ASA Multicast Troubleshooting and Common Problems

Contents

Introduction

This document explains multicast capabilities of the Adaptive Security Appliance (ASA), as well as potential problems that can be encountered when using the feature.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- ASA multicast

Components Used

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

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Feature Information

The ASA Command Line Configuration Guide outlines the multicast-routing feature and how to configure it:

http://www.cisco.com/en/US/docs/security/asa/asa90/configuration/guide/route_multicast.html

Multicast on the ASA can be configured in one of two modes:

- PIM sparse-mode (preferred)
- IGMP Stub-mode (Internet Group Management Protocol, RFC 2236 IGMPv2)

PIM sparse-mode is the preferred choice because the ASA communicates with neighbors using a true multicast routing protocol (PIM). IGMP Stub-mode was the only multicast configuration option before ASA version 7.0 was released, and operated by simply forwarding IGMP reports received from clients towards upstream routers.

PIM Sparse-mode Operation

- The ASA supports PIM sparse-mode and PIM bi-directional mode.
- PIM sparse-mode and IGMP stub-mode commands must not be configured concurrently.
- With PIM sparse-mode all multicast traffic initially flows to the Rendezvous Point (RP), then is forwarded towards the receivers. After some time the multicast flow will go directly from the source to the receivers (bypassing the RP).

The picture below illustrates a common deployment where the ASA has multicast clients on one interface, and PIM neighbors on another:

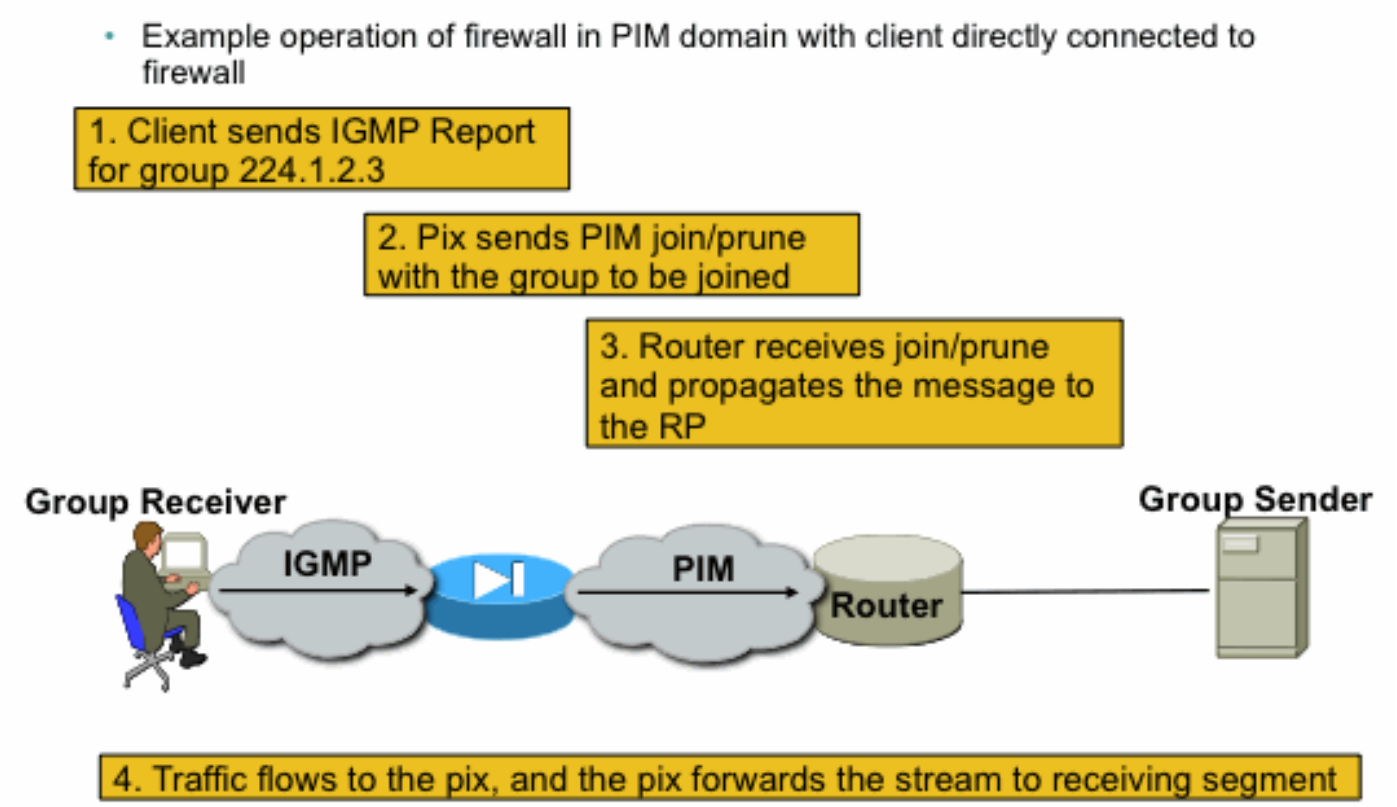

PIM Sparse-mode Sample Configuration

Complete these steps:

