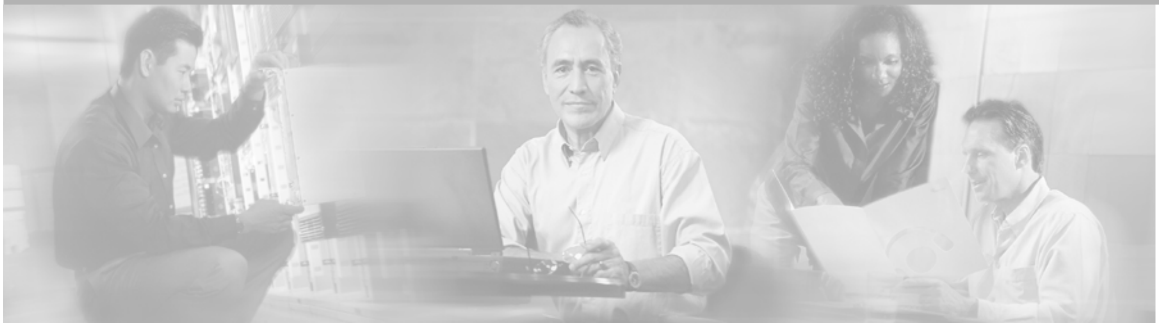

Multicast VPN Troubleshooting Guide

OL-10380-01

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Text Part Number: OL-10380-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Multicast VPN Troubleshooting Guide

CONTENTS

Preface

This guide provides troubleshooting guidelines for mVPN issues that are specific to service provider mVPN setups. This guide can also be used for checking a new mVPN customer setup or adding a new site to an existing mVPN setup. It is assumed that the user has good understanding of MPLS VPNs, multicast VPNs, and their configuration.

To make the troubleshooting easier, this guide is organized as a set of interdependent flowcharts, tables and output examples that guide the user, step-by-step. This guide is not meant to be used for troubleshooting traditional IP multicast issues; for that, refer to the *IP Multicast Troubleshooting Guide* listed in Appendix B "Related Documents."

Note Although this guide is comprehensive, it may not cover every possible mVPN troubleshooting scenario.

Document Purpose

This guide provides mVPN troubleshooting guidelines for the edge PE (MTI, core and customer-facing settings), and basic troubleshooting guidelines for core multicast. It also provides some common configuration errors that should be considered.

Target Audience

This guide is intended for the following users:

- Customers implementing mVPN
- Engineers who work in SP network operation centers (NOCs)

It is assumed that the user has prior understanding of the following:

  - RFC 4364 (formerly RFC 2547)
  - IP multicast
  - Cisco's implementation of mVPN

For details, refer to *IP Multicast Troubleshooting Guide* in Appendix B "Related Documents."

Document Scope

This document focuses on MVPN implementations leveraging the BGP MDT SAFI (SAFI 66), which was implemented in 12.0(29)S, 12.2(7)S and later IOS releases.

Table 1 lists the mVPN options that are covered in this guide:

Table 1 mVPN Options

	Core PIM Mode	Core RP Election Options	VRF PIM Mode	VRF RP Election Options
Default MDT	Bidir Pim-SM with SPT equals to 0	Static AutoRP		
Customer Multicast Groups			PIM-SM with SPT equals to 0	On customer side: static and AutoRP

The guide *does not* cover the following:

- Data MDT
- Inter AS mVPN
- Extranet mVPN
- Multihoming scenarios
- Platform dependent considerations and output
- SSM

It is important to note that although this document does not present troubleshooting information about SSM and Data MDTs, this should not be interpreted as a recommendation to not use such MVPN design improvements.

On the very opposite, Cisco strongly recommends the use of Data MDTs for optimization of an MVPN and SSM is the best PIM mode choice for Data MDTs when possible.

Table 2 lists the main troubleshooting areas for an mVPN setup.

Table 2 Multicast VPN Troubleshooting Areas

Section	Description
Common Configuration Errors, page 1-6	This section lists and describes some common mVPN configuration errors that the user of this guide should consider before getting into the main troubleshooting sections.
Section 1: Verifying the IP Multicast Core Settings for the Default MDT, page 2-3	This section provides troubleshooting guidelines for the core mVPN address range. The following are considered: • PIM modes and source path tree (SPT) threshold related issues • RP settings on P or PE routers.

Table 2 Multicast VPN Troubleshooting Areas (continued)

Section	Description
Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9	This section provides troubleshooting guidelines for the following: • For a given PE, how to ensure that the MTI tunnel interface is UP and associated to the VRF • How to ensure that the BGP and VRF settings for the default mdt are properly configured • How to ensure that the PE is able to establish PIM peerings with remote PEs over the MTI • How to ensure that the (*,GPA) and potential (S,GPA) are in the MRIB, and that the IIL, OIL and flags are appropriate
Section 3: Verifying the Customer-Facing Settings of the PEs, page 2-40	This section provides troubleshooting guidelines for determining whether the issue is on the SP network or the enterprise network. The following are included: • Checking VRF specific PIM and RP settings on the PEs • Analyzing customer specific mroute states on the PEs
Section 4: Verifying mVPN Data Flow, page 2-79	This section describes how to troubleshoot the mVPN data flow. This section only covers data flow on default MDT; not data MDT.
Chapter 3, "Using Device Instrumentation for mVPN Troubleshooting."	This section describes device instrumentation (DI) that are specific to mVPN setups

mVPN Troubleshooting: Getting Started

This chapter describes the following:

- The methodology used throughout the document.
- The first steps to take before troubleshooting an MVPN setup.
- Information that should be taken into account before starting to troubleshoot mVPN setups.
- Common configuration errors that should be considered when troubleshooting mVPN networks.
- Common troubleshooting scenarios.

Methodology of this Document

This section describes the methodology used in this document.

Test Setup

Figure 1-1 shows the test setup that was used for the troubleshooting techniques described in this guide. Note the following in this setup:

- C-RP, P-RP, and RR are outside of the forwarding path.
- OSPF is used as IGP in the core.
- Default MDT: VRF RED 239.0.0.10.
- VRF RED customer groups belong to address range: 225.1.1.1-> 225.1.1.10.
- Cisco IOS Release12.0(31)S1 is running on all routers.
- Traffic on the Customer source was simulated using a Test Tool (SRC1) and Received Traffic was also measured using a Test Tool (RCVR 12, 2, 3).

Figure 1-1 Test Setup

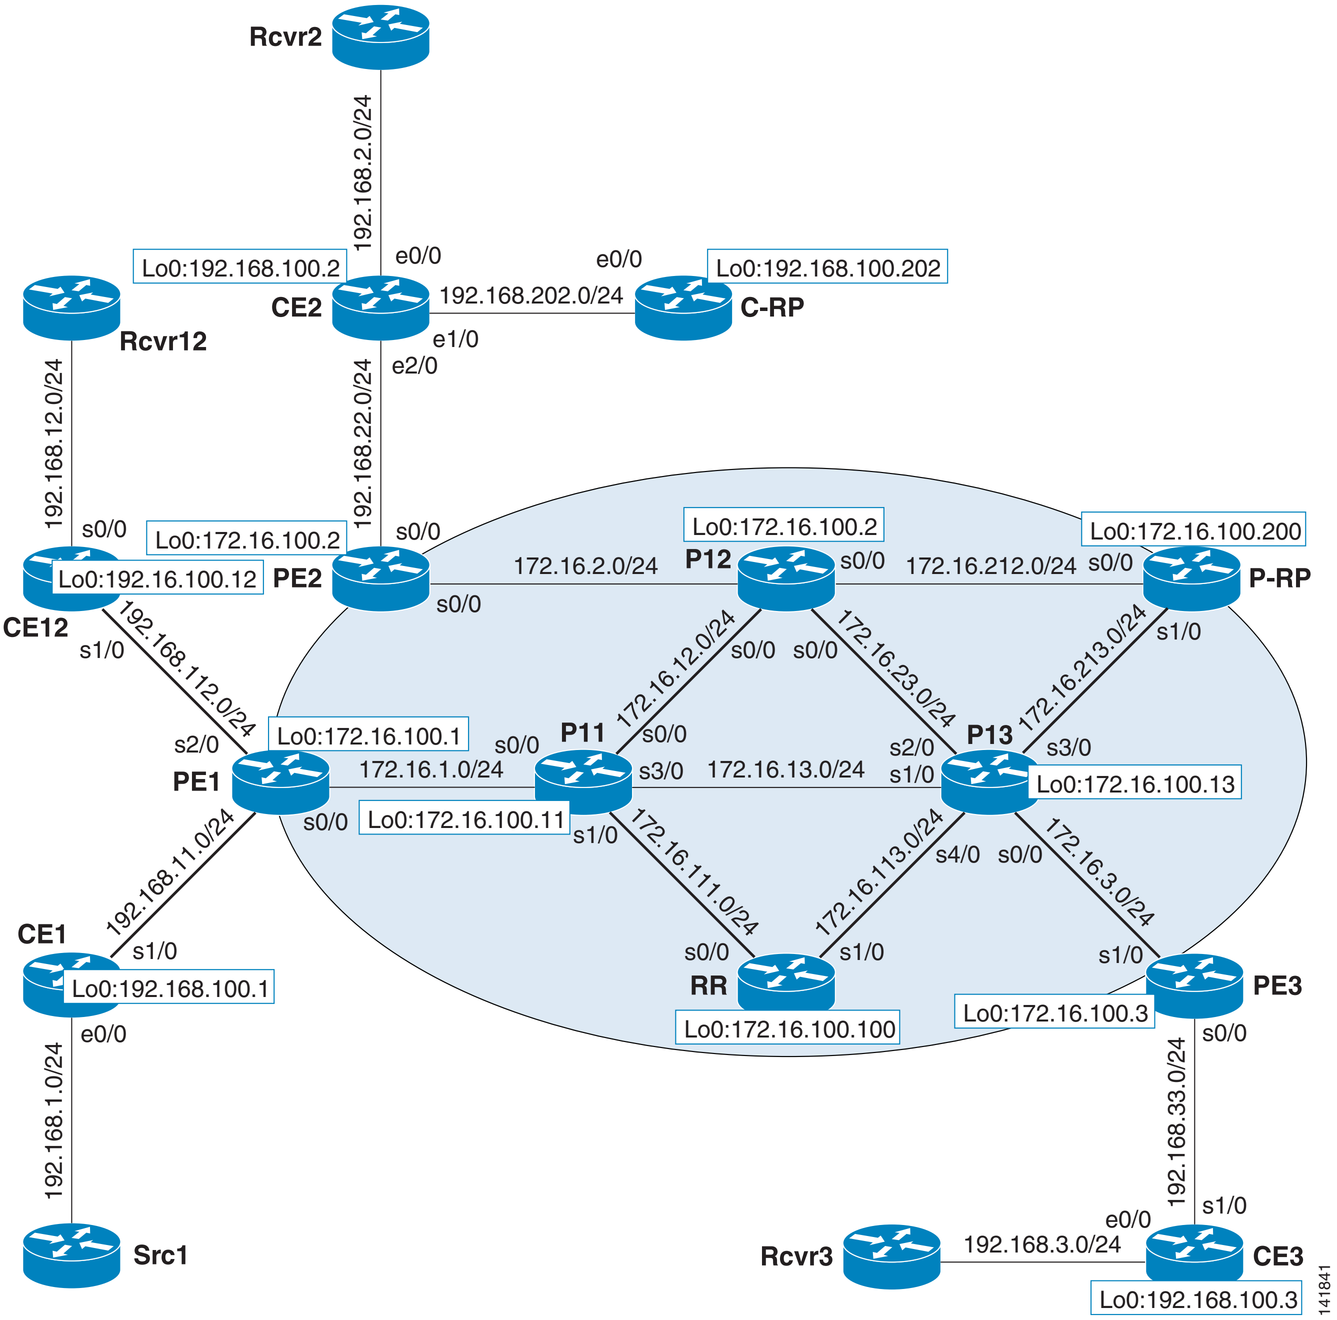

Assumptions and Prerequisites

Before going into the main troubleshooting sections that start with Chapter 2, "Troubleshooting mVPNs," consider the following assumptions:

- Unicast VPN traffic forwarding is working as intended.
- The IP multicast core is properly configured.

- All PE-CE interfaces are enabled for both unicast and multicast (PIM); otherwise, RPF does not work as intended.
- All PE platforms support the mVPN features and are running the appropriate Cisco IOS release.

  For more information, refer to the Cisco Feature navigator:

  http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp
- The IOS release supports the new BGP MDT SAFI.

When dealing with mVPN issues encountered for multiple customers, focus on resolving the issues for one customer at a time.

Notations

For acronyms and notation information, refer to Appendix C, "Terms and Acronyms." Note that some of the acronyms are used for the purpose of this document only.

Gathering Information

Before starting troubleshooting, gather as much information as possible about the mVPN issues and the customer's multicast network characteristics. The following are examples of customers' information to gather:

- C-RP IP address(es)
- C-Sources IP address(es)
- C-receivers IP address(es)
- C-sources traffic rates
- C-GPA

This information can be used to identify the $PE_{C\text{-}RCVR}$, $PE_{C\text{-}SRC}$, $PE_{C\text{-}RP}$, and default MDT GPA.

For the definition of these and other related terms used in this guide, refer to Appendix C, "Terms and Acronyms."

The following are examples of customer issues:

- Customers have problems connecting to a specific site (PE).
- No customer has multicast connectivity.
- Customer's multicast traffic is flowing, but not at the expected rates (partial traffic rates).
- The PEs cannot see each other as PIM neighbors over the MTI.
- One group (S,G) or (*,G) of a customer has issues.

Tests for Enhanced mVPN Troubleshooting

Because troubleshooting SP networks for unicast and multicast VPN involves many elements, the tests in Table 1 can be used to help the troubleshooting steps, if they are incorporated in the troubleshooting process. They provide sanity checks of the core routers and the core-facing PEs.

It is understood that some environments might not allow you to create new configuration states on a live network, so the verification of the core via a dummy vrf might not be feasible.

Note Ensure that the parameters for these tests are correctly configured so as to avoid false positive results.

Note Ensure that the dummy default mdt group address is not used for any customer.

Table 1 High level description of the Sanity Tests

Tests	Description
Test A: Verifying the mVPN Setup for a Dummy Customer	Verifies the core and edge settings for a test customer using all PE sites. Success of this test validates that the multicast core settings for the default mdt address **range** and the BGP settings on all PEs are correctly configured. If this test fails, refer to the troubleshooting guidelines provided in Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9.
Test B: Verifying the IP Multicast Core for the Default MDT Address Range	Verifies the core settings for all PEs sites. Success of this test validates that the multicast core settings for the default mdt address range on all core P, PEs, and RP routers are correctly configured. If this test fails, refer to the troubleshooting guidelines provided in Section 1: Verifying the IP Multicast Core Settings for the Default MDT, page 2-3.

Test A: Verifying the mVPN Setup for a Dummy Customer

The following test configuration creates a “dummy” or “test” mVPN customer. In this test, an mVRF is set up on all PEs for this test customer. Then, a **ping** is performed to check that all PEs in the test mVPN reply, which proves that the core and the generic BGP settings are correct for the mVPN MDT test address.

The test will be useful to quickly isolate MTI and PIM peering issues.

Note The RP used for the test Group address should be the same as the one used for the mVPNs you are troubleshooting.

Note The loopback used for the test VRF should not be a pre-existing loopback. Each PE belonging to the test mVPN will have such a loopback. PE loopback IP addresses can belong to any range but they cannot overlap.

Step 1 The following configuration is recommended on all the PEs, in addition to the standard mVPN configuration, for test A:

```
ip vrf test255
 rd X:Y
 route-target export W:Z
 route-target import W:Z
    mdt default 239.0.0.255
ip multicast-routing vrf test255

interface Loopback255
 ip vrf forwarding test255
 ip address 172.16.255.1 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-dense-mode
 ip igmp join-group 225.1.1.255

router bgp 1
address-family ipv4 vrf test255
redistribute connected
end
```

Note The rd X:Y, route-target W:Z, and the MDT default 239.0.0.255 address from MDT GPA range cannot be assigned to any existing customer. It should be the same on all PEs. One method to ensure availability it to check the Route-Reflector (or multiple Route-Reflectors if partitioning is used).

Step 2 After the above configuration is performed and you enter a **ping** from one of the PEs, a reply is generated and displayed from all remote PEs:

```
PEC-SRC#ping vrf test255 225.1.1.255

Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 225.1.1.255, timeout is 2 seconds:

Reply to request 0 from 172.16.255.1, 24 ms
Reply to request 0 from 172.16.100.2, 72 ms
Reply to request 0 from 172.16.100.3, 72 ms
```

Test B: Verifying the IP Multicast Core for the Default MDT Address Range

This test configuration verifies that the SP core is functional for the mVPN address range.

Step 1 For this test, each PE should be configured with an extra loopback (not the one used as router ID for IGP, LDP, or BGP). The following should be configured on the loopback:

```
ip igmp join-group 239.0.0.254
ip pim sparse mode
```

Note The address from the default MDT GPA range cannot be assigned to any customer MDT.

This example assumes that the default MDT range is 239.x.x.x/8 and 239.0.0.254 is not in use. If this address is used, use a non-assigned address from the MDT default range.

Step 2 After the above configuration is performed and you enter the **ping** command from one of the PEs, a reply is generated and displayed from all remote PEs. Ensure that you enter the ping source address precisely to avoid multiple echo replies:

```
PEC-SRC#ping ip
Target IP address: 239.0.0.254
Repeat count [1]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Interface [All]: loopback 0
Time to live [255]:
Source address:
Type of service [0]:
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 1, 100-byte ICMP Echos to 239.0.0.254, timeout is 2 seconds:

Reply to request 0 from 172.16.100.1, 20 ms
Reply to request 0 from 172.16.2.2, 84 ms
Reply to request 0 from 172.16.3.3, 44 ms
```

- There will be a ping reply from core-facing interfaces of all remote PEs configured for the **igmp join-group 239.0.0.254** command in the global interface.
- There will be a ping reply from the IP address of the local loopback interface configured for the **igmp join-group 239.0.0.254** command.

Common Configuration Errors

Refer to Appendix A and the Design Guides in the References Section for full configurations and design guidance.

Consider the following common errors when troubleshooting an mVPN configuration:

Releases and Platform Support

- Use Feature Navigator to ensure that there is support for mVPN in the platforms and in the Cisco IOS release being used. For more information refer to the following URL (note that a Cisco user ID and password is required).

  http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp
- Ensure that the chosen release supports the latest BGP MDT SAFI implementation.

Default MDT Configuration

- Ensure that the range of default MDT is configured with the same PIM mode and RP election option on all the PEs.
- Ensure that the same MDT default group address is defined on all PEs that belong to a given mVPN.

BGP Configuration

Loopback and BGP Next Hop Configuration

- Ensure that the BGP update source interface from which the MTI takes its properties (normally a loopback) is configured for PIM.
- Ensure that the MTI is unnumbered to the appropriate BGP update-source interface. Normally this is the default, unless the **bgp next-hop** command is used in the VRF context and overrides the default.
- Ensure that the local PE router uses the same update-source/ next hop reference interface to peer with all the remote PEs that belong to the mVPN.

BGP MDT SAFI Configuration

- Configure the Address family MDT SAFI neighbors exactly as the VPNv4 address family (Neighbor addresses and update source values).

Customer Setup

- Ensure that C-SRC is generating multicast stream with the appropriate TTL value.

Enabling Multicast in the Core and Between the PEs and CEs

Checking the Pim Settings

- Ensure that PIM is enabled on the core facing interfaces; use the **show ip pim interface** command to verify.
- Ensure that the **ip multicast-routing** global command is enabled.
- Ensure that the **ip multicast-routing vrf** command is enabled.
- Ensure that the *all* PE-CE interfaces on the PEs are enabled for the appropriate PIM mode. Use the **show ip pim vrf** *vrf-name* **neighbors** command to verify.
- Verify that the PIM SPT threshold value is consistent.

Multicast Boundaries and ACLs

- Ensure that ACLs, multicast-boundary, multicast ttl-threshold, IP PIM neighbor filter statements, and multicast rate-limit configured on the PE-CE interface or on the core links are not blocking any traffic.

Checking the RP Settings

- Ensure that P-RP is known and reachable by all P and PE routers.
- Ensure that Group-to-RP mapping is consistent on all Ps and PEs. If the PE router is, for instance, configured with an RP address that is different from the one configured on another P or PE router, the following error may be generated:%PIM-6-INVALID_RP_JOIN: Received (*, GPA) Join from <PIM-NBR> for invalid RP <RP-ADDR> [NOTE PIM-NBR is printed out all 0s before the update of CSCei28317].
- For PIM SM, ensure that all routers in the network are configured with the same **ip pim spt-threshold** [0|**infinity**] command (function of the Multicast address range).
- If using PIM Bidir with static RP, verify that the **ip pim bidir- enable** command is enabled on all core and edge routers.

- Verify the following RP settings:
  - Verify that for PIM SM and Bidir with static RP the **ip pim rp-address** *address* [**group-list** *acl*] [**override**] [**bidir**] command is enabled on the core and edge routers.
  - The **group-list** allows a group range to be specified.
  - The default is ALL multicast groups or 224.0.0.0/4.
  - The **override** keyword permits the statically defined RP address to take precedence over Auto-RP learned Group-to-RP mapping information. The default behavior without this keyword is that Auto-RP learned information has precedence over static information.
- Check the following potential RP-filters:
  - Check for any filters using the **ip pim accept-rp** command and verify that it is configured to allow the RP-address, including on the RP itself.
  - Check for **ip pim accept-register filters** command on the RP.
  - Check for **ip pim rp-announce-filter rp-list** *group-list* command on the mapping agent for Auto-RP.
  - Check for group-lists or ACL post-pended to commands, such as **ip pim send-rp-announce** and **ip pim rp-address**.

PIM SM with AutoRP

The AutoRP groups (224.0.1.39 and 224.0.1.40) need to be densely switched, which is achieved by one of the following:

- Enabling PIM sparse-dense mode on all PIM interfaces.
- Enabling PIM sparse-mode on all PIM interfaces and enabling the **ip pim autorp listener** command.

Note If a static RP is defined as an RP of last resort, which is most of the time for the router itself, ensure that AutoRP groups are denied by the ACL in the static RP command line.

If both of the above two conditions are met and if the **show ip pim rp mapping** command still does not display any RP mappings learned through AutoRP, then:

- Check the **show ip mroute 224.0.1.39** command on the mapping agent and verify that you see the (S,224.0.1.39), where S equals the RP address.
- Ensure that if there is a filter list on the mapping agent, it allows the valid RP address/group range.
- For more information, refer to the following URL: http://www.cisco.com/en/US/tech/tk828/technologies_configuration_example09186a00801cb923.shtml
- On the PE/P routers, ensure that there are RP discovery messages from the mapping agent. Enter the **show ip mroute 224.0.1.40** command on a router and see if there is the following: (S,224.0.1.40), where S equals the mapping agent.

Common Troubleshooting Scenarios

The following tables provide possible specific mVPN-related scenarios.

- Data Multicast Flow not seen on one PE
- Data Multicast Flow not seen on several PEs
- Control plane issues

Note This section is not meant to replace the more detailed troubleshooting sections, but rather to provide some level of guidance and some precise steps to take in case of common troubleshooting circumstances.

The section assumes that unicast VPN connectivity to and from the customer sites hosted on the corresponding PE is working well.

Data Multicast Flow Not Seen On One PE

Table 1-2 Scenarios for Data Flow Issues on one PE

Scenario	Recommended Steps
No customer connected to one given PE can receive or send multicast data traffic	**1.** Apply the guidelines described in Common Configuration Errors, page 1-6 for each specific C-Src and C-GPA of this customer. **2.** Perform Test A: Verifying the mVPN Setup for a Dummy Customer andTest B: Verifying the IP Multicast Core for the Default MDT Address Range **a.** If Test A: Verifying the mVPN Setup for a Dummy Customer fails, first check Test B: Verifying the IP Multicast Core for the Default MDT Address Range. **b.** If Test B is successful and Test A fails, check the PE edge settings (refer to Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9 for troubleshooting guidelines). **c.** If both Test A and B fail, check the core settings (refer to Section 1: Verifying the IP Multicast Core Settings for the Default MDT, page 2-3 for troubleshooting guidelines). **3.** If you have not implemented Test A and Test B, focus on one customer and apply the guidelines in Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9 for that customer's default mdt-GPA. **4.** Refer to Section 4: Verifying mVPN Data Flow, page 2-79 for any data flow related issues.
One customer has a data issue on a given PE site, while other customers do not	**1.** Refer to Common Configuration Errors, page 1-6 to ensure there is no common error in the configuration. **2.** Focus on the corresponding PE(s). **3.** Ensure unicast VPN connectivity to and from the customer sites hosted on this PE. **4.** Apply the guidelines provided in Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9) on this PE for the customer's default mdt-GPA. **5.** Check the data Plane using Section 4: Verifying mVPN Data Flow, page 2-79

Data Multicast Flow Not Seen on Several PEs

Table 1-3 Scenarios for Data Flow Issues on Multiple PEs

Scenario	Recommended Steps
One customer has a data traffic issue across all of their sites; while other customers don't	1. Refer to Common Configuration Errors, page 1-6 to ensure there is no common error in the configuration. 2. Apply the guidelines in the following: • Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9) for this customer's default mdt-GPA • Section 3: Verifying the Customer-Facing Settings of the PEs, page 2-40) for each specific C-Src and C-GPA of this customer. • Section 4: Verifying mVPN Data Flow, page 2-79.
One particular C-group for a specific customer presents issues, other groups are fine	1. Refer to Common Configuration Errors, page 1-6 to ensure there is no common error in the configuration. 2. Apply the guidelines provided in Section 3: Verifying the Customer-Facing Settings of the PEs, page 2-40. 3. Apply data flow checks provided in Section 4: Verifying mVPN Data Flow, page 2-79.
The customer is only seeing a partial traffic rate for a given group and source	1. Apply the guidelines in Section 4: Verifying mVPN Data Flow, page 2-79 2. Check Section 3: Verifying the Customer-Facing Settings of the PEs, page 2-40 if the issue remains.
One particular C-source of a customer presents issues, others are fine- it has issues with several Groups	1. Refer to Common Configuration Errors, page 1-6 to ensure there is no common error in the configuration. 2. Apply Section 3: Verifying the Customer-Facing Settings of the PEs, page 2-40). 3. Communicate with customer to ensure the settings are correct. 4. Apply the guidelines in Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9) for checking that the $PE_{C\text{-}SRC}$ is located on the site of the C-SRC. 5. Apply data flow checks provided in Section 4: Verifying mVPN Data Flow, page 2-79.

Control Plane Issues on the PE-CE Link(s) or Between PEs

Table 1-4 Scenarios for Control Plane Issues on one or more PE(s)

Scenario	Recommended Steps
The PIM neighborship CE-PE session is down or needs to be checked on a specific Site	**1.** Refer to Common Configuration Errors, page 1-6 to ensure there is no common error in the configuration. **2.** Focus on that specific PE. **3.** Apply the guidelines provided in Section 3: Verifying the Customer-Facing Settings of the PEs, page 2-40). **4.** Communicate with the customer to check the customer-side network setup. **5.** Apply the guidelines provided in Section 4: Verifying mVPN Data Flow, page 2-79.
The PEs do not see each other as PIM neighbors over the MTI	**1.** Refer to Common Configuration Errors, page 1-6 to ensure there is no common error in the configuration. **2.** Perform Test A: Verifying the mVPN Setup for a Dummy Customer andTest B: Verifying the IP Multicast Core for the Default MDT Address Range **a.** If Test A: Verifying the mVPN Setup for a Dummy Customer fails, first check Test B: Verifying the IP Multicast Core for the Default MDT Address Range. **b.** If Test B is successful and Test A fails, check the PE edge setting (refer to Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9 for troubleshooting guidelines). **c.** If both Test A and B fail, check the core settings (refer to Section 1: Verifying the IP Multicast Core Settings for the Default MDT, page 2-3 for troubleshooting guidelines). **3.** If you have not implemented Test A and Test B, apply the guidelines provided in Section 1: Verifying the IP Multicast Core Settings for the Default MDT, page 2-3). **4.** Apply the guidelines provided in Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI, page 2-9) to the PEs.

Troubleshooting mVPNs

This chapter describes how to troubleshoot the mVPN network. The following sections are included:

Reference Diagram for Troubleshooting an mVPN Setup

Figure 2-1 shows the main reference diagram for troubleshooting an mVPN setup. This diagram is the point of reference for all troubleshooting described in this guide.

Figure 2-1 High-level Diagram for Troubleshooting an mVPN Setup

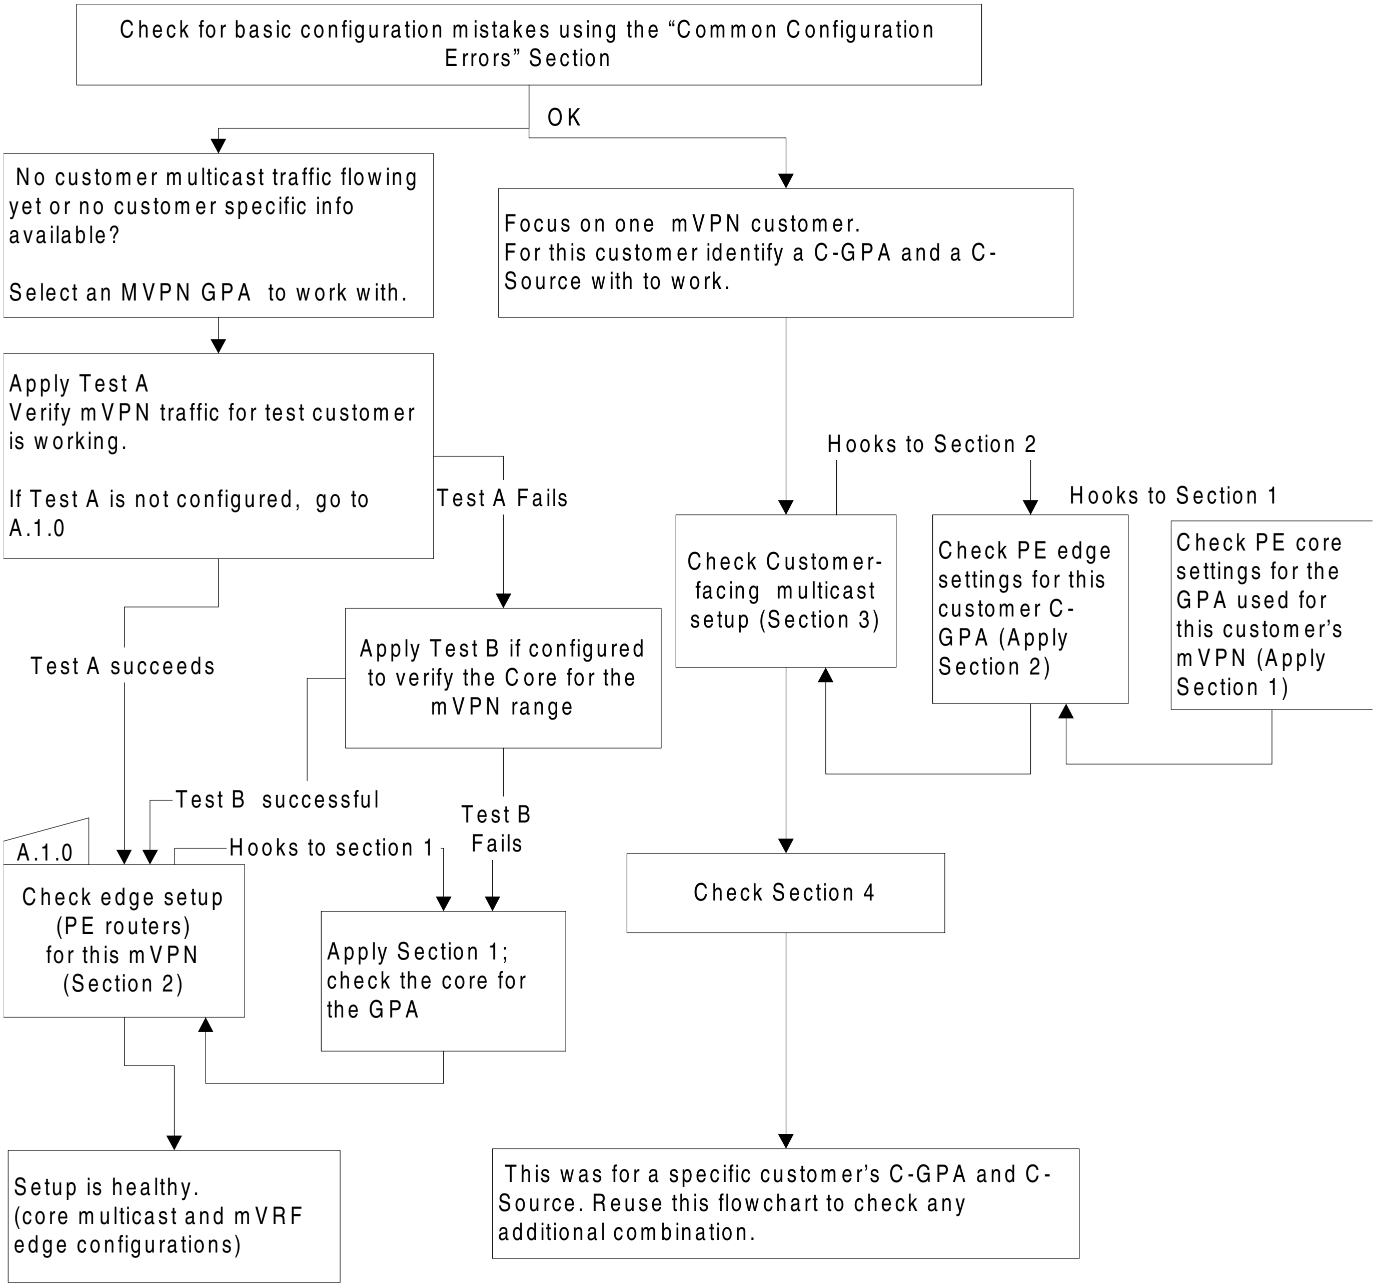
Check for basic configuration mistakes using the "Common Configuration Errors" Section
OK
No customer multicast traffic flowing yet or no customer specific info available?
Select an MVPN GPA to work with.
Focus on one mVPN customer.
For this customer identify a C-GPA and a C-Source with to work.
Apply Test A
Verify mVPN traffic for test customer is working.
If Test A is not configured, go to A.1.0
Test A Fails
Hooks to Section 2
Hooks to Section 1
Check Customer-facing multicast setup (Section 3)
Check PE edge settings for this customer C-GPA (Apply Section 2)
Check PE core settings for the GPA used for this customer's mVPN (Apply Section 1)
Test A succeeds
Apply Test B if configured to verify the Core for the mVPN range
Test B successful
Test B Fails
Hooks to section 1
A.1.0
Check edge setup (PE routers) for this mVPN (Section 2)
Apply Section 1; check the core for the GPA
Check Section 4
Setup is healthy. (core multicast and mVRF edge configurations)
This was for a specific customer's C-GPA and C-Source. Reuse this flowchart to check any additional combination.

Section 1: Verifying the IP Multicast Core Settings for the Default MDT

This section provides a basic outline to help troubleshoot IP multicast in the SP core, with an emphasis on the MDT default address.

Failure of proper operation of multicast in the SP's core network can impact all mVPN customers' traffic. This section does not provide details on traditional IP multicast troubleshooting. For any traditional IP multicast troubleshooting, refer to the I*P Multicast Troubleshooting Guide* at the following URL:

http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080094b55.shtml

The following are possible configuration errors that can be detected:

- Misconfigured PIM configurations on any PE, P, or P-RP routers
- Inconsistent SPT threshold setting
- Inconsistent RP information across the SP core
- RPF check failure for the RP or one of the sources

The following information should be available:

- GPA
- PIM mode (SM or Bidir) for the GPA
- Core RP option (static or Auto RP) and RP IP address
- SPT threshold setting

Ensure that at least two PEs are configured for this mVPN.

Note For this section, we recommend that you implement the test described in Test B: Verifying the IP Multicast Core for the Default MDT Address Range, page 1-5.

Most of the potential issues identified when troubleshooting the IP multicast core for a default mdt group address are IP multicast issues. For this reason, this section focuses only on the IP address range corresponding to the default mdt group address and not on all IP multicast groups that might use the IP multicast core.

Figure 2-2 highlights the areas that are specific to an mVPN scenario where the multicast group address is an default mdt multicast address. Following the diagram are samples of output that can be used for guidance. These output examples correspond to the **show** commands from the diagram.

Figure 2-2 shows a state machine that provides high level guidance. Select an initial core router (PE, RP, or P router) to start with, similar to troubleshooting traditional IP multicast. You can then move, hop-by-hop, as necessary on any additional core router, to verify the IP multicast core for the GPA address configured.

Figure 2-2 Troubleshooting the Multicast SP Core for a given GPA Address

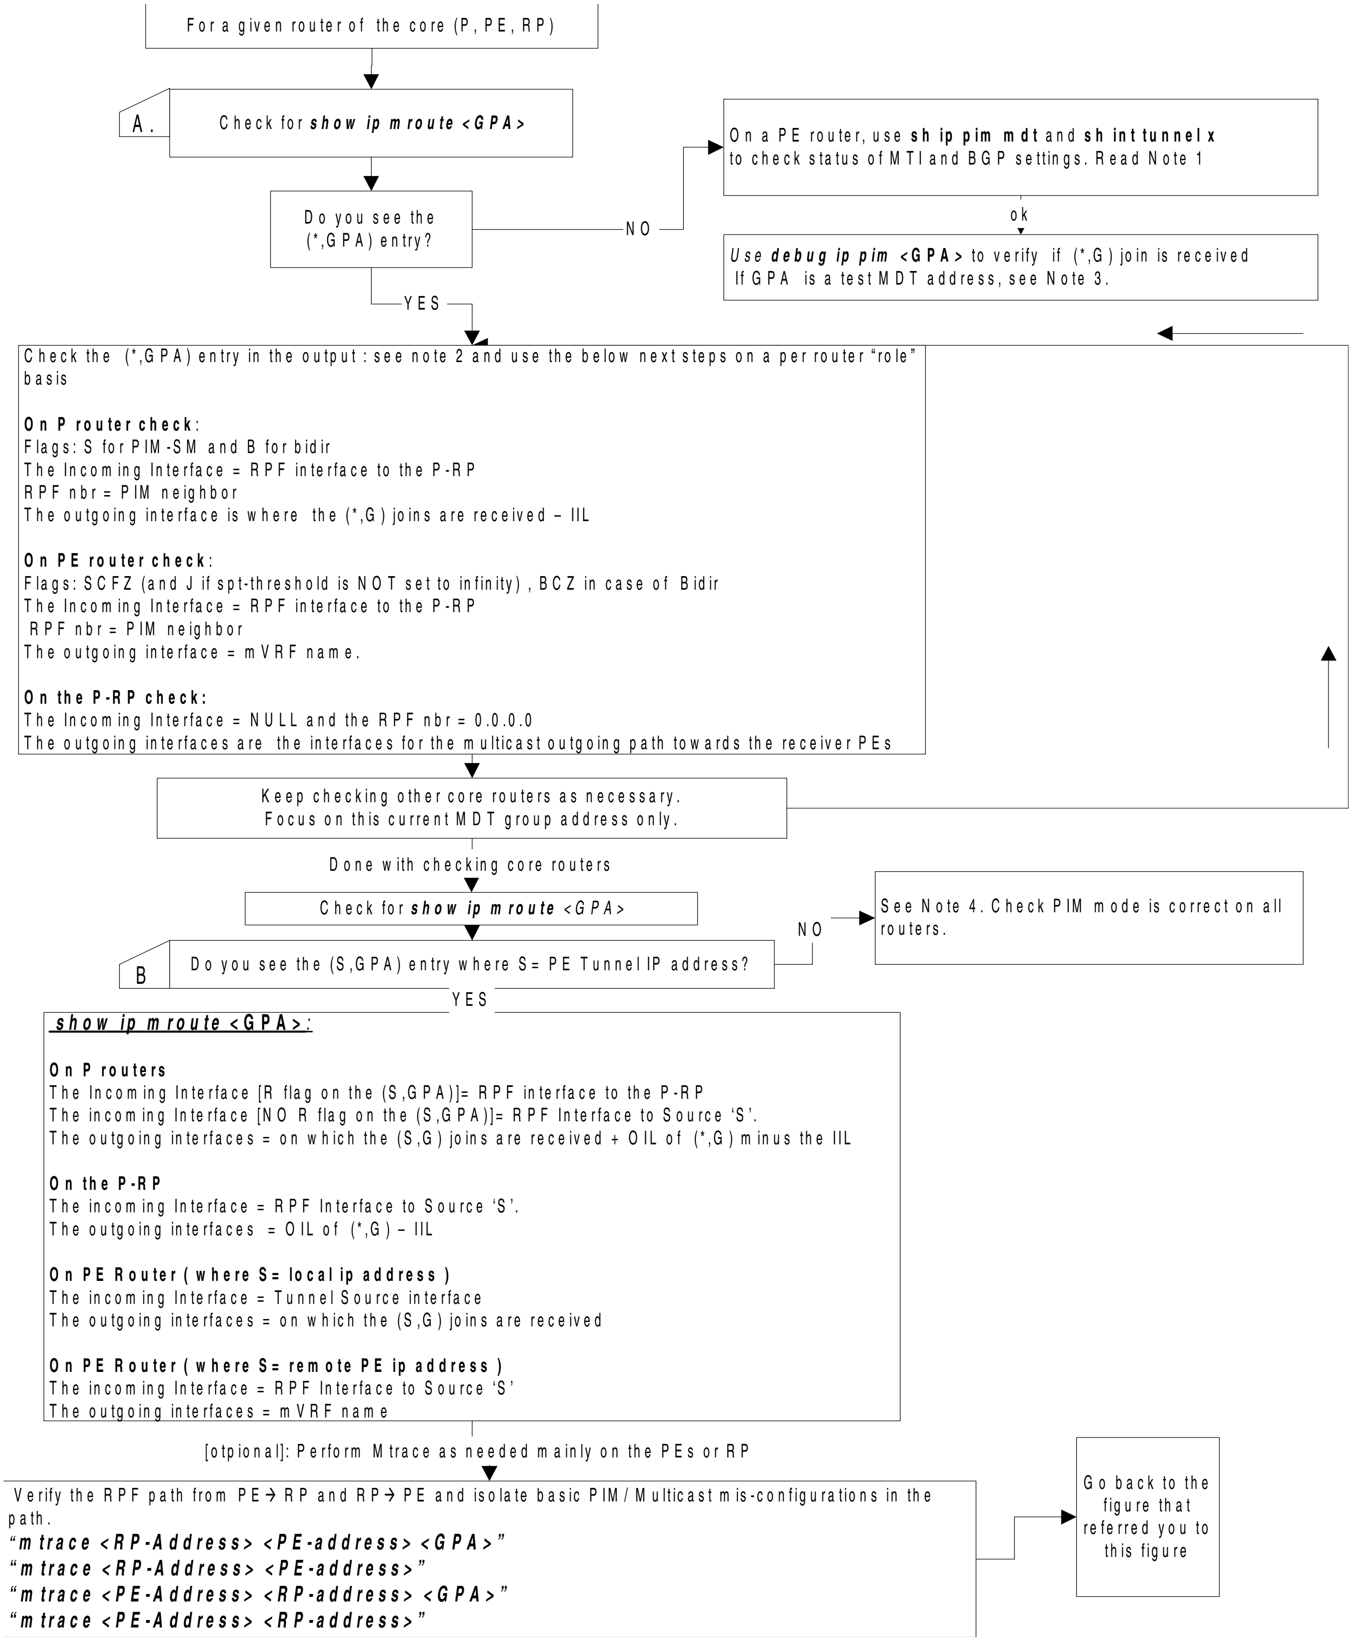

For a given router of the core (P, PE, RP)

Check for ***show ip mroute <GPA>***

Do you see the (*,GPA) entry?

On a PE router, use **sh ip pim mdt** and **sh int tunnel x** to check status of MTI and BGP settings. Read Note 1

Use ***debug ip pim <GPA>*** to verify if (*,G) join is received
If GPA is a test MDT address, see Note 3.

Check the (*,GPA) entry in the output : see note 2 and use the below next steps on a per router "role" basis

On P router check:
Flags: S for PIM-SM and B for bidir
The Incoming Interface = RPF interface to the P-RP
RPF nbr = PIM neighbor
The outgoing interface is where the (*,G) joins are received – IIL

On PE router check:
Flags: SCFZ (and J if spt-threshold is NOT set to infinity) , BCZ in case of Bidir
The Incoming Interface = RPF interface to the P-RP
RPF nbr = PIM neighbor
The outgoing interface = mVRF name.

On the P-RP check:
The Incoming Interface = NULL and the RPF nbr = 0.0.0.0
The outgoing interfaces are the interfaces for the multicast outgoing path towards the receiver PEs

Keep checking other core routers as necessary.
Focus on this current MDT group address only.

Done with checking core routers

Check for ***show ip mroute*** *<GPA>*

See Note 4. Check PIM mode is correct on all routers.

Do you see the (S,GPA) entry where S= PE Tunnel IP address?

show ip mroute <GPA>:

On P routers
The Incoming Interface [R flag on the (S,GPA)]= RPF interface to the P-RP
The incoming Interface [NO R flag on the (S,GPA)]= RPF Interface to Source 'S'.
The outgoing interfaces = on which the (S,G) joins are received + OIL of (*,G) minus the IIL

On the P-RP
The incoming Interface = RPF Interface to Source 'S'.
The outgoing interfaces = OIL of (*,G) – IIL

On PE Router (where S= local ip address)
The incoming Interface = Tunnel Source interface
The outgoing interfaces = on which the (S,G) joins are received

On PE Router (where S= remote PE ip address)
The incoming Interface = RPF Interface to Source 'S'
The outgoing interfaces = mVRF name

[otpional]: Perform Mtrace as needed mainly on the PEs or RP

Verify the RPF path from PE→ RP and RP→ PE and isolate basic PIM/Multicast mis-configurations in the path.
"mtrace <RP-Address> <PE-address> <GPA>"
"mtrace <RP-Address> <PE-address>"
"mtrace <PE-Address> <RP-address> <GPA>"
"mtrace <PE-Address> <RP-address>"

Go back to the figure that referred you to this figure

Notes for Figure 2-2

Note 1

- Ensure that the MTI and BGP configurations are correct.
- If the configurations are incorrect, correct them, and then re-check using the **show ip mroute GPA** command.
- If the router is not a PE or RP, note that the P router might not be part of the multicast core tree for this MDT address.
- Check regular PIM settings for the P or RP router if the (*,GPA) entry is not there and it should be. Refer to the traditional multicast troubleshooting documents for more information.

