

Table of Contents

<u>Multicasting over a GRE Tunnel</u>	1
<u>Introduction</u>	1
<u>Before You Begin</u>	1
<u>Conventions</u>	1
<u>Prerequisites</u>	1
<u>Components Used</u>	1
<u>Configure</u>	1
<u>Network Diagram</u>	2
<u>Configurations</u>	3
<u>Verify</u>	5
<u>Troubleshoot</u>	6
<u>Related Information</u>	7

Multicasting over a GRE Tunnel

Introduction

Before You Begin

- Conventions

- Prerequisites

- Components Used

Configure

- Network Diagram

- Configurations

Verify

Troubleshoot

Related Information

Introduction

This document provides a sample configuration for multicasting over a generic routing encapsulation (GRE) tunnel.

In many network scenarios you will want to configure your network to use GRE tunnels to send Protocol Independent Multicast (PIM) and multicast traffic between routers. Typically, this occurs when the multicast source and receiver are separated by an IP cloud which is not configured for IP multicast routing. In such network scenarios, configuring a tunnel across an IP cloud with PIM enabled transports multicast packets toward the receiver. This document describes the configuration, verification, and related issues pertaining to multicasting over a GRE tunnel.

Before You Begin

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.

Prerequisites

A basic understanding of multicast and PIM would be helpful. Refer to the Multicast Quick–Start Configuration Guide for more information on multicast and PIM.

Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

Configure

As shown in the network diagram below, the multicast source (10.1.1.1) is connected to R102 and is configured for multicast group 239.1.1.20. The multicast receiver (10.2.2.3) is connected to R104 and is

configured to receive multicast packets for group 239.1.1.20. Separating R102 and R104 is an IP cloud, which is not configured for multicast routing.

A tunnel is configured between R102 to R104 sourced with their loopback interfaces. The **ip pim sparse-dense mode** command is configured on tunnel interfaces and multicast-routing is enabled on R102 and R104. Sparse-dense mode configuration on the tunnel interfaces allows sparse-mode or dense-mode packets to be forwarded over the tunnel depending on rendezvous point (RP) configuration for the group.

Note: For dense mode With PIM dense mode configured over the tunnel, an **ip mroute 10.1.1.0 255.255.255.0 tunnel 0** command is configured on R104 to ensure a successful RPF for multicast source address 10.1.1.1. Incoming (10.1.1.1, 239.1.1.20) multicast packets over Tunnel0 (Tu0) are checked for Reverse Path Forwarding (RPF) using this mroute statement. After a successful check, the multicast packets are forwarded to outgoing interface list (OIL) interfaces.

Note: For sparse mode With PIM sparse mode configured over the tunnel, ensure that the following points are addressed:

- For a successful RPF verification of multicast traffic flowing over the shared tree (*,G) from RP, an **ip mroute rp-address nexthop** command needs to be configured for the RP address, pointing to the tunnel interface.

Assuming R102 to be the RP (RP address 2.2.2.2) in this case, the mroute would be the **ip mroute 2.2.2.2 255.255.255.255 tunnel 0** command, which ensures a successful RPF check for traffic flowing over the shared tree.


- For a successful RPF verification of multicast (S,G) traffic flowing over the Shortest Path Tree (SPT), an **ip mroute source-address nexthop** command needs to be configured for the multicast source, pointing to the tunnel interface.

In this case, when SPT traffic is flowing over tunnel interface an **ip mroute 10.1.1.0 255.255.255.0 tunnel 0** command is configured on R104 to ensure a successful RPF verification for incoming (10.1.1.1, 239.1.1.20) multicast packets over the Tu0 interface.

Use the Command Lookup Tool (registered customers only) to find additional information on the commands used in this document.

Network Diagram

This document uses the network setup shown in the diagram below



Configurations

This document uses the configurations shown below.