1. Enable multicast routing (global configuration mode).`ASA(config)# multicast-routing`

2. Define the PIM Rendezvous-point address.`ASA(config)# pim rp-address 172.18.123.3`
3. Allow the multicast packets in on the appropriate interface (necessary only if the security policy of the ASA is blocking the inbound multicast packets).
```
access-list 105 extended permit ip any host 224.1.2.3
access-group 105 in interface outside
```

IGMP Stub-mode Operation

- In IGMP Stub-mode the ASA acts as a multicast client by generating or forwarding IGMP reports (also known as IGMP "joins") towards adjacent routers, to trigger the reception of multicast traffic
- Routers will periodically send queries to the hosts to see if any node on the network wants to continue to receive the multicast traffic.
- IGMP Stub-mode is not recommended because PIM sparse-mode offers many benefits over Stub-mode (including more efficient multicast traffic flows, ability to participate in PIM, etc).

The picture below illustrates the basic operation of an ASA configured for IGMP Stub-mode.

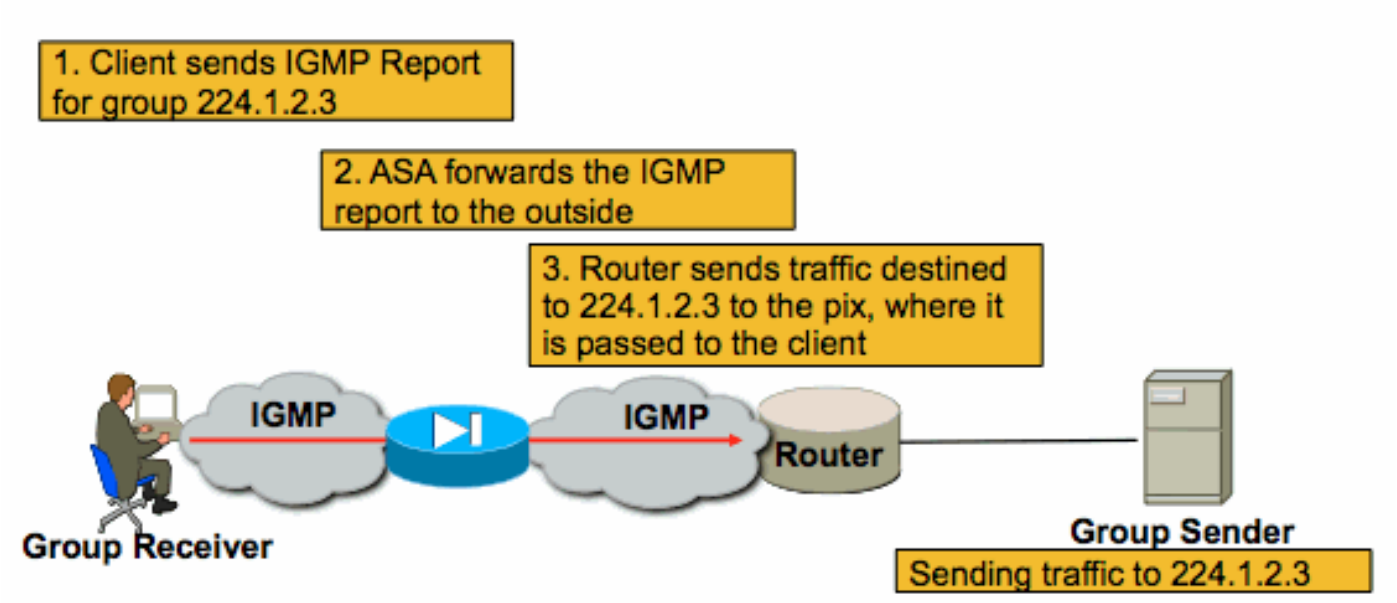

IGMP Stub-mode Configuration

Complete these steps:

1. Enable multicast routing (global config mode).`ASA(config)# multicast-routing`
2. On the interface on which you will receive the igmp reports, configure the igmp forward-interface command. Forward the packets out the interface towards the source of the stream. In the example below, the multicast receivers are directly connected to the inside interface, and the multicast source is beyond the outside interface.!
```
interface Ethernet0
 nameif outside
 security-level 0
 ip address 172.16.1.1 255.255.255.0
 no pim
!
interface Ethernet1
 nameif inside
 security-level 100
 ip address 10.0.0.1 255.255.255.0
```

```
no pim
igmp forward interface outside !
```