Note Refer to the *IP Multicast Troubleshooting Guide* to solve general IP multicast issues: http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080094b55.shtml

- When the issue is corrected, re-check by following this flowchart from the top.

Note 2

- Verify the following for all the core P, PE, and RP routers:
  - The RP address is P-RP.
  - The flag is S for PIM-SM and B for the Bidir case.
- Refer to Sample Output for Figure 2-2, page 2-5 as a reference (as needed).
- As most of the potential issues here are IP multicast-related, correct them using traditional IP multicast troubleshooting guides, then re-check using the **show ip mroute** *GPA* output.

Note 3

It seems that we are not receiving any (*,G) join. If the MDT group address is a test MDT address (for example: address used for Test B setup), ensure that it is not a false positive result.

Note 4

This might be normal if the core is configured for PIM Bidir. In all other cases we should at least have the local PEs (S,GPA) if the MTI/BGP is configured correctly.

Sample Output for Figure 2-2

The following sample output is obtained from the **show ip mroute** command on the PE and P routers (nonRP) in the Bidir mode. Note the B flag. Also note the MDT Z flag.

```
PE1#show ip mroute 239.0.0.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(*, 239.0.0.10), 1w4d/00:02:53, RP 172.16.100.200, flags: BCZ
  Bidir-Upstream: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 1w4d/00:00:00
    Serial0/0, Bidir-Upstream/Sparse, 1w4d/00:00:00

P11#show ip mroute 239.0.0.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 1w4d/00:03:25, RP 172.16.100.200, flags: B
  Bidir-Upstream: Serial3/0, RPF nbr 172.16.13.13
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 1w4d/00:03:04
    Serial3/0, Bidir-Upstream/Sparse, 1w4d/00:00:00
```

Use the following command to check if IP multicast routing is enabled on the router:

```
PE1#show ip multicast
  Multicast Routing: enabled
  Multicast Multipath: disabled
  Multicast Route limit: No limit
  Multicast Triggered RPF check: enabled
  Multicast Fallback group mode: Sparse
```

Use the **show ip pim interface** command to display information about interfaces configured for PIM. Ensure that all the core-facing interfaces and the loopbacks are listed. This command can also be used to verify that the correct PIM mode is configured on the interface, the neighbor count is correct, and the designated router (DR) is correct (which is critical for PIM sparse mode).

```
PE1#show ip pim interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
172.16.100.1     Loopback0                v2/S   0      30     1      172.16.100.1
172.16.100.99    Loopback99               v2/S   0      30     1      172.16.100.99
172.16.1.1       Serial0/0                v2/S   1      30     1      0.0.0.0
172.16.255.1     Loopback254              v2/S   0      30     1      172.16.255.1
```

Use the **show ip pim interface count** command to verify that Fast Switching is enabled on all interfaces, including loopbacks. If this feature is not enabled, use the **ip mroute-cache** command to enable it.

```
PE1#show ip pim interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
172.16.100.1     Loopback0                *   211188/0
172.16.100.99    Loopback99               *   4709/0
172.16.1.1       Serial0/0                *   20136/206330
172.16.255.1     Loopback254              *   6/0
```

Use the **show ip pim neighbor** command to verify that you see PIM neighbors on the core-facing interface. Check for the B flag under Mode if you are running Bidir in the core. The following output is for PIM-SM and Bidir, respectively.

```
PE1#show ip pim neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
172.16.1.11       Serial0/0                19:28:56/00:01:38 v2    1 / P


PE1#show ip pim neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
172.16.1.11       Serial0/0                00:08:00/00:01:38 v2    1 / B P
```

Use the **show ip pim rp mapping** command to check the RP assignment for the MDT-default multicast group range, and to verify that the source of RP learning (static or AutoRP) and the mapping are correct and consistent with rest of the network.

```
PE1#show ip pim rp mapping
PIM Group-to-RP Mappings

Acl: 1, Static-Override
    RP: 172.16.100.200 (?)
RPF checks
```

Use the **show ip rpf** *ip_address* command to display how IP multicast routing does Reverse Path Forwarding (RPF). When you troubleshoot, use it to verify that the RPF information is correct. If it is not, check the unicast routing table for the source address. Also use the **ping** and **trace** commands on the source address to verify that unicast routing works.

```
PE1#show ip rpf 172.16.100.200
RPF information for ? (172.16.100.200)
  RPF interface: Serial0/0
  RPF neighbor: ? (172.16.1.11)
  RPF route/mask: 172.16.100.200/32
  RPF type: unicast (ospf 1)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
```

The following example shows output of RP mapping check:

```
RTR#show ip pim rp mapping
PIM Group-to-RP Mappings

Acl: 1, Static
    RP: 172.16.100.200 (?)
```

Show ip mroute for PIM SM with SPT Threshold 0

The following example shows output of the **show ip mroute** command for the default mdt address.

On the P-RP:

```
P-RP#show ip mroute 239.0.0.10 | begin 239
(*, 239.0.0.10), 18:48:26/00:03:21, RP 172.16.100.200, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 18:48:25/00:02:47
    Serial1/0, Forward/Sparse, 18:48:26/00:03:21
```

```
(172.16.100.1, 239.0.0.10), 18:48:25/00:03:27, flags: T
  Incoming interface: Serial1/0, RPF nbr 172.16.213.13
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 18:48:25/00:02:47

(172.16.100.2, 239.0.0.10), 18:48:25/00:00:54, flags: PT
  Incoming interface: Serial0/0, RPF nbr 172.16.212.12
  Outgoing interface list: Null
```

Note in the above output, the OIL (outgoing interface list) is Null because the (PE2, 239.0.0.10) Tree was pruned (P flag) on the P-RP, as the source path tree was constructed directly between the source PE2 and the other PEs as leaves.

```
(172.16.100.3, 239.0.0.10), 18:48:26/00:03:27, flags: T
  Incoming interface: Serial1/0, RPF nbr 172.16.213.13
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 18:48:25/00:02:47
```

On a P router:

```
P12#show ip mroute 239.0.0.10 | begin 239
(*, 239.0.0.10), 18:48:00/00:02:43, RP 172.16.100.200, flags: S
  Incoming interface: Serial2/0, RPF nbr 172.16.212.200
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 18:48:00/00:02:43

(172.16.100.2, 239.0.0.10), 18:48:00/00:03:24, flags: T
  Incoming interface: Serial0/0, RPF nbr 172.16.2.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 18:47:30/00:02:58
    Serial3/0, Forward/Sparse, 18:48:00/00:02:44
```

On a PE router:

```
PE1#show ip mroute 239.0.0.10 | begin 239
(*, 239.0.0.10), 17:20:07/stopped, RP 172.16.100.200, flags: SJCFZ
  Incoming interface: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 17:20:07/00:00:00

(172.16.100.1, 239.0.0.10), 17:20:07/00:03:26, flags: FT
  Incoming interface: Loopback0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 17:19:15/00:03:20
```

Note The above is the local (S,G) source entry on PE1 (here S= PE1), which explains why there is no Z flag.

```
(172.16.100.2, 239.0.0.10), 17:20:02/00:02:56, flags: JTZ
  Incoming interface: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 17:20:02/00:00:00

(172.16.100.3, 239.0.0.10), 17:19:56/00:02:48, flags: JTZ
  Incoming interface: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 17:19:56/00:00:00
```

Note For a remote PE, there should not be (S,GPA) entry when SPT threshold is set to infinity; but there should be an mstate (S,GPA) entry for the local PE.

Show ip mroute for PIM Bidir

The following example shows output of the **show ip mroute** command for the default mdt group address:

```
PE1#show ip mroute 239.0.0.10 | begin 239

(*, 239.0.0.10), 00:30:00/00:02:51, RP 172.16.100.200, flags: BCZ
  Bidir-Upstream: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 00:30:00/00:00:00
    Serial0/0, Bidir-Upstream/Sparse, 00:30:00/00:00:00


P-RP#show ip mroute | begin 239

(*, 239.0.0.10), 02:34:18/00:02:42, RP 172.16.100.200, flags: B
  Bidir-Upstream: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 02:34:16/00:02:30
    Serial0/0, Forward/Sparse, 02:34:18/00:02:42
```

Using the mtrace Command

The **mtrace** command shows the multicast path from the source to the receiver. This command traces the path between points in the networks, which shows TTL thresholds and delay at each node. When troubleshooting, use the **mtrace** command to find where multicast traffic flow stops, to verify the path of multicast traffic, and to identify sub-optimal paths. In the **mtrace** command, when testing the default mdt, use the MTI or the RP address as source address, not the PE-P core interfaces addresses.

Refer to Using the Mtrace Tool in Troubleshooting Multicast Issues, page 3-7 for details about the **mtrace** command.

Section 2: Verifying the PE-PE PIM Neighbor Relationship over the MTI

This section is divided into several subsections to guide you through the various levels of checks needed on a PE to ensure a healthy interaction with the other PEs that belong to the mVPN, as described below:

- Ensure that the PE settings needed for the MTI interface are valid with respect to PIM parameters, BGP peerings, and IP address.
- Check the mstates for the default mdt in the multicast routing table.

Note that conceptually the PEs are PIM neighbors over the MTI interface because they are one hop away from each other when considering the tunnel interface. Thus, they behave as if they were all directly connected from a PIM point of view and the MTI is the interface towards a virtual shared media.

The high-level flowchart in Figure 2-3 describes how to navigate through the different subsections and should be used as the main reference when troubleshooting the MTI configuration and the PE to PE communication over the MTI.

Consider the following assumptions for the troubleshooting in this section:

- Focus on the mVPN edge, but not the VPN customer multicast traffic.
- Focus on a specific default MDT group.
- Know the number, characteristics, and location of the mVPN PEs and their BGP next hop and BGP update-source information.
- The MPLS VPN unicast connectivity should be operating without any problems.

For each PE, ensure the following:

- The MTI tunnel interface is UP/UP and its configured parameters are the same as for the other PEs.
- The BGP and VRF settings for the default mdt are properly configured.
- The core PIM mode and settings are consistent for all PEs.
- The PE is able to establish PIM peerings with remote PEs over the MTI.
- The (*,GPA) and potential (S,GPA) are in the MRIB and their parameters are configured as intended.

Note that troubleshooting the mVPN data flow (including the core) is described under Section 4: Verifying mVPN Data Flow, page 2-80.

Gathering Information

Before troubleshooting the default mdt settings of the edge PEs, use the following questions to gather as much information as possible about the traffic and issues of the enterprise customers. This will help to select a subset of PEs to start the troubleshooting (unless you already have selected the PEs you want to troubleshoot based on other criteria).

- Is this customer the only enterprise customer that has issues with its multicast traffic?
- Is there a specific POP (PE location) that has multicast issues?
- Which customer groups (C-groups) have multicast issues?
- Are there issues for specific customer multicast groups?
- Which customer sources (C-sources) present issues, and for which groups?
- Where are the sources, receivers, and RPs are located?

Methodology

When dealing with a large number of PEs, it is difficult to verify that all of the PEs see each other as PIM neighbors. It is also difficult to verify that the (S,GPA) state for each PE on the Mroute table of the remote PEs is correct.

Therefore, the methodology followed in this guide is to approach the problem incrementally:

- First, focus on checking only two PEs.
- Then, add one PE at a time to these and check the edge settings for this new subset, focusing on the latest added PE.

The concept of adding a single PE (site) is also useful for troubleshooting setups in which an mVPN service is incrementally added for a new site of a customer VPN.

Note that the above described logic implies that, when adding a given PE to the subset, it is understood that all previously checked PEs are at this point fully operational for the mVPN. The diagrams in this section are based on this assumption.

The total number of PEs considered at a given time is noted *n*.

How to Use this Section

Step 1 Choose a subset of two PEs to focus on first.

Step 2 Refer to Checking the First Two PEs, page 2-11 for the procedure used to check the first two PEs.

Step 3 Choose a PE to incrementally add to the PEs in Step 2.

Step 4 Use Figure 2-3 and the related subsections that the diagram references to check the newly added PE with respect to the initial subset.

Step 5 Keep adding incrementally one PE at a time to the subset of PEs and repeat Steps 3 and 4.

Checking the First Two PEs

As previously explained, the first two PEs to be checked can be considered as a specific case.

Perform the following steps (for each step, use the referenced sections only as guidelines though the process, also use the output for those sections as examples for comparison):

Step 1 Use **show ip pim neighbor vrf** command on each router to determine whether they see each other as PIM neighbors through the tunnel.

Step 2 If one or both of them do not see the other, do the following:

a. Check the PIM settings. You can refer to the information in Checking the PIM Neighborship Information on a Specific PE, page 2-15 (see Figure 2-4 and Figure 2-5).

b. Check the status of the MTI and BGP settings. You can leverage the information provided in Checking the MTI Interface Status and BGP Settings on a Specific PE, page 2-23 (note that this section handles cases where *n is* equal to or greater than 3). Look at one PE and then, if needed, check the other.

c. In addition, possibly enable the **debug ip pim vrf hello** command. Note that this should be done carefully in production networks. Check whether they are sending and receiving hellos.

d. Check the health of the Mroute table for the (*,GPA) mstate for both PEs. Possibly use the guidelines in Figure 2-7 (note that this diagram handles cases where *n is* equal to or greater than 3).

e. Re-check the PIM neighbors; if there is still an issue, check the default mdt group settings over the IP multicast core (Section 1: Verifying the IP Multicast Core Settings for the Default MDT, page 2-3) for the PE that has the issue, or for both PEs as needed.

f. If this does not resolve the PIM neighbors problem, check for ACLs and timers on the PEs.

g. After the PEs are both seeing each other as PIM neighbors over the MDT Tunnel, check the health of the Mroute states for the default MDT group (*,GPA) if not done previously. If applicable, also check the (S,GPA). Possibly use Checking the Mroute State Information pertaining to the Default MDTGPA on a Specific PE, page 2-33 (note that this section handles cases where *n is* equal to or greater than 3).

If issues remain after you have thoroughly checked all of the above, contact the Cisco TAC. It is a good practice to be ready with the output of the above to optimize the assistance from TAC.

High Level Diagram of Steps to Follow when Troubleshooting the PEs Edge Settings of an mVPN for a subset of a Number (*n*) of PEs

The high level diagram in Figure 2-3 and the diagrams in Figure 2-4 through Figure 2-8 focus on incrementally adding one PE to the previously checked PEs. This means that it is assumed that the first two PEs have already been checked using the guidelines in Checking the First Two PEs, page 2-11.

On this last PE added, the following checks are performed:

- Local checks pertaining to PIM neighborship settings and Mroute states.
- Check the remote PEs belonging to the subset to ensure PIM neighborship reciprocity with the last added PE.
- Check the (S,GPA) for the new PE on the remote PEs of the subset (if applicable, depending on the default mdt PIM mode).

Figure 2-3 High-Level Diagram for Troubleshooting the PEs Edge Settings of an mVPN

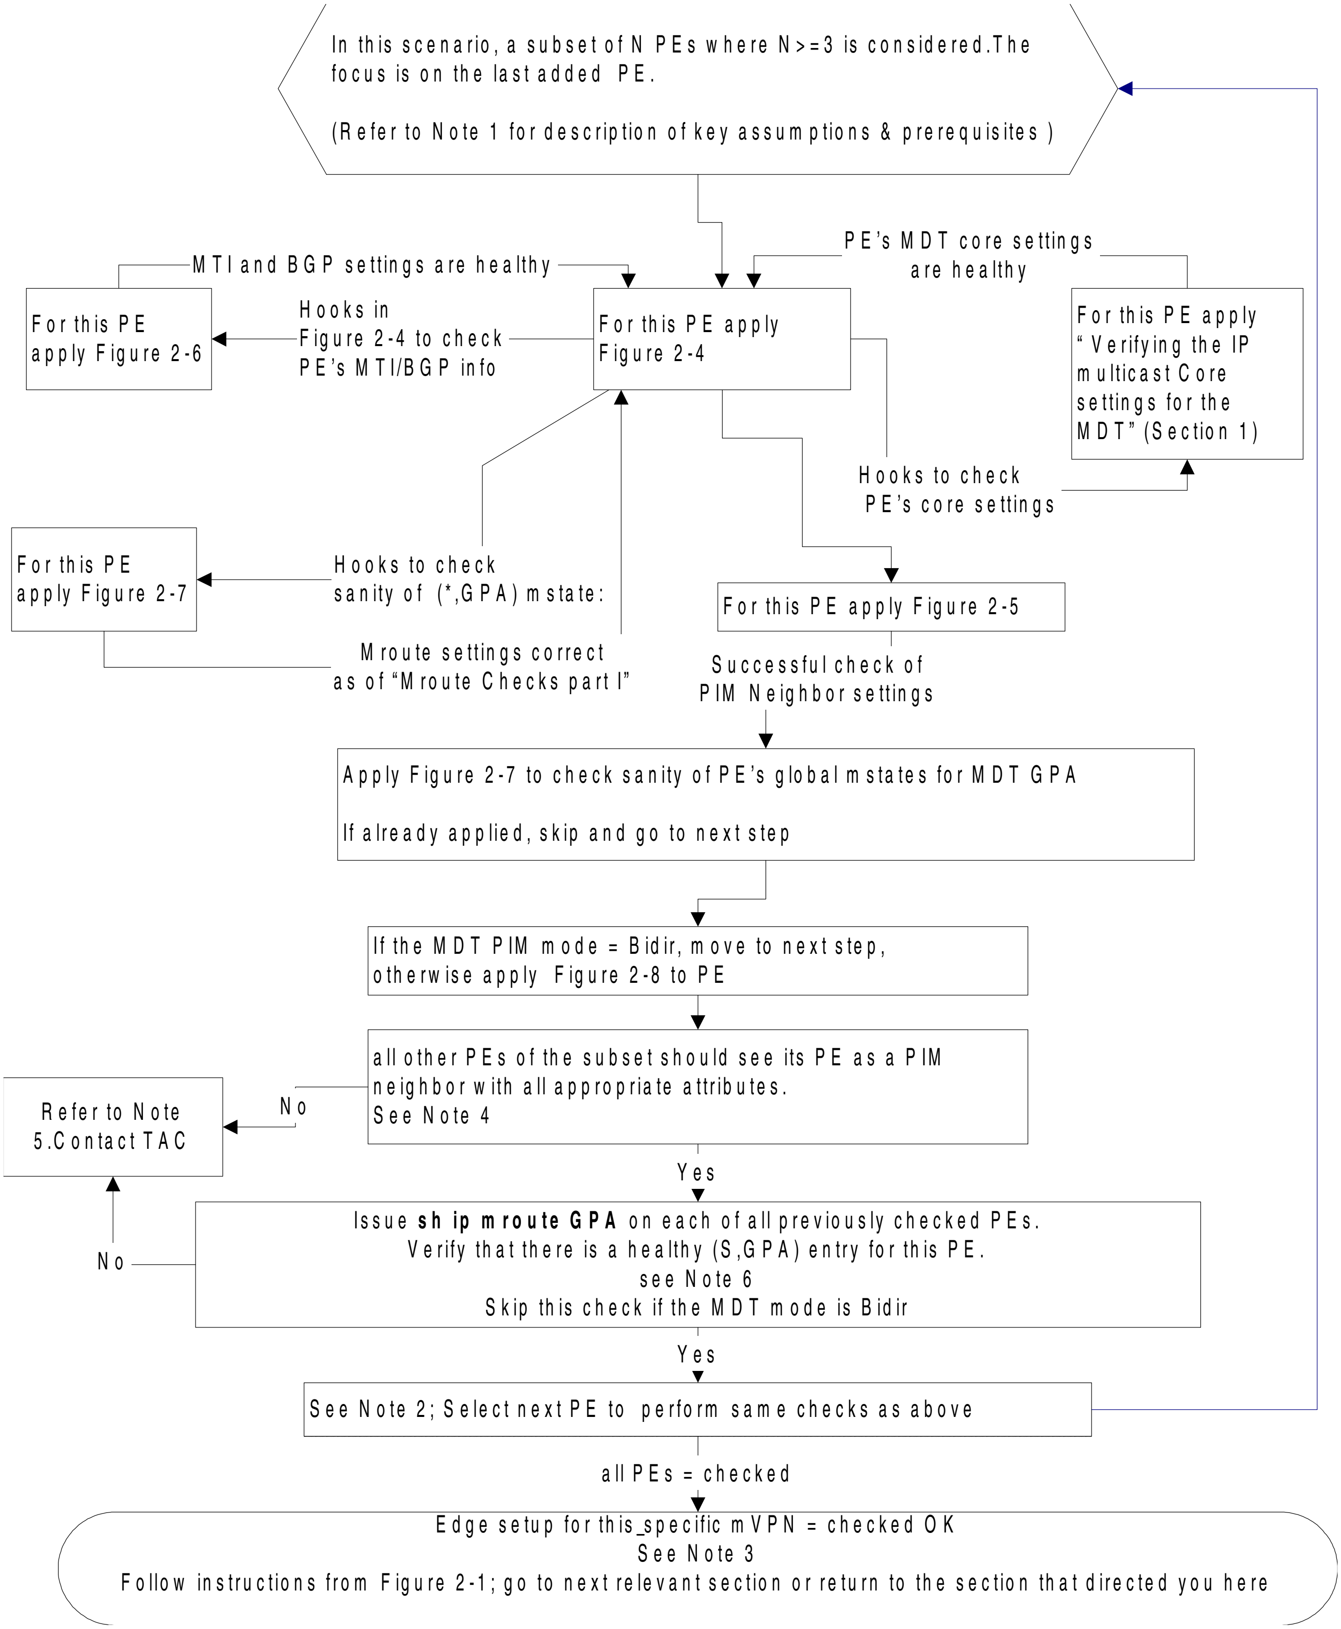
In this scenario, a subset of N PEs where N>=3 is considered.The focus is on the last added PE.
(Refer to Note 1 for description of key assumptions & prerequisites)
MTI and BGP settings are healthy
PE's MDT core settings are healthy
For this PE apply Figure 2-6
Hooks in Figure 2-4 to check PE's MTI/BGP info
For this PE apply Figure 2-4
For this PE apply " Verifying the IP multicast Core settings for the MDT" (Section 1)
Hooks to check PE's core settings
For this PE apply Figure 2-7
Hooks to check sanity of (*,GPA) mstate:
Mroute settings correct as of "Mroute Checks part I"
For this PE apply Figure 2-5
Successful check of PIM Neighbor settings
Apply Figure 2-7 to check sanity of PE's global mstates for MDT GPA
If already applied, skip and go to next step
If the MDT PIM mode = Bidir, move to next step, otherwise apply Figure 2-8 to PE
all other PEs of the subset should see its PE as a PIM neighbor with all appropriate attributes. See Note 4
No
Refer to Note 5.Contact TAC
Yes
Issue **sh ip mroute GPA** on each of all previously checked PEs. Verify that there is a healthy (S,GPA) entry for this PE. see Note 6 Skip this check if the MDT mode is Bidir
No
Yes
See Note 2; Select next PE to perform same checks as above
all PEs = checked
Edge setup for this_specific mVPN = checked OK
See Note 3
Follow instructions from Figure 2-1; go to next relevant section or return to the section that directed you here

Notes for Figure 2-3

Note 1

N = number of PEs in the subset, also known as "PEs for which we are checking the edge settings (PIM neighbors, Mstates)." *N* includes the PE being checked.

Note 2

Check the chosen PE as done for the previously checked PEs. Remember that checking for the PIM neighbors' settings and Mroute table information is a recurring process, adding one PE at a time to the initial subset, until all PEs have been checked.

Note 3

By "Edge setup for this specific mVPN = checked OK," we mean that:

- Each PE of the subset sees all the other PEs as PIM neighbors over the MDT tunnel.
- Each PE of the subset sees all expected Mroutes in its table for the mVPN.

Note 4

- By "all other PEs of the subset," we mean all of the (N-1) PEs that were successfully checked previously.
- Remember that we are applying this diagram in a loop mode, successively for each PE.
- Use the **show ip pim vrf neighbors** command on each remote PE and look for the IP address of the local PE. Verify that each PE is effectively recognized as a PIM neighbor by the current PE. Also check the different PIM neighborship parameters (use guidelines from Figure 2-4 and Figure 2-5 as needed).

Note 5

The described behavior is not normal because the previously checked set of (N-1) PEs was verified as functioning well, and as all settings for the last checked PE are OK.

Note 6

- Verify that the (S,G) entry exists in the other PEs' global Mroute table, where S is this PE's BGP update source address and G is the MDT address (also know as GPA).
- Ensure that the Mroute entry in the tables includes all appropriate parameters (use guidelines from Figure 2-7 as needed).

Checking the PIM Neighborship Information on a Specific PE

This section checks the PIM neighborship between one local PE and a subset of the remote PEs of the MVPN.

The present section is the first section of the PE settings checks to be performed on an **incrementally added PE.** Thus the section diagrams also may link to some other checks such as checks on the MTI and BGP settings (when needed to ensure its health in relation to the PE edge settings for the default mdt).

Ideally, if all local PE parameters had been well configured, the local PE should see all the remote mVPN PEs of the subset as PIM neighbors through the MDT tunnel.

Note Remember that we do incremental checks, thus when checking PIM neighborship on the local PE, only the subset of remote PEs that have previously been checked using this section (Figure 2-4 and Figure 2-5) should be considered, and not all the PEs of the MVPN (this comment is meant to explain references to *expected* neighbors).

Note Remember that you are looking at the PIM neighbor peerings from the view of the *specific* PE being examined. This PE is the new PE incrementally added to the subset. When needed, the flowchart might briefly ask the troubleshooter to log on to other remote PEs to enter some **show** commands, which does not mean that the troubleshooter should check every parameter on these remote PEs (we assume that the remote PEs of the subset have already been checked).

Figure 2-4 and Figure 2-5 show how to troubleshoot PIM neighborship on a given PE newly added to a subset of PEs (that have previously been checked). Note that Figure 2-4 indicates when to use Figure 2-5.

Figure 2-4 Checking PIM Neighbors PART I

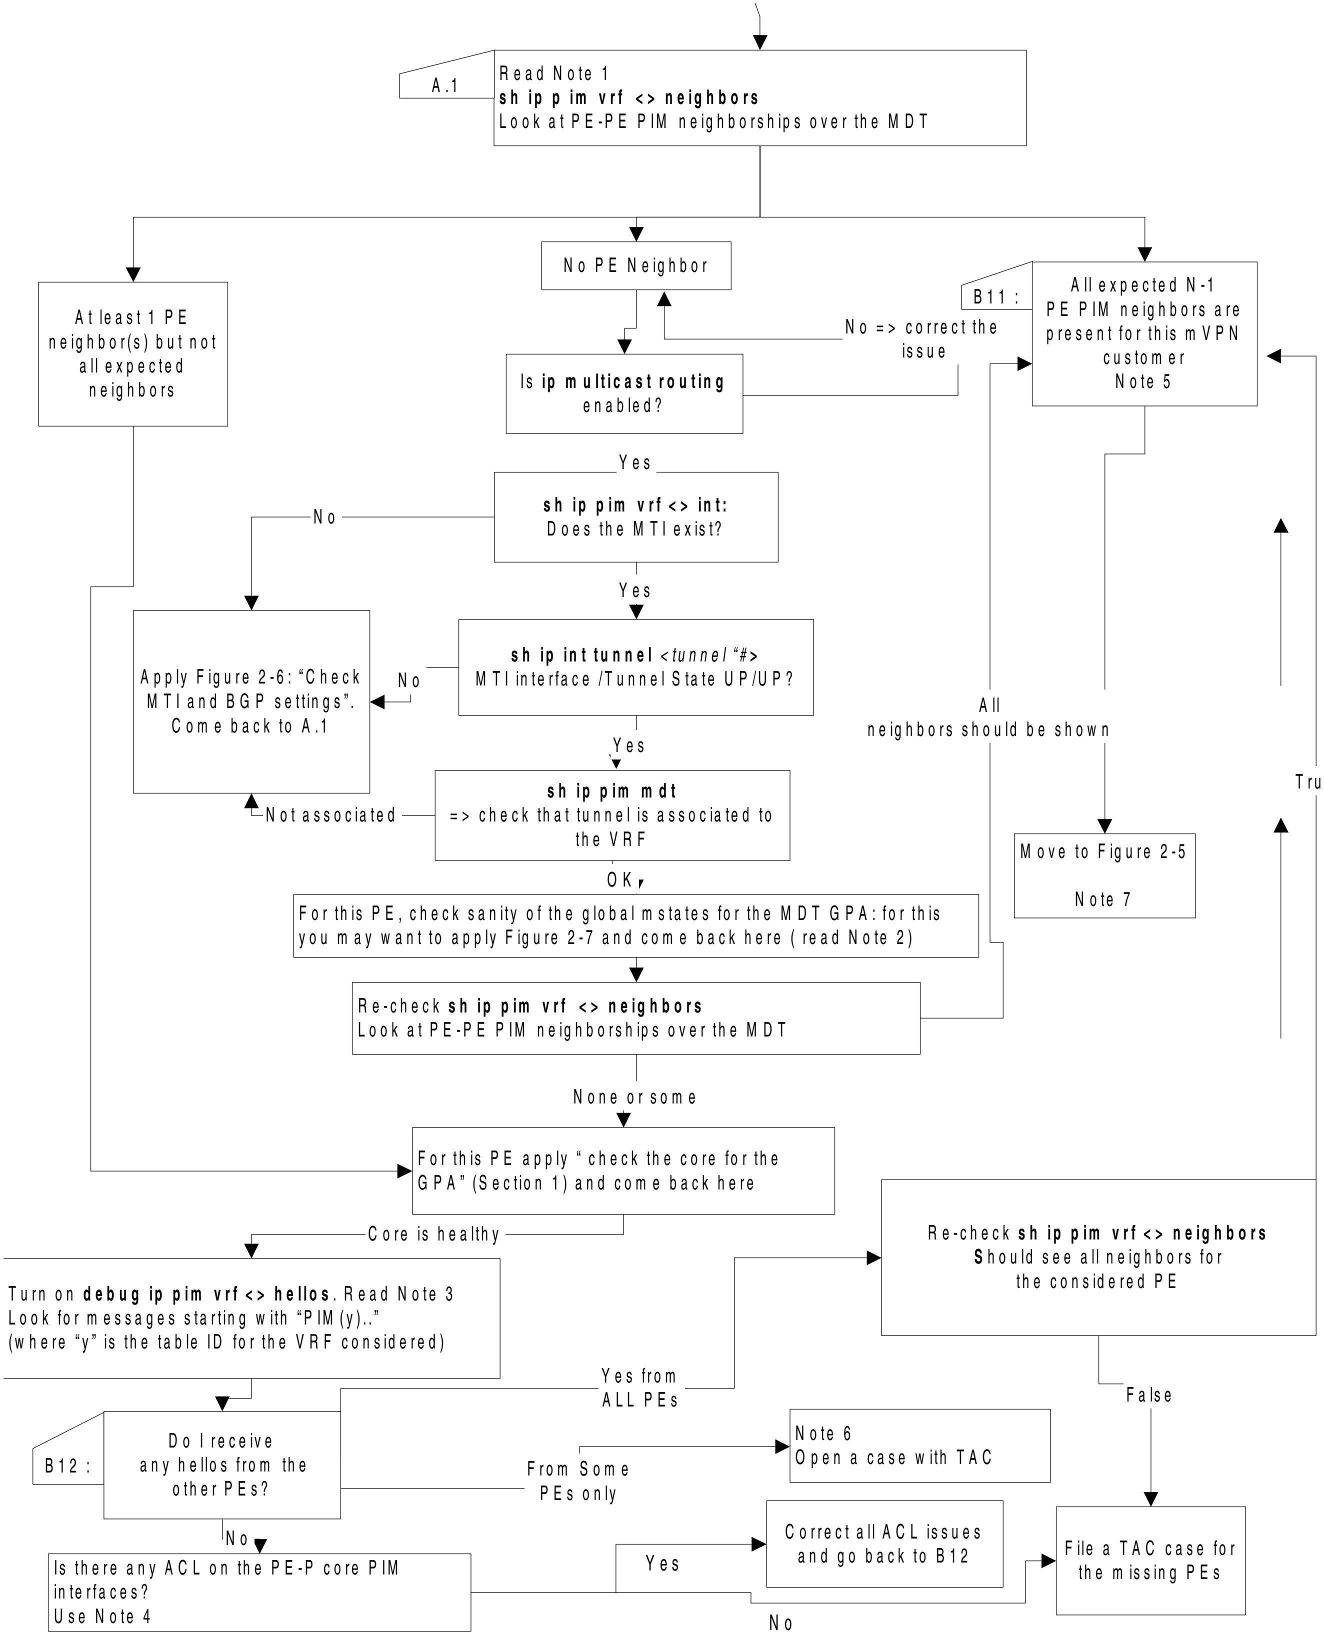
A.1
Read Note 1
sh ip p im vrf <> neighbors
Look at PE-PE PIM neighborships over the MDT
At least 1 PE neighbor(s) but not all expected neighbors
No PE Neighbor
B11 :
All expected N-1 PE PIM neighbors are present for this mVPN customer
Note 5
No => correct the issue
Is ip multicast routing enabled?
Yes
sh ip pim vrf <> int:
Does the MTI exist?
No
Yes
Apply Figure 2-6: "Check MTI and BGP settings". Come back to A.1
No
sh ip int tunnel <tunnel "#>
MTI interface /Tunnel State UP/UP?
All
neighbors should be shown
Yes
sh ip pim mdt
=> check that tunnel is associated to the VRF
Not associated
True
Move to Figure 2-5
Note 7
OK
For this PE, check sanity of the global mstates for the MDT GPA: for this you may want to apply Figure 2-7 and come back here (read Note 2)
Re-check sh ip pim vrf <> neighbors
Look at PE-PE PIM neighborships over the MDT
None or some
For this PE apply " check the core for the GPA" (Section 1) and come back here
Core is healthy
Re-check sh ip pim vrf <> neighbors
Should see all neighbors for the considered PE
Turn on debug ip pim vrf <> hellos. Read Note 3
Look for messages starting with "PIM(y).."
(where "y" is the table ID for the VRF considered)
Yes from ALL PEs
False
B12 :
Do I receive any hellos from the other PEs?
Note 6
Open a case with TAC
From Some PEs only
No
Correct all ACL issues and go back to B12
Yes
File a TAC case for the missing PEs
Is there any ACL on the PE-P core PIM interfaces?
Use Note 4
No

Notes for Figure 2-4

Note 1

- For this section, remember that we consider only the subset of N-1 remote PEs that have previously been checked and we are adding one PE (one being checked currently) to the subset. (we might refer to expected neighbors or considered PEs for the subset of PEs previously checked).
- Remember N>= 3.

Note 2

In Section 1, there are some checks made on the Core Multicast setup for the MDT group address as well.

Note 3

Don’t forget to apply the **term mon command** if you are on a VTY line. Also, use the **show logging** command, being careful to not bring the router down by using a full **debug ip pim debug** command applied on the console.

Note 4

- Use the **show access-list** command and look at the hits on the ACLs, remembering that traffic is encapsulated in GRE.
- Consider ACLs affecting GRE traffic and the MDT GPA address.

Note 5

We are considering N-1 neighbors (because we apply the **show** command on the last PE added to form a subset of N PEs total).

Note 6

Because we previously successfully checked the subset of remote PEs, the local MTI, and local core settings for the MDT, this is not normal.

Note 7

Refer to Figure 2-5 for the second set of checks to be performed on the PIM neighbors for the MTI settings.

Figure 2-5 Checking PIM Neighbors PART II

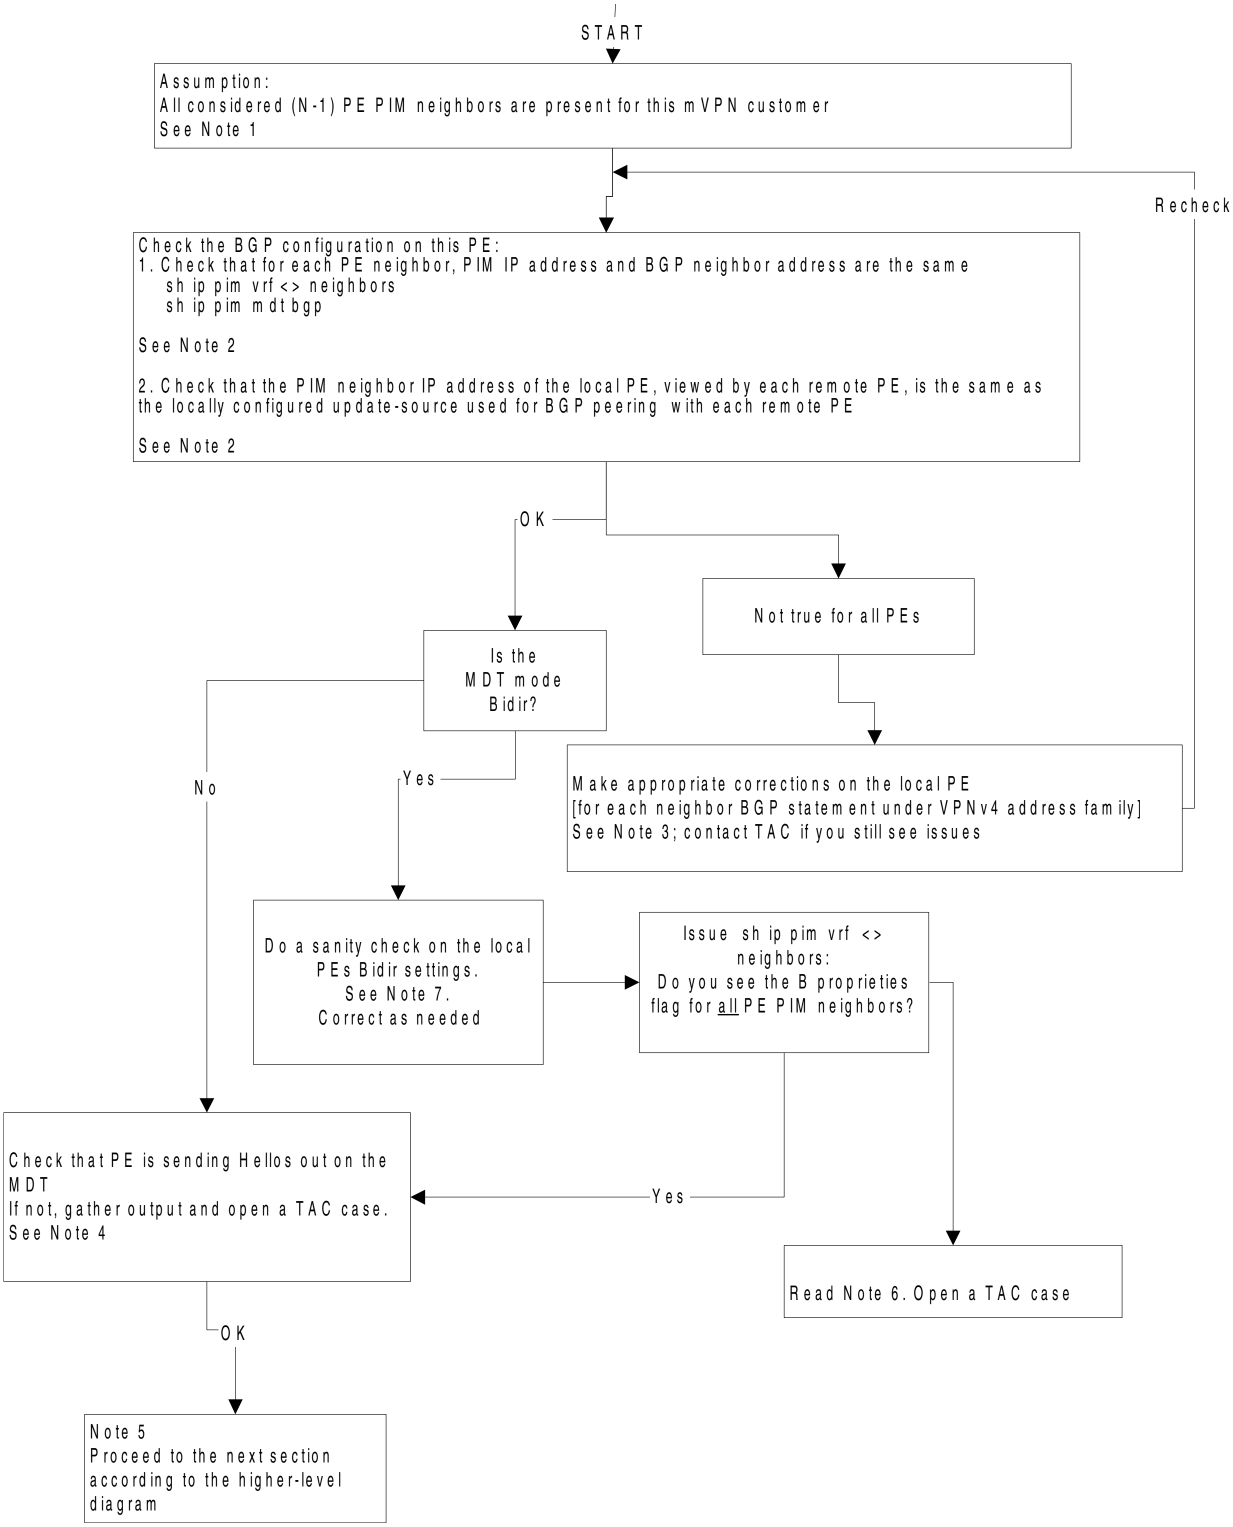
START
Assumption:
All considered (N-1) PE PIM neighbors are present for this mVPN customer
See Note 1
Recheck
Check the BGP configuration on this PE:
1. Check that for each PE neighbor, PIM IP address and BGP neighbor address are the same
sh ip pim vrf <> neighbors
sh ip pim mdt bgp
See Note 2
2. Check that the PIM neighbor IP address of the local PE, viewed by each remote PE, is the same as the locally configured update-source used for BGP peering with each remote PE
See Note 2
OK
Not true for all PEs
Is the MDT mode Bidir?
No
Yes
Make appropriate corrections on the local PE
[for each neighbor BGP statement under VPNv4 address family]
See Note 3; contact TAC if you still see issues
Do a sanity check on the local PEs Bidir settings.
See Note 7.
Correct as needed
Issue sh ip pim vrf <> neighbors:
Do you see the B proprieties flag for all PE PIM neighbors?
Check that PE is sending Hellos out on the MDT
If not, gather output and open a TAC case.
See Note 4
Yes
Read Note 6. Open a TAC case
OK
Note 5
Proceed to the next section according to the higher-level diagram

Notes for Figure 2-5

Note 1

- Figure 2-4 was previously checked for this PE.
- Note that the terminology is the same as defined in Figure 2-5 for checking the PIM neighbors - part I.

Note 2

- Enter the **show ip pim vrf neighbors** command and watch for the IP address of the neighbors. Enter the **show ip pim mdt bgp** command and watch for the next hop information in that output on the remote PEs.

  The two IP addresses representing the local PE in this output should be the same (repeat this on each of the remote PEs).
- Update the source IP address, or the BGP next hop if a next hop was configured for this mVRF.

Note 3

- Keep in mind that BGP peering address changes might restart the BGP peerings.
- Ensure that, on each remote PE, for the PIM neighbor the information is synchronized with the latest BGP update source and BGP next hop address.

Note 4

- You can check this using the **debug ip pim vrf hellos** command.
- This should be the case because the MTI interface and tunnel state is UP/UP and we already checked timers.

Note 5

The PE is healthy with respect to the PIM neighbor settings over the MTI. A higher level diagram would be Figure 2-3. It could be, however, that other sections of this document leverage Figure 2-5. In that case, as you have just successfully completed the checks in this flowchart, return to the figure that referred you to this one.

Note 6

Remember that we have previously checked all (N-1) PEs of the subset.

Note 7

- Check for the RP and the GPA BIDIR settings (check for Bidir globally enabled, RP is bidir enabled) using the following commands:

  show ip pim rp mapping [**in-use**]

  show ip mroute
- When finished, skip to the next step.

Sample Output to Compare your Setup while Leveraging Figures 4 and 5

Use the following output as a reference when using the flowchart diagram. Compare your output to the sample output taken on an appropriate mVPN Bidir and PIM SM setup.

