Troubleshoot ASA Multicast Common Problems

Contents

Introduction

This document describes multicast routing on the Adaptive Security Appliance (ASA), and common problems.

Feature Information

Note: for an updated content about the multicast routing on the Adaptive Security Appliance (ASA), Firepower Threat Defense (FTD) or the Secure Firewall Threat Defense (FTD) refer to these articles:
Troubleshoot Firepower Threat Defense IGMP and Multicast Basics
Troubleshoot Firepower Threat Defense and ASA Multicast PIM

Abbreviations/Acronyms

Acronyms	Explanation
FHR	First-Hop Router â€“ a hop directly connected to the source of the multicast traffic.
LHR	Last-Hop Router â€“ a hop directly connected to the receivers of the multicast traffic.
RP	Rendezvous-Point
DR	Designated Router
SPT	Shortest-Path Tree
RPT	Rendezvous-Point (RP) Tree, share tree
RPF	Reverse Path Forwarding
OIL	Outgoing Interface List
MRIB	Multicast Routing Information Base
MFIB	Multicast Forwarding Information Base
ASM	Any-Source Multicast
BSR	Bootstrap Router
SSM	Source-Specific Multicast
FP	Fast Path
SP	Slow Path

CP	Control Point
PPS	Packet Per Second rate

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

- PIM sparse-mode (Protocol Independent Multicast: RFC 4601)
- IGMP Stub-mode (Internet Group Management Protocol: RFC 2236)

PIM sparse-mode is the preferred choice because the ASA communicates with neighbors via 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.

Components of Multicast

In general a multicast infrastructure is composed of these components:

Sender => Host or Network device that originates the multicast stream. Examples are server that sends video and/or audio stream and network devices running a Routing Protocol such as EIGRP or OSPF.

Receiver => Host or device that receives the multicast stream. This term is more frequently used for hosts activelly interested in the traffic and use IGMP to join or leave the multicast group in question.

Routers / ASA => Network devices responsible to process and forward the multicast stream/traffic to other segments of the network when needed from source to client(s).

Multicast Routing Protocol => Protocol responsible to forward the multicast packets. Most common is PIM (Protocol Independent Multicast), but there are others like MOSPF for example.

Internet Group Management Protocol (IGMP) => Process used by clients to receive a multicast stream from a certain group.

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 goes directly from the source to the receivers (and bypass the RP).

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

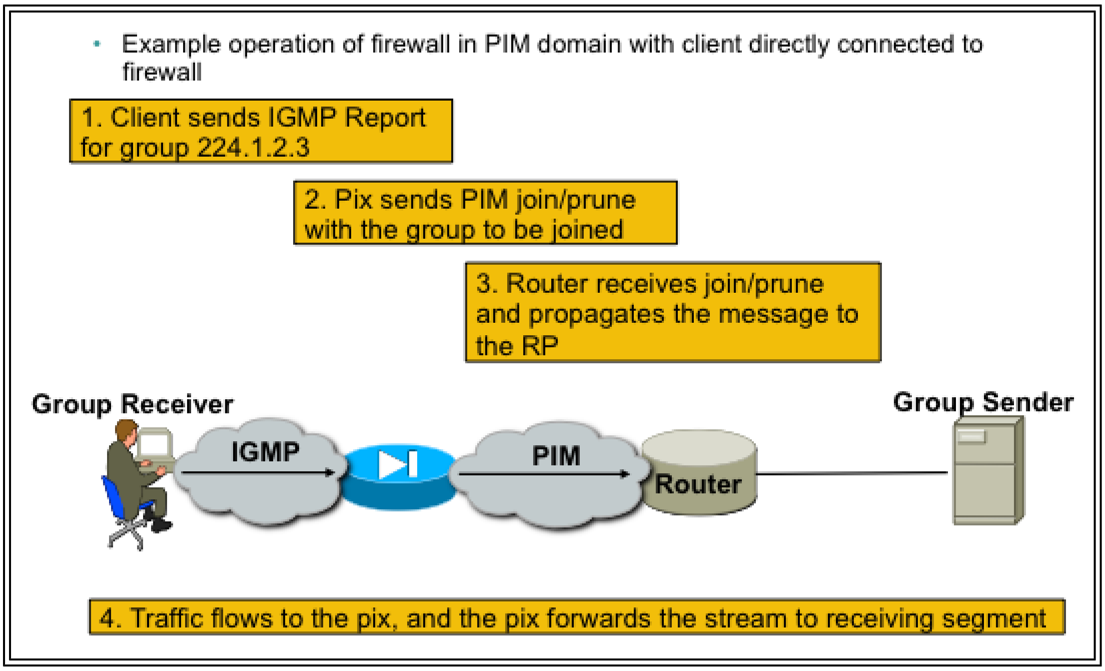

- Example operation of firewall in PIM domain with client directly connected to firewall

1. Client sends IGMP Report for group 224.1.2.3

2. Pix sends PIM join/prune with the group to be joined

3. Router receives join/prune and propagates the message to the RP

4. Traffic flows to the pix, and the pix forwards the stream to receiving segment

PIM Sparse-mode Sample Configuration

1. Enable multicast routing (global configuration mode).

<#root>

```
ASA(config)#

multicast-routing
```

2. Define the PIM Rendezvous-point address.

<#root>

```
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 blocks the inbound multicast packets).

<#root>

