

# Configure Multicast Service Reflection on Nexus 3000

## Contents

[Introduction](#)  
[Prerequisites](#)  
[Requirements](#)  
[Components Used](#)  
[Background Information](#)  
[Supported Cisco Nexus 3k Platforms](#)  
[Supported Methods of Service Reflection](#)  
[Regular Mode Multicast NAT](#)  
[Fast-pass and Fast-pass with No-rewrite Multicast NAT](#)  
[Configure](#)  
[Topology](#)  
[Configuration](#)  
[Switch 1 Config \(Sender\)](#)  
[Switch 2 Config \(Translator\)](#)  
[Switch 3 Config \(Receiver\)](#)  
[Verify](#)  
[Verify Service Reflection Feature](#)  
[Switch 1 Verification](#)  
[Switch 2 Verification](#)  
[Switch 3 Verification](#)  
[Troubleshoot](#)  
[Summary](#)  
[Related Information](#)

## Introduction

This document describes how to configure and verify the Service Reflection feature on the Cisco Nexus 3000 (regular mode) series switches.

## Prerequisites

### Requirements

General recommendations that you have knowledge of these topics:

- Protocol Independent Multicast (PIM)
- Open Shortest Path First (OSPF)
- Network Address Translation (NAT)
- Internet Group Management Protocol (IGMP)

### Components Used

The information in this document is based on these software and hardware versions:

Sw1#	N9K-C93180YC-FX	NXOS: version 9.3(5)
Sw2#	N3K-C3548P-XL	NXOS: version 7.0(3)I7(9)
Sw3#	N3K-C3172TQ-10GT	NXOS: version 7.0(3)I7(9)

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

## Background Information

### Supported Cisco Nexus 3k Platforms

The multicast service reflection feature is supported only on Cisco Nexus 3548-X platforms from Release 7.0(3)I7(2).

### Supported Methods of Service Reflection

#### Regular Mode Multicast NAT


In regular mode, the packets incoming as the S1, G1 interfaces are translated to S2, G2 interfaces and the destination Media Access Control (MAC) address of the outgoing packet is translated as the multicast MAC address of the G2 interface (for example, the translated group).

#### Fast-pass and Fast-pass with No-rewrite Multicast NAT

In fast-pass mode, the S1, G1 interfaces are translated to S2, G2 interfaces and the destination MAC address of the outgoing packet has the multicast MAC address that corresponds to the G1 interface (for example, the MAC address of the pre-translated group).

## Configure

### Topology



Native Group: 239.194.169.1 (G1)

Translated Group: 233.193.40.196 (G2)

Original Source: 10.11.11.1 (S1)

Translated Source: 172.16.0.1. (S2)

## Configuration

### Switch 1 Config (Sender)

```
SW1# show run int eth1/47
```

```
interface Ethernet1/47
no switchport
ip address 10.11.11.1/24
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
```

```
SW1# show run ospf
```

```
feature ospf
router ospf 1
router-id 192.168.1.1
interface Ethernet1/47
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
```

```
SW1# show run pim
```

```
feature pim
ip pim rp-address 10.10.10.10 group-list 239.194.169.1/32
ip pim ssm range 232.0.0.0/8
interface Ethernet1/47
ip pim sparse-mode
```

### Switch 2 Config (Translator)

```
SW2# show run int eth 1/23,eth1/47
```

```
interface Ethernet1/23
no switchport
ip address 10.0.0.1/24
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
no shutdown
```

```
interface Ethernet1/47
```

```
no switchport
ip address 10.11.11.2/24
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
no shutdown
```

```
SW2# show run int lo0,lo411
interface loopback0
ip address 10.10.10.10/32
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode

interface loopback411
ip address 172.16.0.1/32
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
ip igmp join-group 239.194.169.1
```

```
SW2# show run ospf
feature ospf
router ospf 1
router-id 192.168.1.2
```

```
interface loopback0
ip router ospf 1 area 0.0.0.0
```

```
interface loopback411
ip router ospf 1 area 0.0.0.0
```

```
interface Ethernet1/23
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
```

```
interface Ethernet1/47
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
```

```
SW2# show run pim
feature pim
```

```
ip pim rp-address 10.10.10.10 group-list 239.194.169.1/32
ip pim rp-address 172.16.0.1 group-list 233.193.40.196/32
ip pim ssm range 232.0.0.0/8
```

```
interface loopback0
ip pim sparse-mode
```

```
interface loopback411
ip pim sparse-mode
```

```
interface Ethernet1/23
ip pim sparse-mode
```

```
interface Ethernet1/47
ip pim sparse-mode
```

```
ip service-reflect mode regular
ip service-reflect destination 239.194.169.1 to 233.193.40.196 mask-len 32 source 172.16.0.1
hardware profile multicast service-reflect port 7
```