The following shows the output of **show ip pim neighbors vrf RED** command on a correct setup with Bidir as PIM mode in the core for the default MDT, and for PIM SM with SPT 0:

```
PE1# show ip pim vrf RED neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.11.2      Serial1/0                00:08:11/00:01:27 v2    1 / P
192.168.112.2     Serial2/0                00:08:11/00:01:27 v2    1 / P
172.16.100.3      Tunnel0                  00:07:09/00:01:30 v2    1 / DR B P
172.16.100.2      Tunnel0                  00:07:09/00:01:30 v2    1 / B P
```

```
PE1#show ip pim vrf RED neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.11.2      Serial1/0                2w2d/00:01:40     v2    1 / P
192.168.112.2     Serial2/0                2w2d/00:01:41     v2    1 / P
172.16.100.2      Tunnel0                  2w2d/00:01:21     v2    1 / P
172.16.100.3      Tunnel0                  2w2d/00:01:43     v2    1 / DR P
```

The following output of the **show ip mroute** *GPA* **summary** command is taken for Bidir and PIM SM with SPT 0 respectively. Note that in the Bidir case, there is a (*,GPA) state but no (S,GPA) states.

```
PE1#show ip mroute 239.0.0.10 summary
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 00:31:03/00:02:59, RP 172.16.100.200, OIF count: 2, flags: BCZ
```

Note that in the PIM SM case, the (S,GPA) corresponding to S =Local PE does not present a Z flag; this flag is used on the MRIB to indicate a need to decapsulate an incoming packet.

```
PE1#show ip mroute 239.0.0.10 summary
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 2w2d/stopped, RP 172.16.100.200, OIF count: 1, flags: SJCFZ
  (172.16.100.1, 239.0.0.10), 2w2d/00:03:29, OIF count: 1, flags: FT
  (172.16.100.2, 239.0.0.10), 2w2d/00:02:47, OIF count: 1, flags: JTZ
  (172.16.100.3, 239.0.0.10), 2w2d/00:02:55, OIF count: 1, flags: JTZ
```

Following is a sample of the syslog messages that should be seen in the buffer (on the console) sent to a syslog server at startup of a PE if the configuration for the mVPN has been properly added, or when the appropriate commands are added on a PE:

```
PE1#
*Oct  4 01:31:58.151: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed
state to up
*Oct  4 01:31:58.151: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed
state to up
*Oct  4 01:31:58.299: %PIM-5-NBRCHG: neighbor 172.16.100.2 UP on interface Tunnel0 (vrf
RED)
*Oct  4 01:31:58.299: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.2 on
interface Tunnel0 (vrf RED)
*Oct  4 01:31:58.355: %PIM-5-NBRCHG: neighbor 172.16.100.3 UP on interface Tunnel0 (vrf
RED)
*Oct  4 01:31:58.355: %PIM-5-DRCHG: DR change from neighbor 172.16.100.2 to 172.16.100.3
on interface Tunnel0 (vrf RED)
```

Following is a sample output with debugs on for IP PIM hellos in the VRF RED context. Use caution when turning on debugs so as not to bring the device down. Note that the value *x* in the output in such lines as PIM*(x)* identifies the global table (0) or the VRF table ID (here table ID equals 1 for VRF RED).

```
PE1#show debug
IP multicast:
  PIM HELLO debugging is on in vrf RED
PE1#
*Oct  4 02:07:51.779: PIM(0): Send periodic v2 Hello on Serial0/0
*Oct  4 02:07:52.587: PIM(0): Received v2 hello on Serial0/0 from 172.16.1.11
*Oct  4 02:07:55.035: PIM(1): Send periodic v2 Hello on Tunnel0
*Oct  4 02:07:55.095: PIM(0): Send periodic v2 Hello on Loopback0
*Oct  4 02:07:55.127: PIM(0): Received v2 hello on Loopback0 from 172.16.100.1
*Oct  4 02:07:55.375: PIM(1): Received v2 hello on Tunnel0 from 172.16.100.3
*Oct  4 02:07:55.559: PIM(1): Received v2 hello on Serial1/0 from 192.168.11.2
*Oct  4 02:07:57.259: PIM(1): Send periodic v2 Hello on Serial2/0
*Oct  4 02:07:57.691: PIM(1): Received v2 hello on Serial2/0 from 192.168.112.2
*Oct  4 02:07:59.683: PIM(1): Received v2 hello on Tunnel0 from 172.16.100.2
*Oct  4 02:08:02.051: PIM(1): Send periodic v2 Hello on Serial1/0
```

Use the **show ip pim** [**vrf RED**] *interface* **count** command successively to observe the increasing number of packets in or out for a given interface; that is, for the MTI/tunnel interface when using the VRF option. Note that in the absence of customer multicast traffic, the hellos sent over the MTI travel on the default mdt and as such, increment the counters for the tunnel/MTI interface.

Note also that such output allows you to check whether the Fast Switching capability is turned on for a specific interface.

```
PE1#show ip pim interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
172.16.100.1     Loopback0                *   141/0
172.16.1.1       Serial0/0                *   234/134
```

Note It is normal for the number of Outgoing packets to be null on the Loopback0 interface.

```
PE1#sh ip pim interface count

State: * - Fast Switched, D - Distributed Fast Switched
```

```
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
172.16.100.1     Loopback0                *   142/0
172.16.1.1       Serial0/0                *   234/135#

PE1#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
192.168.11.1     Serial1/0                *   0/48
192.168.112.1    Serial2/0                *   0/48
172.16.100.1     Tunnel0                  *   48/0

PE1#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
192.168.11.1     Serial1/0                *   0/49
192.168.112.1    Serial2/0                *   0/49
172.16.100.1     Tunnel0                  *   49/0
```

The following output shows the remote PE next hop information. It can be used to compare with the IP addresses of the PE PIM neighbors for their neighborship with the local PE over the MTI. Note that the output would be the same for Bidir and PIM SM with SPT equals to 0.

```
PE1#show ip pim mdt bgp
MDT (Route Distinguisher + IPv4)               Router ID         Next Hop
  MDT group 239.0.0.10
   2:100:172.16.100.2                          172.16.100.100    172.16.100.2
   3:100:172.16.100.3                          172.16.100.100    172.16.100.3
```

The following is an example of the impact of an ACL placed on a core interface and filtering GRE traffic:

```
PE2# show access-l 101
Extended IP access list 101
    deny gre any host 239.0.0.10
    permit ip any any

PE2#conf t
PEC-RP#(config)#int s0/0
PEC-RP#(config-if)#ip access-group 101 in
PEC-RP#(config-if)#end

PE2#debug ip pim hello
PIM-HELLO debugging is on

PE2#show ip pim  vrf RED  nei
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.22.2      Serial1/0                6d15h/00:01:32    v2    1 / P
172.16.100.3      Tunnel0                  6d12h/00:00:02    v2    1 / DR B P


PE2#
*Sep 15 13:28:03.984: %PIM-5-NBRCHG: neighbor 172.16.100.1 DOWN on interface Tunnel0 (vrf
RED) non DR
******************************************************************************************
******************************************************************************************
```

```
*Sep 15 13:28:07.080: %PIM-5-NBRCHG: neighbor 172.16.100.3 DOWN on interface Tunnel0 (vrf
RED) DR
*Sep 15 13:28:07.080: %PIM-5-DRCHG: DR change from neighbor 172.16.100.3 to 172.16.100.2
on interface Tunnel0 (vrf RED)
*Sep 15 13:28:11.368: PIM(0): Send periodic v2 Hello on Loopback0
*Sep 15 13:28:11.420: PIM(0): Received v2 hello on Loopback0 from 172.16.100.2


PE2#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.22.1     Serial1/0                v2/S   1      30     1      0.0.0.0
172.16.100.2     Tunnel0                  v2/SD  0      30     1      172.16.100.2

*Sep 15 13:28:40.720: PIM(0): Send periodic v2 Hello on Loopback0
*Sep 15 13:28:40.740: PIM(0): Received v2 hello on Loopback0 from 172.16.100.2


*Sep 15 13:28:48.268: PIM(0): Send periodic v2 Hello on Serial0/0

PE2#show ip pim vrf RED neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.22.2      Serial1/0                6d15h/00:01:44    v2    1 / P


PE2#
*Sep 15 13:29:40.148: PIM(0): Send periodic v2 Hello on Loopback0
*Sep 15 13:29:40.176: PIM(0): Received v2 hello on Loopback0 from 172.16.100.2
```

Note These are the locally sent hellos.

```
*Sep 15 13:29:47.500: PIM(0): Send periodic v2 Hello on Serial0/0
*Sep 15 13:29:51.084: PIM(0): Received v2 hello on Serial0/0 from 172.16.2.12
*Sep 15 13:30:09.416: PIM(0): Send periodic v2 Hello on Loopback0
*Sep 15 13:30:09.448: PIM(0): Received v2 hello on Loopback0 from 172.16.100.2
*Sep 15 13:30:17.036: PIM(0): Send periodic v2 Hello on Serial0/0
*Sep 15 13:30:20.876: PIM(0): Received v2 hello on Serial0/0 from 172.16.2.12
*Sep 15 13:30:38.876: PIM(0): Send periodic v2 Hello on Loopback0
*Sep 15 13:30:38.896: PIM(0): Received v2 hello on Loopback0 from 172.16.100.2#
```

Checking the MTI Interface Status and BGP Settings on a Specific PE

Cisco IOS creates an MTI interface for each mVRF customer configured on the PE. It is used to build PE to PE PIM adjacencies and multicast forwarding states for customers. The multicast tunnel interfaces are dynamic interfaces. They are not configurable from CLI, nor are they visible to any unicast routing protocols; they are present only in the corresponding MVRF.

The key items required for MTI to be associated to an mVRF are as follows:

- The **mdt default** *mdt-GPA* command is configured under the VRF definition.
- Appropriate tunnel source loopback interface is selected for the MTI.
- The tunnel source loopback interface is UP/UP and is PIM enabled.
- BGP peers are well configured for the MDT SAFI and should be activated.

The MTI interface has the following characteristics:

Tunnel source loopback interface	Assigned by BGP for each mVRF. Check "How does MTI picks the tunnel source interface?" for details
MTI tunnel destination IP	default mdt group address configured by the operator
MTI IP address	unnumbered to BGP-selected Tunnel source interface
MTI encapsulation	GRE

If the **bgp-next hop** *interface* command is configured in the VRF definition, then use that interface as the tunnel source interface for that VRF.

If the **neighbor** *ip-address* **update-source** *interface* command exists for a peer, the interface specified is used as the tunnel source. Here is the order of peers from which the update-source is determined:

1. The **ibgp neighbor** *ip address* **activate** (numerically lowest IP address) command configured in the address-family IPv4 MDT (even if the peer is not up).
2. The **ibgp neighbor** *ip address* **activate** (numerically lowest IP address) configured in the address-family IPv4 vpnv4 (even if the peer is not up).
3. The **eBGP neighbor** *ip address* (numerically lowest ip-address) command configured in the address-family IPv4 MDT (even if the peer is not up).

The flowchart in Figure 2-6 helps isolate MTI-related issues.

Figure 2-6 Checking the MTI Interface Status and BGP Settings on a Specific PE

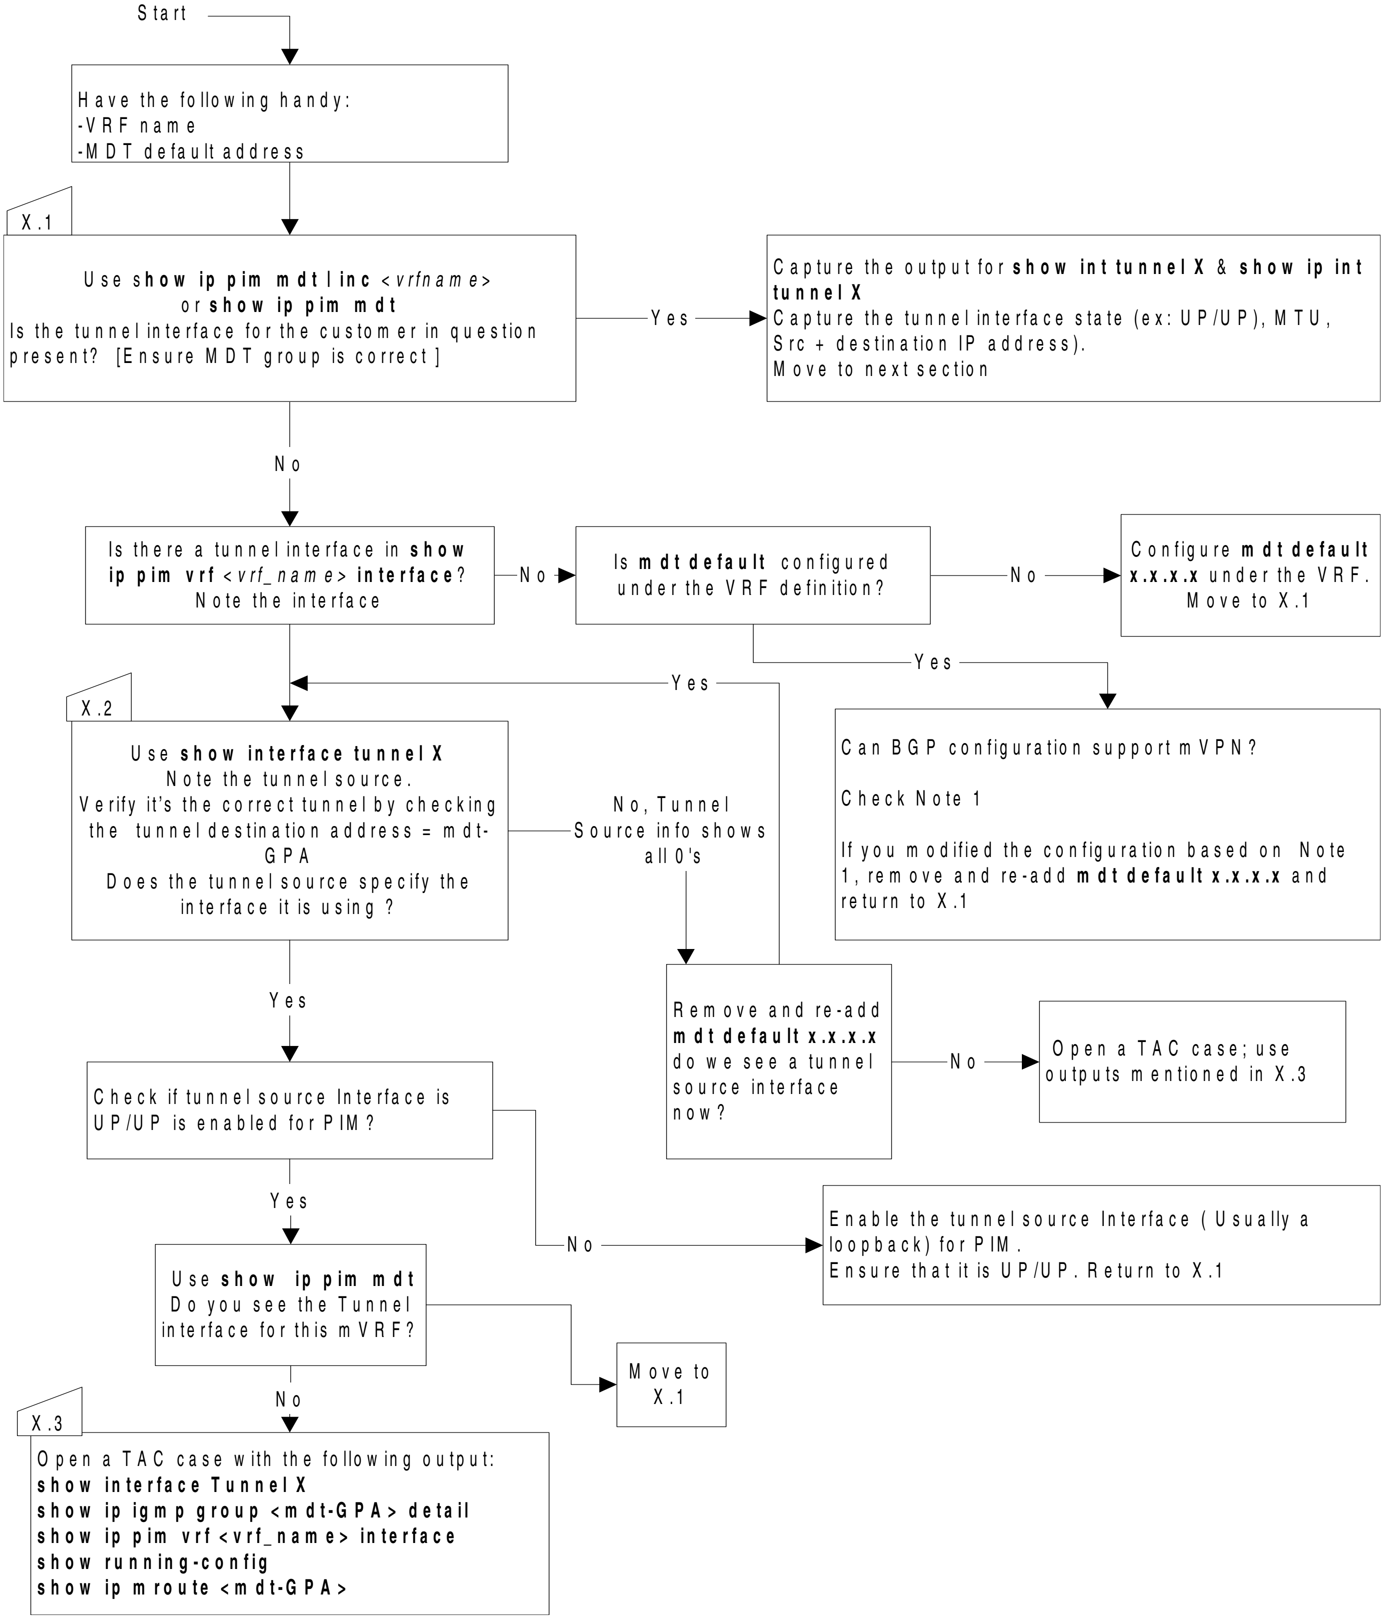
Start
Have the following handy:
-VRF name
-MDT default address
X.1
Use **show ip pim mdt | inc** <*vrfname*> or **show ip pim mdt**
Is the tunnel interface for the customer in question present? [Ensure MDT group is correct]
Yes
Capture the output for **show int tunnel X** & **show ip int tunnel X**
Capture the tunnel interface state (ex: UP/UP), MTU, Src + destination IP address).
Move to next section
No
Is there a tunnel interface in **show ip pim vrf** <*vrf_name*> **interface**? Note the interface
No
Is **mdt default** configured under the VRF definition?
No
Configure **mdt default x.x.x.x** under the VRF. Move to X.1
Yes
Can BGP configuration support mVPN?
Check Note 1
If you modified the configuration based on Note 1, remove and re-add **mdt default x.x.x.x** and return to X.1
Yes
X.2
Use **show interface tunnel X**
Note the tunnel source.
Verify it's the correct tunnel by checking the tunnel destination address = mdt-GPA
Does the tunnel source specify the interface it is using ?
No, Tunnel Source info shows all 0's
Remove and re-add **mdt default x.x.x.x** do we see a tunnel source interface now?
No
Open a TAC case; use outputs mentioned in X.3
Yes
Check if tunnel source Interface is UP/UP is enabled for PIM?
No
Enable the tunnel source Interface (Usually a loopback) for PIM.
Ensure that it is UP/UP. Return to X.1
Yes
Use **show ip pim mdt**
Do you see the Tunnel interface for this mVRF?
Move to X.1
No
X.3
Open a TAC case with the following output:
show interface Tunnel X
show ip igmp group <mdt-GPA> detail
show ip pim vrf <vrf_name> interface
show running-config
show ip mroute <mdt-GPA>

Notes for Figure 2-6

Note 1

- We need to have **neighbor** *ip_addr* **activate** and **neighbor** *ip_addr* **send community both** under the vpnv4 AFI and **neighbor** *ip_addr* **activate** under MDT SAFI.
- Remember that if the PE is running an IOS later than 12.0(29)S that supports MDT SAFI, then it is essential to have the same neighbor and update source parameters in the vpnv4 configurations and in the MDT SAFI for the mVPN tunnel to come up.
- Ensure that you configure the MDT SAFI in BGP (even though we are using Intra AS mVPN). Refer to

  http://www.cisco.com/en/US/products/sw/iosswrel/ps1829/products_feature_guide09186a00802a5bd9.html

Sample Output for Figure 6

The **show ip pim mdt** command is used to verify that the MTI is mapped to the correct mVRF.

The **show ip igmp group** *mdt-GPA* detail can be used to check if the local PE is the leaf and root for the default mdt GPA.

The **show interface tunnel** *X* can be used to check the status, tunnel source loopback, and tunnel destination for the MTI.

Note One of the most common issues encountered when troubleshooting mVPN problems is that the loopback interface that gives its properties to the MTI is not enabled for PIM.

Following is the output from a router for which the MTI and BGP are correctly configured:

```
PE1#show ip pim mdt
  MDT Group        Interface    Source                     VRF
* 239.0.0.10       Tunnel16     Loopback0                  RED
```

In the above output, we can see that the Tunnel interface exists, has an IP address (unnumbered from Loopback0), and was associated to the vrf RED.

```
PE1#show int tunnel 16
Tunnel16 is up, line protocol is up
  Hardware is Tunnel
  Interface is unnumbered.  Using address of Loopback0 (172.16.100.1)
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 172.16.100.1 (Loopback0), destination 239.0.0.10
  Tunnel protocol/transport GRE/IP Multicast, sequencing disabled
  Tunnel TTL 255
  Key disabled
  Checksumming of packets disabled,  fast tunneling enabled
  Last input 00:00:06, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 8000 bits/sec, 1 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     4918 packets input, 4420606 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     196 packets output, 13914 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out

PE1#show ip igmp groups 239.0.0.10 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source

Interface:      Loopback0
Group:          239.0.0.10
Flags:          VG
Uptime:         00:30:55
Group mode:     INCLUDE
Last reporter:  0.0.0.0
Source list is empty

PE1#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.112.1    Serial2/0                v2/S   1      30     1      0.0.0.0
192.168.11.1     Serial1/0                v2/S   1      30     1      0.0.0.0
172.16.100.1     Tunnel16                 v2/SD  2      30     1      172.16.100.3
```

Failure Scenario 1: PIM Disabled on the MTI

In this failure scenario, PIM is disabled on the tunnel source loopback interface:

Step 1 Disabling PIM on the tunnel source loopback interface (loopback 0) in this example.

```
PE1(config)#interface loopback 0
PE1(config-if)#no ip pim
(config-if)#end
PE1#
*Oct 28 05:01:22.028: %SYS-5-CONFIG_I: Configured from console by console
*Oct 28 05:01:25.836: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel16, changed
state to down
*Oct 28 05:01:25.836: %PIM-5-NBRCHG: neighbor 172.16.100.3 DOWN on interface Tunnel16 (vrf
RED) non DR
*Oct 28 05:01:25.836: %PIM-5-NBRCHG: neighbor 172.16.100.2 DOWN on interface Tunnel16 (vrf
RED) non DR
```

Step 2 Let us check the output described above:

```
PE1#show ip pim mdt
  MDT Group       Interface   Source                   VRF
PE1#

PE1c#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.112.1    Serial2/0                v2/S   1      30     1      0.0.0.0
192.168.11.1     Serial1/0                v2/S   1      30     1      0.0.0.0
0.0.0.0          Tunnel16                 v2/SD  0      30     1      0.0.0.0
```

```
PE1#show ip igmp groups 239.0.0.10 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source

PE1#show interface tunnel 16
Tunnel16 is up, line protocol is down
  Hardware is Tunnel
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 0.0.0.0, destination 239.0.0.10
  Tunnel protocol/transport GRE/IP Multicast, sequencing disabled
  Tunnel TTL 255
  Key disabled
  Checksumming of packets disabled,  fast tunneling enabled
  Last input 00:01:31, output 00:01:48, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     5194 packets input, 4666110 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     214 packets output, 15164 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out
```

From the above output it seems that the tunnel interface Up/Down is the cause of the issue.

Step 3 We bring the tunnel interface Down and then Up, without enabling PIM on the loopback interface, as shown in the following example:

```
PE1#ip vrf RED
PE1(config-vrf)#no mdt default
% A new tunnel ID may be used if the default mdt is reconfigured for this VRF.
Tunnel interface was deleted. Partial configuration may reappear on reuse.
PE1(config-vrf)#mdt default 239.0.0.10
*Oct 28 05:07:30.700: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel17, changed
state to up
*Oct 28 05:07:30.700: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel17, changed
state to up
*Oct 28 05:07:31.728: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.1 on
interface Tunnel17 (vrf RED)
```

Step 4 Capturing the output, as shown in the following:

```
PEC-SRC#show ip pim mdt
  MDT Group       Interface   Source                   VRF

PEC-SRC#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.112.1    Serial2/0                v2/S   1      30     1      0.0.0.0
192.168.11.1     Serial1/0                v2/S   1      30     1      0.0.0.0
172.16.100.1     Tunnel17                 v2/SD  0      30     1      172.16.100.1


PE1#show interface tunnel 17
```

```
Tunnel17 is up, line protocol is up
  Hardware is Tunnel
  Interface is unnumbered.  Using address of Loopback0 (172.16.100.1)
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 172.16.100.1 (Loopback0), destination 239.0.0.10
  Tunnel protocol/transport GRE/IP Multicast, sequencing disabled
  Tunnel TTL 255
  Key disabled
  Checksumming of packets disabled,  fast tunneling enabled
  Last input never, output 00:00:17, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     5 packets output, 326 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out


PE1#show ip igmp groups 239.0.0.10 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source
```

Although the tunnel interface is UP/UP, it is **not associated to the VRF**

Note In order to verify that the MTI is associated to the VRF, do not get confused by the output of the **show ip pim vrf RED interface** command. The actual association is shown only by using the **show ip pim mdt** command.

Step 5 We resolve the issue by enabling PIM on the BGP update source interface, from which the MTI takes its properties, as consequence the MTI will come up will all its necessary properties and the PIM neighborships will be created over the MDT tunnel.

```
PE1(config)#interface loopback 0
PE1(config-if)#ip pim sparse-mode
PE1(config-if)#end

*Oct 28 05:10:54.088: %PIM-5-NBRCHG: neighbor 172.16.100.2 UP on interface Tunnel17 (vrf 
RED)
*Oct 28 05:10:54.088: %PIM-5-DRCHG: DR change from neighbor 172.16.100.1 to 172.16.100.2 
on interface Tunnel17 (vrf RED)
*Oct 28 05:10:54.148: %PIM-5-NBRCHG: neighbor 172.16.100.3 UP on interface Tunnel17 (vrf 
RED)
*Oct 28 05:10:54.148: %PIM-5-DRCHG: DR change from neighbor 172.16.100.2 to 172.16.100.3 
on interface Tunnel17 (vrf RED)
*Oct 28 05:10:55.588: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.1 on 
interface Tunnel9 (vrf 2)
```

Failure Scenario 2: BGP Update Source Interface Shutdown

In this scenario, the tunnel source (loopback 0) interface is shut down:

Step 1 Disable the PIM on the tunnel source interface (loopback 0) in the following example:

```
PE1(config)#interface loopback 0
PE1(config-if)#no ip pim
PE1(config-if)#end
PE1#
*Oct 28 05:01:22.028: %SYS-5-CONFIG_I: Configured from console by console
*Oct 28 05:01:25.836: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel16, changed
state to down
*Oct 28 05:01:25.836: %PIM-5-NBRCHG: neighbor 172.16.100.3 DOWN on interface Tunnel16 (vrf
RED) non DR
*Oct 28 05:01:25.836: %PIM-5-NBRCHG: neighbor 172.16.100.2 DOWN on interface Tunnel16 (vrf
RED) non DR

PE1(config)#int loopback 0
PE1(config-if)#shut
*Oct 28 05:13:08.788: %LINK-5-CHANGED: Interface Loopback0, changed state to
administratively down
*Oct 28 05:13:09.808: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed
state to down
*Oct 28 05:13:12.828: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel17, changed
state to down
*Oct 28 05:13:12.828: %PIM-5-NBRCHG: neighbor 172.16.100.3 DOWN on interface Tunnel17 (vrf
RED) DR
*Oct 28 05:13:12.828: %PIM-5-NBRCHG: neighbor 172.16.100.2 DOWN on interface Tunnel17 (vrf
RED) non DR
```

Step 2 Capture output, as shown in the following example:

```
PE1#show ip pim mdt
  MDT Group       Interface   Source                   VRF#

PE1#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.112.1    Serial2/0                v2/S   1      30     1      0.0.0.0
192.168.11.1     Serial1/0                v2/S   1      30     1      0.0.0.0
172.16.100.1     Tunnel17                 v2/SD  0      30     1      0.0.0.0


PE1#show interface tunnel 17
Tunnel17 is up, line protocol is down
  Hardware is Tunnel
  Interface is unnumbered.  Using address of Loopback0 (172.16.100.1)
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 172.16.100.1 (Loopback0), destination 239.0.0.10
  Tunnel protocol/transport GRE/IP Multicast, sequencing disabled
  Tunnel TTL 255
  Key disabled
  Checksumming of packets disabled,  fast tunneling enabled
  Last input 00:01:36, output 00:01:49, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/0 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     161 packets input, 143012 bytes, 0 no buffer
     Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
     20 packets output, 1360 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
     0 output buffer failures, 0 output buffers swapped out

PE1#show ip igmp groups 239.0.0.10 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source
```

Step 3 Bring the tunnel interface form UP/DOWN to UP/UP by issuing removing and re-adding mdt-default:

```
PE1(config)#ip vrf RED
PE1(config-vrf)#no mdt default
% A new tunnel ID may be used if the default mdt is reconfigured for this VRF.
Tunnel interface was deleted. Partial configuration may reappear on reuse.
PE1(config-vrf)#mdt default 239.0.0.10
*Oct 28 05:16:16.748: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel18, changed
state to up
*Oct 28 05:16:16.748: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel18, changed
state to down

PE1#show ip pim mdt
  MDT Group       Interface   Source                   VRF
* 239.0.0.10      Tunnel18    Loopback0                RED

PE1#show running-config interface loopback 0
Building configuration...

Current configuration : 174 bytes
!
interface Loopback0
 ip address 172.16.100.1 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-mode
 ip igmp join-group 239.0.0.100
 load-interval 30
 shutdown
end

PE1#show ip igmp groups 239.0.0.10 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source

PE1#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.112.1    Serial2/0                v2/S   1      30     1      0.0.0.0
192.168.11.1     Serial1/0                v2/S   1      30     1      0.0.0.0
172.16.100.1     Tunnel18                 v2/SD  0      30     1      0.0.0.0#
```

Step 4 Resolving the issue by bring up the loopback interface

```
PE$_{C-SRC}$#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PE$_{C-SRC}$(config)#int loopback 0
```

```
PEC-SRC(config-if)#no shut
PEC-SRC(config-if)#end

*Oct 28 05:18:26.420: %SYS-5-CONFIG_I: Configured from console by console
*Oct 28 05:18:26.660: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.1 on
interface Tunnel18 (vrf RED)
*Oct 28 05:18:27.680: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.1 on
interface Loopback0 (vrf default)
*Oct 28 05:18:27.832: %LINK-3-UPDOWN: Interface Loopback0, changed state to up
*Oct 28 05:18:28.072: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel18, changed
state to up
*Oct 28 05:18:28.852: %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed
state to up

*Oct 28 05:18:47.480: %PIM-5-NBRCHG: neighbor 172.16.100.2 UP on interface Tunnel18 (vrf
RED)
*Oct 28 05:18:47.480: %PIM-5-DRCHG: DR change from neighbor 172.16.100.1 to 172.16.100.2
on interface Tunnel18 (vrf RED)
*Oct 28 05:18:49.272: %PIM-5-NBRCHG: neighbor 172.16.100.3 UP on interface Tunnel18 (vrf
RED)
*Oct 28 05:18:49.272: %PIM-5-DRCHG: DR change from neighbor 172.16.100.2 to 172.16.100.3
on interface Tunnel18 (vrf RED)

PE1#show ip pim mdt
  MDT Group       Interface   Source                   VRF
* 239.0.0.10      Tunnel18    Loopback0                RED

PE1#show ip pim vrf RED interface

Address          Interface                Ver/   Nbr    Query  DR     DR
                                          Mode   Count  Intvl  Prior
192.168.112.1    Serial2/0                v2/S   1      30     1      0.0.0.0
192.168.11.1     Serial1/0                v2/S   1      30     1      0.0.0.0
172.16.100.1     Tunnel18                 v2/SD  2      30     1      172.16.100.3

PE1#show ip igmp groups 239.0.0.10 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source

Interface:      Loopback0
Group:          239.0.0.10
Flags:          VG
Uptime:         00:03:29
Group mode:     INCLUDE
Last reporter:  0.0.0.0
Source list is empty


PE1#show int tunnel 18
Tunnel18 is up, line protocol is up
  Hardware is Tunnel
  Interface is unnumbered.  Using address of Loopback0 (172.16.100.1)
  MTU 1514 bytes, BW 9 Kbit, DLY 500000 usec, rely 255/255, load 1/255
  Encapsulation TUNNEL, loopback not set
  Keepalive not set
  Tunnel source 172.16.100.1 (Loopback0), destination 239.0.0.10
  Tunnel protocol/transport GRE/IP Multicast, sequencing disabled
  Tunnel TTL 255
  Key disabled
  Checksumming of packets disabled,  fast tunneling enabled
  Last input 00:00:09, output 00:00:03, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
```

```
Queueing strategy: fifo
Output queue: 0/0 (size/max)
5 minute input rate 8000 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
   230 packets input, 207952 bytes, 0 no buffer
   Received 0 broadcasts, 0 runts, 0 giants, 0 throttles
   0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
   23 packets output, 1842 bytes, 0 underruns
   0 output errors, 0 collisions, 0 interface resets
   0 output buffer failures, 0 output buffers swapped out
```

Checking the Mroute State Information pertaining to the Default MDTGPA on a Specific PE

This section focuses on a specific PE, as defined by the high level diagram in Figure 2-3. Using the state diagram in Figure 2-7, you can check the health of the default MDT in the multicast core network. This diagram helps to check the mroutes (also known as multicast states) in the global table for the mVPN GPA address, defined by the **vrf mdt default** *GPA* command.

Specifically, this section guides you through checks on the (*, GPA): existence, Z flag, OIL, IIL, and so forth. Depending on the PIM mode for the default mdt, (S,GPA) states is also checked in Figure 2-8.

The following assumptions are made:

- Default MDT mode is PIM SM with SPT equals to 0 or Bidir.
- PIM neighbors have been checked on all PEs for this mVPN.
- Focus is on the mVPN of a specific enterprise customer, and thus on a specific default MDT.

Figure 2-7 Checking the MROUTE Table for the Default mdt Address_Part I

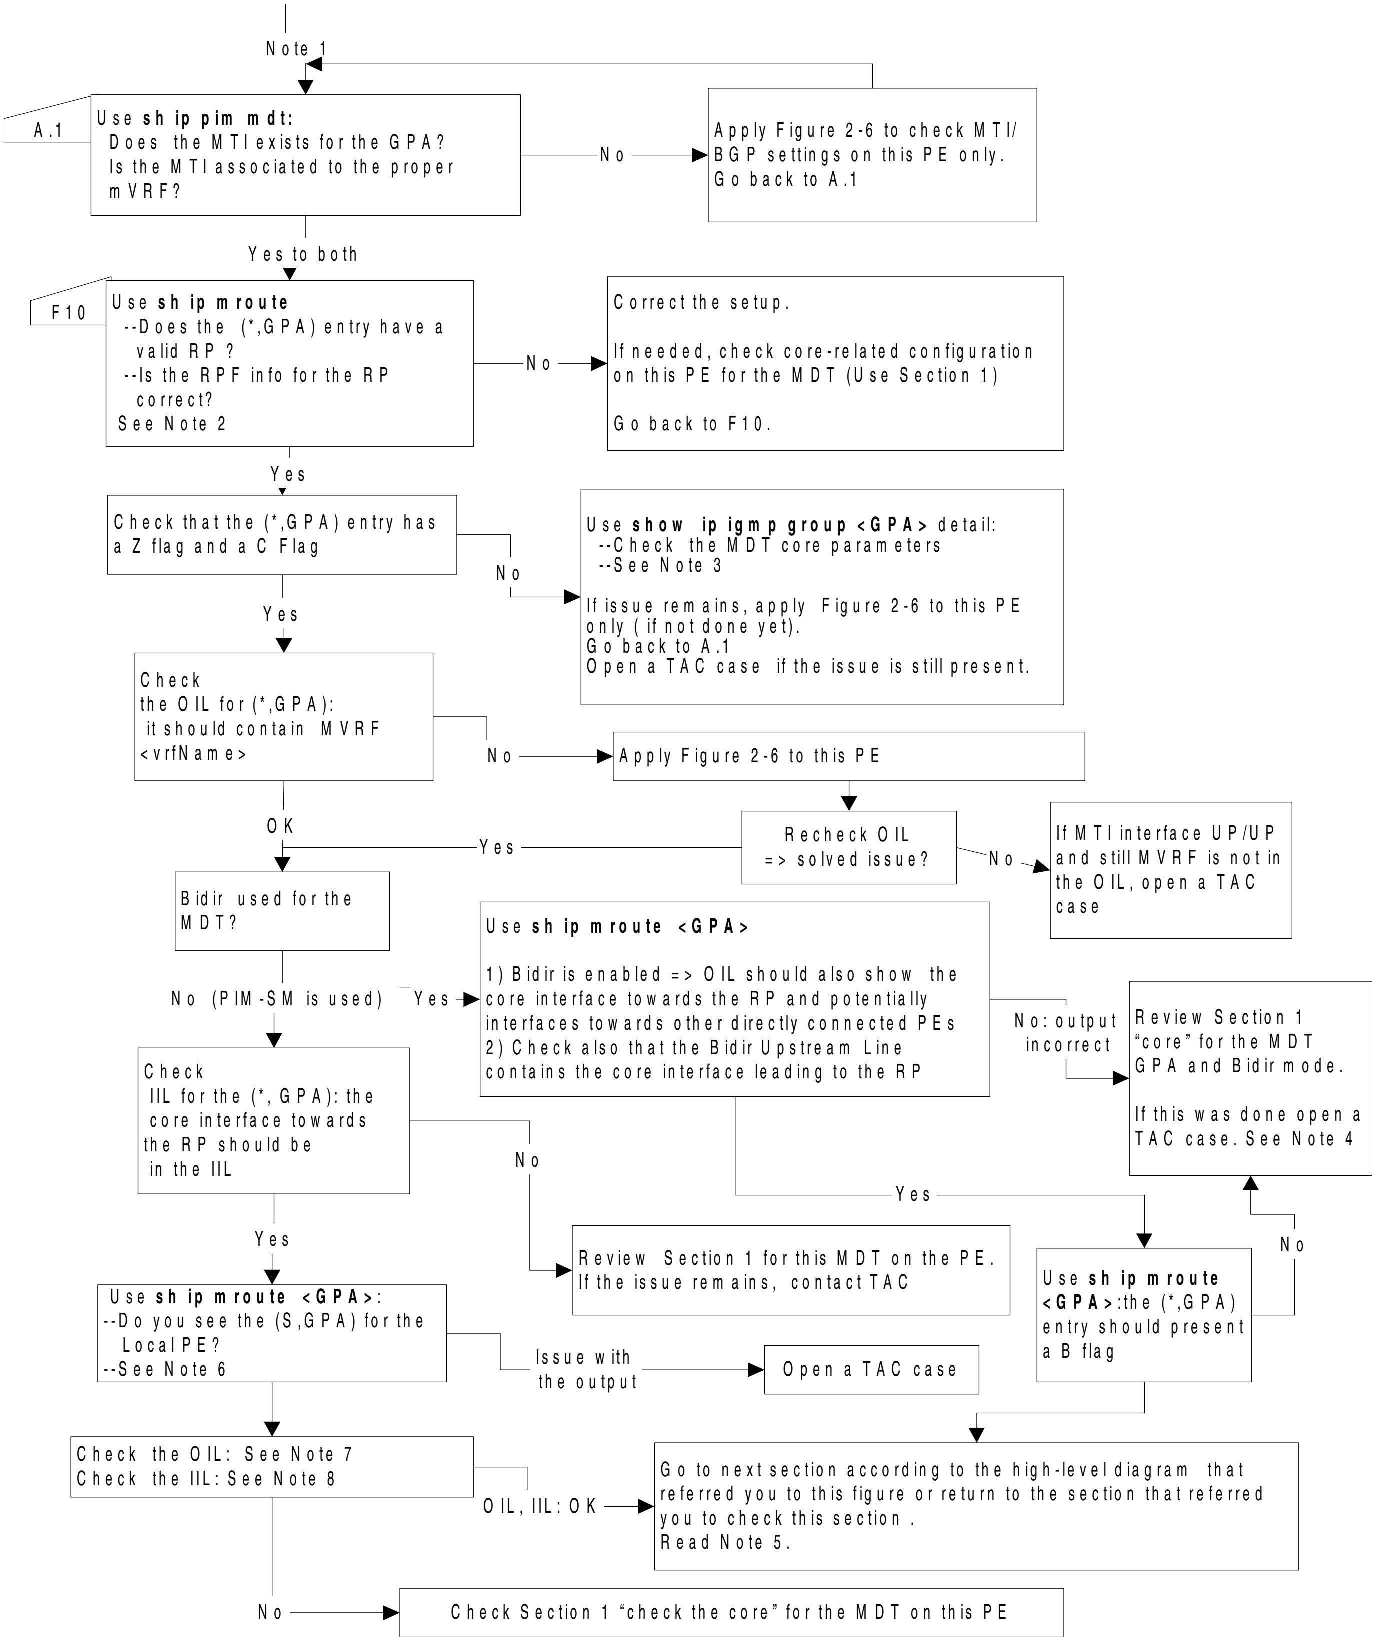
Note 1
A.1
Use **sh ip pim mdt:**
Does the MTI exists for the GPA?
Is the MTI associated to the proper mVRF?
No
Apply Figure 2-6 to check MTI/ BGP settings on this PE only.
Go back to A.1
Yes to both
F10
Use **sh ip mroute**
--Does the (*,GPA) entry have a valid RP ?
--Is the RPF info for the RP correct?
See Note 2
No
Correct the setup.
If needed, check core-related configuration on this PE for the MDT (Use Section 1)
Go back to F10.
Yes
Check that the (*,GPA) entry has a Z flag and a C Flag
No
Use **show ip igmp group <GPA>** detail:
--Check the MDT core parameters
--See Note 3
If issue remains, apply Figure 2-6 to this PE only (if not done yet).
Go back to A.1
Open a TAC case if the issue is still present.
Yes
Check
the OIL for (*,GPA):
it should contain MVRF <vrfName>
No
Apply Figure 2-6 to this PE
OK
Yes
Recheck OIL
=> solved issue?
No
If MTI interface UP/UP and still MVRF is not in the OIL, open a TAC case
Bidir used for the MDT?
No (PIM-SM is used)
Yes
Use **sh ip mroute <GPA>**
1) Bidir is enabled => OIL should also show the core interface towards the RP and potentially interfaces towards other directly connected PEs
2) Check also that the Bidir Upstream Line contains the core interface leading to the RP
No: output incorrect
Review Section 1 "core" for the MDT GPA and Bidir mode.
If this was done open a TAC case. See Note 4
Check
IIL for the (*, GPA): the core interface towards the RP should be in the IIL
No
Yes
Review Section 1 for this MDT on the PE.
If the issue remains, contact TAC
No
Use **sh ip mroute <GPA>**:the (*,GPA) entry should present a B flag
Yes
Use **sh ip mroute <GPA>**:
--Do you see the (S,GPA) for the Local PE?
--See Note 6
Issue with the output
Open a TAC case
Check the OIL: See Note 7
Check the IIL: See Note 8
OIL, IIL: OK
Go to next section according to the high-level diagram that referred you to this figure or return to the section that referred you to check this section .
Read Note 5.
No
Check Section 1 "check the core" for the MDT on this PE

Notes for Figure 2-7

Note 1

- Assumptions are the same as those given for Figure 2-4, with respect to the previously checked set of PEs.
- N>=3.

Note 2

- The valid RP address is the proper RP IP address for the range of addresses to which the MDT belongs.
- Note that in the case of Bidir the RP might be a dummy RP.

Note 3

This is an appropriate PIM type for the range to which the MDT belongs.

Note 4

Take into account that RP validity was checked earlier in this section.

Note 5

Your next step is most likely to proceed to Figure 2-8.

Note 6

Remember that it should not have a Z flag. Refer to the sample output.

Note 7

Do you see the core RPF interface towards the RP and the interfaces towards any directly connected PEs?

Note 8

Do you see the BGP update source interface (or BGP Next Hop if any is defined in the VRF) (for instance Loop 0 in our setup)?