- R102
- R104

R102
<pre>r102# version 12.2 ! hostname r102 ! ! ip subnet-zero no ip domain-lookup !--- It stops IP domain lookup, which improves the show command response time. ! ip multicast-routing !--- Enables IP multicast routing. ! interface Loopback0 ip address 2.2.2.2 255.255.255.255 !--- Tunnel Source interface. ! interface Tunnel0 !--- Tunnel interface configured for PIM and carrying multicast packets to R104. ip address 192.168.24.1 255.255.255.252 ip pim sparse-dense-mode tunnel source Loopback0 tunnel destination 4.4.4.4 ! interface Ethernet0/0 !--- Interface connected to Source. ip address 10.1.1.2 255.255.255.0 ip pim sparse-dense-mode ! ! interface Serial8/0 ip address 192.168.23.1 255.255.255.252 !--- Note IP PIM sparse-dense mode is not configured on Serial interface. ! router ospf 1 log-adjacency-changes network 2.2.2.2 0.0.0.0 area 0 network 10.1.1.0 0.0.0.255 area 0 network 192.168.23.0 0.0.0.255 area 0 ! ip classless</pre>

```
ip pim bidir-enable
!
line con 0
line aux 0
line vty 0 4
  login
!
end
```

R104

```
r104#
version 12.2
!
hostname r104
!
!
ip subnet-zero
no ip domain-lookup

!--- It stops IP domain lookup, which improves the show command response time.

!
ip multicast-routing

!--- Enables IP multicast routing.

!
interface Loopback0
  ip address 4.4.4.4 255.255.255.255

!--- Tunnel Source interface.

!
interface Tunnel0
  ip address 192.168.24.2 255.255.255.252

!--- Tunnel interface configured for PIM and carrying multicast packets.

  ip pim sparse-dense-mode
  tunnel source Loopback0
  tunnel destination 2.2.2.2
!
interface Ethernet0/0
  ip address 10.2.2.2 255.255.255.0
  ip pim sparse-dense-mode
!
interface Serial9/0
  ip address 192.168.34.1 255.255.255.252

!--- Note IP PIM sparse-dense mode is not configured on Serial interface.

!
!
router ospf 1
  log-adjacency-changes
  network 4.4.4.4 0.0.0.0 area 0
  network 10.2.2.0 0.0.0.255 area 0
  network 192.168.34.0 0.0.0.255 area 0
!
ip classless
no ip http server
ip pim bidir-enable
```

```

ip mroute 10.1.1.0 255.255.255.0 Tunnel0

!--- This Mroute ensures a successful RPF check for packets flowing from the source.

!--- 10.1.1.1 over Shared tree in case of Dense mode and SPT in case of Sparse mode.

!
ip mroute 2.2.2.2 255.255.255.255 tunnel 0

!--- This Mroute is required for RPF check when Sparse mode multicast traffic is
!--- flowing from RP (assuming R102 with 2.2.2.2 as RP) towards receiver via tunnel
!--- before the SPT switchover.

line con 0
line aux 0
line vty 0 4
    login
!
end

```

Verify

Complete the following steps to verify your configuration:

1. Use the **show ip igmp group** command to verify that the receiver has sent its IGMP join membership request for group 239.1.1.20 to R104.

```

r104# show ip igmp groups
IGMP Connected Group Membership
Group Address      Interface          Uptime    Expires    Last Reporter
239.1.1.20         Ethernet0/0        00:00:04  00:02:55   10.2.2.3

```

2. Use the **show ip mroute group-address** command to display that when the source 10.1.1.1 starts multicasting packets for the group 239.1.1.20, R102 installs the (*,239.1.1.20) and (10.1.1.1, 239.1.1.20) entries in the R102 mroute table, as shown below.

```

r102# show ip mroute 239.1.1.20
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
Outgoing interface flags: H - Hardware switched
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.20), 00:00:09/00:02:59, RP 0.0.0.0, flags: D
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:09/00:00:00
    Ethernet0/0, Forward/Sparse-Dense, 00:00:09/00:00:00

(10.1.1.1, 239.1.1.20), 00:00:09/00:02:58, flags: T
  Incoming interface: Ethernet0/0, RPF nbr 0.0.0.0
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:09/00:00:00

```

Note: In the (10.1.1.1, 239.1.1.20) entry, the OIL is Tunnel0.

