

Internet Engineering Task Force (IETF)
Request for Comments: 7900
Updates: 6513, 6514, 6625
Category: Standards Track
ISSN: 2070-1721

Y. Rekhter, Ed.
E. Rosen, Ed.
Juniper Networks, Inc.
R. Aggarwal
Arktan
Y. Cai
Alibaba Group
T. Morin
Orange
June 2016

Extranet Multicast in BGP/IP MPLS VPNs

Abstract

Previous RFCs specify the procedures necessary to allow IP multicast traffic to travel from one site to another within a BGP/MPLS IP VPN (Virtual Private Network). However, it is sometimes desirable to allow multicast traffic whose source is in one VPN to be received by systems that are in another VPN. This is known as a "Multicast VPN (MVPN) extranet". This document updates RFCs 6513, 6514, and 6625 by specifying the procedures that are necessary in order to provide extranet MVPN service.

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 <http://www.rfc-editor.org/info/rfc7900>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Terminology	4
1.2. Scope	7
1.2.1. Customer Multicast Control Protocols	7
1.2.2. Provider Multicast Control Protocols	7
1.3. Clarification on Use of Route Distinguishers	8
1.4. Overview	9
2. Extranets and Overlapping Address Spaces	12
2.1. Ambiguity: P-Tunnel with Extranet/Non-extranet Flows	14
2.2. Ambiguity: P-Tunnel with Multiple Extranet Flows	16
2.3. Preventing Misdelivery in These Scenarios	18
2.3.1. Do Not Deliver Packets from the Wrong P-tunnel	18
2.3.2. Policies to Prevent Ambiguity on a P-Tunnel	20
3. Extranet Transmission Models	21
3.1. Transmitting an Extranet C-Flow on a Single PMSI	21
3.1.1. Without Extranet Separation	22
3.1.2. With Extranet Separation	22
3.2. Transmitting an Extranet C-Flow over Multiple PMSIs	23
4. Distribution of Routes That Match C-S/C-RP Addresses	23
4.1. UMH-Eligible Routes	23
4.1.1. Extranet Separation	24
4.2. Distribution of Unicast Routes Matching C-RPs and DRs	25
4.3. Route Targets and Ambiguous UMH-Eligible Routes	26
4.4. Dynamically Marking Extranet Routes	27
4.4.1. The Extranet Source Extended Community	27
4.4.2. Distribution of Extranet Source Extended Community	29
4.5. The Extranet Separation Extended Community	30

5. Origination and Distribution of BGP A-D Routes	30
5.1. Route Targets of UMH-Eligible Routes and A-D Routes	30
5.2. Considerations for Particular Inclusive Tunnel Types	33
5.2.1. RSVP-TE P2MP or Ingress Replication	33
5.2.2. Ingress Replication	34
6. When PIM Is the PE-PE C-Multicast Control Plane	35
6.1. Provisioning VRFs with RTs	36
6.1.1. Incoming and Outgoing Extranet RTs	36
6.1.2. UMH-Eligible Routes and RTs	37
6.1.3. PIM C-Instance Reverse Path Forwarding Determination	37
6.2. "Single PMSI per C-Flow" Model	38
6.2.1. Forming the MI-PMSIs	38
6.2.2. S-PMSIs	41
6.2.3. Sending PIM Control Packets	42
6.2.4. Receiving PIM Control Packets	43
6.2.5. Sending and Receiving Data Packets	43
6.3. "Multiple PMSIs per C-Flow" Model	43
6.3.1. Forming the MI-PMSIs	44
7. When BGP Is the PE-PE C-Multicast Control Plane	46
7.1. Originating C-Multicast Routes	46
7.2. Originating A-D Routes without Extranet Separation	47
7.2.1. Intra-AS I-PMSI A-D Routes	47
7.2.2. S-PMSI A-D Routes	47
7.2.3. Source Active A-D Routes	48
7.2.3.1. When Inter-Site Shared Trees Are Used	48
7.2.3.2. When Inter-Site Shared Trees Are Not Used	49
7.3. Originating A-D Routes with Extranet Separation	49
7.3.1. Intra-AS I-PMSI A-D Routes	49
7.3.2. S-PMSI A-D Routes	50
7.3.3. Source Active A-D Routes	52
7.4. Determining the Expected P-Tunnel for a C-Flow	52
7.4.1. (C-S,C-G) S-PMSI A-D Routes	54
7.4.2. (C-S,C-*) S-PMSI A-D Routes	54
7.4.3. (C-*,C-G) S-PMSI A-D Routes	55
7.4.4. (C-*,C-*) S-PMSI A-D Routes	56
7.4.5. I-PMSI A-D Routes	56
7.5. Packets Arriving from the Wrong P-Tunnel	57
8. Multiple Extranet VRFs on the Same PE	57
9. IANA Considerations	58
10. Security Considerations	59
11. References	61
11.1. Normative References	61
11.2. Informative References	62
Acknowledgments	64
Contributors	64
Authors' Addresses	65

1. Introduction

Previous RFCs [RFC6513] [RFC6514] specify the procedures necessary to allow IP multicast traffic to travel from one site to another within a BGP/MPLS IP VPN (Virtual Private Network). However, it is sometimes desirable to allow multicast traffic whose source is in one VPN to be received by systems that are in another VPN. This is known as an "extranet Multicast VPN (MVPN)". This document specifies the procedures that are necessary in order to provide extranet MVPN functionality.

This document updates RFCs 6513, 6514, and 6625 by specifying the procedures that are necessary in order to provide extranet MVPN service.

1.1. Terminology

This document uses terminology from [RFC6513] and in particular uses the prefixes "C-" and "P-" as specified in Section 3.1 of [RFC6513], and "A-D routes" for "auto-discovery routes".

The term "Upstream Multicast Hop" (UMH) is used as defined in [RFC6513].

The term "UMH-eligible route" is used to mean "route eligible for UMH determination", as defined in Section 5.1.1 of [RFC6513]. We will say that a given UMH-eligible route or unicast route "matches" a given IP address, in the context of a given Virtual Routing and Forwarding table (VRF), if the address prefix of the given route is the longest match in that VRF for the given IP address. We will sometimes say that a route "matches" a particular host if the route matches an IP address of the host.

We follow the terminology of Section 3.2 of [RFC6625] when talking of a "Selective Provider Multicast Service Interface" (S-PMSI) A-D route being "installed". That is, we say that an S-PMSI A-D route is "installed" (in a given VRF) if it has been selected by the BGP decision process as the preferred route for its Network Layer Reachability Information (NLRI). We also follow the terminology of Section 3.2 of [RFC6625] when saying that an S-PMSI A-D route has been "originated by a given PE"; this means that the given Provider Edge's (PE's) IP address is contained in the Originating Router's IP Address field in the NLRI of the route.

We use the following additional terminology and notation:

- o Extranet C-source: a multicast source, in a given VPN, that is allowed by policy to send multicast traffic to receivers that are in other VPNs.
- o Extranet C-receiver: a multicast receiver, in a given VPN, that is allowed by policy to receive multicast traffic from extranet C-sources that are in other VPNs.
- o Extranet C-flow: a multicast flow (with a specified C-source address and C-group address) with the following properties: its source is an extranet C-source, and it is allowed by policy to have extranet C-receivers.
- o Extranet C-group: a multicast group address that is in the "Any-Source Multicast" (ASM) group address range and that is allowed by policy to have extranet C-sources and extranet C-receivers that are not all in the same VPN. Note that we will sometimes refer to "Source-Specific Multicast (SSM) C-group addresses" (i.e., C-group addresses in the SSM group address range) but will never call them "extranet C-groups".

N.B.: Any source of traffic for an extranet C-group is considered to be an extranet C-source, and any receiver of traffic addressed to an extranet C-group is considered to be an extranet C-receiver.

- o Extranet C-RP: a multicast Rendezvous Point (RP) for an extranet C-group; it is allowed by policy to receive PIM Register messages [RFC7761] from outside its VPN and to send multicast data packets to extranet C-receivers outside its VPN.
- o Host(C-S,A): the host (or, if C-S is an "anycast address", the set of hosts) denoted by the address C-S in the context of VPN-A. For example, if a particular C-source in VPN-A has address C-S, then Host(C-S,A) refers to that C-source.
- o "SAFI n" route: a BGP route whose Address Family Identifier (AFI) is either 1 (IPv4) or 2 (IPv6) and whose Subsequent Address Family Identifier (SAFI) is "n".
- o PTA: PMSI Tunnel Attribute [RFC6514].

Note that a given extranet C-source is not necessarily allowed to transmit to every extranet C-receiver; policy determines which extranet C-sources are allowed to transmit to which extranet C-receivers. However, in the case of an extranet (ASM) C-group, all transmitters to the group are allowed to transmit to all the receivers of the group, and all the receivers of the group are allowed to receive from all transmitters to the group.

We say that a given VRF "contains" or "has" a multicast C-source (or that the C-source is "in" the VRF) if that C-source is in a site connected to that VRF and the VRF originates a UMH-eligible route (see Section 4) that matches the address of the C-source.

We say that a given VRF "contains" or "has" a multicast C-receiver (or that the C-receiver is "in" the VRF) if that C-receiver is in a site connected to that VRF.

We say that a given VRF "contains" or "has" the C-RP for a given ASM group (or that the C-RP is "in" the VRF) if that C-RP is in a site connected to that VRF and the VRF originates a unicast route and a (possibly different, possibly the same) UMH-eligible route (see Section 4) whose respective address prefixes match the C-RP address.

[RFC6513] allows a set of "P-tunnels" (defined in Section 3.2 of [RFC6513]) to be aggregated together and transported via an outer P-tunnel; i.e., it allows for the use of hierarchical Label Switched Paths (LSPs) as P-tunnels. A two-level hierarchical LSP, for example, can be thought of as a set of "inner tunnels" aggregated into an outer tunnel. In this document, when we speak of a P-tunnel, we are always speaking of the innermost P-tunnel, i.e., of a P-tunnel at the lowest hierarchical level. P-tunnels are identified in the PMSI Tunnel attributes ("PTAs" in this document) [RFC6514] of BGP auto-discovery (A-D) routes. Two PTAs that have the same Tunnel Type and Tunnel Identifier fields but different MPLS label fields are thus considered to identify two different P-tunnels. (That is, for the purposes of this document, the MPLS label included in the PTA, if any, is considered to be part of the tunnel identifier.)

We say that the NLRI of a BGP S-PMSI A-D route or Source Active A-D route contains (C-S,C-G) if its Multicast Source field contains C-S and its Multicast Group field contains C-G. If either or both of these fields are encoded as a wildcard, we will say that the NLRI contains (C-*,C-*) (both fields encoded as wildcards), (C-*,C-G) (Multicast Source field encoded as a wildcard), or (C-S,C-*) (Multicast Group field encoded as a wildcard).

We use the term "VPN security violation" to refer to any situation in which a packet is delivered to a particular VPN, even though, by policy, it should not be delivered to that VPN.

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 [RFC2119].

1.2. Scope

1.2.1. Customer Multicast Control Protocols

This document presumes that the VPN customer is using PIM - Sparse Mode (PIM-SM) [RFC7761] as the multicast control protocol at the customer sites. PIM-SM may be used in either the ASM service model or the SSM service model; this document covers both cases. Support for other customer IP multicast control protocols (e.g., [RFC5015], PIM - Dense Mode) is outside the scope of this document. Support for the use of MPLS multicast control protocols (e.g., [RFC6388] [RFC4875]) by customer sites is also outside the scope of this document.

When a VPN customer uses ASM, the customer routers need to be able to map from a C-group address to a C-RP address. These mappings can be provisioned in each router, or they can be discovered dynamically through protocols such as the Bootstrap Router (BSR) mechanism [RFC5059]. However, it cannot be assumed that such protocols will automatically work in the context of an extranet. Discussion of the use of such protocols in an extranet is outside the scope of this document.

1.2.2. Provider Multicast Control Protocols

[RFC6513] allows either PIM or BGP to be used as the protocol for distributing customer multicast routing information. Except where otherwise specified, such as in Sections 6 and 7, the procedures of this document cover both cases.

1.3. Clarification on Use of Route Distinguishers

[RFC4364] requires that every VRF be associated with one or more Route Distinguishers (RDs). Each VPN-IPv4 or VPN-IPv6 route that is exported from a particular VRF contains, in its NLRI, an RD that is associated with that VRF.

[RFC4364] allows a given RD to be associated with more than one VRF, as long as all the VRFs associated with that RD belong to the same VPN. However, in the most common deployment model, each RD is associated with one and only one VRF. [RFC6513] and [RFC6514] presuppose this deployment model. That is, [RFC6513] and [RFC6514] presuppose that every RD is associated with one and only one VRF. We will call this the "unique VRF per RD" condition.

[RFC6514] defines the MCAST-VPN address family, which has a number of route types. Each Intra-Autonomous System (Intra-AS) "Inclusive Provider Multicast Service Interface" (I-PMSI) A-D route, S-PMSI A-D route, and Source Active A-D route, when exported from a given VRF, contains, in its NLRI, an RD that is associated with the VRF. [RFC6513] and [RFC6514] also discuss a class of routes known as "UMH-eligible" routes; when a UMH-eligible route is exported from a given VRF, its NLRI contains an RD of the VRF.

[RFC6514] also defines MCAST-VPN routes whose NLRIs do not contain an RD of the VRF from which they are exported: the C-multicast Join routes and the Leaf A-D routes.

Those route types that, when exported from a given VRF, contain (in their NLRIs) an RD of the VRF, will be known in this document as "local-RD routes".

Given the "unique VRF per RD" condition, if one sees that two local-RD routes have the same RD, one can infer that the two routes originated from the same VRF. This inference can be drawn even if the two routes do not have the same SAFI, as long as the two routes are both local-RD routes.

This document builds upon [RFC6513] and [RFC6514]; therefore, the "unique VRF per RD" condition is REQUIRED.

[RFC6514] presupposes a further requirement on the use of RDs in the local-RD routes exported from a given VRF. Suppose that a given VRF exports a Source Active A-D route containing (C-S,C-G). That VRF will also export a UMH-eligible route matching C-S. [RFC6514] presupposes that the UMH-eligible route and the Source Active A-D route have the same RD.

In most cases, not only is a given RD associated with only a single VRF, but a given VRF is associated with only a single RD. We will call this the "unique RD per VRF" condition. When this condition holds, all the local-RD routes exported from a given VRF will have the same RD. This ensures that the presupposition of the previous paragraph will hold, i.e., that the RD in a Source Active A-D route exported from a given VRF will have the same RD as the corresponding UMH-eligible route exported from the same VRF.

Section 7.3 of this document describes a procedure known as "extranet separation". When extranet separation is NOT being used, it is REQUIRED by this document that the "unique RD per VRF" condition hold. This ensures that all the local-RD routes exported from a given VRF will have the same RD.

When extranet separation is used, a VRF that contains both extranet sources and non-extranet sources MUST be configured with two RDs. One of these RDs is known as the "default RD", and the other is known as the "extranet RD". It MUST be known by configuration which RD is the default RD and which is the extranet RD.

When a VRF is configured with only one RD, we will refer to that RD as the "default RD".

In general, local-RD routes exported from a given VRF will contain the default RD. However, when extranet separation is used, some of the local-RD routes exported from the VRF will contain the extranet RD. Details concerning the exported routes that contain the extranet RD can be found in Sections 4.1 and 7.3.