```
access-list 105 extended permit ip any host 224.1.2.3
access-group 105 in interface outside
```

PIM Sparse Mode Example:

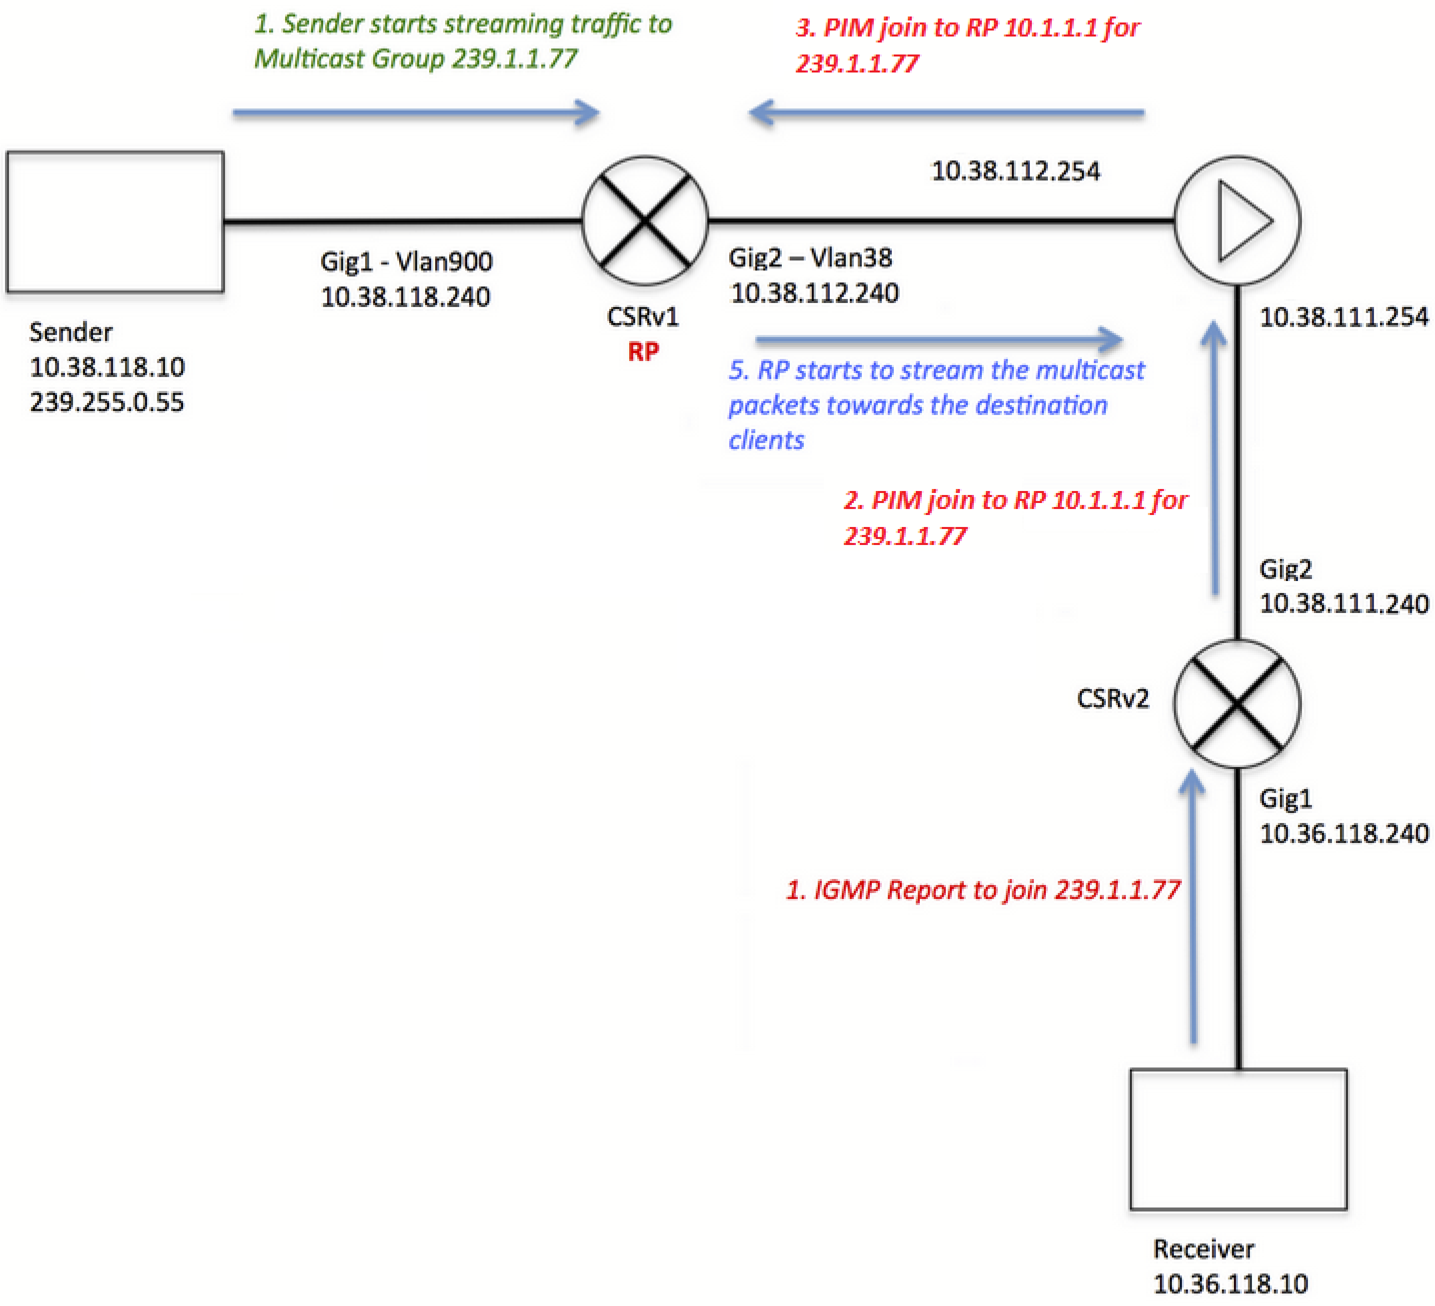
1. Sender starts streaming traffic to Multicast Group 239.1.1.77
3. PIM join to RP 10.1.1.1 for 239.1.1.77
10.38.112.254
Gig1 - Vlan900
10.38.118.240
Gig2 – Vlan38
10.38.112.240
10.38.111.254
Sender
10.38.118.10
239.255.0.55
CSRv1
RP
5. RP starts to stream the multicast packets towards the destination clients
2. PIM join to RP 10.1.1.1 for 239.1.1.77
Gig2
10.38.111.240
CSRv2
Gig1
10.36.118.240
1. IGMP Report to join 239.1.1.77
Receiver
10.36.118.10

Notice that the client IGMP registration (steps in red) and the stream is received by the server (steps in green) have been colored differently, and this was made this way to evidence that both process can occur indenpendently.

Client registration steps (red steps):

1. Client sends an IGMP Report for group 239.1.1.77

2. Router sends an PIM Join message to static RP configured (10.1.1.1) for group 239.1.1.77.

3. ASA sends to RP a PIM Join message for the group 239.1.1.77.

ASA displays PIM *,G entry on the **show mroute** command output:

<#root>