Figure 2-8 Checking the MROUTE Table for the Default mdt Address Part II

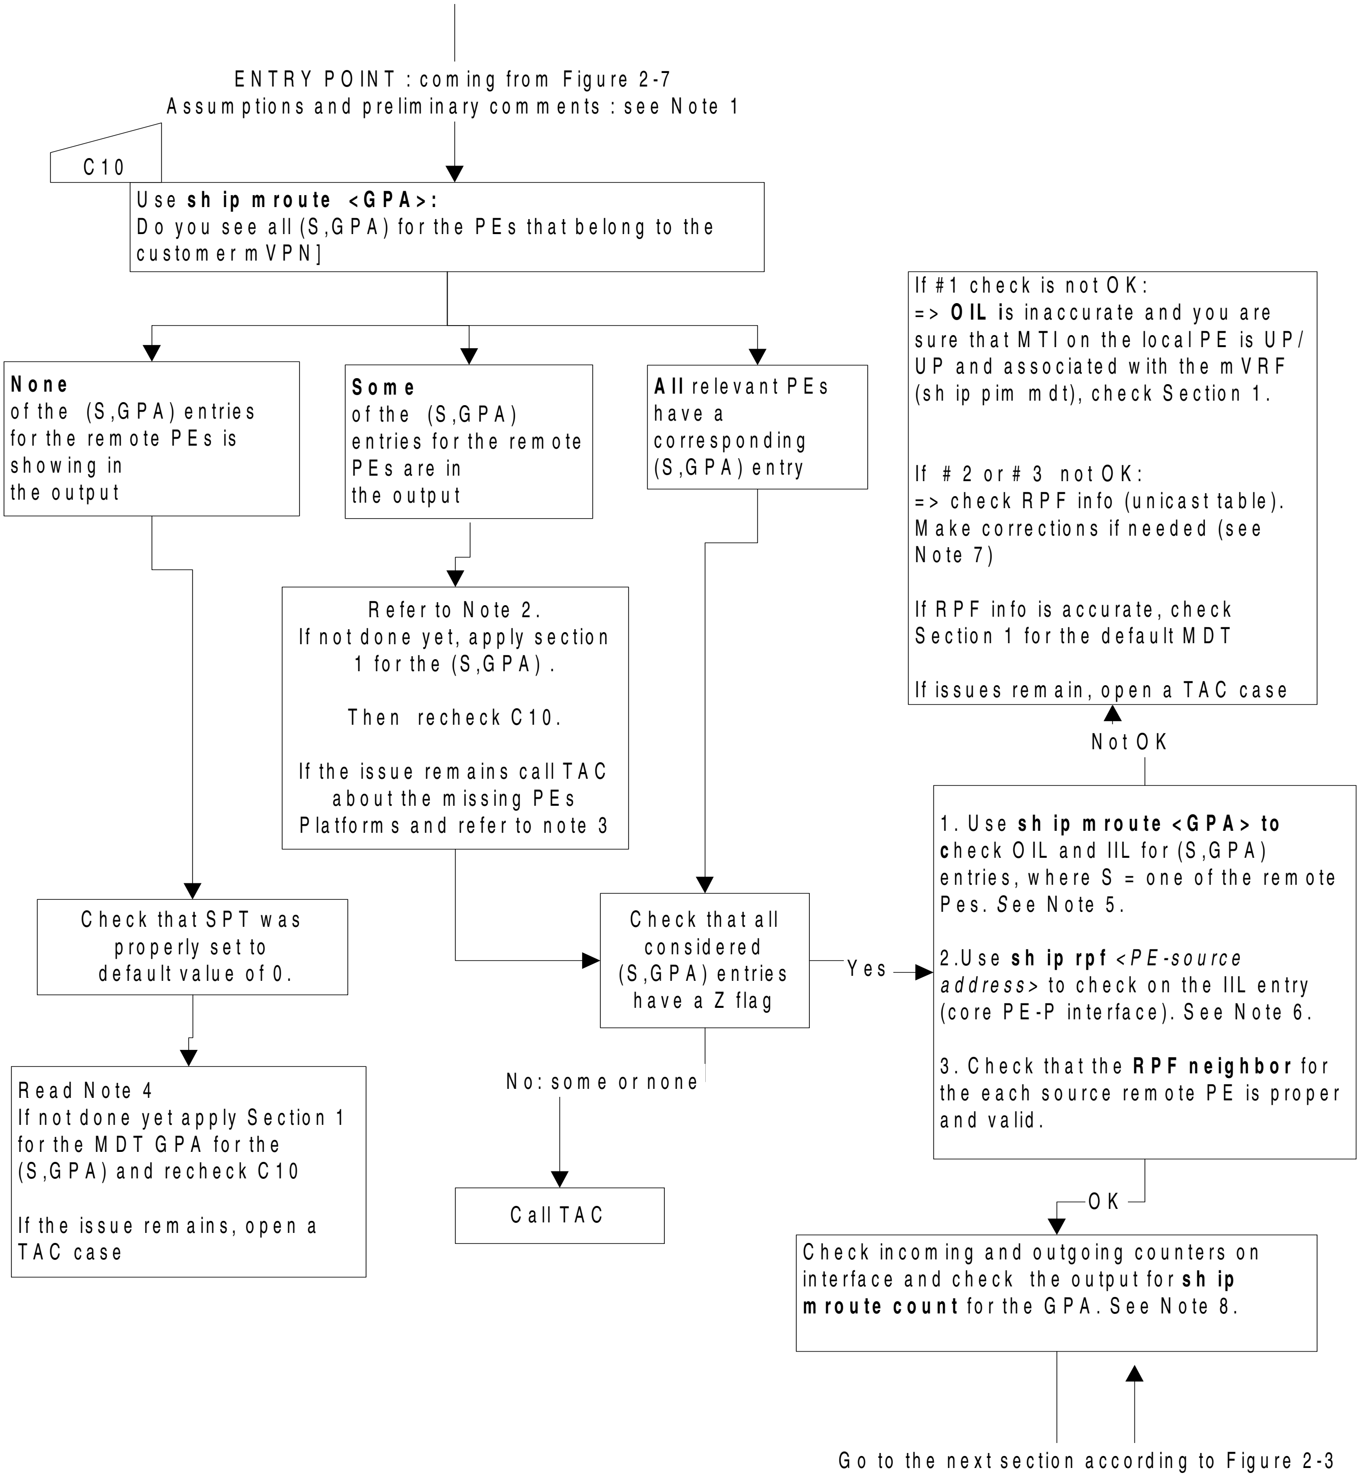
ENTRY POINT : coming from Figure 2-7
Assumptions and preliminary comments : see Note 1
C10
Use **sh ip mroute <GPA>:** Do you see all (S,GPA) for the PEs that belong to the customer mVPN]
None of the (S,GPA) entries for the remote PEs is showing in the output
Some of the (S,GPA) entries for the remote PEs are in the output
All relevant PEs have a corresponding (S,GPA) entry
If #1 check is not OK: => **OIL i**s inaccurate and you are sure that MTI on the local PE is UP/UP and associated with the mVRF (sh ip pim mdt), check Section 1.
If # 2 or # 3 not OK: => check RPF info (unicast table). Make corrections if needed (see Note 7)
If RPF info is accurate, check Section 1 for the default MDT
If issues remain, open a TAC case
Refer to Note 2. If not done yet, apply section 1 for the (S,GPA) .
Then recheck C10.
If the issue remains call TAC about the missing PEs Platforms and refer to note 3
Not OK
1. Use **sh ip mroute <GPA> to c**heck OIL and IIL for (S,GPA) entries, where S = one of the remote Pes. *S*ee Note 5.
2.Use **sh ip rpf** *<PE-source address>* to check on the IIL entry (core PE-P interface). See Note 6.
3. Check that the **RPF neighbor** for the each source remote PE is proper and valid.
Check that SPT was properly set to default value of 0.
Check that all considered (S,GPA) entries have a Z flag
Yes
Read Note 4
If not done yet apply Section 1 for the MDT GPA for the (S,GPA) and recheck C10
If the issue remains, open a TAC case
No: some or none
Call TAC
OK
Check incoming and outgoing counters on interface and check the output for **sh ip mroute count** for the GPA. See Note 8.
Go to the next section according to Figure 2-3

Notes for Figure 2-8

Note 1

- This section applies only if Bidir is not the PIM mode for the MDT default. Its goal is to check the (S,GPA) mstates for the remote PEs.
- Remember that the (S,GPA) for the local PE was checked in Figure 2-7.

Note 2

Remember that we are assuming that all remote PEs in the subset have already been checked and are healthy.

Note 3

If you need to call TAC about some PEs for which you see issues, you might want to keep checking the rest of the section and consider only the subset of PEs for which you see an entry (the PEs that are OK).

Note 4

Remember that we are assuming that all PEs in the subset have already been checked and are healthy.

Note 5

- The MVRF should be in the OIL. There might also be some other interface entries.
- In the OIL in case the PE has some leaves on the (S,G) tree besides the MTI.

Note 6

The core RPF interface towards the Source PE (S) should be in the IIL.

Note 7

Corrections for RPF: Ensure that unicast and multicast paths are congruent.

Note 8

- Verify that both incoming and outgoing counters are increasing (there should be at least the PIM control traffic flowing between PEs over the MDT, for example, PE to PE PIM hellos for the (*,GPA)).
- Use the following command:
  - Look at interface counters for the Tunnel interface and use the **show ip mroute count** command for the GPA.
  - Verify that the FWDING counter is increasing when repeating this command.

Sample Output and Show Commands to use for Checking the Mroutes for the MDT Tree

The **show ip pim rp mapping** command allows you to check that the RP has been defined for the MDT default group address; it also ensures that the PIM mode defined for this RP is the one to be used for the default MDT. The following output is for Bidir and PIM SM SPT 0 modes, respectively.

Note you can also use the **show ip pim rp-hash** *group* command to check which RP is associated to a specific multicast group on a router, when the RP election mode is BSR or auto-rp.

```
PE1#show ip pim rp mapping
PIM Group-to-RP Mappings
```

```
Acl: 1, Static-Override, Bidir Mode
    RP: 172.16.100.200 (?)
```

Use the **show ip rpf** command to check RPF information on a router for a given source (PE) or the RP:

```
PE1#show ip rpf 172.16.100.200
RPF information for ? (172.16.100.200)
  RPF interface: Serial0/0
  RPF neighbor: ? (172.16.1.11)
  RPF route/mask: 172.16.100.200/32
  RPF type: unicast (ospf 1)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
```

Use the following command to check that the default mdt is associated via the tunnel interface:

Note In order to verify that the MTI is associated to the VRF, do not get confused by the output of the **show ip pim vrf RED interface** command. The actual association is shown only by using the **show ip pim mdt** command.

```
PE1#show ip pim mdt
  MDT Group       Interface   Source                   VRF
* 239.0.0.10      Tunnel0     Loopback0                RED
```

Use the following command to check that the local router is subscribed to the default mdt group:

```
PE1#show ip igmp groups 239.0.0.10
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter
239.0.0.10       Loopback0                1w5d      stopped   0.0.0.0
PE1#

PE1#show ip mroute 239.0.0.10 count
IP Multicast Statistics
2 routes using 1388 bytes of memory
2 groups, 0.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 239.0.0.10, Source count: 0, Group pkt count: 165911
  RP-tree: Forwarding: 165911/1/74/0, Other: 165911/0/0
PE1#
```

Following is the standard output of the **show ip mroute** command on a PE for the default mdt group address respectively for Bidir and PIM SM SPT equals 0 modes:

```
PE1#show ip mroute 239.0.0.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 00:02:17/00:02:47, RP 172.16.100.200, flags: BCZ
  Bidir-Upstream: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 00:02:17/00:01:08
    Serial0/0, Bidir-Upstream/Sparse, 00:02:17/00:00:00


PE1#show ip mroute 239.0.0.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 2w2d/stopped, RP 172.16.100.200, flags: SJCFZ
  Incoming interface: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 2w2d/00:00:00

(172.16.100.1, 239.0.0.10), 2w2d/00:03:19, flags: FT
  Incoming interface: Loopback0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 2w2d/00:03:27

(172.16.100.2, 239.0.0.10), 2w2d/00:02:49, flags: JTZ
  Incoming interface: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 2w2d/00:00:00
(172.16.100.3, 239.0.0.10), 2w2d/00:02:58, flags: JTZ
  Incoming interface: Serial0/0, RPF nbr 172.16.1.11
  Outgoing interface list:
    MVRF RED, Forward/Sparse, 2w2d/00:00:00
```

The following is the output on the P-RP in Bidir mode:

```
P-RP#show ip mroute 239.0.0.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 1w5d/00:03:16, RP 172.16.100.200, flags: B
  Bidir-Upstream: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 1w5d/00:02:32
    Serial1/0, Forward/Sparse, 1w5d/00:03:16
```

The following is the output for the RP in PIM SM mode:

```
P-RP#show ip mroute 239.0.0.10
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group
       V - RD & Vector, v - Vector
Outgoing interface flags: H - Hardware switched, A - Assert winner
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.0.0.10), 4w1d/00:03:23, RP 172.16.100.200, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial0/0, Forward/Sparse, 4w1d/00:02:58
    Serial1/0, Forward/Sparse, 4w1d/00:03:23

(172.16.100.1, 239.0.0.10), 4w1d/00:01:18, flags: PT
  Incoming interface: Serial1/0, RPF nbr 172.16.213.13
  Outgoing interface list: Null

(172.16.100.2, 239.0.0.10), 4w1d/00:02:46, flags: PT
  Incoming interface: Serial0/0, RPF nbr 172.16.212.12
  Outgoing interface list: Null

(172.16.100.3, 239.0.0.10), 4w1d/00:00:51, flags: PT
  Incoming interface: Serial1/0, RPF nbr 172.16.213.13
  Outgoing interface list: Null
```

Section 3: Verifying the Customer-Facing Settings of the PEs

This section provides guidelines for isolating issues between the core and enterprise network. It is organized as follows:

- **Main topology**: the recommended troubleshooting steps for checking the customer facing settings on the PEs are explained based on a sufficient subset of the main topology (see Figure 2-9) and using the flow charts in Figures 2-11, 2-12, 2-13 and 2-14.
- We then study **additional topology- cases** (see Figures 2-21,2- 22, 2-23) using the same methodology, and illustrate these with appropriate output.

Assumptions

The following assumptions are made:

- The C-source, C-RP, and C-receivers are *not* behind the same PE.
- The customer mapping agent and customer RP is the same router and is behind the PE$_{C-RP}$
- Although the topology in Figure 2-9 shows only one receiver, there might be several receiving sites, in which case you can repeat the steps pertaining to PE $_{C-RCVR}$.

Have the following information available:

- Customer VRF name.
- Customer source IP address(es) and C-SRC.
- Group address(es) and C-GPA.
- Sites (PE and circuit) that connect to the receiver or receivers' IP addresses.
- Customer RP mode (AutoRP, static RP, bidir) and RP address.
- Customer PIM mode, and if running PIM SM, the threshold configuration specified for switching from shared tree to source tree.

Note This section covers only PIM SM (with static RP or AutoRP).

Main Topology for troubleshooting Customer-Facing Settings on the PEs

Normally, mVPN customers open a trouble ticket with the following problem definition:

- No multicast traffic is received at one or all sites.
- Intermittent traffic issues are occurring, which can include slowness and less throughput for the receiving applications.

For these scenarios, the troubleshooting steps take into account the control and the customer data plane. First, verify the control plane, which includes checking the performance of the multicast route states on the $PE_{C\text{-}SRC}$, $PE_{C\text{-}RCVR,}$ and $PE_{C\text{-}RP\ .}$ When the Mroute states are verified, check the data flow traffic rates and identify traffic drops.

Figure 2-9 shows the topology used for troubleshooting the issues described in this section.

Note that this figure considers a subset of the global topology (Figure 1-1) used throughout this troubleshooting guide. The multicast core network of the service provider is represented virtually by the MTI LAN interface.

Figure 2-9 Topology 1: Troubleshooting Mroute States on Different PEs

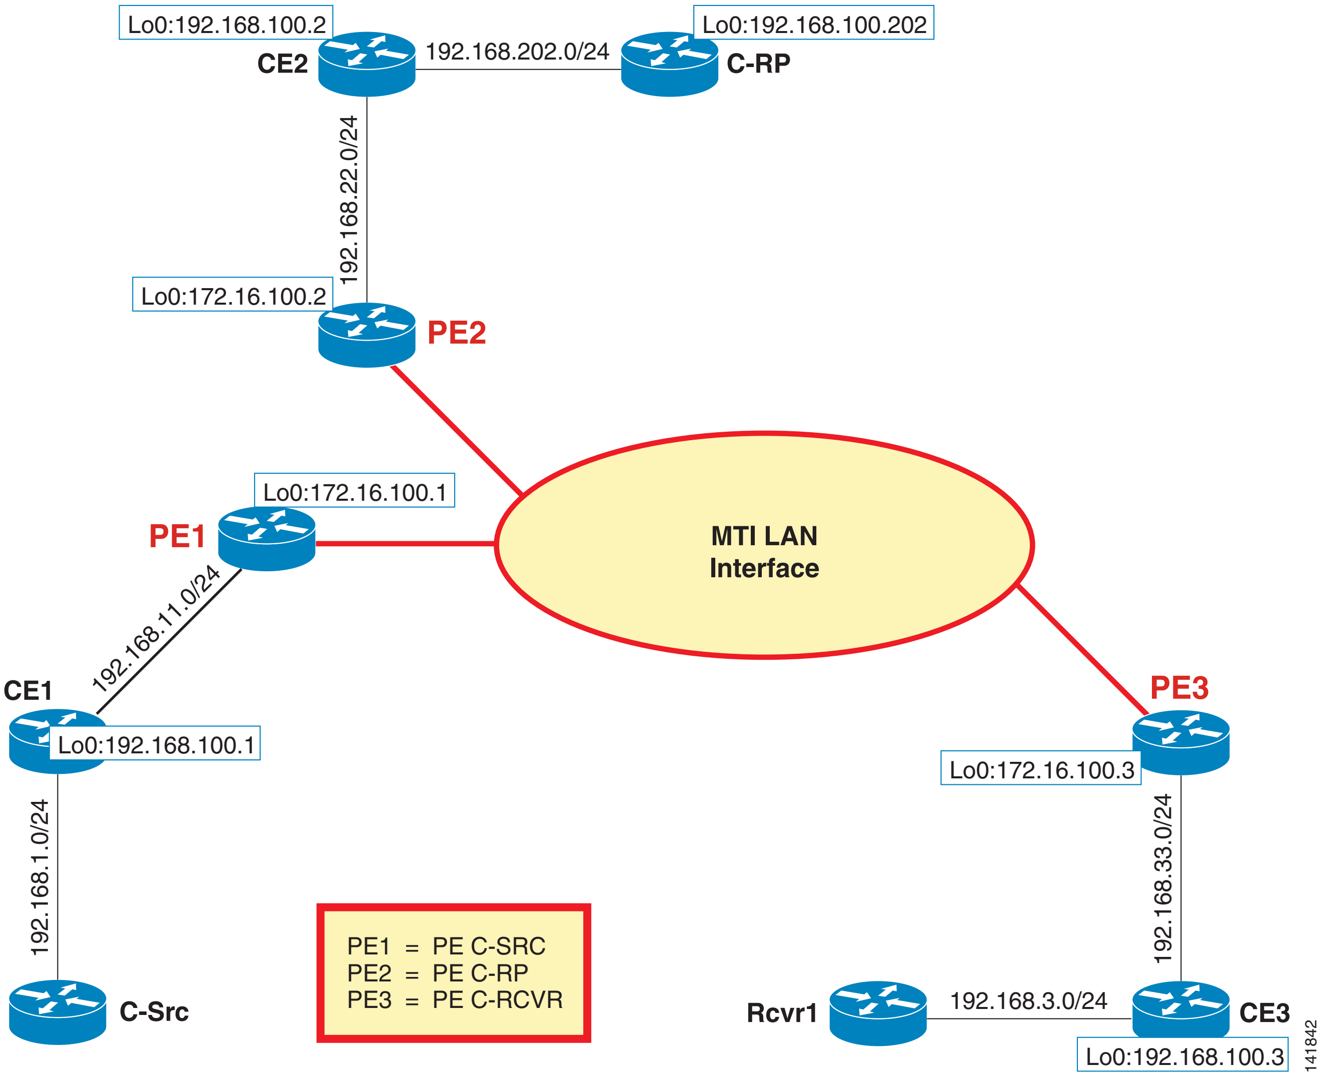

High-Level Diagram for Troubleshooting the Customer-Facing Settings on the PEs

Figure 2-10 provides the high-level diagram for checking the control plane. It is the main figure and should serve as reference point for the other topologies as well.

We start by checking for basic information on the PE$_{C\text{-}RP,}$ PE$_{C\text{-}SRC,}$ and PE$_{C\text{-}RCVR}$. After verification, we move to check the Mroute states for C-GPA and finally verify the data plane for failures.

Assumptions

- Ensure that you have the VRF name, C-GPA, C-SRC, C-RP & PIM mode information ready.
- We repeat the check for each Source and receiver PE for this C-GPA.

Figure 2-10 High-Level Diagram for Troubleshooting the Control Plane

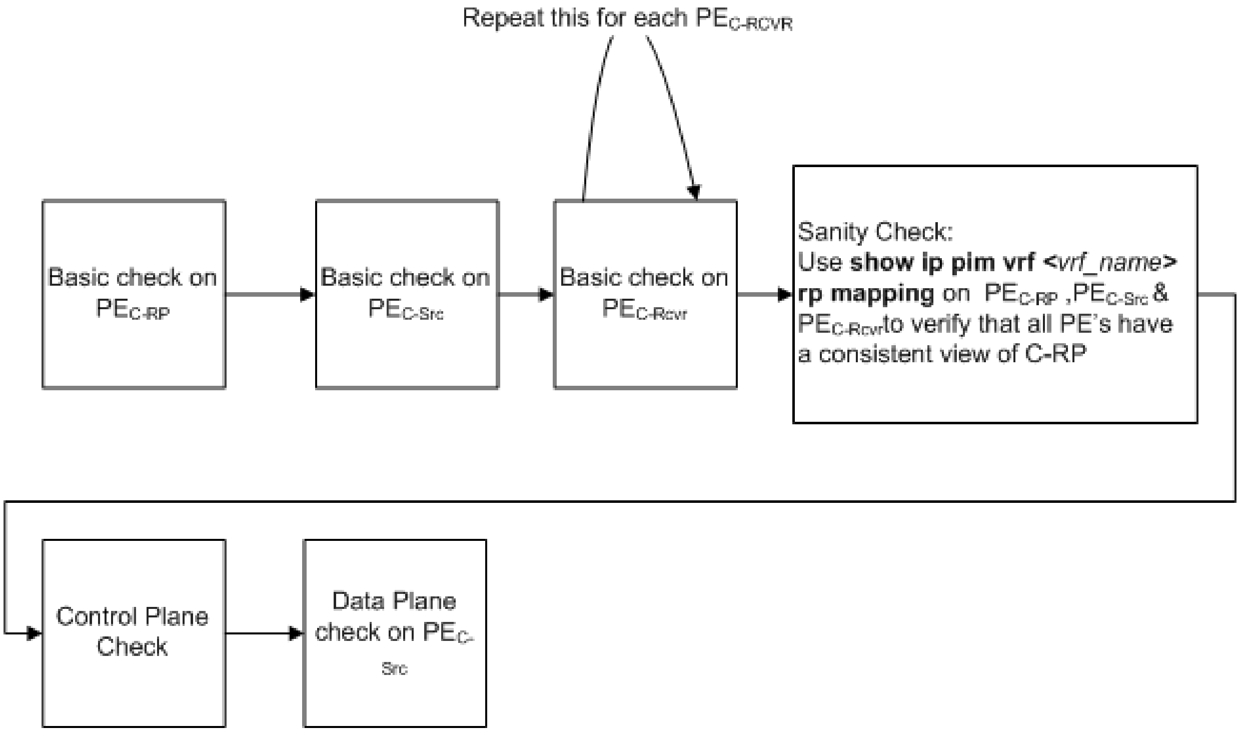
Repeat this for each PE$_{C\text{-}RCVR}$
Basic check on PE$_{C\text{-}RP}$
Basic check on PE$_{C\text{-}Src}$
Basic check on PE$_{C\text{-}Rcvr}$
Sanity Check:
Use **show ip pim vrf** *<vrf_name>* **rp mapping** on PE$_{C\text{-}RP}$,PE$_{C\text{-}Src}$ & PE$_{C\text{-}Rcvr}$to verify that all PE's have a consistent view of C-RP
Control Plane Check
Data Plane check on PE$_{C\text{-}Src}$

Verify the following on the PEs:

- Multicast routing is enabled for the mVRF.
- PEs and CEs are PIM neighbors in the VRF context.
- MTI tunnel interface is up and is associated to the mVRF.
- Fast Switching is enabled on the PE-CE interfaces.
- RPF checks towards C-SRC and C-RP are valid.
- Presence of filters (for example, multicast boundary).
- C-RP information is available, and if AutoRP mode is used, verify that AutoRP information is received on PE$_{C\text{-}RP}$.

Verifying Customer-Facing Settings on the $PE_{C\text{-}RP}$

The flowchart in Figure 2-11 walks through the basic checks on the $PE_{C\text{-}RP}$.

Figure 2-11 Basic Check on PE_{C-RP}

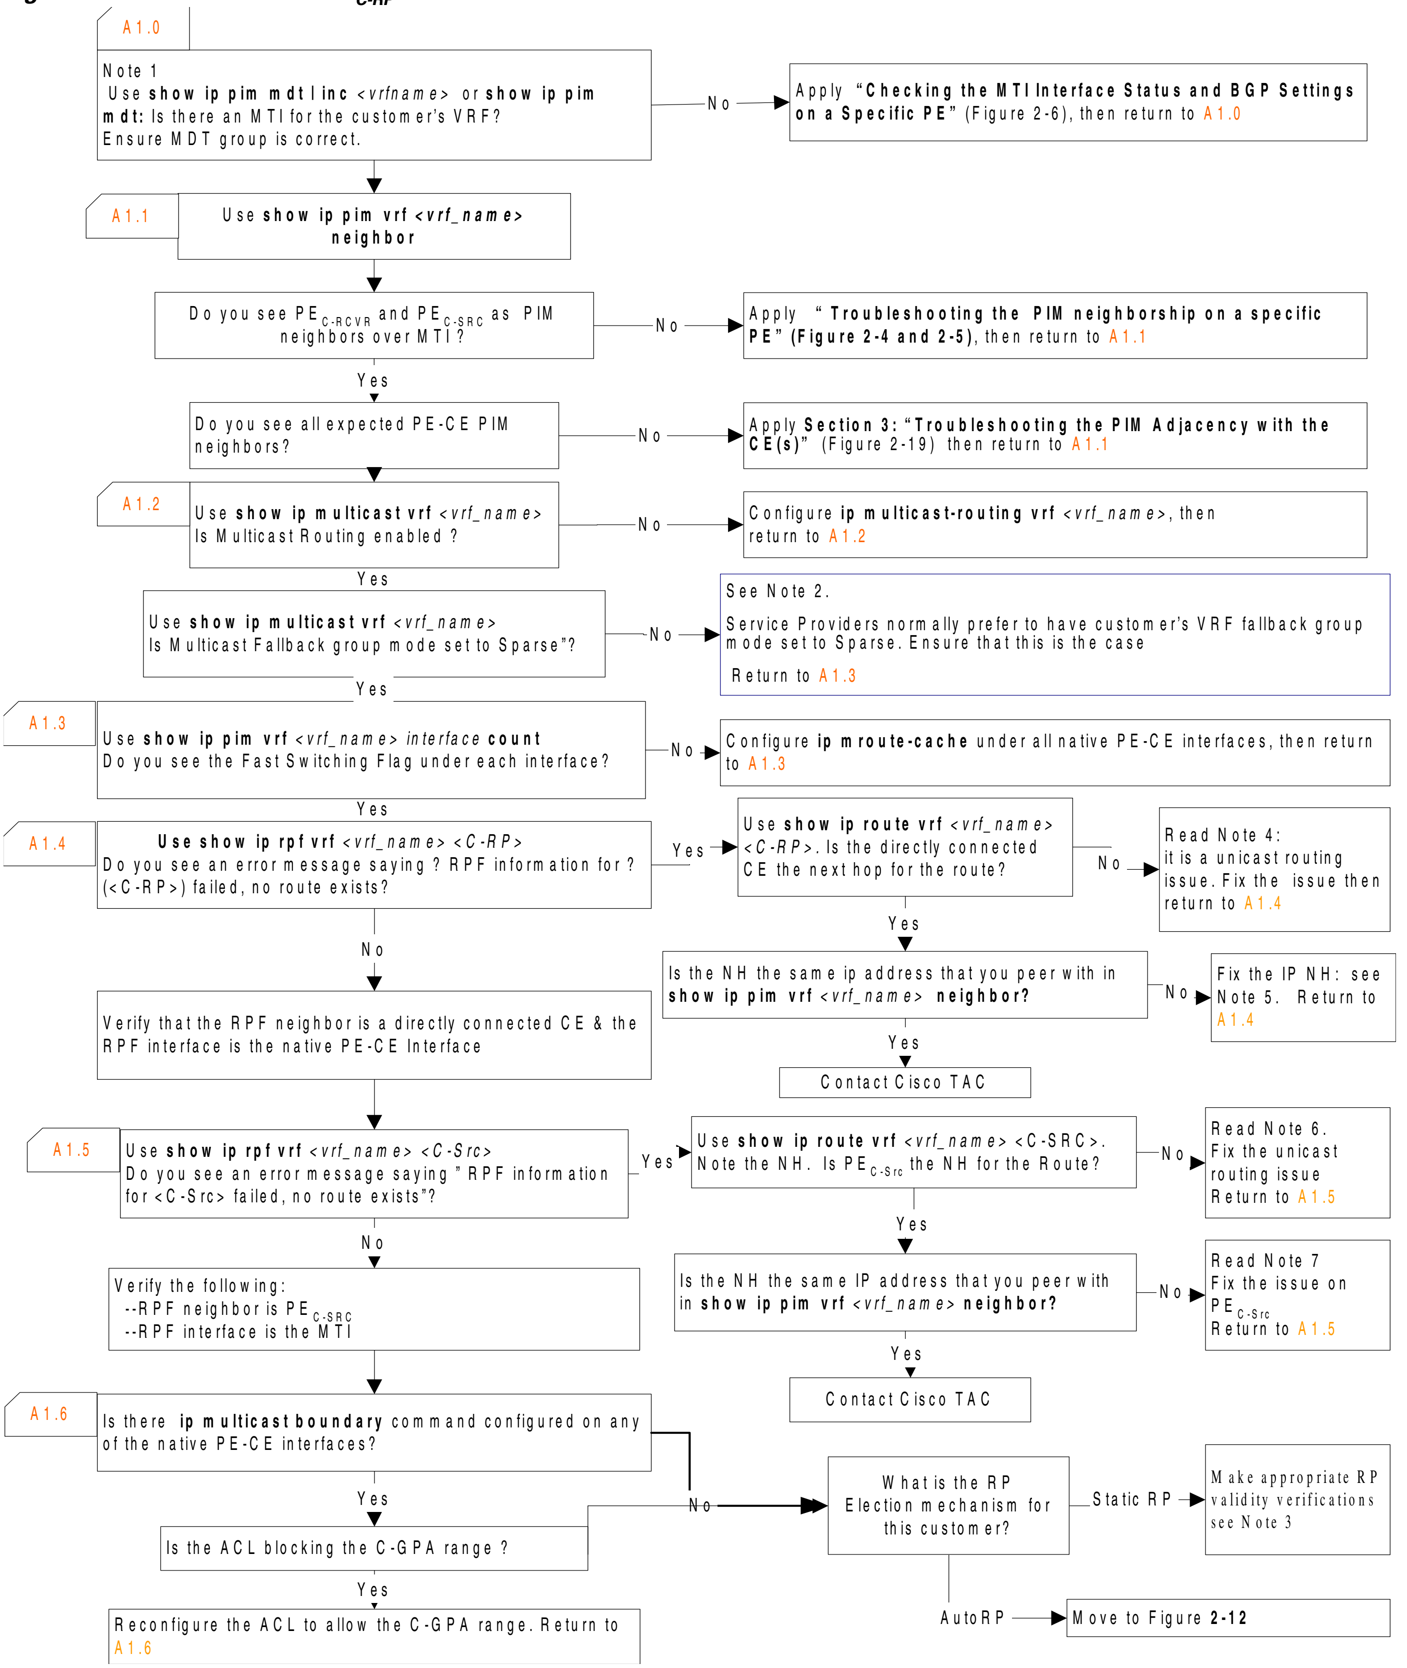
A1.0
Note 1
Use show ip pim mdt I inc <vrfname> or show ip pim mdt: Is there an MTI for the customer's VRF?
Ensure MDT group is correct.
No
Apply "Checking the MTI Interface Status and BGP Settings on a Specific PE" (Figure 2-6), then return to A1.0
A1.1
Use show ip pim vrf <vrf_name> neighbor
Do you see PE C-RCVR and PE C-SRC as PIM neighbors over MTI?
No
Apply "Troubleshooting the PIM neighborship on a specific PE" (Figure 2-4 and 2-5), then return to A1.1
Yes
Do you see all expected PE-CE PIM neighbors?
No
Apply Section 3: "Troubleshooting the PIM Adjacency with the CE(s)" (Figure 2-19) then return to A1.1
A1.2
Use show ip multicast vrf <vrf_name>
Is Multicast Routing enabled ?
No
Configure ip multicast-routing vrf <vrf_name>, then return to A1.2
Yes
Use show ip multicast vrf <vrf_name>
Is Multicast Fallback group mode set to Sparse"?
No
See Note 2.
Service Providers normally prefer to have customer's VRF fallback group mode set to Sparse. Ensure that this is the case
Return to A1.3
Yes
A1.3
Use show ip pim vrf <vrf_name> interface count
Do you see the Fast Switching Flag under each interface?
No
Configure ip mroute-cache under all native PE-CE interfaces, then return to A1.3
Yes
A1.4
Use show ip rpf vrf <vrf_name> <C-RP>
Do you see an error message saying ? RPF information for ? (<C-RP>) failed, no route exists?
Yes
Use show ip route vrf <vrf_name> <C-RP>. Is the directly connected CE the next hop for the route?
No
Read Note 4:
it is a unicast routing issue. Fix the issue then return to A1.4
Yes
Is the NH the same ip address that you peer with in show ip pim vrf <vrf_name> neighbor?
No
Fix the IP NH: see Note 5. Return to A1.4
Yes
Contact Cisco TAC
No
Verify that the RPF neighbor is a directly connected CE & the RPF interface is the native PE-CE Interface
A1.5
Use show ip rpf vrf <vrf_name> <C-Src>
Do you see an error message saying " RPF information for <C-Src> failed, no route exists"?
Yes
Use show ip route vrf <vrf_name> <C-SRC>.
Note the NH. Is PE C-Src the NH for the Route?
No
Read Note 6.
Fix the unicast routing issue
Return to A1.5
Yes
Is the NH the same IP address that you peer with in show ip pim vrf <vrf_name> neighbor?
No
Read Note 7
Fix the issue on PE C-Src
Return to A1.5
Yes
Contact Cisco TAC
No
Verify the following:
--RPF neighbor is PE C-SRC
--RPF interface is the MTI
A1.6
Is there ip multicast boundary command configured on any of the native PE-CE interfaces?
Yes
Is the ACL blocking the C-GPA range ?
Yes
Reconfigure the ACL to allow the C-GPA range. Return to A1.6
No
What is the RP Election mechanism for this customer?
Static RP
Make appropriate RP validity verifications see Note 3
AutoRP
Move to Figure 2-12

Notes for Figure 2-11

Note 1

NH = next hop.

Note 2

For more information, refer to the following URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_4/gtautorp.htm.

Note 3

PIM SM with Static RP: verify that **ip pim rp-address vrf** *vrf_name C-RP-address* [**group-list ACL**] [**override**] exists on the PE and that the ACL reference is correct.

Note 4

If you do not see the prefix, or if the NH is a node other then the directly connected CE, then it is a unicast routing issue.

Note 5

- The IP NH for the prefix to reach the C-RP needs to be a PIM neighbor as well.
- If there are any fancy route maps on the CE or the PE (C-RP), verify that these do not interfere with the control plane and modify accordingly.

Note 6

If you do not see the prefix, or the NH is a node other then the PE$_{\text{C-SRC}}$, then it is a unicast routing issue.

Note 7

- If the NH is one of the IP addresses on PE$_{\text{C-SRC}}$ that is not the same ip address as the PIM neighbor over MTI Interface, then mVPN will break.
- All PEs must advertise prefixes with BGP next hop, which is the tunnel source loopback ip address.

Sample Output for Checks of the PE Settings on the Customer RP Site

Step 1 Performing basic checks on PE$_{\text{C-RP}}$ on the sanity of the MTI:

```
PE2#show ip multicast vrf RED
Multicast Routing: enabled
Multicast Multipath: disabled
Multicast Route limit: No limit
Multicast Triggered RPF check: enabled
Multicast Fallback group mode: Sparse

PE2#show ip pim mdt | include MDT|RED
MDT Group       Interface   Source                   VRF
* 239.0.0.10      Tunnel0     Loopback0                RED


PE2#show ip pim vrf RED neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.22.2      Serial1/0                2d04h/00:01:37    v2    1 / P
172.16.100.1      Tunnel0                  00:04:46/00:01:24 v2    1 / P
172.16.100.3      Tunnel0                  2d04h/00:01:25    v2    1 / DR P
```

```
PE2#show ip pim interface S1/0 detail
Serial1/0 is up, line protocol is up
Internet address is 192.168.22.1/24
Multicast switching: fast
Multicast packets in/out: 5614/161535
Multicast boundary: not set
Multicast TTL threshold: 0
PIM: enabled
PIM version: 2, mode: sparse
PIM DR: 0.0.0.0
PIM neighbor count: 1
PIM Hello/Query interval: 30 seconds
PIM NBMA mode: disabled
PIM ATM multipoint signalling: disabled
PIM domain border: disabled
PIM neighbors rpf proxy capable: TRUE
Multicast Tagswitching: disabled
```

Step 2 Ensure that the RP address is known in the VRF at the PE level:

```
PE2#show ip rpf vrf RED 192.168.100.202 → C-RP
RPF information for ? (192.168.100.202)
RPF interface: Serial1/0
RPF neighbor: ? (192.168.22.2)
RPF route/mask: 192.168.100.202/32
RPF type: unicast (bgp 1)
RPF recursion count: 1
Doing distance-preferred lookups across tables

PE2#show ip rpf vrf RED 192.168.11.1     → C-Src
RPF information for ? (192.168.11.1)
RPF interface: Tunnel0
RPF neighbor: ? (172.16.100.1)
RPF route/mask: 192.168.11.0/24
RPF type: unicast (bgp 1)
RPF recursion count: 0
Doing distance-preferred lookups across tables
BGP originator: 172.16.100.1
```

Step 3 After the above is verified, the next step is to check the RP mapping on $\text{PE}_{\text{C-RP .}}$

Use the **show ip pim rp mapping** command for this purpose. Note the following configuration checks:

- Static RP mapping

  If the customer is running static RP, verify that all the PEs in the mVRF are configured with the correct static RP:

  ip pim rp-address *rp-address* [**group_acl**] [**override**]

  It is recommended to use the **override** keyword to override the dynamically learned RP mappings. Also verify that the value specified by *group_acl* is correctly configured in the ACL and correctly referenced in the above command.

- AutoRP mapping using Cisco AutoRP

  For AutoRP, verify the following:

  - $\text{PE}_{\text{C-RP}}$ learns the RP mapping from the CE side through AutoRP (**show ip pim rp-hash** in the vrf context)
  - AutoRP advertises the RP mapping to the rest of the PEs connected on the MTI

The flowchart in Figure 2-12 walks through the essential configuration of AutoRP on $PE_{C-RP.}$

Figure 2-12 AutoRP Troubleshooting on PE_{C-RP}

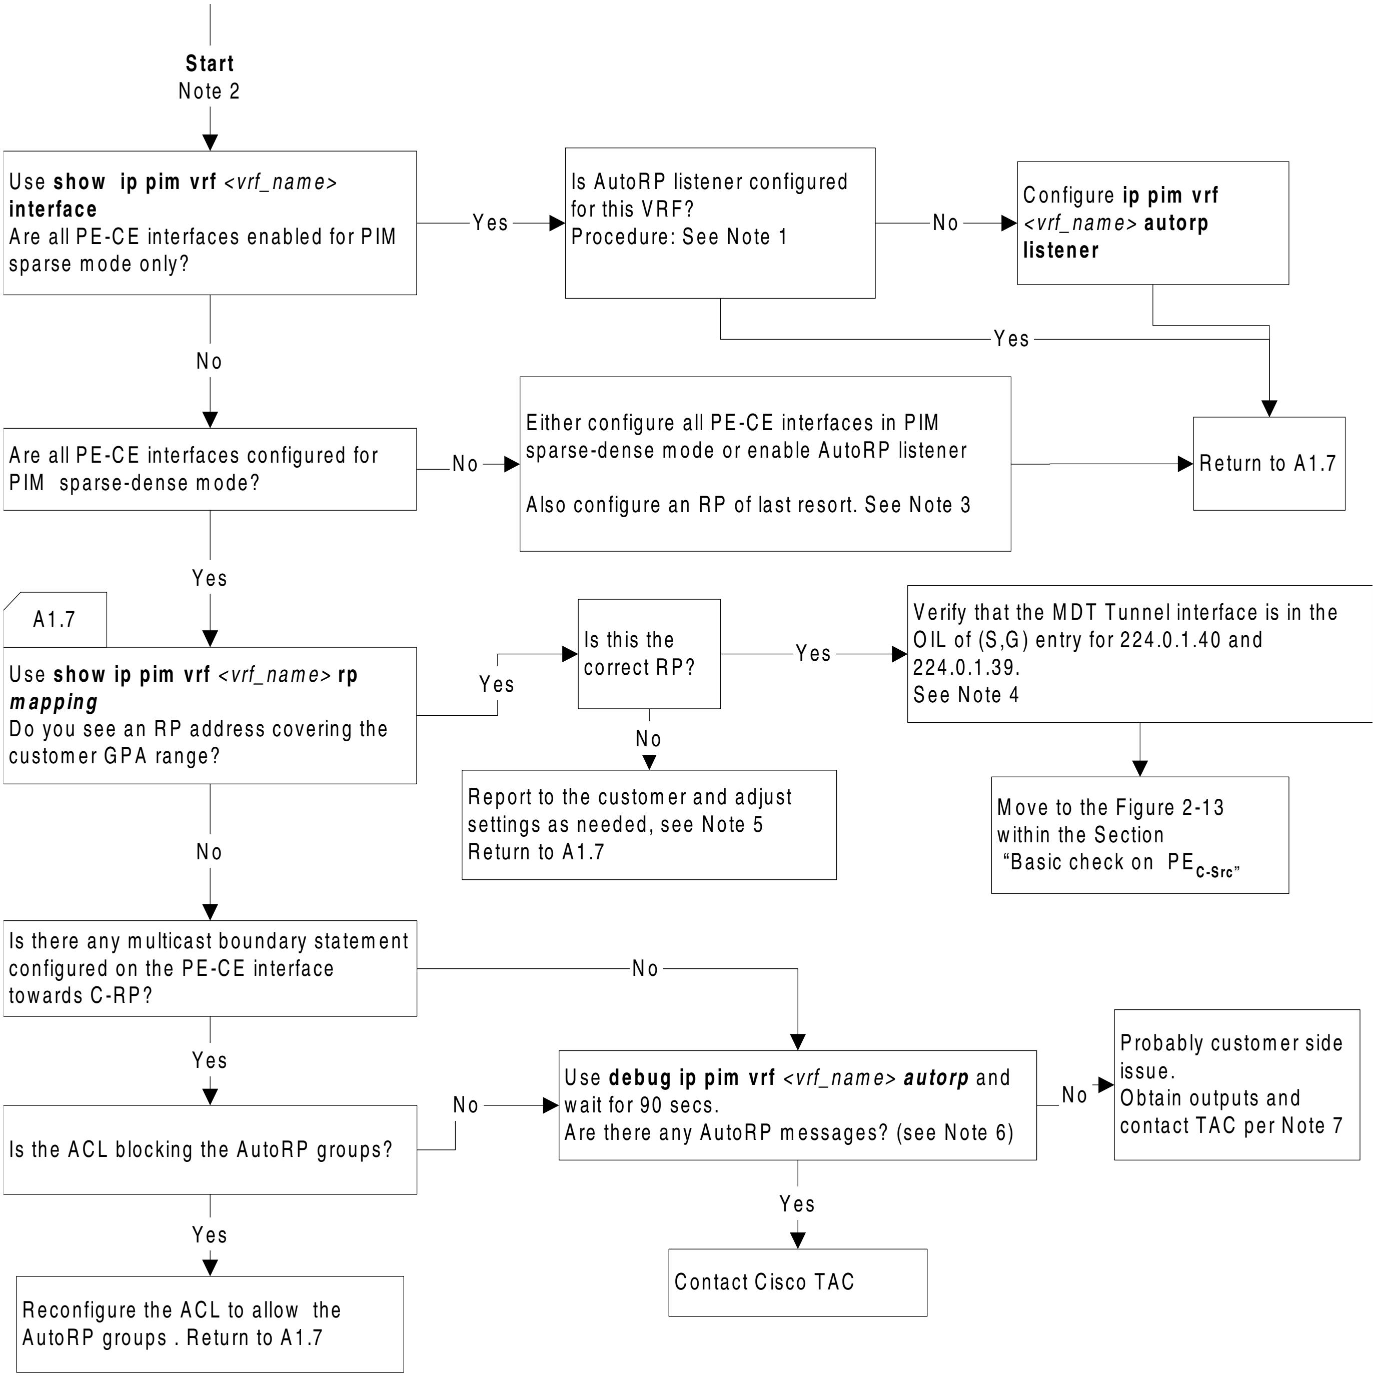
Start
Note 2
Use show ip pim vrf <vrf_name> interface
Are all PE-CE interfaces enabled for PIM sparse mode only?
Yes
Is AutoRP listener configured for this VRF?
Procedure: See Note 1
No
Configure ip pim vrf <vrf_name> autorp listener
Yes
No
Are all PE-CE interfaces configured for PIM sparse-dense mode?
No
Either configure all PE-CE interfaces in PIM sparse-dense mode or enable AutoRP listener
Also configure an RP of last resort. See Note 3
Return to A1.7
Yes
A1.7
Use show ip pim vrf <vrf_name> rp mapping
Do you see an RP address covering the customer GPA range?
Yes
Is this the correct RP?
Yes
Verify that the MDT Tunnel interface is in the OIL of (S,G) entry for 224.0.1.40 and 224.0.1.39.
See Note 4
No
Report to the customer and adjust settings as needed, see Note 5
Return to A1.7
Move to the Figure 2-13 within the Section "Basic check on PE C-Src"
No
Is there any multicast boundary statement configured on the PE-CE interface towards C-RP?
No
Yes
Is the ACL blocking the AutoRP groups?
No
Use debug ip pim vrf <vrf_name> autorp and wait for 90 secs.
Are there any AutoRP messages? (see Note 6)
No
Probably customer side issue.
Obtain outputs and contact TAC per Note 7
Yes
Contact Cisco TAC
Yes
Reconfigure the ACL to allow the AutoRP groups . Return to A1.7

Notes for Figure 2-12

Note 1

- Do you see "AutoRP groups over sparse mode interface is enabled" under **show ip pim vrf** *vrf_name* **autorp**?

  OR

- Does the **show running-config | inc autorp list** command display "ip pim vrf *vrf_name* autorp listener?"

Note 2

Auto RP groups are 224.0.1.39 and 224.0.1.40.

Note 3

- We either need to configure all PE and CE interfaces in PIM sparse-dense mode, or enable AutoRP listener to densely switch the AutoRP groups to allow AutoRP work successfully. AutoRP is the recommended option.
- Also, configure an RP of last resort to avoid all customer traffic getting densely switched.
- Read the comment on Auto-RP at

http://www.cisco.com/warp/public/732/Tech/multicast/docs/multicastdesign.pdf

Note 4

For a complete verification of the flags and interface lists, refer to Table 2-1.

Note 5

Report to the customer and ensure that they advertise the correct RP to group mapping and we are in sync with the correct RP address. We might be seeing messages like %PIM-6-INVALID_RP_JOIN in the logs.

Note 6

Verify using the **show logging** command or check the console session.

Note 7

- Use the **show ip mroute 224.0.1.39** and **show ip mroute 224.0.1.40** commands from the CE to verify that the interface facing the PE is in the OIL.
- If the CE has the Interface in the OIL and we are still not seeing the RP mapping, then open a case with Cisco TAC, supplying the relevant output.