Note that the "unique VRF per RD" condition applies to the extranet RD as well as the default RD. That is, a given extranet RD is associated with a unique VRF.

1.4. Overview

Consider two VPNs, VPN-S and VPN-R, each of which supports MVPN functionality as specified in [RFC6513] and/or [RFC6514]. In the simplest configuration, VPN-S is a collection of VRFs, each of which is configured with a particular Route Target (RT) value (call it "RT-S") as its import RT and as its export RT. Similarly, VPN-R is a collection of VRFs, each of which is configured with a particular RT value (call it "RT-R") as its import RT and as its export RT.

In this configuration, multicast C-receivers contained in a VPN-R VRF cannot receive multicast data traffic from multicast C-sources contained in a VPN-S VRF. If it is desired to allow this, one needs to create an MVPN "extranet". Creating an extranet requires procedures in addition to those specified in [RFC6513], [RFC6514], and [RFC6625]; this document specifies these additional procedures.

In the example above, the additional procedures will allow a selected set of routes exported from the VPN-S VRFs (i.e., from the VRFs containing extranet C-sources) to be imported into the VPN-R VRFs (i.e., into the VRFs containing extranet C-receivers). These routes include the routes that are to be eligible for use as UMH routes (see Section 5.1 of [RFC6513]) in the extranet, as well as a selected set of BGP A-D routes (Intra-AS I-PMSI A-D routes, S-PMSI A-D routes, and Source Active A-D routes). Importing these routes into the VPN-R VRFs makes it possible to determine, in the context of a VPN-R VRF, that a particular C-multicast Join needs to be delivered to a particular VPN-S VRF. It also makes it possible to determine, in the context of a VPN-R VRF, the P-tunnel through which the aforementioned VPN-S VRF sends a particular C-flow.

Depending on the type of P-tunnel used, it may also be necessary for Leaf A-D routes to be exported by one or more VPN-R VRFs and imported into a VPN-S VRF.

There are no extranet-specific procedures governing the use and distribution of BGP C-multicast routes.

If PIM is used as the PE-PE protocol for distributing C-multicast routing information, additional BGP A-D routes must be exported from the VPN-R VRFs and imported into the VPN-S VRFs, so that the VPN-S VRFs can join the P-tunnels that the VPN-R VRFs use for sending PIM control messages. Details can be found in Section 6.

The simple example above describes an extranet created from two MVPNs, one of which contains extranet C-sources and one of which contains extranet C-receivers. However, the procedures described in this document allow for much more complicated scenarios.

For instance, an extranet may contain extranet C-sources and/or extranet C-receivers from an arbitrary number of VPNs, not just from two VPNs. An extranet C-receiver in VPN-R may be allowed to receive multicast traffic from extranet C-sources in VPN-A, VPN-B, and VPN-C. Similarly, extranet C-sources in VPN-S may be allowed to send multicast traffic to multicast C-receivers that are in VPN-A, VPN-B, VPN-C, etc.

A given VPN customer may desire that only some of its multicast C-sources be treated as extranet C-sources. This can be accomplished by appropriate provisioning of the import and export RTs of that customer's VRFs (as well as the VRFs of other VPNs that contain extranet C-receivers for extranet C-flows of the given customer).

A given VPN customer may desire that some of its extranet C-sources can transmit only to a certain set of VPNs while other of its extranet C-sources can transmit only to a different set of VPNs. This can be accomplished by provisioning the VRFs to export different routes with different RTs.

In all these cases, the VPN customers set the policies, and the Service Provider (SP) implements the policies by the way it provisions the import and export RTs of the VRFs. It is assumed that the customer communicates to the SP the set of extranet C-source addresses and the set of VPNs to which each C-source can transmit. (Recall that every C-source that can transmit to an extranet C-group is an extranet C-source and must be able to transmit to any VPN that has receivers for that group. This must be taken into account when the provisioning is done.) This customer/SP communication is part of the service provisioning process and is outside the scope of this document.

It is possible that an extranet C-source will transmit both extranet C-flows and non-extranet C-flows. However, if extranet C-receiver C-R can receive extranet C-flows from extranet C-source C-S, the procedures of this document do not prevent C-R from requesting and receiving the non-extranet flows that are transmitted by C-S. Therefore, allowing an extranet C-source to transmit non-extranet C-flows is NOT RECOMMENDED. However, the SP has no control over the set of C-flows transmitted by a given C-source and can do no more than communicate this recommendation to its customers. (Alternatively, the customer and SP may coordinate on setting up filters to prevent unauthorized flows from being sent to a customer site; such a procedure is outside the scope of this document.) See Section 10 ("Security Considerations") for additional discussion of this issue.

Whenever a VPN is provisioned, there is a risk that errors in provisioning may result in an unintended cross-connection of VPNs. This would create a security problem for the customers. When provisioning an extranet, attention to detail is particularly important, as an extranet intentionally cross-connects VPNs. Care must always be taken to ensure that the cross-connections are according to the policy agreed upon by the SP and its customers.

Additionally, if one is connecting two VPNs that have overlapping address spaces, one has to be sure that the inter-VPN traffic neither originates from nor is destined to the part of the address space that is in the overlap. Other problems that can arise due to overlapping address spaces are discussed in Section 2.

2. Extranets and Overlapping Address Spaces

As specified in [RFC4364], the address space of one VPN may overlap with the address space of another. A given address may be "ambiguous" in that it denotes one system within VPN-A and a different system within VPN-B. In the notation of Section 1.1, if an address C-S is ambiguous between VPN-A and VPN-B, then $\text{Host}(C-S,A) \neq \text{Host}(C-S,B)$. However, any given address C-S MUST be unambiguous (i.e., MUST denote a single system) in the context of a given VPN.

When a set of VRFs belonging to different VPNs are combined into an extranet, it is no longer sufficient for an address to be unambiguous only within the context of a single VPN:

1. Suppose that C-S is the address of a given extranet C-source contained in VPN-A. Now consider the set of VPNs {VPN-B, VPN-C, ...} containing extranet C-receivers that are allowed by policy to receive extranet C-flows from VPN-A's C-S. The address C-S MUST be unambiguous among this entire set of VPNs {VPN-A, VPN-B, VPN-C, ...}; i.e., $\text{Host}(C-S,A) == \text{Host}(C-S,B) == \text{Host}(C-S,C)$.

The implication is that C-S in VPN-A is not necessarily an extranet C-source for all VPNs that contain extranet C-receivers; policy MUST be used to ensure that C-S is an extranet C-source for a given VPN, say VPN-B, only if C-S is unambiguous between VPN-A and VPN-B.

2. If a given VRF contains extranet C-receivers for a given extranet C-source, then the address of this C-source MUST be unambiguous among all the extranet C-sources for which there are C-receivers in the VRF. This is true whether or not C-sources are in VRFs that belong to the same VPN or different VPNs.

The implication is that if C-S in VRF-X is ambiguous with C-S in VRF-Y, then there MUST NOT be any VRF, say VRF-Z, containing C-receivers that are allowed by policy to receive extranet C-flows from both C-S in VRF-X and C-S in VRF-Y.

Note: A VPN customer may be using anycast addresses. An anycast address is intentionally ambiguous, as it denotes a set of systems rather than a single system. In this document, we will consider an anycast address to be unambiguous in a given context as long as it denotes the same set of systems whenever it occurs in that context.

A multicast C-group address, say C-G, may also be ambiguous in that it may be used for one multicast group in VPN-A and for an entirely different multicast group in VPN-B. If a set of MVPNs are combined into an extranet and C-G is an extranet C-group, it is necessary to ensure that C-G is unambiguous among the entire set of VPNs whose VRFs contain extranet C-sources, C-RPs, and/or extranet C-receivers for that C-group. This may require, as part of the provisioning process, customer/SP communication that is outside the scope of this document.

Subject to these restrictions, the SP has complete control over the distribution of routes in an MVPN. This control is exerted by provisioning either (1) the export RTs on the VRFs that originate the routes (i.e., the VRFs that contain the extranet C-sources) or (2) the import RTs on the VRFs that receive the routes (i.e., the VRFs that contain the extranet C-receivers), or both.

Some of the rules and restrictions on provisioning the RTs are applicable to all extranets; these are specified in Section 4. Sections 6 and 7 list additional rules and restrictions that are applicable only to particular extranet scenarios.

Even if all the RTs are provisioned according to the above rules and restrictions, it is still possible for a single P-tunnel to contain multicast data packets whose source and/or group addresses are ambiguous in the context of the set of PEs that receive data from the P-tunnel. That is, the above rules and restrictions are necessary, but not sufficient, to prevent address ambiguity from causing misdelivery of traffic. To prevent such misdelivery, additional procedures or policies must be used.

Sections 2.1 and 2.2 describe scenarios in which a given P-tunnel may carry data packets with ambiguous addresses. The additional procedures and policies needed to prevent misdelivery of data in those scenarios are outlined in Section 2.3. (The detailed procedures described in Sections 6 and 7 incorporate the considerations discussed in Section 2.3.)

2.1. Ambiguity: P-Tunnel with Extranet/Non-extranet Flows

In the following, we will use the notation "VRF A-n" to mean "VRF n of VPN-A".

If VPN-A and VPN-B have overlapping address spaces and are part of the same extranet, then the following problem may exist, as illustrated in Figure 1.




Figure 1: Ambiguity of Extranet and Non-extranet Source Address

Suppose that:

- o C-G is an SSM C-group used in VPN-A and VPN-B.
- o VRF A-1, on PE1, contains an extranet C-source, with IP address C-S1, that is allowed to have receivers in VPN-B. VRF A-1 thus exports to VPN-B a UMH-eligible route matching C-S1.
- o In addition, VRF A-1 contains a non-extranet C-source with IP address C-S2. VRF A-1 exports a UMH-eligible route matching C-S2 to other VPN-A VRFs but NOT to VPN-B.
- o VRF B-1, also on PE1, contains a non-extranet C-source with IP address C-S2. A UMH-eligible route matching C-S2 is thus exported from VRF B-1 to other VRFs in VPN-B.
- o Host(C-S2,A) != Host(C-S2,B). That is, C-S2 is an ambiguous address in any extranet that contains both VPN-A VRFs and VPN-B VRFs.
- o VRF B-2, on some other PE, say PE2, requests the multicast flow (C-S1,C-G). In the context of VRF B-2, C-S1 matches the route exported from VRF A-1. Thus, B-2's request to receive the (C-S1,C-G) flow is transmitted to VRF A-1.

- o VRF A-1 responds to VRF B-2's request for (C-S1,C-G) traffic by transmitting that traffic on P-tunnel P1.
- o VRF B-2 joins P-tunnel P1 in order to receive the (C-S1,C-G) traffic.
- o VRF A-2, on PE2, requests the (non-extranet) multicast flow (C-S2,C-G). In the context of VRF A-2, C-S2 matches the route exported from VRF A-1. Thus, A-2's request to receive the (C-S2,C-G) traffic is transmitted to VRF A-1.
- o VRF A-1 responds to VRF A-2's request for (C-S2,C-G) traffic by transmitting that traffic on P-tunnel P1.
- o VRF A-2 joins P-tunnel P1 in order to receive the (C-S2,C-G) traffic.
- o VRF B-2 requests the (non-extranet) multicast flow (C-S2,C-G). In the context of VRF B-2, C-S2 matches the route exported from VRF B-1. Thus, B-2's request to receive the (C-S2,C-G) flow is transmitted to VRF B-1.
- o VRF B-1 responds to VRF B-2's request for (C-S2,C-G) traffic by transmitting that traffic on P-tunnel P2.
- o VRF B-2 joins P-tunnel P2.

Since VRF B-2 has joined P-tunnel P1 and P-tunnel P2, it will receive (C-S2,C-G) traffic on both P-tunnels. The (C-S2,C-G) traffic that VRF B-2 needs to receive is traveling on P-tunnel P2; this (C-S2,C-G) traffic must be forwarded by B-2 to any attached customer sites that have C-receivers for it. But B-2 MUST discard the (C-S2,C-G) traffic that it receives on P1, as this is not the traffic that it has requested. If the (C-S2,C-G) traffic arriving on P1 were forwarded to B-2's customer sites, the C-receivers would not be able to distinguish the two flows, and the result would be a corrupted data stream.

Note that the procedures of Section 9.1.1 of [RFC6513] ("Discarding Packets from Wrong PE") will not cause VRF B-2 to discard the (C-S2,C-G) traffic that arrives on tunnel P1, because P1 and P2 have the same upstream PE.

Therefore, it is necessary to EITHER (1) prevent the above scenario from occurring OR (2) ensure that multicast data packets will be discarded if they arrive on the wrong P-tunnel (even if they arrive from the expected PE). See Section 2.3 for further discussion of this issue.

2.2. Ambiguity: P-Tunnel with Multiple Extranet Flows

Figure 2 illustrates another example of how overlapping address spaces may cause a problem.




Figure 2: Ambiguity of Extranet Source Addresses

Suppose that:

- o C-G is an SSM C-group address that is used in VPN-A, VPN-B, VPN-C, and VPN-D.
- o VRF A-1, on PE1, contains an extranet C-source, with IP address C-S1, that is allowed by policy to have C-receivers in VPN-C (but not in VPN-D). VRF A-1 thus exports a UMH-eligible route matching C-S1 to VPN-C.
- o In addition, VRF A-1 contains an extranet C-source, with IP address C-S2, that is allowed by policy to have C-receivers in VPN-D (but not in VPN-C). VRF A-1 thus exports a UMH-eligible route matching C-S2 to VPN-D.
- o VRF B-1, also on PE1, contains an extranet C-source, with IP address C-S2, that is allowed by policy to have C-receivers in VPN-C (but not in VPN-D). VRF B-1 thus exports a UMH-eligible route matching C-S2 to VPN-C.
- o Host(C-S2,A) != Host(C-S2,B). That is, C-S2 is an ambiguous address in any extranet that contains both VPN-A VRFs and VPN-B VRFs.

- o VRF C-1, on some other PE, say PE2, requests the extranet multicast flow (C-S1,C-G). In the context of VRF C-1, C-S1 matches the route exported from VRF A-1. Thus, C-1's request to receive the (C-S1,C-G) flow is transmitted to VRF A-1.
- o VRF A-1 responds to VRF C-1's request for (C-S1,C-G) traffic by transmitting that traffic on P-tunnel P1.
- o VRF C-1 joins P-tunnel P1 in order to receive the (C-S1,C-G) traffic.
- o VRF C-1 requests the extranet multicast flow (C-S2,C-G). In the context of VRF C-1, C-S2 matches the route exported from VRF B-1. Thus, C-1's request to receive the (C-S2,C-G) flow is transmitted to VRF B-1.
- o VRF B-1 responds by transmitting its (C-S2,C-G) traffic on P-tunnel P2.
- o VRF C-1 joins P-tunnel P2 in order to receive the (C-S2,C-G) traffic.
- o VRF D-1, on PE2, requests the extranet multicast flow (C-S2,C-G). In the context of VRF D-1, C-S2 matches the route exported from VRF A-1. Thus, D-1's request to receive the (C-S2,C-G) flow is transmitted to VRF A-1.
- o VRF A-1 responds by transmitting its (C-S2,C-G) traffic on P-tunnel P1.
- o VRF D-1 joins P-tunnel P1 in order to receive the (C-S2,C-G) traffic.

