Solução de problemas do Firepower Threat Defense, Conceitos básicos de IGMP e multicast

Contents

Introdução

Este documento descreve os conceitos básicos do multicast e como o Firepower Threat Defense (FTD) implementa o Internet Group Management Protocol (IGMP).

Pré-requisitos

Requisitos

Conhecimento básico de roteamento IP.

Componentes Utilizados

As informações neste documento foram criadas a partir de dispositivos em um ambiente de laboratório específico. Todos os dispositivos utilizados neste documento foram iniciados com uma configuração (padrão) inicial. Se a rede estiver ativa, certifique-se de que você entenda o impacto potencial de qualquer comando.

O conteúdo deste artigo também se aplica ao software Adaptive Security Appliance (ASA).

As informações neste documento são baseadas nestas versões de software e hardware:

- Defesa contra ameaças do Cisco Firepower 4125 versão 7.1.0.
- Firepower Management Center (FMC) versão 7.1.0.
- ASA versão 9.19.1.

Informações de Apoio

Definições

- Unicast = de um único host para outro host (um para um).
- Broadcast = de um único host para TODOS os hosts possíveis (um para todos).
- **Multicast = de um host de um grupo de hosts para um grupo de hosts (um para muitos ou muitos para muitos).**
- Anycast = de um host para o host mais próximo de um grupo (um para um de muitos).

Conceitos básicos

- O Multicast RFC 988 foi escrito em 1986 por Steve Deering.
- O Multicast IPv4 usa o intervalo 224.0.0.0/4 (primeiros 4 bits 1110) - 224.0.0.0 - 239.255.255.255.
- Para IPv4, o endereço MAC de L2 deriva do IP multicast de L3: 01005e (24 bits) + 25º bit sempre 0 + 23 bits inferiores do endereço IPv4 multicast.
- O Multicast IPv6 usa o intervalo FF00::/8 e é mais flexível que o multicast IPv4, pois pode incorporar o IP do ponto de encontro (RP).
- Para IPv6, o endereço MAC de L2 deriva do multicast de L3: 3333 + 32 bits inferiores do endereço IPv6 multicast.
- Vantagens do multicast: eficiência devido à carga reduzida na origem. Desempenho, pois evita a duplicação ou a inundação de tráfego.
- Desvantagens de multicast: transporte não confiável (baseado em UDP), sem prevenção de congestionamento, entrega fora de sequência.
- O multicast não é suportado na Internet pública, pois requer todos os dispositivos no caminho para ativá-lo. Normalmente, usado quando todos os dispositivos estão sob uma autoridade administrativa comum.
- Aplicações típicas de multicast: fluxo de vídeo interno, videoconferência.

Multicast versus unicast replicado

No Unicast Replicado, a origem cria várias cópias do mesmo pacote unicast (réplicas) e as envia para vários hosts de destino. O multicast move a carga do host de origem para a rede, enquanto no Unicast Replicado todo o trabalho é feito no host de origem.

Configurar

Conceitos básicos de IGMP

- O IGMP é a "linguagem" falada entre os receptores multicast e o dispositivo L3 local (normalmente um roteador).
- O IGMP é um protocolo da camada 3 (como o ICMP) e usa o **número 2 do protocolo IP.**
- Existem atualmente 3 versões de IGMP. A versão padrão do IGMP no firewall é a versão 2. **Somente as versões 1 e 2 são suportadas atualmente.**
- Entre IGMPv1 e IGMPv2, as principais diferenças são:
    - IGMPv1 não tem mensagem de Grupo de Saída.
    - O IGMPv1 não tem uma consulta específica de grupo (usada pelo firewall quando um host sai de um grupo multicast).

- IGMPv1 não tem processo de eleição de consultante.

- **O IGMPv3 não é suportado atualmente** no ASA/FTD, mas como referência, a diferença importante entre o IGMPv2 e o IGMPv3 é a inclusão de uma consulta específica de grupo e origem no IGMPv3, que é usada no Source-Specific Multicast (SSM).
- Consultas IGMPv1/IGMPv2/IGMPv3 = **224.0.0.1**
  Licença de IGMPv2 = **224.0.0.2**
  Relatório de associação IGMPv3 = **224.0.0.22**
- Se um host desejar ingressar pode enviar uma mensagem de **relatório de associação IGMP não solicitado**:

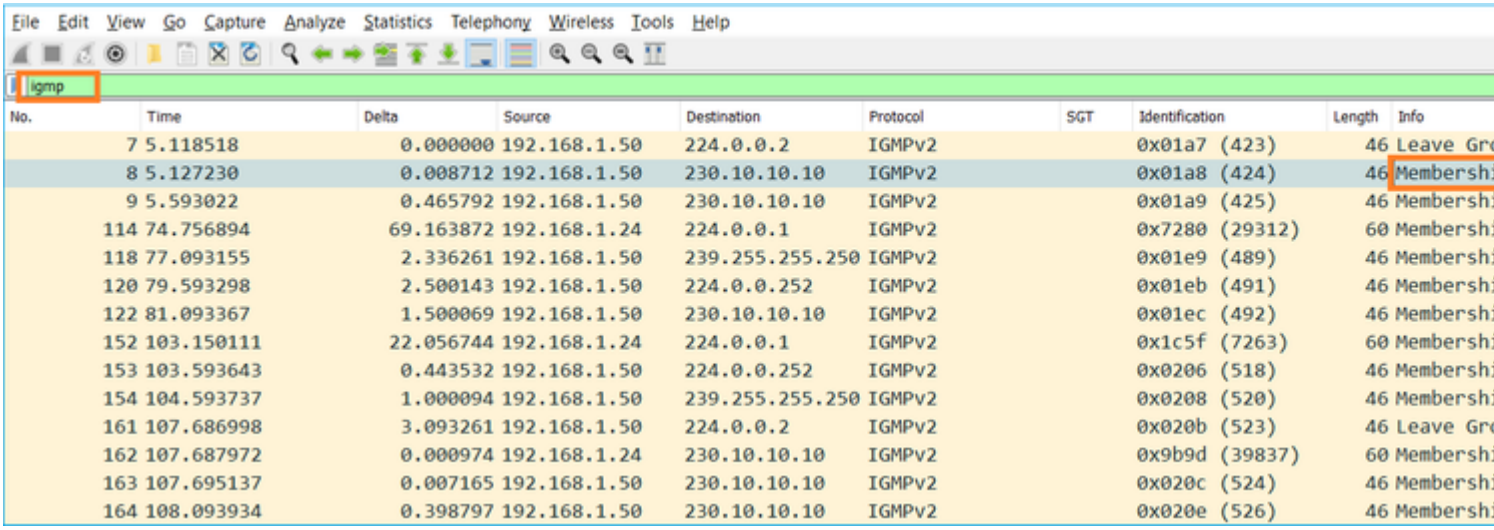

- Do ponto de vista do firewall, há **2 tipos de consultas IGMP: consultas gerais** e **consultas específicas de grupo**
- Quando o firewall recebe uma mensagem IGMP Leave Group, ele precisa verificar se há outros membros desse grupo na sub-rede. Por esse motivo, o firewall envia uma **consulta específica ao grupo:**

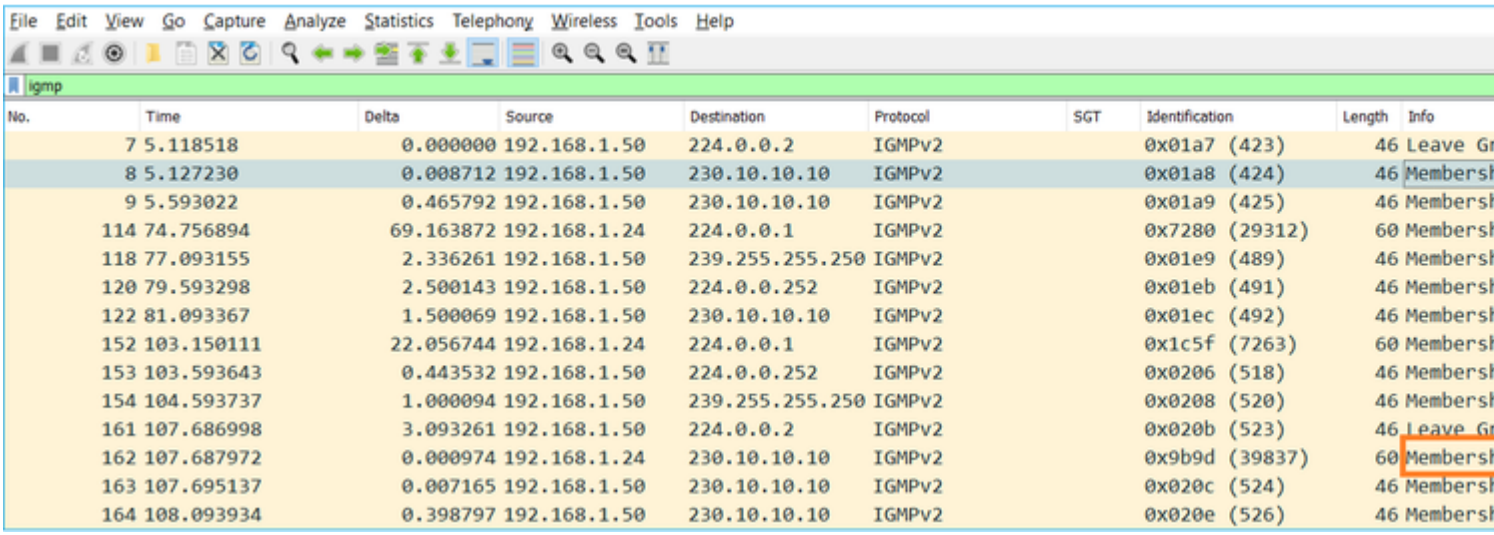

- Em sub-redes onde há vários roteadores/firewalls, um **consultante** (um dispositivo que envia todas as consultas IGMP) é escolhido:

```
firepower#
```

```
show igmp interface INSIDE
```

```
INSIDE is up, line protocol is up
  Internet address is 192.168.1.97/24
  IGMP is enabled on interface
  Current IGMP version is 2
  IGMP query interval is 125 seconds
  IGMP querier timeout is 60 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: 2
  Cumulative IGMP activity: 21 joins, 20 leaves
```

```
 IGMP querying router is 192.168.1.97 (this system)
```

```
<-- IGMP querier
```

- No FTD, semelhante a um ASA clássico, você pode habilitar **debug igmp** para ver mensagens relacionadas ao IGMP:

<#root>

```
firepower#
```

```
debug igmp
```

```
IGMP debugging is on
IGMP: Received v2 Query on DMZ from 192.168.6.1
```

```
IGMP: Received v2 Report on INSIDE from 192.168.1.50 for 239.255.255.250
```

```
<-- Received an IGMP packet
IGMP: group_db: add new group 239.255.255.250 on INSIDE
IGMP: MRIB updated (*,239.255.255.250) : Success
IGMP: Switching to EXCLUDE mode for 239.255.255.250 on INSIDE
IGMP: Updating EXCLUDE group timer for 239.255.255.250
IGMP: Received v2 Report on INSIDE from 192.168.1.50 for 230.10.10.10
IGMP: group_db: add new group 230.10.10.10 on INSIDE
IGMP: MRIB updated (*,230.10.10.10) : Success
IGMP: Switching to EXCLUDE mode for 230.10.10.10 on INSIDE
IGMP: Updating EXCLUDE group timer for 230.10.10.10
IGMP: Send v2 general Query on INSIDE
IGMP: Received v2 Query on INSIDE from 192.168.1.97
IGMP: Send v2 general Query on OUTSIDE
IGMP: Received v2 Query on OUTSIDE from 192.168.103.91
IGMP: Received v2 Report on INSIDE from 192.168.1.50 for 239.255.255.250
IGMP: Updating EXCLUDE group timer for 239.255.255.250
IGMP: Received v2 Report on INSIDE from 192.168.1.50 for 230.10.10.10
IGMP: Updating EXCLUDE group timer for 230.10.10.10
```

- Um host normalmente deixa um grupo multicast com uma mensagem **Leave Group** (IGMPv2).

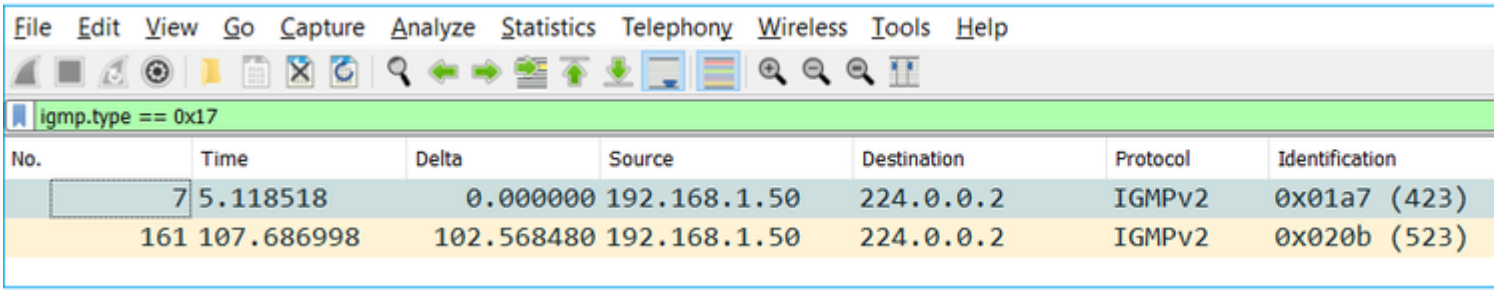

Tarefa 1 - Tráfego de multicast do plano de controle

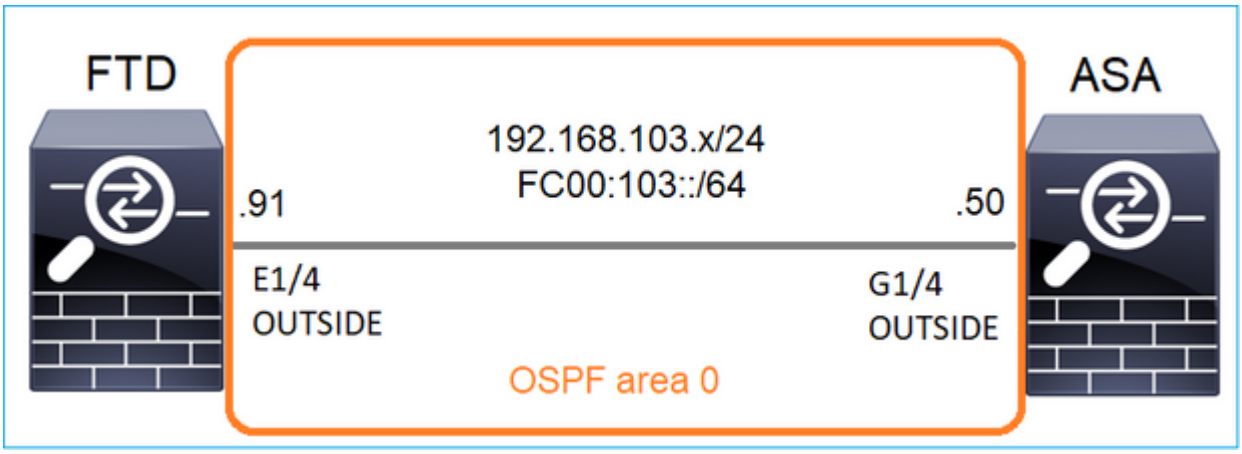

Configure um OSPFv2 e um OSPFv3 entre o FTD e o ASA. Verifique como os 2 dispositivos lidam com o tráfego multicast de L2 e L3 gerado pelo OSPF.