```
ciscoasa#

show mroute 239.1.1.77


Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(*, 239.1.1.77), 00:03:43/00:02:41, RP 10.1.1.1, flags: S
  Incoming interface: outside
  RPF nbr: 10.38.111.240
  Immediate Outgoing interface list:
    inside, Forward, 00:03:43/00:02:41
```

But as the source server have not started any stream, the "show mfib" output on the ASA does not display any received packets:

<#root>

```
ciscoasa#

show mfib 239.1.1.77


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

(*,239.1.1.77) Flags: C K
   Forwarding: 0/0/0/0, Other: 0/0/0
   outside Flags: A
   inside Flags: F NS
     Pkts: 0/0
```

Before the server starts to send any traffic to the multicast group, the RP only displays a "*.G" entry with no incoming interface on the list, as for example:

<#root>

```
CRSv#

show ip mroute 239.1.1.77

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, E - Extranet,
       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,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.77), 00:00:02/00:03:27, RP 10.1.1.1, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet2, Forward/Sparse-Dense, 00:00:02/00:03:27
```

Once the server starts to stream to the multicast group, the RP creates a "S,G" entry and puts the interface faced to the sender on the incoming interface list and starts to send the traffic downstream to the ASA:

<#root>

```
CRSv#

show ip mroute 239.1.1.77


...

(*, 239.1.1.77), 00:03:29/stopped, RP 10.1.1.1, flags: SF
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet2, Forward/Sparse-Dense, 00:03:29/00:02:58

(10.38.118.10, 239.1.1.77), 00:00:07/00:02:52, flags: FT
  Incoming interface: GigabitEthernet1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet2, Forward/Sparse-Dense, 00:00:07/00:03:22
```

Use these commands for verifications:

- **show mroute** command displays a "S,G" entry

- **show mfib** command displays forward packet counters

- **show conn** command displays connection related to the multicast group ip

<#root>

```
ciscoasa#

show mroute 239.1.1.77
```

```
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(*, 239.1.1.77), 00:06:22/00:02:50, RP 10.1.1.1, flags: S
  Incoming interface: outside
  RPF nbr: 10.38.111.240
  Immediate Outgoing interface list:
    inside, Forward, 00:06:22/00:02:50

(10.38.118.10, 239.1.1.77), 00:03:00/00:03:28, flags: ST
  Incoming interface: outside
  RPF nbr: 10.38.111.240
  Immediate Outgoing interface list:
    inside, Forward, 00:03:00/00:03:26

ciscoasa#
```

show mfib 239.1.1.77

```
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

(*,239.1.1.77) Flags: C K
   Forwarding: 15/0/1271/0, Other: 0/0/0
   outside Flags: A
   inside Flags: F NS
     Pkts: 0/15
(10.38.118.10,239.1.1.77) Flags: K
   Forwarding: 7159/34/1349/360, Other: 0/0/0
   outside Flags: A
   inside Flags: F NS
     Pkts: 7159/5

ciscoasa#
```

show conn all | i 239.1.1.77

```
UDP outside  10.38.118.10:58944 inside  239.1.1.77:5004, idle 0:00:00, bytes 10732896, flags -
UDP outside  10.38.118.10:58945 inside  239.1.1.77:5005, idle 0:00:01, bytes 2752, flags -
UDP outside  10.38.118.10:58944 NP Identity Ifc  239.1.1.77:5004, idle 0:00:00, bytes 0, flags -
UDP outside  10.38.118.10:58945 NP Identity Ifc  239.1.1.77:5005, idle 0:00:01, bytes 0, flags -
```

Note: Once the client closes the multicast client application, the host sends a IGMP Query message.

In case this is the only host known by the router as a client wants to receive the stream, the router sends a IGMP Prune message to the RP.

IGMP Stub-mode Operation

- In IGMP Stub-mode the ASA acts as a multicast client, and generates or forwards IGMP reports (also known as IGMP "joins") towards adjacent routers, to trigger the reception of multicast traffic
- Routers periodically sends 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 (with more efficient multicast traffic flows, ability to participate in PIM, etc).

This picture illustrates the basic operation of an ASA configured for IGMP Stub-mode:

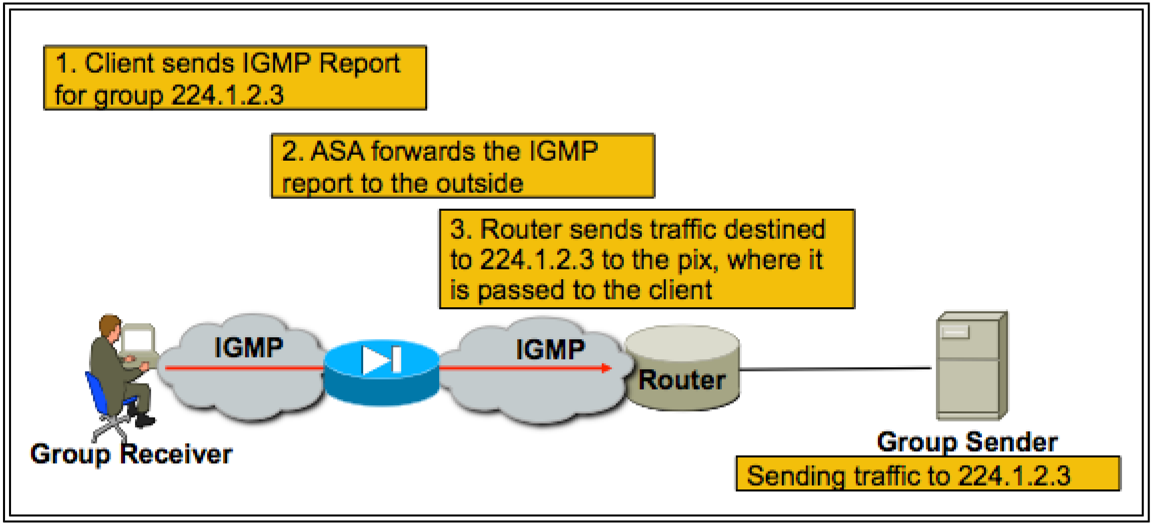
1. Client sends IGMP Report for group 224.1.2.3
2. ASA forwards the IGMP report to the outside
3. Router sends traffic destined to 224.1.2.3 to the pix, where it is passed to the client
IGMP
IGMP
Router
Group Receiver
Group Sender
Sending traffic to 224.1.2.3

IGMP Stub-mode Configuration

1. Enable multicast routing (global config mode).

<#root>