3. Allow the multicast packets in on the appropriate interface (only necessary if the security policy of the ASA denies the inbound multicast traffic).`access-list 105 extended permit ip any host 224.1.2.3`
   `access-group 105 in interface outside`Often there is confusion around the different **igmp interface sub-mode** commands, and the diagram below attempts to describe when to use each:

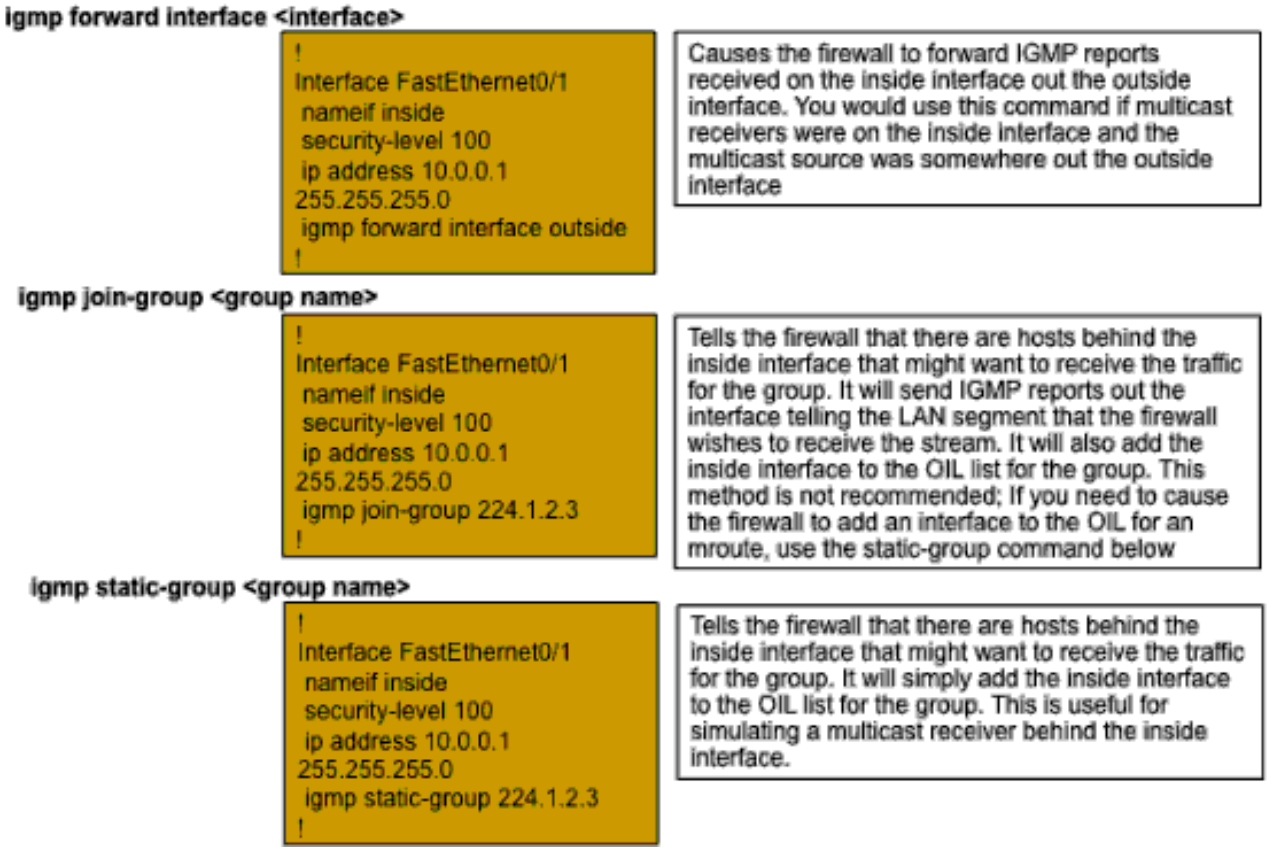

Troubleshooting Methodology

Information To Collect When Troubleshooting Multicast Problems

In order to completely understand and diagnose a multicast forwarding problem on the ASA, some or all of this information might be needed:

- A description of the network topology, including location fo the multicast senders, receivers, and rendezvous-point.
- The specific group IP address the traffic is using, as well as the ports and protocols employed.
- Syslogs generated by the ASA at the time the multicast stream has trouble.
- Specific show command output from the ASA command line interface, including:`show mroute`
  `show mfib`
  `show pim neighbor`
  `show route`
  `show tech-support`
- Packet captures to show if the multicast data arrives at the ASA, and if the packets are forwarded through the ASA.
- Packet captures showing IGMP and/or PIM packets.
- Information from adjacent multicast devices (routers) such as ‘show mroute’ and ‘show mfib’.

- Packet captures and/or show commands to determine if the ASA is dropping the multicast packets. The 'show asp drop' command can be used to determine if the ASA is dropping the packets. Additionally, packet captures of type 'asp-drop' can be used to capture all packets the ASA drops, then examined to see if the multicast packets are present in the drop capture.

