Risoluzione dei problemi di base di Firepower Threat Defense per IGMP e multicast

Sommario

Introduzione

Questo documento descrive le basi del multicast e il modo in cui Firepower Threat Defense (FTD) implementa il protocollo IGMP (Internet Group Management Protocol).

Prerequisiti

Requisiti

Conoscenze base di routing IP.

Componenti usati

Le informazioni discusse in questo documento fanno riferimento a dispositivi usati in uno specifico ambiente di emulazione. Su tutti i dispositivi menzionati nel documento la configurazione è stata

ripristinata ai valori predefiniti. Se la rete è operativa, valutare attentamente eventuali conseguenze derivanti dall'uso dei comandi.

Il contenuto di questo articolo è applicabile anche al software Adaptive Security Appliance (ASA).

Le informazioni fornite in questo documento si basano sulle seguenti versioni software e hardware:

- Cisco Firepower 4125 Threat Defense versione 7.1.0.
- Firepower Management Center (FMC) versione 7.1.0.
- ASA versione 9.19.1.

Premesse

Definizioni

- Unicast = da un singolo host a un altro host (uno a uno).
- Trasmissione = da un singolo host a TUTTI gli host possibili (uno a tutti).
- Multicast = da un host di un gruppo di host a un gruppo di host (uno-a-molti o molti-a-molti).
- Anycast = da un host all'host più vicino di un gruppo (uno-a-uno-di-molti).

Nozioni di base

- Multicast RFC 988 è stato scritto nel 1986 da Steve Deering.
- Il multicast IPv4 utilizza l'intervallo 224.0.0.0/4 (primi 4 bit 1110) - 224.0.0.0 - 239.255.255.255.
- Per IPv4, l'indirizzo MAC L2 deriva da IP multicast L3: 01005e (24 bit) + 25^{esimo} bit sempre 0 + 23 bit inferiori dell'indirizzo IPv4 multicast.
- Il multicast IPv6 utilizza l'intervallo FF00::/8 ed è più flessibile del multicast IPv4 in quanto può incorporare IP di Rendezvous Point (RP).
- Per IPv6 l'indirizzo MAC L2 deriva dal multicast L3: 3333 + 32 bit inferiori dell'indirizzo IPv6 multicast.
- Vantaggi del multicast: efficienza dovuta alla riduzione del carico sull'origine. Prestazioni, in quanto evita la duplicazione del traffico o l'effetto flooding.
- Svantaggi del multicast: trasporto inaffidabile (basato su UDP), nessuna prevenzione delle congestioni, consegna fuori sequenza.
- Il multicast non è supportato nell'Internet pubblica perché per abilitarlo sono necessari tutti i dispositivi nel percorso. In genere viene utilizzato quando tutti i dispositivi sono sottoposti a un'autorità amministrativa comune.
- Applicazioni multicast tradizionali: streaming video interno, videoconferenza.

Confronto tra multicast e unicast replicato

Nell'unicast replicato l'origine crea più copie dello stesso pacchetto unicast (repliche) e le invia a più host di destinazione. Il multicast sposta il carico dall'host di origine alla rete, mentre in unicast replicato tutto il lavoro viene eseguito sull'host di origine.

Configurazione

Nozioni di base su IGMP

- IGMP è la "lingua" parlata tra i ricevitori multicast e il dispositivo L3 locale (in genere un router).
- IGMP è un protocollo di layer 3 (come ICMP) e usa il protocollo IP numero 2.
- Attualmente sono disponibili 3 versioni IGMP. La versione IGMP predefinita sul firewall è la versione 2. Al momento sono supportate solo le versioni 1 e 2.
- Le differenze principali tra IGMPv1 e IGMPv2 sono:

    - IGMPv1 non ha un messaggio di uscita dal gruppo.
    - IGMPv1 non dispone di query specifiche del gruppo (utilizzate dal firewall quando un host lascia un gruppo multicast).
    - IGMPv1 non dispone di un processo di selezione tramite query.