Solução

configuração de OSPFv2

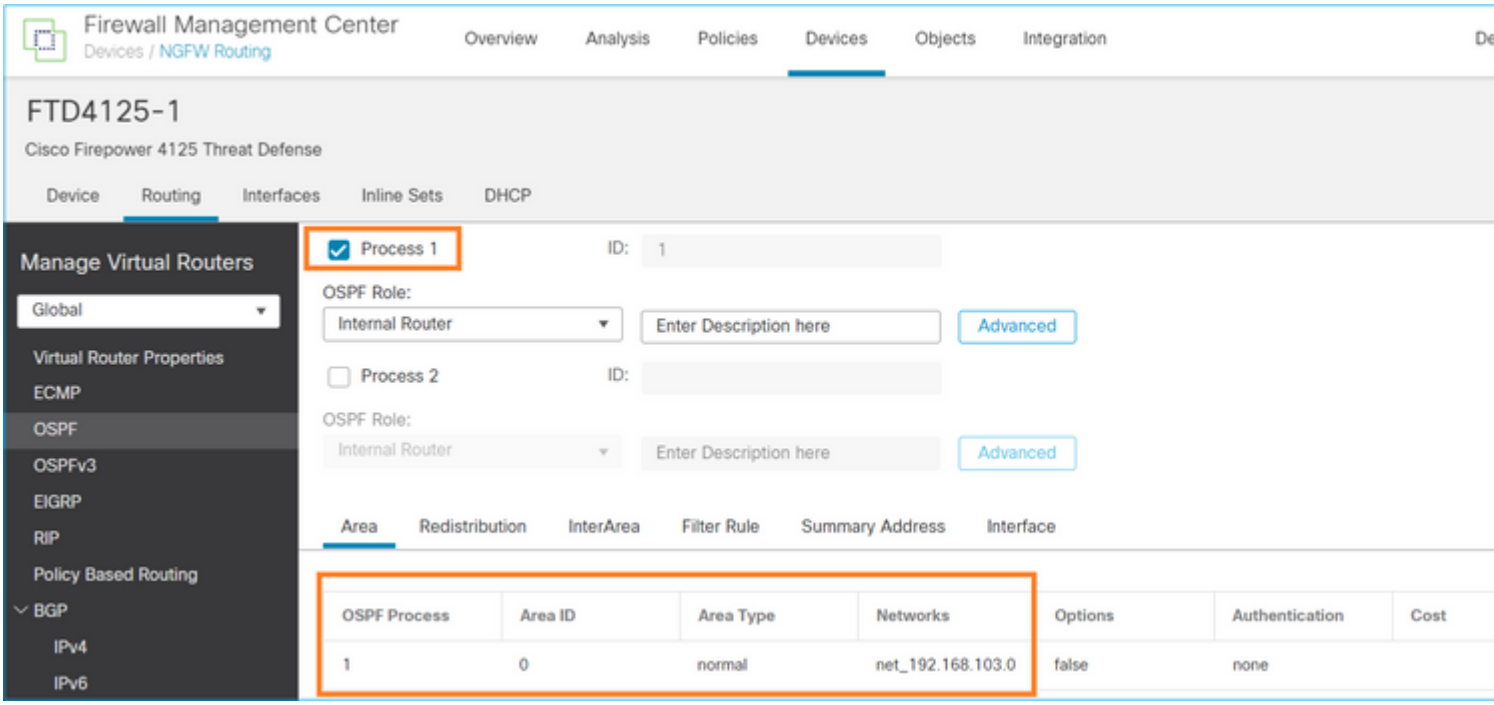

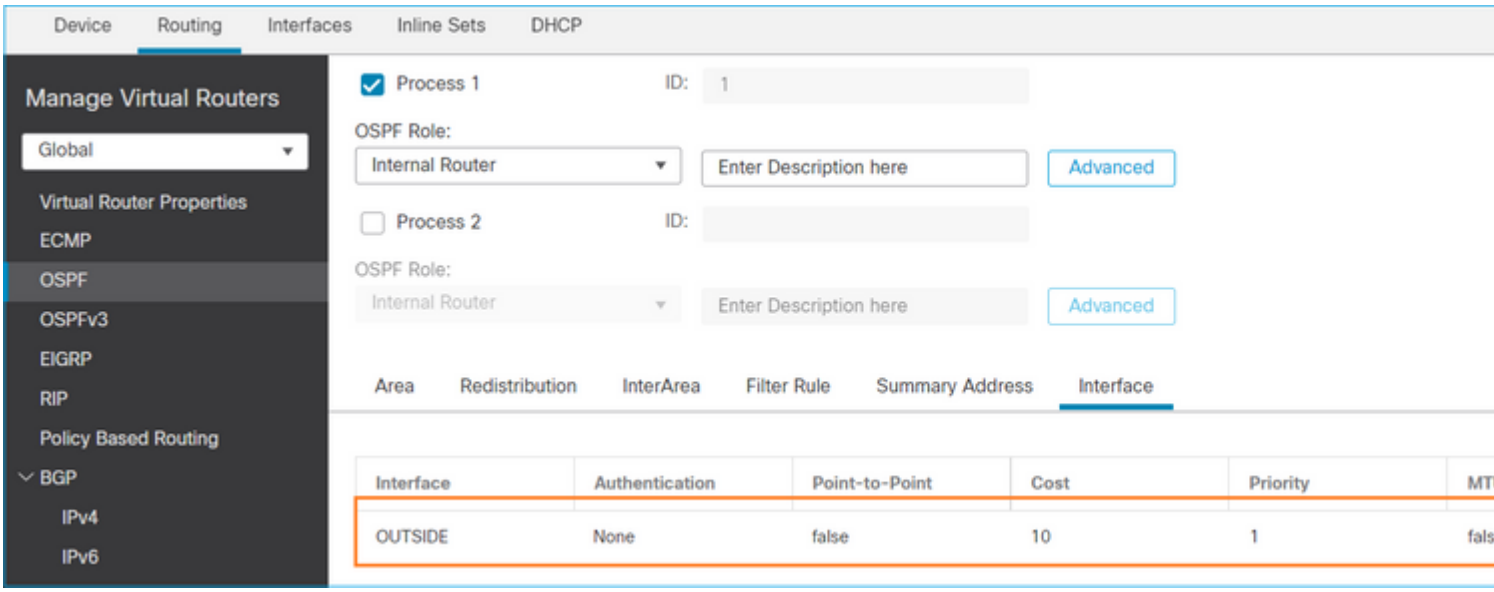

Da mesma forma, para OSPFv3

Configuração na CLI do FTD:

<#root>

```
router ospf 1

 network 192.168.103.0 255.255.255.0 area 0

 log-adj-changes
!

ipv6 router ospf 1

 no graceful-restart helper
 log-adjacency-changes
!
interface Ethernet1/4
nameif OUTSIDE
security-level 0
ip address 192.168.103.91 255.255.255.0
ipv6 address fc00:103::91/64
ospf authentication null

ipv6 ospf 1 area 0
```

A configuração cria essas entradas nas tabelas de permissão do Caminho de Segurança Acelerado (ASP) de FTD para que o tráfego multicast de entrada não seja bloqueado:

<#root>

```
firepower#

show asp table classify domain permit

...
in id=0x14f922db85f0, priority=13,

domain=permit, deny=false
```

```
<-- permit the packets
        hits=1, user_data=0x0, cs_id=0x0, reverse, flags=0x0, protocol=89
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any


dst ip/id=224.0.0.5, mask=255.255.255.255,

 port=0, tag=any, dscp=0x0, nsg_id=none     <-- OSPF for IPv4


input_ifc=OUTSIDE

(vrfid:0), output_ifc=identity(vrfid:0)     <-- ingress interface
in id=0x14f922db9350, priority=13,

domain=permit, deny=false

<-- permit the packets
        hits=0, user_data=0x0, cs_id=0x0, reverse, flags=0x0, protocol=89
        src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any


 dst ip/id=224.0.0.6, mask=255.255.255.255

, port=0, tag=any, dscp=0x0, nsg_id=none     <-- OSPF for IPv4


input_ifc=OUTSIDE

(vrfid:0), output_ifc=identity(vrfid:0)     <-- ingress interface
```

Para IPv6:

<#root>

```
...
in id=0x14f923fb16f0, priority=13,

domain=permit, deny=false

 <-- permit the packets
        hits=1, user_data=0x0, cs_id=0x0, reverse, flags=0x0, protocol=89
        src ip/id=::/0, port=0, tag=any


dst ip/id=ff02::5/128

, port=0, tag=any, , nsg_id=none     <-- OSPF for IPv6


input_ifc=OUTSIDE

(vrfid:0), output_ifc=identity(vrfid:0)  <-- ingress interface
in id=0x14f66e9d4780, priority=13,

domain=permit, deny=false

<-- permit the packets
        hits=0, user_data=0x0, cs_id=0x0, reverse, flags=0x0, protocol=89
        src ip/id=::/0, port=0, tag=any


dst ip/id=ff02::6/128
```

```
, port=0, tag=any, , nsg_id=none     <-- OSPF for IPv6
```

input_ifc=OUTSIDE

```
(vrfid:0), output_ifc=identity(vrfid:0)  <-- ingress interface
...
```

As adjacências de OSPFv2 e OSPFv3 são UP:

<#root>

firepower#

show ospf neighbor

```
Neighbor ID Pri State Dead Time Address Interface
192.168.103.50 1
```

FULL/BDR

```
 0:00:35 192.168.103.50 OUTSIDE   <-- OSPF neighbor is up
```

firepower#

show ipv6 ospf neighbor

```
Neighbor ID Pri State Dead Time Interface ID Interface
192.168.103.50 1
```

FULL/BDR

```
 0:00:34 3267035482 OUTSIDE       <-- OSPF neighbor is up
```

Estas são as sessões OSPF multicast terminadas na caixa:

<#root>

firepower#

show conn all | include OSPF

```
OSPF OUTSIDE fe80::2be:75ff:fef6:1d8e NP Identity Ifc ff02::5, idle 0:00:09, bytes 5924, flags
OSPF OUTSIDE 192.168.103.50 NP Identity Ifc 224.0.0.5, idle 0:00:03, bytes 8904, flags
OSPF OUTSIDE ff02::5 NP Identity Ifc fe80::f6db:e6ff:fe33:442e, idle 0:00:01, bytes 6304, flags
OSPF OUTSIDE 224.0.0.5 NP Identity Ifc 192.168.103.91, idle 0:00:00, bytes 25220, flags
```

Como teste, habilite a captura para IPv4 e limpe as conexões com o dispositivo:

<#root>

firepower#

```
capture CAP interface OUTSIDE trace

firepower#

clear conn all

12 connection(s) deleted.
firepower#

clear capture CAP

firepower# !
```

Aviso: isso causa uma interrupção! O exemplo é mostrado apenas para fins de demonstração!

Os pacotes OSPF capturados:

<#root>

```
firepower# show capture CAP | include proto-89

1: 12:25:33.142189 192.168.103.50 > 224.0.0.5 ip-proto-89, length 60
2: 12:25:33.702691 192.168.103.91 > 224.0.0.5 ip-proto-89, length 60
7: 12:25:36.317000 192.168.206.100 > 224.0.0.5 ip-proto-89, length 56
8: 12:25:36.952587 fe80::2be:75ff:fef6:1d8e > ff02::5 ip-proto-89 40 [flowlabel 0xe] [hlim 1]
12: 12:25:41.282608 fe80::f6db:e6ff:fe33:442e > ff02::5 ip-proto-89 40 [flowlabel 0xe] [hlim 1]
```

Veja como o pacote multicast OSPFv2 é tratado pelo firewall:

<#root>

```
firepower#

show capture CAP packet-number 1 trace

115 packets captured


1: 12:25:33.142189 192.168.103.50 > 224.0.0.5 ip-proto-89, length 60

<-- The first packet of the flow
Phase: 1
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 6344 ns
Config:
Additional Information:
MAC Access list

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 6344 ns
Config:
Implicit Rule
```

```
Additional Information:
MAC Access list

Phase: 3
Type: ROUTE-LOOKUP
Subtype: No ECMP load balancing
Result: ALLOW
Elapsed time: 10736 ns
Config:
Additional Information:
Destination is locally connected. No ECMP load balancing.
Found next-hop 192.168.103.50 using egress ifc OUTSIDE(vrfid:0)

Phase: 4
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 5205 ns
Config:
Implicit Rule
Additional Information:

Phase: 5
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 5205 ns
Config:
Additional Information:

Phase: 6
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 5205 ns
Config:
Additional Information:

Phase: 7
Type: CLUSTER-REDIRECT
Subtype: cluster-redirect
Result: ALLOW
Elapsed time: 29280 ns
Config:
Additional Information:

Phase: 8
Type: MULTICAST
Subtype:
Result: ALLOW
Elapsed time: 976 ns
Config:
Additional Information:


Phase: 9



Type: OSPF

<-- The OSPF process

Subtype: ospf
```

Result: ALLOW

Elapsed time: 488 ns

Config:

Additional Information:

Phase: 10
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Elapsed time: 13176 ns
Config:
Additional Information:
New flow created with id 620, packet dispatched to next module

Result:
input-interface: OUTSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 82959 ns

É assim que o pacote multicast do OSPFv3 é tratado pelo firewall:

<#root>

firepower#

show capture CAP packet-number 8 trace

274 packets captured

8: 12:25:36.952587 fe80::2be:75ff:fef6:1d8e > ff02::5 ip-proto-89 40 [flowlabel 0xe] [hlim 1]

<-- The first packet of the flow
Phase: 1
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 7564 ns
Config:
Additional Information:
MAC Access list

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 7564 ns
Config:
Implicit Rule
Additional Information:
MAC Access list

Phase: 3
Type: ROUTE-LOOKUP
Subtype: No ECMP load balancing
Result: ALLOW
Elapsed time: 8296 ns
Config:
Additional Information:
Destination is locally connected. No ECMP load balancing.
Found next-hop ff02::5 using egress ifc identity(vrfid:0)

Phase: 4
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 8784 ns
Config:
Implicit Rule
Additional Information:

Phase: 5
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 8784 ns
Config:
Additional Information:

Phase: 6
Type: CLUSTER-REDIRECT
Subtype: cluster-redirect
Result: ALLOW
Elapsed time: 27816 ns
Config:
Additional Information:

Phase: 7

Type: OSPF

<-- The OSPF process

Subtype: ospf

Result: ALLOW

Elapsed time: 976 ns

Config:

Additional Information:

```
Phase: 8
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Elapsed time: 13664 ns
Config:
Additional Information:
New flow created with id 624, packet dispatched to next module

Result:
input-interface: OUTSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: NP Identity Ifc
Action: allow
Time Taken: 83448 ns
```

Tarefa 2 - Configurar multicast básico

Topologia

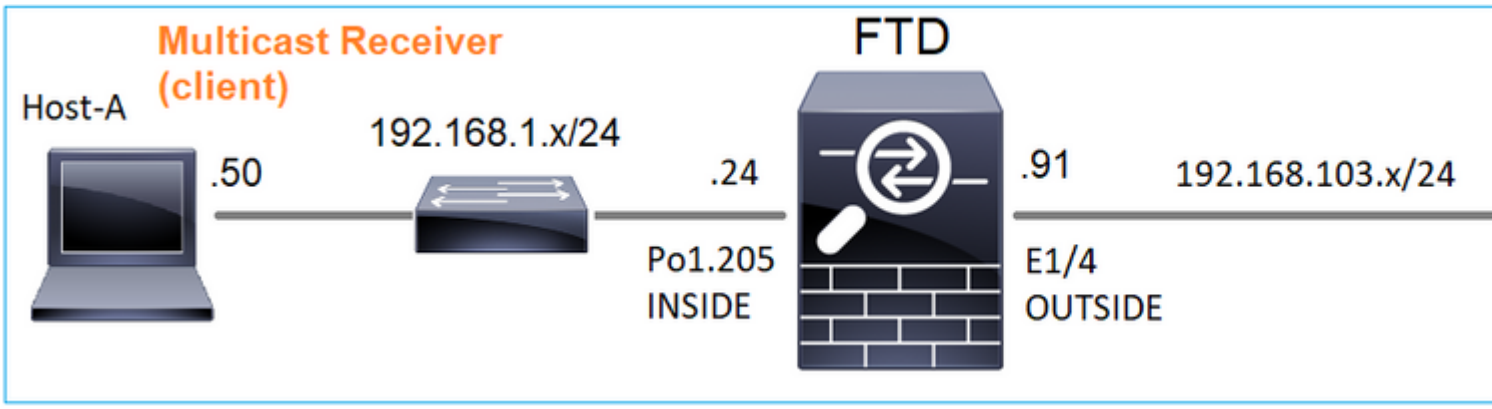

Requisitos

Configure o firewall de modo que o tráfego multicast do servidor seja transmitido para o cliente multicast no IP 230.10.10.10

Solução

Do ponto de vista do firewall, a configuração mínima é ativar o roteamento multicast globalmente. Isso ativa o IGMP e o PIM em segundo plano em todas as interfaces de firewall.