```
ASA(config)#
```

```
multicast-routing
```

2. On the interface on which the firewall receives the igmp reports, configure the igmp forward-interface command. Forward the packets out the interface towards the source of the stream. In this example, the multicast receivers are directly connected to the inside interface, and the multicast source is beyond the outside interface.

<#root>

```
!
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).

<#root>

```
ASA(config)#
access-list 105 extended permit ip any host 224.1.2.3
ASA(config)#
access-group 105 in interface outside
```

Often there is confusion around the different igmp interface sub-mode commands, and this diagram describes when to use each:

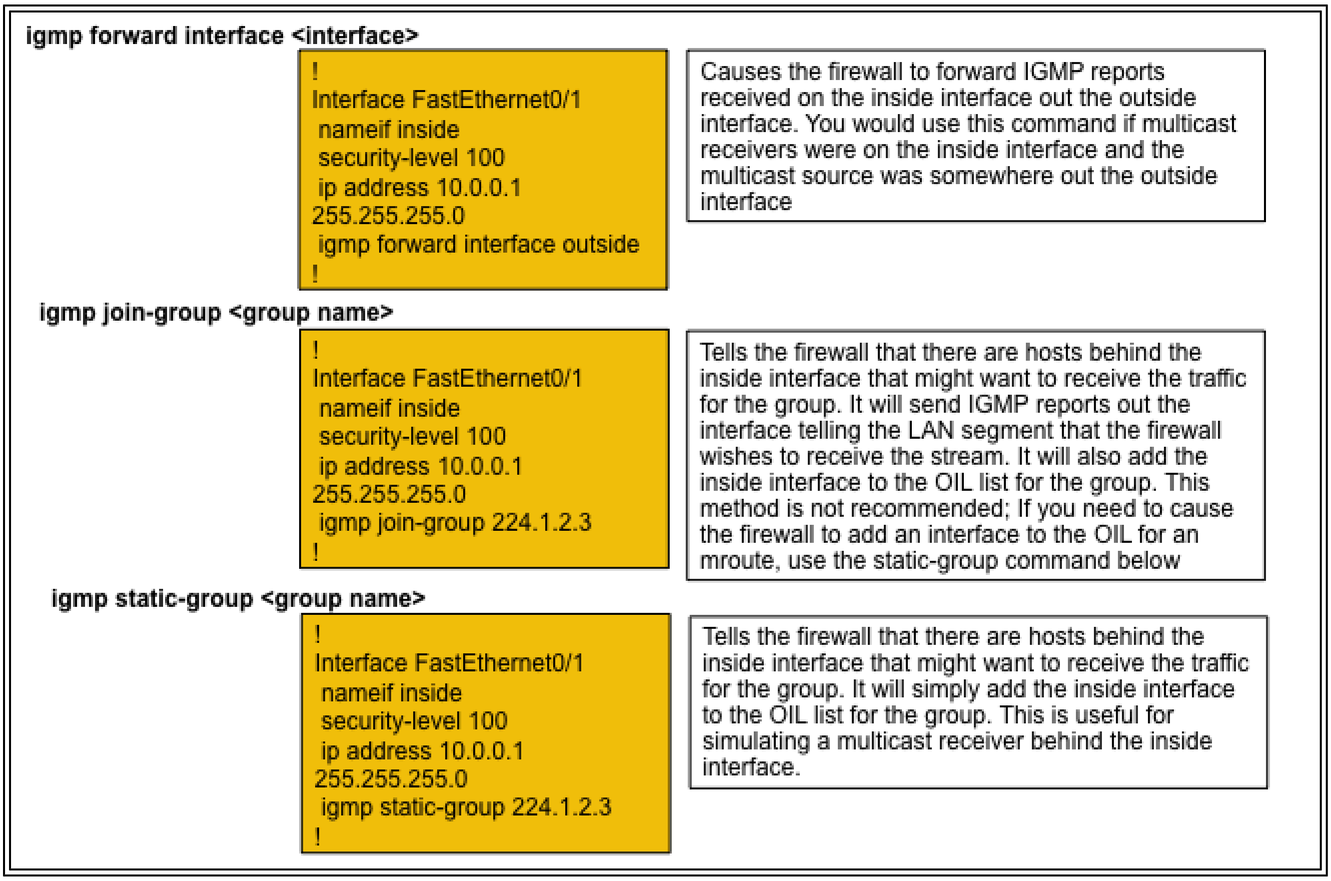

igmp forward interface <interface>

```
!
Interface FastEthernet0/1
 nameif inside
 security-level 100
 ip address 10.0.0.1
255.255.255.0
 igmp forward interface outside
!
```