- IGMPv3 non è attualmente supportato su ASA/FTD, ma come riferimento, la differenza importante tra IGMPv2 e IGMPv3 è l'inclusione di una query specifica di gruppo e origine in IGMPv3, che viene utilizzata in SSM (Source-Specific Multicast).
- Query IGMPv1/IGMPv2/IGMPv3 = 224.0.0.1
  Uscita IGMPv2 = 224.0.0.2
  Rapporto appartenenza IGMPv3 = 224.0.0.22
- Se un host desidera partecipare, può inviare un messaggio di rapporto appartenenza IGMP non richiesto:

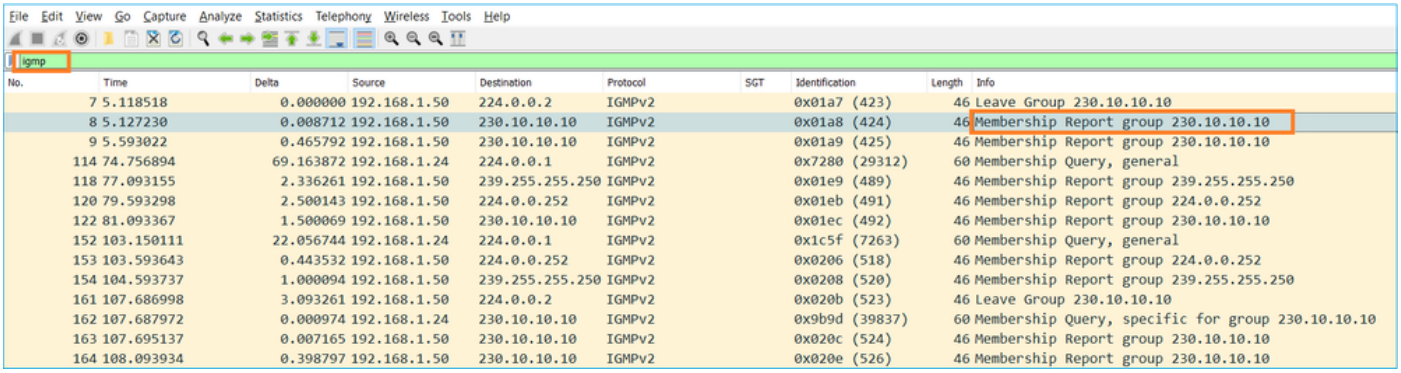

No.	Time	Delta	Source	Destination	Protocol	SGT	Identification	Length	Info
7	5.118518	0.000000	192.168.1.50	224.0.0.2	IGMPv2		0x01a7 (423)	46	Leave Group 230.10.10.10
8	5.127230	0.008712	192.168.1.50	230.10.10.10	IGMPv2		0x01a8 (424)	46	Membership Report group 230.10.10.10
9	5.593022	0.465792	192.168.1.50	230.10.10.10	IGMPv2		0x01a9 (425)	46	Membership Report group 230.10.10.10
114	74.756894	69.163872	192.168.1.24	224.0.0.1	IGMPv2		0x7280 (29312)	60	Membership Query, general
118	77.093155	2.336261	192.168.1.50	239.255.255.250	IGMPv2		0x01e9 (489)	46	Membership Report group 239.255.255.250
120	79.593298	2.500143	192.168.1.50	224.0.0.252	IGMPv2		0x01eb (491)	46	Membership Report group 224.0.0.252
122	81.093367	1.500069	192.168.1.50	230.10.10.10	IGMPv2		0x01ec (492)	46	Membership Report group 230.10.10.10
152	103.150111	22.056744	192.168.1.24	224.0.0.1	IGMPv2		0x1c5f (7263)	60	Membership Query, general
153	103.593643	0.443532	192.168.1.50	224.0.0.252	IGMPv2		0x0206 (518)	46	Membership Report group 224.0.0.252
154	104.593737	1.000094	192.168.1.50	239.255.255.250	IGMPv2		0x0208 (520)	46	Membership Report group 239.255.255.250
161	107.686998	3.093261	192.168.1.50	224.0.0.2	IGMPv2		0x020b (523)	46	Leave Group 230.10.10.10
162	107.687972	0.000974	192.168.1.24	230.10.10.10	IGMPv2		0x9b9d (39837)	60	Membership Query, specific for group 230.10.10.10
163	107.695137	0.007165	192.168.1.50	230.10.10.10	IGMPv2		0x020c (524)	46	Membership Report group 230.10.10.10
164	108.093934	0.398797	192.168.1.50	230.10.10.10	IGMPv2		0x020e (526)	46	Membership Report group 230.10.10.10