In this example, VRF A-1 has chosen to use the same P-tunnel, P1, to carry both its (C-S2,C-G) traffic and the (C-S1,C-G) traffic. VRF C-1 has joined tunnel P1 in order to receive the (C-S1,C-G) traffic from VRF A-1, which means that VRF C-1 will also receive the unwanted (C-S2,C-G) traffic from P1. VRF C-1 is also expecting (C-S2,C-G) traffic from VRF B-1; this traffic will be received from P2. Thus, VRF C-1 is receiving (C-S2,C-G) traffic on both tunnels, and both C-flows arrive from the expected PE, PE1.

Therefore, it is necessary to EITHER (1) prevent the above scenario from occurring OR (2) ensure that VRF C-1 discards any (C-S,C-G) traffic that arrives from the wrong P-tunnel. See Section 2.3 for further discussion of this issue.

Note that the ambiguity described in this section (Section 2.2) would not occur if C-G were an (ASM) extranet C-group. In that case, the scenario would violate the rule, given previously in Section 2, requiring that all sources sending to a particular ASM extranet C-group must have addresses that are unambiguous over all the MVPNs receiving traffic for that C-group.

2.3. Preventing Misdelivery in These Scenarios

There are two ways to prevent the scenarios discussed in Sections 2.1 and 2.2 from resulting in misdelivery of data; these techniques are discussed in Sections 2.3.1 and 2.3.2, respectively.

2.3.1. Do Not Deliver Packets from the Wrong P-tunnel

Consider a particular C-flow that has receivers in a particular VRF. Sections 6 and 7 describe a set of procedures that enable an egress PE to determine the "expected P-tunnel" for that C-flow in the context of that VRF. If a PE receives packets of the C-flow (as determined by the IP source and/or destination address of the packet), it checks to see if the packet was received on the expected P-tunnel for that VRF. If so, the packet is delivered to the VRF (and thus to the C-flow's receivers in that VRF). If not, the packet is not delivered to the VRF.

Note that at a given egress PE, the wrong P-tunnel for one VRF may be the correct P-tunnel for another.

These procedures, if applied at every PE that joins a given P-tunnel, are sufficient to prevent misdelivery of traffic in the scenarios discussed in Sections 2.1 and 2.2.

IF these procedures cannot be applied by every PE that is attached to a given extranet, then the policies of Section 2.3.2 MUST be applied at every VRF containing C-sources for that extranet.

In some cases, however, it may be safe to deliver packets that arrive from other than the expected P-tunnel. Suppose that it is known that every packet gets transmitted on only a single P-tunnel. (This will be the case if the "single PMSI per C-flow" transmission model, discussed in Section 3.1, is being used.) Suppose also that it is known that T1 and T2 carry only packets that arrived at the same ingress PE, over one or more VRF interfaces that are associated with the same VRF (i.e., that there is a particular VRF that is the ingress VRF for ALL the packets carried by T1 or T2). In this case, if T1 is the expected P-tunnel for a given (C-S,C-G), it is NOT necessary to discard (S,G) packets that arrive over T2.

It is not always possible to determine whether two P-tunnels are carrying packets from the same ingress VRF. However, in some cases, this can be determined by examination of the A-D routes in which the tunnels have been advertised.

Consider the following example:

- o Tunnel T1 is a Point-to-Multipoint (P2MP) multipoint Label Distribution Protocol (mLDP) or RSVP-TE P-tunnel advertised in an Intra-AS I-PMSI A-D route (call it "R1").
- o Tunnel T2 is a P2MP mLDP or RSVP-TE P-tunnel advertised in an S-PMSI A-D route (call it "R2").
- o The respective NLRIs of R1 and R2 contain the same RD value.
- o The MPLS Label field of R1's PTA is zero, and the MPLS label value of R2's PTA is zero.

In this example, it can be concluded that T1 and T2 are carrying packets from the same ingress VRF. Thus, if T1 is the expected P-tunnel for a (C-S,C-G) flow, (S,G) packets from T2 can be safely delivered to the egress VRF; they do not need to be discarded. Similarly, if T2 is the expected P-tunnel for a (C-S,C-G) flow, (S,G) packets from T1 can be safely delivered to the egress VRF.

Another example is the following:

- o Tunnel T3 is a P2MP mLDP or RSVP-TE P-tunnel advertised in a (C-*,C-*) S-PMSI A-D route (call it "R3").
- o Tunnel T4 is a P2MP mLDP or RSVP-TE P-tunnel advertised in a (C-S,C-G) S-PMSI A-D route (call it "R4").
- o The respective NLRIs of R3 and R4 contain the same RD value.
- o The MPLS Label field of R3's PTA is zero, and the MPLS label value of R4's PTA is zero.

In this example, it can be concluded that T3 and T4 are carrying packets from the same ingress VRF. Thus, if T3 is the expected P-tunnel for a (C-S,C-G) flow, (S,G) packets from T4 can be safely delivered to the egress VRF; they do not need to be discarded. Similarly, if T4 is the expected P-tunnel for a (C-S,C-G) flow, (S,G) packets from T3 can be safely delivered to the egress VRF.

When Ingress Replication (IR) P-tunnels are being used, please see [MVPN-IR], especially Section 7 ("The PTA's 'MPLS Label' Field") for a discussion of how to determine when packets from other than the expected P-tunnel must be discarded.

2.3.2. Policies to Prevent Ambiguity on a P-Tunnel

For P-tunnels that are advertised in S-PMSI A-D routes whose NLRI contains (C-S,C-G) or (C-S,C-*), the ambiguities described in Sections 2.1 and 2.2 can be prevented by provisioning a policy that assigns, to such P-tunnels, only flows from the same C-source.

However, it is not always possible to determine, through inspection of the control messages, whether this policy has been deployed. For instance, suppose that (1) a given VRF has imported a set of S-PMSI A-D routes, (2) each route in the set has bound only a single (C-S1,C-G1) to a single P-tunnel, and (3) each route in the set identifies a different P-tunnel in its PTA than the P-tunnel identified by the PTA of any other route in the set. One cannot infer from this that there is no ambiguity, as the same P-tunnel may also have been advertised in an S-PMSI A-D route that is not imported by the given VRF, and that S-PMSI A-D route may have bound (C-S2,C-G2) to the P-tunnel, where C-S1 != C-S2.

Therefore, in order to determine that a given P-tunnel (advertised in a (C-S,C-G) or (C-S,C-*) S-PMSI A-D route) carries only C-flows from a single C-source, a PE must have a priori knowledge (through provisioning) that this policy has been deployed. In the remainder of this document, we will refer to this policy as the "single C-source per (C-S,C-G) or (C-S,C-*) P-tunnel" policy. Note that this policy is only applicable to P-tunnels that are advertised only in (C-S,C-G) or (C-S,C-*) S-PMSI A-D routes.

Of course, if a P-tunnel is advertised in (a) an I-PMSI A-D route, (b) an S-PMSI A-D route whose NLRI contains (C-*,C-*), or (c) an S-PMSI A-D route whose NLRI contains (C-*,C-G), then it is always possible for the P-tunnel to contain traffic from multiple C-sources; there is no policy that can prevent that.

However, if a P-tunnel advertised in a (C-*,C-G) S-PMSI A-D route contains only traffic addressed to a single C-G, the address uniqueness rules of Section 2 prevent the C-source addresses from being ambiguous; the set of C-sources transmitting to a particular extranet C-group address must be unambiguous over the set of MVPNs that have receivers for that C-group. So, for P-tunnels that are advertised in (C-*,C-G) S-PMSI A-D routes, the ambiguities described in Sections 2.1 and 2.2 can be prevented by provisioning a policy

that assigns to such P-tunnels only flows to the same extranet C-group. We will refer to this policy as the "single C-group per (C-*,C-G) P-tunnel" policy.

These considerations can be summarized as follows. IF the procedures referenced in Section 2.3.1 cannot be applied, then the PEs MUST be provisioned so that all of the following conditions hold true for the VRFs that contain extranet C-sources:

- o the "single C-source per (C-S,C-G) or (C-S,C-*) P-tunnel" policy is provisioned,
- o either no (C-*,C-G) S-PMSI A-D routes are advertised or the "single C-group per (C-*,C-G) P-tunnel" policy is provisioned,
- o no P-tunnels are advertised in I-PMSI A-D routes, and
- o no (C-*,C-*) S-PMSI A-D routes are advertised.

Section 3 of this document describes a procedure known as "extranet separation". When extranet separation is used, the ambiguity described in Section 2.1 is prevented. However, the ambiguity described in Section 2.2 is not prevented by extranet separation. Therefore, the use of extranet separation is not a sufficient condition for avoiding the use of the procedures discussed in Section 2.3.1. Extranet separation is, however, implied by the policies discussed in this section (Section 2.3.2).

3. Extranet Transmission Models

This document specifies several "extranet transmission" models. A given VRF containing extranet C-sources or C-receivers MUST use only one of these models. Further, if VRF-S contains extranet C-sources, VRF-R contains extranet C-receivers, and it is allowed by policy for an extranet C-receiver in VRF-R to receive a C-flow from an extranet C-source in VRF-S, then VRF-S and VRF-R MUST use the same extranet transmission model. The model used by a given VRF is determined by provisioning.

3.1. Transmitting an Extranet C-Flow on a Single PMSI

In one extranet transmission model, which we call the "transmitting an extranet C-flow on a single PMSI" model or, more simply, the "single PMSI per C-flow" model, a PE transmitting a packet of an extranet C-flow transmits it on only a single PMSI. If the PMSI is instantiated by a multicast P-tunnel, this means that the PE transmits the packet on a single P-tunnel. Of course, if the PE is a replication point for that multicast P-tunnel, the packet is

transmitted more than once by the PE. Similarly, if the PMSI is instantiated by IR, each packet may be transmitted multiple times. It is still the case, though, that the packet is transmitted only on one PMSI.

This document provides procedures for supporting this transmission model using either BGP or PIM as the PE-PE C-multicast control protocol.

There are two variants of this transmission model: "without extranet separation" and "with extranet separation".

3.1.1. Without Extranet Separation

In this variant, multicast data traffic from extranet C-sources and from non-extranet C-sources may be carried in the same P-tunnel.

This document provides procedures for supporting this variant using either BGP or PIM as the PE-PE C-multicast control protocol.

3.1.2. With Extranet Separation

In this variant, multicast data traffic from extranet C-sources and from non-extranet C-sources are never carried in the same P-tunnel. Under certain circumstances, this can reduce the amount of multicast data traffic that is delivered unnecessarily to certain PE routers. It also eliminates the ambiguity discussed in Section 2.1.

By definition, when extranet separation is used, the following rule MUST be applied:

Traffic from extranet C-sources MUST NOT be carried in the same P-tunnel as traffic from non-extranet C-sources.

This rule does not impact those VRFs that contain only non-extranet C-sources, nor does it impact those VRFs that contain only extranet C-sources. However, if a particular VRF contains both kinds of C-sources, it will need to advertise some P-tunnels that are used for carrying only extranet C-flows and some that are used only for carrying non-extranet C-flows.

This document provides procedures for supporting extranet separation when BGP is used as the PE-PE C-multicast control protocol. Support for extranet separation using PIM as the PE-PE C-multicast control protocol is outside the scope of this document.

3.2. Transmitting an Extranet C-Flow over Multiple PMSIs

The second extranet transmission model is called the "transmitting an extranet C-flow over multiple PMSIs" model or, more simply, the "multiple PMSIs per C-flow" model. In this model, a PE may transmit the packets of an extranet C-flow on several different PMSIs.

Support for extranet separation with this model is outside the scope of this document.

This document provides procedures for supporting this transmission model when PIM is used as the PE-PE C-multicast control protocol. Support for this transmission model when BGP is used as the PE-PE C-multicast control protocol is outside the scope of this document.

4. Distribution of Routes That Match C-S/C-RP Addresses

4.1. UMH-Eligible Routes

As described in Section 5.1 of [RFC6513], in order for a C-flow (C-S,C-G) to be carried across the SP backbone, a VRF that has multicast receivers for that C-flow must import a route that matches C-S, and this route must be "eligible for UMH selection". In this document, we will refer to these routes as "UMH-eligible extranet C-source routes".

The UMH-eligible extranet C-source routes do not necessarily have to be unicast routes; they MAY be SAFI 129 routes (see Section 5.1.1 of [RFC6513]). For example, suppose that one wants a VPN-R C-receiver to be able to receive extranet C-flows from C-sources in VPN-S but does not want any VPN-R system to be able to send unicast traffic to those C-sources. One can achieve this by using SAFI 129 routes as the UMH-eligible routes exported from VPN-S and imported by VPN-R. Since SAFI 129 routes are used only for UMH determination and not for unicast routing, this allows the multicast traffic to be forwarded properly but does not create unicast routes to the C-sources.

If a customer is using PIM-SM in the ASM model and one or more customer sites have C-receivers that are allowed by policy to join a (C-*,C-G) tree, where C-G is an extranet C-group, then any VRF with C-receivers for that group MUST import a UMH-eligible route that matches C-RP, where C-RP is the Rendezvous Point (RP) address for C-G.

The UMH-eligible extranet C-source and C-RP routes do not have to be "host routes". That is, they can be routes whose IPv4 address prefixes are not 32 bits in length or whose IPv6 address prefixes are not 128 bits in length. So, it is possible for a UMH-eligible

extranet C-source route to match the address of an extranet C-source and to also match the address of a non-extranet C-source. However, if such a route is exported from a VPN-S VRF and imported by a VPN-R VRF, VPN-R receivers will be able to receive C-flows from any non-extranet C-sources whose addresses match that route. To prevent this, the VPN-S VRF SHOULD be provisioned such that it will NOT export a UMH-eligible route that matches (in the context of the VPN-R VRF) both extranet C-sources and non-extranet C-sources. Failure to follow this rule may result in a VPN security violation. (See Section 10.)

In general, one does not want ALL the routes from the VPN-S VRFs to be exported to all the VPN-R VRFs, as only a subset of the routes in the VPN-S VRFs will be UMH-eligible extranet C-source routes. Route distribution is, as is always the case for a BGP/MPLS IP VPN [RFC4364], controlled by Route Targets (RTs). A variety of route distribution policies can be created by appropriately provisioning the import and export RTs of the various VRFs.

For example, the VPN-S VRFs that contain extranet C-sources could be configured to apply an export RT whose value is "RT-A-extranet" to the routes that match the extranet C-sources. The VPN-R VRFs that contain extranet C-receivers allowed to receive extranet C-flows from VPN-S extranet C-sources could then be configured with "RT-A-extranet" as an import RT.

Arbitrarily complex policies can be created by suitable manipulation of the import and export RTs.

4.1.1. Extranet Separation

If extranet separation is being used and a given VRF is exporting UMH-eligible routes for both extranet C-sources and non-extranet C-sources, then the VRF MUST be configured not only with its default RD but also with an extranet RD. The exported UMH-eligible routes MUST contain the extranet RD in their NLRIs.

4.2. Distribution of Unicast Routes Matching C-RPs and DRs

Consider a C-source, C-S, that may transmit to a particular extranet C-group, C-G.

In order to follow the procedures of [RFC7761],

- o The "first-hop designated router (DR)" for C-S needs to be able to unicast PIM Register messages to a C-RP that services C-G.
- o The C-RPs servicing C-G need to be able to unicast PIM Register-Stop messages to the DR for C-S.

