN1/24.
- * Since the RT-5 for SN1/24 had a GW IP = ESI = 0, a non-zero label, and a next hop, and since the model is interface-less, DGW1 will not need a recursive lookup to resolve the route.
- * The IP packet destined to IPx is encapsulated with: inner source MAC = DGW1 MAC, inner destination MAC = M1, outer source IP (tunnel source IP) = DGW1 IP, and outer destination IP (tunnel destination IP) = NVE1 IP. The source and inner destination MAC addresses are not needed if IP NVO tunnels are used.

(4) When the packet arrives at NVE1:

- * NVE1 will identify the IP-VRF for an IP lookup based on the label (the inner destination MAC is not needed to identify the IP-VRF).
- * An IP lookup is performed in the routing context, where SN1

turns out to be a local subnet associated with BD-2. A subsequent lookup in the ARP table and the BD FIB will provide the forwarding information for the packet in BD-2.

The model described above is called an "interface-less" model since the IP-VRFs are connected directly through tunnels, and they don't require those tunnels to be terminated in SBDs instead, as in Sections 4.4.2 or 4.4.3.

4.4.2. Interface-ful IP-VRF-to-IP-VRF with SBD IRB

Figure 9 depicts the Interface-ful IP-VRF-to-IP-VRF with SBD IRB model.

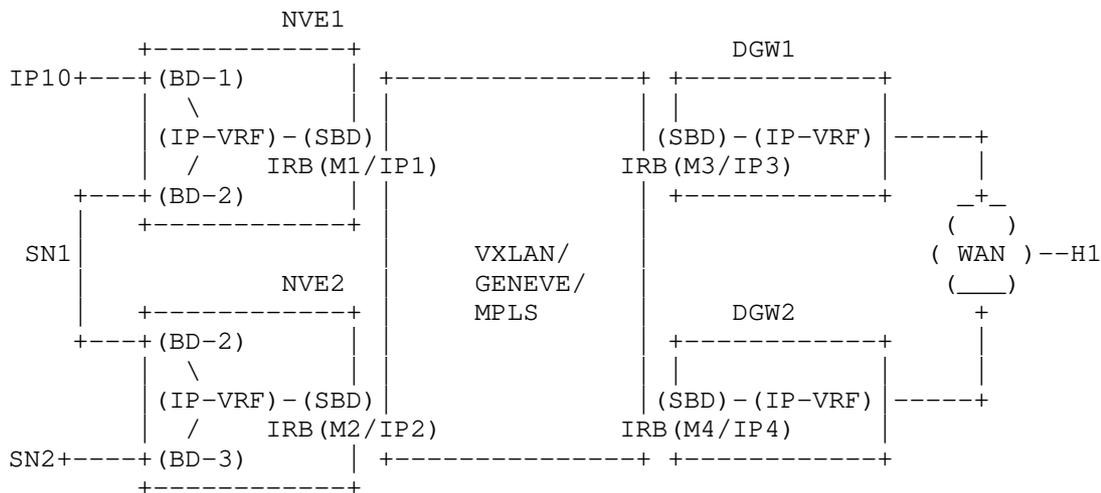


Figure 9: Interface-ful with SBD IRB Model

In this model:

- a) As in Section 4.4.1, the NVEs and DGWs must provide connectivity between hosts in SN1, SN2, and IP10 and in hosts sitting at the other end of the WAN.
- b) However, the NVE/DGWs are now connected through Ethernet NVO tunnels terminated in the SBD instance. The IP-VRFs use IRB interfaces for their connectivity to the SBD.
- c) Each SBD IRB has an IP and a MAC address, where the IP address must be reachable from other NVEs or DGWs.
- d) The SBD is attached to all the NVE/DGWs in the tenant domain BDs.
- e) The solution must provide Layer 3 connectivity for Ethernet NVO tunnels -- for instance, VXLAN or GENEVE (with Ethernet payload).