Table 2-1 VRF Mroute states on PE$_{C\text{-}RP}$ for AutoRP Groups

For (*,224.0.1.39)	RP 0.0.0.0, Flags = D
Incoming interface	Null, RPF neighbor 0.0.0.0
Outgoing interfaces	Default mdt tunnel interface plus all native PE-CE interfaces
For (S,224.0.1.39)	Flags PT (flag P is there because the assumption is that the customer mapping agent /RP candidate is the same router and on the local site)
Incoming interface	PE-CE interface closest to the RP and RPF neighbor CE router
Outgoing interfaces	Default mdt tunnel interface plus any additional native PE-CE interfaces, excluding the IIL
For (*,224.0.1.40)	RP 0.0.0.0, flags: DCL
Incoming interface	Null, RPF neighbor 0.0.0.0
Outgoing interfaces	Default mdt tunnel interface plus all native PE-CE interfaces
For (S,224.0.1.40)	Flags: LT
Incoming interface	PE-CE interface closest to the RP and RPF neighbor CE router
Outgoing interfaces	Default mdt tunnel interface plus any additional native PE-CE interfaces, excluding the IIL

Sample Output for Checks on the PE Located Close to the Customer RP, with AutoRP in the Customer Network

The following example shows output from PE$_{C\text{-}RP}$ for mroutes of AutoRP groups:

```
(*, 224.0.1.39), 00:14:20/stopped, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:14:20/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:14:20/00:00:00

(192.168.100.202, 224.0.1.39), 00:02:21/00:01:09, flags: PT
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Prune/Sparse-Dense, 00:02:15/00:01:08

PE2#show ip mroute vrf RED 224.0.1.40 | begin 224
(*, 224.0.1.40), 00:15:59/stopped, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:15:59/00:00:00
    Tunnel0, Forward/Sparse-Dense, 00:15:59/00:00:00
```

```
(192.168.100.202, 224.0.1.40), 00:15:06/00:02:48, flags: LT
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:15:06/00:00:00
```

Verifying Customer-Facing Settings on the PE$_{C-SRC}$

The flowchart in Figure 2-13 walks through the basic checks on the PE$_{C-SRC.}$

Figure 2-13 Basic Check on PE_{C-SRC}

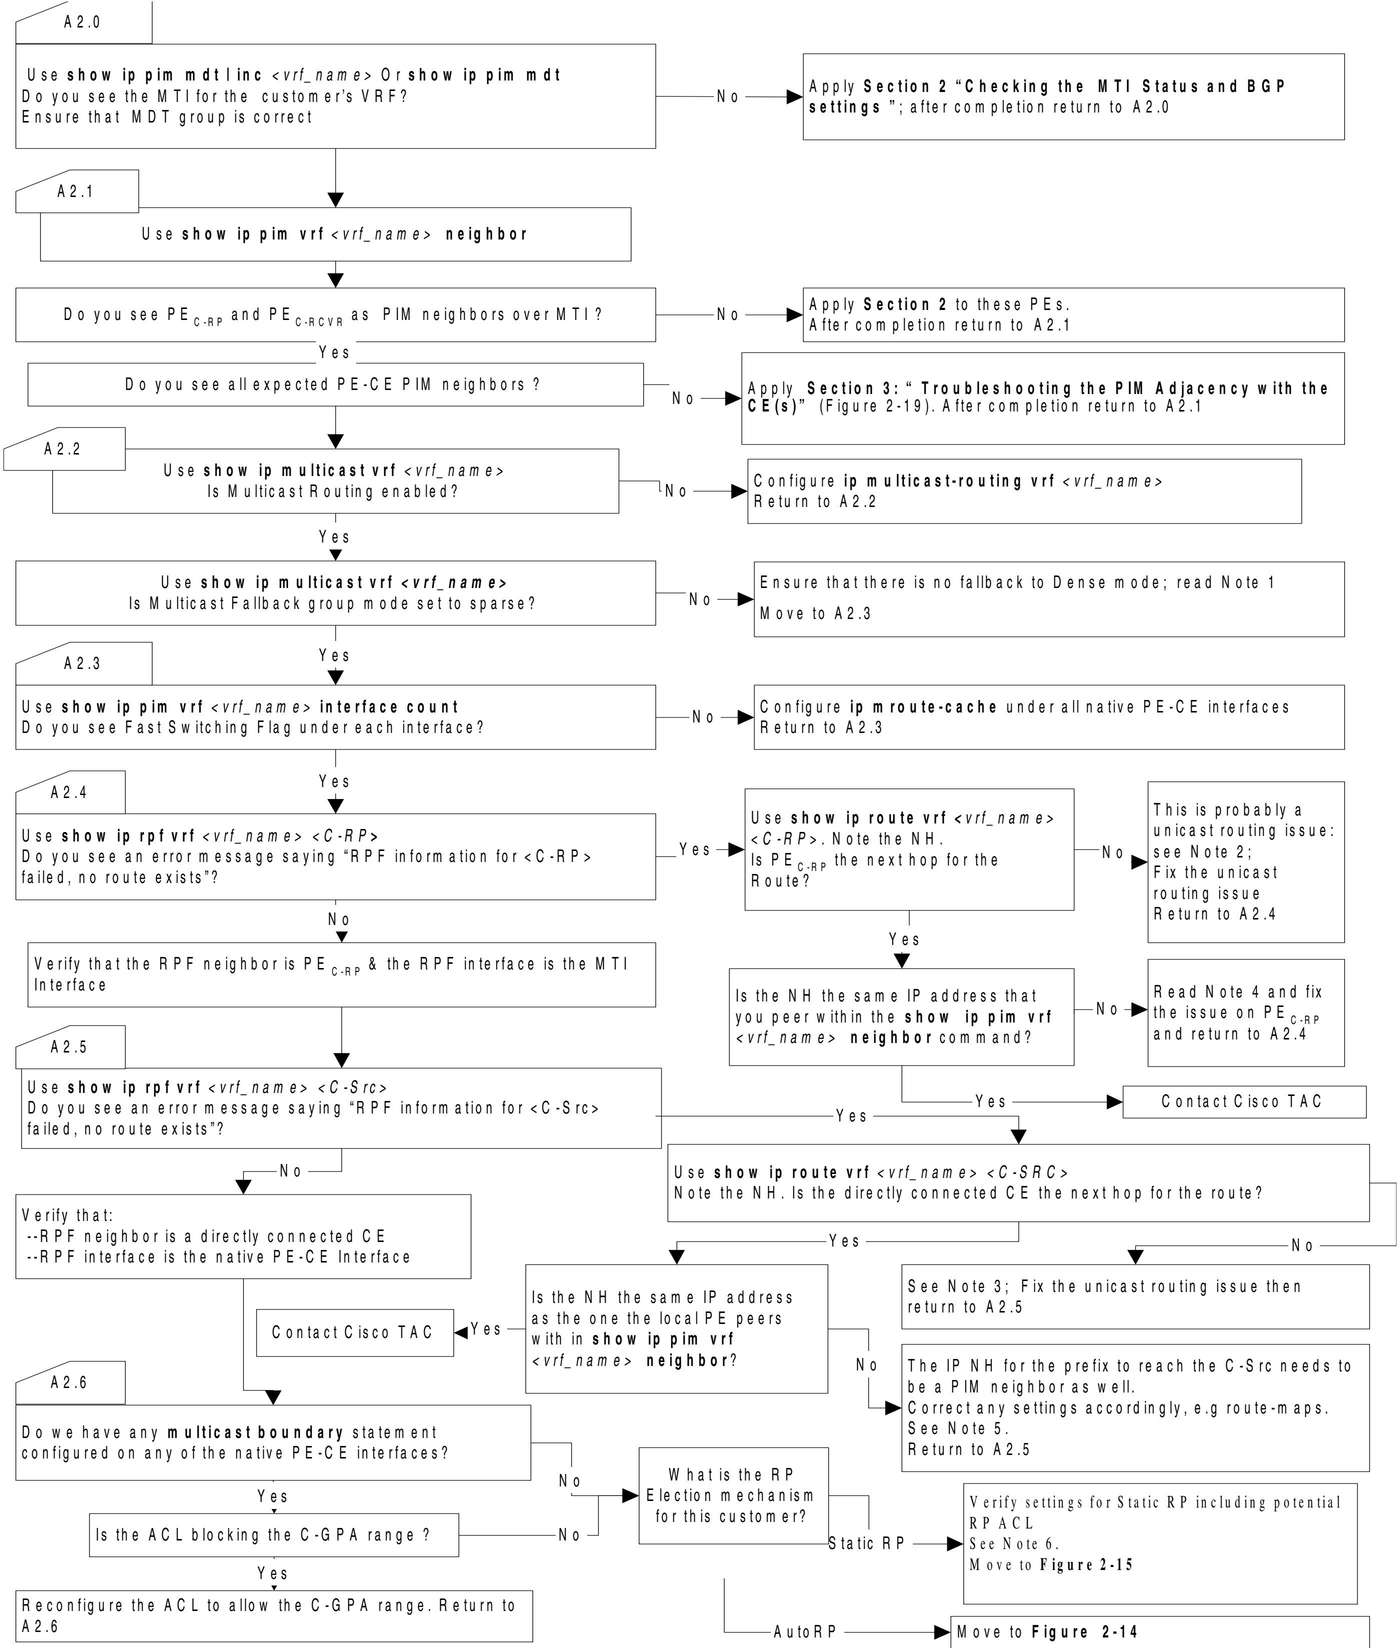

A2.0

Use **show ip pim mdt | inc** *<vrf_name>* Or **show ip pim mdt**
Do you see the MTI for the customer's VRF?
Ensure that MDT group is correct

Apply **Section 2 "Checking the MTI Status and BGP settings**"; after completion return to A2.0

A2.1

Use **show ip pim vrf** *<vrf_name>* **neighbor**

Do you see PE_{C-RP} and PE_{C-RCVR} as PIM neighbors over MTI?

Apply **Section 2** to these PEs.
After completion return to A2.1

Do you see all expected PE-CE PIM neighbors ?

Apply **Section 3: "Troubleshooting the PIM Adjacency with the CE(s)"** (Figure 2-19). After completion return to A2.1

A2.2

Use **show ip multicast vrf** *<vrf_name>*
Is Multicast Routing enabled?

Configure **ip multicast-routing vrf** *<vrf_name>*
Return to A2.2

Use **show ip multicast vrf *<vrf_name>***
Is Multicast Fallback group mode set to sparse?

Ensure that there is no fallback to Dense mode; read Note 1

Move to A2.3

A2.3

Use **show ip pim vrf** *<vrf_name>* **interface count**
Do you see Fast Switching Flag under each interface?

Configure **ip mroute-cache** under all native PE-CE interfaces
Return to A2.3

A2.4

Use **show ip rpf vrf** *<vrf_name> <C-RP>*
Do you see an error message saying "RPF information for <C-RP> failed, no route exists"?

Use **show ip route vrf** *<vrf_name>* *<C-RP>*. Note the NH.
Is PE_{C-RP} the next hop for the Route?

This is probably a unicast routing issue: see Note 2;
Fix the unicast routing issue
Return to A2.4

Verify that the RPF neighbor is PE_{C-RP} & the RPF interface is the MTI Interface

Is the NH the same IP address that you peer with in the **show ip pim vrf** *<vrf_name>* **neighbor** command?

Read Note 4 and fix the issue on PE_{C-RP} and return to A2.4

Contact Cisco TAC

A2.5

Use **show ip rpf vrf** *<vrf_name> <C-Src>*
Do you see an error message saying "RPF information for <C-Src> failed, no route exists"?

Use **show ip route vrf** *<vrf_name> <C-SRC>*
Note the NH. Is the directly connected CE the next hop for the route?

Verify that:
--RPF neighbor is a directly connected CE
--RPF interface is the native PE-CE Interface

Is the NH the same IP address as the one the local PE peers with in **show ip pim vrf** *<vrf_name>* **neighbor**?

Contact Cisco TAC

See Note 3; Fix the unicast routing issue then return to A2.5

The IP NH for the prefix to reach the C-Src needs to be a PIM neighbor as well.
Correct any settings accordingly, e.g route-maps.
See Note 5.
Return to A2.5

A2.6

Do we have any **multicast boundary** statement configured on any of the native PE-CE interfaces?

What is the RP Election mechanism for this customer?

Verify settings for Static RP including potential RP ACL
See Note 6.
Move to **Figure 2-15**

Is the ACL blocking the C-GPA range ?

Reconfigure the ACL to allow the C-GPA range. Return to A2.6

Move to **Figure 2-14**

Notes for Figure 2-13

NH = next hop.

Note 1

- Generally, service providers prefer to have customer VRFs Fallback group mode set to Sparse.
- Review the document at http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_4/gtautorp.htm
- Move to A2.3.

Note 2

If you do not see the prefix, or the next hop is a node other then the PE$_{C\text{-}RP}$, then it is a unicast routing issue.

Note 3

If you do not see the prefix, or the next hop is a node other than the directly connected CE, then it is a unicast routing issue.

Note 4

- If the NH is one of the ip addresses on PE$_{C\text{-}RP}$ that is not the same ip address as the PIM neighbor over MTI Interface, then mVPN will break.
- All PEs must advertise prefixes with BGP next hop, which is the tunnel source loopback ip address.

Note 5

If there are any fancy route maps that modify the next hop on either the CE or the PE$_{C\text{-}SRC}$, modify and correct them.

Note 6

- If using PIM SM with Static RP then verify that **ip pim rp-address vrf** *vrf_name C-RP-address* [**group-list ACL**] [**override**] exists on the PE.
- Verify that the referenced ACL is correct.

Sample Output for Checks on the PE-Facing the Customer Source

The following example shows output from PE$_{C\text{-}SRC}$ for basic checks:

```
PE1#show ip multicast vrf RED
  Multicast Routing: enabled
  Multicast Multipath: disabled
  Multicast Route limit: No limit
  Multicast Triggered RPF check: enabled
  Multicast Fallback group mode: Sparse

PE1#show ip pim mdt | inc MDT|RED
  MDT Group        Interface    Source                   VRF
* 239.0.0.10       Tunnel7      Loopback0                RED

PE1#show ip pim vrf RED neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.11.2      Serial1/0                03:56:44/00:01:41 v2    1 / P
```

```
172.16.100.2      Tunnel7                  03:59:21/00:01:35 v2    1 / P
172.16.100.3      Tunnel7                  03:59:21/00:01:25 v2    1 / DR P

PE1#show ip pim vrf RED interface S1/0 detail
Serial1/0 is up, line protocol is up
  Internet address is 192.168.11.1/24
  Multicast switching: fast
  Multicast packets in/out: 1967476/1928
  Multicast boundary: not set
  Multicast TTL threshold: 0
  PIM: enabled
    PIM version: 2, mode: sparse
    PIM DR: 0.0.0.0
    PIM neighbor count: 1
    PIM Hello/Query interval: 30 seconds
    PIM NBMA mode: disabled
    PIM ATM multipoint signalling: disabled
    PIM domain border: disabled
    PIM neighbors rpf proxy capable: TRUE
    Multicast Tagswitching: disabled

PE1#show ip rpf vrf RED 192.168.100.202   → C-RP
RPF information for ? (192.168.100.202)
  RPF interface: Tunnel7
  RPF neighbor: ? (172.16.100.2)
  RPF route/mask: 192.168.100.202/32
  RPF type: unicast (bgp 1)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
  BGP originator: 172.16.100.2

PE1#show ip rpf vrf RED 192.168.1.1     →C-Src
RPF information for ? (192.168.1.1)
  RPF interface: Serial1/0
  RPF neighbor: ? (192.168.11.2)
  RPF route/mask: 192.168.1.0/24
  RPF type: unicast (bgp 1)
  RPF recursion count: 1
  Doing distance-preferred lookups across tables
```

After the above is verified, check the RP mapping on the $PE_{C\text{-}SRC.}$ Also, verify that the ACL referenced is not misconfigured.

- Static RP Mapping

  If the customer is running static RP, verify that all the PEs in the mVRF are configured with the correct static RP:

  ip pim rp-address *rp-address* [**group_acl**] [**override**]

  It is recommended to use the **override** keyword to override the dynamically learned RP mappings. Also, verify that the value specified by *group_acl* is correctly configured in the ACL and correctly referenced in the above command.

- AutoRP Mapping Using Cisco AutoRP

  For AutoRP, verify the following:

  – $PE_{C\text{-}SRC}$ learns the RP mapping from the $PE_{C\text{-}RP}$ over the MTI

  – $PE_{C\text{-}SRC}$ advertises the RP mapping to the directly connected CEs

  The flowchart in Figure 2-14 walks through the essential configurations for AutoRP on the $PE_{C\text{-}SRC.}$

Figure 2-14 AutoRP Troubleshooting on $PE_{C\text{-}SRC}$

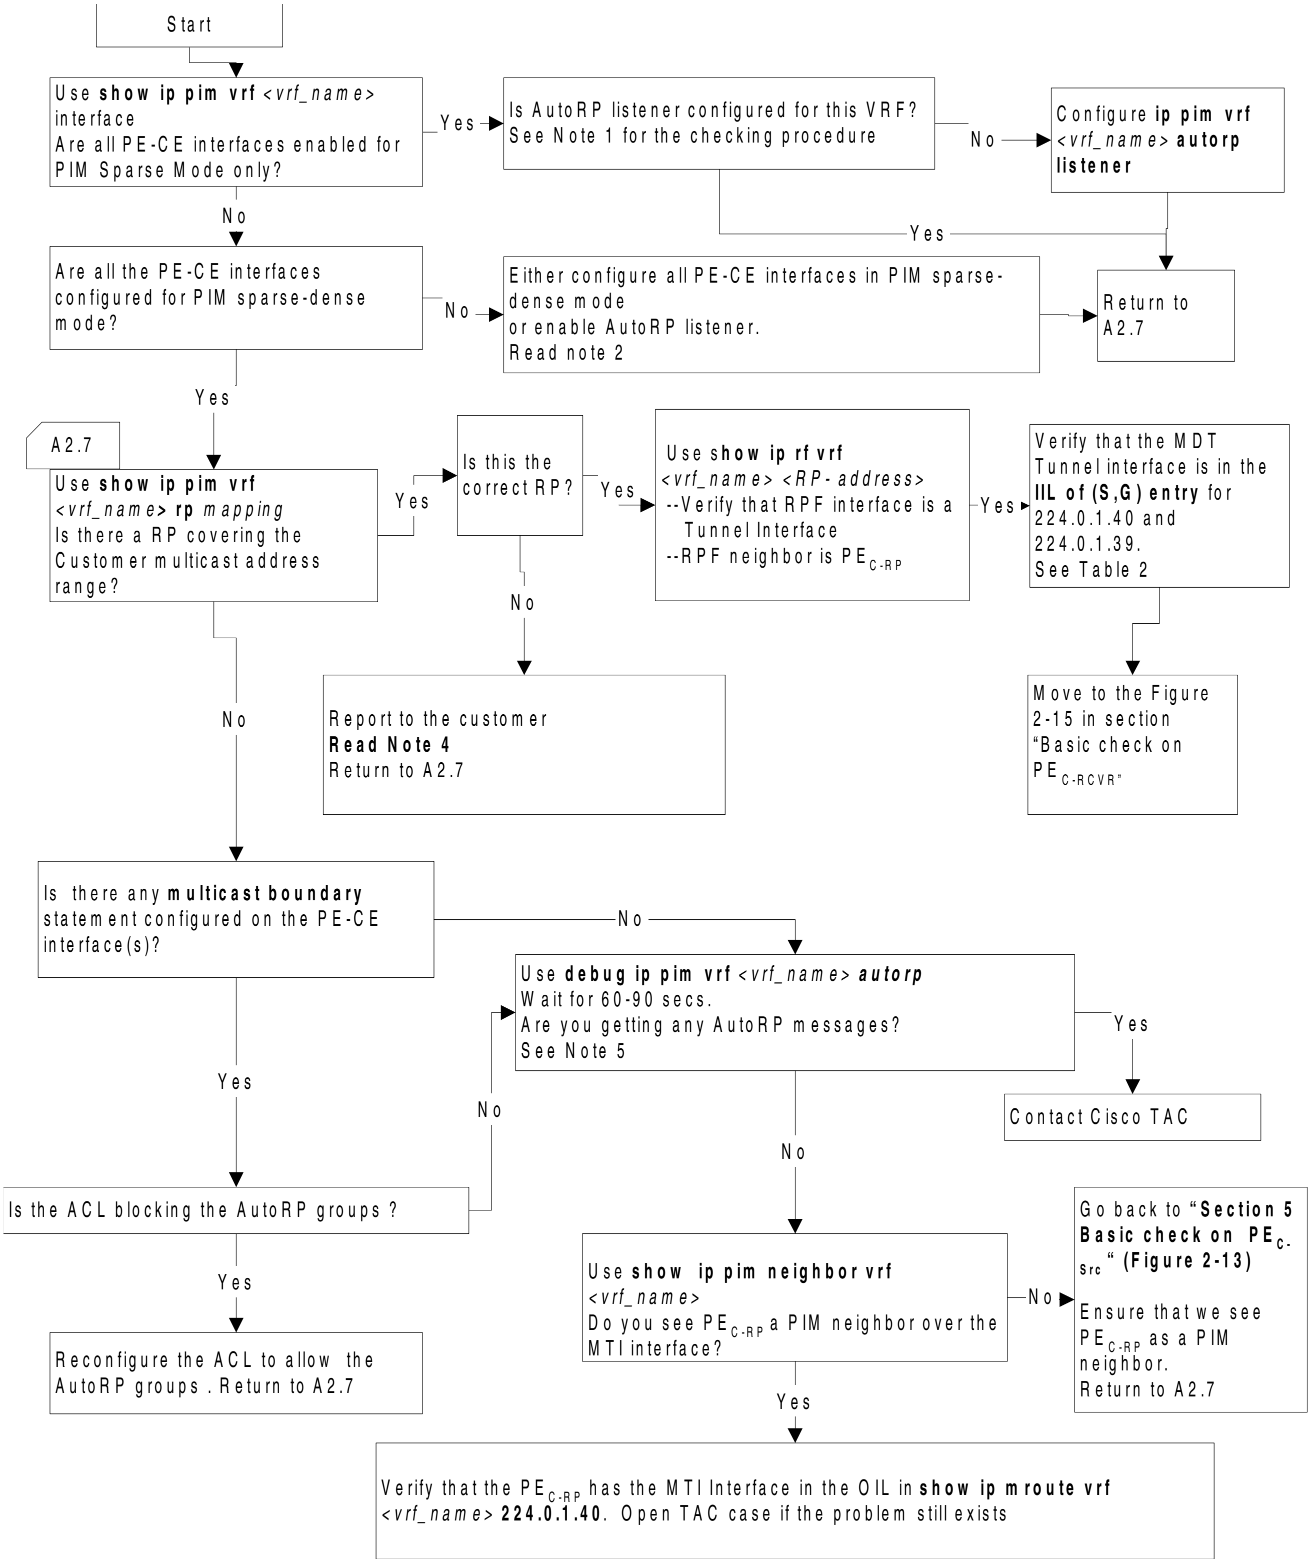
Start
Use **show ip pim vrf** <*vrf_name*> interface
Are all PE-CE interfaces enabled for PIM Sparse Mode only?
Yes
Is AutoRP listener configured for this VRF? See Note 1 for the checking procedure
No
Configure **ip pim vrf** <*vrf_name*> **autorp listener**
Yes
No
Are all the PE-CE interfaces configured for PIM sparse-dense mode?
No
Either configure all PE-CE interfaces in PIM sparse-dense mode
or enable AutoRP listener.
Read note 2
Return to A2.7
Yes
A2.7
Use **show ip pim vrf** <*vrf_name*> **rp** *mapping*
Is there a RP covering the Customer multicast address range?
Yes
Is this the correct RP?
Yes
Use **show ip rf vrf** <*vrf_name*> <*RP-address*>
--Verify that RPF interface is a Tunnel Interface
--RPF neighbor is PE C-RP
Yes
Verify that the MDT Tunnel interface is in the **IIL of (S,G) entry** for 224.0.1.40 and 224.0.1.39.
See Table 2
No
Report to the customer
Read Note 4
Return to A2.7
Move to the Figure 2-15 in section "Basic check on PE C-RCVR"
No
Is there any **multicast boundary** statement configured on the PE-CE interface(s)?
No
Use **debug ip pim vrf** <*vrf_name*> ***autorp***
Wait for 60-90 secs.
Are you getting any AutoRP messages?
See Note 5
Yes
Contact Cisco TAC
Yes
No
No
Is the ACL blocking the AutoRP groups ?
Yes
Reconfigure the ACL to allow the AutoRP groups . Return to A2.7
Use **show ip pim neighbor vrf** <*vrf_name*>
Do you see PE C-RP a PIM neighbor over the MTI interface?
No
Go back to "**Section 5 Basic check on PE C-Src**" **(Figure 2-13)**
Ensure that we see PE C-RP as a PIM neighbor.
Return to A2.7
Yes
Verify that the PE C-RP has the MTI Interface in the OIL in **show ip mroute vrf** <*vrf_name*> **224.0.1.40**. Open TAC case if the problem still exists

Notes for Figure 2-14

Note 1

- Do you see "AutoRP groups over sparse mode interface is enabled" under **show ip pim vrf** *vrf_name* **autorp**?

  OR

- Does the **show running-config | inc autorp list** command display "ip pim vrf *vrf_name* autorp listener?"

Note 2

- Enable AutoRP listener to densely switch the AutoRP groups to allow AutoRP to work successfully. AutoRP is the recommended option. Also, configure an RP of last resort to avoid all customer traffic getting densely switched.
- Read the comment on AutoRP at:

http://www.cisco.com/warp/public/732/Tech/multicast/docs/multicastdesign.pdf

Note 3

For a complete verification of the flags and interface lists, refer to Table 2-2.

Note 4

Ensure that they advertise the correct RP to group mapping and that it is the correct RP address. We might be seeing messages like %PIM-6-INVALID_RP_JOIN in the logs.

Note 5

Look for log messages in the buffer (using the **show logging** command) or check the console.

Table 2-2 lists the Mroute states for the AutoRP groups on the $PE_{C\text{-}SRC}$.

Table 2-2 VRF Mroute states on $PE_{C\text{-}SRC}$ for AutoRP groups

For (*,224.0.1.39)	RP 0.0.0.0, Flags = D
Incoming interface	Null, RPF neighbor 0.0.0.0
Outgoing interfaces	Default MDT tunnel interface plus all native PE-CE interfaces
For (S,224.0.1.39)	Flags PT (Flag P is there because the assumption is that the customer mapping agent and RP candidate are on the same site)
Incoming interface	Default mdt tunnel interface and RPF neighbor $PE_{C\text{-}RP}$
Outgoing interfaces	All native PE-CE interfaces
For (S,224.0.1.40)	Flags: LT

Table 2-2 VRF Mroute states on PE$_{C-SRC}$ for AutoRP groups

Incoming interface	Default mdt tunnel interface and RPF neighbor PE$_{C-RP}$
Outgoing interfaces	All native PE-CE interfaces
For (*,224.0.1.40)	RP 0.0.0.0, flags: DCL
Incoming interface	Null, RPF neighbor 0.0.0.0
Outgoing interfaces	Default mdt tunnel interface plus all native PE-CE interfaces

Sample Output on the PE-Facing the Customer Source, with AutoRP in the Customer Network

The following example shows output from PE$_{C-SRC}$ for mroutes of AutoRP groups:

```
PE1#show ip mroute vrf RED 224.0.1.39 | be 224
(*, 224.0.1.39), 04:54:14/stopped, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 04:52:19/00:00:00
    Tunnel7, Forward/Sparse-Dense, 04:54:14/00:00:00

(192.168.100.202, 224.0.1.39), 00:00:14/00:02:54, flags: PT
  Incoming interface: Tunnel7, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Prune/Sparse, 00:00:14/00:02:48

PE1#show ip mroute vrf RED 224.0.1.40 | be 224
(*, 224.0.1.40), 2d08h/stopped, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 13:54:09/00:00:00
    Tunnel7, Forward/Sparse-Dense, 04:55:07/00:00:00

(192.168.100.202, 224.0.1.40), 04:55:02/00:02:18, flags: LT
  Incoming interface: Tunnel7, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 04:55:02/00:00:00
```

Verifying Customer-Facing Settings on the PE$_{C-RCVR}$

The flow chart in Figure 2-15 walks through the basic checks on the PE$_{C-RCVR.}$ Repeat this procedure for each PE$_{C-RCVR.}$

Figure 2-15 Troubleshooting Guidelines for $PE_{C\text{-}RCVR}$.

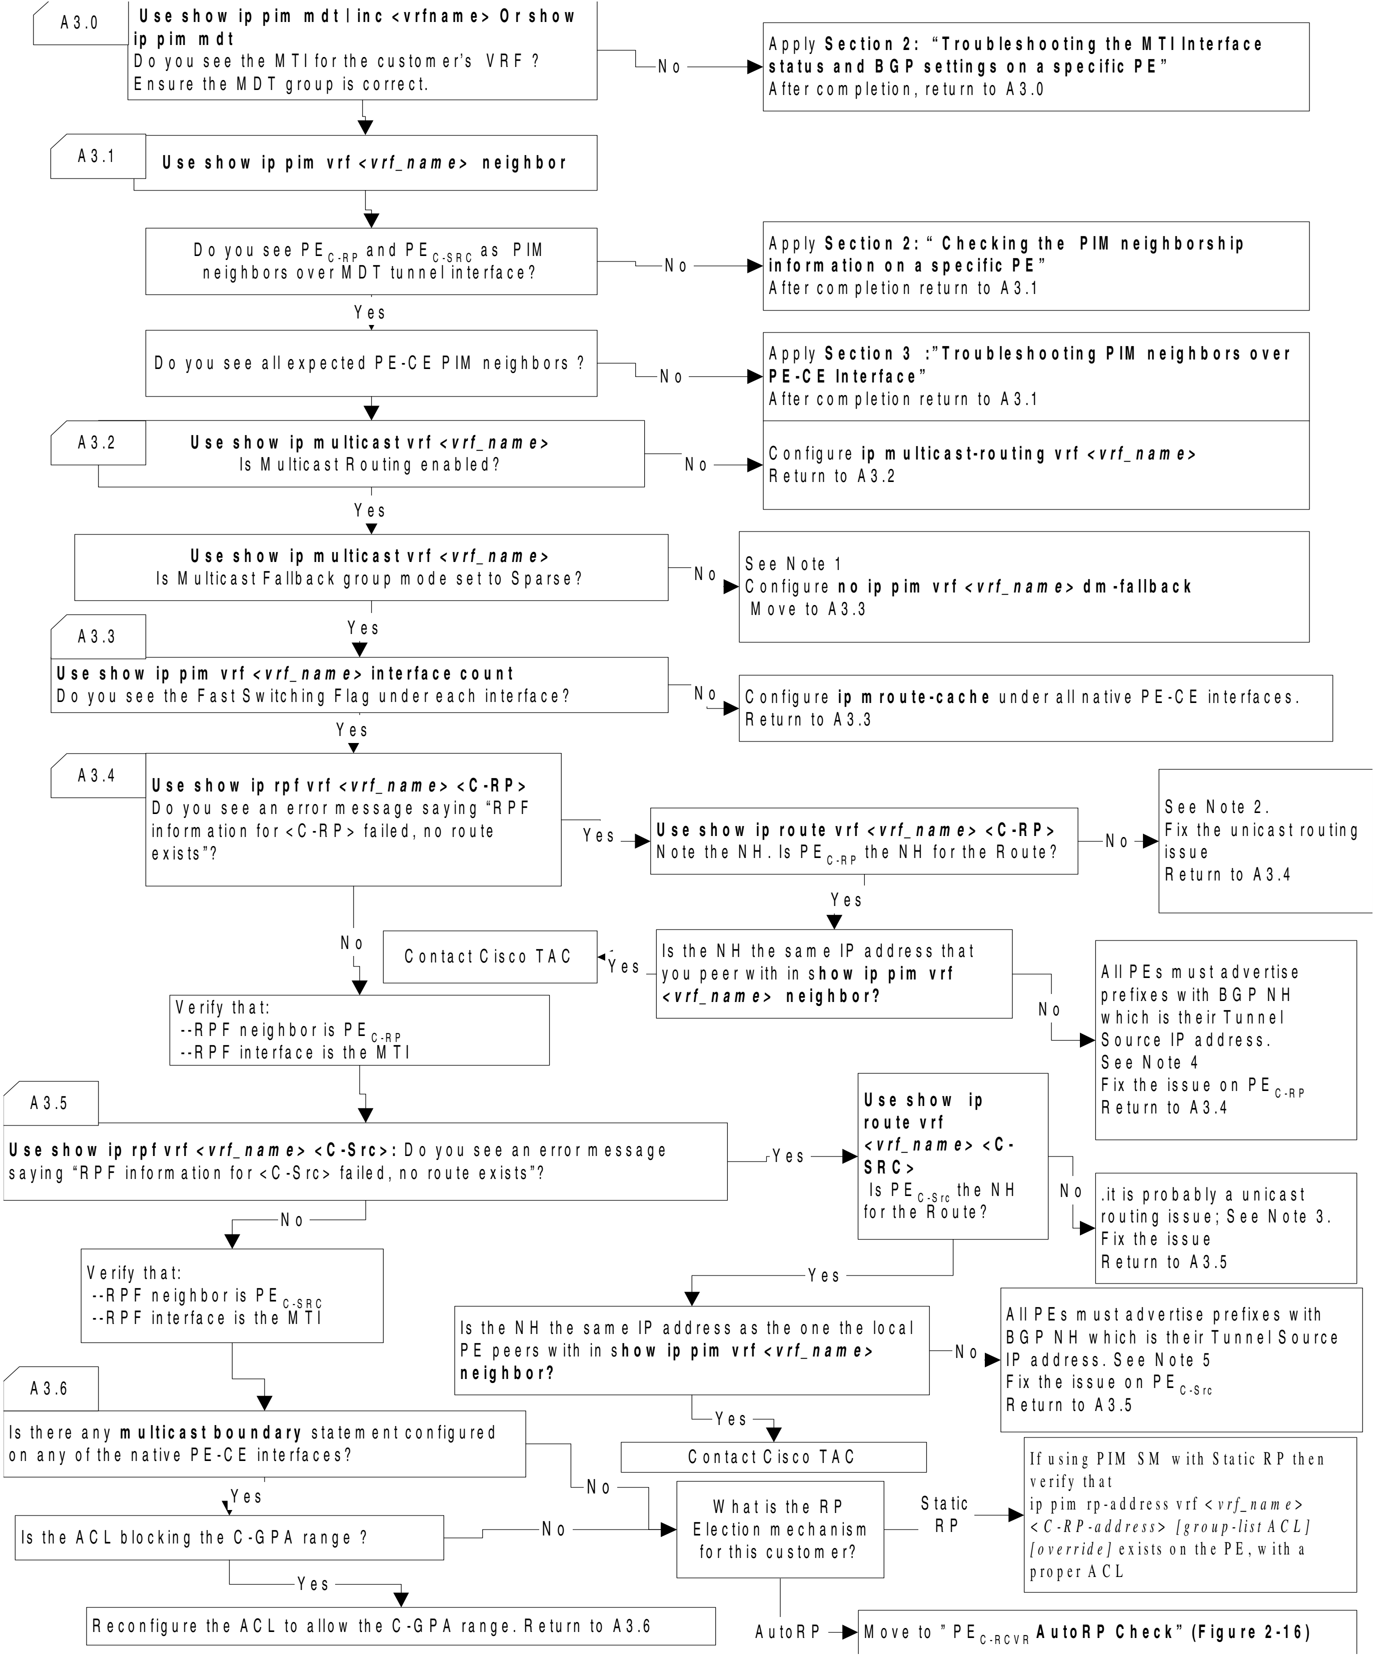
A3.0
Use show ip pim mdt | inc <vrfname> Or show ip pim mdt
Do you see the MTI for the customer's VRF ? Ensure the MDT group is correct.
No
Apply Section 2: "Troubleshooting the MTI Interface status and BGP settings on a specific PE"
After completion, return to A3.0
A3.1
Use show ip pim vrf <vrf_name> neighbor
Do you see PE C-RP and PE C-SRC as PIM neighbors over MDT tunnel interface?
No
Apply Section 2: " Checking the PIM neighborship information on a specific PE"
After completion return to A3.1
Yes
Do you see all expected PE-CE PIM neighbors ?
No
Apply Section 3 :"Troubleshooting PIM neighbors over PE-CE Interface"
After completion return to A3.1
A3.2
Use show ip multicast vrf <vrf_name>
Is Multicast Routing enabled?
No
Configure ip multicast-routing vrf <vrf_name>
Return to A3.2
Yes
Use show ip multicast vrf <vrf_name>
Is Multicast Fallback group mode set to Sparse?
No
See Note 1
Configure no ip pim vrf <vrf_name> dm-fallback
Move to A3.3
Yes
A3.3
Use show ip pim vrf <vrf_name> interface count
Do you see the Fast Switching Flag under each interface?
No
Configure ip mroute-cache under all native PE-CE interfaces.
Return to A3.3
Yes
A3.4
Use show ip rpf vrf <vrf_name> <C-RP>
Do you see an error message saying "RPF information for <C-RP> failed, no route exists"?
Yes
Use show ip route vrf <vrf_name> <C-RP>
Note the NH. Is PE C-RP the NH for the Route?
No
See Note 2.
Fix the unicast routing issue
Return to A3.4
Yes
Is the NH the same IP address that you peer with in show ip pim vrf <vrf_name> neighbor?
Yes
Contact Cisco TAC
No
All PEs must advertise prefixes with BGP NH which is their Tunnel Source IP address.
See Note 4
Fix the issue on PE C-RP
Return to A3.4
No
Verify that:
--RPF neighbor is PE C-RP
--RPF interface is the MTI
A3.5
Use show ip rpf vrf <vrf_name> <C-Src>: Do you see an error message saying "RPF information for <C-Src> failed, no route exists"?
Yes
Use show ip route vrf <vrf_name> <C-SRC>
Is PE C-Src the NH for the Route?
No
.it is probably a unicast routing issue; See Note 3.
Fix the issue
Return to A3.5
Yes
Is the NH the same IP address as the one the local PE peers with in show ip pim vrf <vrf_name> neighbor?
No
All PEs must advertise prefixes with BGP NH which is their Tunnel Source IP address. See Note 5
Fix the issue on PE C-Src
Return to A3.5
Yes
Contact Cisco TAC
No
Verify that:
--RPF neighbor is PE C-SRC
--RPF interface is the MTI
A3.6
Is there any multicast boundary statement configured on any of the native PE-CE interfaces?
No
What is the RP Election mechanism for this customer?
Static RP
If using PIM SM with Static RP then verify that ip pim rp-address vrf <vrf_name> <C-RP-address> [group-list ACL] [override] exists on the PE, with a proper ACL
AutoRP
Move to "PE C-RCVR AutoRP Check" (Figure 2-16)
Yes
Is the ACL blocking the C-GPA range ?
No
Yes
Reconfigure the ACL to allow the C-GPA range. Return to A3.6

Notes for Figure 2-15

NH = Next hop.

Note 1

Generally, service providers prefer to have customer VRFs Fallback group mode to Sparse

- Read the document at

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123newft/123t/123t_4/gtautorp.htm
- Move to A3.3.

Note 2

If you do not see the prefix or the next hop is a node other then the PE$_{\text{C-RP}}$ then it is a unicast routing issue

Note 3

If you do not see the prefix, or the next hop is a node other than the PE$_{\text{C-SRC}}$, then it is a unicast routing issue.

Note 4

If the NH (from the **show ip route vrf** *vrf_name C-SRC* command) is one of the IP addresses on PE$_{\text{C-RP}}$ that is not the same IP address as the PIM neighbor over the MTI Interface, then mVPN will break.

Note 5

If the NH (from the **show ip route vrf** *vrf_name C-SRC* command) is one of the IP addresses on PE$_{\text{C-SRC}}$ that is not the same ip address as the PIM neighbor over MTI Interface, then mVPN will break.

Sample Output for the PEs Facing one or more Customer Receivers

The following example shows output from PE$_{\text{C-RCVR}}$ for basic checks:

```
PE3#show ip multicast vrf RED
  Multicast Routing: enabled
  Multicast Multipath: disabled
  Multicast Route limit: No limit
  Multicast Triggered RPF check: enabled
  Multicast Fallback group mode: Sparse

PE3#show ip pim mdt | inc MDT|RED
  MDT Group       Interface   Source                   VRF
* 239.0.0.10      Tunnel0     Loopback0                RED

PE3#show ip pim vrf RED neighbor
PIM Neighbor Table
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Priority/Mode
192.168.33.2      Serial1/0                2d09h/00:01:34    v2    1 / P
172.16.100.1      Tunnel0                  04:56:48/00:01:16 v2    1 / P
172.16.100.2      Tunnel0                  2d09h/00:01:19    v2    1 / P

PE3#show ip pim vrf RED interface Serial 1/0 detail
Serial1/0 is up, line protocol is up
  Internet address is 192.168.33.1/24
  Multicast switching: fast
  Multicast packets in/out: 0/344922
  Multicast boundary: not set
  Multicast TTL threshold: 0
  PIM: enabled
    PIM version: 2, mode: sparse
    PIM DR: 0.0.0.0
    PIM neighbor count: 1
    PIM Hello/Query interval: 30 seconds
    PIM NBMA mode: disabled
    PIM ATM multipoint signalling: disabled
    PIM domain border: disabled
    PIM neighbors rpf proxy capable: TRUE
  Multicast Tagswitching: disabled

PE3#show ip rpf vrf RED 192.168.100.202   → C-RP
RPF information for ? (192.168.100.202)
  RPF interface: Tunnel0
  RPF neighbor: ? (172.16.100.2)
  RPF route/mask: 192.168.100.202/32
  RPF type: unicast (bgp 1)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
  BGP originator: 172.16.100.2

PE3#show ip rpf vrf RED 192.168.1.1    →C-Src
RPF information for ? (192.168.1.1)
  RPF interface: Tunnel0
  RPF neighbor: ? (172.16.100.1)
  RPF route/mask: 192.168.1.0/24
  RPF type: unicast (bgp 1)
  RPF recursion count: 0
  Doing distance-preferred lookups across tables
  BGP originator: 172.16.100.1
```

After the above is verified, check the RP mapping on PE$_{\text{C-RCVR.}}$

- Static RP Mapping

  If the customer is running static RP, verify that all the PEs in the mVRF are configured with the correct static RP:

  ip pim rp-address *rp-address* [**group_acl**] [**override**]

  It is recommended to use the **override** keyword to override the dynamically learned RP mappings. Also verify that the value specified by *group_acl* is correctly configured in the ACL and correctly referenced in the above command.

- AutoRP Mapping using Cisco AutoRP

  For AutoRP, verify the following:

  – PE$_{\text{C-RCVR}}$ learns the RP mapping from the PE$_{\text{C-RP}}$ over the MTI
  – Advertises the RP mapping to the directly connected CEs

  The flowchart in Figure 2-16 walks through the essential AutoRP configurations on the PE$_{\text{C-RCVR.}}$

Figure 2-16 AutoRP Troubleshooting Guidelines on PE-$_{C-RVCR}$

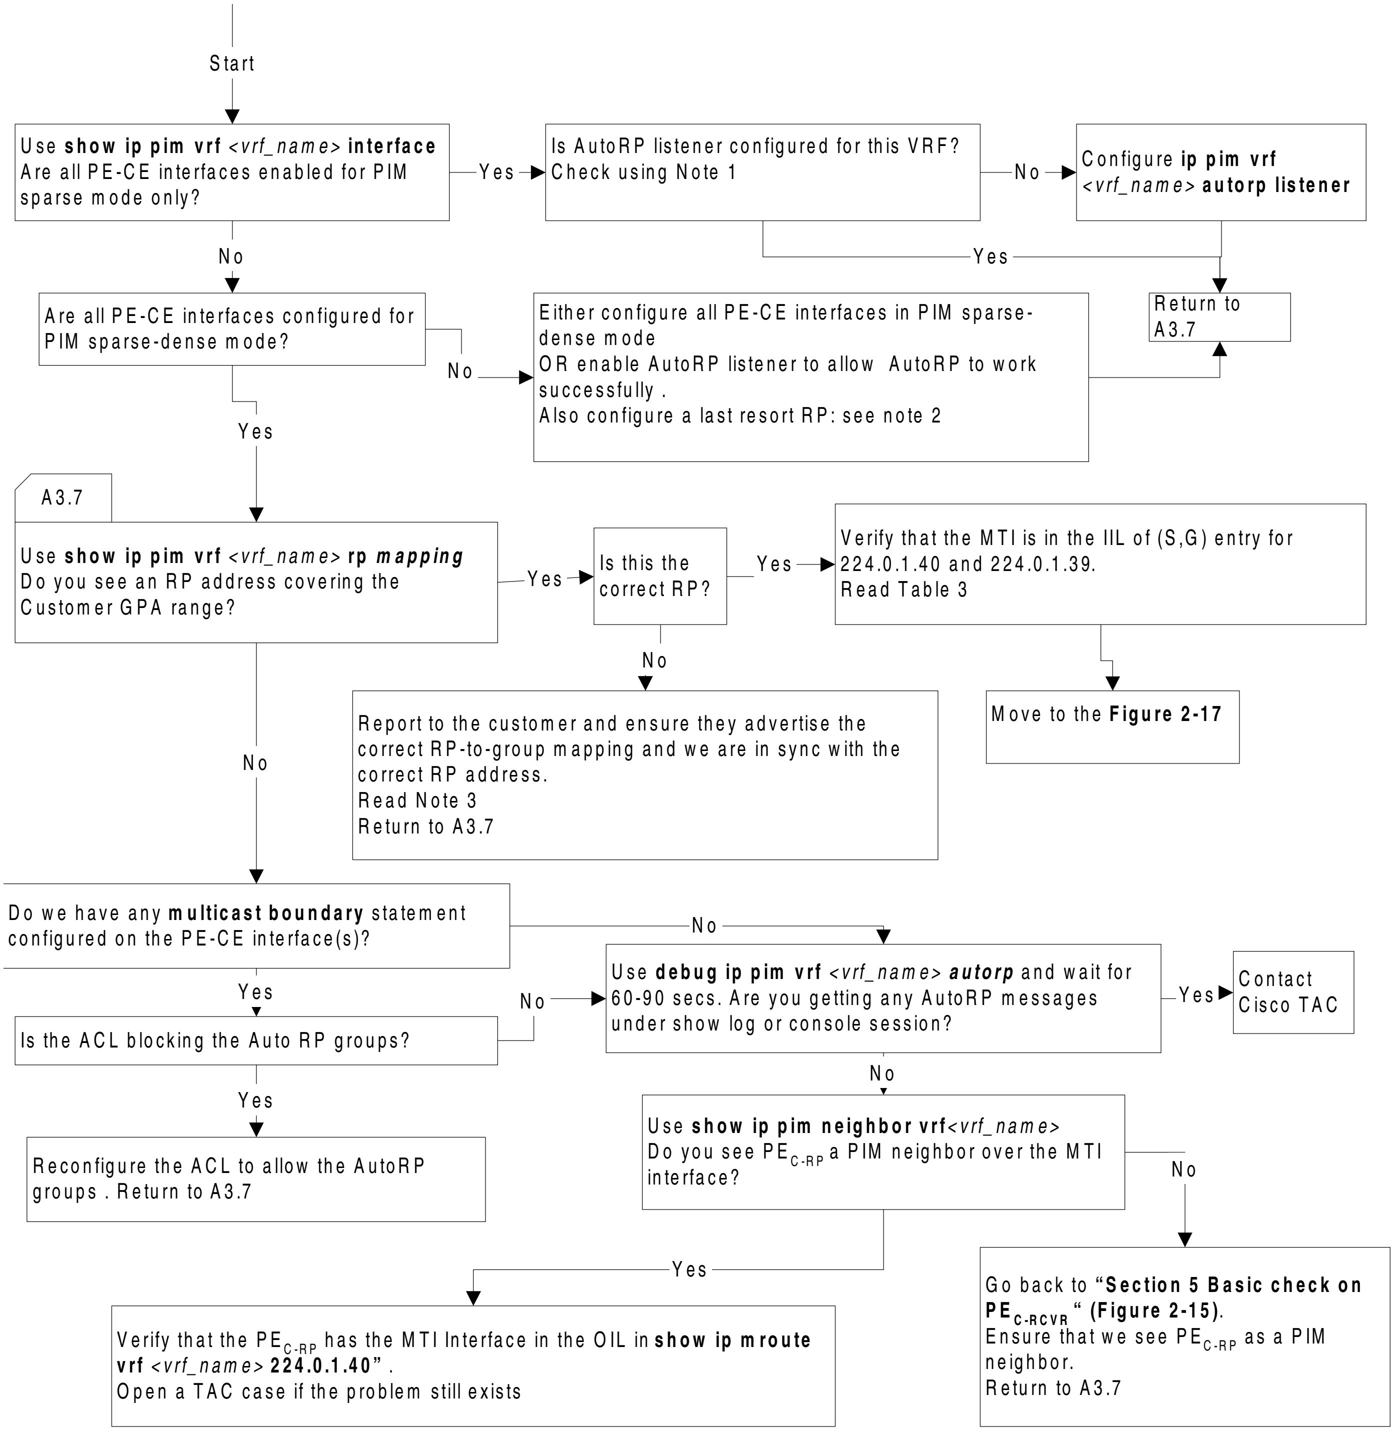

Use **show ip pim vrf** <*vrf_name*> **interface**
Are all PE-CE interfaces enabled for PIM sparse mode only?