It follows that if a VRF contains C-S but does not contain a C-RP for C-G, then the VRF MUST import a unicast route matching a C-RP for C-G. Note that the unicast route matching the C-RP is needed whether or not the VRF has also imported a SAFI 129 route matching the C-RP. (If the VRF also contains receivers for C-G and UMH determination is being done using SAFI 129 routes, both a unicast route and a SAFI 129 matching C-RP route are needed.)

Similarly, if a VRF contains a C-RP for C-G but does not contain C-S, the VRF MUST import a unicast route matching the DR for C-S. Note that the unicast route matching the DR for C-S is needed even if UMH determination is being done using SAFI 129 routes; in that case, if the VRF also contains receivers for C-G, it needs to import a SAFI 129 route matching C-S and a unicast route matching the DR for C-S.

If, for a particular extranet C-group, C-G, the customer is using "anycast-RP" [RFC3446] [RFC4610] or the Multicast Source Discovery Protocol (MSDP) [RFC3618], then all the C-RPs serving C-G need to send unicast messages to each other. Thus, any VRF that contains a C-RP for C-G needs to import unicast routes matching ALL the other C-RPs that serve C-G.

The need to distribute these unicast routes is usually not a problem as long as all the C-sources and C-RPs for C-G are in the same MVPN. If, however, the C-sources are not all in the same MVPN, great care must be taken to ensure that the unicast routes mentioned above are properly distributed.

There may be scenarios in which all the C-sources for C-G are in the same MVPN, but there are receivers in different VPNs, and some or all of the VPNs with receivers have their own C-RPs for C-G. In this case, care must be taken to ensure that the C-RPs can all unicast to each other.

4.3. Route Targets and Ambiguous UMH-Eligible Routes

This section imposes a constraint on the way RTs are assigned to (a) UMH-eligible routes and (b) the BGP A-D routes that advertise P-tunnels (i.e., BGP A-D routes that contain a PTA). The constraint specified here applies to any extranet for which the ambiguity described in Section 2.2 is possible. (The conditions under which such ambiguity is possible are also described in Section 2.2.)

We want to ensure that, in any given VRF, the UMH-eligible route matching a given extranet C-source has an RT in common with every BGP A-D route that advertises a P-tunnel that may be used to carry extranet multicast traffic from that C-source. We also want to ensure that the UMH-eligible route matching a given extranet C-source does not have any RT in common with any BGP A-D route that advertises a P-tunnel that may be used to carry any multicast traffic from a different C-source that has the same IP address. This enables us to determine whether traffic that appears to be from the given C-source is really arriving on the wrong P-tunnel and hence is really from a different C-source with the same IP address.

Suppose that an IP address C-S is used in VPN-A as the address of one system and used in VPN-B as the address of a different system. In this case, one or more VPN-A VRFs may export a VPN-IP route whose NLRI is $\langle RD1, S \rangle$, and one or more VPN-B VRFs may export a VPN-IP route whose NLRI is $\langle RD2, S \rangle$, where $RD1 \neq RD2$. Consider two routes -- R1 and R2 -- for which the following conditions all hold:

- o R1 and R2 are UMH-eligible extranet C-source or C-RP routes, or are unicast routes matching a C-RP.
- o R1 is exported from a VRF of VPN-A, while R2 is exported from a VRF of a different VPN, say VPN-B.
- o R1's NLRI specifies IP address prefix S/n.
- o R2's NLRI specifies IP address prefix S/m.
- o $m \geq n$ (S/m is either the same as or more specific than S/n).
- o There is some host address H such that:
 - * H denotes a different system in VPN-A than in VPN-B, and
 - * $H/m == S/m$ (so either S/m or S/n might be a longest match for H in some VRF).

We impose the following constraint: RTs MUST be assigned in such a way that R1 and R2 do not have any RT in common.

(This constraint is not as onerous as it may seem. Typically, R1 and R2 would not have an RT in common, as that might result in their being imported into the same VRF, making the address H ambiguous in that VRF.)

Sections 6 and 7 specify procedures for determining if a received C-flow has been received over the expected P-tunnel. Those procedures will not work if this constraint is violated. (The constraint described in this section is necessary, but not sufficient, for the procedures of Sections 6 and 7 to work; additional constraints that cover the assignment of RTs to BGP A-D routes are given in subsequent sections.)

4.4. Dynamically Marking Extranet Routes

4.4.1. The Extranet Source Extended Community

Sections 4.1, 4.2, and 4.3 place specific requirements on the way in which certain VPN-IP routes are distributed. In order to ensure that these requirements are met, a VPN customer must tell its SP which routes are the matching routes for extranet C-sources and C-RPs. This may be done as part of the provisioning process. Note that this does not necessarily require customer/provider interaction every time the customer adds a new extranet C-source or C-RP, but only when the IP address of the new C-source or C-RP does not match an existing route that is already being distributed as a VPN-IP extranet route. Nevertheless, it seems worthwhile to support an OPTIONAL mechanism that allows a customer to dynamically mark certain routes as being extranet routes.

To facilitate this, we define a new Transitive Opaque Extended Community (see [RFC4360], [RFC7153], and Section 9 of this document): the Extranet Source Extended Community. When a Customer Edge (CE) router advertises (via BGP) a route to a PE router and the AFI/SAFI of the route is 1/1, 1/2, 1/4, 2/1, 2/2, or 2/4, the Extranet Source Extended Community MAY be attached to the route. The value field of the Extended Community MUST be set to zero. By placing this Extended Community on a particular route, a CE router indicates to a PE router that the procedures of Sections 4.1, 4.2, and 4.3 are to be applied to that route. That is, the CE router may use this Extended Community to indicate to the PE router that a particular route is to be treated as a route that matches the address of an extranet source and is to be exported accordingly to other VPNs. A PE router that interprets this Extended Community MUST ignore the contents of the value field.

Whether a CE router uses the Extranet Source Extended Community is determined by the configuration of the CE router. If used, the set of routes to which the Extended Community is attached is also determined by configuration of the CE. Note that a particular PE router may or may not support the use of the Extranet Source Extended Community by a particular CE router; this is determined by the service agreement between the SP and its customer.

If a CE is advertising SAFI 2 routes to the PE as the UMH-eligible extranet C-source and C-RP routes and the CE is using the Extranet Source Extended Community, it is important that the CE attach that Extended Community to the SAFI 2 routes, rather than just to the corresponding SAFI 1 routes. Otherwise, extranet receivers may not be able to join the (C-S,C-G) or (C-*,C-G) multicast trees.

However, if the C-sources and the C-RPs for a given extranet C-group are not all in the same VPN, the Extended Community would also have to be attached to the SAFI 1 routes that match the C-RP addresses and to the SAFI 1 routes that match the addresses of the first-hop designated routers for all the C-sources. Otherwise, the first-hop routers might not be able to send PIM Register messages to the C-RPs, and the C-RPs might not be able to send PIM Register-Stop messages to the first-hop routers.

While this Extended Community allows a customer to inform the SP dynamically that certain routes are "extranet routes", it does not allow a customer to control the set of RTs that the route will carry when it is redistributed as a VPN-IP route. Thus, it is only useful when all the extranet routes from a given VRF are exported with exactly the same set of RTs. (cf. Section 4.3.1 of [RFC4364], which does provide a mechanism that, if properly supported by the SP, allows the customer to determine the set of RTs carried by a VPN-IP route.) A CE SHOULD NOT attach the Extranet Source Extended Community to any route for which it uses another method of specifying the RTs to be carried by that route. A CE SHOULD NOT attach the Extranet Source Extended Community to a route unless all the extranet routes from the CE's VPN are intended to carry the same set of RTs.

A PE SHOULD ignore the Extranet Source Extended Community if it appears on a route that the CE should not have put it on. A PE that ignores the Extranet Source Extended Community SHOULD NOT follow the procedures of Section 4.4.2.

Note that misconfiguration on the CE router can result in the Extranet Source Extended Community being mistakenly attached to a route that is not intended to be exported as an extranet route. This could result in a VPN security violation.

4.4.2. Distribution of Extranet Source Extended Community

Suppose that a PE receives from a CE a route (call it "R") with the Extranet Source Extended Community. The PE must determine (via the considerations discussed in Section 4.4.1) whether it should ignore that Extended Community on route R; if it should ignore the Extended Community, the procedures described in this section are not followed.

Otherwise, when the PE originates a VPN-IP route corresponding to route R, the PE MUST attach this Extended Community to that route.

A Route Reflector MUST NOT add or remove the Extranet Source Extended Community from the VPN-IP routes reflected by the Route Reflector, including the case where VPN-IP routes received via Internal BGP (IBGP) are reflected to External BGP (EBGP) peers (inter-AS option (c); see Section 10 of [RFC4364]). The value of the Extended Community MUST NOT be changed by the Route Reflector.

When re-advertising VPN-IP routes, Autonomous System Border Routers (ASBRs) MUST NOT add/remove the Extranet Source Extended Community from these routes. This includes inter-AS options (b) and (c) (see Section 10 of [RFC4364]). The value of the Extended Community MUST NOT be changed by the ASBRs.

When a PE advertises (via BGP) IP routes to a CE, these routes MUST NOT carry the Extranet Source Extended Community unless the PE-CE connection is actually an inter-AS option (a) connection (see Section 10 of [RFC4364]). When the PE-CE connection is not an inter-AS option (a) connection, a CE that receives an IP route with the Extranet Source Extended Community MUST remove it from the route before re-advertising the route.

The rules for attaching the Extranet Source Extended Community to a VPN-IP route, and the rules for propagating that Extended Community, are needed in order to support the scenario in which a VPN contains an option (a) interconnect (see Section 10 of [RFC4364]). At the option (a) interconnect, the VPN-IP route gets translated back to an IP route, and the RTs are stripped off before the IP route is propagated. If the Extranet Source Extended Community has also been stripped off, there is no way for the router at the other end of the option (a) interconnect to know that the route represents an extranet source. Thus, the technique of using the Extranet Source Extended Community to dynamically signal that a particular route represents an extranet source will not work correctly across an option (a) interconnect unless the rules in this section are followed.

4.5. The Extranet Separation Extended Community

We define a new Transitive Opaque Extended Community: the Extranet Separation Extended Community (see [RFC4360], [RFC7153], and Section 9 of this document). This Extended Community is used only when extranet separation is being used. Its value field MUST be set to zero upon origination, MUST be ignored upon reception, and MUST be passed unchanged by intermediate routers. A Route Reflector MUST NOT add or remove the Extranet Separation Extended Community from the routes it reflects, including the case where routes received via IBGP are reflected to EBGPeers (inter-AS option (c); see Section 10 of [RFC4364]).

If a VRF has been provisioned to use extranet separation and that VRF has been provisioned to transmit any extranet C-flows on a P-tunnel that it advertises in an I-PMSI A-D route or a (C-*,C-*) S-PMSI A-D route, then any UMH-eligible routes that are exported from that VRF following the procedures of Sections 4.1, 4.2, and 4.3 MUST carry the Extranet Separation Extended Community. In addition, if an I-PMSI A-D route and/or (C-*,C-*) S-PMSI A-D route exported from that VRF is used to carry extranet traffic, that A-D route MUST also carry the Extranet Separation Extended Community. Further details may be found in Sections 7.3, 7.4.4, and 7.4.5.

5. Origination and Distribution of BGP A-D Routes

Except where otherwise specified, this section describes procedures and restrictions that are independent of the PE-PE C-multicast control protocol.

5.1. Route Targets of UMH-Eligible Routes and A-D Routes

Suppose that there is an extranet C-flow such that:

- o The extranet C-source of that C-flow is in VRF A-1.
- o One or more extranet C-receivers of that C-flow are in VRF B-1.

In this case, VRF A-1 MUST export a UMH-eligible route that matches the extranet C-source address, and VRF B-1 MUST import that route. In addition, VRF A-1 MUST export an Intra-AS I-PMSI A-D route or an S-PMSI A-D route specifying the P-tunnel through which it will send the data traffic of the given extranet C-flow, and VRF B-1 MUST import that route. If BGP is the PE-PE C-multicast control protocol, then under certain conditions (as specified in [RFC6514]), VRF A-1 may also need to export a Source Active A-D route specifying that it contains a source of the given C-flow, and VRF B-1 must import that Source Active A-D route. That is, in order for VRF B-1 to receive a

C-flow from a given extranet C-source contained in VRF A-1, VRF A-1 MUST export a set of A-D routes that are "about" that source, and VRF B-1 MUST import them.

One way to ensure this is to provision an RT that is carried by all the routes exported from VRF A-1 that are "about" a given extranet C-source and also provision this RT as an import RT at any VRF (such as VRF B-1) that is allowed to receive extranet flows from that source.

If the "single PMSI per C-flow" transmission model is being used (with or without extranet separation), there is an additional requirement, stated below, regarding the way RTs are provisioned, as the RTs carried by a UMH-eligible route that matches a given extranet C-source may need to be used to identify the A-D routes that are "about" that source.

Consider the following scenario:

- o IP address S is the address of one system in VPN-A and the address of a different system in VPN-B.
- o VRF A-1 on PE1 exports UMH-eligible route R1, which is a matching route for S.
- o VRF A-1 on PE1 exports an A-D route P1 whose PTA identifies a P-tunnel through which VRF A-1 may send traffic whose C-source is S, where one of the following conditions holds:
 - * P1 is an I-PMSI A-D route, OR
 - * P1 is an S-PMSI A-D route whose NLRI contains (C-*,C-*) or (C-*,C-G), OR
 - * P1 is an S-PMSI A-D route whose NLRI contains (C-S,C-G) or (C-S,C-*), BUT the "single C-source per (C-S,C-G) or (C-S,C-*) P-tunnel" policy is not provisioned, OR
 - * P1 is a Source Active A-D route whose NLRI contains (C-S,C-G).

- o VRF B-1 on PE1 exports a UMH-eligible route R2, which is a matching route for S.
- o VRF B-1 on PE1 exports an A-D route P2 whose PTA identifies a P-tunnel on which VRF B-1 may send traffic whose C-source is S, where one of the following conditions holds:
 - * P2 is an I-PMSI A-D route, OR
 - * P2 is an S-PMSI A-D route whose NLRI specifies (C-*,C-*) or (C-*,C-G), OR
 - * P2 is an S-PMSI A-D whose NLRI specifies (C-S,C-G) or (C-S,C-*), BUT the "single C-source per (C-S,C-G) or (C-S,C-*) P-tunnel" policy is not provisioned, OR
 - * P2 is a Source Active A-D route whose NLRI contains (C-S,C-G).

As implied by the rules of Section 4.1, there MUST NOT be any RT that is common to both R1 and R2. In addition, the following set of rules for RT assignment MUST be followed when extranets are supported. These rules support all the extranet transmission models described in this specification:

- o There MUST NOT be any RT that is carried by both P1 and P2.
- o The intersection of the set of RTs carried by P1 and the set of RTs carried by R1 MUST be non-null, and any VRF that imports both P1 and R1 MUST be configured with an import RT from this intersection.
- o The intersection of the set of RTs carried by P2 and the set of RTs carried by R2 MUST be non-null, and any VRF that imports both P2 and R2 MUST be configured with an import RT from this intersection.

Suppose that VRF C-1 on PE2 imports P1 and R1 from VRF A-1 while also importing P2 from VRF B-1. Since

- o R1 is VRF C-1's route to S,
- o R1 has an RT in common with P1, and
- o R1 has no RT in common with P2,

it can be concluded that VRF C-1 should expect that multicast traffic from S will arrive on the P-tunnel specified in P1. See Sections 6 and 7 for more details on determining the expected P-tunnel for a given extranet C-flow.