Causes the firewall to forward IGMP reports received on the inside interface out the outside interface. You would use this command if multicast receivers were on the inside interface and the multicast source was somewhere out the outside interface

igmp join-group <group name>

```
!
Interface FastEthernet0/1
 nameif inside
 security-level 100
 ip address 10.0.0.1
255.255.255.0
 igmp join-group 224.1.2.3
!
```

Tells the firewall that there are hosts behind the inside interface that might want to receive the traffic for the group. It will send IGMP reports out the interface telling the LAN segment that the firewall wishes to receive the stream. It will also add the inside interface to the OIL list for the group. This method is not recommended; If you need to cause the firewall to add an interface to the OIL for an mroute, use the static-group command below

igmp static-group <group name>

```
!
Interface FastEthernet0/1
 nameif inside
 security-level 100
 ip address 10.0.0.1
255.255.255.0
 igmp static-group 224.1.2.3
!
```

Tells the firewall that there are hosts behind the inside interface that might want to receive the traffic for the group. It will simply add the inside interface to the OIL list for the group. This is useful for simulating a multicast receiver behind the inside interface.

Bidir PIM

In bidirectional PIM, there is no shared tree (SPT). This means three things:

1. The first hop router (connected to sender) does not send PIM Register packets to the RP.

2. The RP does not send PIM JOIN messages to join the source tree.

3. Routers in the path to the receiver send PIM join messages towards the RP to join the RPT.

This means that the ASA does not generate an (S,G) because devices do not join the SPT. All multicast traffic goes through the RP. The ASA forwards all multicast traffic as long as there is an (*,G). If there is no (*,G), that means that the ASA never received a PIM join packet. If that is the case, the ASA must not forward multicast packets.

Bidir PIM Configuration

1. Enable multicast routing (global configuration mode).

<#root>

```
ASA(config)#

 multicast-routing
```

2. Define the PIM Rendezvous-point address.

<#root>

```
ASA(config)#

pim rp-address 172.18.123.3 bidir
```

3. Allow the multicast packets in on the appropriate interface (necessary only if the security policy of the ASA blocks the inbound multicast packets).

<#root>

```
access-list 105 extended permit ip any host 224.1.2.3
access-group 105 in interface outside
```

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 is needed:

â€¢ A description of the network topology, the location fo the multicast senders, receivers, and rendezvous-point.
â€¢ The specific group IP address, as well as the ports and protocols employed.
â€¢ Syslogs generated by the ASA at the time the multicast stream has trouble.
â€¢ Specific show command output from the ASA command line interface:

```
<#root>

show mroute
show mfib
show pim neighbor
show route
show tech-support
```

â€¢ Packet captures to show if the multicast data arrives at the ASA, and if the packets are forwarded through the ASA (**take a note of the IP Time to Live (TTL) of the packet**. This can be seen with the command 'show capture x detail')
â€¢ Packet captures for IGMP and/or PIM packets. Example:

```
<#root>

capture cap1 interface outside match ip any host 239.1.1.77

  >>> This captures the multicast traffic itself

capture cappim1 interface inside match pim any any

         >>> This captures PIM Join/Prune messages

capture capigmp interface inside match igmp any any

          >>> This captures IGMP Report/Query messages
```

â€¢ Information from adjacent multicast devices (routers) such as "show mroute" and "show mfib".
â€¢ Packet captures and/or show commands to determine if the ASA drops the multicast packets. The 'show asp drop' command can be used to determine if the ASA drops 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):

```
<#root>

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 the test:

<#root>

```
ciscoasa#

clear mfib counters
```

Packet Captures

The onboard packet capture utility is very useful to troubleshoot multicast problems. In this example, all ingress packets at the DMZ interface, destined to 239.17.17.17 is captured:

<#root>

```
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
....
```

The output of the **show capture x detail** command shows the TTL of the packets, which is quite useful. In this output, the TTL of the packet is 1 (and the ASA passes this packet since it does not decrement the TTL of IP packets by default) but a downstream router would drop the packets:

<#root>

```
ASA#
```

```
show cap capout detail
```

```
453 packets captured
...
   1: 14:40:39.427147 c062.6baf.8dc3 0100.5e7f.02c3 0x8100 Length: 1362
      802.1Q vlan#1007 P0 10.4.2.95.1806 > 239.255.2.195.5000:  [udp sum ok] udp 1316 (DF) [ttl 1] (id
```

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

<#root>

```
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#
```

Note that the TTL of the packets can be seen with the 'show capture x detail' command.

Here we can see ASP drop captures taken that show the dropped multicast packets and the reason for the drops (punt-rate-limit):

<#root>