- Dal punto di vista del firewall, sono disponibili 2 tipi di query IGMP: query generali e query specifiche di gruppo
- Quando il firewall riceve un messaggio IGMP Abbandona gruppo, deve verificare se nella subnet sono presenti altri membri di tale gruppo. Per questo motivo, il firewall invia una query specifica del gruppo:

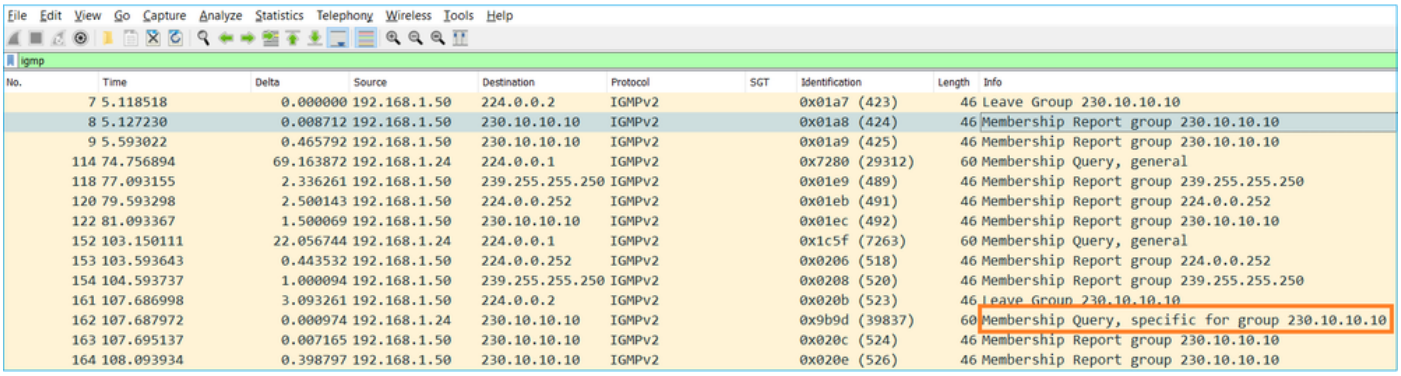

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

igmp

No.	Time	Delta	Source	Destination	Protocol	SGT	Identification	Length	Info
7	5.118518	0.000000	192.168.1.50	224.0.0.2	IGMPv2		0x01a7 (423)	46	Leave Group 230.10.10.10
8	5.127230	0.008712	192.168.1.50	230.10.10.10	IGMPv2		0x01a8 (424)	46	Membership Report group 230.10.10.10
9	5.593022	0.465792	192.168.1.50	230.10.10.10	IGMPv2		0x01a9 (425)	46	Membership Report group 230.10.10.10
114	74.756894	69.163872	192.168.1.24	224.0.0.1	IGMPv2		0x7280 (29312)	60	Membership Query, general
118	77.093155	2.336261	192.168.1.50	239.255.255.250	IGMPv2		0x01e9 (489)	46	Membership Report group 239.255.255.250
120	79.593298	2.500143	192.168.1.50	224.0.0.252	IGMPv2		0x01eb (491)	46	Membership Report group 224.0.0.252
122	81.093367	1.500069	192.168.1.50	230.10.10.10	IGMPv2		0x01ec (492)	46	Membership Report group 230.10.10.10
152	103.150111	22.056744	192.168.1.24	224.0.0.1	IGMPv2		0x1c5f (7263)	60	Membership Query, general
153	103.593643	0.443532	192.168.1.50	224.0.0.252	IGMPv2		0x0206 (518)	46	Membership Report group 224.0.0.252
154	104.593737	1.000094	192.168.1.50	239.255.255.250	IGMPv2		0x0208 (520)	46	Membership Report group 239.255.255.250
161	107.686998	3.093261	192.168.1.50	224.0.0.2	IGMPv2		0x020b (523)	46	Leave Group 230.10.10.10
162	107.687972	0.000974	192.168.1.24	230.10.10.10	IGMPv2		0x9b9d (39837)	60	Membership Query, specific for group 230.10.10.10
163	107.695137	0.007165	192.168.1.50	230.10.10.10	IGMPv2		0x020c (524)	46	Membership Report group 230.10.10.10
164	108.093934	0.398797	192.168.1.50	230.10.10.10	IGMPv2		0x020e (526)	46	Membership Report group 230.10.10.10