### Switch 3 Config (Receiver)

```
SW3# show run int eth 1/24
interface Ethernet1/24
ip address 10.0.0.2/24
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
ip igmp join-group 233.193.40.196
no shutdown
```

```
SW3# show run ospf
feature ospf
router ospf 1
router-id 192.168.1.3

interface Ethernet1/24
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
```

```
SW3# show run pim

feature pim
ip pim rp-address 172.16.0.1 group-list 233.193.40.196/32
ip pim ssm range 232.0.0.0/8

interface Ethernet1/24
ip pim sparse-mode
```

## Verify

Use this section in order to confirm that your configuration works properly.

### Verify Service Reflection Feature

#### Switch 1 Verification

```
SW1# show ip mroute
IP Multicast Routing Table for VRF "default"

(*, 232.0.0.0/8), uptime: 3w6d, pim ip
  Incoming interface: Null, RPF nbr: 0.0.0.0
  Outgoing interface list: (count: 0)

(10.11.11.1/32, 239.194.169.1/32), uptime: 00:06:57, pim ip
  Incoming interface: Ethernet1/47, RPF nbr: 10.11.11.1
  Outgoing interface list: (count: 1)
  Ethernet1/47, uptime: 00:06:57, pim, (RPF)
```

#### Switch 2 Verification

<#root>

SW2# show ip mroute  
IP Multicast Routing Table for VRF "default"

(\* , 232.0.0.0/8), uptime: 00:04:39, pim ip  
  Incoming interface: Null, RPF nbr: 0.0.0.0  
  Outgoing interface list: (count: 0)

(\* , 233.193.40.196/32), uptime: 00:04:11, pim ip

**Incoming interface: loopback411**

, RPF nbr: 172.16.0.1 <--

**Translation (ingress) Loopback interface**

  Outgoing interface list: (count: 1)  
  Ethernet1/23, uptime: 00:03:59, pim <--

**Egress interface for S2,G2**

(172.16.0.1/32, 233.193.40.196/32), uptime: 00:00:15, ip mrib pim  
  Incoming interface: loopback411, RPF nbr: 172.16.0.1  
  Outgoing interface list: (count: 1)  
  Ethernet1/23, uptime: 00:00:15, pim

(\* , 239.194.169.1/32), uptime: 00:04:34, static pim ip <-- (The NAT router would pull the traffic by us)  
  Incoming interface: loopback0, RPF nbr: 10.10.10.10  
  Outgoing interface list: (count: 1)

**loopback411,**

uptime: 00:04:34, static <--

**Translation (egress) Loopback interface**

(10.11.11.1/32, 239.194.169.1/32), uptime: 00:00:17, ip mrib pim  
  Incoming interface: Ethernet1/47, RPF nbr: 10.11.11.1, internal <--

**Ingress interface for S1,G1**

  Outgoing interface list: (count: 1)  
  loopback411, uptime: 00:00:17, mrib

SW2# show ip mroute sr <--

(Only SR nat routes)

IP Multicast Routing Table for VRF "default"

(

\* , 239.194.169.1/32

```
), uptime: 00:09:29, static pim ip
  NAT Mode: Ingress
  NAT Route Type: Pre
  Incoming interface:
```

#### loopback0

```
, RPF nbr: 10.10.10.10
  Translation list: (count: 1)
  SR: (
```

```
172.16.0.1, 233.193.40.196
```

```
)
```

```
(
```

```
10.11.11.1/32, 239.194.169.1/32
```

```
), uptime: 00:05:12, ip mrib pim
  NAT Mode: Ingress
  NAT Route Type: Pre
  Incoming interface:
```

#### Ethernet1/47

```
, RPF nbr: 10.11.11.1, internal
  Translation list: (count: 1)
  SR: (
```

```
172.16.0.1, 233.193.40.196
```

```
)
```

## Switch 3 Verification

```
SW3# show ip mroute
```

```
IP Multicast Routing Table for VRF "default"
```

```
(* , 232.0.0.0/8), uptime: 02:45:09, pim ip
```

```
Incoming interface: Null, RPF nbr: 0.0.0.0
```

```
Outgoing interface list: (count: 0)
```

```
(* , 233.193.40.196/32), uptime: 01:47:02, ip pim igmp
```

```
Incoming interface: Ethernet1/24, RPF nbr: 10.0.0.1
```

```
Outgoing interface list: (count: 1)
```

```
Ethernet1/24, uptime: 01:43:27, igmp, (RPF)
```

```
(172.16.0.1/32, 233.193.40.196/32), uptime: 00:02:59, ip mrib pim
```

```
Incoming interface: Ethernet1/24, RPF nbr: 10.0.0.1
```

```
Outgoing interface list: (count: 1)
```

```
Ethernet1/24, uptime: 00:02:59, mrib, (RPF)
```

## Troubleshoot

This section provides information you can use in order to troubleshoot your configuration.