```
ASA#
```

```
show cap capasp det
```

```
 12: 14:37:26.538332 c062.6baf.8dc3 0100.5e7f.02c3 0x8100 Length: 1362
      802.1Q vlan#1007 P0 10.76.4.95.1806 > 239.255.2.195.5000:  [udp sum ok] udp 1316 (DF) [ttl 1] (id
  13: 14:37:26.538439 c062.6baf.8dc3 0100.5e7f.02c3 0x8100 Length: 1362
      802.1Q vlan#1007 P0 10.76.4.95.1806 > 239.255.2.195.5000:  [udp sum ok] udp 1316 (DF) [ttl 1] (id
```

Example ASA PIM Sparse-Mode Multicast Deployment

This diagrams illustrate how the ASA interacts with neighbor devices in PIM sparse-mode.

Understand the network topology

Determine exactly the location of the senders and receivers of the specific multicast stream. Also, determine the multicast group IP address, as well as the location of the rendezvous point.

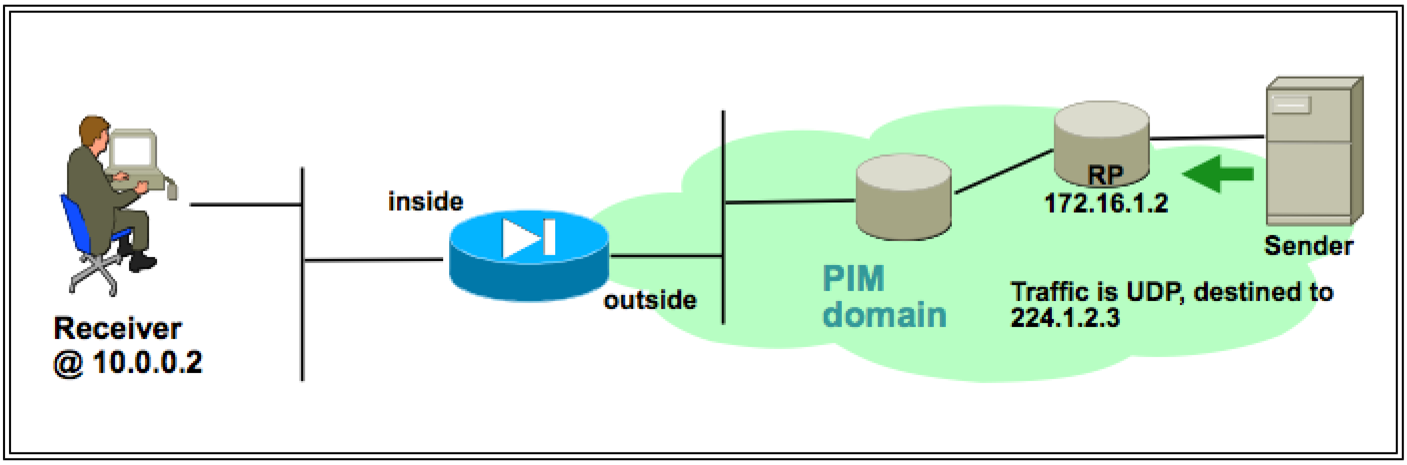

In this case, the data can be received at the outside interface of the ASA, and forwarded to the multicast receiver on the inside interface. Because the receiver is in the same IP subnet as the inside interface of the ASA, expect to see an IGMP Report received at the inside interface when the client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

Verify that 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.

Verify that 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.

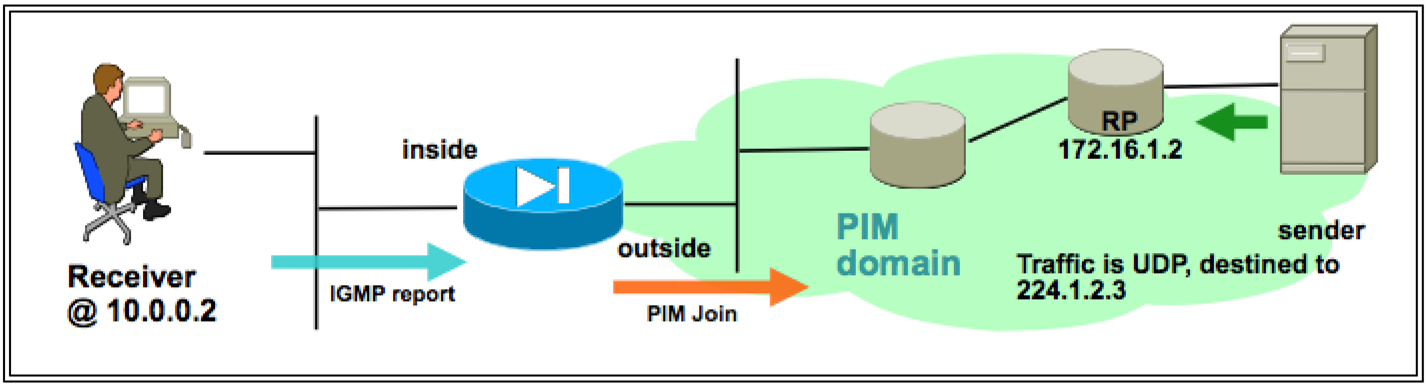

This output is from debug pim group 224.1.2.3 and shows that the ASA successfully sends the PIM join message. The sender of the multicast stream is 192.168.1.50.

```
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 192.168.1.50
```

Verify that the ASA receives and forwards the multicast stream

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

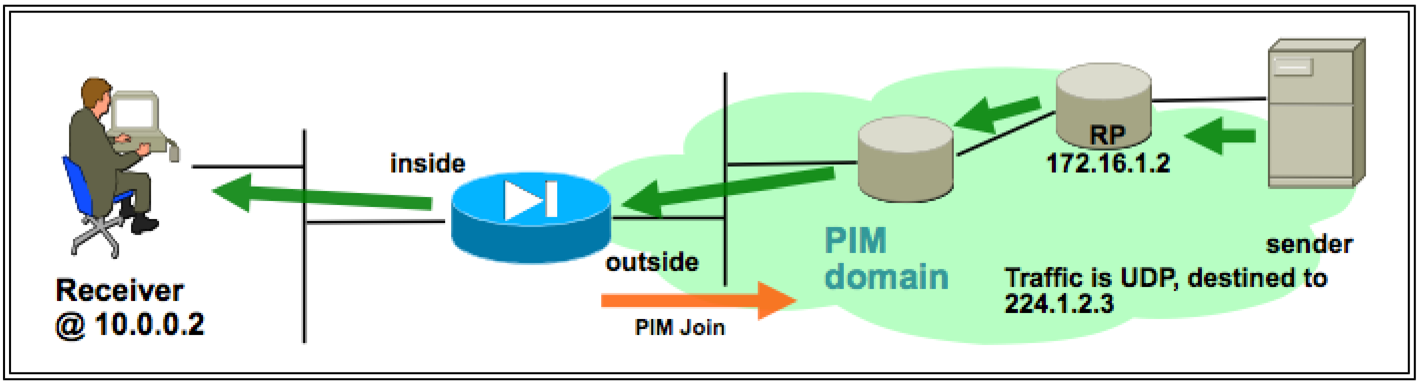

The **show mroute** and **show mfib** commands, as well as packet captures, can be used to verify the ASA receives and forwards the multicast packets.

A connection is built in the connection table to represent the multicast stream:

<#root>

```
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

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. This diagram 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.

