Configure Multicast Filtering on Nexus 7K/N9K

Contents

Introduction

This document describes the different ways to configure the possible ways to block or filter certain multicast traffic on Nexus 7000/9000 switches. It can also be used to conserve multicast resources. One of the common examples is Microsoft's implementation of Universal plug and play operation which uses SSDP to communicate between the servers.

Prerequisites

Requirements

Cisco recommends that you have knowledge of how Any-Source Multicast (ASM) with the use of PIM Sparse mode works on the Nexus platform.

Components Used

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

- Nexus 7K with F3/M3 LC running NXOS 7.3(4)D1(1)
- Nexus N9K-C93180YC-EX/FX with 7.0(3)I7(9) or 9.3(5)

Note: Results can vary if SW/HW is different.

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

Here is the list of the acronyms used:

RP – Rendezvous Point

FHR – First Hop Router

LHR – Last Hop Router

SRC – Multicast Source

REC – Multicast Receiver

PACL – Port Access-List

RACL – Routed Access-List

SVI – Switched Virtual Interface

ACL – Access Control List

Configure

Generic Topology

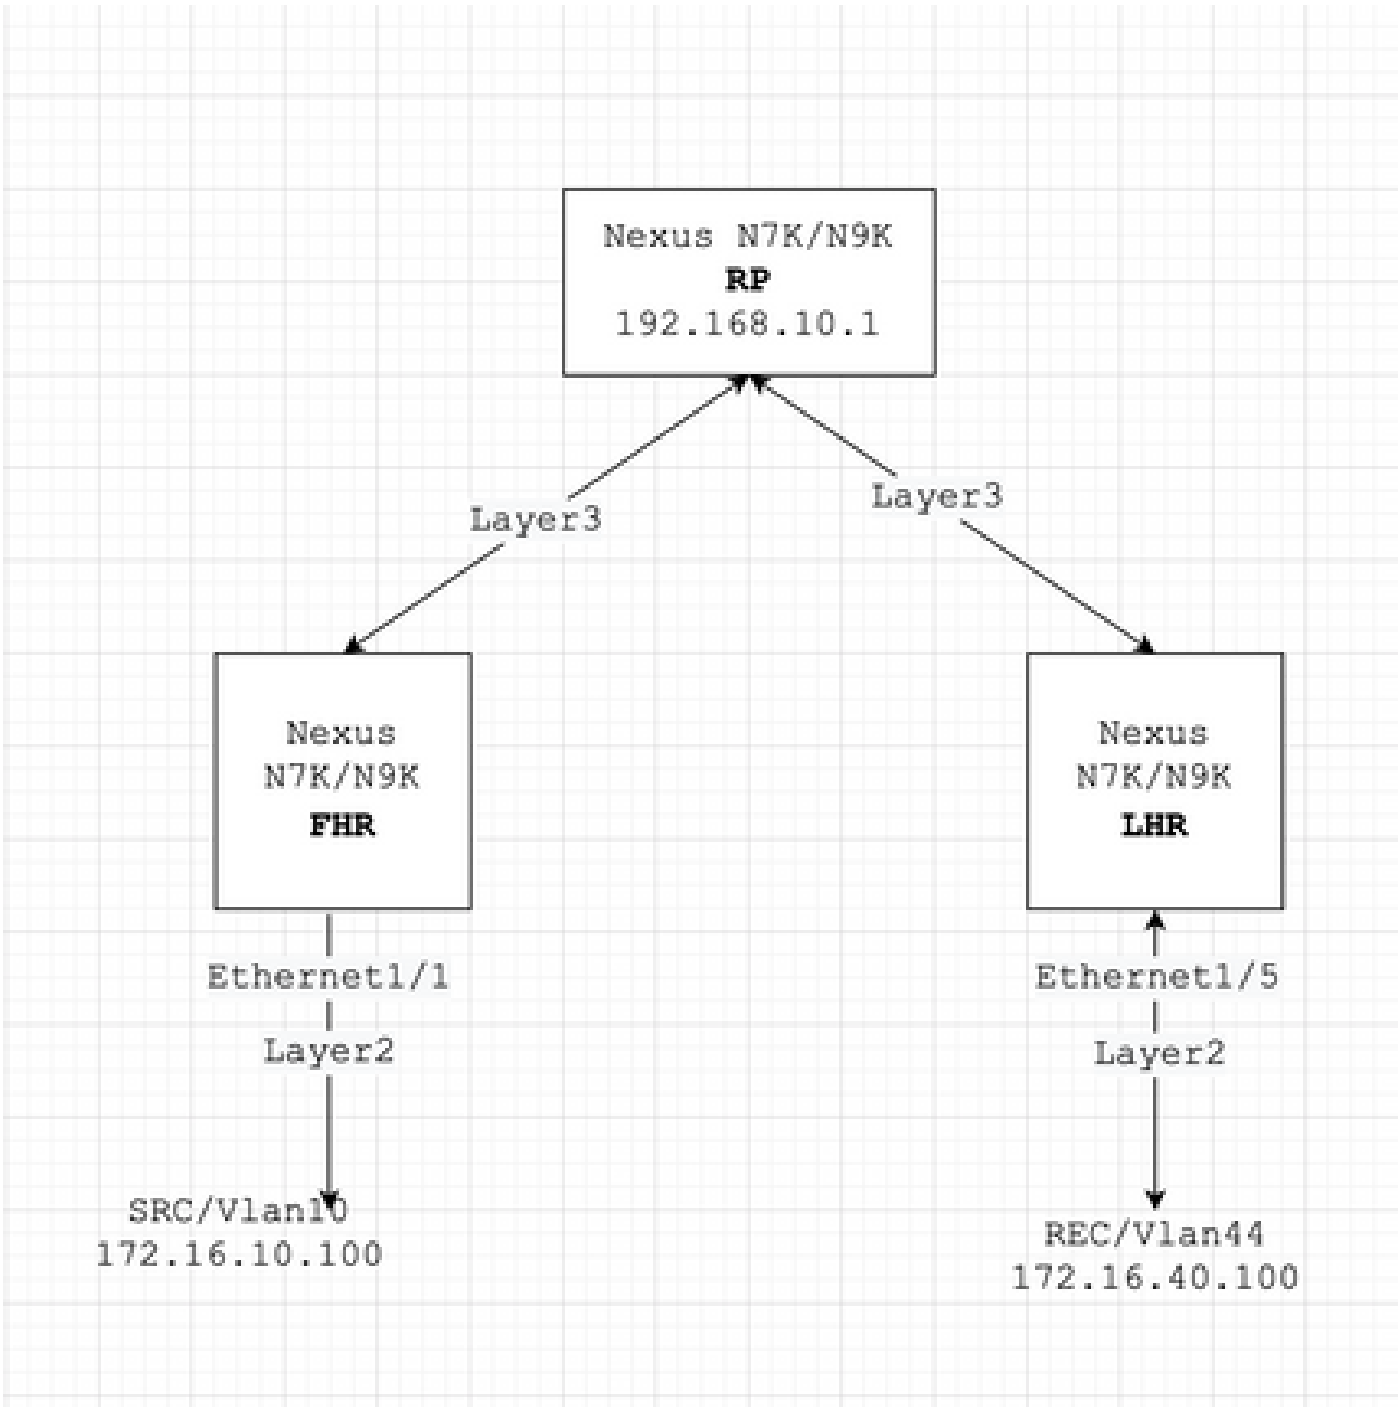

Configuration Examples

Let’s assume this:

RP’s IP address is 192.168.10.1

SRC’s IP address is 172.16.10.100/32

SSDP Group: 239.255.255.250/239.255.255.253

Now, let us discuss the configuration based on the role of the device. For example, FHR, LHR, RP, and more.