Is AutoRP listener configured for this VRF?
Check using Note 1

Configure **ip pim vrf** <*vrf_name*> **autorp listener**

Return to A3.7

Are all PE-CE interfaces configured for PIM sparse-dense mode?

Either configure all PE-CE interfaces in PIM sparse-dense mode
OR enable AutoRP listener to allow AutoRP to work successfully .
Also configure a last resort RP: see note 2

Use **show ip pim vrf** <*vrf_name*> **rp *mapping***
Do you see an RP address covering the Customer GPA range?

Is this the correct RP?

Verify that the MTI is in the IIL of (S,G) entry for 224.0.1.40 and 224.0.1.39.
Read Table 3

Move to the **Figure 2-17**

Report to the customer and ensure they advertise the correct RP-to-group mapping and we are in sync with the correct RP address.
Read Note 3
Return to A3.7

Do we have any **multicast boundary** statement configured on the PE-CE interface(s)?

Is the ACL blocking the Auto RP groups?

Reconfigure the ACL to allow the AutoRP groups . Return to A3.7

Use **debug ip pim vrf** <*vrf_name*> ***autorp*** and wait for 60-90 secs. Are you getting any AutoRP messages under show log or console session?

Contact Cisco TAC

Use **show ip pim neighbor vrf**<*vrf_name*>
Do you see PE$_{C-RP}$ a PIM neighbor over the MTI interface?

Verify that the PE$_{C-RP}$ has the MTI Interface in the OIL in **show ip mroute vrf** <*vrf_name*> **224.0.1.40"** .
Open a TAC case if the problem still exists

Go back to **"Section 5 Basic check on PE$_{C-RCVR}$ " (Figure 2-15)**.
Ensure that we see PE$_{C-RP}$ as a PIM neighbor.
Return to A3.7

Notes for Figure 2-16

Note 1

- Do you see “AutoRP groups over sparse mode interface is enabled” under **show ip pim vrf** *vrf_name* **autorp**?

  OR

- Does **show running-config | inc autorp list** display “ip pim vrf *vrf_name* autorp listener?”

Note 2

- Last resort RP will avoid all customer traffic getting densely switched.
- Read the comment on Auto-RP at

http://www.cisco.com/warp/public/732/Tech/multicast/docs/multicastdesign.pdf

Note 3

We might be seeing messages like %PIM-6-INVALID_RP_JOIN in the logs.

Table 2-3 VRF Mroute states on PE$_{C\text{-}RCVR}$ for AutoRP groups

For (*,224.0.1.39)	RP 0.0.0.0, Flags = D
Incoming interface	Null, RPF neighbor 0.0.0.0
Outgoing interfaces	Default mdt tunnel interface plus all native PE-CE interfaces
For (S,224.0.1.39)	Flags PT (flag P is there because the assumption is that the customer mapping agent and RP candidate are on the same site)
Incoming interface	Default mdt tunnel interface and RPF neighbor PE$_{C\text{-}RP}$
Outgoing interfaces	All native PE-CE interfaces
For (*,224.0.1.40)	RP 0.0.0.0, flags: DCL
Incoming interface	Null, RPF neighbor 0.0.0.0
Outgoing interfaces	Default mdt tunnel interface plus all native PE-CE interfaces
For (S,224.0.1.40)	Flags: LT
Incoming interface	Default mdt tunnel interface and RPF neighbor PE$_{C\text{-}RP}$
Outgoing interfaces	All native PE-CE interfaces

Sample Output on the PE$_{C\text{-}RCVR}$ for Mroutes of AutoRP Groups

```
PE3#show ip mroute vrf RED 224.0.1.39 | be 224
(*, 224.0.1.39), 2d09h/stopped, RP 0.0.0.0, flags: D
```

```
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 2d09h/00:00:00
    Tunnel0, Forward/Sparse-Dense, 2d09h/00:00:00

(192.168.100.202, 224.0.1.39), 00:03:24/00:00:10, flags: PT
  Incoming interface: Tunnel0, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Prune/Sparse, 00:03:24/00:00:09

PE3#show ip mroute vrf RED 224.0.1.40 | be 224
(*, 224.0.1.40), 2d09h/stopped, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 2d09h/00:00:00
    Tunnel0, Forward/Sparse-Dense, 2d09h/00:00:00

(192.168.100.202, 224.0.1.40), 2d09h/00:02:50, flags: LT
  Incoming interface: Tunnel0, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 2d09h/00:00:00
```

Verifying Consistency of the Customer RP Settings

After the individual PEs have been checked, verify that the RP mapping is consistent across all PEs involved.

Verifying the Multicast Routes on the PEs in the VRF Context

This section describes how to check the customer-specific Mroute states on the PEs.

Here is a high level description of the steps:

- Start by checking the (*,C-GPA) entry on $\text{PE}_{\text{C-RP.}}$
- If the (*,C-GPA) entry does not exists on $\text{PE}_{\text{C-RP,}}$ check the $\text{PE}_{\text{C-RCVR}}$ to verify that it has received the PIM joins from the CE and has forwarded them to the $\text{PE}_{\text{C-RP}}$.
- After the (*,C-GPA) entry is verified with Table 2-4, check the (C-SRC, C-GPA) entry on $\text{PE}_{\text{C-RP}}$, $\text{PE}_{\text{C-SRC}}$ and all the $\text{PE}_{\text{C-RCVR}}$.

The flowchart in Figure 2-17 provides the step-by-step procedure for troubleshooting this aspect of the control plane.

Figure 2-17 Troubleshooting Guidelines for the Control Plane

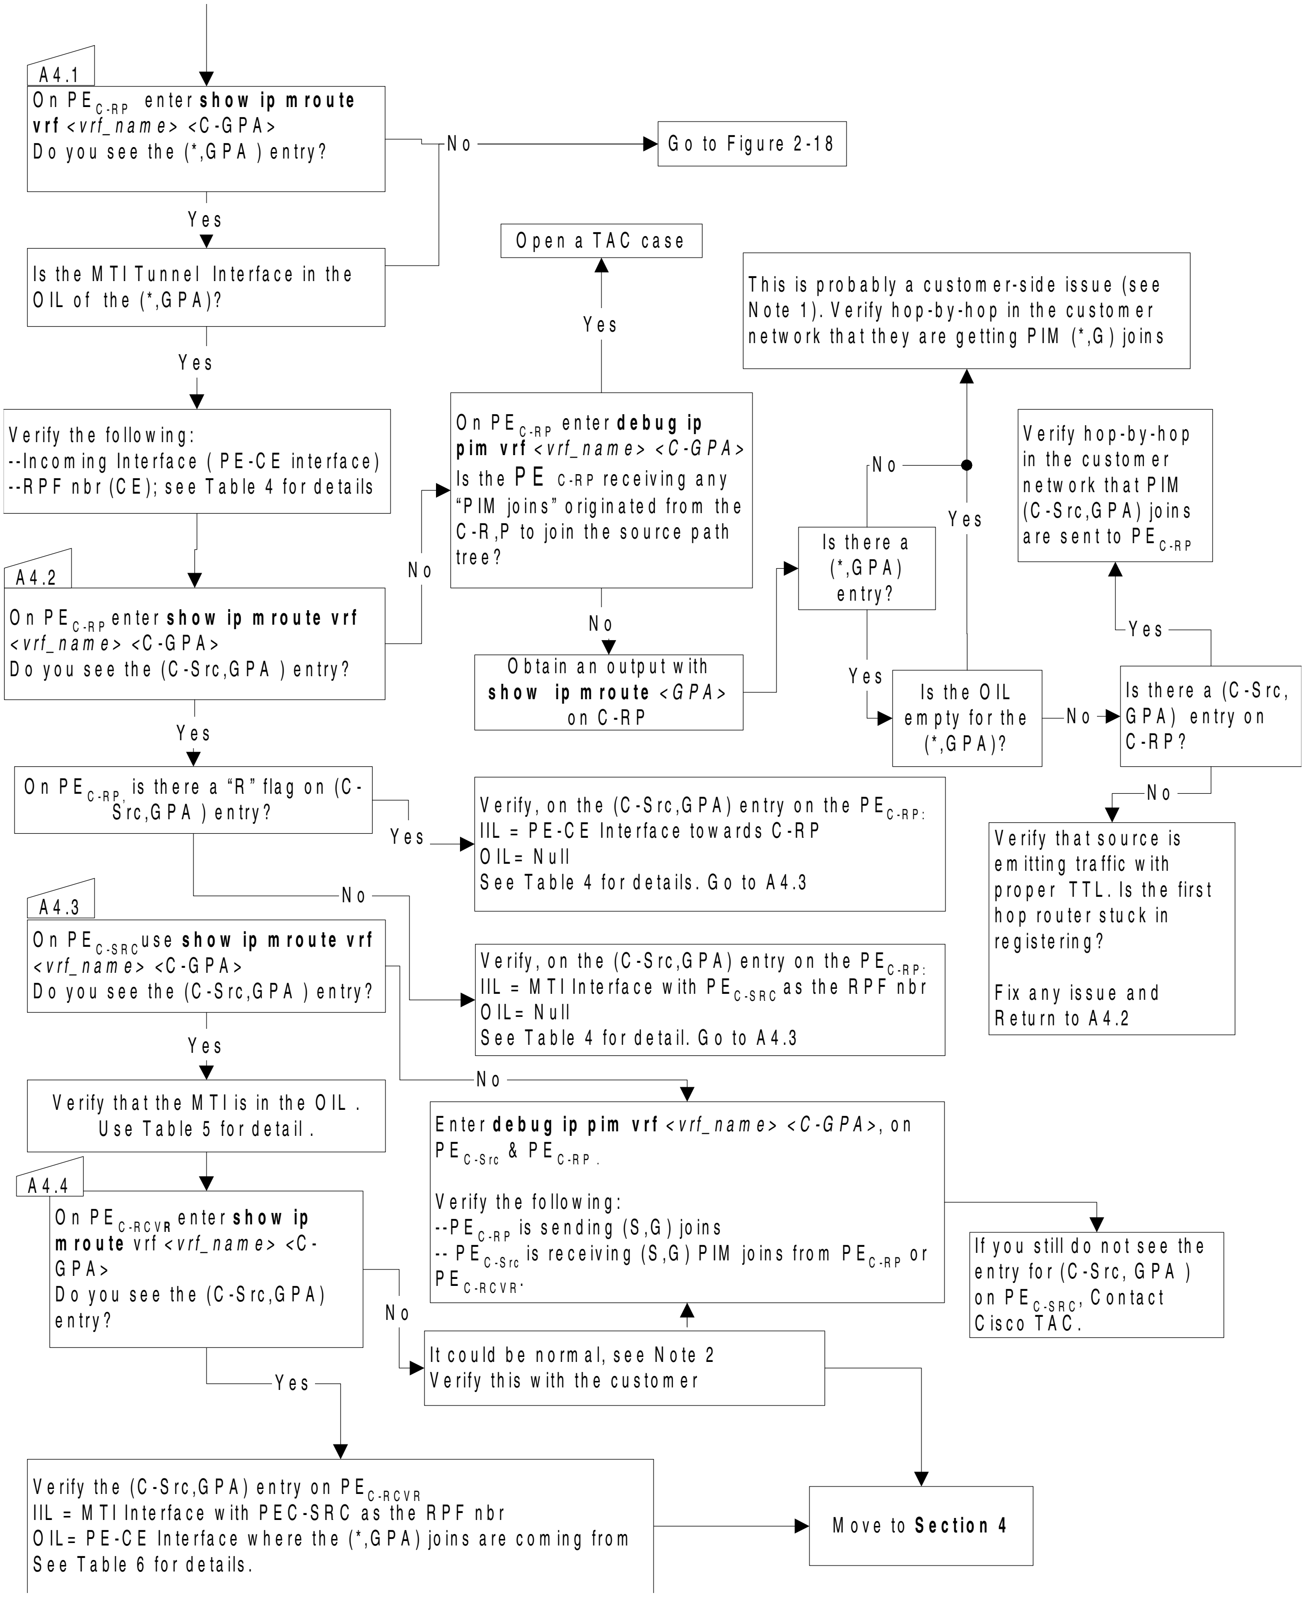
A4.1
On PE_{C-RP} enter **show ip mroute vrf** <*vrf_name*> <C-GPA> Do you see the (*,GPA) entry?
No
Go to Figure 2-18
Yes
Is the MTI Tunnel Interface in the OIL of the (*,GPA)?
Yes
Verify the following: --Incoming Interface (PE-CE interface) --RPF nbr (CE); see Table 4 for details
A4.2
On PE_{C-RP} enter **show ip mroute vrf** <*vrf_name*> <C-GPA> Do you see the (C-Src,GPA) entry?
No
Open a TAC case
Yes
On PE_{C-RP} enter **debug ip pim vrf** <*vrf_name*> <*C-GPA*> Is the PE $_{C-RP}$ receiving any "PIM joins" originated from the C-R,P to join the source path tree?
No
Obtain an output with **show ip mroute** <*GPA*> on C-RP
Is there a (*,GPA) entry?
No
Yes
This is probably a customer-side issue (see Note 1). Verify hop-by-hop in the customer network that they are getting PIM (*,G) joins
Yes
Is the OIL empty for the (*,GPA)?
No
Is there a (C-Src, GPA) entry on C-RP?
Yes
Verify hop-by-hop in the customer network that PIM (C-Src,GPA) joins are sent to PE_{C-RP}
No
Verify that source is emitting traffic with proper TTL. Is the first hop router stuck in registering? Fix any issue and Return to A4.2
Yes
On $PE_{C-RP,}$ is there a "R" flag on (C-Src,GPA) entry?
Yes
Verify, on the (C-Src,GPA) entry on the $PE_{C-RP:}$ IIL = PE-CE Interface towards C-RP OIL= Null See Table 4 for details. Go to A4.3
No
A4.3
On PE_{C-SRC}use **show ip mroute vrf** <*vrf_name*> <C-GPA> Do you see the (C-Src,GPA) entry?
Verify, on the (C-Src,GPA) entry on the $PE_{C-RP:}$ IIL = MTI Interface with PE_{C-SRC} as the RPF nbr OIL= Null See Table 4 for detail. Go to A4.3
No
Yes
Verify that the MTI is in the OIL . Use Table 5 for detail .
Enter **debug ip pim vrf** <*vrf_name*> <*C-GPA*>, on PE_{C-Src} & $PE_{C-RP.}$ Verify the following: --PE_{C-RP} is sending (S,G) joins -- PE_{C-Src} is receiving (S,G) PIM joins from PE_{C-RP} or PE_{C-RCVR}.
If you still do not see the entry for (C-Src, GPA) on PE_{C-SRC}, Contact Cisco TAC.
A4.4
On PE_{C-RCVR} enter **show ip mroute** vrf <*vrf_name*> <C-GPA> Do you see the (C-Src,GPA) entry?
No
It could be normal, see Note 2 Verify this with the customer
Yes
Verify the (C-Src,GPA) entry on PE_{C-RCVR} IIL = MTI Interface with PEC-SRC as the RPF nbr OIL= PE-CE Interface where the (*,GPA) joins are coming from See Table 6 for details.
Move to **Section 4**

Notes for Figure 2-17

Note 1

This is probably a customer issue. The C-RP has not received the PIM (*,GPA) join from the PE$_{\text{C-RP.}}$

Note 2

This would be normal if the last hop router (normally a C or CE router closest to the receiver IGMP host) was configured with a finite SPT threshold value, and the traffic received on the shared tree was below the threshold value.

Note In Figure 2-17, if (*,C-GPA) entry is not present on PE$_{\text{C-RP,}}$ the flowchart in Figure 2-18 can be used to further troubleshoot the control plane.

Figure 2-18 Troubleshooting Guidelines for (, GPA) Entry*

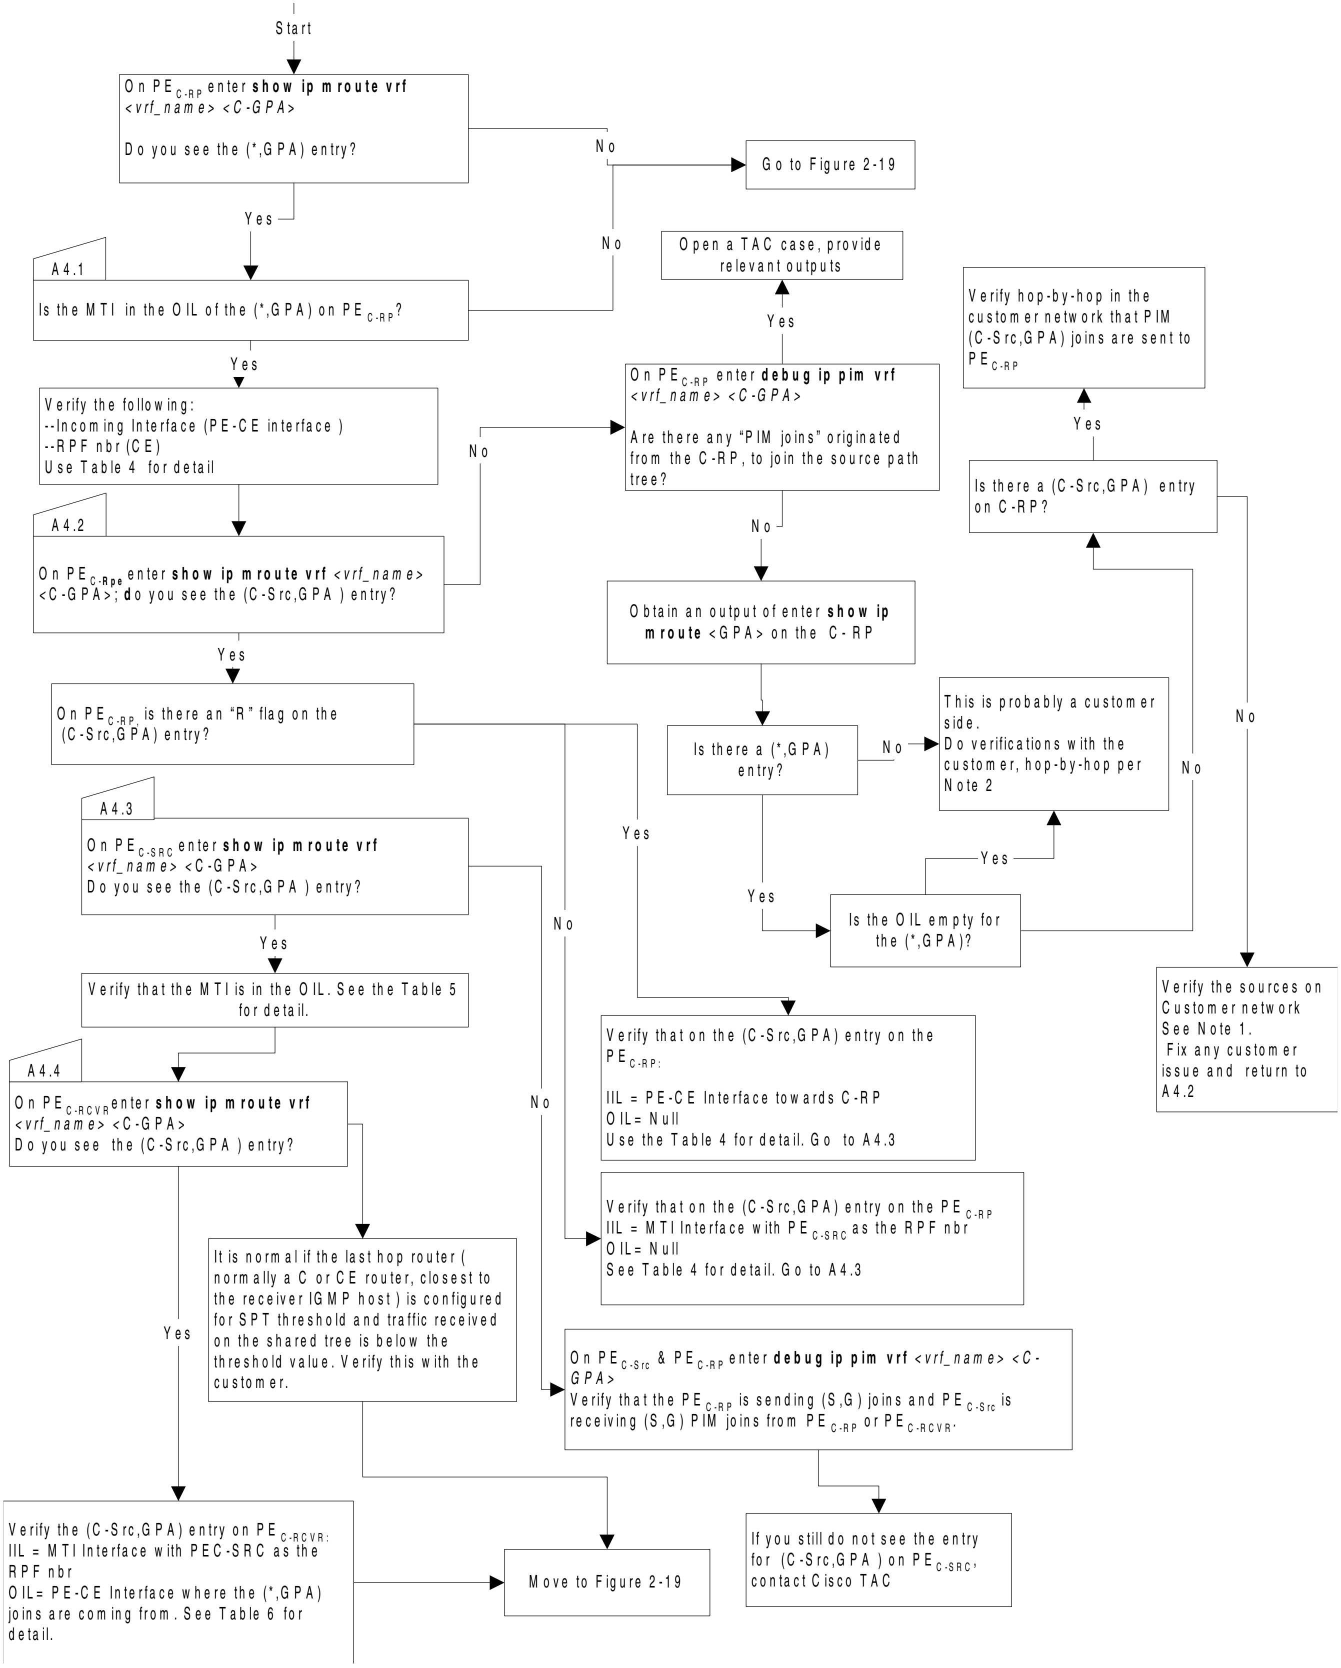

On PE_{C-RP} enter **show ip mroute vrf** *<vrf_name> <C-GPA>*
Do you see the (*,GPA) entry?

Go to Figure 2-19

Is the MTI in the OIL of the (*,GPA) on PE_{C-RP}?

Open a TAC case, provide relevant outputs

Verify hop-by-hop in the customer network that PIM (C-Src,GPA) joins are sent to PE_{C-RP}

Verify the following:
--Incoming Interface (PE-CE interface)
--RPF nbr (CE)
Use Table 4 for detail

On PE_{C-RP} enter **debug ip pim vrf** *<vrf_name> <C-GPA>*
Are there any "PIM joins" originated from the C-RP, to join the source path tree?

Is there a (C-Src,GPA) entry on C-RP?

On PE_{C-Rpe} enter **show ip mroute vrf** *<vrf_name>* <C-GPA>; **d**o you see the (C-Src,GPA) entry?

Obtain an output of enter **show ip mroute** <GPA> on the C-RP

On $PE_{C-RP,}$ is there an "R" flag on the (C-Src,GPA) entry?

Is there a (*,GPA) entry?

This is probably a customer side.
Do verifications with the customer, hop-by-hop per Note 2

On PE_{C-SRC} enter **show ip mroute vrf** *<vrf_name>* <C-GPA>
Do you see the (C-Src,GPA) entry?

Is the OIL empty for the (*,GPA)?

Verify that the MTI is in the OIL. See the Table 5 for detail.

Verify the sources on Customer network
See Note 1.
Fix any customer issue and return to A4.2

Verify that on the (C-Src,GPA) entry on the $PE_{C-RP:}$
IIL = PE-CE Interface towards C-RP
OIL= Null
Use the Table 4 for detail. Go to A4.3

On PE_{C-RCVR} enter **show ip mroute vrf** *<vrf_name>* <C-GPA>
Do you see the (C-Src,GPA) entry?

Verify that on the (C-Src,GPA) entry on the PE_{C-RP}
IIL = MTI Interface with PE_{C-SRC} as the RPF nbr
OIL= Null
See Table 4 for detail. Go to A4.3

It is normal if the last hop router (normally a C or CE router, closest to the receiver IGMP host) is configured for SPT threshold and traffic received on the shared tree is below the threshold value. Verify this with the customer.

On PE_{C-Src} & PE_{C-RP} enter **debug ip pim vrf** *<vrf_name> <C-GPA>*
Verify that the PE_{C-RP} is sending (S,G) joins and PE_{C-Src} is receiving (S,G) PIM joins from PE_{C-RP} or PE_{C-RCVR}.

Verify the (C-Src,GPA) entry on $PE_{C-RCVR:}$
IIL = MTI Interface with PEC-SRC as the RPF nbr
OIL= PE-CE Interface where the (*,GPA) joins are coming from. See Table 6 for detail.

Move to Figure 2-19

If you still do not see the entry for (C-Src,GPA) on PE_{C-SRC}, contact Cisco TAC

Notes for Figure 18

Note 1

Verify that the source is emitting traffic with the proper TTL. Also, check to see if the first hop router is stuck in registering.

Note 2

This is a customer issue. The C-RP has not received the PIM (*,GPA) join from the $PE_{C\text{-}RP.}$ In the customer's network, verify hop by hop that they are getting PIM (*,G) joins.

Table 2-4 shows the VRF Mroute states for the control plane on the $PE_{C\text{-}RP.}$

Table 2-4 VRF Mroute states on $PE_{C\text{-}RP}$

For (*,G)	Flags: S
Incoming interface towards RP	PE-CE interface
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	Default mdt tunnel interface
For (S,G) with no R flag set, where the last-hop router stays on source path tree	Flags: PTX
Incoming interface towards source	Default mdt tunnel interface (RPF neighbor = $PE_{C\text{-}SRC}$)
Outgoing interfaces	Null
For (S,G) with R bit set, when last hop router joins the source path tree	Flags: PR
Incoming interface towards RP	PE-CE interface towards RP
Outgoing interfaces	Null

Table 2-5 shows the VRF Mroute states for the control plane on the $PE_{C\text{-}SRC.}$

Table 2-5 From $PE_{C\text{-}SRC}$

For (*,G)	Flags: SP
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor equals $PE_{C\text{-}RP}$)
Outgoing interfaces	Null
For (S,G)	Flags: T
Incoming interface towards source	PE-CE interface
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	Default mdt tunnel interface

Table 2-6 shows the VRF Mroute states for the control plane on the PE$_{C\text{-}RCVR.}$

Table 2-6 From PE$_{C\text{-}RCVR}$

For (*,G)	Flags: S
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor equals PE$_{C\text{-}RP}$)
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Note The following entry exists only if the last hop router joins the source path tree.

For (S,G)	Flags: T
Incoming interface towards source	Default mdt tunnel interface (RPF neighbor equals PE$_{C\text{-}SRC}$)
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Sample Output for the Mroutes in the VRF Context

The following example shows output when the last hop router joins the shared path tree:

```
PE3#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 2d09h/00:02:56, RP 192.168.100.202, flags: S
 Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 2d09h/00:02:56
(192.168.1.1, 225.1.1.3), 06:01:53/00:02:26, flags: PR
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list: Null

PE3#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 2d09h/00:02:52, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel0, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 2d09h/00:02:52
(192.168.1.1, 225.1.1.3), 06:02:32/00:03:24, flags: T
  Incoming interface: Tunnel0, RPF nbr 172.16.100.1
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 06:02:32/00:02:52

PE1#show ip mro vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 06:03:06/stopped, RP 192.168.100.202, flags: SP
  Incoming interface: Tunnel7, RPF nbr 172.16.100.2
  Outgoing interface list: Null

(192.168.1.1, 225.1.1.3), 06:03:06/00:03:29, flags: T
  Incoming interface: Serial1/0, RPF nbr 192.168.11.2
  Outgoing interface list:
    Tunnel7, Forward/Sparse-Dense, 06:03:06/00:02:56
```

The following example shows the last hop router staying on the shared path tree:

```
PE2#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 00:05:01/00:03:11, RP 192.168.100.202, flags: S
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:04:04/00:03:11
(192.168.1.1, 225.1.1.3), 00:04:04/00:02:55, flags: PTX
  Incoming interface: Tunnel0, RPF nbr 172.16.100.1
  Outgoing interface list: Null
```

On $\text{PE}_{\text{C-SRC}}$

```
PE1#show ip mro vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 00:04:27/stopped, RP 192.168.100.202, flags: SP
  Incoming interface: Tunnel7, RPF nbr 172.16.100.2
  Outgoing interface list: Null
(192.168.1.1, 225.1.1.3), 00:03:42/00:03:29, flags: T
  Incoming interface: Serial1/0, RPF nbr 192.168.11.2
  Outgoing interface list:
    Tunnel7, Forward/Sparse-Dense, 00:03:42/00:03:24
```

```
PE1#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 00:05:25/00:03:24, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel0, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:04:20/00:02:51
```

Troubleshooting the PIM Adjacency with the CEs

If there is an issue of PIM adjacency on the PE- CE link, the flowchart in Figure 2-19 should be used for troubleshooting or for checking the PIM adjacency.

Repeat the steps of Figure 2-19 as many times as needed, depending on the number of PE-CE links.

Figure 2-19 Troubleshooting PIM Neighbors over the PE-CE Interface

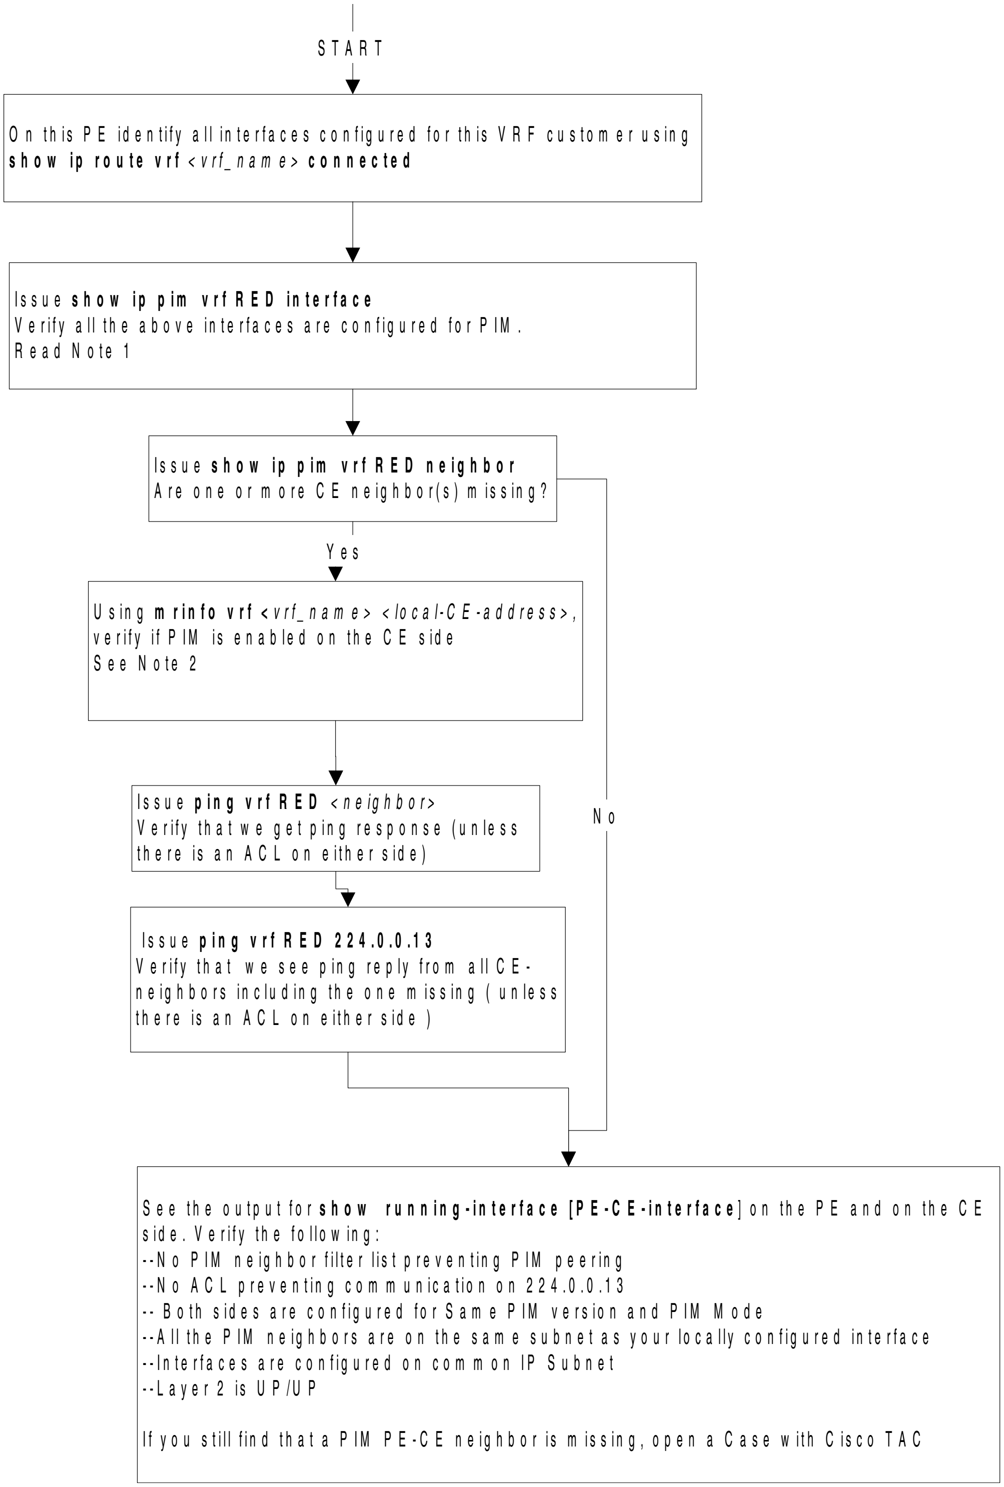

On this PE identify all interfaces configured for this VRF customer using **show ip route vrf** <*vrf_name*> **connected**

Issue **show ip pim vrf RED interface**
Verify all the above interfaces are configured for PIM.
Read Note 1

Issue **show ip pim vrf RED neighbor**
Are one or more CE neighbor(s) missing?

Using **mrinfo vrf** <*vrf_name*> <*local-CE-address*>,
verify if PIM is enabled on the CE side
See Note 2

Issue **ping vrf RED** <*neighbor*>
Verify that we get ping response (unless there is an ACL on either side)

Issue **ping vrf RED 224.0.0.13**
Verify that we see ping reply from all CE-neighbors including the one missing (unless there is an ACL on either side)

See the output for **show running-interface [PE-CE-interface**] on the PE and on the CE side. Verify the following:
--No PIM neighbor filter list preventing PIM peering
--No ACL preventing communication on 224.0.0.13
-- Both sides are configured for Same PIM version and PIM Mode
--All the PIM neighbors are on the same subnet as your locally configured interface
--Interfaces are configured on common IP Subnet
--Layer 2 is UP/UP

If you still find that a PIM PE-CE neighbor is missing, open a Case with Cisco TAC

Notes for Figure 2-19

Note 1

We might see some additional interfaces, such as MTI and customer interfaces, that are in a shutdown state.

Note 2

It is possible that we do not get any response (if disabled by mrinfo filters or an ACL for igmp).

Sample Output for Troubleshooting the PIM Adjacency on the PE-CE Link

PE1#**mrinfo vrf RED 192.168.112.2**

```
192.168.112.2 [version  12.0] [flags: PMA]:
  192.168.12.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
  192.168.0.2 -> 0.0.0.0 [1/0/pim/disabled/down/leaf]            → Interface down
  192.168.212.2 -> 0.0.0.0 [1/0/pim/disabled/down/leaf]
  192.168.112.2 -> 192.168.112.1 [1/0/pim]    → With PIM enabled
```

```
PE1#mrinfo vrf RED 192.168.112.2
192.168.112.2 [version  12.0] [flags: PMA]:
  192.168.12.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
  192.168.0.2 -> 0.0.0.0 [1/0/pim/disabled/down/leaf]
  192.168.212.2 -> 0.0.0.0 [1/0/pim/disabled/down/leaf]
```

We are missing the line for 192.168.112.2 because PIM is not enabled on the CE.

Topology 2

In this topology we have added a receiver on the PE$_{C\text{-}SRC}$, on top of what is depicted in Figure 2-9. The pages that follow provide troubleshooting guidelines based on this difference.

Figure 2-20 Topology 2: Troubleshooting Customer-Facing Settings of the PEs

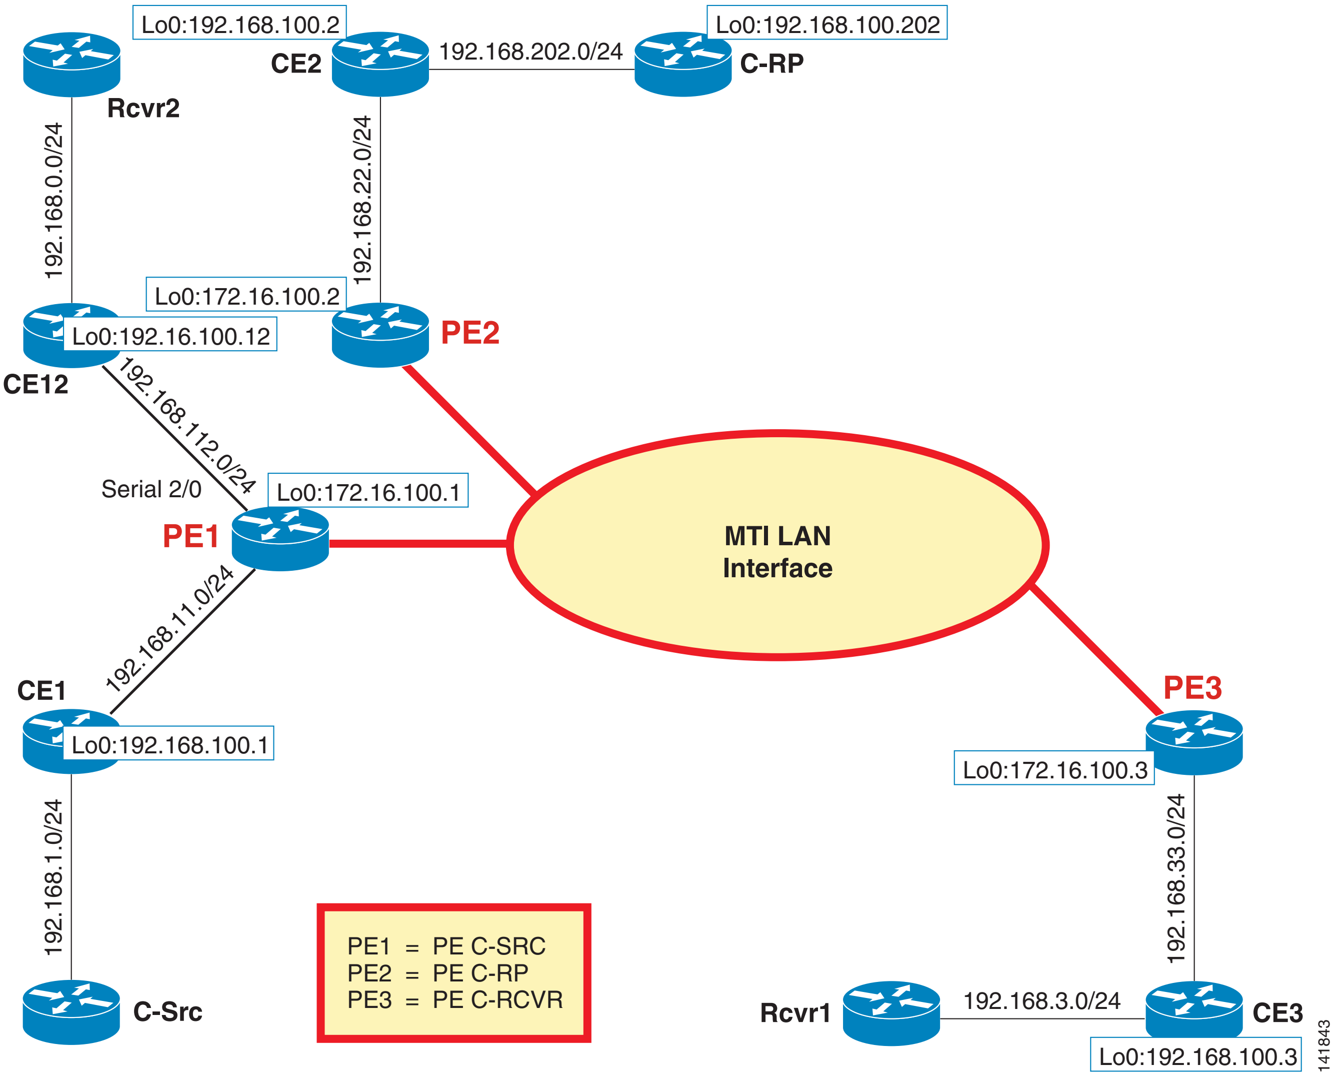

If we consider the Mroute states for (*,C-GPA) and (C-SRC, C-GPA) for the routers, then $PE_{C\text{-}RP,}$ and $PE_{C\text{-}RCVR}$ will look the same as the tables provided for the topology in Figure 2-9.

The only difference between the two topologies shown in Figure 2-9 and Figure 2-20 is the Mroute state for $PE_{C\text{-}SRC.}$

Table 2-7 shows the Mroute state for $PE_{C\text{-}RP}$ for the scenario shown in Figure 2-20.

Table 2-7 From $PE_{C\text{-}RP}$

For (*,C-GPA)	Flags: S
Incoming Interface towards RP	PE-CE interface
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	Default mdt tunnel interface

Table 2-7 From $\textit{PE}_{\textit{C-RP}}$

For (C-SRC, C-GPA)	Flags: PR
Incoming Interface towards RP	PE-CE interface towards RP
Outgoing Interfaces	Null

Table 2-8 shows the Mroute state for $\text{PE}_{\text{C-SRC}}$ for the scenario shown in Figure 2-20.

Table 2-8 Mroute for $\textit{PE}_{\textit{C-SRC}}$

For (*,G)	Flags: S
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor equals $\text{PE}_{\text{C-RP}}$)
Outgoing interfaces towards Receivers PE-CE interface	PE-CE Interface towards Rcvr 2
For (S,G)	Flags: T
Incoming interface towards source	PE-CE interface
Outgoing interfaces towards receivers PE-CE interface	Default mdt tunnel interface plus PE-CE interfaces towards Rcvr 2

The following tables show the Mroute state for $\text{PE}_{\text{C-RCVR}}$ for the scenario shown in Figure 2-20.

Table 2-9 Mroute for $\textit{PE}_{\textit{C-RCVR}}$

For (*,G)	Flags: S
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor = $\text{PE}_{\text{C-RP}}$)
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Note The following entry exists only if the last hop router joins the source path tree.

For (S,G)	Flags: T
Incoming Interface towards source	Default mdt tunnel interface (RPF neighbor =$\text{PE}_{\text{C-SRC}}$)
Outgoing interfaces towards receivers, where IGMP/PIM Joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Sample Output of the Mroute

The following example shows output on PE2, taken when the last hop router joins the source path tree:

```
PE2#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 01:23:52/00:02:35, RP 192.168.100.202, flags: S
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:03:32/00:02:35

(192.168.1.1, 225.1.1.3), 01:23:52/00:02:05, flags: PR
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list: Null

PE3#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 01:21:30/00:02:32, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel3, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:01:10/00:02:32

(192.168.1.1, 225.1.1.3), 00:02:50/00:03:24, flags: T
  Incoming interface: Tunnel3, RPF nbr 172.16.100.1
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:01:10/00:02:32

PE1#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 3d14h/00:03:02, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel14, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial2/0, Forward/Sparse, 00:05:23/00:03:02

(192.168.1.1, 225.1.1.3), 3d14h/00:03:28, flags: T
  Incoming interface: Serial1/0, RPF nbr 192.168.11.2
  Outgoing interface list:
    Serial2/0, Forward/Sparse, 00:05:23/00:03:02
    Tunnel14, Forward/Sparse-Dense, 00:05:23/00:03:03
```

Topology 3

In the topology shown in Figure 2-21, an additional receiver is included on the site of PE$_{C\text{-}RP;}$ providing a total of three receivers.