While the assignment of import and export RTs to routes is a deployment and provisioning issue rather than a protocol issue, it should be understood that failure to follow these rules is likely to result in VPN security violations.

5.2. Considerations for Particular Inclusive Tunnel Types

An Inclusive Tunnel (sometimes referred to as an "Inclusive Tree"; see Section 2.1.1 of [RFC6513]) is a tunnel that, by default, carries all the multicast traffic of a given MVPN that enters the backbone network via a particular PE. An Inclusive Tunnel is advertised in the PTA of an I-PMSI A-D route.

5.2.1. RSVP-TE P2MP or Ingress Replication

This section applies when Inclusive Tunnels are created using either RSVP-TE P2MP or IR.

Suppose that a VRF, say VRF-S, contains a given extranet C-source C-S, and VRF-S advertises in its Intra-AS I-PMSI A-D route either a P2MP RSVP-TE P-tunnel or an IR P-tunnel to carry extranet traffic.

In order for VRF-S to set up the P2MP RSVP-TE or IR P-tunnel, it must know all the PEs that are leaf nodes of the P-tunnel, and to learn this it must import an Intra-AS I-PMSI A-D route from every VRF that needs to receive data through that tunnel.

Therefore, if VRF-R contains an extranet C-receiver that is allowed by policy to receive extranet flows from C-S, the RT(s) carried by the Intra-AS I-PMSI A-D routes originated by VRF-R MUST be such that those Intra-AS I-PMSI A-D routes will be imported into VRF-S.

In the case of IR, this has the following consequence: if an egress PE has n VRFs with receivers for a flow that VRF-S transmits on its I-PMSI, that egress PE will receive n copies of the same packet, one for each of the n VRFs.

Note that Section 9.1.1 of [RFC6514] prohibits the "Leaf Information Required" flag from being set in the PTA of an Intra-AS I-PMSI A-D route. If this prohibition is ever removed, the requirement of this section will apply only if VRF-S does not set that flag.

5.2.2. Ingress Replication

This section applies only when Inclusive Tunnels are created via IR.

[RFC6513] and [RFC6514] specify procedures that allow I-PMSIs to be instantiated by IR. The concept of an IR P-tunnel, and the procedures for supporting IR P-tunnels, are explained more fully in [MVPN-IR]. An IR P-tunnel can be thought of as a P2MP tree in which a packet is transmitted from one node on the tree to another by being encapsulated and sent through a unicast tunnel.

As discussed in Section 2, when I-PMSIs are used to support extranets, egress PEs MUST have the ability to discard customer multicast data packets that arrive on the wrong P-tunnel. When I-PMSIs are instantiated by IR, this implies that the following two procedures MUST be followed:

1. One of the following three procedures MUST be followed:
 - a. the "Single Forwarder Selection" procedures of Section 9.1.2 of [RFC6513]
 - b. the "native PIM methods" of Section 9.1.3 of [RFC6513]
 - c. the unicast encapsulation used to transmit packets along the IR P-tunnel is such as to enable the receiving node to identify the transmitting node (note that this would not be the case if, for example, the unicast tunnels were MP2P LSPs)

and

2. If a PE assigns an MPLS label value in the PTA of an Intra-AS or Inter-AS I-PMSI A-D route that it originates, that label value MUST NOT appear in the PTA of any other I-PMSI or S-PMSI A-D route originated by the same PE.

Failure to follow these procedures would make it impossible to discard packets that arrive on the wrong P-tunnel and thus could lead to duplication of data.

If it is desired to support extranets while also using IR to instantiate the PMSIs, an alternative is to use (C-*,C-*) S-PMSIs instead of I-PMSIs. (See [RFC6625], as well as Sections 7.2.2, 7.3.2, and 7.4.4 of this document.) This has much the same effect in the data plane, and there are no restrictions on the type of unicast tunnel that can be used for instantiating S-PMSIs.

Section 6.4.5 of [RFC6513] describes a way to support VPNs using I-PMSIs that are instantiated by IR, using no S-PMSIs, but using "explicit tracking" to ensure that a C-flow goes only to egress PEs that have receivers for it. This document does not provide procedures to support extranets using that model.

6. When PIM Is the PE-PE C-Multicast Control Plane

As specified in [RFC6513], when PIM is used as the PE-PE C-multicast control plane for a particular MVPN, there is a "Multidirectional Inclusive Provider Multicast Service Interface" (MI-PMSI) for that MVPN, and all the PEs of that MVPN must be able to send and receive on that MI-PMSI. Associated with each VRF of the MVPN is a PIM C-instance, and the PIM C-instance treats the MI-PMSI as if it were a LAN interface. That is, the "ordinary" PIM procedures run over the MI-PMSI just as they would over a real LAN interface, except that the data-plane and control-plane "Reverse Path Forwarding (RPF) checks" need to be modified. Section 5.2 of [RFC6513] specifies the RPF check modifications for non-extranet MVPN service.

For example, suppose that there are two VPNs: VPN-S and VPN-R. In the absence of extranet support, all the VRFs of VPN-S are connected via one MI-PMSI (call it "the VPN-S MI-PMSI"), and all the VRFs of VPN-R are connected via another ("the VPN-R MI-PMSI"). If we want to provide extranet service in which the extranet C-sources are attached to some set of VPN-S VRFs while the extranet C-receivers are attached to some set of VPN-R VRFs, then we have two choices:

1. either the VPN-R VRFs need to join the VPN-S MI-PMSI, or
2. the VPN-S VRFs need to join the VPN-R MI-PMSI.

The first choice is used to support the "single PMSI per C-flow" transmission model. The second choice is used to support the "multiple PMSIs per C-flow" transmission model.

Procedures for both models are described below.

To support these models, it must be possible to determine which I-PMSI A-D routes are associated with the VPN-S I-PMSI and which I-PMSI A-D routes are associated with the VPN-R I-PMSI. Procedures are given for assigning RTs to these routes in a way that makes this determination possible.

Both models allow the use of S-PMSIs to carry multicast data traffic. If a VRF containing receivers can receive from multiple MI-PMSIs, each S-PMSI must be uniquely associated with a particular MI-PMSI. Procedures are given for assigning RTs to these routes in a way that makes this determination possible.

All the procedures specified in Sections 3, 4, and 5 still apply.

Note that there are no special extranet procedures for Inter-AS I-PMSI A-D routes or for Leaf A-D routes. Source Active A-D routes are not used when PIM is the PE-PE C-multicast protocol.

6.1. Provisioning VRFs with RTs

6.1.1. Incoming and Outgoing Extranet RTs

In the absence of extranet service, suppose that each VRF of a given VPN (call it "VPN-S") is configured with RT-S as its import and export RT, and that each VRF of a second VPN (call it "VPN-R") is configured with RT-R as its import and export RT. We will refer to RT-S and RT-R as "non-extranet RTs".

Now suppose that VPN-S contains some extranet C-sources and VPN-R contains some extranet C-receivers that are allowed by policy to receive extranet C-flows from the VPN-S extranet C-sources.

To set up this S-to-R extranet, provisioning an additional RT (call it "RT-S-to-R") whose value is, in general, distinct from RT-S and RT-R is REQUIRED.

A VPN-S VRF that contains extranet C-sources allowed to transmit to VPN-R MUST be configured with RT-S-to-R as an "Outgoing Extranet RT".

A VPN-R VRF that contains extranet C-receivers allowed to receive packets from VPN-S MUST be configured with RT-S-to-R as an "Incoming Extranet RT".

Note that the terms "Incoming" and "Outgoing" in this context refer to the direction of multicast data packets relative to the VRF.

The Incoming Extranet RTs and Outgoing Extranet RTs that are configured for a given VRF serve as import RTs for that VRF. They also serve as export RTs, but only for specific routes as specified in Section 6.1.2 below.

Note that any VRF that contains both extranet C-sources and extranet C-receivers MUST be configured with both Outgoing Extranet RTs and Incoming Extranet RTs.

A VRF MAY be configured with more than one Incoming Extranet RT and/or Outgoing Extranet RT.

If it happens to be the case that all C-sources in VPN-S are extranet C-sources allowed to transmit to VPN-R, then VPN-S VRFs MAY be configured such that RT-S is both a non-extranet RT and an Outgoing Extranet RT, and VPN-R VRFs MAY be configured such that RT-S is an Incoming Extranet RT.

6.1.2. UMH-Eligible Routes and RTs

Suppose that R1 is a route exported from a VPN-S VRF and matching an extranet C-source that is allowed by policy to transmit to VPN-R. In that case, R1 MUST carry the Outgoing Extranet RT used for the S-to-R extranet. This will cause the route to be imported into the VPN-R VRFs that have extranet C-receivers that are allowed by policy to receive from VPN-S.

The rules of Section 4 regarding RTs and ambiguous addresses still apply.

6.1.3. PIM C-Instance Reverse Path Forwarding Determination

Suppose that a PIM control message (call it "M") is received by a given VRF (call it "VRF-V") from a particular P-tunnel T. In order to process control message M, the PIM C-instance associated with VRF-V may need to do an "RPF determination" (see Section 5.2.2 of [RFC6513]) for a particular IP prefix S. RPF determination is based upon the rules for UMH selection as specified in Section 5.1 of [RFC6513].

This document specifies an additional constraint on the UMH selection procedure. When doing RPF determination for a PIM control message received over a P-tunnel, a route matching prefix S is not considered to be eligible for UMH selection unless there is an RT (call it "RT1"), configured as one of VRF-V's Outgoing Extranet RTs, such that the following two conditions both hold:

1. The route matching S is exported from VRF-V carrying RT1, and
2. An I-PMSI A-D route advertising P-tunnel T (in its PTA) has been imported into VRF-V, and that I-PMSI A-D route carries RT1.

6.2. "Single PMSI per C-Flow" Model

In this model, if a VPN-S VRF has extranet multicast C-sources and a VPN-R VRF has extranet multicast C-receivers allowed by policy to receive from the C-sources in the VPN-S VRF, then the VPN-R VRF joins the MI-PMSI that VPN-S uses for its non-extranet traffic.

6.2.1. Forming the MI-PMSIs

Consider a VPN-S VRF that has extranet C-sources. Per [RFC6513], each VPN-S VRF must originate an Intra-AS I-PMSI A-D route containing a PTA specifying the P-tunnel to be used as part of the VPN-S MI-PMSI. In the absence of extranet service, this route carries the VRF's non-extranet RT, RT-S. When extranet service is provided (using the "single PMSI per C-flow" model), this route MUST also carry each of the VRF's Outgoing Extranet RTs.

Consider a VPN-R VRF that has extranet C-receivers. Per [RFC6513], each VPN-R VRF must originate an Intra-AS I-PMSI A-D route containing a PTA specifying the P-tunnel to be used as part of the VPN-R MI-PMSI. This route carries the VRF's non-extranet RT, RT-R. When extranet service is provided (using the "single PMSI per C-flow" model), the VPN-R VRF MUST also originate one or more additional Intra-AS I-PMSI A-D routes. It MUST originate one additional Intra-AS I-PMSI A-D route for each Incoming Extranet RT with which it has been configured; each such route will carry exactly one of the configured Incoming Extranet RTs.

Note that when a VRF originates more than one Intra-AS I-PMSI A-D route, each of them MUST contain a different RD in its NLRI. In addition, we add the requirement that any pair of such routes MUST NOT contain an RT in common.

A VRF with extranet C-sources MUST join the P-tunnels advertised in the imported I-PMSI A-D routes that carry its non-extranet RT or any of its Outgoing Extranet RTs. This set of P-tunnels will be treated as instantiating a single MI-PMSI; the associated PIM C-instance will treat that MI-PMSI as a single LAN and will run PIM procedures on that LAN, as specified in [RFC6513]. The fact that the MI-PMSI attaches to VRFs of different VPNs is not known to the PIM C-instance of the VRF containing the sources.

A VRF with extranet C-receivers MUST join the P-tunnels advertised in all the imported I-PMSI A-D routes. The set of P-tunnels advertised in the I-PMSI A-D routes that carry a particular Incoming Extranet RT are treated as instantiating a particular MI-PMSI. So, a VRF with C-receivers will "see" several MI-PMSIs, one corresponding to the non-extranet, and as many as one for each configured Incoming Extranet RT. The PIM C-instance associated with the VRF will treat each of these MI-PMSIs as a separate LAN interface.

As an example, suppose that:

- o All VPN-R VRFs are configured with RT-R as a non-extranet import and export RT, and
- o VPN-R VRFs with extranet receivers are configured with RT-S-to-R as an Incoming Extranet RT, and
- o VPN-S VRFs with extranet transmitters are configured with:
 - * RT-S as a non-extranet import and export RT
 - * a list of IP addresses that are the addresses of the extranet sources
 - * RT-S-to-R as an Outgoing Extranet RT

VPN-S VRFs will then export UMH-eligible routes matching extranet C-sources, and these routes will carry both RT-S and RT-S-to-R. Each VPN-S VRF will also export an Intra-AS I-PMSI A-D route that carries both RT-S and RT-S-to-R.

VPN-R VRFs will originate and export two Intra-AS I-PMSI A-D routes: one carrying RT-R and one carrying RT-S-to-R. The Intra-AS I-PMSI A-D route with RT-S-to-R will be imported into the VPN-S VRFs.

VPN-R will regard all the I-PMSI A-D routes it has exported or imported with RT-S-to-R as part of a single MI-PMSI. VPN-R will regard all the I-PMSI A-D routes it has exported or imported with RT-R as part of a second MI-PMSI. The PIM C-instance associated with

a VPN-R VRF will treat the two MI-PMSIs as two separate LAN interfaces. However, the VPN-S VRFs will regard all the I-PMSI A-D routes imported with RT-S or RT-S-to-R as establishing only a single MI-PMSI. One can think of this as follows: the VPN-R VRFs have joined the VPN-S MI-PMSI as well as the VPN-R MI-PMSI.

Extranets consisting of more than two VPNs are easily supported as follows. Suppose that there are three VPNs: VPN-A, VPN-B, and VPN-C. VPN-A and VPN-B have extranet C-sources, and VPN-C contains receivers for both VPN-A extranet C-sources and VPN-B extranet C-sources. In this case, the VPN-C VRFs that have receivers for both VPN-A and VPN-B sources may be provisioned as follows. These VPN-C VRFs may be provisioned with RT-C as a non-extranet RT, and with RT-A-to-C and RT-B-to-C as Incoming Extranet RTs. In this case, the VPN-C VRFs that are so provisioned will originate three Intra-AS I-PMSI A-D routes (each with a different RD in its NLRI), each of which carries exactly one of the three RTs just mentioned. The VPN-B VRFs with extranet C-sources will be provisioned with RT-B-to-C as an Outgoing Extranet RT, and the VPN-A VRFs will be provisioned with RT-A-to-C as an Outgoing Extranet RT. The result will be that the PIM C-instance associated with a VPN-C VRF will see three LAN interfaces: one for the non-extranet and one for each of the two extranets. This generalizes easily to the case where there are VPN-C receivers in n different extranets (i.e., receiving extranet flows whose sources are in n different VPNs).

Suppose again that there are three VPNs -- VPN-A, VPN-B, and VPN-C -- but in this example VPN-A is the only one with extranet sources while VPN-B and VPN-C both have receivers for the VPN-A extranet sources. This can be provisioned as either one extranet or two extranets.