- Nelle subnet in cui sono presenti più router/firewall, viene selezionato un interrogante (un dispositivo che invia tutte le query IGMP):

<#root>

```
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
```

- Su FTD, simile a una appliance ASA classica, è possibile abilitare il comando debug igmp per visualizzare i messaggi relativi a 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
```

- Un host in genere lascia un gruppo multicast con un messaggio Leave Group (IGMPv2).

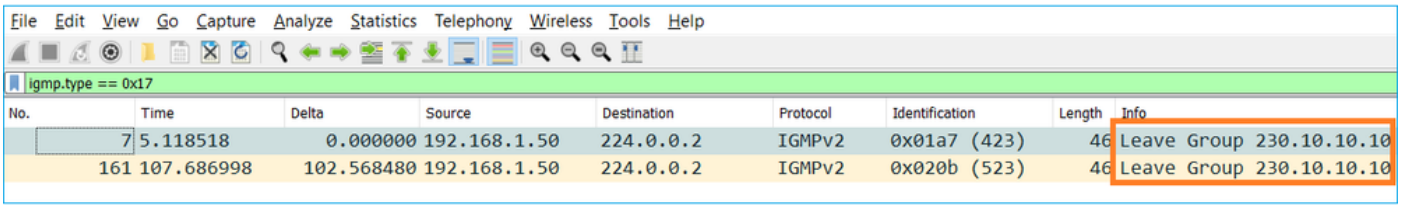