FHR – Typically Multicast SRC is Directly Connected Here

1. Filter Registration towards the existing RP.

```
ip pim rp-address 192.168.10.1 route-map filter-registration
!
Route-map filter-registration deny 5
 match ip multicast source 172.16.10.100/32 group 239.255.255.250/32
// Above line is specific to SRC/GROUP pair

Route-map filter-registration deny 7
 match ip multicast group 239.255.255.250/32
// Above line is for any SRC and specific group
!
Route-map filter-registration permit 100
 Match ip multicast group 224.0.0.0/4
```

2. Filter Registration towards the RP by defining a bogus RP (which does not exist (For example, 1.1.1.1) for SSDP groups; FHR, in this case, assumes the role of RP.

```
ip route 1.1.1.1/32 Null0
!
ip pim rp-address 1.1.1.1 route-map SSDP_groups
!
Route-map SSDP_groups permit 5
 match ip multicast group 239.255.255.250/32
Route-map SSDP_groups permit 10
 match ip multicast group 239.255.255.253/32
Route-map SSDP_groups deny 20
 match ip multicast group 224.0.0.0/4
!
ip pim rp-address 192.168.10.1 route-map all_other_groups
!
Route-map all_other_groups deny 5
 match ip multicast group 239.255.255.250/32
Route-map all_other_groups deny 10
 match ip multicast group 239.255.255.253/32
Route-map all_other_groups permit 20
 match ip multicast group 224.0.0.0/4
```

Verify:

```
Nexus9K_OR_N7K# show ip pim rp

PIM RP Status Information for VRF "default"
BSR disabled
Auto-RP disabled
BSR RP Candidate policy: None
BSR RP policy: None
Auto-RP Announce policy: None
Auto-RP Discovery policy: None
RP: 192.168.10.1, (0),
 uptime: 00:00:27   priority: 0,
```

```
 RP-source: (local), group-map: Filter-registration,
 group rangs:
 224.0.0.0/4
 239.255.255.253/32 (deny)
 239.255.255.250/32 (deny)

Nexus9K_OR_N7K# show ip mroute
IP Multicast Routing Table for VRF "default"
(172.16.10.100/32, 239.255.255.250/32), uptime: 00:04:12, ip pim
  Incoming interface: Vlan10, RPF nbr: 172.16.10.100
  Outgoing interface list: (count: 0)

Nexus9K_OR_N7K# show system internal mfwd event-history pkt
 pkt events for MCASTFWD process
2021 Jan  1 11:11:41.792316 mcastfwd [21914]: [21933]: Create state for (172.16.10.100, 239.255.255.250

Nexus9K_OR_N7K # show ip pim internal event-history null-register
2021 Jan 01 11:15:19.095711: E_DEBUG    pim [21935]:  Null Register not sent for (172.16.10.100/32, 239
```

This output confirms that FHR is not registering the stream to RP.

LHR – Typically Multicast REC is Directly Connected Here

3. Applying IGMP policy on ingress SVI (where REC resides). The idea here is to filter the IGMP membership reports for SSDP groups from REC.

```
ip pim rp-address 192.168.10.1 group-list 224.0.0.0/4
!
route-map filter-SSDP-joins deny 5
 match ip multicast group 239.255.255.250/32
route-map filter-SSDP-joins deny 6
 match ip multicast group 239.255.255.253/32
route-map filter-SSDP-joins permit 100
 match ip multicast group 224.0.0.0/4
!
Interface VlanXX
ip igmp report-policy filter-SSDP-joins
```

Verify:

```
Nexus9K_OR_N7K (config)# show ip mroute 239.255.255.250
IP Multicast Routing Table for VRF "default"
Group not found
!
Nexus9K_OR_N7K (config)# show ip igmp snooping groups vlan 44
Type: S - Static, D - Dynamic, R - Router port, F - Fabricpath core port
Vlan  Group Address      Ver  Type  Port list
44    */*                -    R     Vlan44
44    239.255.255.250    v2   D     Eth1/5
!
Nexus9K_OR_N7K (config)# show ip igmp internal event-history debugs
debugs events for IGMP process
2021 Jan  1 11:52:21.277915 igmp [1125]: : Filtered group 239.255.255.250
```

```
2021 Jan  1 11:52:21.277903 igmp [1125]: : Received v2 Report for 239.255.255.250 from 172.16.44.100 (V
```

This output confirms IGMP membership report is filtered and (*,G) join is not sent to RP.