Note that the figure for this topology is the same as the main figure for this document (Figure 1-1).

Figure 2-21 Topology 3: Additional Receiver Included on Site $PE_{C\text{-}RP}$

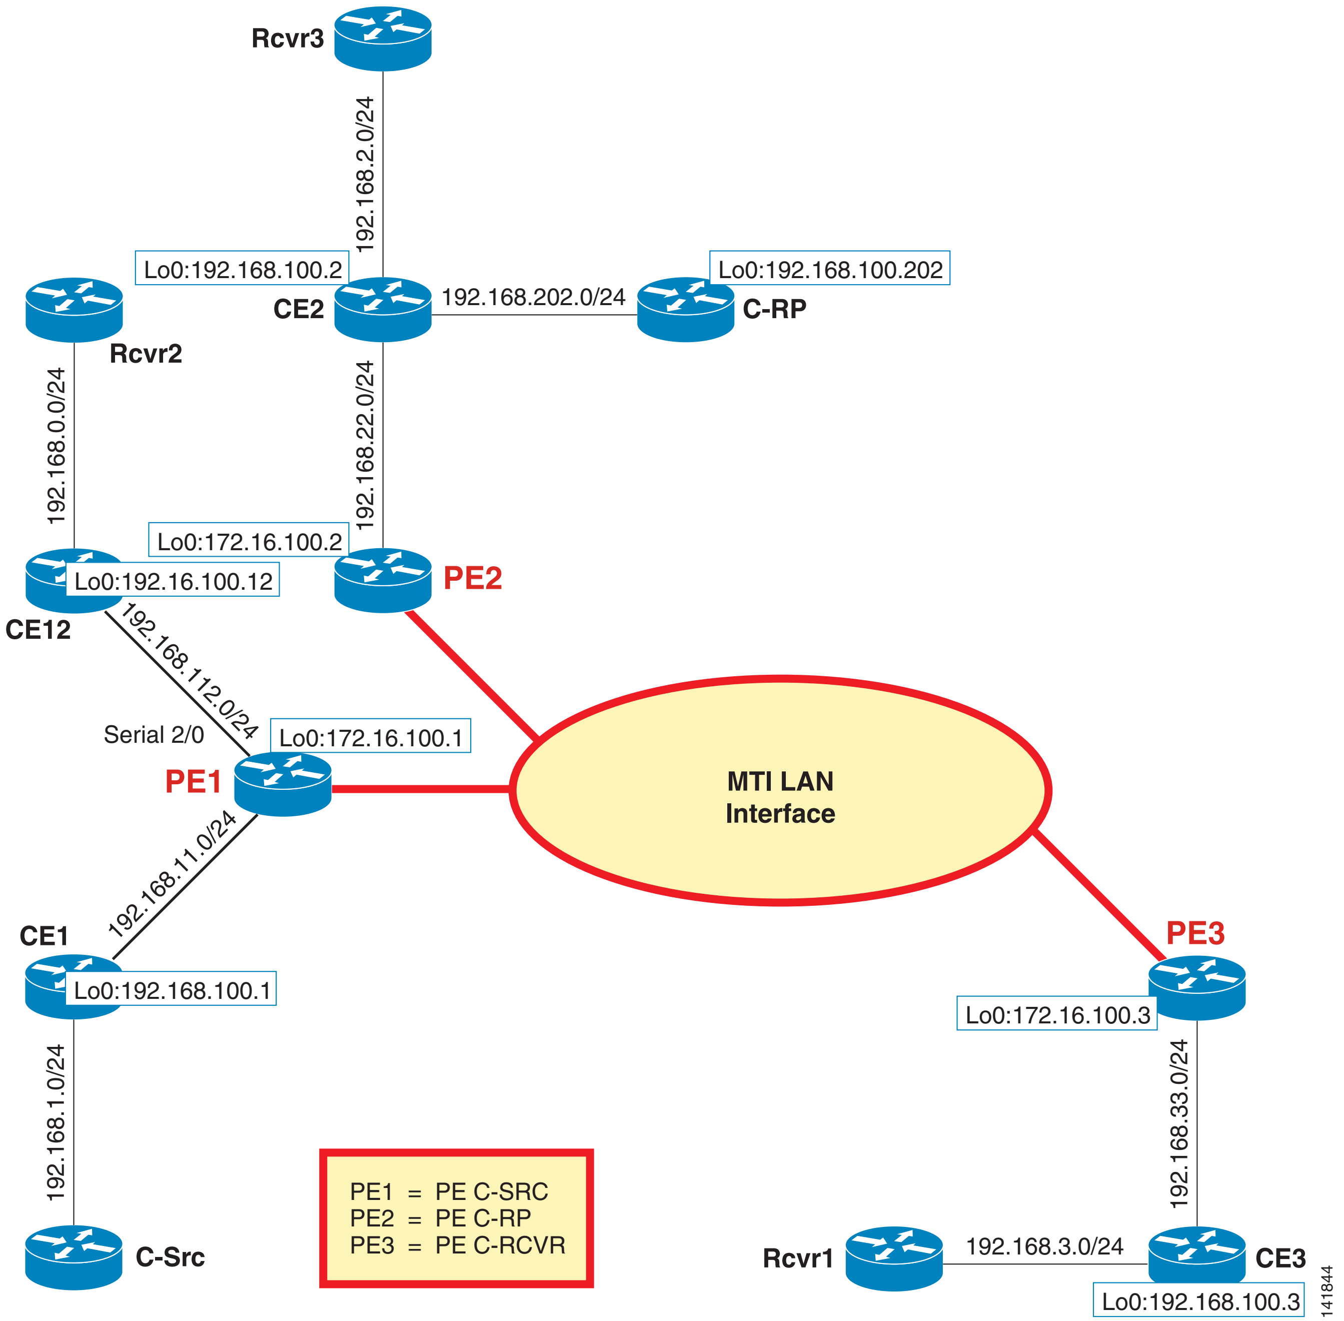

The Mroute states for (*,C-GPA) and (C-SRC, C-GPA) are shown in the following tables. The only difference between the tables shown for the main topology and topologies 2, and 3 is the Mroute state for $PE_{C\text{-}SRC.}$

VRF Mroute States on PE$_{C\text{-}RP}$.

Table 2-10 From PE$_{C\text{-}RP}$

For (*,C-GPA)	Flags: S
Incoming interface towards RP	PE-CE interface
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	Default mdt tunnel interface
For (C-SRC, C-GPA)	Flags: T
Incoming interface towards source	Default mdt tunnel interface (RPF neighbor equals PE$_{C\text{-}SRC}$)
Outgoing interfaces towards receivers, where IGMP/PIM Joins are received)	PE-CE interfaces where the PIM/IGMP joins are received

Table 2-11 VRF Mroute States on PE$_{C\text{-}SRC}$ from PE$_{C\text{-}SRC}$

For (*,G)	Flags: S
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor = PE$_{C\text{-}RP}$)
Outgoing interfaces towards receivers PE-CE interface	PE-CE Interface (Towards Rcvr 2)
For (S,G)	Flags: T
Incoming interface towards source	PE-CE interface
Outgoing interfaces towards receivers PE-CE interface	Default mdt tunnel interface plus PE-CE interfaces (towards Rcvr 2)

Table 2-12 VRF Mroute states on PE$_{C\text{-}RCVR}$ From PE$_{C\text{-}RCVR}$

For (*,G)	Flags: S
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor = PE$_{C\text{-}RP}$)
Outgoing interfaces towards receivers where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Note The following entry exists only if the last hop router joins the source path tree.

Table 2-12 VRF Mroute states on PE_{C-RCVR} From PE_{C-RCVR}

For (S,G)	Flags: T
Incoming interface towards source	Default mdt tunnel interface (RPF neighbor =PE_{C-SRC})
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Sample Output of Multicast Routes for Topology 3

The following example shows output when the last hop router joins the source path tree:

```
PE2#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 01:40:58/00:02:53, RP 192.168.100.202, flags: S
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:45/00:02:53

(192.168.1.1, 225.1.1.3), 01:40:58/00:03:29, flags: T
  Incoming interface: Tunnel0, RPF nbr 172.16.100.1
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:00:45/00:02:53

PE3#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 01:41:12/00:02:42, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel3, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:00:59/00:02:42

(192.168.1.1, 225.1.1.3), 00:22:31/00:03:23, flags: T
  Incoming interface: Tunnel3, RPF nbr 172.16.100.1
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:00:59/00:02:42

PE1#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 3d14h/00:03:08, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel14, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial2/0, Forward/Sparse, 00:00:26/00:03:08

(192.168.1.1, 225.1.1.3), 3d14h/00:03:24, flags: T
  Incoming interface: Serial1/0, RPF nbr 192.168.11.2
  Outgoing interface list:
    Serial2/0, Forward/Sparse, 00:00:26/00:03:08
    Tunnel14, Forward/Sparse-Dense, 00:00:26/00:03:10
```

Topology 4

The topology in Figure 2-22 shows that the PE$_{C\text{-}SRC}$ and the PE$_{C\text{-}RP}$ are on the same site.

Figure 2-22 PE$_{C\text{-}SRC}$ and the PE$_{C\text{-}RP}$ On the Same Site

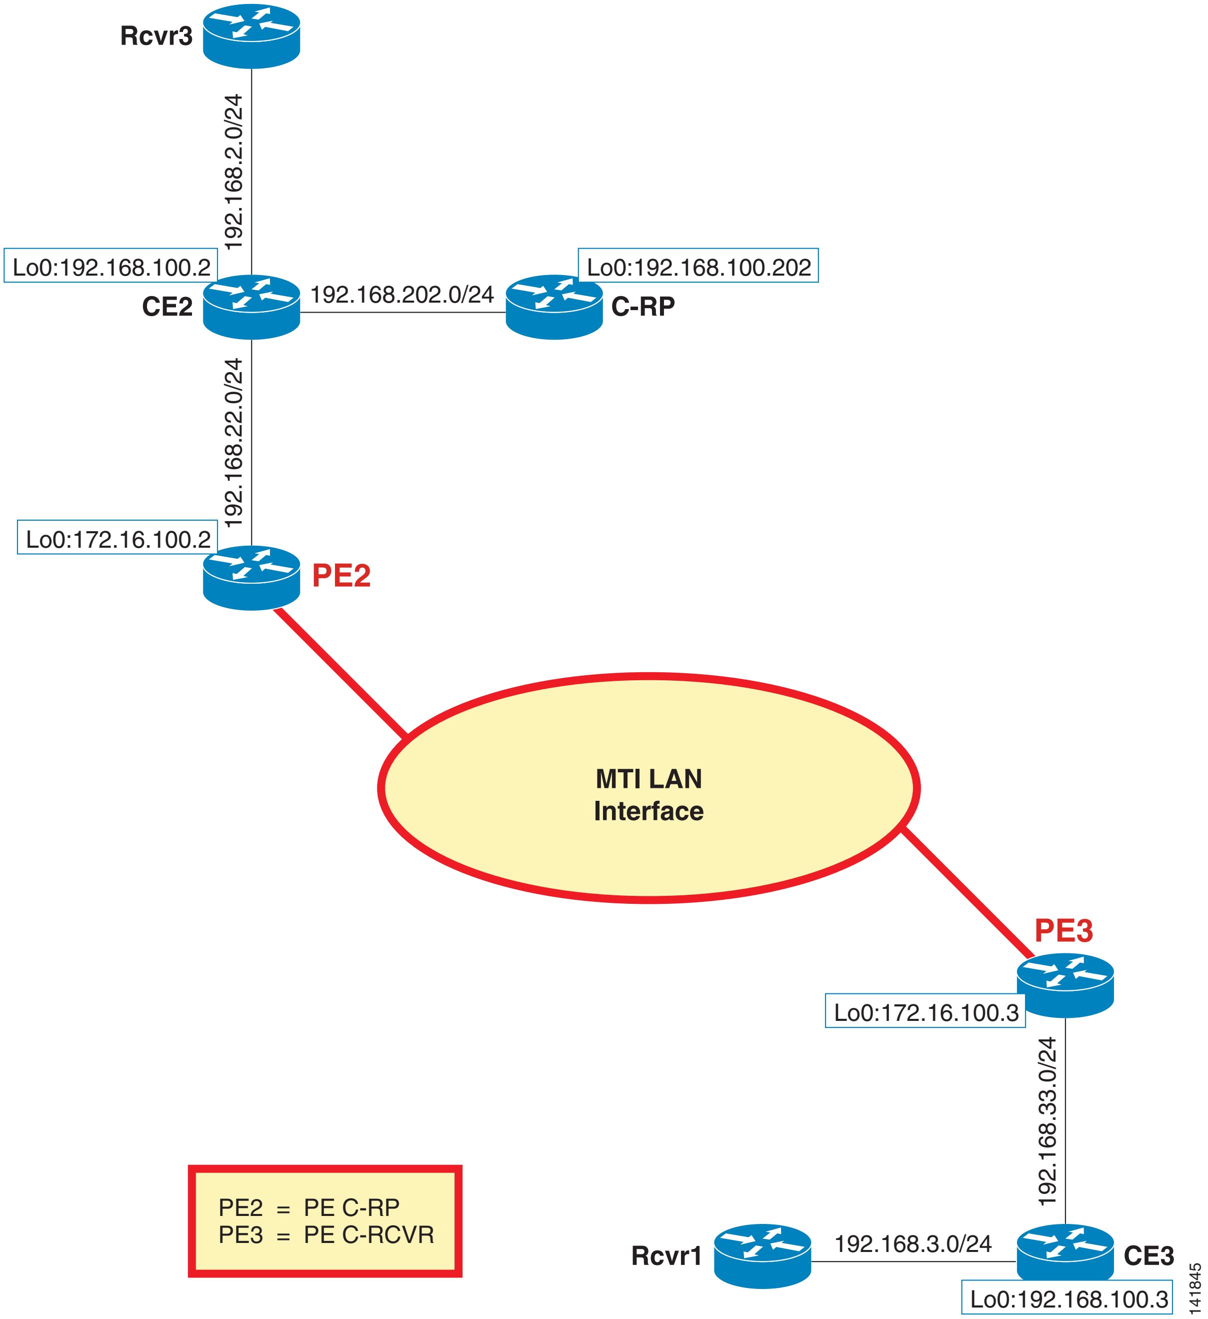

The following tables show the Mroute states for (*,C-GPA) and (C-SRC, C-GPA) on PE$_{\text{/C-SRC}}$ (see Table 2-13) and PE$_{\text{C-RCVR}}$ (see Table 2-14).

Table 2-13 Mroute States from PE$_{\text{C-SRC}}$

For (*,C-GPA)	Flags: S
Incoming interface towards RP	PE-CE interface
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	Default mdt tunnel interface
For (C-SRC,C-GPA)	Flags: T
Incoming interface towards source	MDT tunnel interface (RPF neighbor equals PE$_{\text{C-SRC}}$)
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP Joins are received

Table 2-14 Mroute States from PE$_{\text{C-RCVR}}$

For (*,G)	Flags: S
Incoming interface towards RP	Default mdt tunnel interface (RPF neighbor equals PE$_{\text{C-SRC}}$)
Outgoing interfaces towards receivers, where IGMP/PIM Joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Note The following entry exists only if the last hop router joins the source path tree.

For (S,G)	Flags: T
Incoming interface towards source	Default mdt tunnel interface (RPF neighbor equals PE$_{\text{C-SRC}}$)
Outgoing interfaces towards receivers, where IGMP/PIM joins are received	PE-CE interfaces where the PIM/IGMP joins are received

Sample Output of Multicast Routes for Topology 4

The following example shows output when the last hop router joins the source path tree:

```
PE2#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 00:03:24/00:02:42, RP 192.168.100.202, flags: S
```

```
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:03:01/00:02:42

(192.168.2.1, 225.1.1.3), 00:03:01/00:03:26, flags: T
  Incoming interface: Serial1/0, RPF nbr 192.168.22.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:03:01/00:02:42

PE3#show ip mroute vrf RED 225.1.1.3 | begin 225
(*, 225.1.1.3), 02:07:48/00:02:34, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel3, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:27:34/00:02:34

(192.168.2.1, 225.1.1.3), 00:12:40/00:03:21, flags: T
  Incoming interface: Tunnel3, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 00:12:40/00:02:34
```

Section 4: Verifying mVPN Data Flow

This section provides the steps that can be used to troubleshoot the mVPN data flow. The topology shown in Figure 2-23 is used to cover the different show commands that can be used to monitor the data flow. Note that this section covers only the data flow on the default MDT and not on the data MDT.

Assumptions

The following assumptions are made:

- Mroute state entries for (*,C-GPA) and (C-SRC, C-GPA) are present with valid incoming interface list (IIL) and outgoing interface list (OIL) in the mVRF table on $\text{PE}_{\text{C-SRC}}$ and $\text{PE}_{\text{C-RCVR}}$.
- The last hop routers (CE3 and CE12 for the topology in Figure 2-23),where the receivers are connected, are configured to join the shortest path tree (for example, SPT equals 0 on last hop routers).
- There is no backdoor connectivity between the VPN sites.
- The topology in Figure 2-23 shows one $\text{PE}_{\text{C-RCVR}}$. If there is more than one, the same concept can be used to analyze and troubleshoot the additional data flow.
- There are no multicast-boundary, ACLs, multicast-threshold settings in the SP core and edge that are dropping the packet.
- This section focuses on one flow (C-SRC, C-GPA); the same steps can be used for checking additional flows.

Topology for Data Flow Troubleshooting steps in the Default MDT

Figure 2-23 Troubleshooting Topology for mVPN Data Flow

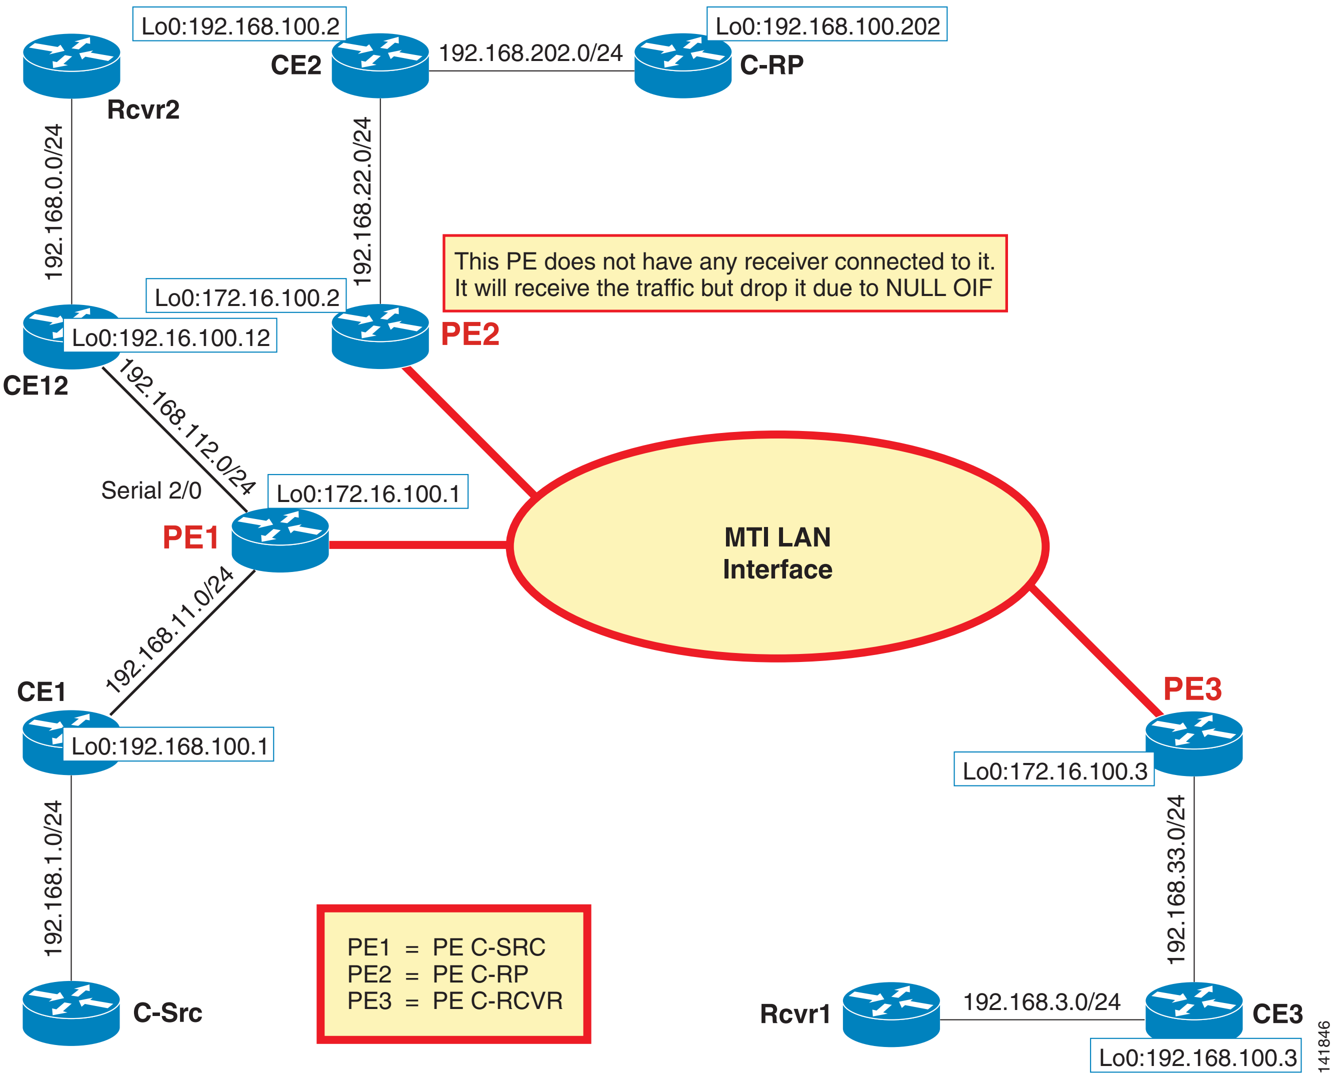

Customer Output

It is extremely helpful to obtain two consecutive sets of output (at 30 to 60 secs interval) from CEs, C-RP, first-hop router and, the last-hop routers using the following commands:

```
show interface  <Facing_PE>
show ip pim interface stats
show ip pim interface counts
show ip mroute <C-GPA>
show ip mroute <C-GPA> active 1
show ip mroute <C-GPA> count
show ip traffic | include bad hop
show ip traffic | include fragment
```

Verifying Mroute States on $PE_{C\text{-}SRC}$ and $PE_{C\text{-}RCVR}$

The control plane states on $PE_{C\text{-}SRC}$ and $PE_{C\text{-}RCVR}$ will provide output that is similar to that shown in the following examples.

$PE_{C\text{-}SRC}$ receives the traffic from one customer site and replicates it to other local sites where there are interested receivers. If there are interested sites on other PEs, then $PE_{C\text{-}SRC}$ replicates the stream over the MTI interface for this mVRF. The following example shows output from the PEs:

```
PE1#show ip mroute vrf RED 225.1.1.10 | begin 225
(*, 225.1.1.10), 07:16:49/00:03:07, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel7, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial2/0, Forward/Sparse, 07:04:28/00:03:07

(192.168.1.1, 225.1.1.10), 02:28:47/00:03:24, flags: T
  Incoming interface: Serial1/0, RPF nbr 192.168.11.2
  Outgoing interface list:
    Serial2/0, Forward/Sparse, 02:28:47/00:03:07
    Tunnel7, Forward/Sparse-Dense, 02:28:47/00:03:27


PE3# show ip mroute vrf RED 225.1.1.10 | begin 225
(*, 225.1.1.10), 02:39:18/00:03:06, RP 192.168.100.202, flags: S
  Incoming interface: Tunnel0, RPF nbr 172.16.100.2
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 02:39:18/00:03:06

(192.168.1.1, 225.1.1.10), 02:28:24/00:03:22, flags: T
  Incoming interface: Tunnel0, RPF nbr 172.16.100.1
  Outgoing interface list:
    Serial1/0, Forward/Sparse, 02:28:24/00:03:06
```

The flowchart in Figure 2-24 shows how to monitor the mVPN data flow.

Figure 2-24 Troubleshooting mVPN Data Flow

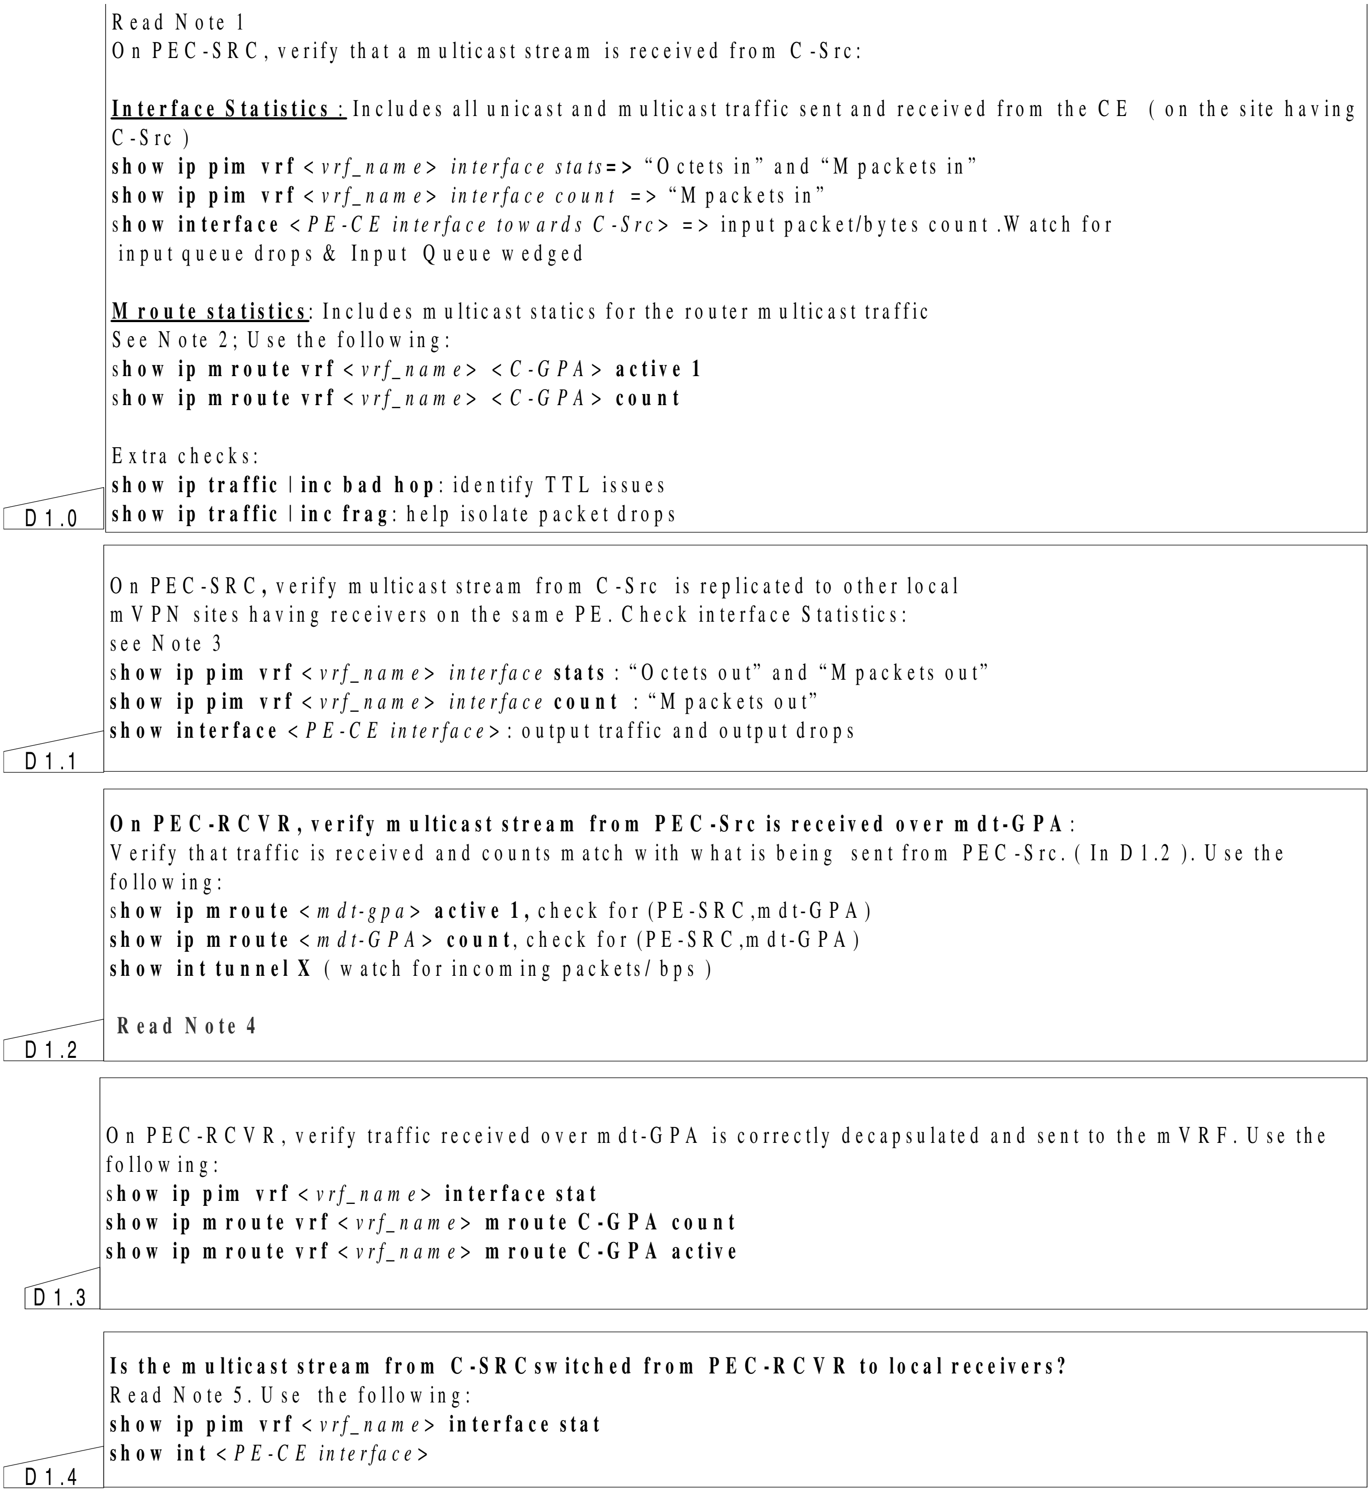

Read Note 1
On PEC-SRC, verify that a multicast stream is received from C-Src:

<u>**Interface Statistics :**</u> Includes all unicast and multicast traffic sent and received from the CE (on the site having C-Src)
show ip pim vrf <*vrf_name*> *interface stats***=>** "Octets in" and "M packets in"
show ip pim vrf <*vrf_name*> *interface count* => "M packets in"
show interface <*PE-CE interface towards C-Src*> => input packet/bytes count .Watch for input queue drops & Input Queue wedged

<u>**Mroute statistics**</u>: Includes multicast statics for the router multicast traffic
See Note 2; Use the following:
show ip mroute vrf <*vrf_name*> <*C-GPA*> **active 1**
show ip mroute vrf <*vrf_name*> <*C-GPA*> **count**

Extra checks:
show ip traffic | inc bad hop: identify TTL issues
show ip traffic | inc frag: help isolate packet drops

D1.0

On PEC-SRC**,** verify multicast stream from C-Src is replicated to other local mVPN sites having receivers on the same PE. Check interface Statistics:
see Note 3
show ip pim vrf <*vrf_name*> *interface* **stats** : "Octets out" and "M packets out"
show ip pim vrf <*vrf_name*> *interface* **count** : "M packets out"
show interface <*PE-CE interface*>: output traffic and output drops

D1.1

On PEC-RCVR, verify multicast stream from PEC-Src is received over mdt-GPA:
Verify that traffic is received and counts match with what is being sent from PEC-Src. (In D1.2). Use the following:
show ip mroute <*mdt-gpa*> **active 1,** check for (PE-SRC,mdt-GPA)
show ip mroute <*mdt-GPA*> **count**, check for (PE-SRC,mdt-GPA)
show int tunnel X (watch for incoming packets/ bps)

Read Note 4

D1.2

On PEC-RCVR, verify traffic received over mdt-GPA is correctly decapsulated and sent to the mVRF. Use the following:
show ip pim vrf <*vrf_name*> **interface stat**
show ip mroute vrf <*vrf_name*> **mroute C-GPA count**
show ip mroute vrf <*vrf_name*> **mroute C-GPA active**

D1.3

Is the multicast stream from C-SRC switched from PEC-RCVR to local receivers?
Read Note 5. Use the following:
show ip pim vrf <*vrf_name*> **interface stat**
show int <*PE-CE interface*>

D1.4

Notes for Figure 2-24

Note 1

- We assume that multicast stream in question is the traffic flowing from C-SRC to C-GPA.
- Note that most of the output needs to be captured twice, at 30 to 60 second intervals, so we can compare the evolution of the traffic flow.

Note 2

- Interface Statistics: Includes all unicast and multicast traffic sent and received from the CE (on the site having C-SRC):

  show ip pim vrf *vrf_name* **interface stats**

  Look for Octets in and Mpackets in count and verify it is increasing.

  show ip pim vrf *vrf_name* **interface count**

  Look for Mpackets in count and verify it is increasing.

  s**how interface** *PE-CE interface towards C-SRC*

  Check for the input packet/bytes count and verify it is increasing, watch for input queue drops and Input Queue Wedged.

- Mroute statistics: Includes multicast statics for:

  s**how ip mroute vrf** *vrf_name C-GPA* **active 1**

  Check for multicast stream from C-SRC, and note the kbps.

  s**how ip mroute vrf** *vrf_name C-GPA* **count**

  Check for multicast stream from C-SRC, note the Forwarding Counts, and watch for RPF or other drops.

- Extra checks:

  show ip traffic | inc bad hop to identify TTL issues (it is a per-box counter).

  show ip traffic | inc frag to help isolate packet drops due to DF bit set on the C-SRC and low MTU on MTI. (It is a per-box counter). This is less likely to be an issue in production.

- Includes all unicast and multicast traffic sent and received from the CE (on the site having C-RCVR).

  show ip pim vrf *vrf_name* **interface stats**

  Look for the Octets out and Mpackets out count and verify it is increasing.

  show ip pim vrf *vrf_name* **interface count**

  Look for the Mpackets out count and verify it is increasing.

  show interface *PE-CE interface*

  Check for output traffic and output drops.

- If the PE-CE link is a frame relay interface, then watch for counter broadcasts sent or dropped x/y under the **show interface** *PE-CE_interface* command.

Note 3

- On $\text{PE}_{\text{C-RCVR}}$, verify that the multicast stream from $\text{PE}_{\text{C-SRC}}$ is received over mdt-GPA.
- Use the following output to verify that the traffic is received and the counts match with what is being sent from $\text{PE}_{\text{C-SRC}}$:

show ip mroute *mdt-gpa* **active 1**

Check for (PE-SRC,mdt-GPA) and note the kbps value.

show ip mroute *mdt-GPA* **count**

Check for (PE-SRC,mdt-GPA) and note the kbps value.

Note 4

We might not see the (PE$_{C\text{-}SRC}$, mdt-GPA) entry on PE$_{C\text{-}RCVR}$ if the multicast protocol in the SP-Core is set for Bidir or **spt threshold-infinity** is configured on all the PEs. In such a case, the best method would be to track the counters on the shared tree (RProoted tree).

Note 5

Use the following output of two consecutive captures (with a 60 second interval) on PE$_{C\text{-}RCVR}$ to help us relate the data being received and being sent over to another local receiver.

show ip pim vrf *vrf_name* interface stat

show int *PE-CE interface*

Check for output traffic, and watch for any drops.

The following describes the troubleshooting steps for Figure 2-24.

Verify that the Multicast Stream is Received on PE$_{C\text{-}SRC}$ from C-SRC

Q. Is PE$_{C\text{-}SRC}$ receiving any multicast traffic from the C-SRC?

A. The output below shows two consecutive captures (at 30 to 60 secs interval) on PE$_{C\text{-}SRC:}$

- **Interface Statistics**: Includes all unicast and multicast traffic sent and received from the CE (on the site having C-SRC)
- Enter the **show ip pim vrf** *vrf_name* **interface stats** command, look for **Octets in** and **Mpackets in** count, and verify that it is increasing.
- Enter the **show ip pim vrf** *vrf_name* **interface** *count* command, look for **Mpackets in** count and verify it is increasing.
- Enter the **show interface** *PE-CE interface towards C-SRC* command; check for **input packet/bytes** count and verify that it is increasing. Look for input queue drops and wedged input queue.

```
PE1#show ip pim vrf RED interface stats

Interface        Mpackets In    Mpackets Out        Octets In        Octets Out
Serial2/0                  0           63106                0          64018536
Serial1/0            1586098            6049       1626265744            304800
Tunnel14                1175          126143            59208         129675004
PE1#
PE1#
PE1#show ip pim vrf RED interface stats

Interface        Mpackets In    Mpackets Out        Octets In        Octets Out
Serial2/0                  0           63106                0          64018536
Serial1/0            1586122            6049       1626290416            304800
Tunnel14                1175          126167            59208         129699676
```

```
PE1#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface              FS  Mpackets In/Out
192.168.112.1    Serial2/0              *   0/63106
192.168.11.1     Serial1/0              *   1586135/6049
172.16.100.1     Tunnel14               *   1175/126180

PE1#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface              FS  Mpackets In/Out
192.168.112.1    Serial2/0              *   0/63106
192.168.11.1     Serial1/0              *   1586149/6049
172.16.100.1     Tunnel14               *   1175/126194
```

Note These counters represent aggregated multicast traffic coming from the customer. The multicast traffic can belong to any GPA and can originate from any C-SRC address.

Q. Is PE$_{\text{C-SRC}}$ receiving (C-SRC, C-GPA) multicast traffic from the C-SRC?

A. The following sample output shows how to verify if the PE$_{\text{C-SRC}}$ is receiving traffic.

- **Mroute statistics**: Includes multicast statics for the (C-SRC, C-GPA)
- Enter the **show ip mroute vrf** *vrf_name C-GPA* **active 1** command to check for multicast stream from C-SRC (note the kpbps)
- Enter the **show ip mroute vrf** *vrf_name C-GPA* **count** command to check for multicast stream from C-SRC. Note the forwarding counts and look for RPF or other drops.

  The output can help in verifying that the **Forwarding Counts** are increasing for the (C-SRC, C-GPA) entry.

```
PE1#show ip mroute vrf RED 225.1.1.10 active 1
Active IP Multicast Sources - sending >= 1 kbps

Group: 225.1.1.10, (?)
   Source: 192.168.1.1 (?)
     Rate: 1 pps/8 kbps(1sec), 8 kbps(last 30 secs), 7 kbps(life avg)
PE1#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
18 routes using 10408 bytes of memory
12 groups, 0.50 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 1893
  RP-tree: Forwarding: 2/0/1028/0, Other: 3/1/0
  Source: 192.168.1.1/32, Forwarding: 1891/1/1028/8, Other: 1892/0/1

PE1#show ip mroute vrf RED 225.1.1.10 active 1
Active IP Multicast Sources - sending >= 1 kbps

Group: 225.1.1.10, (?)
   Source: 192.168.1.1 (?)
     Rate: 1 pps/9 kbps(1sec), 9 kbps(last 10 secs), 8 kbps(life avg)

PE1#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
```

```
10 routes using 5144 bytes of memory
5 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 42557
  RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0
  Source: 192.168.1.1/32, Forwarding: 42557/1/1028/9, Other: 42557/0/0

PE1#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
10 routes using 5144 bytes of memory
5 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 42562
  RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0
  Source: 192.168.1.1/32, Forwarding: 42562/1/1028/8, Other: 42562/0/0
```

Traffic Failure: No Multicast Traffic Received on PE$_{\text{C-SRC}}$

If there is no multicast traffic for (C-SRC, C-GPA) from the C-SRC, check for the following:

- Input queue drops on PE$_{\text{C-SRC}}$ (under the **show interface** command)
- Bad hop count increasing under the **show ip traffic | inc bad hop** command, which identifies TTL issues on the traffic received:

```
PE1#show ip pim vrf RED interface stats

Interface        Mpackets In    Mpackets Out          Octets In         Octets Out
Serial2/0                  0           28868                  0           29119072
Serial1/0             614933            1608          627908124              81024
Tunnel7                15381          536051           14239668          548262428
PE1#show ip pim vrf RED interface stats

Interface        Mpackets In    Mpackets Out          Octets In         Octets Out
Serial2/0                  0           28869                  0           29119120
Serial1/0             614933            1609          627908124              81072
Tunnel7                15382          536051           14239716          548262428

PE1#show ip traffic | include bad hop
         0 format errors, 0 checksum errors, 2373 bad hop count
PE1#show ip traffic | include bad hop
         0 format errors, 0 checksum errors, 2389 bad hop count
```

- Check the multicast-boundary statement and inbound ACLs configuration under the PE-CE interface and verify that the boundary/ACL threshold is not dropping the traffic.
- Check for multicast rate-limit statement on the PE$_{\text{C-SRC}}$ and the CE-facing interface, and verify that the customer multicast stream does not exceed the specified rate-limit.
- At this point, all of the output mentioned in Customer Output, page 2-81 should be obtained from the CE router. Ensure that the Mroute entry exists on the CE with correct OIL and traffic is being sent to the PE$_{\text{C-SRC}}$ router. If you still cannot determine the cause of the problem, then open a case with Cisco TAC.

Note If the PE-CE link is a Frame-Relay interface, check for broadcasts sent/dropped 9/0 under the **show interface** *Facing PE C-SRC* command on the CE router.

- RPF Drops

  If there is multicast traffic for (C-SRC, C-GPA) from the CE, but instead of incrementing the forward count, the traffic is dropped because of RPF failure or other drops, then consider the following.

  – If there are RPF drops, verify the RPF interface and RPF neighbor by entering the following commands: **show ip rpf vrf** *vrf_name C-SRC.*

  – Check for available paths in the routing table to reach C-SRC by entering the **show ip route vrf** *vrf_name C-SRC* command.

  The **show ip mroute vrf** *c-GPA* **count** command can show an RPF failed counter, and indicate a potential RPF issue, as shown in the following example:

```
PE1#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
17 routes using 8392 bytes of memory
12 groups, 0.41 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 0
  RP-tree: Forwarding: 0/0/0/0, Other: 344/344/0
  Source: 192.168.1.1/32, Forwarding: 0/-1/0/0, Other: 174/0/174
```

  Do you see multiple paths? If the answer is Yes, the detailed troubleshooting is out of the scope of the document, but the following information might help.

  Normally, there should be a directly connected CE as the RPF neighbor, but if there are multiple paths in the route table, without the **ip multicast vrf** *RED* **multipath** command enabled, Cisco IOS selects the PIM neighbor with the highest IP address as the RPF neighbor. If the **ip multicast vrf** *RED* **multipath** command is enabled, the selection of the RPF neighbor is random and any PIM neighbor can be chosen as the RPF neighbor, based on the multipath hash algorithm towards the C-SRC. Refer to the following URL for more information:

  http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a00804454e1.html

  For example, if the customer has enabled the **maximum-path eibgp** command in the VRF context and the following conditions are met, then multicast stream will fail because of RPF drops:

  – A Cisco IOS release that does not have the update based on **CSCef11357**is being used.

  – Another PE is chosen as the RPF neighbor with tunnel interface as the RPF interface.

  For this case or any other cases where RPF neighbor may be incorrectly chosen, use the static mroutes in the mVRF context to get around the situation. The command is as follows:

  ip mroute vrf *vrf-name ip-address mask NH* | *Interface* |...]

  For a permanent solution, consider the following factors:

  – The available paths in the route table to *ip-address mask.*

  – The IP multicast multipath feature that might be enabled on your router.

  – The Cisco IOS Release that you are running and if need be contact Cisco TAC to get a permanent solution.

- Fragmentation (multicast stream with DF bit set)

  MTU issues are outside the scope of the document. If the PE is unable to fragment the packet because the DF bit is set on the multicast stream sent by the customer, then the **Other drops** counter increment should be shown. Use the following command to do this:

  show ip traffic | inc frag <-Two consecutive sets of output with 30 second intervals

  Check for **couldn't fragment count** in the display when entering this command.

```
PE1#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
17 routes using 8392 bytes of memory
12 groups, 0.41 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 0
  RP-tree: Forwarding: 0/0/0/0, Other: 344/344/0
  Source: 192.168.1.1/32, Forwarding: 0/0/0/0, Other: 174/0/174

PE1#show ip traffic | include fragment
         414 fragmented, 12936 couldn't fragment

PE1#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
17 routes using 8392 bytes of memory
12 groups, 0.41 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 0
  RP-tree: Forwarding: 0/0/0/0, Other: 344/344/0
  Source: 192.168.1.1/32, Forwarding: 0/0/0/0, Other: 178/0/178
PE1#show ip traffic | include fragment
         414 fragmented, 12948 couldn't fragment
```

- OIF NULL

  Enter the **show ip mroute vrf** *vrf_name C-GPA* command

  Verify that the OIL has a non-NULL value for the (C-SRC, C-GPA) entry. This ensures that there are receivers on local or remote PEs for this traffic stream. If there are none, then it is normal for the value to be NULL.

Verify that the Multicast stream is switched to other local sites on PE$_{C\text{-}SRC}$

On PE$_{C\text{-}SRC}$, verify that multicast stream from C-SRC is replicated to other local mVPN sites having receivers on the same PE.

The following output shows two consecutive captures (with 60 secs interval) on PE$_{C\text{-}SRC}$ to help relate the data being received and being sent over to another local receiver.