Na interface do usuário do FMC:

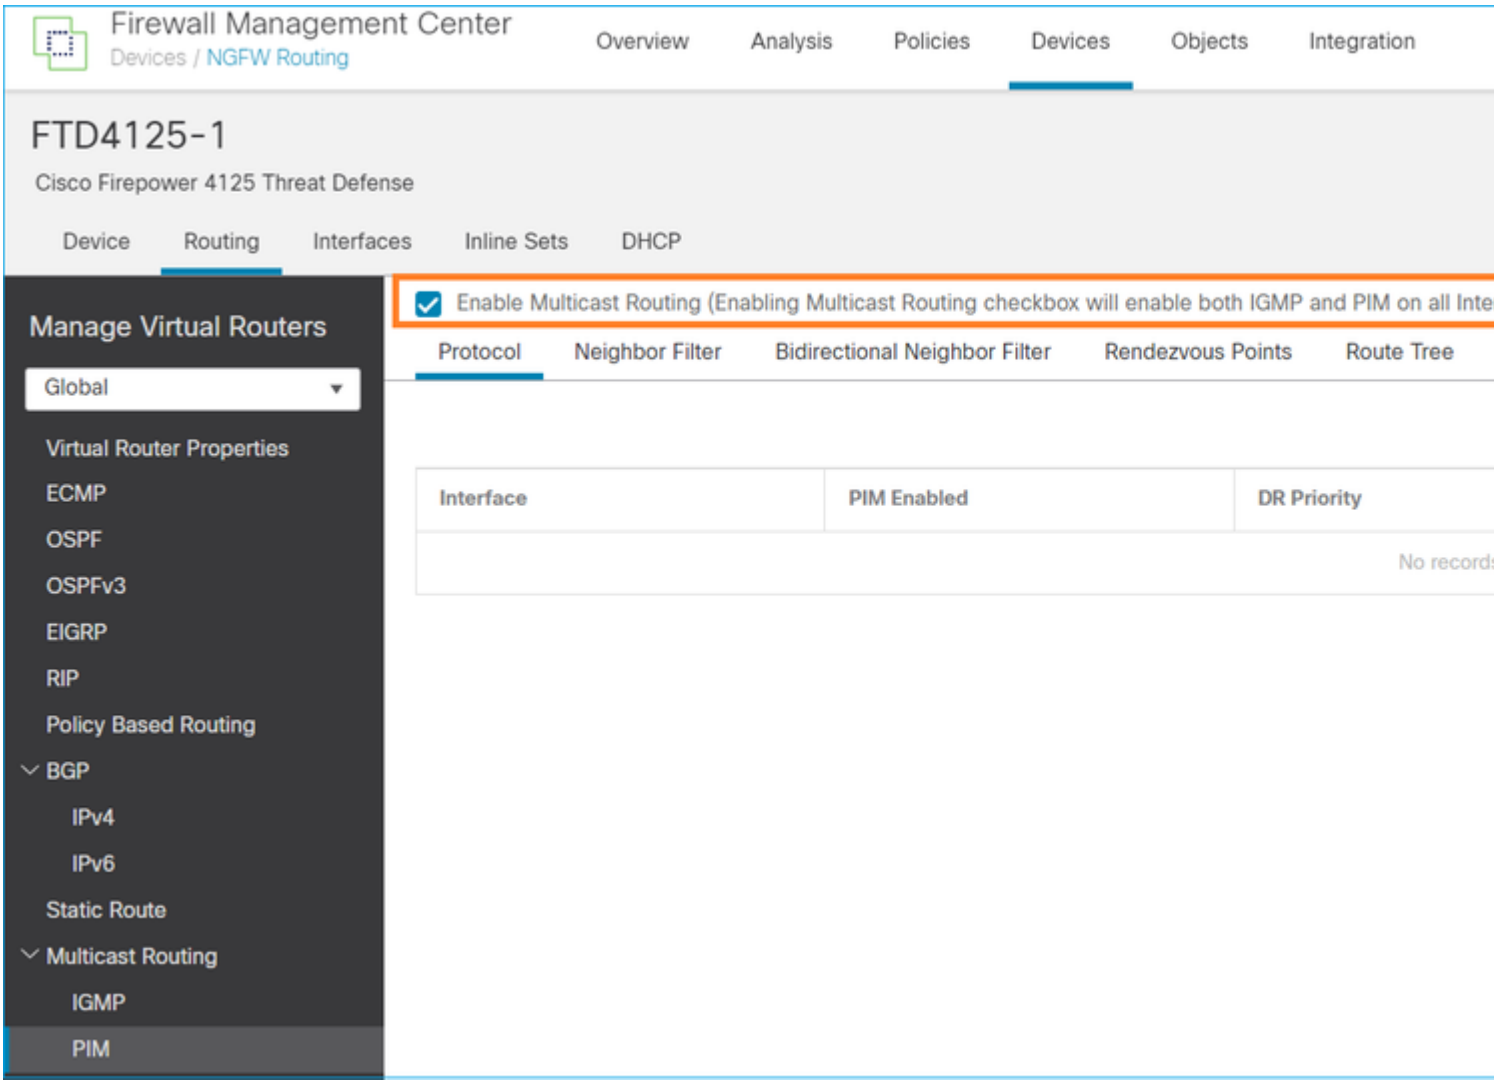

Na CLI do firewall, esta é a configuração enviada por push:

<#root>

```
firepower#
```

`show run multicast-routing`

`multicast-routing`

```
<-- Multicast routing is enabled
```

Verificação de IGMP

<#root>

```
firepower#
```

`show igmp interface`

```
diagnostic is up, line protocol is up
  Internet address is 0.0.0.0/0
  IGMP is disabled on interface
```

INSIDE is up, line protocol is up

<-- The interface is UP
  Internet address is 192.168.1.24/24

  IGMP is enabled on interface

<-- IGMP is enabled on the interface

  Current IGMP version is 2

<-- IGMP version
  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: 1
  Cumulative IGMP activity: 4 joins, 3 leaves
  IGMP querying router is 192.168.1.24 (this system)

OUTSIDE is up, line protocol is up

<-- The interface is UP
  Internet address is 192.168.103.91/24

  IGMP is enabled on interface

<-- IGMP is enabled on the interface

  Current IGMP version is 2

<-- IGMP version
  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: 1
  Cumulative IGMP activity: 1 joins, 0 leaves
  IGMP querying router is 192.168.103.91 (this system)

<#root>

firepower#

show igmp group

IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
239.255.255.250 INSIDE 00:09:05 00:03:19 192.168.1.50
239.255.255.250 OUTSIDE 00:06:01 00:02:33 192.168.103.60

<#root>

firepower#

show igmp traffic

IGMP Traffic Counters
Elapsed time since counters cleared: 03:40:48 Received Sent

```
                            Received     Sent
Valid IGMP Packets          21           207
Queries                     0            207
Reports                     15           0       <-- IGMP Reports received and sent
Leaves                      6            0
Mtrace packets              0            0
DVMRP packets               0            0
PIM packets                 0            0

Errors:
Malformed Packets           0
Martian source              0
Bad Checksums               0
```

Verificação de PIM

<#root>

firepower#

show pim interface

```
Address          Interface          PIM  Nbr   Hello  DR         DR
                                       Count Intvl  Prior

0.0.0.0          diagnostic         off  0     30     1          not elected
192.168.1.24     INSIDE             on   0     30     1          this system
192.168.103.91   OUTSIDE            on   0     30     1          this system
```

Verificação de MFIB

<#root>

firepower#

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

(*,224.0.1.39) Flags: S K
```

Forwarding: 0/0/0/0

, Other: 0/0/0 <-- The Forwarding counters are: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per secon

```
(*,224.0.1.40) Flags: S K
   Forwarding: 0/0/0/0,

Other: 8/8/0

   <-- The Other counters are: Total/RPF failed/Other drops
(*,232.0.0.0/8) Flags: K
   Forwarding: 0/0/0/0, Other: 0/0/0
```

Tráfego multicast através do firewall

Nesse caso, o aplicativo media player do VLC é usado como um servidor multicast e um cliente para testar o tráfego multicast:

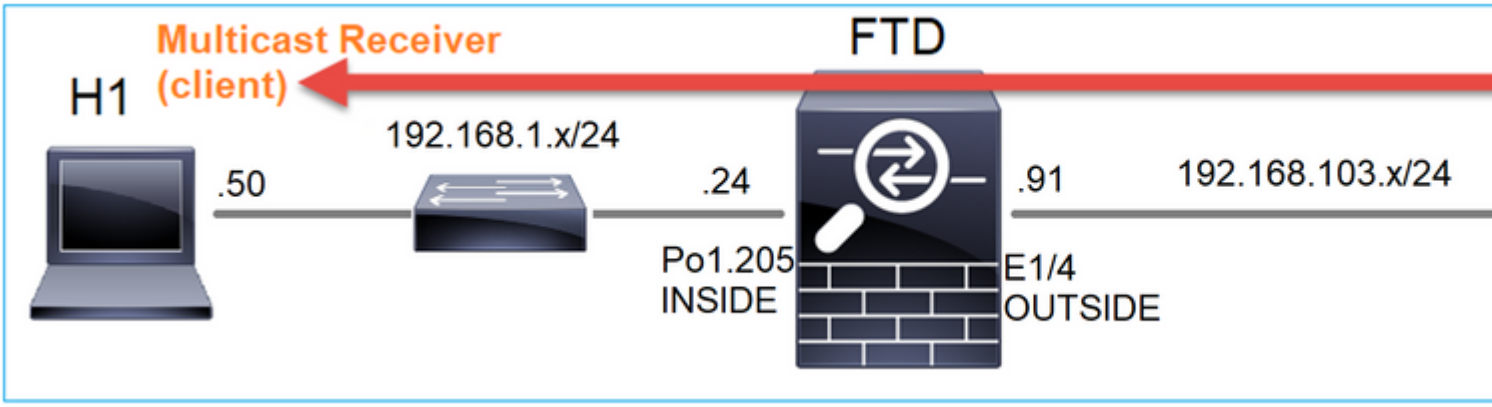

Configuração do servidor multicast VLC:

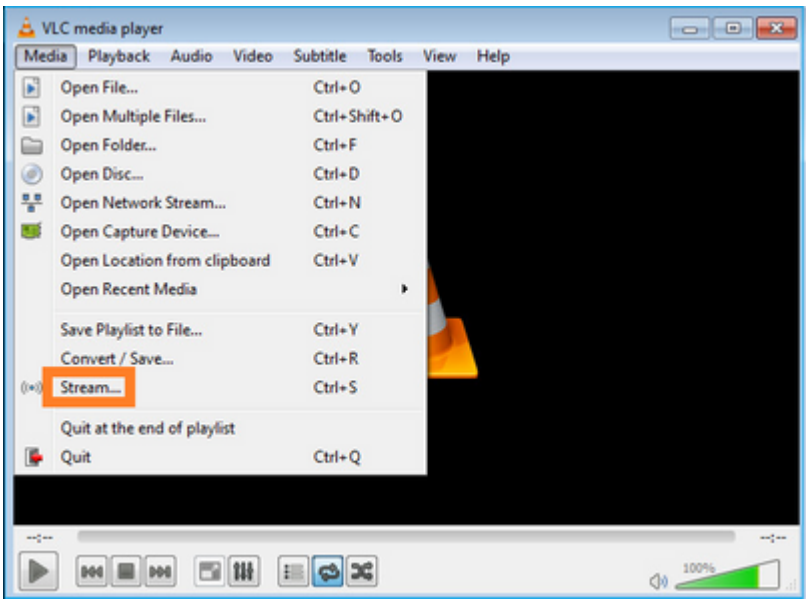

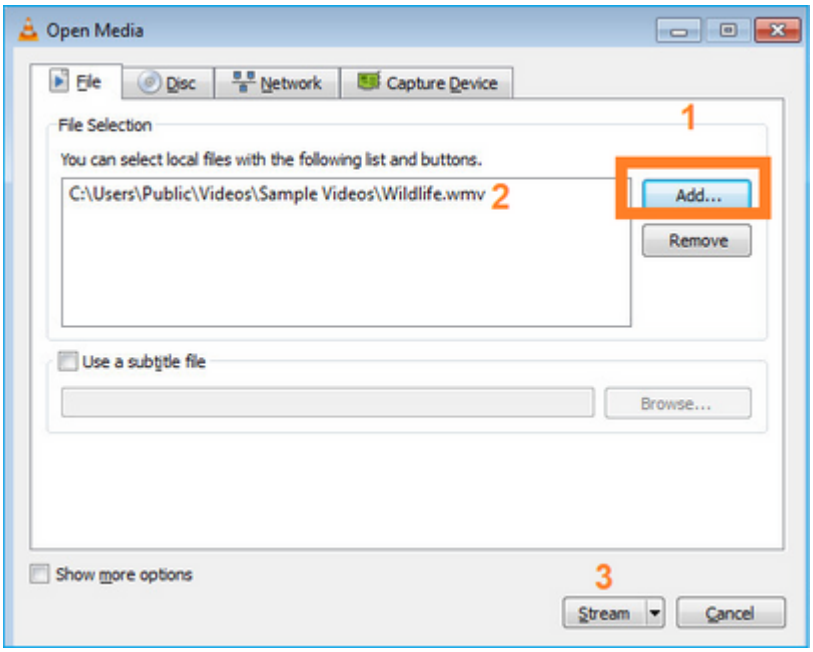

Na próxima tela, basta selecionar **Avançar.**

Selecione o formato:

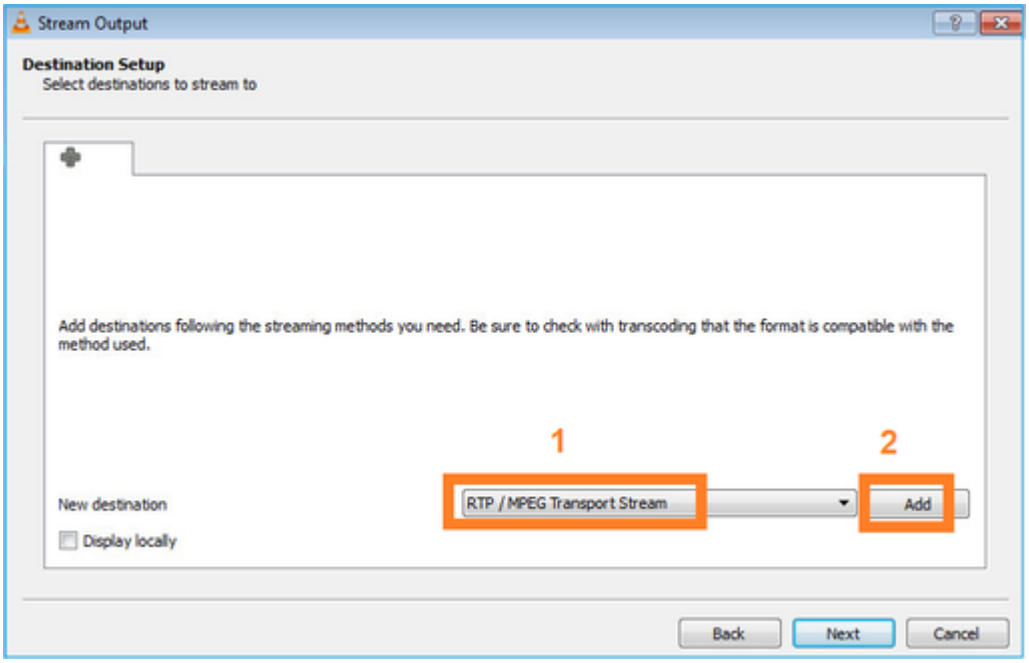

Especifique o IP e a porta multicast:

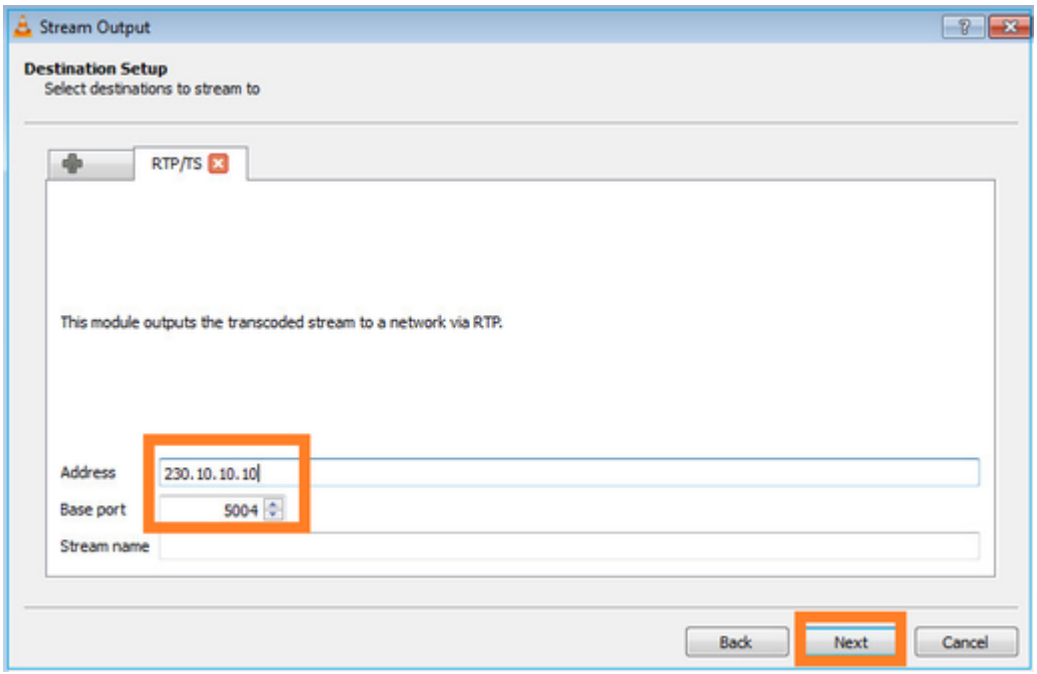

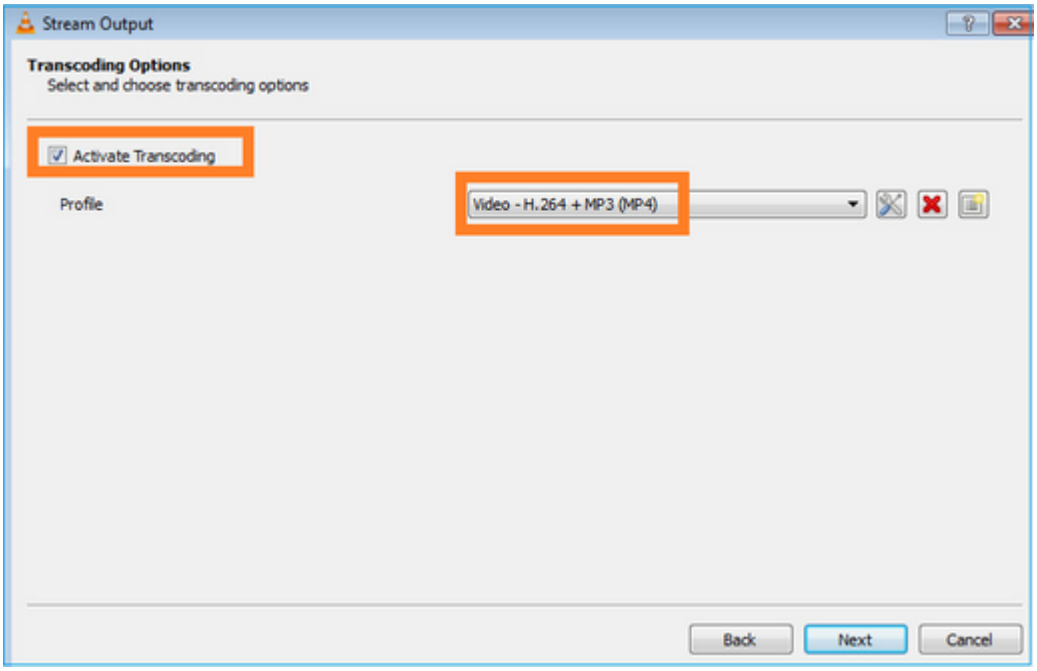

Ativar capturas LINA no firewall FTD:

<#root>