If S2 and G2 do not get created or the user faces random translation issues, you can check these points:

1. Once the traffic is received (pre-translated), post-translated entries are created based on pkt **punted in mcastfwd**.
2. If you do not see pkt punted in mcastfwd, you can check if you get the requested traffic on the **ingress interface via ACL**.
3. If you see increased counters in ACL, check same **traffic hits CPU via ethalyzer**.
4. Can also check translation in **MRIB event-history**:

<#root>

```
SW2# show system internal mfwd ip mroute -->
```

**Packets Punted in Mcast Forwarding.**

MCASFWD Multicast Routing Table for VRF "default"

(0.0.0.0/0, 232.0.0.0/8)

Software switched packets: 0, bytes: 0

RPF fail packets: 0, bytes: 0

(0.0.0.0/0, 233.193.40.196/32)

Software switched

**packets: 1**

, bytes: 84

RPF fail packets: 0, bytes: 0

(172.16.0.1/32, 233.193.40.196/32), data-alive

Software switched

**packets: 1**

, bytes: 84

RPF fail packets: 8, bytes: 672

(0.0.0.0/0, 239.194.169.1/32)

Software switched

**packets: 1**

, bytes: 84

RPF fail packets: 0, bytes: 0

(10.11.11.1/32, 239.194.169.1/32), data-alive

Software switched

**packets: 10**

, bytes: 840

RPF fail packets: 0, bytes: 0

<#root>

```
SW2# show ip access-lists test
```

```
IP access list test
```



```
statistics per-entry
10 permit ip any 239.194.169.1/32 [match=105] <--
```

Interested traffic hitting ingress interface

```
20 permit ip any any [match=11]
```

```
interface Ethernet1/47
no switchport
ip access-group test in <--
```

ACL applied on ingress interface

```
ip address 10.11.11.2/24
ip ospf network point-to-point
ip router ospf 1 area 0.0.0.0
ip pim sparse-mode
no shutdown
```

<#root>

```
SW2# ethanalyzer loca int inband display-filter "ip.addr == 239.194.169.1" limit-captured-frames 0
```

```
--> Confirm (S1,G1) seen on CPU
```

Capturing on inband

```
wireshark-cisco-mtc-dissector: ethertype=0xde09, devicetype=0x0
2022-09-18 04:21:37.840227 10.11.11.1 -> 239.194.169.1 ICMP Echo (ping) request
2022-09-18 04:21:37.841275 10.11.11.1 -> 239.194.169.1 ICMP Echo (ping) request
2022-09-18 04:21:37.860153 10.11.11.1 -> 239.194.169.1 ICMP Echo (ping) request
2022-09-18 04:21:37.861199 10.11.11.1 -> 239.194.169.1 ICMP Echo (ping) request
2022-09-18 04:21:37.880072 10.11.11.1 -> 239.194.169.1 ICMP Echo (ping) request
2022-09-18 04:21:37.881113 10.11.11.1 -> 239.194.169.1 ICMP Echo (ping) request
```

```
SW2# ethanalyzer local interface inband capture-filter "host 172.16.0.1" limit-captured-frames 0
```

```
--> Confirm (S2,G2) seen on CPU
```

Capturing on inband

```
wireshark-cisco-mtc-dissector: ethertype=0xde09, devicetype=0x0
2022-09-18 03:12:51.423484 172.16.0.1 -> 233.193.40.196 ICMP Echo (ping) request
2022-09-18 03:12:51.423978 10.0.0.2 -> 172.16.0.1 ICMP Echo (ping) reply
2022-09-18 03:12:53.425754 172.16.0.1 -> 233.193.40.196 ICMP Echo (ping) request
2022-09-18 03:12:53.425761 10.0.0.2 -> 172.16.0.1 ICMP Echo (ping) reply
2022-09-18 03:12:55.426719 172.16.0.1 -> 233.193.40.196 ICMP Echo (ping) request
2022-09-18 03:12:55.426726 10.0.0.2 -> 172.16.0.1 ICMP Echo (ping) reply
2022-09-18 03:12:57.428669 172.16.0.1 -> 233.193.40.196 ICMP Echo (ping) request
2022-09-18 03:12:57.429175 10.0.0.2 -> 172.16.0.1 ICMP Echo (ping) reply
2022-09-18 03:12:59.429890 172.16.0.1 -> 233.193.40.196 ICMP Echo (ping) request
2022-09-18 03:12:59.430386 10.0.0.2 -> 172.16.0.1 ICMP Echo (ping) reply
10 packets captured
```