Useful show Command Output

The **show mroute** command output shows the various groups and forwarding information, and is very similar to the IOS **show mroute** command. The **show mfib** command displays the forwarding status of the various multicast groups. It is especially important to observe the *Forwarding* packet counter, as well as *Other* (which indicates drops):

```
ciscoasa# show mfib
Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
             AR - Activity Required, K - Keepalive
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts: Total/RPF failed/Other drops
Interface Flags: A - Accept, F - Forward, NS - Negate Signalling
             IC - Internal Copy, NP - Not platform switched
             SP - Signal Present
Interface Counts: FS Pkt Count/PS Pkt Count
(*,224.1.2.3) Flags: S K
   Forwarding: 0/0/0/0, Other: 0/0/0
   inside Flags: F
     Pkts: 0/0
(192.168.1.100,224.1.2.3) Flags: K
   Forwarding: 6749/18/1300/182, Other: 690/0/690
   outside Flags: A
   inside Flags: F
     Pkts: 6619/8
(*,232.0.0.0/8) Flags: K
   Forwarding: 0/0/0/0, Other: 0/0/0
ciscoasa#
```

The **clear mfib counters** command can be used to clear the counters, which is very useful during testing:

```
ciscoasa# clear mfib counters
ciscoasa#
```

Using Packet Captures to Capture Multicast Traffic

The ASA's onboard packet capture utility is very useful for troubleshooting multicast problems. In the example below, all packets arriving at the ASA's DMZ interface, destined to 239.17.17.17 will be captured:

```
ciscoasa# capture dmzcap interface dmz
ciscoasa# capture dmzcap match ip any host 239.17.17.17
ciscoasa# show cap dmzcap

324 packets captured

   1: 17:13:30.976618       802.1Q vlan#301 P0 10.1.123.129.2000 >
239.17.17.17.16384:
  udp 172
   2: 17:13:30.976679       802.1Q vlan#301 P0 10.1.123.129.2000 >
239.17.17.17.16384:
```

```
  udp 172
   3: 17:13:30.996606       802.1Q vlan#301 P0 10.1.123.129.2000 >
239.17.17.17.16384:
  udp 172
   4: 17:13:30.996652       802.1Q vlan#301 P0 10.1.123.129.2000 >
239.17.17.17.16384:
  udp 172
   5: 17:13:31.016676       802.1Q vlan#301 P0 10.1.123.129.2000 >
239.17.17.17.16384:
  udp 172
   6: 17:13:31.016722       802.1Q vlan#301 P0 10.1.123.129.2000 >
239.17.17.17.16384:
  udp 172
....
```

Packet captures are also useful for capturing PIM and IGMP traffic. The capture below shows the inside interface has received an IGMP packet (IP protocol 2) sourced from 10.0.0.2:

```
ciscoasa# capture capin interface inside
ciscoasa# capture capin match igmp any any
ciscoasa# show cap capin
1 packets captured
1: 10:47:53.540346 802.1Q vlan#15 P0 10.0.0.2 > 224.1.2.3:
  ip-proto-2, length 8
ciscoasa#
```

Example ASA PIM Sparse-Mode Multicast Deployment

The diagrams below illustrate how the ASA interacts with neighbor devices to get multicast traffic flowing with PIM sparse-mode. In this specific example, the ASA receives.

Understanding the network topology

Determine exactly where the sender and receiver of the specific multicast stream you are testing reside. Also, determine the multicast group IP address being used, as well as the location of the rendezvous point.

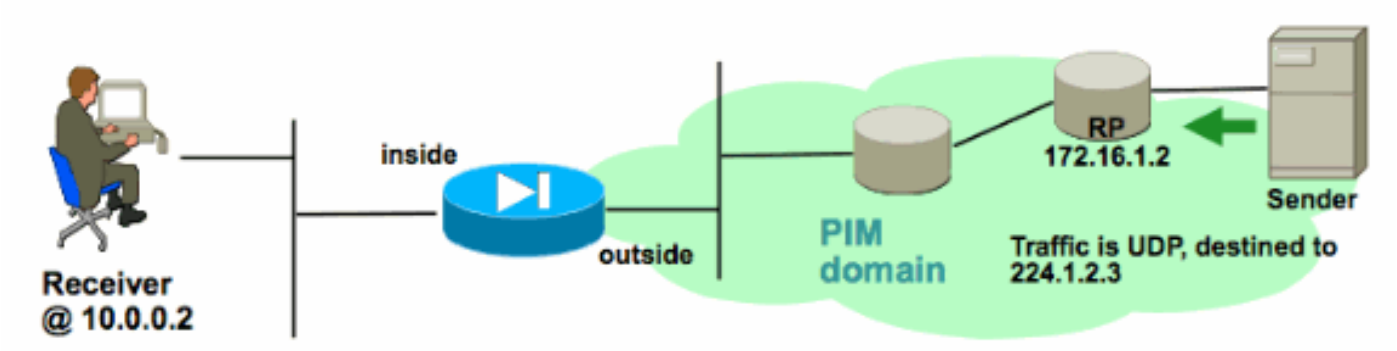

In this case, the data should be received at the outside interface of the ASA, and forwarded to the multicast receiver on the inside interface. Because the receiver is in the same IP subnet as the inside interface of the ASA, expect to see an IGMP Report received at the ASA's inside interface when the client requests to receive the stream. The IP address of the sender is 192.168.1.50.