```
firepower#
capture INSIDE interface INSIDE match ip host 192.168.103.60 host 230.10.10.10
firepower#
capture OUTSIDE interface OUTSIDE trace match ip host 192.168.103.60 host 230.10.10.10
```

Selecione o botão **Stream** para o dispositivo iniciar o fluxo multicast:

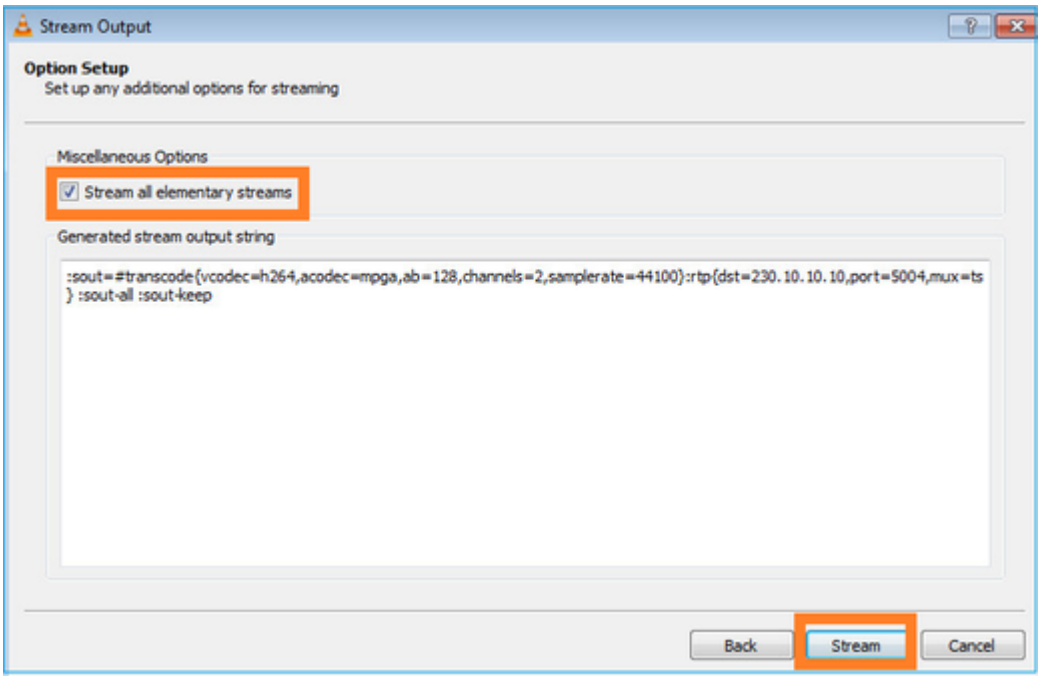

Ative a opção de â€˜loopâ€™ para que o fluxo seja enviado continuamente:

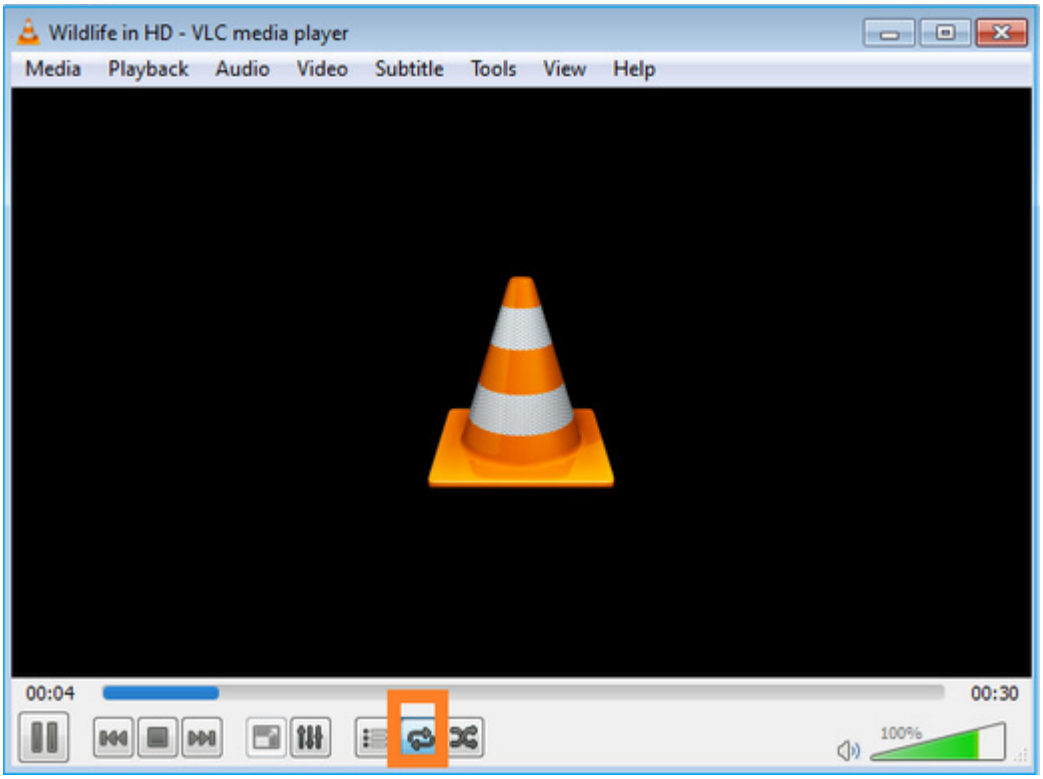

Verificação (cenário não operacional)

Este cenário é uma demonstração de um cenário não operacional. O objetivo é demonstrar o comportamento do firewall.

O dispositivo de firewall obtém o fluxo de multicast, mas não o encaminha:

<#root>

```
firepower#
```

```
show capture


capture INSIDE type raw-data interface INSIDE

[Capturing - 0 bytes]

<-- No packets sent or received
match ip host 192.168.103.60 host 230.10.10.10
capture OUTSIDE type raw-data trace interface OUTSIDE

[Buffer Full - 524030 bytes]

<-- The buffer is full
match ip host 192.168.103.60 host 230.10.10.10
```

As quedas LINA ASP de firewall mostram:

```
<#root>

firepower#

clear asp drop

firepower#

show asp drop


Frame drop:

Punt rate limit exceeded (punt-rate-limit)                          232

<-- The multicast packets were dropped
  Flow is denied by configured rule (acl-drop)                        2
  FP L2 rule drop (l2_acl)                                            2

Last clearing: 18:38:42 UTC Oct 12 2018 by enable_15

Flow drop:

Last clearing: 08:45:41 UTC May 17 2022 by enable_15
```

Para rastrear um pacote, é necessário capturar o primeiro pacote do fluxo de multicast. Por esse motivo, limpe os fluxos atuais:

```
<#root>

firepower#

clear capture OUTSIDE


firepower#

clear conn all addr 230.10.10.10

2 connection(s) deleted.
```

```
firepower#
show capture OUTSIDE


379 packets captured

1: 08:49:04.537875 192.168.103.60.54100 > 230.10.10.10.5005: udp 64
2: 08:49:04.537936 192.168.103.60.54099 > 230.10.10.10.5004: udp 1328
3: 08:49:04.538027 192.168.103.60.54099 > 230.10.10.10.5004: udp 1328
4: 08:49:04.538058 192.168.103.60.54099 > 230.10.10.10.5004: udp 1328
5: 08:49:04.538058 192.168.103.60.54099 > 230.10.10.10.5004: udp 1328
6: 08:49:04.538073 192.168.103.60.54099 > 230.10.10.10.5004: udp 1328
...
```

A opção â€˜detailâ€™ revela o endereço MAC multicast:

```
<#root>

firepower#
show capture OUTSIDE detail


379 packets captured

1: 08:49:04.537875 0050.569d.344a
0100.5e0a.0a0a
 0x0800 Length: 106
192.168.103.60.54100 > 230.10.10.10.5005: [udp sum ok] udp 64 (ttl 100, id 19759)
2: 08:49:04.537936 0050.569d.344a
0100.5e0a.0a0a
 0x0800 Length: 1370
192.168.103.60.54099 > 230.10.10.10.5004: [udp sum ok] udp 1328 (ttl 100, id 19760)
3: 08:49:04.538027 0050.569d.344a 0100.5e0a.0a0a 0x0800 Length: 1370
192.168.103.60.54099 > 230.10.10.10.5004: [udp sum ok] udp 1328 (ttl 100, id 19761)
...
```

O rastreamento de um pacote real mostra que o pacote é permitido, mas isso não é o que realmente acontece:

```
<#root>

firepower#
show capture OUTSIDE packet-number 1 trace


379 packets captured

1: 08:49:04.537875 192.168.103.60.54100 > 230.10.10.10.5005: udp 64
Phase: 1
```

```
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 11712 ns
Config:
Additional Information:
MAC Access list

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 11712 ns
Config:
Implicit Rule
Additional Information:
MAC Access list

Phase: 3
Type: ROUTE-LOOKUP
Subtype: No ECMP load balancing
Result: ALLOW
Elapsed time: 7808 ns
Config:
Additional Information:
Destination is locally connected. No ECMP load balancing.
Found next-hop 192.168.103.60 using egress ifc OUTSIDE(vrfid:0)

Phase: 4
Type: ACCESS-LIST
Subtype: log
Result: ALLOW
Elapsed time: 5246 ns
Config:
access-group CSM_FW_ACL_ global
access-list CSM_FW_ACL_ advanced permit ip any any rule-id 268434432
access-list CSM_FW_ACL_ remark rule-id 268434432: ACCESS POLICY: mzafeiro_empty - Default
access-list CSM_FW_ACL_ remark rule-id 268434432: L4 RULE: DEFAULT ACTION RULE
Additional Information:
This packet will be sent to snort for additional processing where a verdict will be reached

Phase: 5
Type: CONN-SETTINGS
Subtype:
Result: ALLOW
Elapsed time: 5246 ns
Config:
class-map class-default
match any
policy-map global_policy
class class-default
set connection advanced-options UM_STATIC_TCP_MAP
service-policy global_policy global
Additional Information:

Phase: 6
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 5246 ns
Config:
Additional Information:
```

```
Phase: 7
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 5246 ns
Config:
Additional Information:

Phase: 8
Type: CLUSTER-REDIRECT
Subtype: cluster-redirect
Result: ALLOW
Elapsed time: 31232 ns
Config:
Additional Information:

Phase: 9
```

Type: MULTICAST

```
<-- multicast process
Subtype:
Result: ALLOW
Elapsed time: 976 ns
Config:
Additional Information:

Phase: 10
```

Type: FLOW-CREATION

```
<-- the packet belongs to a new flow
Subtype:
Result: ALLOW
Elapsed time: 20496 ns
Config:
Additional Information:
New flow created with id 3705, packet dispatched to next module

Result:
input-interface: OUTSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE(vrfid:0)
output-status: up
output-line-status: up
```

Action: allow

```
<-- The packet is allowed
Time Taken: 104920 ns
```

Com base nos contadores mroute e mfib, os pacotes são descartados porque a Outgoing Interface List (OIL) está vazia:

<#root>

firepower#

show mroute

```
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


(192.168.103.60, 230.10.10.10), 00:01:33/00:01:56, flags: SPF



Incoming interface: OUTSIDE


RPF nbr: 192.168.103.60

Outgoing interface list: Null

<-- The OIL is empty!

(*, 239.255.255.250), 00:01:50/never, RP 0.0.0.0, flags: SCJ
Incoming interface: Null
RPF nbr: 0.0.0.0
Immediate Outgoing interface list:
INSIDE, Forward, 00:01:50/never
```

Os contadores MFIB mostram falhas de RPF que, neste caso, não é o que realmente acontece:

<#root>

```
firepower#

 show mfib 230.10.10.10

Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,
             AR - Activity Required, K - Keepalive
firepower# show mfib 230.10.10.10
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

<-- Multicast forwarding counters

Other counts: Total/RPF failed

/Other drops            <-- Multicast drop counters
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