EVPN type 5 routes will be used to advertise the IP prefixes, whereas EVPN RT-2 routes will advertise the MAC/IP addresses of each SBD IRB interface. Each NVE/DGW will advertise an RT-5 for each of its prefixes with the following fields:

- * RD as per [RFC7432].
- * Ethernet Tag ID = 0.
- * IP prefix length and IP address, as explained in the previous sections.
- * GW IP address = IRB-IP of the SBD (this is the Overlay Index that will be used for the recursive route resolution).
- * ESI = 0.
- * Label value should be zero since the RT-5 route requires a recursive lookup resolution to an RT-2 route. It is ignored on

reception, and the MPLS label or VNI from the RT-2's MPLS Label1 field is used when forwarding packets.

Each RT-5 will be sent with a Route Target identifying the tenant (IP-VRF). The EVPN Router's MAC Extended Community should not be sent in this case.

The following example illustrates the procedure to advertise and forward packets to SN1/24 (IPv4 prefix advertised from NVE1):

(1) NVE1 advertises the following BGP routes:

- * Route type 5 (IP Prefix route) containing:
 - IPL = 24, IP = SN1, Label = SHOULD be set to 0.
 - GW IP = IP1 (SBD IRB's IP).
 - Route Target identifying the tenant (IP-VRF).
- * Route type 2 (MAC/IP Advertisement route for the SBD IRB) containing:
 - ML = 48, M = M1, IPL = 32, IP = IP1, Label = 10.
 - A BGP Encapsulation Extended Community [RFC9012].
 - Route Target identifying the SBD. This Route Target may be the same as the one used with the RT-5.

(2) DGW1 imports the received routes from NVE1:

- * DGW1 installs SN1/24 in the IP-VRF identified by the RT-5 Route Target.
 - Since GW IP is different from zero, the GW IP (IP1) will be used as the Overlay Index for the recursive route resolution to the RT-2 carrying IP1.

(3) When DGW1 receives a packet from the WAN with destination IPx, where IPx belongs to SN1/24:

- * A destination IP lookup is performed on the DGW1 IP-VRF table. The lookup yields SN1/24, which is associated with the Overlay Index IP1. The forwarding information is derived from the RT-2 received for IP1.
- * The IP packet destined to IPx is encapsulated with: inner source MAC = M3, inner destination MAC = M1, outer source IP (source VTEP) = DGW1 IP, and outer destination IP (destination VTEP) = NVE1 IP.

(4) When the packet arrives at NVE1:

- * NVE1 will identify the IP-VRF for an IP lookup based on the label and the inner MAC DA.
- * An IP lookup is performed in the routing context, where SN1 turns out to be a local subnet associated with BD-2. A subsequent lookup in the ARP table and the BD FIB will provide the forwarding information for the packet in BD-2.

The model described above is called an "interface-ful with SBD IRB" model because the tunnels connecting the DGWs and NVEs need to be terminated into the SBD. The SBD is connected to the IP-VRFs via SBD IRB interfaces, and that allows the recursive resolution of RT-5s to GW IP addresses.

4.4.3. Interface-ful IP-VRF-to-IP-VRF with Unnumbered SBD IRB

Figure 10 depicts the Interface-ful IP-VRF-to-IP-VRF with unnumbered

SBD IRB model. Note that this model is similar to the one described in Section 4.4.2, only without IP addresses on the SBD IRB interfaces.