PIM – Enabled Router Acting as FHR/LHR

You can use a combination of options 1, 2, and 3, depending on your requirements.

For example:

4. Filter Registration towards the existing RP (FHR role):

```
ip pim rp-address 192.168.10.1 route-map filter-registration
!
Route-map filter-registration deny 5
 match ip multicast source 172.16.10.100/32 group 239.255.255.250/32
Route-map filter-registration deny 7
 match ip multicast group 239.255.255.250/32
!
Route-map filter-registration permit 100
 Match ip multicast group 224.0.0.0/4
```

5. IGMP policy to filter IGMP membership reports from REC (LHR role).

```
ip pim rp-address 192.168.10.1 group-list 224.0.0.0/4
!
route-map filter-SSDP-joins deny 5
 match ip multicast group 239.255.255.250/32
route-map filter-SSDP-joins deny 6
 match ip multicast group 239.255.255.253/32
route-map filter-SSDP-joins permit 100
 match ip multicast group 224.0.0.0/4
!
Interface VlanXX
ip igmp report-policy filter-igmp-joins
```

Verify:

Pretty much the same as verification done in points C and D mentioned previously.

```
Show ip mroute
Show ip pim rp
Show ip pim internal event-history join-prune
Show ip igmp internal event-history debugs
```

RP – This is Rendezvous Point

6. Registration policy to block the registration of SSDP group from FHR.

```
ip pim rp-address 192.168.10.1 group-list 224.0.0.0/4
ip pim register-policy all_groups
!
Route-map all_groups deny 5
 match ip multicast group 239.255.255.250/32
Route-map all_groups deny 10
 match ip multicast group 239.255.255.253/32
Route-map all_groups permit 20
 match ip multicast group 224.0.0.0/4
```

Verify:

```
Nexus9K_OR_N7K (config)# show ip mroute 239.255.255.250
IP Multicast Routing Table for VRF "default"
Group not found
!
Nexus9K_OR_N7K (config)# show ip pim internal event-history data-register-receive
2021 Jan 08 03:33:06.353951: E_DEBUG    pim [1359]:  Register disallowed by policy
2021 Jan 08 03:33:06.353935: E_DEBUG    pim [1359]:  Received DATA Register from 172.16.10.1 for (172.1
2021 Jan 08 03:29:42.602744: E_DEBUG    pim [1359]:  Add new route (172.16.10.100/32, 239.1.1.1/32) to

F241.01.13-C9318OYC-EX-1(config)# show ip pim internal event-history null-register
2021 Jan 08 03:35:40.966617: E_DEBUG    pim [1359]:  Send Register-Stop to 172.16.10.1 for (172.16.10.1
2021 Jan 08 03:35:40.966613: E_DEBUG    pim [1359]:  Register disallowed by policy
2021 Jan 08 03:35:40.966597: E_DEBUG    pim [1359]:  Received NULL Register from 172.16.10.1 for (172.1
```

This output confirms RP is blocking registration for group 239.255.255.250.

7. Applying the join-prune policy on the RP - both pim (*,G) join and (S,G) join for the SSDP group only.

```
ip pim rp-address 192.168.10.1 group-list 224.0.0.0/4
ip pim register-policy all_groups
!
Route-map all_groups deny 5
 match ip multicast group 239.255.255.250/32
Route-map all_groups deny 10
 match ip multicast group 239.255.255.253/32
Route-map all_groups permit 20
 match ip multicast group 224.0.0.0/4
!
Interface Ethernet/Y
 ip pim sparse-mode
 ip pim jp-policy all_groups
```

Verify:

```
Nexus9K_OR_N7K # show ip mroute 239.255.255.253
IP Multicast Routing Table for VRF "default"
Group not found
!
F241.01.13-C9318OYC-EX-1# show ip pim internal event-history join-prune
2021 Jan 08 03:53:41.643419: E_DEBUG    pim [1359]:  Join disallowed by inbound JP policy
```

This output confirms (*,G) PIM join is blocked by RP.