(192.168.103.60,230.10.10.10) Flags: K


Forwarding: 0/0/0/0

,
```

Other: 650/650

/0   <-- Allowed and dropped multicast packets

Falhas de RPF semelhantes na saída 'show mfib count':

<#root>

firepower#

show mfib count

IP Multicast Statistics
8 routes, 4 groups, 0.25 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts:

Total/RPF failed

/Other drops(OIF-null, rate-limit etc)
Group: 224.0.1.39
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0
Group: 224.0.1.40
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0

Group: 230.10.10.10

  Source: 192.168.103.60,
    Forwarding: 0/0/0/0,

Other: 1115/1115

/0   <-- Allowed and dropped multicast packets
  Tot. shown: Source count: 1, pkt count: 0
Group: 232.0.0.0/8
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0
Group: 239.255.255.250
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0

Configure o receptor multicast VLC:

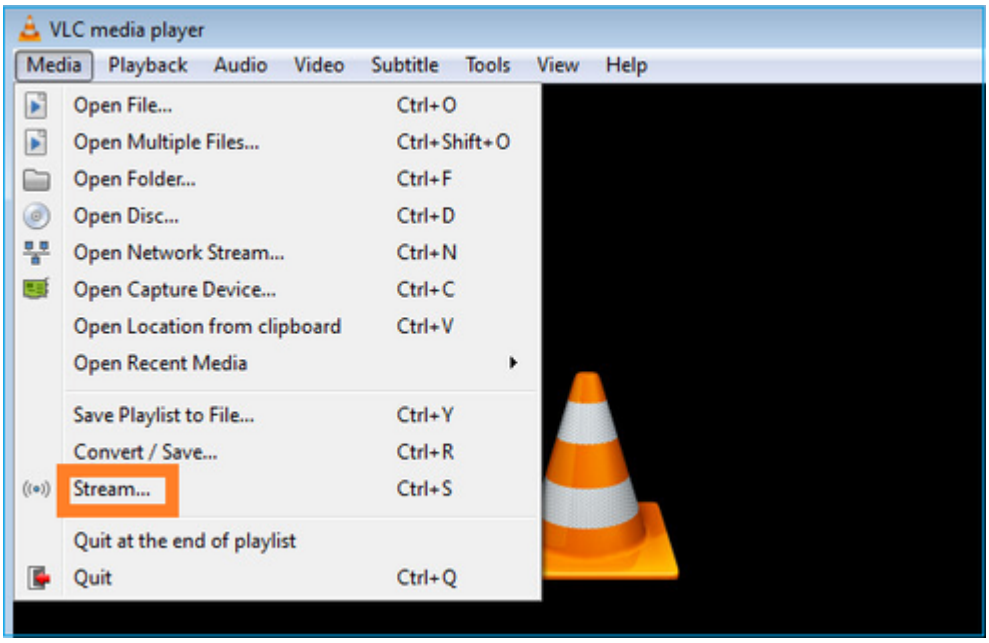

Especifique o IP de origem de multicast e selecione **Reproduzir:**

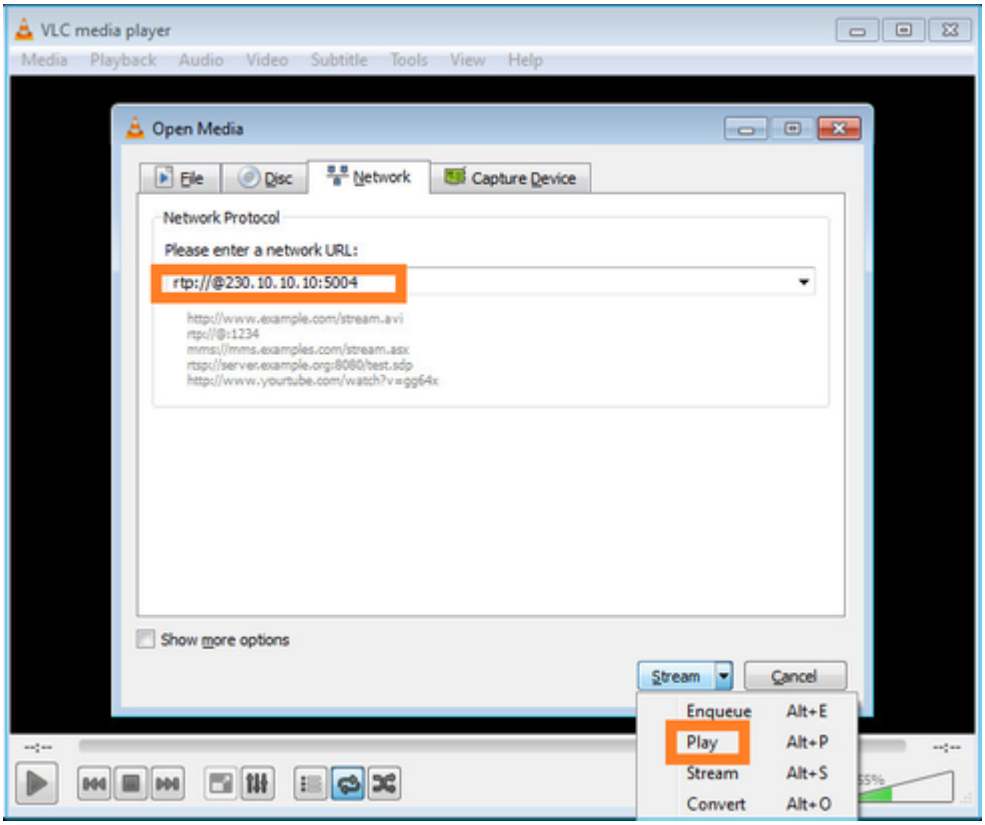

No back-end, assim que você seleciona **Play**, o host anuncia sua vontade de se juntar ao grupo multicast específico e envia uma mensagem **IGMP Report**:

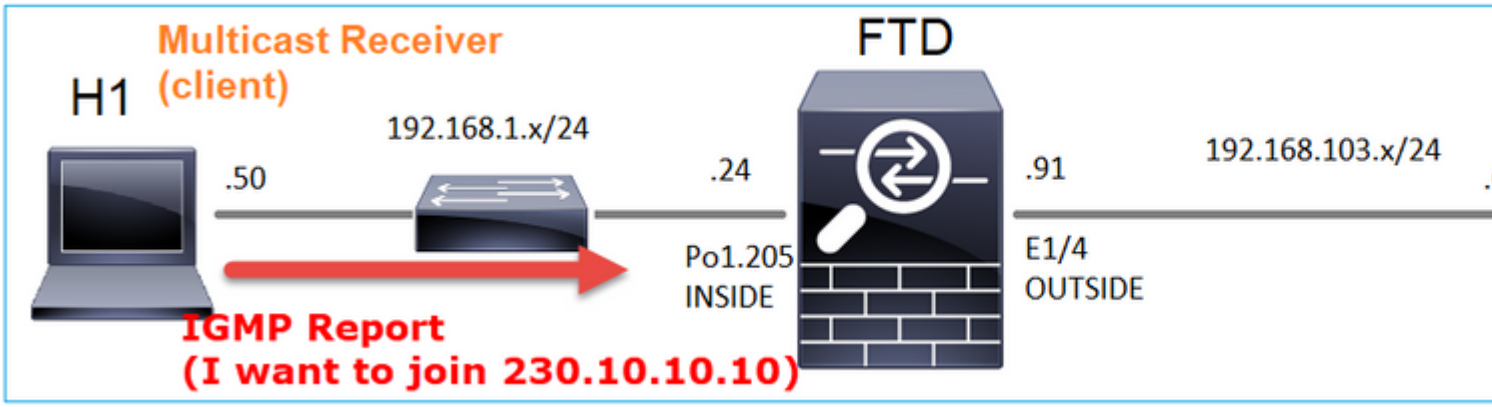

Se você habilitar uma depuração, poderá ver as mensagens de relatório IGMP:

<#root>

```
firepower#
```

debug igmp group 230.10.10.10

IGMP: Received v2 Report on INSIDE from 192.168.1.50 for 230.10.10.10

```
<-- IGMPv2 Report received
IGMP: group_db: add new group 230.10.10.10 on INSIDE
IGMP: MRIB updated (*,230.10.10.10) : Success
IGMP: Switching to EXCLUDE mode for 230.10.10.10 on INSIDE
IGMP: Updating EXCLUDE group timer for 230.10.10.10
```

O fluxo inicia:

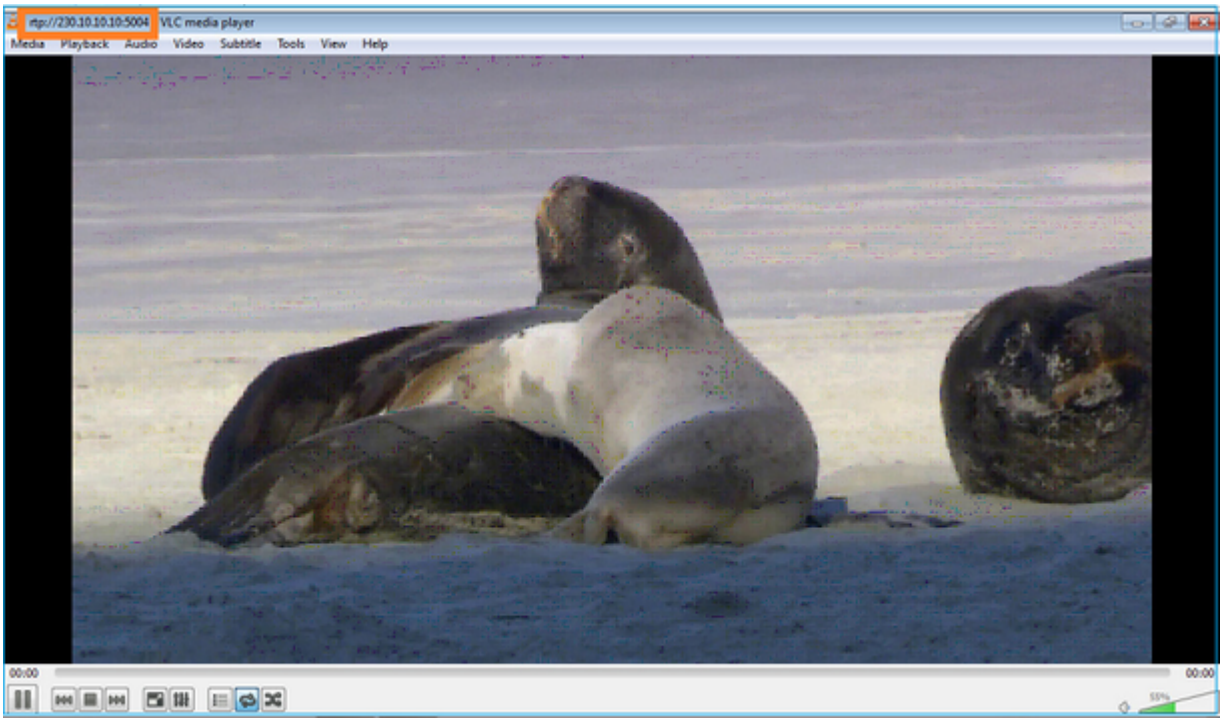

Verificação (cenário operacional)

<#root>

firepower#

show capture

capture INSIDE type raw-data interface INSIDE

[Buffer Full - 524156 bytes]

<-- Multicast packets on the egress interface
match ip host 192.168.103.60 host 230.10.10.10
capture OUTSIDE type raw-data trace interface OUTSIDE

[Buffer Full - 524030 bytes]

<-- Multicast packets on the ingress interface
match ip host 192.168.103.60 host 230.10.10.10

A tabela mroute do firewall:

<#root>

firepower#

show mroute

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

(*, 230.10.10.10), 00:00:34/never, RP 0.0.0.0, flags: SCJ
  Incoming interface: Null
  RPF nbr: 0.0.0.0
  Immediate Outgoing interface list:
    INSIDE, Forward, 00:00:34/never

(192.168.103.60, 230.10.10.10), 00:01:49/00:03:29, flags: SFJT

Incoming interface: OUTSIDE

RPF nbr: 192.168.103.60

Inherited Outgoing interface list:

INSIDE, Forward, 00:00:34/never

<-- The OIL shows an interface

<#root>

firepower#

show mfib 230.10.10.10

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

(*,230.10.10.10) Flags: C K
Forwarding: 0/0/0/0, Other: 0/0/0
INSIDE Flags: F NS
Pkts: 0/0

(192.168.103.60,230.10.10.10) Flags: K

Forwarding: 6373/0/1354/0,

Other: 548/548/0 <-- There are multicast packets forwarded

OUTSIDE Flags: A

INSIDE Flags: F NS

Pkts: 6373/6

contadores mfib:

<#root>

firepower#

show mfib count

IP Multicast Statistics
10 routes, 5 groups, 0.40 average sources per group

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second

Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)
Group: 224.0.1.39

```
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0
Group: 224.0.1.40
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0
```

Group: 230.10.10.10

RP-tree:

Forwarding: 0/0/0/0, Other: 0/0/0

Source: 192.168.103.60,

Forwarding: 7763/0/1354/0,

```
Other: 548/548/0   <-- There are multicast packets forwarded
  Tot. shown: Source count: 1, pkt count: 0
Group: 232.0.0.0/8
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0
Group: 239.255.255.250
  RP-tree:
    Forwarding: 0/0/0/0, Other: 0/0/0
  Source: 192.168.1.50,
    Forwarding: 7/0/500/0, Other: 0/0/0
  Tot. shown: Source count: 1, pkt count: 0
```

Espionagem de IGMP

- O Snooping IGMP é um mecanismo usado em switches para evitar inundação de multicast.
- O switch monitora os relatórios IGMP para determinar onde estão localizados os hosts (receptores).
- O switch monitora as Consultas IGMP para determinar onde os roteadores/firewalls (remetentes) estão localizados.
- O rastreamento de IGMP é ativado por padrão na maioria dos switches Cisco. Consulte os guias de comutação relacionados para obter mais detalhes. Este é um exemplo de saída de um switch Catalyst L3:

<#root>

switch#

show ip igmp snooping statistics

```
  Current number of Statistics entries    : 15
  Configured Statistics database limit    : 32000
  Configured Statistics database threshold: 25600
  Configured Statistics database limit    : Not exceeded
  Configured Statistics database threshold: Not exceeded
```

```
Snooping statistics for Vlan204
#channels: 3
#hosts   : 5

Source/Group                    Interface       Reporter        Uptime   Last-Join Last-Leave
0.0.0.0/230.10.10.10            Vl204:Gi1/48    192.168.1.50    2d13h        -     2d12h
0.0.0.0/230.10.10.10            Vl204:Gi1/48    192.168.1.97    2d13h    2d12h         -
0.0.0.0/230.10.10.10            Vl204:Gi2/1     192.168.1.50    2d10h    02:20:05  02:20:00
0.0.0.0/239.255.255.250         Vl204:Gi2/1     192.168.1.50    2d11h    02:20:05  02:20:00
0.0.0.0/239.255.255.250         Vl204:Gi2/1     192.168.2.50    2d14h    2d13h         -
0.0.0.0/239.255.255.250         Vl204:Gi2/1     192.168.6.50    2d13h        -     2d13h
0.0.0.0/224.0.1.40              Vl204:Gi2/26    192.168.2.1     2d14h    00:00:39  2d13h

Snooping statistics for Vlan206
#channels: 4
#hosts   : 3

Source/Group                    Interface       Reporter        Uptime   Last-Join Last-Leave
0.0.0.0/230.10.10.10            Vl206:Gi1/48    192.168.6.91    00:30:15 2d13h     2d13h
0.0.0.0/239.10.10.10            Vl206:Gi1/48    192.168.6.91    2d14h    2d13h         -
0.0.0.0/239.255.255.250         Vl206:Gi2/1     192.168.6.50    2d12h    00:52:49  00:52:45
0.0.0.0/224.0.1.40              Vl206:Gi2/26    192.168.6.1     00:20:10 2d13h     2d13h
0.0.0.0/230.10.10.10            Vl206:Gi2/26    192.168.6.1     2d13h    2d13h         -
0.0.0.0/230.10.10.10            Vl206:Gi2/26    192.168.6.91    2d13h        -     2d13h
0.0.0.0/239.10.10.10            Vl206:Gi2/26    192.168.6.1     2d14h    2d14h         -
0.0.0.0/239.10.10.10            Vl206:Gi2/26    192.168.6.91    2d14h        -     2d14h
```

Tarefa 3 - Grupo estático IGMP versus grupo de junção IGMP

Overview

	ip igmp static-group	**ip igmp join-group**
Aplicado na interface FTD?	Yes	Yes
O FTD atrai um fluxo multicast?	Sim, um PIM Join é enviado para o dispositivo upstream. a origem ou em direção ao ponto de encontro (RP). Isso ocorrerá somente se o FTD com esse comando for o Roteador Designado (DR) PIM nessa interface.	Sim, um PIM Join é enviado para o dispositivo upstream. a origem ou em direção ao ponto de encontro (RP). Isso ocorrerá somente se o FTD com esse comando for o Roteador Designado (DR) PIM nessa interface.
O FTD encaminha o tráfego multicast para fora da interface?	Yes	Yes
O FTD consome e responde ao tráfego multicast?	No	Sim, o FTD envia o fluxo de multicast para a CPU, consome-o e responde à origem.
impacto de CPU	Mínimo, pois o pacote não é enviado para a CPU.	Pode afetar a CPU do FTD, já que cada pacote multicast que pertence ao grupo é enviado para a CPU do FTD.

Requisito da tarefa

Considere esta topologia:

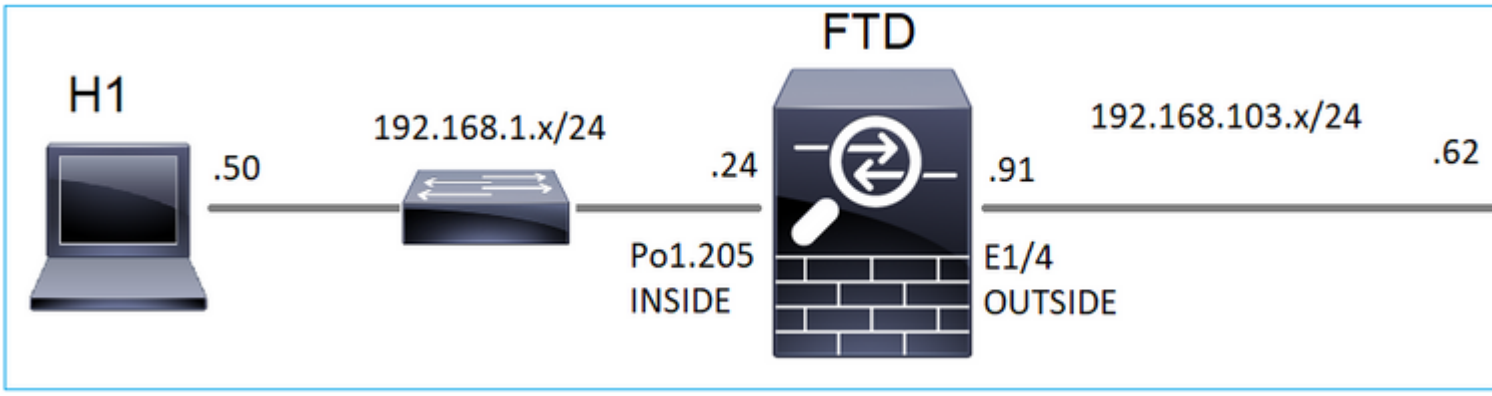

No firewall, habilite estas capturas:

<#root>