<#root>

```
SW2# show ip pim event-history mrib
```

```
--> Event history to confirm that the translation is being done
```

```
2022 Sep 18 04:28:39.970688: E_DEBUG pim [19433]: Sending ack: xid: 0xeeee00d2
2022 Sep 18 04:28:39.970255: E_DEBUG pim [19433]: MRIB Join notify for (10.11.11.1/32, 239.194.169.1)
2022 Sep 18 04:28:39.968875: E_DEBUG pim [19433]: MRIB sr route type notif for (10.11.11.1/32, 239.194.169.1)
2022 Sep 18 04:28:39.968859: E_DEBUG pim [19433]: pim_process_mrib_rpf_notify: MRIB RPF notify for (
: 0.0.0.0, route-type 1
2022 Sep 18 04:28:39.968307: E_DEBUG pim [19433]: Copied the flags from MRIB for route (10.11.11.1/32, 239.194.169.1)
2022 Sep 18 04:28:39.968301: E_DEBUG pim [19433]: MRIB Join notify for (10.11.11.1/32, 239.194.169.1)
2022 Sep 18 04:28:39.968294: E_DEBUG pim [19433]: Received a notify message from MRIB xid: 0xeeee00cc
2022 Sep 18 04:28:35.904652: E_DEBUG pim [19433]: Sending ack: xid: 0xeeee00cc
2022 Sep 18 04:28:35.904625: E_DEBUG pim [19433]: pim_process_mrib_rpf_notify: MRIB RPF notify for (
e RLOC address: 0.0.0.0, route-type 0
2022 Sep 18 04:28:35.904484: E_DEBUG pim [19433]: pim_process_mrib_rpf_notify: After copying the val
ype 0
2022 Sep 18 04:28:35.904476: E_DEBUG pim [19433]: pim_process_mrib_rpf_notify: MRIB RPF notify for (
.0.0.0, route-type 0
2022 Sep 18 04:28:35.904400: E_DEBUG pim [19433]: MRIB Join notify for (172.16.0.1/32, 233.193.40.196/32)
2022 Sep 18 04:28:35.904343: E_DEBUG pim [19433]: MRIB Join notify for (0.0.0.0/32, 233.193.40.196/32)
2022 Sep 18 04:27:49.862827: E_DEBUG pim [19433]: pim_process_mrib_rpf_notify: After copying the val
2022 Sep 18 04:27:49.862812: E_DEBUG pim [19433]: pim_process_mrib_rpf_notify: MRIB RPF notify for (
type 0
2022 Sep 18 04:27:49.862798: E_DEBUG pim [19433]: MRIB Join notify for (*, 239.194.169.1/32)
2022 Sep 18 04:27:49.862795: E_DEBUG pim [19433]: MRIB Join notify for (172.16.0.1/32, 233.193.40.196/32)
2022 Sep 18 04:27:49.862789: E_DEBUG pim [19433]: MRIB Join notify for (0.0.0.0/32, 233.193.40.196/32)
2022 Sep 18 04:27:49.861870: E_DEBUG pim [19433]: Creating PIM route for (*, 239.194.169.1/32)
2022 Sep 18 04:27:49.861868: E_DEBUG pim [19433]: MRIB Join notify for (*, 239.194.169.1/32)
```

## Summary

- In Regular mode, traffic hits the original S, G entry in the first pass and recirculates due to Outgoing Interface List (OIFL) which has only the loopback port. In the second pass, it derives the destination MAC for the rewrite.
- In the third pass, the multicast route lookup happens on the translated S, G and the packet is forwarded to the corresponding translated group OIFL ports.

- Added static join on loopback to force the traffic to be received on the NAT box.
- When first packet is received for (s1, g1), switch would program (s1,g1) with new SR flag (s1, g ---> s2 ,g2) .
- The switch would use this metadata to do recircle of the packet and punt the packet for g2. Once (S2, G2) packet is punted to sup, FHR (first-hop router) functionality would be triggered on NAT box for s2,g2.
- Once the traffic is received ie pre-translated and post-translated entries would be created based on pkt punted in mcastfwd.
- If you donâ€™t see the packet punted in mcastfwd for respective group, you can use the mentioned troubleshoot process to confirm if interested traffic hits the switch

## **Related Information**

- [Cisco Technical Support & Downloads](#)