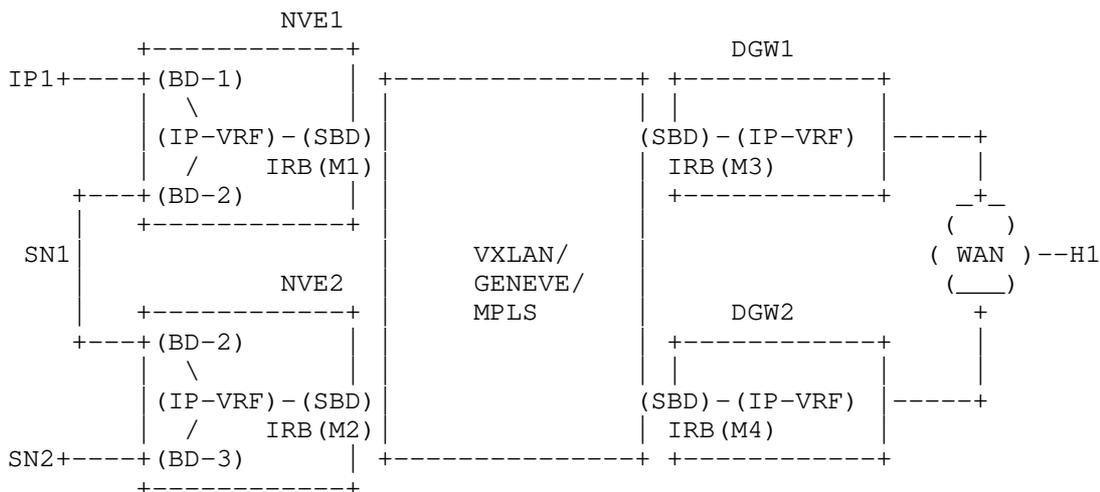


Figure 10: Interface-ful with Unnumbered SBD IRB Model

In this model:

- As in Sections 4.4.1 and 4.4.2, the NVEs and DGWs must provide connectivity between hosts in SN1, SN2, and IP1 and in hosts sitting at the other end of the WAN.
- As in Section 4.4.2, the NVE/DGWs are connected through Ethernet NVO tunnels terminated in the SBD instance. The IP-VRFs use IRB interfaces for their connectivity to the SBD.
- However, each SBD IRB has a MAC address only and no IP address (which is why the model refers to an "unnumbered" SBD IRB). In this model, there is no need to have IP reachability to the SBD IRB interfaces themselves, and there is a requirement to limit the number of IP addresses used.
- As in Section 4.4.2, the SBD is composed of all the NVE/DGW BDs of the tenant that need inter-subnet forwarding.
- As in Section 4.4.2, the solution must provide Layer 3 connectivity for Ethernet NVO tunnels -- for instance, VXLAN or GENEVE (with Ethernet payload).

This model will also make use of the RT-5 recursive resolution. EVPN type 5 routes will advertise the IP prefixes along with the EVPN Router's MAC Extended Community used for the recursive lookup, whereas EVPN RT-2 routes will advertise the MAC addresses of each SBD IRB interface (this time without an IP).

Each NVE/DGW will advertise an RT-5 for each of its prefixes with the same fields as described in Section 4.4.2, except:

* GW IP address = set to 0.

Each RT-5 will be sent with a Route Target identifying the tenant (IP-VRF) and the EVPN Router's MAC Extended Community containing the MAC address associated with the SBD IRB interface. This MAC address may be reused for all the IP-VRFs in the NVE.

The example is similar to the one in Section 4.4.2:

(1) NVE1 advertises the following BGP routes:

* Route type 5 (IP Prefix route) containing the same values as in the example in Section 4.4.2, except:

- GW IP = SHOULD be set to 0.

- EVPN Router's MAC Extended Community containing M1 (this will be used for the recursive lookup to an RT-2).

- * Route type 2 (MAC route for the SBD IRB) with the same values as in Section 4.4.2, except:

- ML = 48, M = M1, IPL = 0, Label = 10.

(2) DGW1 imports the received routes from NVE1:

- * DGW1 installs SN1/24 in the IP-VRF identified by the RT-5 Route Target.

- The MAC contained in the EVPN Router's MAC Extended Community sent along with the RT-5 (M1) will be used as the Overlay Index for the recursive route resolution to the RT-2 carrying M1.

(3) When DGW1 receives a packet from the WAN with destination IPx, where IPx belongs to SN1/24:

- * A destination IP lookup is performed on the DGW1 IP-VRF table. The lookup yields SN1/24, which is associated with the Overlay Index M1. The forwarding information is derived from the RT-2 received for M1.

- * The IP packet destined to IPx is encapsulated with: inner source MAC = M3, inner destination MAC = M1, outer source IP (source VTEP) = DGW1 IP, and outer destination IP (destination VTEP) = NVE1 IP.

(4) When the packet arrives at NVE1:

- * NVE1 will identify the IP-VRF for an IP lookup based on the label and the inner MAC DA.

- * An IP lookup is performed in the routing context, where SN1 turns out to be a local subnet associated with BD-2. A subsequent lookup in the ARP table and the BD FIB will provide the forwarding information for the packet in BD-2.