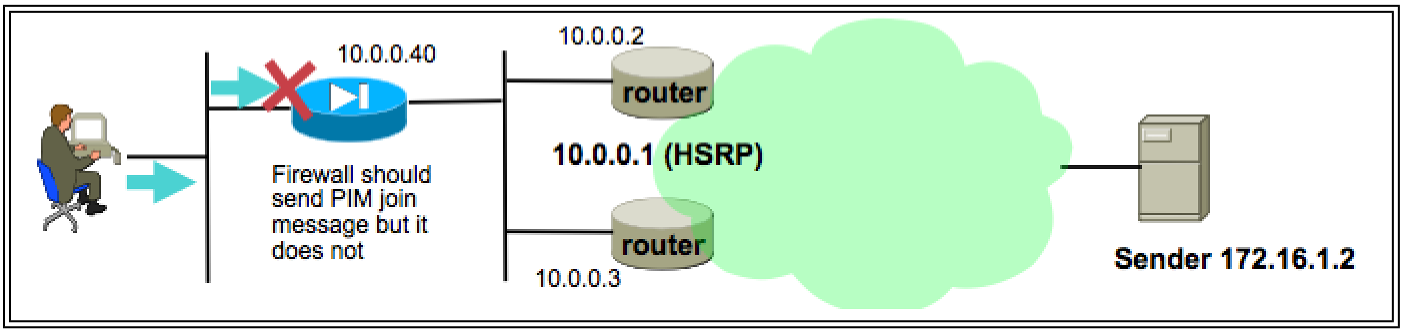

The output of **debug pim** command shows that the ASA cannot send the PIM message to the upstream next-hop router:

```
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 routing table configuration and the HSRP configuration used by the PIM neighbors.

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

<#root>

```
ciscoasa#

show 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:

<#root>

```
ciscoasa#

show 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 interface IP addresses (10.0.0.2 or 10.0.0.3). In this case, this command allows the ASA to send PIM messages directed towards the multicast sender at 172.16.1.2:

<#root>

```
ciscoasa(config)#

mroute 172.16.1.2 255.255.255.255 10.0.0.3
```

Once this is done the multicast routing table overrides the unicast routing table of the ASA, and the ASA sends 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 is generated and the packet is simply dropped, and stream reception fails.

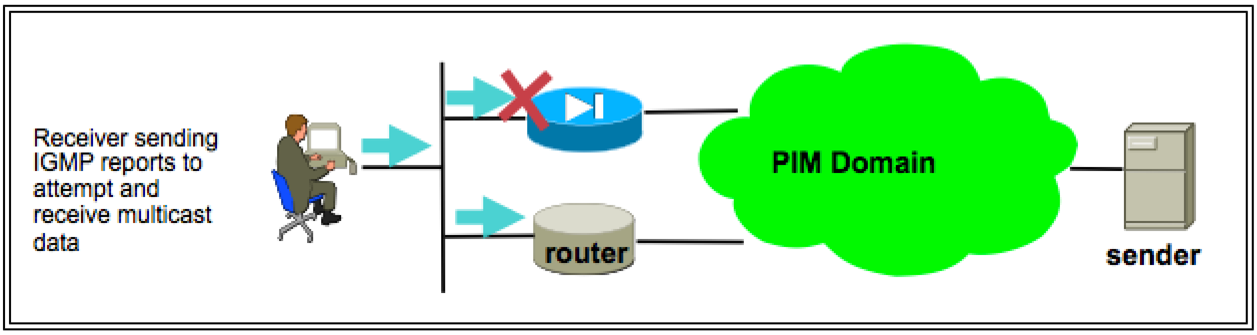
Receiver sending IGMP reports to attempt and receive multicast data
PIM Domain
router
sender

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

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

<#root>

```
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 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 drop IGMP reports. The solution is to disable PIM on the interface (with the *no pim* command on the interface involved), or to make the ASA the DR for the segment via the **pim dr-priority** interface command.

IGMP Reports are Denied by the Firewall when IGMP Interface Limit is Exceeded

By default, the ASA allows for 500 current active joins (reports) tracked on an interface. This is the maximum value that is configurable. If a large number of multicast streams are requested by clients off of an interface, the maximum of 500 active joins can be encountered, and the ASA could ignore additional incoming IGMP reports from the multicast receivers.

To confirm if this is the cause of a multicast failure, issue the command 'show igmp interface *interfacename*' and look for the 'IGMP limit' information for the interface.

<#root>

```
ASA#

show igmp interface inside

Hosting-DMZ is up, line protocol is up
  Internet address is 10.11.27.13/24
  IGMP is enabled on interface
  Current IGMP version is 2
  IGMP query interval is 125 seconds
  IGMP querier timeout is 255 seconds
  IGMP max query response time is 10 seconds
  Last member query response interval is 1 seconds
  Inbound IGMP access group is:


IGMP limit is 500, currently active joins: 500

  Cumulative IGMP activity: 7018 joins, 6219 leaves
  IGMP querying router is 10.11.27.13 (this system)
```

```
DEBUG - IGMP: Group x.x.x.x limit denied on outside
```

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 the **debug igmp** command shows 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*:

<#root>

```
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
```

One option is to add an mroute for the sender of the traffic. In this example, the mroute command is used to satisfy the RPF check for multicast traffic sourced from 172.16.1.2 received on the outside interface:

<#root>

```
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 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 attempts to receive traffic along the source-specific tree. This router generates a new PIM join for the group and send it towards the sender of the multicast stream (and not towards the RP, as before).

The sender of the multicast traffic can 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 are dropped and this message is 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, that points 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 fails 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 test 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 can experience high CPU and the multicast stream can experience packet drops if all of these are true about the multicast topology:

1. The ASA acts 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 mentioned symptoms are encountered, then due do a design limitation the ASA is 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 has this problem, complete these steps:

Step 1: Check if the ASA is the RP:

<#root>

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

Step 2: Check if the ASA is the last hop router:

<#root>

```
show igmp group
```

```
<mcast_group_IP>
```

Step 3: Check if the ASA is the first hop router:

<#root>