To provision it as one extranet, the VPN-A VRFs are configured with one Outgoing Extranet RT (call it "RT-A-extranet"). The VPN-B and VPN-C VRFs with extranet receivers will be provisioned with RT-A-extranet as the Incoming Extranet RT. Thus, the VPN-B and VPN-C VRFs will each originate two Intra-AS I-PMSI A-D routes: one for the non-extranet and one for the extranet. From a given VRF, the Intra-AS I-PMSI A-D route for the extranet will carry RT-A-extranet but will not share any RT with the non-extranet A-D routes exported from the same VRF.

The result is that the VPN-B and VPN-C VRFs each belong to two MI-PMSIs: one for the extranet and one for the intranet. The MI-PMSI for the extranet attaches VPN-A VRFs, VPN-B VRFs, and VPN-C VRFs.

Alternatively, one could provision the VPN-A VRFs so that some UMH-eligible extranet source routes carry an RT that we will call "RT-A-to-B" and some carry an RT that we will call "RT-A-to-C". The VPN-A VRFs would be configured with both of these as Outgoing Extranet RTs. To allow an extranet flow from a VPN-A source to have both VPN-B and VPN-C receivers, the UMH-eligible route for that source would carry both RTs. VPN-B VRFs (but not VPN-C VRFs) would be provisioned with RT-A-to-B as an Incoming Extranet RT. VPN-C VRFs (but not VPN-B VRFs) would be provisioned with RT-A-to-C as an Incoming Extranet RT.

Following the rules above, if any VPN-A extranet source is to have both VPN-B and VPN-C receivers, the VPN-B and VPN-C VRFs will each originate two I-PMSI A-D routes: one for the extranet and one for the non-extranet. The single Intra-AS I-PMSI A-D route originated by the VPN-A VRFs will have both RT-A-to-B and RT-A-to-C among its RTs (as well as VPN-A's non-extranet RT). The extranet I-PMSI A-D route originated from a VPN-B VRF would have RT-A-to-B, and the extranet I-PMSI A-D route originated from a VPN-C VRF would have RT-A-to-C.

If a given VRF contains both extranet C-receivers and extranet C-sources, the procedures described above still work, as the VRF will be configured with both Incoming Extranet RTs and Outgoing Extranet RTs; the VRF functions as both a VPN-S VRF and a VPN-R VRF.

6.2.2. S-PMSIs

When PIM is used as the PE-PE C-multicast control plane, every S-PMSI is considered to be part of the "emulated LAN" that "corresponds" to a particular MI-PMSI.

When the bindings of C-flows to particular S-PMSIs are announced via S-PMSI Join messages (Section 7 of [RFC6513]) sent on the MI-PMSI, the S-PMSI is considered to be part of the same LAN interface as the corresponding MI-PMSI.

When the bindings of C-flows to particular S-PMSIs are announced via S-PMSI A-D routes, any S-PMSI A-D route exported from that VRF MUST have an RT in common with exactly one of the Intra-AS A-D routes exported from that VRF, and this MUST be one of the VRF's Outgoing Extranet RTs. Further, the S-PMSI A-D route MUST NOT have an RT in common with any other Intra-AS A-D route exported from a VRF on the same PE. A given S-PMSI A-D route will be considered to "correspond" to the MI-PMSI of the Intra-AS I-PMSI A-D route (originated from the same PE) with which it shares an RT.

The MI-PMSI that corresponds to a given S-PMSI is determined as follows:

- o If (1) there is an Intra-AS I-PMSI A-D route originated by the same PE that originated the S-PMSI A-D route, (2) those two routes have an RT in common, and (3) that RT is one of the VRF's Incoming Extranet RTs, then the S-PMSI corresponds to the I-PMSI associated with that Intra-AS I-PMSI A-D route.
- o Otherwise, if (1) there is an Inter-AS I-PMSI A-D route originated in the same AS as the S-PMSI A-D route, (2) those two routes have an RT in common, and (3) that RT is one of the VRF's Incoming Extranet RTs, then the S-PMSI corresponds to the I-PMSI associated with that Inter-AS I-PMSI A-D route.
- o Otherwise, there must be a configuration error (a violation of the requirements of Sections 3, 4, and 5 of this document).

When wildcard S-PMSIs are used, the rules given in [RFC6625] for determining whether a given S-PMSI A-D route is a "match for reception" to a given (C-S,C-G) or (C-*,C-G) are modified as follows:

A given S-PMSI A-D route MUST NOT be considered to be a "match for reception" for a given (C-S,C-G) or (C-*,C-G) state UNLESS that S-PMSI A-D route "corresponds" (as defined above) to the MI-PMSI that is the incoming interface for the given state.

The rules given in [RFC6625] for determining whether a given S-PMSI A-D route is a "match for transmission" are unchanged.

6.2.3. Sending PIM Control Packets

Suppose that a PE, say PE1, receives a PIM Join(S,G) from a CE, over a VRF interface that is associated with a VPN-R VRF. The PE does the RPF check for S by looking up S in the VPN-R VRF. The PIM C-instance associated with that VRF must determine the correct P-tunnel over which to send a PIM Join(S,G) to other PEs.

To do this, PE1 finds, in the VRF associated with the interface over which the Join was received, the selected UMH route for S, following the procedures of Section 5.1 of [RFC6513]. PE1 determines the set of RTs carried by that route. PE1 then checks to see if there is an Intra-AS I-PMSI A-D route, currently originated by PE1, that has an RT in common with the selected UMH route for S.

If the rules of Sections 3, 4, and 5 have been followed, each of PE1's selected UMH routes will share an RT with a single one of PE1's currently originated Intra-AS I-PMSI A-D routes. If this is so, the Join is sent on the P-tunnel advertised in the PTA of that route. Otherwise, the Join MUST NOT be sent.

In essence, this procedure makes the RPF check for C-S resolve to the MI-PMSI that is serving as the next-hop "interface" to C-S.

If a PE receives a PIM Join(*,G) from a CE, the procedure for doing the RPF check is the same, except that the selected UMH route will be a route to the C-RP associated with the C-G group.

6.2.4. Receiving PIM Control Packets

When a PIM C-instance receives a PIM control message from a P-tunnel, it needs to identify the message's incoming interface. This incoming interface is the MI-PMSI of which the P-tunnel is a part.

6.2.5. Sending and Receiving Data Packets

The rules for choosing the PMSI on which to send a multicast data packet are as specified in [RFC6513] and [RFC6625], with one new restriction: a VPN-S VRF always transmits a multicast data packet either on the VPN-S MI-PMSI or on an S-PMSI that corresponds to the VPN-S MI-PMSI. From the perspective of the PIM C-instance, there is only one outgoing interface.

When a PIM C-instance receives a multicast data packet from a given P-tunnel and that P-tunnel is being used to instantiate an MI-PMSI, the MI-PMSI of which the P-tunnel is a part (see Sections 6.2.1 and 6.2.2) is considered to be the packet's incoming interface. If the packet is received on a P-tunnel that was advertised in an S-PMSI A-D route, the packet's incoming interface is the MI-PMSI to which that S-PMSI route corresponds, as defined in Section 6.2.2. Ordinary PIM rules for data-plane RPF checks apply.

Following ordinary PIM procedures, packets arriving from an unexpected incoming interface are discarded. This eliminates any problems due to the ambiguities described in Sections 2.1 and 2.2.

6.3. "Multiple PMSIs per C-Flow" Model

In this model, if a VPN-S VRF has extranet multicast C-sources and a VPN-R VRF has extranet multicast C-receivers allowed by policy to receive from the C-sources in the VPN-S VRF, then the VPN-S VRF joins the MI-PMSI that VPN-R uses for its non-extranet traffic.

In the "single PMSI per C-flow" transmission model (as described in Section 6.2), a PE that needs to transmit a multicast data packet to a set of other PEs transmits the packet on a single PMSI. This means that if a packet needs to be transmitted from a VPN-A VRF and received at a VPN-B VRF and a VPN-C VRF, there must be some P-tunnel from which the VPN-B and VPN-C VRFs can both receive packets.

In the "multiple PMSIs per C-flow" transmission model, a PE that needs to transmit a multicast data packet to a set of other PEs may transmit the packet on several different PMSIs. (Of course, any given packet is transmitted only once on a given P-tunnel.) For example, if a C-flow (C-S,C-G) has a VPN-A C-source, a VPN-B receiver, and a VPN-C receiver, there could be one PMSI that the VPN-A VRF uses to transmit the packet to the VPN-B VRFs and another PMSI that the VPN-A VRF uses to transmit the packet to the VPN-C VRFs.

6.3.1. Forming the MI-PMSIs

Consider a VPN-R VRF that has extranet C-receivers. Per [RFC6513], each VPN-R VRF must originate an Intra-AS I-PMSI A-D route containing a PTA specifying the P-tunnel to be used as part of the VPN-R MI-PMSI. In the absence of extranet service, this route carries the VRF's non-extranet RT, RT-R. When extranet service is provided (using the "single PMSI per C-flow" model), this route MUST also carry each of the VRF's Incoming Extranet RTs.

Consider a VPN-S VRF that has extranet C-sources. Per [RFC6513], each VPN-S VRF must originate an Intra-AS I-PMSI A-D route containing a PTA specifying the P-tunnel to be used as part of the VPN-S MI-PMSI. This route carries the VRF's non-extranet RT, RT-S. When extranet service is provided using the "multiple PMSIs per C-flow" model, the VPN-S VRF MUST also originate one or more additional Intra-AS I-PMSI A-D routes. It MUST originate one additional Intra-AS I-PMSI A-D route for each Outgoing Extranet RT with which it has been configured; each such route will have a distinct RD and will carry exactly one of the configured Outgoing Extranet RTs.

As with the "single PMSI per C-flow" transmission model, VRFs containing extranet C-receivers need to import UMH-eligible extranet C-source routes from VRFs containing C-sources. This is ensured by the rules of Sections 3, 4, and 5.

However, in the "multiple PMSIs per C-flow" model, a VRF containing only C-receivers originates only a single Intra-AS I-PMSI A-D route carrying the non-extranet RT and all the Incoming Extranet RTs.

When a VRF containing C-receivers imports Intra-AS I-PMSI A-D routes that carry the non-extranet RT or one of the Incoming Extranet RTs, the P-tunnels specified in the PTA of all such routes are considered to be part of the same MI-PMSI. That is, the associated PIM C-instance will treat them as part of a single interface.

In this model, it is the VRF containing extranet C-sources that MUST originate multiple Intra-AS I-PMSI A-D routes. Each such route MUST have a distinct RD, and the set of RTs carried by any one of these routes MUST be disjoint from the set carried by any other. There MUST be one such route for each of the VRF's Outgoing Extranet RTs, and each such route MUST carry exactly one of the VRF's Outgoing Extranet RTs. The VRFs containing extranet C-sources MUST also import all the A-D routes originated by the VRFs containing extranet C-receivers. If a set of originated and/or imported Intra-AS I-PMSI A-D routes have an RT in common and that RT is one of the VRF's Outgoing Export RTs, then those routes are considered to be "about" the same MI-PMSI. The PIM C-instance of the VRF treats each MI-PMSI as a LAN interface.

In effect, if VPN-S has only extranet C-sources and VPN-R has only extranet C-receivers, this model has the VPN-S VRFs join the VPN-R MI-PMSI. The VPN-S VRFs will thus be attached to multiple MI-PMSIs, while the VPN-R VRFs are attached to only one. The fact that the VPN-R MI-PMSI is attached to VPN-S VRFs is not known to the PIM C-instance at the VPN-R VRFs.

If a VPN-A VRF has extranet C-sources allowed to send to C-receivers in a VPN-B VRF and the VPN-B VRF has C-sources allowed to send to C-receivers in the VPN-A VRF, the above procedures still work as specified.

Following normal PIM procedures, when the PIM C-instance at a VRF with extranet C-sources receives a Join(C-S,C-G) or a Join(C-*,C-G) over an MI-PMSI, it may create (C-S,C-G) or (C-*,C-G) state, and the MI-PMSI over which the Join was received may be added to the set of outgoing interfaces for that multicast state. If n MI-PMSIs are added to the outgoing interface list for a particular multicast state, a multicast data packet may need to be replicated n times and transmitted once on each of the n MI-PMSIs.

Since all of the multicast data packets received from another PE are received over a single emulated LAN, it is not necessary to have any special procedures to determine a packet's incoming interface. The ambiguities described in Sections 2.1 and 2.2 do not occur, because a VPN-R VRF can only receive multicast data traffic that has been requested by a VPN-R VRF.

7. When BGP Is the PE-PE C-Multicast Control Plane

This document assumes that if BGP is used as the PE-PE C-multicast control plane, the "single PMSI per C-flow" model is used. Procedures for providing the "multiple PMSIs per C-flow" model with BGP C-multicast are outside the scope of this document.

When BGP is used as the C-multicast control plane, the "single PMSI per C-flow" model may be used either with or without extranet separation. (Recall that "extranet separation" means that no P-tunnel can carry traffic from both extranet sources and non-extranet sources.) In either case, the data traffic may be carried on Inclusive Tunnels only, Selective Tunnels only (known as the "S-PMSI only" model), or a combination of Inclusive Tunnels and Selective Tunnels. This is determined by provisioning. The procedures specified below support all three choices.

Note that there are no special extranet procedures for Inter-AS I-PMSI A-D routes or for Leaf A-D routes.

7.1. Originating C-Multicast Routes

This section applies whether extranet separation is used or not.

When it is necessary to originate a C-multicast Source Tree Join for (C-S,C-G), a PE must follow the procedures of Section 11.1.3 ("Constructing the Rest of the C-Multicast Route") of [RFC6514] to find the selected UMH route for C-S. When it is necessary to originate a C-multicast Shared Tree Join for (C-*,C-G), where C-G is an ASM group, a PE must follow the procedures of that section to find the selected UMH route for C-G's C-RP.

Section 11.1.3 of [RFC6514] specifies how information from the selected UMH route is used to find an Intra-AS I-PMSI A-D route or an Inter-AS I-PMSI A-D route. Information from that I-PMSI A-D route is then used to construct part of the C-multicast route.

For extranets, the rules given in Section 7.4.5 of this document are used to find the Inter-AS I-PMSI A-D route or an Intra-AS I-PMSI A-D route that "corresponds" to the selected UMH route; the rules in Section 7.4.5 replace the rules given in Section 11.1.3 of [RFC6514] for finding the Inter-AS or Intra-AS I-PMSI A-D route.

Information from this I-PMSI A-D route is then used, as specified in Section 11.1.3 of [RFC6514], to construct the C-multicast route.

7.2. Originating A-D Routes without Extranet Separation

7.2.1. Intra-AS I-PMSI A-D Routes

Consider a VRF (call it "VRF-S") that contains extranet C-sources and exports UMH-eligible routes matching those C-sources. The VRF may also originate and export an Intra-AS I-PMSI A-D route.

As specified in [RFC6514], if exactly one Intra-AS I-PMSI A-D route is originated by and exported from VRF-S, the RTs carried by that route MUST be chosen such that every VRF that imports a UMH-eligible route from VRF-S also imports this Intra-AS I-PMSI A-D route.

If inclusive P-tunnels are being used to carry extranet C-flows, there are additional requirements on the way the RTs carried by the Intra-AS I-PMSI A-D routes must be chosen, as specified in the following paragraph.