The model described above is called an "interface-ful with unnumbered SBD IRB" model (as in Section 4.4.2) but without the SBD IRB having an IP address.

5. Security Considerations

This document provides a set of procedures to achieve inter-subnet forwarding across NVEs or PEs attached to a group of BDs that belong to the same tenant (or VPN). The security considerations discussed in [RFC7432] apply to the intra-subnet forwarding or communication within each of those BDs. In addition, the security considerations in [RFC4364] should also be understood, since this document and [RFC4364] may be used in similar applications.

Contrary to [RFC4364], this document does not describe PE/CE route distribution techniques but rather considers the CEs as TSs or VAs that do not run dynamic routing protocols. This can be considered a security advantage, since dynamic routing protocols can be blocked on the NVE/PE ACs, not allowing the tenant to interact with the infrastructure's dynamic routing protocols.

In this document, the RT-5 may use a regular BGP next hop for its resolution or an Overlay Index that requires a recursive resolution to a different EVPN route (an RT-2 or an RT-1). In the latter case, it is worth noting that any action that ends up filtering or modifying the RT-2 or RT-1 routes used to convey the Overlay Indexes will modify the resolution of the RT-5 and therefore the forwarding of packets to the remote subnet.

6. IANA Considerations

IANA has registered value 5 in the "EVPN Route Types" registry [EVPNRouteTypes] defined by [RFC7432] as follows:

Value	Description	Reference
5	IP Prefix	RFC 9136

Table 3

7. References

7.1. Normative References

- [EVPNRouteTypes] IANA, "EVPN Route Types", <<https://www.iana.org/assignments/evpn>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018, <<https://www.rfc-editor.org/info/rfc8365>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9135] Sajassi, A., Salam, S., Thoria, S., Drake, J., and J. Rabadan, "Integrated Routing and Bridging in Ethernet VPN (EVPN)", RFC 9135, DOI 10.17487/RFC9135, October 2021, <<https://www.rfc-editor.org/info/rfc9135>>.

7.2. Informative References

- [IEEE-802.1Q] IEEE, "IEEE Standard for Local and Metropolitan Area Networks -- Bridges and Bridged Networks", DOI 10.1109/IEEESTD.2018.8403927, IEEE Std 802.1Q, July 2018, <https://standards.ieee.org/standard/802_1Q-2018.html>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC5227] Cheshire, S., "IPv4 Address Conflict Detection", RFC 5227, DOI 10.17487/RFC5227, July 2008, <<https://www.rfc-editor.org/info/rfc5227>>.
- [RFC5798] Nadas, S., Ed., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, DOI 10.17487/RFC5798, March 2010,

<<https://www.rfc-editor.org/info/rfc5798>>.

- [RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", RFC 7348, DOI 10.17487/RFC7348, August 2014, <<https://www.rfc-editor.org/info/rfc7348>>.
- [RFC7365] Lasserre, M., Balus, F., Morin, T., Bitar, N., and Y. Rekhter, "Framework for Data Center (DC) Network Virtualization", RFC 7365, DOI 10.17487/RFC7365, October 2014, <<https://www.rfc-editor.org/info/rfc7365>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.

Acknowledgments

The authors would like to thank Mukul Katiyar, Jeffrey Zhang, and Alex Nichol for their valuable feedback and contributions. Tony Przygienda and Thomas Morin also helped improve this document with their feedback. Special thanks to Eric Rosen for his detailed review, which really helped improve the readability and clarify the concepts. We also thank Alvaro Retana for his thorough review.

Contributors

In addition to the authors listed on the front page, the following coauthors have also contributed to this document:

Senthil Sathappan
Florin Balus
Aldrin Isaac
Senad Palislamovic
Samir Thoria

Authors' Addresses

Jorge Rabadan (editor)
Nokia
777 E. Middlefield Road
Mountain View, CA 94043
United States of America

Email: jorge.rabadan@nokia.com

Wim Henderickx
Nokia

Email: wim.henderickx@nokia.com

John Drake
Juniper

Email: jdrake@juniper.net

Wen Lin
Juniper

Email: wlin@juniper.net

Ali Sajassi
Cisco

Email: sajassi@cisco.com