```
firepower#
 capture CAPI interface OUTSIDE trace match icmp host 192.168.103.62 any
firepower#
 capture CAPO interface INSIDE match icmp host 192.168.103.62 any
```

1. Use o ping ICMP do switch L3 para enviar tráfego multicast para o IP 230.11.11.11 e verifique como isso é tratado pelo firewall.
2. Ative o comando **igmp static-group** na interface INSIDE do firewall e verifique como o fluxo multicast (IP 230.11.11.11) é tratado pelo firewall.
3. Ative o comando **igmp static-group** na interface INSIDE do firewall e verifique como o fluxo multicast (IP 230.11.11.11) é tratado pelo firewall.

Solução

O firewall não tem nenhuma rota mpara o IP 230.11.11.11:

<#root>

```
firepower#
show mroute

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

(*, 239.255.255.250), 00:43:21/never, RP 0.0.0.0, flags: SCJ
  Incoming interface: Null
```

```
  RPF nbr: 0.0.0.0
  Immediate Outgoing interface list:
    OUTSIDE, Forward, 00:05:41/never
    INSIDE, Forward, 00:43:21/never
```

Uma maneira simples de testar o multicast é usar a ferramenta de ping ICMP. Nesse caso, inicie um ping do R2 para o endereço IP multicast 230.11.11.11:

<#root>

```
L3-Switch#
```

ping 230.11.11.11 re 100

```
Type escape sequence to abort.
Sending 100, 100-byte ICMP Echos to 230.11.11.11, timeout is 2 seconds:
..............................
```

No firewall, um mroute é criado dinamicamente e o OIL está vazio:

<#root>

```
firepower#
```

show mroute

```
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
```

(192.168.103.62, 230.11.11.11), 00:02:33/00:00:56, flags: SPF

```
<-- The mroute is added
```

** Incoming interface: OUTSIDE**

** RPF nbr: 192.168.103.62**

Outgoing interface list: Null

```
<-- The OIL is empty
```

A captura no firewall mostra:

<#root>

```
firepower# show capture

capture CAPI type raw-data trace interface OUTSIDE

[Capturing - 1040 bytes]

<-- There are ICMP packets captured on ingress interface
match icmp host 192.168.103.62 any
capture CAPO type raw-data interface INSIDE

[Capturing - 0 bytes]

<-- There are no ICMP packets on egress
match icmp host 192.168.103.62 any
```

O firewall cria conexões para cada ping, mas descarta os pacotes silenciosamente:

<#root>

```
firepower#

show log | include 230.11.11.11


May 17 2022 11:05:47: %FTD-7-609001:

Built local-host identity:230.11.11.11

<-- A new connection is created
May 17 2022 11:05:47: %FTD-6-302020: Built inbound ICMP connection for faddr 192.168.1.99/6 gaddr 230.11
May 17 2022 11:05:47: %FTD-6-302020: Built inbound ICMP connection for faddr 192.168.103.62/6 gaddr 230.
May 17 2022 11:05:49: %FTD-6-302021: Teardown ICMP connection for faddr 192.168.1.99/6 gaddr 230.11.11.1
May 17 2022 11:05:49: %FTD-6-302021: Teardown ICMP connection for faddr 192.168.103.62/6 gaddr 230.11.11
May 17 2022 11:05:49: %FTD-7-609002:

Teardown local-host identity:230.11.11.11 duration 0:00:02

<-- The connection is closed
May 17 2022 11:05:51: %FTD-7-609001:

Built local-host identity:230.11.11.11

<

--

A new connection is created
May 17 2022 11:05:51: %FTD-6-302020: Built inbound ICMP connection for faddr 192.168.1.99/6 gaddr 230.11
May 17 2022 11:05:51: %FTD-6-302020: Built inbound ICMP connection for faddr 192.168.103.62/6 gaddr 230.
May 17 2022 11:05:53: %FTD-6-302021: Teardown ICMP connection for faddr 192.168.1.99/6 gaddr 230.11.11.1
May 17 2022 11:05:53: %FTD-6-302021: Teardown ICMP connection for faddr 192.168.103.62/6 gaddr 230.11.11
May 17 2022 11:05:53: %FTD-7-609002:

Teardown local-host identity:230.11.11.11 duration 0:00:02

<-- The connection is closed
```

Observação: a captura de queda LINA ASP não mostra os pacotes descartados

---

A principal indicação de quedas de pacotes multicast é:

<#root>

```
firepower#
```

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

(*,224.0.1.39) Flags: S K
  Forwarding: 0/0/0/0, Other: 0/0/0
(*,224.0.1.40) Flags: S K
  Forwarding: 0/0/0/0, Other: 0/0/0
```

(192.168.103.62,230.11.11.11)

```
 Flags: K         <-- The multicast stream
  Forwarding: 0/0/0/0,
```

Other: 27/27/0

```
<-- The packets are dropped
```

igmp static-group

No FMC, configure um grupo IGMP estático:

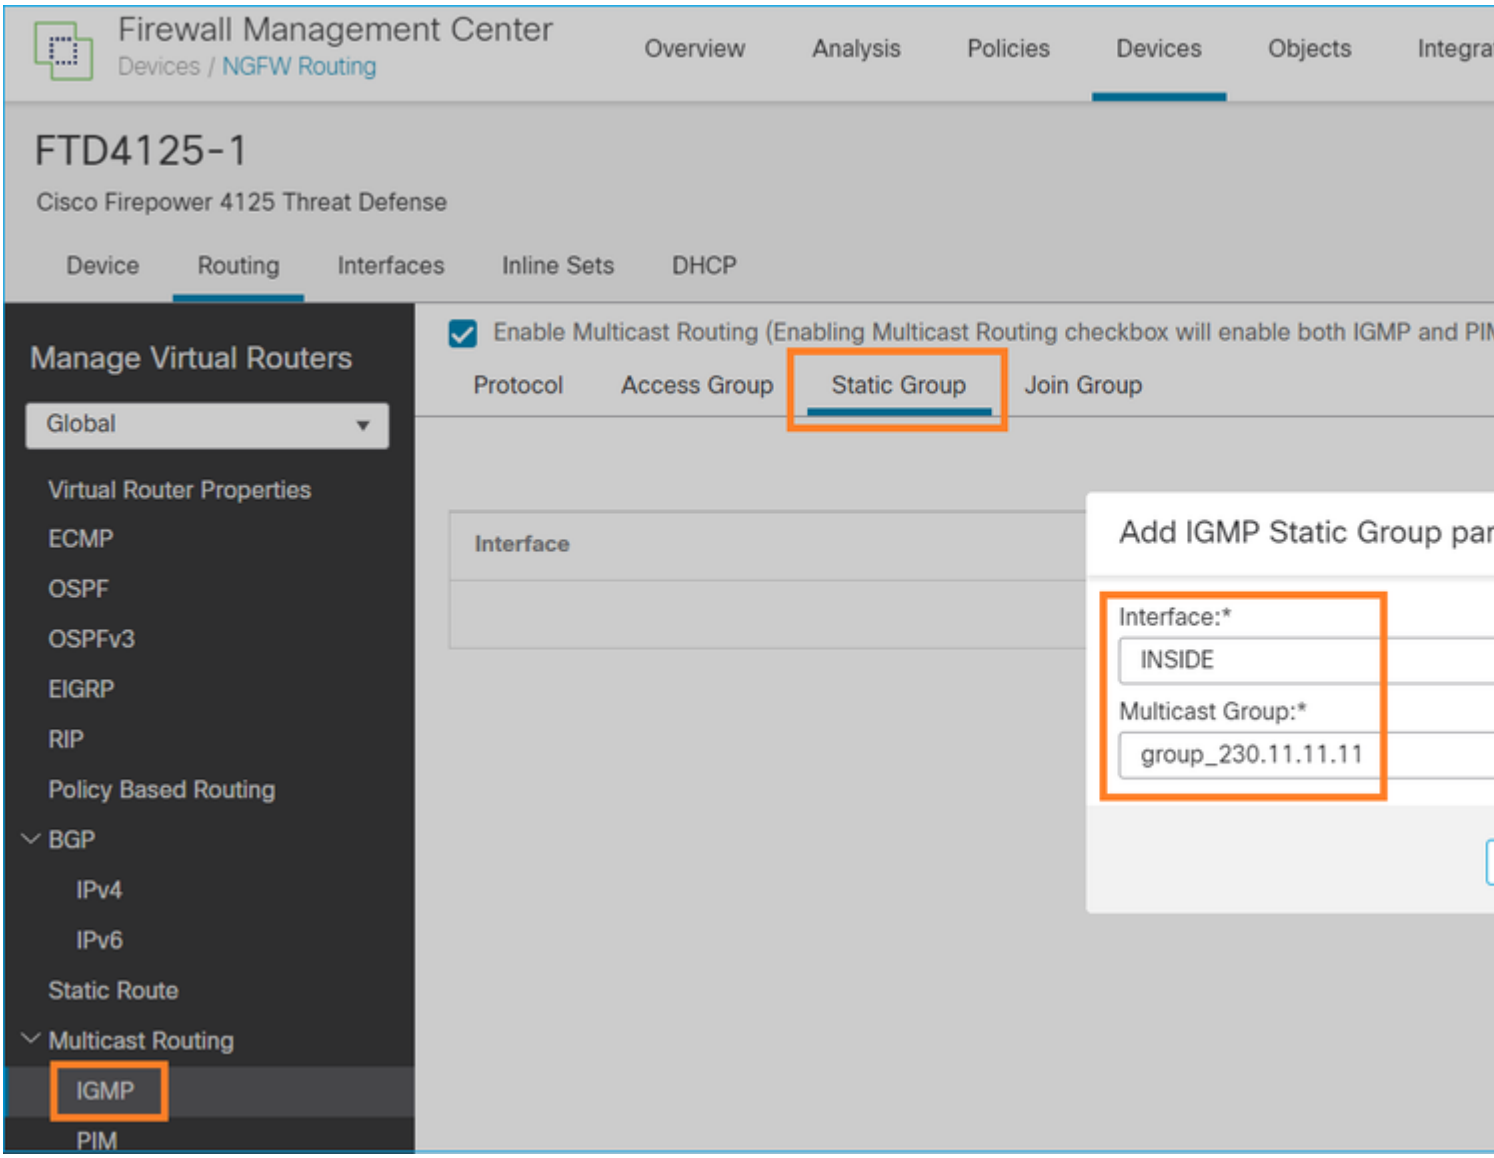

Isso é o que é implantado em segundo plano:

<#root>

```
interface Port-channel1.205
 vlan 205
 nameif INSIDE
 cts manual
  propagate sgt preserve-untag
  policy static sgt disabled trusted
 security-level 0
 ip address 192.168.1.24 255.255.255.0

 igmp static-group 230.11.11.11

<-- IGMP static group is enabled on the interface
```

O ping falha, mas o tráfego multicast ICMP é encaminhado agora através do firewall:

<#root>

```
L3-Switch#
ping 230.11.11.11 re 10000
Type escape sequence to abort.
Sending 10000, 100-byte ICMP Echos to 230.11.11.11, timeout is 2 seconds:
.........................
```

<#root>

```
firepower#
show capture

capture CAPI type raw-data trace interface OUTSIDE
[Capturing - 650 bytes]
<-- ICMP packets are captured on ingress interface
match icmp host 192.168.103.62 any
capture CAPO type raw-data interface INSIDE
[Capturing - 670 bytes]
<-- ICMP packets are captured on egress interface
match icmp host 192.168.103.62 any
```

<#root>

```
firepower#
show capture CAPI
8 packets captured
1: 11:31:32.470541 192.168.103.62 > 230.11.11.11 icmp: echo request
2: 11:31:34.470358 192.168.103.62 > 230.11.11.11 icmp: echo request
3: 11:31:36.470831 192.168.103.62 > 230.11.11.11 icmp: echo request
4: 11:31:38.470785 192.168.103.62 > 230.11.11.11 icmp: echo request
...
firepower#
show capture CAPO

11 packets captured
1: 11:31:32.470587 802.1Q vlan#205 P0 192.168.103.62 > 230.11.11.11 icmp: echo request
2: 11:31:34.470404 802.1Q vlan#205 P0 192.168.103.62 > 230.11.11.11 icmp: echo request
3: 11:31:36.470861 802.1Q vlan#205 P0 192.168.103.62 > 230.11.11.11 icmp: echo request
4: 11:31:38.470816 802.1Q vlan#205 P0 192.168.103.62 > 230.11.11.11 icmp: echo request
```

Observação: o rastreamento do pacote mostra uma saída incorreta (a interface de entrada é igual à de saída). Para obter mais detalhes, verifique a ID de bug da Cisco CSCvm89673.

<#root>

```
firepower#
```

show capture CAPI packet-number 1 trace

1: 11:39:33.553987 192.168.103.62 > 230.11.11.11 icmp: echo request

```
Phase: 1
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 3172 ns
Config:
Additional Information:
MAC Access list

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 3172 ns
Config:
Implicit Rule
Additional Information:
MAC Access list

Phase: 3
Type: ROUTE-LOOKUP
Subtype: No ECMP load balancing
Result: ALLOW
Elapsed time: 9760 ns
Config:
Additional Information:
Destination is locally connected. No ECMP load balancing.
Found next-hop 192.168.103.62 using egress ifc OUTSIDE(vrfid:0)

Phase: 4
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 5368 ns
Config:
Implicit Rule
Additional Information:

Phase: 5
Type: CONN-SETTINGS
Subtype:
Result: ALLOW
Elapsed time: 5368 ns
Config:
class-map class-default
match any
policy-map global_policy
class class-default
set connection advanced-options UM_STATIC_TCP_MAP
service-policy global_policy global
Additional Information:

Phase: 6
Type: NAT
```

Subtype: per-session
Result: ALLOW
Elapsed time: 5368 ns
Config:
Additional Information:

Phase: 7
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 5368 ns
Config:
Additional Information:

Phase: 8
Type: CLUSTER-REDIRECT
Subtype: cluster-redirect
Result: ALLOW
Elapsed time: 31720 ns
Config:
Additional Information:

Phase: 9
Type: INSPECT
Subtype: np-inspect
Result: ALLOW
Elapsed time: 488 ns
Config:
class-map inspection_default
match default-inspection-traffic
policy-map global_policy
class inspection_default
inspect icmp
service-policy global_policy global
Additional Information:

Phase: 10
Type: INSPECT
Subtype: np-inspect
Result: ALLOW
Elapsed time: 2440 ns
Config:
Additional Information:

Phase: 11

Type: MULTICAST

<-- The packet is multicast

Subtype:

Result: ALLOW

Elapsed time: 976 ns

Config:

Additional Information:

Phase: 12

Type: FLOW-CREATION

<-- A new flow is created
Subtype:
Result: ALLOW
Elapsed time: 56120 ns
Config:
Additional Information:
New flow created with id 5690, packet dispatched to next module

Phase: 13
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 10248 ns
Config:
Additional Information:
MAC Access list

Result:

input-interface: OUTSIDE(vrfid:0)

input-status: up
input-line-status: up

output-interface: OUTSIDE(vrfid:0)

output-status: up
output-line-status: up

Action: allow

<-- The packet is allowed
Time Taken: 139568 ns

Dica: você pode fazer ping com timeout 0 a partir do host de origem e pode verificar os contadores mfib do firewall:

<#root>

L3-Switch#

ping 230.11.11.11 re 500 timeout 0

Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 230.11.11.11, timeout is 0 seconds:
.....
.....
.....

```
.....................
```

<#root>

```
firepower# clear mfib counters

firepower# !ping from the source host.

firepower#

show mfib 230.11.11.11


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

(*,230.11.11.11) Flags: C K
  Forwarding: 0/0/0/0, Other: 0/0/0
  INSIDE Flags: F NS
    Pkts: 0/0
(192.168.103.62,230.11.11.11) Flags: K


Forwarding: 500/0/100/0, Other: 0/0/0

<-- 500 multicast packets forwarded. The average size of each packet is 100 Bytes
  OUTSIDE Flags: A
  INSIDE Flags: F NS
    Pkts: 500/0
```

igmp join-group

No FMC remote, configure a configuração de grupo estático previamente configurada e configure um grupo de união IGMP:

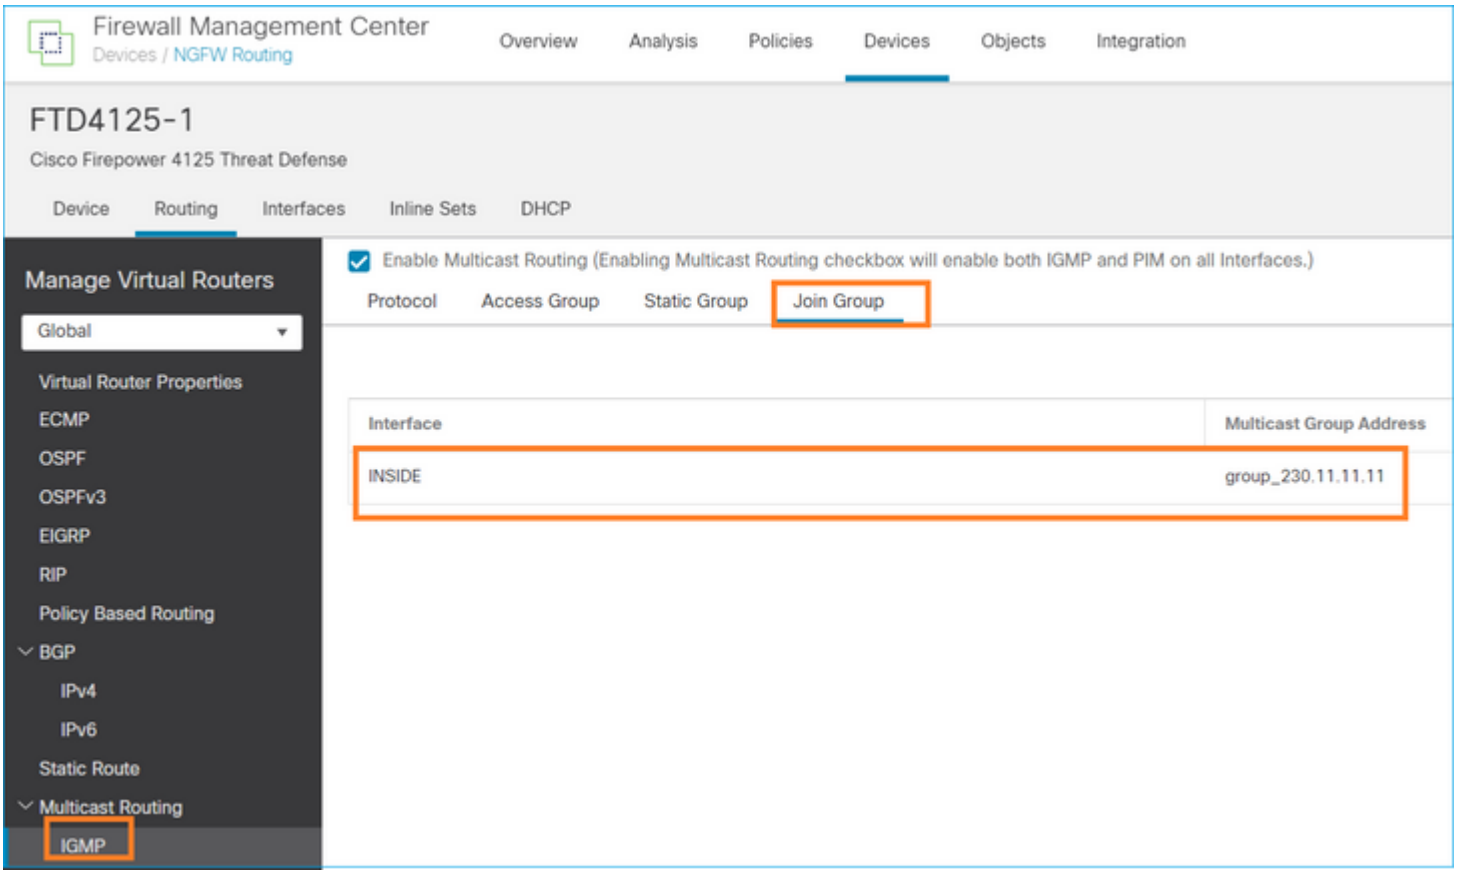

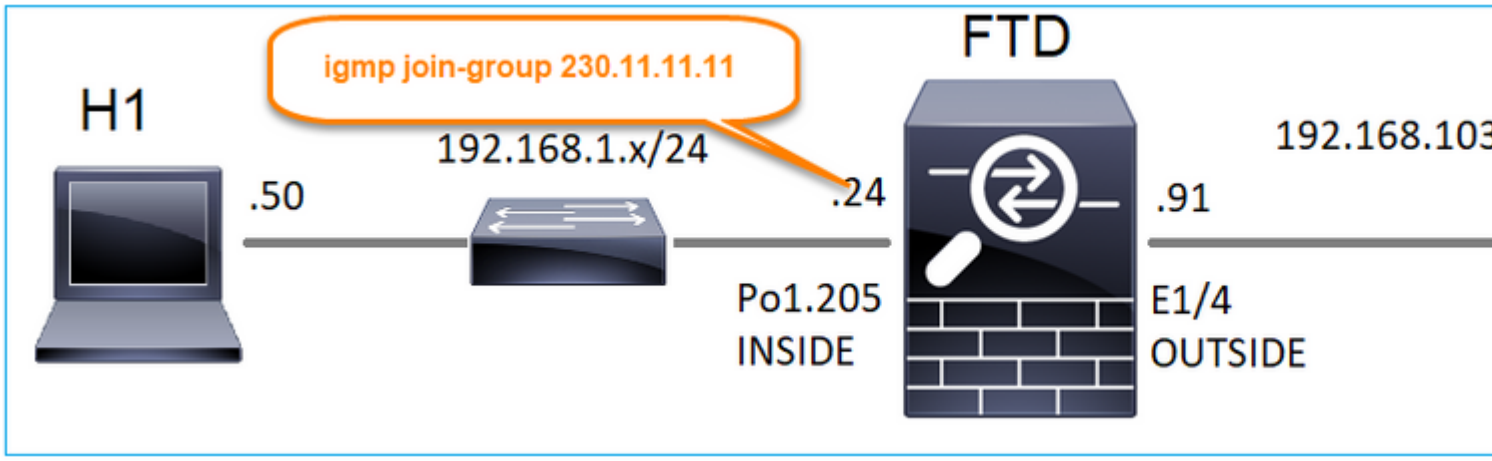

A configuração implantada:

<#root>

firepower#

show run interface Port-channel1.205

```
!
interface Port-channel1.205
vlan 205
nameif INSIDE
cts manual
propagate sgt preserve-untag
policy static sgt disabled trusted
security-level 0
```

```
ip address 192.168.1.24 255.255.255.0
```

```
igmp join-group 230.11.11.11
```

```
<-- The interface joined the multicast group
```

O grupo IGMP:

<#root>

```
firepower#
```

```
show igmp group
```

```
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
```

```
230.11.11.11 INSIDE 00:30:43 never 192.168.1.24
```

```
<-- The group is enabled on the interface
```

A partir do host de origem, tente o primeiro teste multicast ICMP em direção ao IP 230.11.11.11:

<#root>

```
L3-Switch#
```

```
ping 230.11.11.11 repeat 10
```

```
Type escape sequence to abort.
Sending 10, 100-byte ICMP Echos to 230.11.11.11, timeout is 2 seconds:

Reply to request 0 from 192.168.1.24, 12 ms
Reply to request 1 from 192.168.1.24, 8 ms
Reply to request 2 from 192.168.1.24, 8 ms
Reply to request 3 from 192.168.1.24, 8 ms
Reply to request 4 from 192.168.1.24, 8 ms
Reply to request 5 from 192.168.1.24, 12 ms
Reply to request 6 from 192.168.1.24, 8 ms
Reply to request 7 from 192.168.1.24, 8 ms
Reply to request 8 from 192.168.1.24, 8 ms
Reply to request 9 from 192.168.1.24, 8 ms
```

Observação: se você não vir todas as respostas, verifique a ID de bug da Cisco CSCvm90069.

Tarefa 4 - Configurar o roteamento multicast stub IGMP

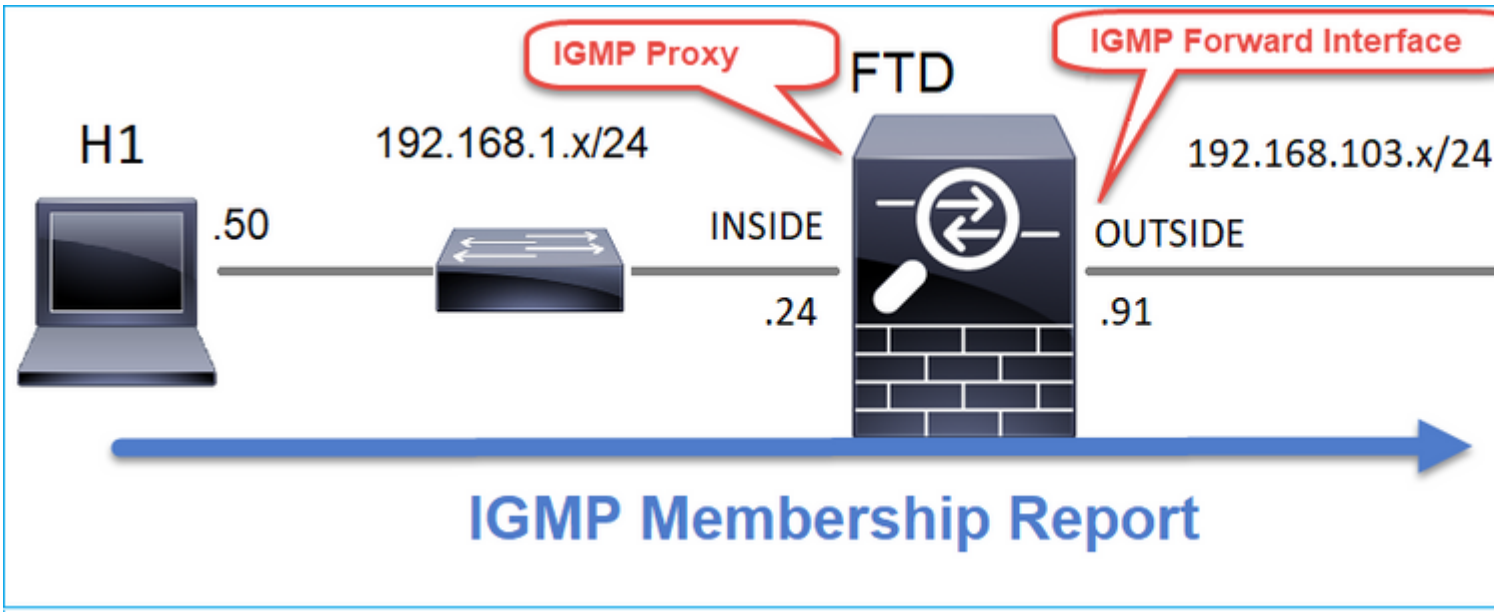

Configure o roteamento multicast stub no FTD para que as mensagens do Relatório de Associação IGMP recebidas na interface INSIDE sejam encaminhadas para a interface EXTERNA.

Solução

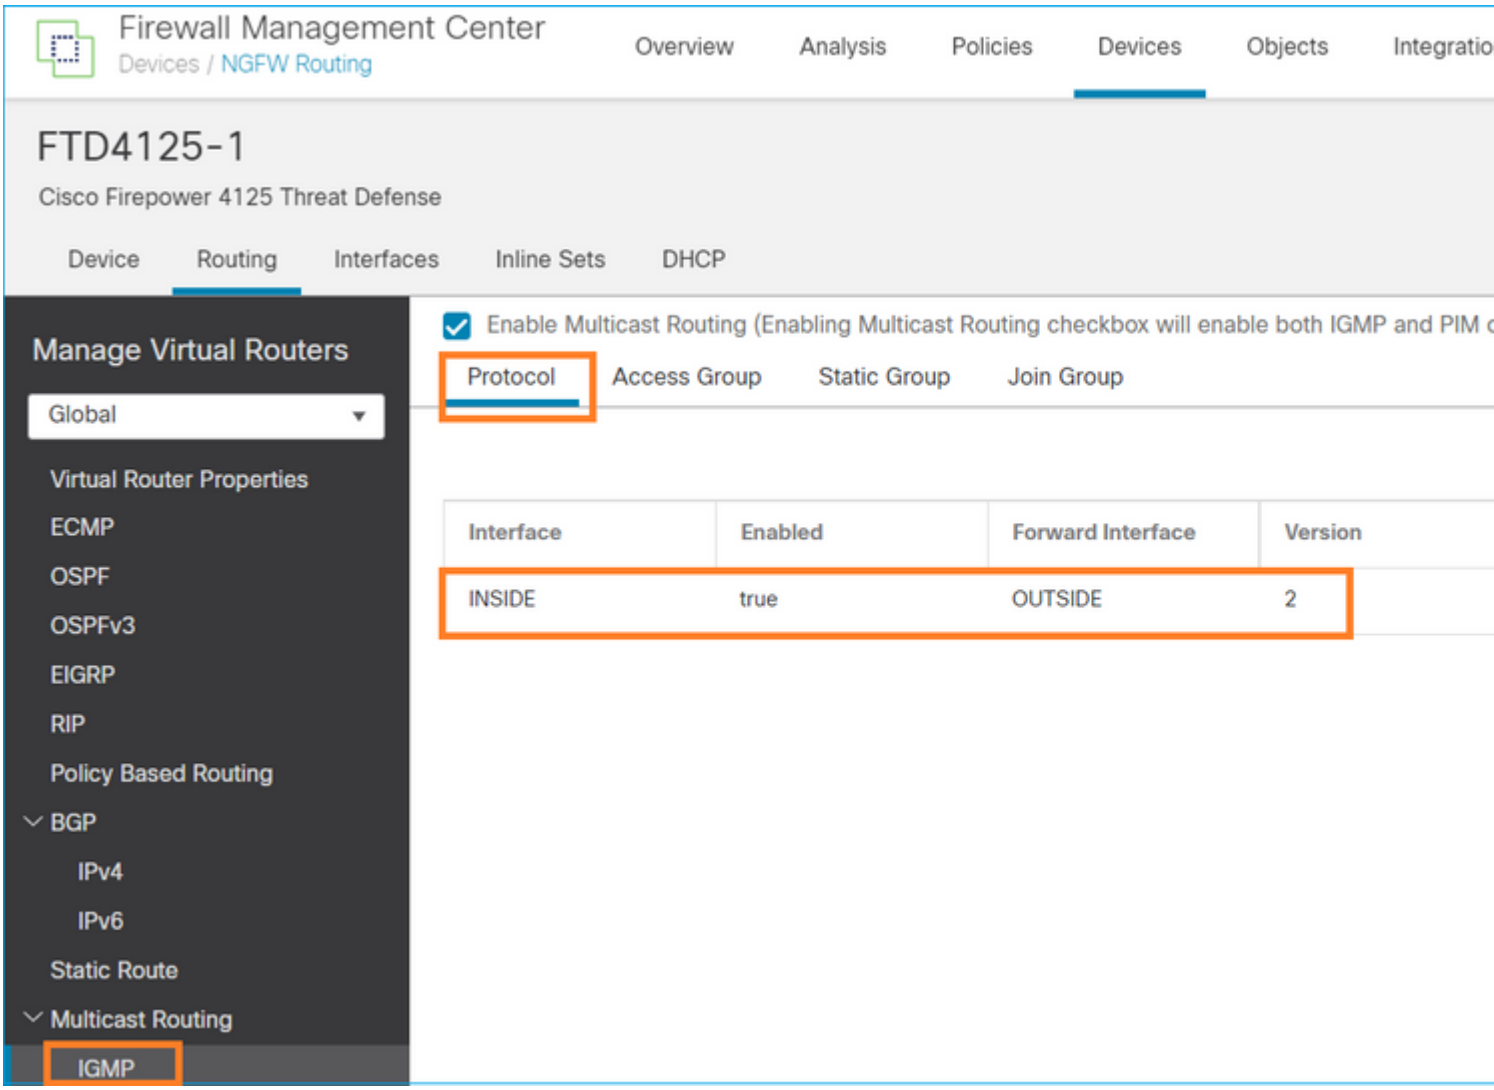

A configuração implantada:

<#root>

```
firepower#
```

show run multicast-routing

multicast-routing

```
<-- Multicast routing is enabled
firepower#
```

show run interface Port-channel1.205

```
!
interface Port-channel1.205
 vlan 205
 nameif INSIDE
 cts manual
  propagate sgt preserve-untag
  policy static sgt disabled trusted
 security-level 0
 ip address 192.168.1.24 255.255.255.0
```

igmp forward interface OUTSIDE

```
<-- The interface does stub multicast routing
```

Verificação

Habilitar capturas no FTD:

<#root>

```
firepower#
```

capture CAPI interface INSIDE trace match igmp any host 230.10.10.10

```
firepower#
```

capture CAPO interface OUTSIDE match igmp any host 230.10.10.10

Verificação

Para forçar um Relatório de Associação IGMP, você pode usar um aplicativo como o VLC:

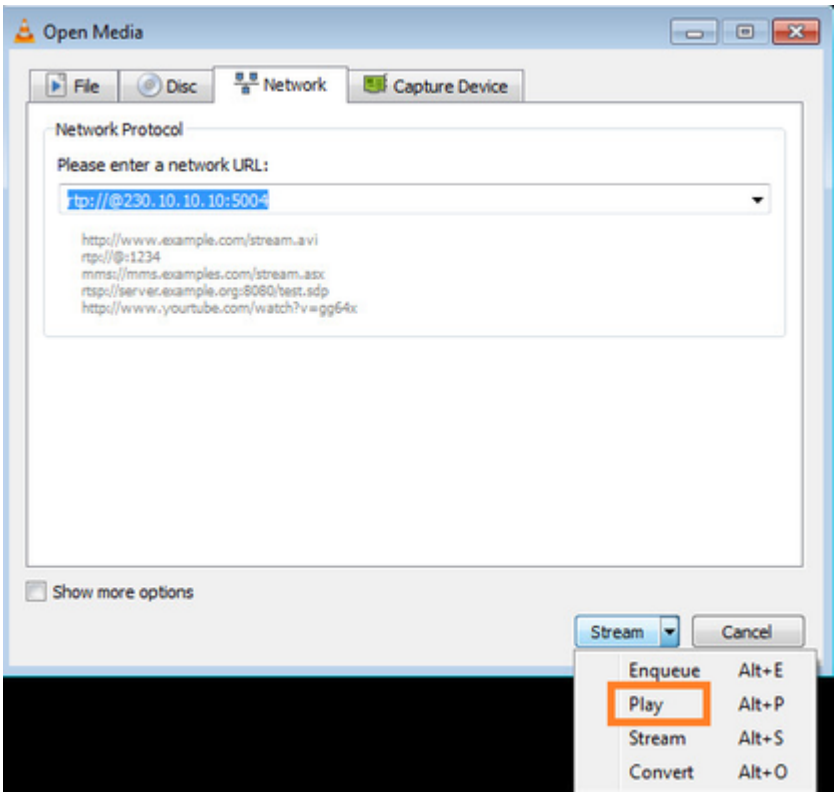

O FTD faz o proxy dos pacotes IGMP:

<#root>

```
firepower#
```

show capture

```
capture CAPI type raw-data trace interface INSIDE
```

[Capturing - 66 bytes]

```
<-- IGMP packets captured on ingress
match igmp any host 230.10.10.10
capture CAPO type raw-data interface OUTSIDE
```

[Capturing - 62 bytes]

```
<-- IGMP packets captured on egress
match igmp any host 230.10.10.10
```

O FTD altera o IP de origem:

<#root>

```
firepower#
```

show capture CAPI

```
1 packet captured
```

```
   1: 12:21:12.820483 802.1Q vlan#205 P6
```

192.168.1.50

```
 > 230.10.10.10 ip-proto-2, length 8    <-- The source IP of the packet on ingress interface
1 packet shown
firepower#
```

show capture CAPO

```
1 packet captured

   1: 12:21:12.820743
```

192.168.103.91

```
 > 230.10.10.10 ip-proto-2, length 8  <-- The source IP of the packet on egress interface
1 packet shown
```

Se você verificar o pcap no Wireshark, poderá ver que o pacote é completamente regenerado pelo firewall (a identificação de IP é alterada).

Uma entrada de grupo é criada no FTD:

<#root>

```
firepower#
```

show igmp group

```
IGMP Connected Group Membership
Group Address    Interface            Uptime    Expires   Last Reporter
```

230.10.10.10 INSIDE 00:15:22 00:03:28 192.168.1.50

```
<-- IGMP group is enabled on the ingress interface
239.255.255.250  INSIDE               00:15:27  00:03:29  192.168.1.50
```

O firewall FTD cria 2 conexões de plano de controle:

<#root>

```
firepower#
```

show conn all address 230.10.10.10

```
9 in use, 28 most used
Inspect Snort:
preserve-connection: 0 enabled, 0 in effect, 0 most enabled, 0 most in effect
```

IGMP INSIDE 192.168.1.50 NP Identity Ifc 230.10.10.10, idle 0:00:09, bytes 8, flags

```
<-- Connection terminated on the ingress interface
```

IGMP OUTSIDE 230.10.10.10 NP Identity Ifc 192.168.103.91, idle 0:00:09, bytes 8, flags

```
<-- Connection terminated on the egress interface
```

Rastreamento do primeiro pacote:

<#root>

```
firepower#
```

show capture CAPI packet-number 1 trace

```
6 packets captured
```

1: 12:21:12.820483 802.1Q vlan#205 P6 192.168.1.50 > 230.10.10.10 ip-proto-2, length 8

```
<-- The first packet of the flow
Phase: 1
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 5124 ns
Config:
Additional Information:
MAC Access list

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 5124 ns
Config:
Implicit Rule
Additional Information:
MAC Access list

Phase: 3
Type: ROUTE-LOOKUP
Subtype: No ECMP load balancing
Result: ALLOW
Elapsed time: 7808 ns
Config:
Additional Information:
Destination is locally connected. No ECMP load balancing.
Found next-hop 192.168.1.50 using egress ifc INSIDE(vrfid:0)

Phase: 4
Type: CLUSTER-DROP-ON-SLAVE
Subtype: cluster-drop-on-slave
Result: ALLOW
Elapsed time: 5368 ns
Config:
Additional Information:

Phase: 5
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 5368 ns
Config:
```

Implicit Rule
Additional Information:

Phase: 6
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 5368 ns
Config:
Additional Information:

Phase: 7
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 5368 ns
Config:
Additional Information:

Phase: 8
Type: CLUSTER-REDIRECT
Subtype: cluster-redirect
Result: ALLOW
Elapsed time: 40504 ns
Config:
Additional Information:

Phase: 9

Type: MULTICAST

<-- The packet is multicast

Subtype:

Result: ALLOW

Elapsed time: 976 ns

Config:

Additional Information:

Phase: 10

Type: FLOW-CREATION

<-- A new flow is created

Subtype:

Result: ALLOW

Elapsed time: 17568 ns

Config:

Additional Information:

New flow created with id 5945, packet dispatched to next module

Phase: 11

Type: FLOW-CREATION

<-- A second flow is created

Subtype:

Result: ALLOW

Elapsed time: 39528 ns

Config:

Additional Information:

New flow created with id 5946, packet dispatched to next module

Phase: 12
Type: NEXTHOP-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Lookup Nexthop on interface
Result: ALLOW
Elapsed time: 6344 ns
Config:
Additional Information:
Found next-hop 230.10.10.10 using egress ifc OUTSIDE(vrfid:0)

Phase: 13
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 9760 ns

```
Config:
Additional Information:
MAC Access list

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: INSIDE(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 154208 ns
```

Problemas conhecidos

Filtrar tráfego multicast em zonas de destino

Você não pode especificar uma zona de segurança de destino para a regra da Política de Controle de Acesso que corresponde ao tráfego multicast:

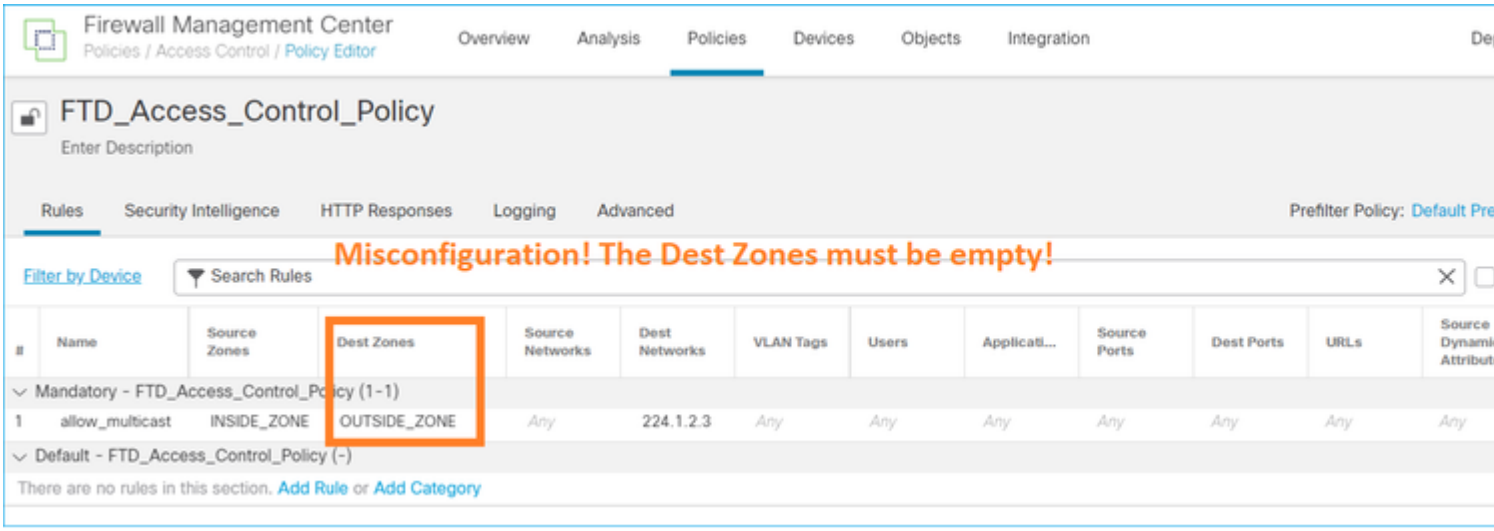

Este fato está igualmente documentado no guia do utilizador do FMC:

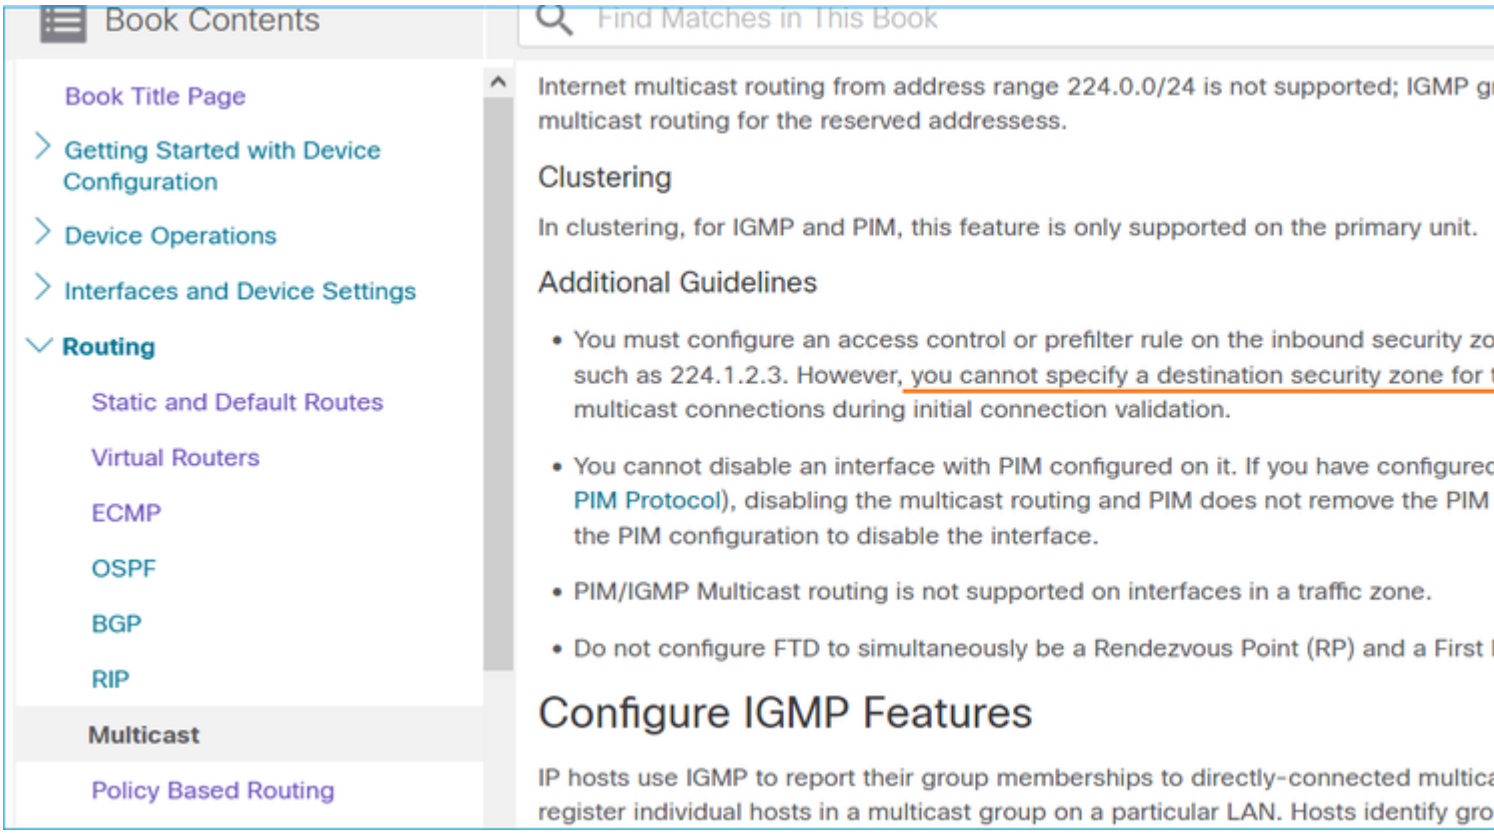

Os relatórios IGMP são negados pelo firewall quando o limite de interface IGMP é excedido

Por padrão, o firewall permite no máximo 500 junções ativas atuais (relatórios) em uma interface. Se esse limite for excedido, o firewall ignorará os relatórios IGMP de entrada adicionais dos receptores multicast.

Para verificar o limite de IGMP e as junções ativas, execute o comando **show igmp interface *nameif***:

<#root>

```
asa#

show igmp interface inside

inside is up, line protocol is up
  Internet address is 10.10.10.1/24
  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: 0 joins, 0 leaves
  IGMP querying router is 10.10.10.1 (this system)
```

O comando de depuração IGMP **debug igmp** mostra esta saída:

<#root>

asa#

debug igmp

Apr 20 2023 09:37:10: %ASA-7-711001: IGMP: Group 230.1.2.3 limit denied on inside

ID de bug da Cisco CSCuw84390 rastreia o aprimoramento para aumentar o limite de IGMP.

O Firewall ignora os relatórios IGMP para o intervalo de endereço 232.x.x.x/8

O intervalo de endereços 232.x.x.x/8 deve ser usado com o Source Specific Multicast (SSM). O firewall não oferece suporte à funcionalidade SSM (Source Specific Multicast, envio múltiplo específico de origem) do PIM e à configuração relacionada.

O comando de depuração IGMP **debug igmp** mostra esta saída:

<#root>

asa#

debug igmp

Apr 20 2023 09:37:10: %ASA-7-711001: IGMP: Received v2 Report on inside from 10.10.10.11 for 232.179.89.
Apr 20 2023 09:37:10: %ASA-7-711001: IGMP: group_db: add new group 232.179.89.253 on inside

Apr 20 2023 09:37:10: %ASA-7-711001: IGMP: Exclude report on inside ignored for SSM group 232.179.89.253

ID de bug da Cisco CSCsr53916 rastreia o aprimoramento para suportar o intervalo do SSM.

Informações Relacionadas

- Roteamento multicast para defesa contra ameaças do Firepower
- Solucionar problemas do Firepower Threat Defense e do ASA Multicast PIM

Sobre esta tradução
A Cisco traduziu este documento com a ajuda de tecnologias de tradução automática e humana para oferecer conteúdo de suporte aos seus usuários no seu próprio idioma, independentemente da localização.

Observe que mesmo a melhor tradução automática não será tão precisa quanto as realizadas por um tradutor profissional.

A Cisco Systems, Inc. não se responsabiliza pela precisão destas traduções e recomenda que o documento original em inglês (link fornecido) seja sempre consultado.