If VRF-S is using inclusive P-tunnels but is not using extranet separation, there is one inclusive P-tunnel rooted at VRF-S, and this tunnel carries both extranet and non-extranet C-flows. This Inclusive Tunnel is identified in the PTA of the Intra-AS I-PMSI A-D route originated from VRF-S. The set of RTs carried by this Intra-AS I-PMSI A-D route MUST be chosen so as to ensure that every VRF that imports a UMH-eligible route from this VRF-S also imports this Intra-AS I-PMSI A-D route. Further, the set of RTs carried by this Intra-AS I-PMSI A-D route MUST be chosen such that it has at least one RT in common with every UMH-eligible route that is exported from the VRF.

7.2.2. S-PMSI A-D Routes

Let R-SP be an S-PMSI A-D route that is exported from VRF-S. Suppose that R-SP is used to bind some or all of the extranet C-flows from a given extranet C-source to a given selective P-tunnel. Let R-UMH be a UMH-eligible route that is exported from VRF-S and matches the given extranet C-source. In that case, R-SP and R-UMH MUST have at least one RT in common. Further, the RTs carried by these two routes MUST be such that every VRF that imports R-UMH also imports R-SP. These rules apply whether or not R-SP uses wildcards [RFC6625].

An implementation MUST allow the set of RTs carried by the S-PMSI A-D routes to be specified by configuration. In the absence of such configuration, an S-PMSI A-D route originated by a given VRF, say VRF-X, MUST carry a default set of RTs, as specified by the following rules:

1. By default, an S-PMSI A-D route originated by VRF-X for a given (C-S,C-G) or (C-S,C-*) carries the same RT(s) as the UMH-eligible route originated by VRF-X that matches C-S.
2. By default, an S-PMSI A-D route originated by VRF-X for a given (C-*,C-G) carries as its RTs a set union of all RT(s) of the UMH-eligible route(s) matching the multicast C-sources contained in VRF-X that could originate traffic for that C-G. Moreover, if the VRF contains (as defined in Section 1.1) the C-RP of C-G, then this set union also includes the RT(s) of the UMH-eligible route matching C-RP and the RT(s) of the unicast VPN-IP route matching C-RP.
3. By default, if a (C-*,C-*) S-PMSI A-D route originated by VRF-X is to be used for both extranet and non-extranet traffic, it carries the same RTs that would be carried (as specified in Section 7.2.1) by an I-PMSI A-D route originated by VRF-X if that I-PMSI A-D route were advertising an inclusive P-tunnel for carrying both extranet and non-extranet traffic. In general, a given VRF would not originate both (a) an S-PMSI A-D route advertising a (C-*,C-*) selective P-tunnel for both extranet and non-extranet traffic and (b) an I-PMSI A-D route advertising an inclusive P-tunnel for both extranet and non-extranet traffic, as the inclusive P-tunnel would not get used in that case.

7.2.3. Source Active A-D Routes

7.2.3.1. When Inter-Site Shared Trees Are Used

This section applies when inter-site shared trees are used, as specified in Section 13 of [RFC6514].

If VRF-S exports a Source Active A-D route that contains C-S in the Multicast Source field of its NLRI and VRF-S also exports a UMH-eligible route matching C-S, the Source Active A-D route MUST carry at least one RT in common with the UMH-eligible route. The RT MUST be chosen such that the following condition holds: if a VRF, say VRF-R, contains an extranet C-receiver allowed by policy to receive extranet traffic from C-S, then VRF-R imports both the UMH-eligible route and the Source Active A-D route.

By default, a Source Active A-D route for a given (C-S,C-G), exported by a given VRF, carries the same set of RTs as the UMH-eligible route matching C-S that is exported from that VRF.

7.2.3.2. When Inter-Site Shared Trees Are Not Used

This section applies when inter-site shared trees are not used, as specified in Section 14 of [RFC6514].

Suppose that a VRF, say VRF-X, contains the C-RP for a given extranet C-group, say C-G. If C-S is an active source for C-G, then, following the procedures of Section 14.1 of [RFC6514], VRF-X may export a Source Active A-D route that contains C-S in the Multicast Source field of its NLRI. With the following text, this document replaces the rule specified in Section 14.1 of [RFC6514] for constructing the RT(s) carried by such a route: VRF-X MUST be configured such that the Source Active A-D route for (C-S,C-G) carries the same set of RTs as the UMH-eligible route matching C-S that is exported from the VRF(s) containing C-S. This way, if a VRF, say VRF-R, contains an extranet C-receiver allowed by policy to receive extranet traffic from C-S, then VRF-R imports both the UMH-eligible route and the Source Active A-D route.

7.3. Originating A-D Routes with Extranet Separation

If a VRF contains both extranet C-sources and non-extranet C-sources, it MUST be configured with both a default RD and an extranet RD (see Section 1.3). The use of these RDs is explained in the following subsections.

7.3.1. Intra-AS I-PMSI A-D Routes

This section applies when VRF-S is using extranet separation AND when VRF-S is using an inclusive P-tunnel to carry some or all of the extranet C-flows that it needs to transmit to other VRFs.

If VRF-S contains both extranet C-sources and non-extranet C-sources, and inclusive P-tunnels are used to carry both extranet C-flows and non-extranet C-flows, then there MUST be two Inclusive Tunnels from VRF-S, one of which is to be used only to carry extranet C-flows (the "extranet inclusive P-tunnel") and one of which is to be used only to carry non-extranet C-flows (the "non-extranet inclusive P-tunnel").

In this case, the VRF MUST originate two Intra-AS I-PMSI A-D routes. Their respective NLRIs MUST of course have different RDs. One of the Intra-AS I-PMSI A-D routes identifies the extranet inclusive P-tunnel in its PTA. This route MUST have the VRF's extranet RD in its NLRI. The other route identifies the non-extranet inclusive P-tunnel in its PTA. This route MUST have the VRF's default RD in its PTA.

If VRF-S uses an inclusive P-tunnel for carrying extranet traffic but does not use an inclusive P-tunnel for carrying non-extranet traffic, then of course only a single Intra-AS I-PMSI A-D route need be originated. The PTA of this route identifies the "extranet inclusive P-tunnel". The NLRI of that route MUST contain the VRF's extranet RD.

An Intra-AS I-PMSI A-D route whose PTA identifies an extranet inclusive P-tunnel MUST carry the Extranet Separation Extended Community defined in Section 4.5.

The RTs carried by an Intra-AS I-PMSI A-D route whose PTA identifies the "extranet inclusive P-tunnel" MUST be chosen such that the following condition holds: if a VRF (call it "VRF-R") imports a UMH-eligible route from VRF-S and that route matches an extranet C-source, then VRF-R also imports that Intra-AS I-PMSI A-D route.

Note that when extranet separation is used, it is possible to use an inclusive P-tunnel for non-extranet traffic while using only selective P-tunnels for extranet traffic. It is also possible to use an inclusive P-tunnel for extranet traffic while using only selective P-tunnels for non-extranet traffic.

7.3.2. S-PMSI A-D Routes

Let R-SP be an S-PMSI A-D route that is exported from VRF-S. Suppose that R-SP is used to bind some or all of the extranet C-flows from a given extranet C-source to a given selective P-tunnel. Let R-UMH be a UMH-eligible route that is exported from VRF-S and matches the given extranet C-source. In that case, R-SP and R-UMH MUST have at least one RT in common. Further, the RTs carried by these two routes MUST be such that every VRF that imports R-UMH also imports R-SP. These rules apply whether or not R-SP uses wildcards [RFC6625].

The following rules, specific to the use of extranet separation, apply:

- o A selective P-tunnel MUST NOT carry C-flows from both extranet and non-extranet C-sources.
- o If it is desired to use a (C-*,C-*) S-PMSI to carry extranet traffic and also use a (C-*,C-*) S-PMSI to carry non-extranet traffic, then two (C-*,C-*) S-PMSI A-D routes MUST be originated. These two routes MUST have different RDs in their respective NLRI fields, and their respective PTAs MUST identify different P-tunnels. If the route advertises a P-tunnel that carries only non-extranet traffic, the route's NLRI MUST contain the VRF's default RD. If the route advertises a P-tunnel that carries only extranet traffic, the route's NLRI MUST contain the VRF's extranet RD.
- o In the following cases, an S-PMSI A-D route exported from the VRF MUST have the VRF's extranet RD in its NLRI:
 - * The S-PMSI A-D route is a (C-S,C-G) or a (C-S,C-*) S-PMSI A-D route, and C-S is an extranet C-source.
 - * The S-PMSI A-D route is a (C-*,C-G) S-PMSI A-D route, and C-G is an extranet C-group.

In all other cases, a (C-S,C-G), (C-S,C-*), or (C-*,C-G) S-PMSI A-D route MUST have the VRF's default RD in its NLRI.

- o A (C-*,C-*) S-PMSI A-D route advertising a P-tunnel that is used to carry extranet traffic MUST carry the Extranet Separation Extended Community defined in Section 4.5.

An implementation MUST allow the set of RTs carried by the S-PMSI A-D routes to be specified by configuration. In the absence of such configuration, an S-PMSI A-D route originated by a given VRF, say VRF-X, MUST carry a default set of RTs, as specified by the following rules:

1. Rule 1 of Section 7.2.2 applies.
2. By default, if C-G is an extranet C-group, rule 2 of Section 7.2.2 applies.
3. By default, if a (C-*,C-*) S-PMSI A-D route originated by VRF-X is to be used for extranet traffic, it carries the same RTs that would be carried (as specified in Section 7.3.1) by an I-PMSI A-D route originated by VRF-X if that I-PMSI A-D route were

advertising an inclusive P-tunnel for carrying extranet traffic. In general, a given VRF would not originate both an S-PMSI A-D route advertising a (C-*,C-*) selective P-tunnel for extranet traffic and an I-PMSI A-D route advertising an inclusive P-tunnel for extranet traffic, as the inclusive P-tunnel would not get used in that case.

7.3.3. Source Active A-D Routes

The procedures of Section 7.2.3 apply.

However, if a Source Active A-D route is exported from a given VRF and the route contains C-S, where C-S is an extranet C-source, then the RD of the route's NLRI MUST be the extranet RD of the VRF. Otherwise, the RD is the default RD of the VRF.

7.4. Determining the Expected P-Tunnel for a C-Flow

This section applies whether extranet separation is used or not.

In the context of a VRF with receivers for a particular C-flow, a PE must determine the P-tunnel over which packets of that C-flow are expected to arrive. This is done by finding an I-PMSI or S-PMSI A-D route that "matches" the flow. The matching A-D route will contain a PTA that specifies the P-tunnel being used to carry the traffic of that C-flow. We will refer to this P-tunnel as the "expected P-tunnel" for the C-flow. (Note that, per [MVPN-IR], if the PTA specifies a tunnel of type "Ingress Replication" (IR), the identifier of the P-tunnel is actually the NLRI of the I-PMSI or S-PMSI A-D route. If the PTA specifies a tunnel type other than IR, the identifier of the P-tunnel is found in the Tunnel Identifier field of the PTA.)

A PE that needs to receive a given (C-S,C-G) or (C-*,C-G) C-flow MUST join the expected P-tunnel for that C-flow, and the PE MUST remain joined to the P-tunnel as long as (1) the PE continues to need to receive the given C-flow and (2) the P-tunnel continues to remain the expected P-tunnel for that C-flow. Procedures for joining and leaving a tunnel depend, of course, on the tunnel type.

If a PTA specifies a non-zero MPLS label for a tunnel that is not an IR tunnel, then the PE originating the A-D route containing that PTA is advertising an aggregate P-tunnel. The aggregate P-tunnel can be thought of as an outer P-tunnel multiplexing some number of inner P-tunnels. The inner P-tunnels are demultiplexed by means of the MPLS label in the PTA. In this document, when we talk of the "expected P-tunnel" in the context of an aggregate P-tunnel, we refer

to a particular inner P-tunnel, not to the outer P-tunnel. It is this "inner P-tunnel" that is the expected P-tunnel for a given C-flow.

In order to find the expected P-tunnel for a given C-flow, the upstream PE of the C-flow is first determined. Then, the S-PMSI A-D routes originated by that PE are examined, and their NLRIs compared to the (C-S/C-RP,C-G) of the flow, to see if there is a "match for reception". (If there is no S-PMSI A-D route that matches a given C-flow, the expected P-tunnel for that C-flow may have been advertised in an I-PMSI A-D route; see Section 7.4.5.)

The rules for determining, in non-extranet cases, whether a given C-flow is a "match for reception" for a given S-PMSI A-D route are given in Section 3.2 of [RFC6625]. Note that we use the terms "installed" and "originated" as they are defined in Section 3.2 of [RFC6625]. (See also Section 1.1 of this document.)

This specification provides additional rules for determining whether a given S-PMSI A-D route is a "match for reception" for a given (C-S/C-RP,C-G). Note that these rules all assume the context of a particular VRF into which the A-D route has been imported.

The rules given in [RFC6625] for determining whether a given S-PMSI A-D route is a "match for transmission" remain unchanged.

Suppose that a PE has originated a C-multicast Shared Tree Join for (C-*,C-G) but has not originated a C-multicast Source Tree Join for (C-S,C-G). Suppose also that the PE has received and installed a Source Active A-D route for (C-S,C-G). As described in Section 13.2 of [RFC6514], the PE must receive the (C-S,C-G) traffic from the tunnel the originator of the installed Source Active A-D route uses for sending (C-S,C-G).

The originator of the installed Source Active A-D route is determined as follows:

1. Look at the "UMH Route Candidate Set" for C-S, as defined in Section 5.1.3 of [RFC6513].
2. From that set, select a subset of UMH routes to C-S, such that each route in the subset has at least one RT in common with the Source Active A-D route and at least one of the RTs in common is an import RT of the VRF.
3. From that subset, find the route whose RD is the same as the RD from the NLRI of the Source Active A-D route.

4. The upstream PE is the PE identified in the VRF Route Import Extended Community of that route.
5. The upstream AS is the AS identified in the Source AS Extended Community of that route.

If step 2 results in an empty set or step 3 fails to find a route, then the upstream PE of the Source Active A-D route cannot be determined, and it is necessary to act as if the Source Active A-D route had not been installed. (A subsequent change to the UMH Route Candidate Set for C-S may require that a new attempt be made to determine the upstream PE.)

Once the upstream PE is determined, the P-tunnel over which the flow is expected is determined according to the procedures already described in this section.

7.4.1. (C-S,C-G) S-PMSI A-D Routes

When extranet functionality is being provided, an S-PMSI A-D route whose NLRI contains (C-S,C-G) is NOT considered to be a "match for reception" for a given C-flow (C-S,C-G) unless one of the following conditions holds (in addition to the conditions specified in [RFC6625]):

- o the "single C-source per (C-S,C-G) or (C-S,C-*) P-tunnel" is provisioned, or
- o the selected UMH route for C-S has at least one RT in common with the S-PMSI A-D route, and at least one of the common RTs is an import RT of the VRF.

7.4.2. (C-S,C-*) S-PMSI A-D Routes

When extranet functionality is being provided, an S-PMSI A-D route whose NLRI contains (C-S,C-*) is NOT considered to be a "match for reception" for a given C-flow (C-S,C-G) unless one of the following conditions holds (in addition to the conditions specified in [RFC6625]):

- o the "single C-source per (C-S,C-G) or (C-S,C-*) P-tunnel" is provisioned, or
- o the selected UMH route for C-S has at least one RT in common with the S-PMSI A-D route, and at least one of the common RTs is an import RT of the VRF.