Verifying the ASA receives the IGMP report from the receiver

In this example, the IGMP report is generated by the receiver and processed by the ASA.

Packet captures and the output of debug igmp can be used to verify that the ASA received, and successfully processed the IGMP message.

Verifying the ASA sends a PIM join message towards the rendezvous point

The ASA interprets the IGMP report and generates a PIM join message, then sends it out the interface towards the RP.

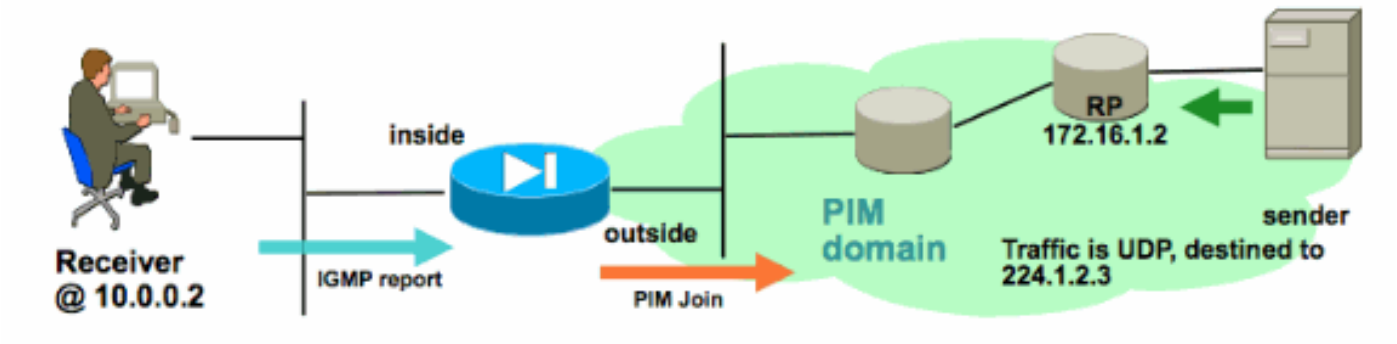

The output below is from debug pim group 224.1.2.3 and shows the ASA successfully sending the PIM join message. The sender of the multicast stream is 192.168.1.50

```
IPv4 PIM: (*,224.1.2.3) J/P processing
IPv4 PIM: (*,224.1.2.3) Periodic J/P scheduled in 50 secs
IPv4 PIM: (*,224.1.2.3) J/P adding Join on outside
IPv4 PIM: (*,224.1.2.3) inside Processing timers
IPv4 PIM: Sending J/P message for neighbor 10.2.3.2 on outside for 1 groups
IPv4 PIM: [0] (192.168.1.50,224.1.2.3/32) MRIB update (a=0,f=0,t=1)
IPv4 PIM: [0] (192.168.1.50,224.1.2.3/32) outside MRIB update (f=20,c=20)
IPv4 PIM: [0] (192.168.1.50,224.1.2.3) Signal present on outside
IPv4 PIM: (192.168.1.50,224.1.2.3) Create entry
IPv4 PIM: [0] (192.168.1.50,224.1.2.3/32) outside MRIB modify NS
IPv4 PIM: Adding monitor for 192.168.1.5
```

Verifying the ASA receives and forwards the multicast stream

The ASA begins receiving multicast traffic on the outside interface (illustrated by the green arrows), and forwarding it to the receivers on the inside.