```
show mroute
```

```
<mcast_group_IP>
```

These steps can be taken to mitigate this problem:

- Modify the topology so that the ASA is not the RP. Or, make the sender or receiver not directly connected to the ASA

- Instead of PIM, use IGMP stub-mode for mulitcast forwarding.

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 in the process of creation, some multicast packets can 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 are no longer be rate limited.

Packets dropped for this reason have the ASP drop reason of "(punt-rate-limit) Punt rate limit exceeded". This 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:

<#root>

```
ASA #

show 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 r
   2: 16:14:49.919172 10.23.2.2.810 > 239.255.123.123.890:  udp 32 Drop-reason: (punt-rate-limit) Punt r
2 packets shown
```

A Disconnecting Multicast Receiver Interrupts Multicast Group Reception On Other Interfaces

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

The issue is identified by the Cisco bug ID CSCeg48235 IGMP Leave on one interface interrupts multicast traffic 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 multi

The problem is triggered when the firewall forwards the IGMP leave for the group towards the upstream de

Conditions:
The PIX or ASA must be configured for IGMP stub mode multicast. IGMP stub mode is a legacy multicast for

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, so their IGMP
```

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

With this specific issue the ASA drops 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 drops 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 are dropped with the ASP drop counter "FP no mcast output intrf (no-mcast-intrf)".

The ASA continuously drops some packets (but not all) in a multicast stream due to Control point rate limiting

The traffic is most likely is rate limited by the control point due to punt-rate-limit. Look at the asp drop output and captures to confirm:

<#root>

```
ASA#

show asp drop


Frame drop:
  Punt rate limit exceeded (punt-rate-limit)                               1492520




ASA#  show cap capasp det
 12: 14:37:26.538332 c062.6baf.8dc3 0100.5e7f.02c3 0x8100 Length: 1362
      802.1Q vlan#1007 P0 10.76.4.95.1806 > 239.255.2.195.5000:  [udp sum ok] udp 1316 (DF) [ttl 1] (id
```

The mfib entry shows that all traffic is process switched:

<#root>

```
ASA(config)#

show mfib  239.255.2.1195

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

(*,239.255.2.195) Flags: C K
   Forwarding: 4278/50/1341/521, Other: 0/0/0
   Outside-1007 Flags: A
   RDEQ-to-Corporate Flags: F NS
     Pkts: 0/4278                               <---- HERE
```

The multicast routing table shows an (*,G) but no (S,G).

<#root>

```
ASA(config)#
```

```
show mroute  239.255.2.1195
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(*, 239.255.2.195), 00:44:03/00:02:44, RP 10.1.135.10, flags: S
  Incoming interface: Outside-1007
  RPF nbr: 10.100.254.18
  Immediate Outgoing interface list:
    RDEQ-to-Corporate, Forward, 00:44:03/00:02:44
```

The problem here is that the TTL of the packets multicast data packets that arrive at the ASA is 1. The ASA is forwarding these packets to the downstream device (because it does not decrement TTL), but the router downstream drops the packets. As a result, the downstream router does not send a PIM (S,G) join (a source-specific join) to the ASA towards the sender. The ASA does not build an (S,G) entry until it receives this PIM join. Because the (S,G) is not built, all the multicast traffic is process switched that results in rate limit.

The resolution to this problem is to ensure the TTL of the packets is not 1, that allows the downstream device to send the source-specific join towards the sender; this causes the ASA to install a source-specific mroute in the table, and then all the packets are fast-switched (instead of processed switched) and the traffic must flow through the ASA without a problem.

The multicast stream is halted due to a PIM ASSERT message

If two network devices forward the same multicast packets onto the same subnet then ideally, one of them must stop forwarding the packets (since it is a waste to duplicate the stream). If routers running PIM detect they receive the same packets they also generate on that same interface, they generate ASSERT messages on that LAN to elect which network device stops forwarding the stream.

More information about this message can be seen in a section of RFC 4601 related to the ASSERT process.

The debugs show that the ASA receives an IGMP report for group 239.1.1.227, but it ignores the report due to the assert message it receives from a neighboring router:

```
IPv4 PIM: (*,239.1.1.227) Periodic J/P scheduled in 50 secs
IPv4 PIM: (*,239.1.1.227) J/P adding Join on outside
IPv4 PIM: (10.99.41.205,239.1.1.227)RPT J/P adding Prune on outside
IPv4 PIM: (10.99.41.253,239.1.1.227)RPT J/P adding Prune on outside
IGMP: Received v2 Report on inside from 10.20.213.204 for 239.1.1.227
IGMP: Updating EXCLUDE group timer for 239.1.1.227
IPv4 PIM: (10.99.41.253,239.1.1.227) Received [15/110] Assert from 10.20.13.2 on inside
IPv4 PIM: (10.99.41.253,239.1.1.227) Assert processing message wins
IPv4 PIM: (10.99.41.253,239.1.1.227) inside Update assert timer (winner 10.20.13.2)
```

This problem was observed in a production network where two sites were accidently bridged at layer-2, such that the LAN where the multicast receivers was on had two devices forwarding multicast traffic towards them. Due to another network problem, the ASA and another device could not detect each other via PIM hellos, and therefore they both assumed the Designated Router role for the LAN. This caused the multicast traffic to work for a while, and then fail when the ASSERT messages were sent by the devices. To resolve the problem, the incorrect connection that bridged the devices at layer 2 was disabled, and then the problem was resolved.

ASA sends PIM Join but it is not processed by Neighbour due to size of packet greater than MTU

This was observed in 629575899. The ASA was configured for Jumbo Frames and the 4900 was not. When the client requested more than 73 multicast streams, certain multicast streams would not work. 73 SGs would create a PIM Join message of size 1494, which was still within MTU. 74SGs would create a PIM Join message larger than 1500, which caused the 4900M to drop the packet inbound.

The fix for this issue was:

1. Ensure Jumbo Frames are enabled globally on the 4900M

2. Configure both the physical interface and SVI with an MTU of 9216