- **Interface Statistics:** Includes all unicast and multicast traffic sent and received from the CE (on the site having C-RCVR)
- Enter the **show ip pim vrf** *vrf_name* **interface stats** command and look for **Octets out** and **Mpackets out** count and verify that they are increasing in value.
- Enter the **show ip pim vrf** *vrf_name* **interface count** command and look for **Mpackets out** count and verify that it is increasing in value.

- Enter the **show interface** *PE-CE interface* command and check for output traffic and output drops counts.

```
PE1#show ip pim vrf RED interface stats
Interface        Mpackets In    Mpackets Out         Octets In       Octets Out
Serial2/0                  0           29894                 0         30136644
Serial1/0             619881            1646         632994668            82932
Tunnel7                15419          537525          14241576        549777700

PE1#show ip pim vrf RED interface stats
Interface        Mpackets In    Mpackets Out         Octets In       Octets Out
Serial2/0                  0           29978                 0         30222996
Serial1/0             620295            1646         633420260            82932
Tunnel7                15419          537650          14241576        549906200
```

Note The above output can help to verify that multicast traffic is being sent from the PE$_{\text{C-SRC}}$ to CE12. The multicast traffic can belong to any GPA and from any C-SRC address.

```
PE1#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
192.168.112.1    Serial2/0                *   0/63106
192.168.11.1     Serial1/0                *   1590419/6079
172.16.100.1     Tunnel14                 *   1205/130464

PE1#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
192.168.112.1    Serial2/0                *   0/63106
192.168.11.1     Serial1/0                *   1590427/6079
172.16.100.1     Tunnel14                 *   1205/130472

PE1#show ip pim vrf RED interface stats
Interface        Mpackets In    Mpackets Out         Octets In       Octets Out
Serial2/0                  0           63106                 0         64018536
Serial1/0            1590439            6079        1630728292           306312
Tunnel14                1205          130484             60720        134137552

PE1#show ip pim vrf RED interface stats

Interface        Mpackets In    Mpackets Out         Octets In       Octets Out
Serial2/0                  0           63106                 0         64018536
Serial1/0            1590446            6079        1630735488           306312
Tunnel14                1205          130491             60720        134144748
```

Note The above output verifies that multicast traffic is being sent from the PE$_{\text{C-SRC}}$ to CE12. The multicast traffic can belong to any C-CPA and can originate from any C-SRC address.

Note If the PE-CE link is a Frame Relay interface, then look for **broadcasts sent/dropped 9/0** under the **show interface** *FacingCE12* command on the PE$_{\text{C-SRC}}$ router. There have been numerous cases in which output drops occur in a Frame Relay broadcast queue.

Verifying that the Traffic is being Received on PE$_{\text{C-RCVR}}$ over the Global Default MDT Group Entry

Enter the **show ip mroute** *mdt-GPA* command on PE$_{\text{C-RCVR}}$.Verify that **Z** flag is set on the (PE$_{\text{C-SRC}}$, mdt-GPA) entry if available or (*,C-GPA) entry. Use the following output to verify that traffic is received and the counts correspond to what is being sent from PE$_{\text{C-SRC}}$:

s**how ip mrout***e mdt-GPA* **active 1**

Two consecutive captures 60 seconds apart

show ip mroute *mdt-GPA* **count**

show interface tunnel *X*

The following example shows the output of the **show ip mroute** command:

```
PE3#show ip mroute 239.0.0.10 active 1
Active IP Multicast Sources - sending >= 1 kbps
Group: 239.0.0.10, (?)
   Source: 172.16.100.1 (?)
     Rate: 3 pps/26 kbps(1sec), 26 kbps(last 0 secs), 5 kbps(life avg)

PE3#show ip mroute 239.0.0.10 count
IP Multicast Statistics
11 routes using 5776 bytes of memory
5 groups, 1.20 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 239.0.0.10, Source count: 3, Group pkt count: 3991374
  RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0
  Source: 172.16.100.1/32, Forwarding: 3861857/3/1042/26, Other: 3861857/0/0
  Source: 172.16.100.2/32, Forwarding: 87006/1/242/0, Other: 87006/0/0
  Source: 172.16.100.3/32, Forwarding: 42511/0/76/0, Other: 42515/0/4

PE3#show ip mroute 239.0.0.10 count
IP Multicast Statistics
11 routes using 5776 bytes of memory
5 groups, 1.20 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 239.0.0.10, Source count: 3, Group pkt count: 3991415
  RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0
  Source: 172.16.100.1/32, Forwarding: 3861896/3/1042/26, Other: 3861896/0/0
  Source: 172.16.100.2/32, Forwarding: 87006/0/242/0, Other: 87006/0/0
  Source: 172.16.100.3/32, Forwarding: 42513/0/76/0, Other: 42517/0/4

PE3#show int tunnel 0 | include drop|packet
  Checksumming of packets disabled,  fast tunneling enabled
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  5 minute input rate 25000 bits/sec, 3 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     4198720 packets input, 19321738 bytes, 0 no buffer
     42526 packets output, 3262530 bytes, 0 underruns

PE3#show int tunnel 0 | include drop|packet
  Checksumming of packets disabled,  fast tunneling enabled
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0
  5 minute input rate 25000 bits/sec, 3 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     4198789 packets input, 19388954 bytes, 0 no buffer
     42526 packets output, 3262530 bytes, 0 underruns
```

Note The (PE C-SRC, mdt-GPA) entry might not be seen on $PE_{C\text{-}RCVR}$ if multicast in the SP-Core is set for bidir or spt threshold-infinity is configured on all the PEs. In such a case, the best option is to track the counters on the RP-tree.

Failure Scenario: Traffic Not Being Received on $PE_{C\text{-}RCVR}$ over the Global default MDT Entry

If we do not receive the multicast traffic stream over the default mdt-GPA, trace the physical path between $PE_{C\text{-}SRC}$ and $PE_{C\text{-}RCVR}$ identify all the hops (P or PE routers) and find out if any of the intermediate hops is dropping the stream.

Use the **mtrace** command for troubleshooting:

```
PE3# mtrace <PEC-Src> < PEC-Rcvr> [mdt-GPA]

PE3# mtrace <PEC-Src>
```

For more information about the **mtrace** command, refer to Using the Mtrace Tool in Troubleshooting Multicast Issues, page 3-7.

Note Use default mdt source IP addresses for $PE_{C\text{-}SRC}$ and $PE_{C\text{-}RCVR}$.

The following example shows output of the **mtrace** command:

```
PE3#mtrace 172.16.100.1 172.16.100.3 239.0.0.10
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.100.3 via group 239.0.0.10
From source (?) to destination (?)
Querying full reverse path...
 0  172.16.100.3
-1  172.16.100.3 PIM  [172.16.100.1/32]
-2  172.16.3.13 PIM  [172.16.100.1/32]
-3  172.16.13.11 PIM  [172.16.100.1/32]
-4  172.16.1.1 PIM  [172.16.100.1/32]
-5  172.16.100.1
```

Jump on each hop from $PE_{C\text{-}SRC}$ to $PE_{C\text{-}RCVR}$ for verifying that the forwarding counts are increasing in **show ip mroute** *mdt-GPA* command (two consecutive sets of output 60 seconds apart).

The following example shows output from a P router (172.16.13.11):

```
P11#show ip mroute 239.0.0.10 count
IP Multicast Statistics
15 routes using 7704 bytes of memory
7 groups, 1.14 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 239.0.0.10, Source count: 3, Group pkt count: 3962805
  RP-tree: Forwarding: 3/0/74/0, Other: 3/0/0
  Source: 172.16.100.1/32, Forwarding: 3869629/3/1042/25, Other: 3869629/0/0
  Source: 172.16.100.2/32, Forwarding: 65000/1/301/0, Other: 65656/656/0
  Source: 172.16.100.3/32, Forwarding: 28173/1/76/0, Other: 28173/0/0

P11#show ip mroute 239.0.0.10 count
IP Multicast Statistics
15 routes using 7704 bytes of memory
7 groups, 1.14 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)
```

```
Group: 239.0.0.10, Source count: 3, Group pkt count: 3962823
  RP-tree: Forwarding: 3/0/74/0, Other: 3/0/0
  Source: 172.16.100.1/32, Forwarding: 3869646/3/1042/26, Other: 3869646/0/0
  Source: 172.16.100.2/32, Forwarding: 65001/1/301/0, Other: 65657/656/0
  Source: 172.16.100.3/32, Forwarding: 28173/0/76/0, Other: 28173/0/0
```

Note The (PE C-SRC, mdt-GPA) entry might not be seen on PE$_{\text{C-RCVR}}$ if multicast in the SP-Core is set for bidir or spt threshold-infinity is configured on all the PEs. In this case, the best option is to track the counters on the RP-tree.

Verifying that the (C-SRC, C-GPA) Traffic is Decapsulated on PE$_{\text{C-RCVR}}$ and sent to the mVRF

Use the following commands to check if the multicast stream switched from PE$_{\text{C-RCVR}}$ to CE3:

show ip pim vrf *vrf_name* **interface stat**

show ip mroute vrf *vrf_name* **mroute C-GPA count** (Check for forwarding counts)
show ip mroute vrf *vrf_name* **mroute C-GPA active 1**

show interface *PE-CE interface* (check for output traffic and look for any drops)

The following example shows output of the mVRF multicast table:

```
PE3#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
9 routes using 5024 bytes of memory
5 groups, 0.80 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 123
  RP-tree: Forwarding: 9/0/804/0, Other: 9/0/0
  Source: 192.168.1.1/32, Forwarding: 114/2/524/9, Other: 114/0/0


PE3#
PE3#show ip mroute vrf RED 225.1.1.10 count
IP Multicast Statistics
9 routes using 5024 bytes of memory
5 groups, 0.80 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.10, Source count: 1, Group pkt count: 127
  RP-tree: Forwarding: 9/0/804/0, Other: 9/0/0
  Source: 192.168.1.1/32, Forwarding: 118/2/524/9, Other: 118/0/0
```

The following output shows two consecutive captures (with 60 secs interval) on PE3 (PE$_{\text{C-RCVR}}$) to help relate the data being sent and received from another local receiver.

```
PE3#show ip pim vrf RED interface stats
Interface        Mpackets In    Mpackets Out         Octets In        Octets Out
Serial1/0                  0         3998866                 0        4095804344
Tunnel0              4098008               0        4197609496                 0

PE3#show ip pim vrf RED interface stats
Interface        Mpackets In    Mpackets Out         Octets In        Octets Out
Serial1/0                  0         3998874                 0        4095812568
Tunnel0              4098020               0        4197621832                 0
```

Note The above output verifies that multicast traffic is being sent from PE3 (PE$_{\text{C-RCVR)}}$ to CE3. The multicast traffic can belong to any C- GPA and can originate from any C-SRC address.

Note If the PE-CE link is a Frame-Relay interface, check for **broadcasts sent/dropped 9/0** under the **show interface** *interface facing CE12* command on the PE$_{\text{C-SRC}}$ router. There have been numerous cases in which the output drops occur in the Frame Relay broadcast queue.

Sample Failure Scenario

If the traffic counters are not incrementing, the router might be dropping the packets due to TTL issues.

Step 1 Use the **show ip pim vrf** *vrf_name* **interface stats** command to view the traffic counters and check for packet drops. The (C-SRC, C-GPA) entry might be missing in **show ip mroute vrf** *vrf_name CGPA* **count**.

```
PE3#show ip mroute vrf RED 225.1.1.3 count
IP Multicast Statistics
8 routes using 4580 bytes of memory
4 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.3, Source count: 1, Group pkt count: 3934
  RP-tree: Forwarding: 1/0/1028/0, Other: 8/0/7
  Source: 192.168.2.1/32, Forwarding: 3933/0/1028/0, Other: 3996/0/63

PE3#show ip mroute vrf RED 225.1.1.3 count
IP Multicast Statistics
8 routes using 4580 bytes of memory
4 groups, 1.00 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.3, Source count: 1, Group pkt count: 3934
  RP-tree: Forwarding: 1/0/1028/0, Other: 8/0/7
  Source: 192.168.2.1/32, Forwarding: 3933/0/1028/0, Other: 4024/0/91


PE3#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
192.168.33.1     Serial1/0                *   0/4774568
172.16.100.3     Tunnel3                  *   949805/0

PE3#show ip pim vrf RED interface count

State: * - Fast Switched, D - Distributed Fast Switched
       H - Hardware Switching Enabled
Address          Interface                FS  Mpackets In/Out
192.168.33.1     Serial1/0                *   0/4774568
172.16.100.3     Tunnel3                  *   949863/0
```

```
PE3#show ip pim vrf RED interface stat

Interface        Mpackets In    Mpackets Out         Octets In        Octets Out
Serial1/0                  0         4774569                 0         586481216
Tunnel3               949884               0         968708852                 0

PE3#show ip pim vrf RED interface stat

Interface        Mpackets In    Mpackets Out         Octets In        Octets Out
Serial1/0                  0         4774569                 0         586481216
Tunnel3               949900               0         968725300                 0
PE3#
```

Step 2 Check for TTL issues by issuing the **show ip traffic | inc bad hop** command:

```
PE3#show ip traffic | include bad hop
         0 format errors, 0 checksum errors, 23456 bad hop count

PE3#show ip traffic | include bad hop
         0 format errors, 0 checksum errors, 23456 bad hop count
```

Step 3 Note that you could soon lose the (C-SRC, C-GPA) entry from PE3:

```
PE3#show ip mroute vrf RED 225.1.1.3 count
IP Multicast Statistics
6 routes using 3676 bytes of memory
4 groups, 0.50 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 225.1.1.3, Source count: 0, Group pkt count: 1
  RP-tree: Forwarding: 1/0/1028/0, Other: 136/0/135

PE3#show ip pim vrf RED interface ser1/0 detail
Serial1/0 is up, line protocol is up
  Internet address is 192.168.33.1/24
  Multicast switching: fast
  Multicast packets in/out: 0/4774577
  Multicast boundary: not set
  Multicast TTL threshold: 2
  PIM: enabled
    PIM version: 2, mode: sparse
    PIM DR: 0.0.0.0
    PIM neighbor count: 1
    PIM Hello/Query interval: 30 seconds
    PIM NBMA mode: disabled
    PIM ATM multipoint signalling: disabled
    PIM domain border: disabled
    PIM neighbors rpf proxy capable: TRUE
  Multicast Tagswitching: disabled
```

CHAPTER 3

Using Device Instrumentation for mVPN Troubleshooting

Device Instrumentation (DI) is the set of Cisco IOS embedded tools that allow for an intelligent, remote, and, in some cases, proactive management of network devices. This section presents the various DI features you can use when troubleshooting mVPNs.

Note that the focus in this section is on mVPN-specific features and not on the whole set of DI tools available for IP multicast.

The following features are covered:

- Syslogs
- mVPN Management Information Base (MIB)
- VPN-MIB
- Other PIM, IGMP, IGP, interface, and BGP related MIBs
- Multicast NetFlow (egress and ingress) cache and data export
- mtrace tool
- mrinfo tool
- Ping tool

The range of possible mVPN issues are often caused either by misconfigurations, changes of design, or in some cases by platform, or Cisco IOS-specific behaviors.

Deploying and leveraging DI allows you to passively monitor and proactively retrieve information when necessary about the mVPN setup and potential configuration changes. The following sections describe the benefits available by using the DI.

Using DI When Adding an mVPN Customer Site

When you fully configure an mVPN for the first time or add a customer site, you can watch for Syslog messages, along with the **show** command output. This enables you to monitor events and look for mVPN-related messages (for example, MTI interface UP, PE-PE neighbors UP over the MTI, PE-CE neighbor UP over the mVRF interface, and DR election over the MTI).

```
*Aug 31 19:30:17.331: %BGP-5-ADJCHANGE: neighbor 172.16.100.100 Up
*Aug 31 19:32:06.523: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.3 on
interface Tunnel0 (vrf RED)
```

```
*Aug 31 19:32:06.523: %PIM-5-NBRCHG: neighbor 172.16.100.1 UP on interface Tunnel0 (vrf
RED)
```

If the above messages are not generated when adding a new mVPN customer, it is an indication that the configuration is not working as intended.

At the end of a set of configurations, as a last check after provisioning a service, you can poll specific values and check their accuracy automatically and proactively by using the mVPN MIB and other available MIBs, including the following:

- IF-MIB, namely the interface UP/DOWN notification (unicast and multicast).
- VPN-MIB, namely the VRF UP/DOWN notification (unicast and multicast).
- PIM MIBs (for instance IPMROUTE-STD-MIB, CISCO-IPMROUTE-MIB, PIM-MIB, CISCO-PIM-MIB, and IGMP-MIB/IGMP-STD-MIB).

For instance, the mVPN MIB provides mVPN specific information located on a PE, such as Multicast Distribution Tree (MDT) source and group addresses for an mVPN, BGP next hop information, and state of the local mVRFs.

Caution We recommend logging the debug and syslog messages in the buffer rather than displaying them on the console. We also recommend leveraging an SNMP and Syslog server for logging the SNMP traps and syslog messages, respectively. Be careful when turning on debugs in a production environment so as to not bring the router down.

Using SNMP for Monitoring an Existing mVPN

When an mVPN setup is up and running, We highly recommend that you keep monitoring the various events, turning on and logging SNMP notifications for the appropriate MIBs, and logging syslog messages for later retrieval.

Monitoring the events allows for early detection of potential issues, and logging events allows you to consult the history of events and changes, and to track the potential causes of a problem.

Note It is good practice to set the clock of your routers and enable service timestamps to have a time reference on all your devices for syslog and debug messages. It is also strongly recommended to sync up all routers clocks using NTP.

Note Remember that not all events represent immediate problems; many alarm messages have a warning-only level or represent events that are normal.

The CiscoMvpnMvrfChange notification from the mVPN MIB is triggered mainly for events such as the following:

- Tunnel interface UP.
- Tunnel interface DOWN (for example, at the time of creation or deletion of the multicast VPN routing and forwarding instance (MVRF), or a change of the MDT IP address for the MVRF).

The following command allows SNMP traps to be sent out to an SNMP server:

```
PE1(config)#snmp-server host 172.16.100.100 version 2c mvpnTG mvpn
```

The following output results were obtained by turning on **debug snmp** on a device. This is not the normal procedure to capture SNMP notifications, because such SNMP messages would be sent to an SNMP server. We illustrate, by the following example, the deletion of an MVRF by mistake.

```
PE1(config-vrf)#no mdt default 239.11.11.11
% A new tunnel ID may be used if the default mdt is reconfigured for this VRF.
Tunnel interface was deleted. Partial configuration may reappear on reuse.
PE1(config-vrf)#
*Sep  2 03:38:56.331: %LINK-5-CHANGED: Interface Tunnel0, changed state to
administratively down
*Sep  2 03:38:57.359: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel0, changed
state to down
PE1(config-vrf)#
*Sep  2 03:38:59.403: SNMP: Queuing packet to 172.16.100.100
*Sep  2 03:38:59.403: SNMP: V2 Trap, reqid 2, errstat 0, erridx 0
 sysUpTime.0 = 83040
 internet.6.3.1.1.4.1.0 = ciscoMvpnMvrfChange
 ciscoMvpnGenericEntry.1.11.116.101.115.116.109.86.80.78.77.73.66 = 3
*Sep  2 03:38:59.683: SNMP: Packet sent via UDP to 172.16.100.100
PE1(config-vrf)# mdt default 239.11.11.12
PE1(config-vrf)#
*Sep  2 03:39:14.491: SNMP: Queuing packet to 172.16.100.100
*Sep  2 03:39:14.491: SNMP: V2 Trap, reqid 2, errstat 0, erridx 0
 sysUpTime.0 = 84549
 internet.6.3.1.1.4.1.0 = ciscoMvpnMvrfChange
 ciscoMvpnGenericEntry.1.11.116.101.115.116.109.86.80.78.77.73.66 = 3
*Sep  2 03:39:14.767: SNMP: Packet sent via UDP to 172.16.100.100
*Sep  2 03:39:15.419: %LINEPROTO-5-UPDOWN: Line protocol on Interface Tunnel1, changed
state to up
*Sep  2 03:39:16.459: %PIM-5-DRCHG: DR change from neighbor 0.0.0.0 to 172.16.100.1 on
interface Tunnel1 (vrf testmVPNMIB)
```

Other information for monitoring the mVPN pertains to BGP adjacencies, such as CISCO-BGP4-MIB, and BGP4-MIB; and interface, link and IGP events, such as interface MIB, OSPF, RIP, EIGRP MIBs, and VRF events such as VPN-MIB.

Multicast NetFlow can also be used for monitoring. The ingress and egress multicast flows and RPF failure information are stored in the Netflow cache and can be exported for further processing and monitoring. These captured flows also constitute a good troubleshooting tool because they represent a great source of information, as described in more detail in the following sections.

Overview of the Benefits of the mVPN MIB

The tables provided as part of the mVPN MIB allow for the following query capabilities:

- The state of the MVRFs, including the name of the MVRF, whether it is active, and the number of active multicast-enabled interfaces.
- MDT default group address and encapsulation information.
- Next hop information used to receive Border Gateway Protocol (BGP) MDT updates for Source Specific Multicast (SSM) mode.
- Traffic threshold that determines switchover to an MDT data group.
- Type of MDT group being used for a given (S,G) multicast route entry that exists on each configured MVRF, source address, and group address of the multicast route entry.
- Source and group address used for encapsulation.
- Information on MDT data groups currently joined.

- Information on mVPN-specific MDT tunnels present in the device.
- Trap notification enabled on the router.

Using DI to Optimize Troubleshooting Steps

Multicast Netflow, mtrace, mrinfo, and ping can be very helpful in understanding what is happening in the network, because they allow the users to visualize the packet path and determine on which router the problem exists.

Naturally, as mentioned throughout this document, **show** commands are also key to troubleshooting an mVPN setup, especially the packet counters and mrouter counters, which allow us to visualize the data flows and their evolution.

Multicast NetFlow

You can leverage multicast NetFlow for monitoring the various multicast flows within the core or at the mVRF level. Retrieving the information while troubleshooting may help solve problems more quickly, because the stored information contain the history of events, and the cached information is constantly updated.

In general, NetFlow is used for providing usage and billing information for security reasons (monitoring). Multicast NetFlow allows for visualizing the various flows and capturing Reverse Path Forwarding (RPF) failure statistics, which allows you to check on the traffic and visualize where the flow(s) stopped.

Information can be exported if version 9 of NetFlow is used. With version 5, the information can be visualized locally in the cache of the router.

- For an ingress flow, the cache logs the flows as source IP address, group address, and average of replicated packets, which is the average duplication factor (in packets and in bytes, noted IPM X Packets, and Y bytes)
- The outgoing flow is accounted for as a separate flow.

```
routerA#sh ip cache verbose flow
IP packet size distribution (192805 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .572 .149 .039 .022 .011 .007 .012 .027 .003 .003 .004 .003 .000  .003

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .002 .001 .000 .026 .106 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 4456704 bytes
  9 active, 65527 inactive, 20 added
  3652 ager polls, 0 flow alloc failures
  Active flows timeout in 30 minutes
  Inactive flows timeout in 15 seconds
IP Sub Flow Cache, 270472 bytes
  1 active, 16383 inactive, 0 added, 0 added to flow
  0 alloc failures, 0 force free
  1 chunk, 0 chunks added
  last clearing of statistics 00:06:11

Protocol         Total    Flows   Packets Bytes  Packets Active(Sec)  Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec     /Flow      /Flow
TCP-other           12      0.0         1    48      0.0       0.0      15.5
```

```
UDP-other            9      0.0         1   127      0.0       0.0      15.4
Total:              21      0.0         1    74      0.0       0.0      15.5

SrcIf          SrcIPaddress    DstIf          DstIPaddress    Pr TOS Flgs   Pkts
Port Msk AS                    Port Msk AS    NextHop              B/Pk     Active
IPM: OPkts    OBytes

Se4/0:0        1.192.3.100     Fa1/0          1.1.160.100     06 68  18      78K
0017 /0  0                     84E5 /0  0     0.0.0.0                53     920.3
Se4/0:0        1.199.240.2     Null           224.0.0.10      58 C0  10      201
0000 /0  0                     0000 /0  0     0.0.0.0               100     928.3
IPM:     0       0
```

The IPM header lines in the preceding example confirm that Multicast NetFlow is supported. Also note that in this case the router does not need to duplicate the multicast traffic, which is why the IPM counters show a value of 0. The above last entry shows one multicast flow seen from 1.199.240.2 received on Se4 destined for 224.0.0.10 that was terminated in the router.

Multicast Netflow: RPF Monitoring

You can obtain statistics about RPF failures by leveraging Multicast NetFlow, which provides statistics to account for all packets that fail the RPF check and that are dropped. This helps to identify patterns and to narrow down the list of possible problems.

The following commands turn on multicast egress and ingress NetFlow. Note that ingress NetFlow does not appear on the running configuration after being enabled.

- Multicast NetFlow

```
(Global config)#ip multicast netflow rpf-failure
```

- Configuration per interface:

```
#ip route-cache flow
#ip multicast netflow egress
sh ip cache verbose flow
```

Note NetFlow accounts for only fast switching or multicast distributed fast switching (MDFS)-switched packets, but not for process-switched flows.

Multicast Netflow: Performance and Memory Considerations

Be aware that egress NetFlow slightly affects the performance of a router device. There is some memory impact because information is stored in the cache of the router. Note that Multicast NetFlow adds two more fields to the traditional flow entry. For more information, refer to the following URL: http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a0080438df4.html

Using mrinfo for Troubleshooting Multicast Issues

The **mrinfo** command shows multicast neighbor router information, router capabilities and code version, multicast interface information, TTL thresholds, metrics, protocol, and status. It is useful when you need to verify multicast neighbors, and confirm that bi-directional neighbor adjacency exists.

The syntax for the **mrinfo** command is as follows:

mrinfo [*hostname-or-address*] [*source-address-or-interface*]

The **mrinfo** command requests the configuration information from the multicast router, which can be either an IP address or a system name. The **mrinfo** command sends out the ASK_NEIGHBORS IGMP message to the specified multicast router. When the router receives the request, it responds with the version number and a list of their neighboring multicast router addresses. Depending on the **mrinfo** version number, the response to the **mrinfo** request may contain additional information such as metrics, thresholds, and flags from the multicast router sending back its configuration information. If the multicast router is not specified, the request is sent to the local router. If *source-address-or-interface* is specified, this becomes the source address used on mrinfo requests. When it is omitted, the source address is based on the outbound interface for the destination.

The configuration information for each interface is displayed in the following format:

multicast-router [*IOS version mainline*] [**FLAG**]

interface_addr -> *neighbor_addr* (*neighbor_name*) [metrics/thresh/*flags*]

If there are multiple neighbor routers on one interface, they are all reported in the output.

The possible **FLAG** values are as follows:

- P—Prune-capable
- M—mtrace-capable
- S—SNMP-capable
- A—Auto-RP-capable

The possible interface flag values are as follows:

- Leaf—There is no PIM neighbor on the interface.
- Querier—The local router is the querier of the subnet.
- PIM—If the neighbor is a PIM neighbor.
- Tunnel—Neighbors are reached via tunnel.

The following example shows output for the **mrinfo** command (refer to Figure 1-1 on page 1-2 for the topology):

PE1 querying CE1:

```
PE$_{C-SRC}$#mrinfo vrf RED 192.168.11.2 serial 1/0
192.168.11.2 [version 12.0] [flags: PMA]:
192.168.1.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
192.168.0.1 -> 0.0.0.0 [1/0/pim/querier/leaf]
192.168.11.2 -> 192.168.11.1 [1/0/pim]
```

PE1 querying CE2:

```
PE$_{C-SRC}$#mrinfo vrf RED 192.168.22.2 serial 1/0
192.168.22.2 [version 12.0] [flags: PMA]:
192.168.2.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
192.168.202.2 -> 192.168.202.202 [1/0/pim]
192.168.22.2 -> 192.168.22.1 [1/0/pim]
```

PE1 querying PE2:

```
PE$_{C-SRC}$#mrinfo vrf RED 192.168.22.1 serial 1/0
192.168.22.1 [version 12.0] [flags: PMA]:
192.168.22.1 -> 192.168.22.2 [1/0/pim]
192.168.212.1 -> 0.0.0.0 [1/0/pim/down/leaf]
```

```
172.16.100.2 -> 172.16.100.1 [1/0/tunnel/pim]
172.16.100.2 -> 172.16.100.3 [1/0/tunnel/pim]
```

CE1 querying PE1:

```
CE1#mrinfo 192.168.11.1
192.168.11.1 [version 12.0] [flags: PMA]:
192.168.11.1 -> 192.168.11.2 [1/0/pim]
192.168.112.1 -> 192.168.112.2 [1/0/pim]
172.16.100.1 -> 172.16.100.3 [1/0/tunnel/pim]
172.16.100.1 -> 172.16.100.2 [1/0/tunnel/pim]
```

CE1 querying CE2 and CE3:

```
CE1#mrinfo 192.168.100.3
192.168.100.3 [version 12.0] [flags: PMA]:
192.168.3.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
192.168.33.2 -> 192.168.33.1 [1/0/pim]
CE1#mrinfo 192.168.100.2
192.168.100.2 [version 12.0] [flags: PMA]:
192.168.2.2 -> 0.0.0.0 [1/0/pim/querier/leaf]
192.168.202.2 -> 192.168.202.202 [1/0/pim]
192.168.22.2 -> 192.168.22.1 [1/0/pim]
```

CE1 querying PE2:

```
CE1#mrinfo 192.168.22.1
192.168.22.1 [version 12.0] [flags: PMA]:
192.168.22.1 -> 192.168.22.2 [1/0/pim]
192.168.212.1 -> 0.0.0.0 [1/0/pim/down/leaf]
172.16.100.2 -> 172.16.100.1 [1/0/tunnel/pim]
172.16.100.2 -> 172.16.100.3 [1/0/tunnel/pim]
```

Service providers can use the **ip multicast mrinfo-filter** command to prevent the enterprise customers from querying the service provider and finding all PEs in the mVPN.

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_command_reference_chapter09186a008010ea27.html#wp1093149

Using the Mtrace Tool in Troubleshooting Multicast Issues

The **mtrace** command shows the multicast path from the source to the receiver. This command traces the path between points in the networks, which shows TTL thresholds and delay at each node. When troubleshooting, use the **mtrace** command to find where multicast traffic flow stops, to verify the path of multicast traffic, and to identify sub-optimal paths.

The syntax for the **mtrace** command is as follows: **mtrace** *src IP addr dest IP addr grp addr.*

- Multicast packets to the group and traces RPF path from destination back to the source.
- The **mtrace** command can be used at a completely different router. The source will send the trace results to the issuing router.
- If the group-address is not specified, then the group 244.2.0.1 is used (MBone Audio). In this case, it is called a weak mtrace.

The following examples show how the **mtrace** command can be used to troubleshoot issues with an mVPN network (refer to Figure 1-1 on page 1-2 for the topology used).

```
P-RP#mtrace 172.16.100.1 239.0.0.10
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via group 239.0.0.10
```

```
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM Reached RP/Core [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM [172.16.100.1/32]
-4 172.16.1.1 PIM [172.16.100.1/32]

P-RP#mtrace 172.16.100.1
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM [172.16.100.1/32]
-4 172.16.1.1 PIM [172.16.100.1/32]
-5 172.16.100.1
```

Traffic Failure Scenarios

The following subsections show traffic failure scenarios on various interfaces.

PIM Disabled on PE1-P11 (on P11 interface)

```
P-RP#mtrace 172.16.100.1
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 None No route

P-RP#mtrace 172.16.100.1 239.0.0.10
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via group 239.0.0.10
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM Reached RP/Core [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM No route
```

PIM Static Configuration Missing on P11

```
P-RP#mtrace 172.16.100.1 239.0.0.10
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via group 239.0.0.10
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM Reached RP/Core [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM [172.16.100.1/32]
-4 172.16.1.1 PIM [172.16.100.1/32]

P-RP#mtrace 172.16.100.1
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via RPF
```

```
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM [172.16.100.1/32]
-4 172.16.1.1 PIM [172.16.100.1/32]
-5 172.16.100.1
```

IP Multicast Routing Disabled on P11

```
P-RP#mtrace 172.16.100.1 239.0.0.10
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via group 239.0.0.10
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM Reached RP/Core [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 None Multicast disabled [172.16.100.1/32]

P-RP#mtrace 172.16.100.1
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM Multicast disabled [172.16.100.1/32]
-4 172.16.1.1 PIM [172.16.100.1/32]
```

PIM Disabled on PE1-P11 (on PE1 Interface)

```
P-RP#mtrace 172.16.100.1 239.0.0.10
Type escape sequence to abort.
Mtrace from 172.16.100.1 to 172.16.213.200 via group 239.0.0.10
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM Reached RP/Core [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM [172.16.100.1/32]
-4 172.16.1.1 None Multicast disabled [172.16.100.1/32]

P-RP#mtrace 172.16.100.1
Type escape sequence to abort
Mtrace from 172.16.100.1 to 172.16.213.200 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.213.200
-1 172.16.213.200 PIM [172.16.100.1/32]
-2 172.16.213.13 PIM [172.16.100.1/32]
-3 172.16.13.11 PIM [172.16.100.1/32]
-4 172.16.1.1 PIM Multicast disabled [172.16.100.1/32]
```

PE$_{C-SRC}$#**mtrace 172.16.100.200**

```
Type escape sequence to abort.
Mtrace from 172.16.100.200 to 172.16.1.1 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.1.1
-1 172.16.1.1 PIM [172.16.100.200/32]
-2 172.16.1.11 PIM [172.16.100.200/32]
```

```
-3 172.16.13.13 PIM [172.16.100.200/32]
-4 172.16.213.200 PIM [172.16.100.200/32]
-5 172.16.100.200
```

```
PE$_{C-SRC}$#mtrace 172.16.100.2 172.16.100.3
Type escape sequence to abort.
Mtrace from 172.16.100.2 to 172.16.100.3 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.100.3
-1 172.16.3.3 PIM [172.16.100.2/32]
-2 172.16.3.13 PIM [172.16.100.2/32]
-3 172.16.23.12 PIM [172.16.100.2/32]
-4 172.16.2.2 PIM [172.16.100.2/32]
-5 172.16.100.2
```

```
PE$_{C-SRC}$#mtrace 172.16.100.3 172.16.100.2
Type escape sequence to abort.
Mtrace from 172.16.100.3 to 172.16.100.2 via RPF
From source (?) to destination (?)
Querying full reverse path...
0 172.16.100.2
-1 172.16.2.2 PIM [172.16.100.3/32]
-2 172.16.2.12 PIM [172.16.100.3/32]

-3 172.16.23.13 PIM [172.16.100.3/32]
-4 172.16.3.3 PIM [172.16.100.3/32]
-5 172.16.100.3
```

```
*Aug 29 19:22:31.254: IGMP(0): Send Mtrace request (g: 0.0.0.0, s: 172.16.100.3, d:
172.16.100.2)
*Aug 29 19:22:31.254: to 172.16.100.2, nhops 32, ttl 64
*Aug 29 19:22:31.362: IGMP(0): Received Mtrace response from 172.16.100.3 (Serial0/0)
```

APPENDIX A

Sample Configurations for Testbed

Below are sample configurations for the PE, P-RP, C-RP, and RR.

The first set of output is taken for the MDT Bidir setup with P-RP set statically and AutoRP used on the customer network with PIM SM mode.

The second set of output is taken for the MDT PIM SM setup with the same RP and Customer parameters as for BIDIR otherwise.

Note that the RR, C-RP and CEs configurations are the same in the Bidir and PIM-SM test setups.

BIDIR

```
PE1#sh run
Building configuration...

Current configuration : 2112 bytes
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname PE1
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
ip vrf RED
 rd 1:100
 route-target export 1:100
 route-target import 1:100
 mdt default 239.0.0.10
!
ip multicast-routing
ip multicast-routing vrf RED
mpls label protocol ldp
tag-switching tdp router-id Loopback0 force
!
!
!
interface Loopback0
```

```
 ip address 172.16.100.1 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Serial0/0
 ip address 172.16.1.1 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
 no fair-queue
!
interface Serial1/0
 ip vrf forwarding RED
 ip address 192.168.11.1 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Serial2/0
 ip vrf forwarding RED
 ip address 192.168.112.1 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
router ospf 1
 router-id 172.16.100.1
 log-adjacency-changes
 mpls ldp autoconfig
 network 172.16.0.0 0.0.255.255 area 0
!
router bgp 1
 bgp router-id 172.16.100.1
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 172.16.100.100 remote-as 1
 neighbor 172.16.100.100 update-source Loopback0
 !
 address-family ipv4 mdt
neighbor 172.16.100.100 activate
 neighbor 172.16.100.100 send-community both
 exit-address-family
 !
 address-family vpnv4
 neighbor 172.16.100.100 activate
 neighbor 172.16.100.100 send-community both
 exit-address-family
 !
 address-family ipv4 vrf RED
 neighbor 192.168.11.2 remote-as 65000
 neighbor 192.168.11.2 activate
 neighbor 192.168.11.2 as-override
 neighbor 192.168.112.2 remote-as 65000
 neighbor 192.168.112.2 activate
 neighbor 192.168.112.2 as-override
 no auto-summary
 no synchronization
 exit-address-family
!
ip classless
!
ip pim bidir-enable
ip pim rp-address 172.16.100.200 1 override bidir
ip pim vrf RED autorp listener
!
access-list 1 permit 239.0.0.0 0.0.0.255
!
```

```
!
control-plane
!
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 privilege level 15
 no login
!
no cns aaa enable
end
```

```
P-RP#sh run
Building configuration...

Current configuration : 1199 bytes
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname P-RP
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
ip multicast-routing
mpls label protocol ldp
tag-switching tdp router-id Loopback0 force
!
!
!
interface Loopback0
 ip address 172.16.100.200 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Serial0/0
 ip address 172.16.212.200 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
 no fair-queue
!
interface Serial1/0
 ip address 172.16.213.200 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
router ospf 1
 router-id 172.16.100.200
 log-adjacency-changes
 mpls ldp autoconfig
 network 172.16.0.0 0.0.255.255 area 0
!
ip classless
!
ip pim bidir-enable
ip pim rp-address 172.16.100.200 1 override bidir
!
access-list 1 permit 239.0.0.0 0.0.0.255
!
!
control-plane
!
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 privilege level 15
 no login
!
no cns aaa enable
end

P-RP#
```

```
RR#sh run
Building configuration...

Current configuration : 1744 bytes
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname RR
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
!
!
!
interface Loopback0
 ip address 172.16.100.100 255.255.255.255
 no ip directed-broadcast
!
interface Serial0/0
 ip address 172.16.111.100 255.255.255.0
 no ip directed-broadcast
 no fair-queue
!
interface Serial1/0
```

```
 ip address 172.16.113.100 255.255.255.0
 no ip directed-broadcast
!
router ospf 1
 max-metric router-lsa
 log-adjacency-changes
 network 172.16.0.0 0.0.255.255 area 0
!
router bgp 1
 bgp router-id 172.16.100.100
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor ibgp peer-group
 neighbor ibgp remote-as 1
 neighbor ibgp update-source Loopback0
 neighbor 172.16.100.1 peer-group ibgp
 neighbor 172.16.100.2 peer-group ibgp
 neighbor 172.16.100.3 peer-group ibgp
 !
 address-family ipv4 mdt
 neighbor ibgp activate
 neighbor ibgp send-community both
 neighbor ibgp route-reflector-client
 neighbor 172.16.100.1 peer-group ibgp
 neighbor 172.16.100.2 peer-group ibgp
 neighbor 172.16.100.3 peer-group ibgp
 exit-address-family
 !
 address-family vpnv4
 neighbor ibgp activate
 neighbor ibgp send-community both
 neighbor ibgp route-reflector-client
 neighbor 172.16.100.1 peer-group ibgp
 neighbor 172.16.100.2 peer-group ibgp
 neighbor 172.16.100.3 peer-group ibgp
 exit-address-family
!
ip classless
!
!
!
!
control-plane
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 privilege level 15
 no login
!
no cns aaa enable
end

RR#
```

```
CE1#sh run
Building configuration...

Current configuration : 1011 bytes
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname CE1
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
ip multicast-routing
!
!
!
interface Loopback0
 ip address 192.168.100.1 255.255.255.255
 no ip directed-broadcast
!
interface Ethernet0/0
 ip address 192.168.1.2 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Serial1/0
 ip address 192.168.11.2 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
 no fair-queue
!
router bgp 65000
 no synchronization
 bgp log-neighbor-changes
 redistribute connected
 neighbor 192.168.11.1 remote-as 1
 no auto-summary
!
ip classless
!
ip pim autorp listener
!
!
!
control-plane
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 login
!
no cns aaa enable
end

CE1#
```

```
C-RP#sh run
Building configuration...

Current configuration : 1082 bytes
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname C-RP
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
ip multicast-routing
!
!
!
interface Loopback0
 ip address 192.168.100.202 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Ethernet0/0
 ip address 192.168.202.202 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
router eigrp 65000
 !
 address-family ipv4
 network 192.168.0.0 0.0.255.255
 no auto-summary
 exit-address-family
!
ip classless
!
ip pim send-rp-announce Loopback0 scope 16 group-list 1
ip pim send-rp-discovery Loopback0 scope 16
!
access-list 1 deny   224.0.1.39
access-list 1 deny   224.0.1.40
access-list 1 permit 225.0.0.0 0.255.255.255
!
!
control-plane
!
line con 0
line aux 0
line vty 0 4
 privilege level 15
 no login
!
no cns aaa enable
end
```

PIM SM with SPT 0

```
PE2#sh run
Building configuration...

Current configuration : 2000 bytes
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname PE2
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
ip vrf RED
 rd 2:100
 route-target export 1:100
 route-target import 1:100
 mdt default 239.0.0.10
!
ip multicast-routing
ip multicast-routing vrf RED
mpls label protocol ldp
tag-switching tdp router-id Loopback0 force
!
!
!
interface Loopback0
 ip address 172.16.100.2 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Serial0/0
 ip address 172.16.2.2 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
 no fair-queue
!
interface Serial1/0
 ip vrf forwarding RED
 ip address 192.168.22.1 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
router ospf 1
 router-id 172.16.100.2
 log-adjacency-changes
 mpls ldp autoconfig
 network 172.16.0.0 0.0.255.255 area 0
!
router bgp 1
 bgp router-id 172.16.100.2
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 172.16.100.100 remote-as 1
 neighbor 172.16.100.100 update-source Loopback0
 !
 address-family ipv4
 neighbor 172.16.100.100 activate
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family ipv4 mdt
 neighbor 172.16.100.100 activate
 neighbor 172.16.100.100 send-community both
 exit-address-family
 !
 address-family vpnv4
 neighbor 172.16.100.100 activate
 neighbor 172.16.100.100 send-community both
 exit-address-family
 !
 address-family ipv4 vrf RED
 neighbor 192.168.22.2 remote-as 65000
 neighbor 192.168.22.2 activate
 neighbor 192.168.22.2 as-override
 no auto-summary
 no synchronization
 exit-address-family
!
ip classless
!
ip pim rp-address 172.16.100.200 1
ip pim ssm default
ip pim vrf RED autorp listener
ip pim vrf RED spt-threshold 0
!
access-list 1 permit 239.0.0.0 0.0.0.255
!
!
control-plane
!
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 login
!
no cns aaa enable
end

PE2#
```

```
P-RP#sh run
Building configuration...

Current configuration : 1046 bytes
```

```
!
version 12.0
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname P-RP
!
boot-start-marker
boot-end-marker
!
!
clock timezone PST -8
ip subnet-zero
ip cef
no ip domain-lookup
ip multicast-routing
mpls label protocol ldp
tag-switching tdp router-id Loopback0 force
!
!
!
interface Loopback0
 ip address 172.16.100.200 255.255.255.255
 no ip directed-broadcast
 ip pim sparse-mode
!
interface Serial0/0
 ip address 172.16.212.200 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
 no fair-queue
!
interface Serial1/0
 ip address 172.16.213.200 255.255.255.0
 no ip directed-broadcast
 ip pim sparse-mode
!
router ospf 1
 router-id 172.16.100.200
 log-adjacency-changes
 mpls ldp autoconfig
 network 172.16.0.0 0.0.255.255 area 0
!
ip classless
!
ip pim rp-address 172.16.100.200 1
ip pim ssm default
!
access-list 1 permit 239.0.0.0 0.0.0.255
!
!
control-plane
!
!
line con 0
 exec-timeout 0 0
line aux 0
line vty 0 4
 login
!
no cns aaa enable
end
```

Related Documents

- Basic Multicast Troubleshooting Tools

  http://www.cisco.com/warp/public/105/57.html
- IP Multicast Troubleshooting Guide

  http://www.cisco.com/en/US/tech/tk828/technologies_tech_note09186a0080094b55.shtml
- IP Multicast MRM tool

  http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fipr_c/ipcpt3/1cftools.htm
- Configuring and troubleshooting Bidir

  http://www.cisco.com/en/US/products/sw/iosswrel/ps1835/products_configuration_guide_chapter09186a00800ca796.html
- mVPN MIB
  http://www.cisco.com/en/US/partner/products/sw/iosswrel/ps1829/products_feature_guide09186a008028e3ee.html#wp1046445
- mVPN Design Guide

  http://www.cisco.com/en/US/partner/tech/tk828/tech_digest09186a00801a64a3.html

Terms and Acronyms

The following table lists and defines the key terms used in the guide.

Acronyms	Definition
BGP SAFI	BGP Sub-sequent Address Family
BIDIR	Bidirectional PIM
BSR	Bootstrap Router
C-GPA	Designates a customer multicast group address, in the context of a specific VPN that is being troubleshooted.
C-RP	Customer Rendez-vous Point
C-SRC	Customer Source
DI	Device Instrumentation
GPA	Designates the default MDT multicast group address for the mVPN that is being troubleshooted
MDT	Multicast Distribution Tree (MDT) group address
MTI	Multicast Tunnel Interface
NH	Next hop
SSM	Source Specific Multicast
MRIB	Multicast Routing Information Base
$PE_{C\text{-}RCVR}$	PE connected to the site where there is a receiver for the customer multicast traffic of a considered customer multicast source
$PE_{C\text{-}SRC}$	PE connected to the site where there is a Customer multicast source for the considered customer Multicast group
$PE_{C\text{-}RP}$	PE connected to the site where there is the Customer RP for the considered customer Multicast group
P-RP	Provider Rendez-vous Point
SP	Service provider
VRF	VPN routing/forwarding (VRF) instance
VPN	Virtual Private Network