7.4.3. (C-*,C-G) S-PMSI A-D Routes

When extranet functionality is being provided, an S-PMSI A-D route whose NLRI contains (C-*,C-G) is NOT considered to be a "match for reception" for a given C-flow (C-S,C-G) in a given VRF unless either condition 1 or condition 2 below holds (in addition to the conditions specified in [RFC6625]):

1. The given VRF has currently originated a C-multicast Shared Tree Join route for (C-*,C-G), and
 - a. (C-*,C-G) matches an installed (C-*,C-G) S-PMSI A-D route (according to [RFC6625]) in the given VRF, and
 - b. either
 - i. the "single C-group per (C-*,C-G) P-tunnel" policy has been provisioned, or
 - ii. the RTs of that S-PMSI A-D route form a non-empty intersection with the RTs carried in the VRF's selected UMH route for C-RP of that C-G, or
 - iii. installed in the VRF is at least one (C-S,C-G) Source Active A-D route that was originated by the same PE as the (C-*,C-G) S-PMSI A-D route.
2. The given VRF does not have a currently originated C-multicast Shared Tree Join for (C-*,C-G), but
 - a. there are one or more values for C-S for which the VRF has a currently originated Source Tree Join C-multicast route for (C-S,C-G), and
 - b. the (C-* C-G) S-PMSI A-D route matches (according to [RFC6625]) each such (C-S,C-G), and
 - c. either
 - i. the "single C-group per (C-*,C-G) P-tunnel" policy has been provisioned, or
 - ii. the RTs of that S-PMSI A-D route form a non-empty intersection with the RTs carried in the VRF's selected UMH routes for each such C-S

If a VRF has an installed (C-*,C-G) S-PMSI A-D route but does not have a (C-S,C-G) or (C-*,C-G) multicast state that matches that route for reception, the procedures of Section 12.3 ("Receiving S-PMSI A-D Routes by PEs") of [RFC6514] are not invoked for that route. If those multicast states are created at some later time when the route is still installed, the procedures of Section 12.3 of [RFC6514] are invoked at that time.

7.4.4. (C-*,C-*) S-PMSI A-D Routes

A (C-*,C-*) S-PMSI A-D route (call it "R-AD") is NOT considered to be a "match for reception" for a given C-flow (C-S,C-G) or (C-*,C-G) unless the following conditions hold (in addition to the conditions specified in [RFC6625]):

- o The selected UMH route (call it "R-UMH") for C-S or for C-G's C-RP, respectively, has at least one RT in common with R-AD, and at least one of the common RTs is an import RT of the VRF.
- o Either R-AD and R-UMH both carry the Extranet Separation Extended Community or neither carries the Extranet Separation Extended Community.

7.4.5. I-PMSI A-D Routes

If a particular egress VRF in a particular egress PE contains no matching S-PMSI A-D routes for a particular C-flow, then the C-flow is expected to arrive (at that egress VRF) on an inclusive P-tunnel.

Suppose that an egress PE has originated a (C-S,C-G) C-multicast Source Tree Join. Let R-UMH be the selected UMH route (in the given egress VRF) for C-S. As specified in [RFC6514], the selected upstream PE for (C-S,C-G) is determined from the VRF Route Import Extended Community of R-UMH, and the "selected upstream AS" for the flow is determined from the Source AS Extended Community of R-UMH.

Suppose that an egress PE has originated a (C-*,C-G) C-multicast Shared Tree Join but has not originated a (C-S,C-G) C-multicast Source Tree Join. If the egress VRF does not have a (C-S,C-G) Source Active A-D route installed, the selected upstream PE is determined from the VRF Route Import Extended Community of the installed UMH-eligible route matching C-RP, where C-RP is the RP for the group C-G. The selected upstream AS for the flow is determined from the Source AS Extended Community of that route. If the egress VRF does have a (C-S,C-G) Source Active A-D route installed, the selected upstream PE and upstream AS are determined as specified in Section 7.4. In either case, let R-UMH be the installed UMH-eligible route matching C-S.

The inclusive P-tunnel that is expected to be carrying a particular C-flow is found as follows:

- o If the selected upstream AS is the local AS or segmented Inter-AS P-tunnels are not being used to instantiate I-PMSIs, then look in the VRF for an installed Intra-AS I-PMSI A-D route, R-AD, such that (a) R-AD is originated by the selected upstream PE, (b) R-AD has at least one RT in common with R-UMH, (c) at least one of the common RTs is an import RT of the local VRF, and (d) either R-AD and R-UMH both carry the Extranet Separation Extended Community or neither carries the Extranet Separation Extended Community.

The PTA of R-AD specifies the P-tunnel over which the traffic of the given C-flow is expected.

- o If the selected upstream AS is not the local AS and segmented Inter-AS P-tunnels are being used to instantiate I-PMSIs, then look in the VRF for an installed Inter-AS I-PMSI A-D route, R-AD, such that (a) the Source AS field of R-AD's NLRI contains the AS number of the selected upstream AS, (b) R-AD has at least one RT in common with R-UMH, (c) at least one of the common RTs is an import RT of the local VRF, and (d) either R-AD and R-UMH both carry the Extranet Separation Extended Community or neither carries the Extranet Separation Extended Community.

The PTA of R-AD specifies the P-tunnel over which the traffic of the given C-flow is expected.

7.5. Packets Arriving from the Wrong P-Tunnel

Any packets that arrive on a P-tunnel other than the expected P-tunnel (as defined in Section 7.4) MUST be discarded unless it is known that all the packets carried by both P-tunnels are from the same ingress VRF. (See Section 2.3.1 for a more detailed discussion of when to discard packets from other than the expected P-tunnel.) Note that packets arriving on the wrong P-tunnel are to be discarded even if they are arriving from the expected PE.

8. Multiple Extranet VRFs on the Same PE

When multiple VRFs that contain extranet receivers for a given extranet source are present on the same PE, this PE becomes a single leaf of the P-tunnel used for sending (multicast) traffic from that source to these extranet receivers. The PE MUST be able to replicate this traffic to the multiple VRFs. Specific procedures for doing so are local to the PE and are outside the scope of this document.

Two or more VRFs on the same PE may import the same S-PMSI A-D route. If this S-PMSI A-D route contains a PTA that has its "Leaf Information Required" flag set, it may be necessary for the PE to originate a Leaf A-D route whose NLRI is computed from the NLRI of the S-PMSI A-D route. (Details are provided in [RFC6514].) Note that for a given S-PMSI A-D route, the PE can originate only one corresponding Leaf A-D route, even if the S-PMSI A-D route is imported into multiple VRFs. This Leaf A-D route can thus be thought of as originating from several VRFs. It MUST NOT be withdrawn by the PE until there are no longer any VRFs originating it.

[RFC6514] specifies conditions under which a PE originates a C-multicast Source Tree Join or a C-multicast Shared Tree Join, based on the (*,G) and (S,G) states associated with a given VRF. It also specifies the procedure for computing the NLRI of each such route. While a given PE may contain two or more VRFs that have (extranet) receivers for the same extranet C-flow, the PE cannot originate more than one BGP route with a given NLRI. If there are multiple VRFs, each of which has state that is sufficient to cause a given C-multicast route to be originated, the route can be thought of as originating from several VRFs. It MUST NOT be withdrawn by the PE until there is no longer any VRF with multicast state sufficient to cause the route to be originated.

For a given extranet, the site(s) that contains the extranet source(s) and the site(s) that contains the extranet receiver(s) may be connected to the same PE. In this scenario, the procedures by which (multicast) traffic from these sources is delivered to these receivers are local to the PE and are outside the scope of this document.

An implementation MUST support multiple extranet VRFs on a PE.

9. IANA Considerations

IANA has allocated two new codepoints from the "First Come First Served" [RFC5226] range of the "Transitive Opaque Extended Community Sub-Types" registry (under the top-level registry "Border Gateway Protocol (BGP) Extended Communities" registry).

- o Extranet Source Extended Community (0x04)
- o Extranet Separation Extended Community (0x05)

10. Security Considerations

The security considerations of [RFC6513] and [RFC6514] are applicable.

As is the case with any application of technology based upon [RFC4364], misconfiguration of the RTs may result in VPN security violations (i.e., may result in a packet being delivered to a VPN where, according to policy, it is not supposed to go).

In those cases where the set of extranet sources of a particular VRF are manually configured, improper configuration of the VRF can result in VPN security violations -- traffic from a host that is not an extranet source may be treated as if it were traffic from an extranet source.

Section 4.4 specifies the optional use of a new Extended Community -- the Extranet Source Extended Community. Security considerations regarding the use and distribution of that Extended Community are discussed in that section.

The procedures of this document do not provide encryption of the data flows that are sent across the SP backbone network. Hence, these procedures do not by themselves ensure the privacy or integrity of the data against attacks on the backbone network.

In general, different VPNs are allowed to have overlapping IP address spaces; i.e., a host in one VPN may have the same IP address as a host in another. This is safe because the customer routes from a given VPN do not pass into other VPNs. Even if there are overlapping address spaces among VPNs, the routes that are known at any given VPN site are unambiguous, as long as the address space of that VPN is unambiguous. However, this is not necessarily true when extranet service is provided. If an extranet C-receiver in VPN-R is to be able to receive multicast traffic from an extranet C-source in VPN-S, then the address of the VPN-S extranet C-source must be imported into one or more VPN-R VRFs. If that address is also the address of a VPN-R non-extranet C-source, then a system attempting to receive an extranet C-flow from the VPN-R extranet C-source may instead receive a non-extranet C-flow from the VPN-S C-source. Otherwise, a VPN security violation may result.

That is, when provisioning an extranet between two VPNs that have overlapping address spaces, one must ensure that the IP addresses of the extranet sources and the extranet receivers are not from the overlapping part of the address space. This document specifies that if a route is imported into a given VRF, all addresses that match that route must be unambiguous in the context of that VRF. Improper

provisioning of the extranet source addresses or improper provisioning of the RTs may cause this rule to be violated and may result in a VPN security violation.

It is possible that a given multicast C-source is the source of multiple flows, some of which are intended to be extranet C-flows and some of which are intended to be non-extranet flows. However, the procedures of this document will allow any C-receiver that is able to receive the extranet C-flows from a given C-source to also receive the non-extranet C-flows from that source. As a result, VPN security violations may result if any system is a C-source for both extranet and non-extranet C-flows. However, the set of C-flows transmitted by a given C-source is not under the control of the SP. SPs who offer the extranet MVPN service must make sure that this potential for VPN security violations is clearly understood by the customers who administer the C-sources.

This specification does not require that UMH-eligible routes be "host routes"; they may be less specific routes. So, it is possible for the NLRI of a UMH-eligible route to contain an address prefix that matches the address of both an extranet C-source and a non-extranet C-source. If such a route is exported from a VPN-S VRF and imported by a VPN-R VRF, C-receivers contained in VPN-R will be able to receive C-flows from the non-extranet C-sources whose addresses match that route. This may result in VPN security violations. Service providers who offer the extranet MVPN service must make sure that this is clearly understood by the customers who administer the distribution of routes from CE routers to PE routers.

If the address ambiguities described in Sections 2.1 and 2.2 are not prohibited by deployment of the policies described in Section 2.3.2, VRFs must be able to discard traffic that arrives on the wrong P-tunnel (as specified in Sections 2.3.1 and 7.5). Otherwise, VPN security violations may occur.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<http://www.rfc-editor.org/info/rfc4364>>.
- [RFC6513] Rosen, E., Ed., and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<http://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<http://www.rfc-editor.org/info/rfc6514>>.
- [RFC6625] Rosen, E., Ed., Rekhter, Y., Ed., Hendrickx, W., and R. Qiu, "Wildcardcards in Multicast VPN Auto-Discovery Routes", RFC 6625, DOI 10.17487/RFC6625, May 2012, <<http://www.rfc-editor.org/info/rfc6625>>.
- [RFC7153] Rosen, E. and Y. Rekhter, "IANA Registries for BGP Extended Communities", RFC 7153, DOI 10.17487/RFC7153, March 2014, <<http://www.rfc-editor.org/info/rfc7153>>.
- [RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, RFC 7761, DOI 10.17487/RFC7761, March 2016, <<http://www.rfc-editor.org/info/rfc7761>>.

11.2. Informative References

- [MVPN-IR] Rosen, E., Ed., Subramanian, K., and Z. Zhang, "Ingress Replication Tunnels in Multicast VPN", Work in Progress, draft-ietf-bess-ir-03, April 2016.
- [RFC3446] Kim, D., Meyer, D., Kilmer, H., and D. Farinacci, "Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)", RFC 3446, DOI 10.17487/RFC3446, January 2003, <<http://www.rfc-editor.org/info/rfc3446>>.
- [RFC3618] Fenner, B., Ed., and D. Meyer, Ed., "Multicast Source Discovery Protocol (MSDP)", RFC 3618, DOI 10.17487/RFC3618, October 2003, <<http://www.rfc-editor.org/info/rfc3618>>.
- [RFC4610] Farinacci, D. and Y. Cai, "Anycast-RP Using Protocol Independent Multicast (PIM)", RFC 4610, DOI 10.17487/RFC4610, August 2006, <<http://www.rfc-editor.org/info/rfc4610>>.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<http://www.rfc-editor.org/info/rfc4875>>.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", RFC 5015, DOI 10.17487/RFC5015, October 2007, <<http://www.rfc-editor.org/info/rfc5015>>.
- [RFC5059] Bhaskar, N., Gall, A., Lingard, J., and S. Venaas, "Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)", RFC 5059, DOI 10.17487/RFC5059, January 2008, <<http://www.rfc-editor.org/info/rfc5059>>.

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC6388] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, DOI 10.17487/RFC6388, November 2011, <<http://www.rfc-editor.org/info/rfc6388>>.

Acknowledgments

The authors wish to thank DP Ayyadevara, Robert Kebler, Padmini Misra, Rayen Mohanty, Maria Napierala, Karthik Subramanian, and Kurt Windisch for their contributions to this work.

We also wish to thank Lizhong Jin and Rishabh Parekh for their reviews and comments.

Special thanks to Jeffrey (Zhaohui) Zhang for his careful review and for providing the ASCII art appearing in Section 2.

Contributors

Below is a list of other contributing authors, in alphabetical order:

Wim Henderickx
Nokia
Copernicuslaan 50
Antwerp 2018
Belgium

Email: wim.henderickx@nokia.com

Praveen Muley
Nokia

Email: Praveen.Muley@nokia.com

Ray Qiu
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
United States

Email: rqiujuniper.net

IJsbrand Wijnands
Cisco Systems, Inc.
De Kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

Authors' Addresses

Yakov Rekhter (editor)
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
United States

Eric C. Rosen (editor)
Juniper Networks, Inc.
10 Technology Park Drive
Westford, Massachusetts 01886
United States

Email: erosen@juniper.net

Rahul Aggarwal
Arktan

Email: ragnarwa_1@yahoo.com

Yiqun Cai
Alibaba Group
400 S El Camino Real #400
San Mateo, CA 94402
United States

Email: yiqun.cai@alibaba-inc.com

Thomas Morin
Orange
2 Avenue Pierre-Marzin
22307 Lannion Cedex
France

Email: thomas.morin@orange.com