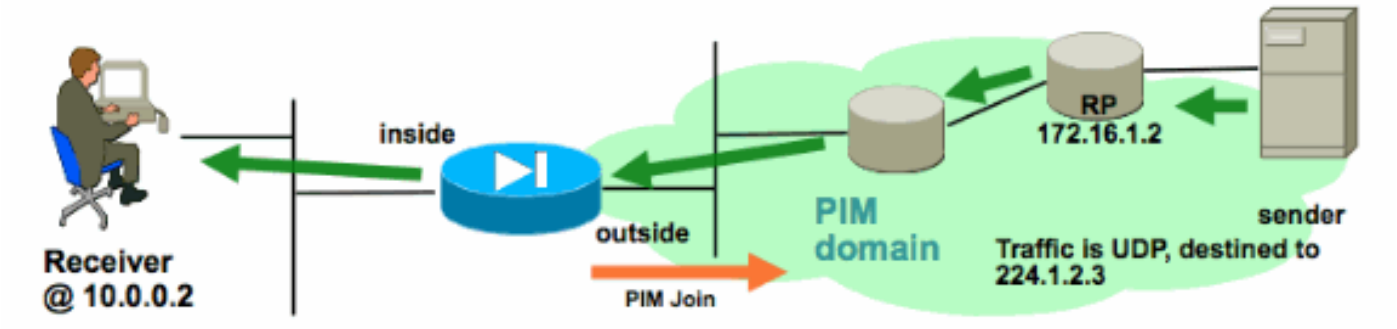

The **show mroute** and **show mfib** commands, as well as packet captures, can be used to verify the ASA receives and forwards the multicast packets.

A connection will be built in the ASA's connection table to represent the multicast stream:

```
ciscoasa# show conn
59 in use, 29089 most used
...
```

```
UDP outside:192.168.1.50/52075 inside:224.1.2.3/1234 flags -
...
```

Data Analysis

Common Problems

This section provides a series of real-world ASA multicast related problems that network administrators have encountered in the past.

The ASA Fails To Send PIM Messages Toward Upstream Routers Due To HSRP

When this problem is encountered, the ASA fails to send any PIM messages out an interface. The diagram below shows that the ASA cannot send PIM messages towards the sender, but the same problem can be seen when the ASA needs to send a PIM message towards the RP.

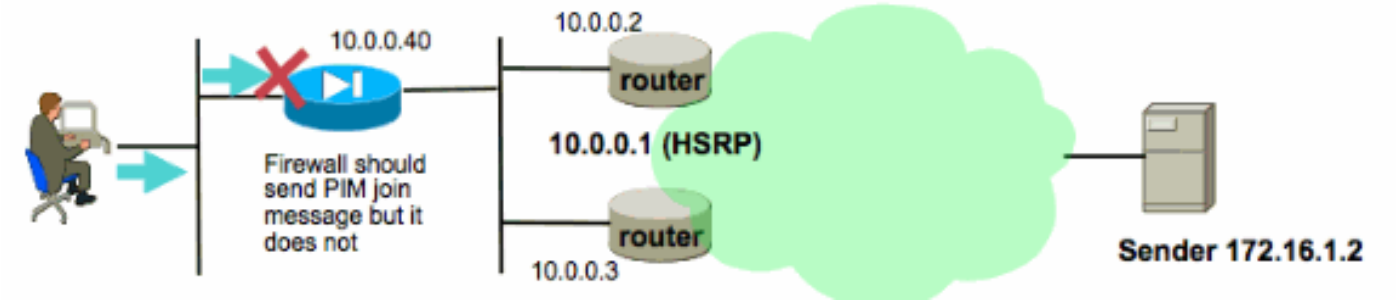

The output of **debug pim** shows that the ASA cannot send the PIM message to the upstream next-hop router:

```
IPv4 PIM: Sending J/P to an invalid neighbor: outside 10.0.0.1
```

This issue is not specific to the ASA, and also affects routers. The problem is triggered by the combination of the ASA's routing table configuration and the HSRP configuration used by the PIM neighbors.

The ASA's routing table points to the HSRP IP 10.0.0.1 as the next-hop device:

```
ciscoasa# sh run route
route outside 0.0.0.0 0.0.0.0 10.0.0.1 1
```

However, the PIM neighbor relationship is formed between the physical interface IP addresses of the routers, and not the HSRP IP:

```
ciscoasa# sh pim neighbor
Neighbor Address  Interface          Uptime    Expires DR pri Bidir
10.0.0.2          outside            01:18:27  00:01:25 1
10.0.0.3          outside            01:18:03  00:01:29 1 (DR)
```

Refer to Why Doesn't PIM Sparse Mode Work with a Static Route to an HSRP Address? for more information.

An excerpt from the document:

"Why is the router not sending the Join/Prune message? RFC 2362 states that "a router sends a

periodic Join/Prune message to each distinct RPF neighbor associated with each (S,G), (,G) and (*,*,RP) entry. Join/Prune messages are sent only if the RPF neighbor is a PIM neighbor."*

In order to mitigate the problem, add a static mroute entry on the ASA for the traffic in question. Make sure that it points to one of the two router's interface IP addresses (10.0.0.2 or 10.0.0.3 in the example above). In this case, the following command allows the ASA to send PIM messages directed towards the multicast sender at 172.16.1.2:

```
ciscoasa(config)# mroute 172.16.1.2 255.255.255.255 10.0.0.3
```

Once this is done the multicast routing table will override the unicast routing table of the ASA, and the ASA will send the PIM messages directly to the 10.0.0.3 neighbor.

The ASA Ignores IGMP Reports Because It Is Not The Designated Router On The LAN Segment

For this problem, the ASA receives an IGMP report from a directly connected multicast receiver, yet it ignores it. No debug output will be generated and the packet is simply dropped, and stream reception fails.