3. Use the **show ip mroute group-address** command to verify that R104 has the (*,239.1.1.20) and (10.1.1.1, 239.1.1.20) entries while it is forwarding multicast packets for group 239.1.1.20 sourced from 10.1.1.1, as shown below.

```
r104# show ip mroute 239.1.1.20
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
Outgoing interface flags: H - Hardware switched
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.20), 00:07:10/00:00:00, RP 0.0.0.0, flags: DCL
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  Tunnel0, Forward/Sparse-Dense, 00:07:10/00:00:00
  Ethernet0/0, Forward/Sparse-Dense, 00:07:10/00:00:00

(10.1.1.1, 239.1.1.20), 00:01:13/00:02:24, flags: CLT
Incoming interface: Tunnel0, RPF nbr 192.168.24.1, Mroute
Outgoing interface list:
  Ethernet0/0, Forward/Sparse-Dense, 00:01:13/00:00:00
```

Note: In (10.1.1.1, 239.1.1.20), the incoming interface is Tunnel0 and the RPF neighbor is 192.168.24.1 the Tunnel head end on R102. The RPF verification is done based on the Mroute configured on R104, and the multicast packets are pushed out to the OIL to the receiver connected on the Ethernet 0/0 interface.

4. Use the **show ip rpf ip-address** command to perform an RPF verification for packets sourced from 10.1.1.1. The following example confirms that RPF for 10.1.1.1 is via Tunnel 0, on which we are receiving the multicast (S,G) packets.

```
r104> show ip rpf 10.1.1.1
RPF information for ? (10.1.1.1)
RPF interface: Tunnel0
RPF neighbor: ? (192.168.24.1)
RPF route/mask: 10.1.1.1/24
RPF type: static
RPF recursion count: 0
Doing distance-preferred lookups across tables
```

Certain **show** commands are supported by the Output Interpreter (registered customers only) tool, which allows you to view an analysis of **show** command output.

Troubleshoot

If your multicast over GRE tunnel is not working, one of the following could be the cause:

- Tunnel not UP/UP The tunnel source and destination do not match on each end of the tunnel. For example, if the tunnel destination on R102 was changed to the IP address 10.2.2.2 instead of 2.2.2.2 while the configuration on R104 remained the same, the tunnel would not come up. To verify the status of the tunnel, use the **show interface tunnel 0** command.
- RPF Failure Use the **show ip mroute count** command to verify that the multicast packets are dropped because of RPF failure. A sample output of the **show ip mroute count** command and its

increasing counters for RPF failure is shown in bold below:

```
r104# show ip mroute count
IP Multicast Statistics
3 routes using 1642 bytes of memory
2 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: 224.0.1.40, Source count: 0, Packets forwarded: 0, Packets received: 0

Group: 239.1.1.20, Source count: 1, Packets forwarded: 11, Packets received: 45
Source: 10.1.1.1/32, Forwarding: 11/0/100/0, Other: 25/14/0
```

*!--- After some time, the **show ip mroute count** command is issued again.
!--- You can see the RPF failed counter increasing:*

```
r104# show ip mroute count
IP Multicast Statistics
3 routes using 1642 bytes of memory
2 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: 224.0.1.40, Source count: 0, Packets forwarded: 0, Packets received: 0

Group: 239.1.1.20, Source count: 1, Packets forwarded: 11, Packets received: 50
Source: 10.1.1.1/32, Forwarding: 11/0/100/0, Other: 30/19/0
r104#
```

Use the **show ip rpf source** command to ensure that the RPF interface is the same as that on which the source multicast packets are received Tunnel 0 in this example. Refer to the IP Multicast Troubleshooting Guide for more information about RPF failures.

- PIM Neighbors Router R102 is not forwarding over the Tunnel0 interface because it is not seeing a PM neighbor R104. You can use the **show ip pim neighbor** command on R102 to show the neighbor R104 over the tunnel. You can also use the **show ip pim int** command to show that there is a neighbor. Verify that the interface level **ip pim sparse-dense-mode** command is configured on both ends of the tunnel and that IP multicast-routing is enabled.

Related Information

- [Tunneling IP Multicast Packets Through a PIX Firewall](#)
 - [Multicast Quick-Start Configuration Guide](#)
 - [Using GRE and DVMRP Tunnels to Provide Multicast Transit](#)
 - [IP Multicast Troubleshooting Guide](#)
 - [Basic Multicast Troubleshooting Tools](#)
 - [TCP/IP Multicast Support Page](#)
 - [Technical Support – Cisco Systems](#)
-

All contents are Copyright © 1992–2003 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.