Configure Conserve HW Entries for Multicast

Although all the options discussed in sections A, B or C; prevent either FHR, LHR or FHR/LHR from registering the stream at RP or prevent sending PIM Join (*,G) towards the RP respectively; a mroute or snooping entry can still be created and it consumes multicast HW entries.

> **Note**: You can use RACL or PACL on ingress SVI or Layer2 interfaces/port-channels/VPC port-channels in case VPC is configured. If SRC/REC are sprayed out in different VLAN or L2 interfaces, then it also means RACL or PACL needs to be applied on all of those. But, depending HW/SW (mostly due to HW limitation) results can vary.

PACL

Configure PACL on ingress Layer2 port or port-channel or VPC port-channel to block SSDP traffic or creation of (S, G) entry on FHR.

> **Note**: Depending on the HW used (for example, Nexus N9000), TCAM can needs to be carved prior (which requires reloading) to apply the PACL.

For example:

```
ip access-list BlockAllSSDP
 Statistics per-entry
 10 deny ip any 239.255.255.250/32
 20 deny ip any 239.255.255.253/32
 30 permit ip any any
!
Interface Ethernet X/Y
Or
Interface port-channel XX
ip port-access group BlockAllSSDP in
```

Verify:

```
```

```
F241.01.13-C9318OYC-EX-1#  sh ip mroute  239.255.255.250
IP Multicast Routing Table for VRF "default"
Group not found
!
show ip access-lists BlockAllSSDP
IP access list BlockAllSSDP
        statistics per-entry
        10 deny ip any 239.255.255.250/32 [match=3] -> Drop counters
        20 deny ip any 239.255.255.253/32 [match=0]
        30 permit ip any any [match=0]
```

Since both multicast traffic/IGMP membership ports are blocked via PACL, you do not see any snooping, mroute entry. Essentially PACL is dropping them both.

RACL

You can configure RACL on ingress SVI where SRC exist but depending on SW/HW used; (S, G) entry can still be created or traffic can be forwarded to other local VLANs.

```
ip access-list BlockAllSSDP
 Statistics per-entry
 10 deny ip any 239.255.255.250/32
 20 deny ip any 239.255.255.253/32
 30 permit ip any any
!
Interface VlanXX
ip port-access group BlockAllSSDP in
```

Verify:

It is pretty much the same as PACL but the RACL option can not provide the same results as PACL; mostly its HW limitation is mentioned earlier as well.

COPP

You can also block SSDP at COPP. This is an example configuration:

```
class-map type control-plane match-any nossdp
  match access-group name nossdp
policy-map type control-plane nossdp
  class nossdp
    police cir 0 bps bc 0 bytes conform transmit violate drop control-plane dynamic
  service-policy-dyn input nossdp
!
ip access-list nossdp
  statistics per-entry
  10 permit ip any 239.255.255.250/32
  20 permit igmp any 239.255.255.250/32
  30 permit pim any 239.255.255.250/32
```

Global Multicast Boundary

Beginning with Cisco NX-OS Release 10.2(1), Global Boundary Multicast configuration is supported.

You need to configure the {ip | ipv6} multicast group-range prefix-list <prefix-list-name> command in VRF configuration mode to define a global range of IP multicast groups and channels to be permitted or denied for the global multicast boundary. This command is used to disable multicast protocol actions and traffic forwarding for unauthorized groups or channels for all interfaces on a router. The prefix list configures the boundary. A sample configuration is provided below:

```
vrf context enterprise
ip multicast group-range prefix-list test
```

https://www.cisco.com/c/en/us/td/docs/dcn/nx-os/nexus9000/102x/configuration/multicast-routing/cisco-nexus-9000-series-nx-os-multicast-routing-configuration-guide-release-102x/m-overview.html#concept_29A33F30E7F84F7AA20C8D7D1A22ED98