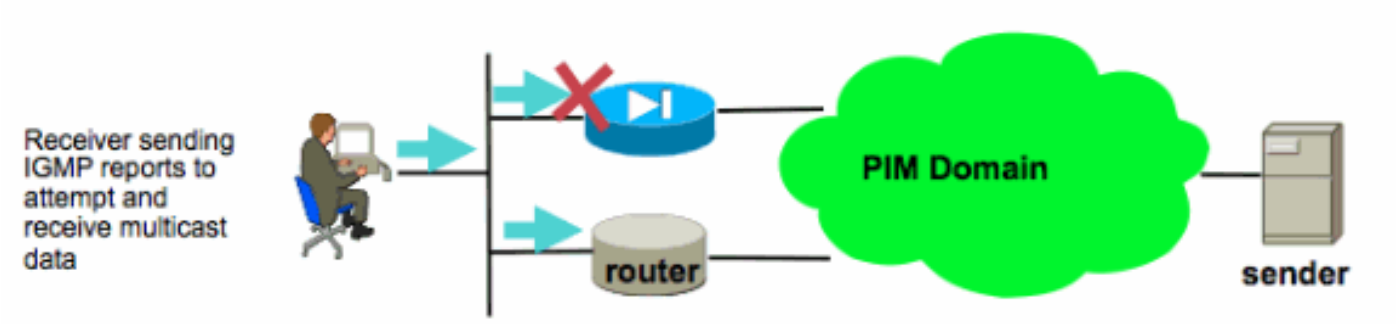

For this problem, the ASA is ignoring the packet because it is not the PIM elected designated router on the LAN segment where the clients reside.

The ASA CLI output below shows that a different device is the Designated Router (denoted by "DR") on the inside interface network:

```
ciscoasa#show pim neighbor

 Neighbor Address  Interface          Uptime    Expires DR pri Bidir
 192.168.1.2       outside         01:18:27  00:01:25 N/A>
 10.0.0.2          inside             01:18:03  00:01:29 1 (DR)
```

By default, PIM is enabled on all ASA interfaces when the **multicast-routing** command is added to the ASA's configuration. If there are other PIM neighbors (other routers or ASAs) on the inside interface of the ASA (where the clients reside) and one of those neighbors were elected because the DR for that segment, then other, non-DR routers will drop IGMP reports. The solution is to disable PIM on the ASA's interface (with the **no pim** command on the interface involved), or to make the ASA the DR for the segment using the **pim dr-priority interface** command.

The ASA Fails To Forward Multicast Traffic In The 232.x.x.x/8 range

This address range is for use with Source Specific Multicast (SSM) which the ASA does not currently support.

The output of **debug igmp** will show this error:

```
IGMP: Exclude report on inside ignored for SSM group 232.179.89.253
```

The ASA Drops Multicast Packets Due To Reverse Path Forwarding Check

In this case, the ASA receives multicast traffic on an interface, but it is not forwarded on to the receiver. Packets are dropped by the ASA because they fail the Reverse Path Forwarding (RPF) security check. The RPF is enabled on all interfaces for multicast traffic and cannot be disabled (for unicast packets the check is not on by default, and is enabled with the **ip verify reverse-path interface** command).

Due to the RPF check, when multicast traffic is received at an interface, the ASA checks to see that it has a route back to the source of the multicast traffic traffic (it checks the unicast and multicast routing table) on that interface. If it does not have a route to the sender, it drops the packet. These drops can be seen as a counter in the output of **show asp drop**:

```
ciscoasa(config)# show asp drop

Frame drop:
  Invalid UDP Length                                    2
  No valid adjacency                                   36
  No route to host                                   4469
  Reverse-path verify failed                       121012
```

This problem can be mitigated by adding a specific multicast routing table entry to the ASA for the sender of the traffic. In the example below, the mroute command is used to satisfy the RPF check for multicast traffic sourced from 172.16.1.2 received on the outside interface:

```
ciscoasa(config)# mroute 172.16.1.2 255.255.255.255 outside
```

The ASA Does Not Generate PIM Join Upon PIM Switchover to Source-tree

Initially, PIM sparse-mode multicast packets will flow from the multicast sender to the RP, then from the RP to the receiver via a shared multicast tree. However, once the aggregate bit rate reaches a certain threshold, the router closest to the multicast receiver will attempt to receive traffic along the source-specific tree. This router will generate a new PIM join for the group and send it towards the sender of the multicast stream (and not towards the RP, as before).

Depending on the network topology, the sender of the multicast traffic might reside on a different ASA interface than the RP. When the ASA receives the PIM join to switch to the source specific tree, the ASA must have a route to the IP address of the sender. If this route is not found, the PIM join packet will be dropped and the following message will be seen in the output of **debug pim**:

```
NO RPF Neighbor to send J/P
```

The solution for this problem is to add a static mroute entry for the sender of the stream, pointing out the ASA interface off of which the sender resides.

The ASA Drops Multicast Packets Due To Time To Live (TTL) Exceeded

In this case, multicast traffic is failing because the TTL of the packets is too low. This causes the ASA, or some other device in the network, to drop them.

Often multicast packets have the IP TTL value set very low by the application that sent them. Sometimes this is done by default to help ensure that the multicast traffic does not travel too far though the network. For example, by default the Video LAN Client application (a popular multicast transmitter and testing tool) sets the TTL in the IP packet to 1 by default.

The ASA Experiences High CPU Usage And Dropped Packets Due To Specific Multicast Topology

The ASA might experience high CPU and the multicast stream might experience packet drops if all of the following are true about the multicast topology:

1. The ASA is acting as the RP.
2. The ASA is the first hop receiver of the multicast stream. This means that the multicast sender is in the same IP subnet an ASA interface.
3. The ASA is the last hop router of the multicast stream. This means that a multicast receiver is in the same IP subnet as an ASA interface.

If all of the above are true, then due do a design limitation the ASA will be forced to process switch the multicast traffic. This results in high data rate multicast streams to experience packet drops. The show asp drop counter that increments when these packets are dropped is punt-rate-limit.

In order to determine if an ASA is experiencing this problem, complete these steps:

Step 1: Check if the ASA is the RP by using the two commands:

```
show run pim
show pim tunnel
```

Step 2: Check if the ASA is the last hop router by using this command:

```
show igmp group <mcast_group_IP>
```

Step 3: Check if the ASA is the first hop router by using this command:

```
show mroute <mcast_group_IP>
```

A Disconnecting Multicast Receiver Interrupts Multicast Group Reception On Other Interfaces

Only ASAs operating in IGMP Stub-mode experience this problem. ASAs that participate in PIM multicast routing are not affected.

The issue is identified by the bug CSCeg48235 - IGMP: Stopping group rcvr interrupts group reception on other interfaces

This is the release note from the bug, which explains the problem:

```
Symptom:
When a PIX or ASA firewall is configured for IGMP stub mode multicast reception and
traffic from a multicast group is forwarded to more than one interface, if a host
behind a receiving interface sends an IGMP Leave message for the group, it could
temporarily interrupt the reception for that group on other interfaces of the
firewall.
```

```
The problem is triggered when the firewall forwards the IGMP leave for the group
towards the upstream device; that device then sends a IGMP query to determine if any
other receivers exist out that interface towards the firewall, but the firewall does
not report that it still has valid receivers.

Conditions:
The PIX or ASA must be configured for IGMP stub mode multicast. IGMP stub mode is a
legacy multicast forwarding technique, whereby IGMP packets from receivers are
forwarded
through the firewall towards the source of the stream. It is recommended to use PIM
multicast routing instead of stub igmp forwarding.

Workarounds:
1) Use PIM multicast routing instead of IGMP stub mode.
2) Decrease multicast IGMP query timers so that the receivers are queried more
frequently,
causing their IGMP reports to be forwarded towards the sender more frequently, thus
restarting the stream quicker.
```

The ASA Drops Multicast Packets Due to Security Policy Of Outbound Access-list

With this specific issue the ASA is correctly dropping multicast packets (per the configured security policy). However, it is difficult for the network administrator to identify the reason for the packet drops. In this case, the ASA is dropping packets due to the outbound access-list configured for an interface. The workaround is to permit the multicast stream in the outbound access-list.

When this occurs, multicast packets will be dropped and the ASP drop counter will be "FP no mcast output intrf (no-mcast-intrf)".

The ASA Drops The First Few Packets When A Multicast Stream Is First Started

When the first packets of a multicast stream arrive at the ASA, the ASA must build that particular multicast connection and the associated mroute entry to forward the packets. While the entry is being created some multicast packets might be dropped until the mroute and connections have been established (usually this takes less than a second). Once the multicast stream setup is complete, the packets will no longer be rate limited.

Packets dropped for this reason will have the ASP drop reason of "(punt-rate-limit) Punt rate limit exceeded". Below is the output of **show capture asp** (where asp is an ASP drop capture configured on the ASA to capture dropped packets) and you can see the multicast packets that were dropped for this reason:

```
ASA # sh capture asp
2 packets captured
   1: 16:14:49.419091 10.23.2.2.810 > 239.255.123.123.890:  udp 32 Drop-reason:
 (punt-rate-limit) Punt rate limit exceeded
   2: 16:14:49.919172 10.23.2.2.810 > 239.255.123.123.890:  udp 32 Drop-reason:
 (punt-rate-limit) Punt rate limit exceeded
2 packets shown
```

Related Information

- Technical Support & Documentation - Cisco Systems