No.	Time	Delta	Source	Destination	Protocol	Identification	Length	Info
7	5.118518	0.000000	192.168.1.50	224.0.0.2	IGMPv2	0x01a7 (423)	46	Leave Group 230.10.10.10
161	107.686998	102.568480	192.168.1.50	224.0.0.2	IGMPv2	0x020b (523)	46	Leave Group 230.10.10.10

Attività 1 - Traffico multicast Control-Plane

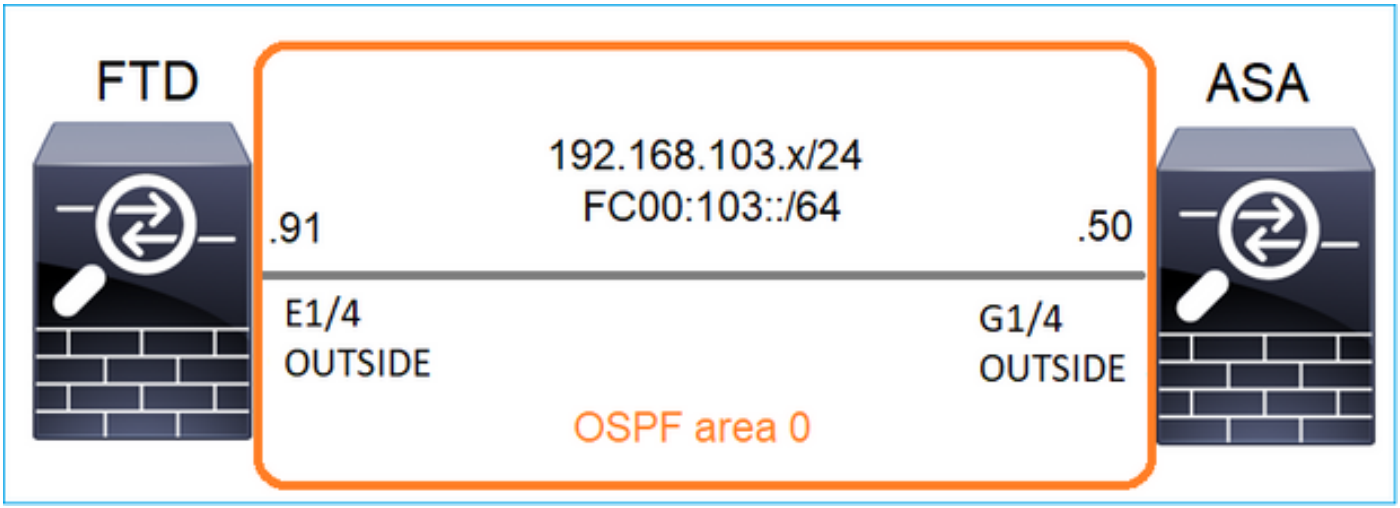

Configurare un OSPFv2 e un OSPFv3 tra FTD e ASA. Controllare come i 2 dispositivi gestiscono il traffico L2 e L3 multicast generato da OSPF.

Soluzione

Configurazione OSPFv2

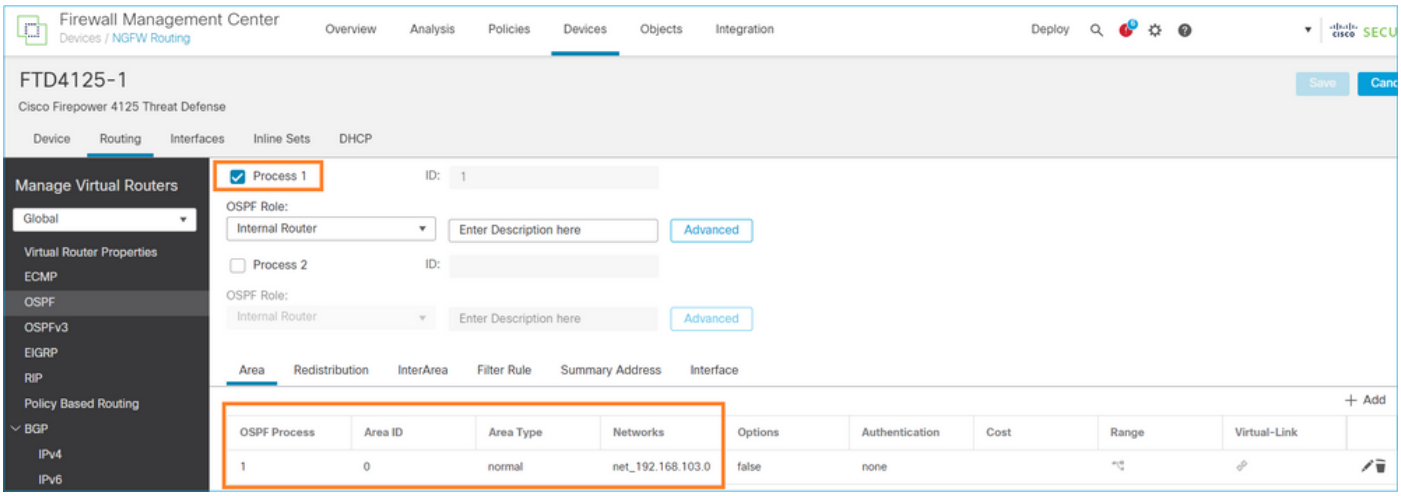

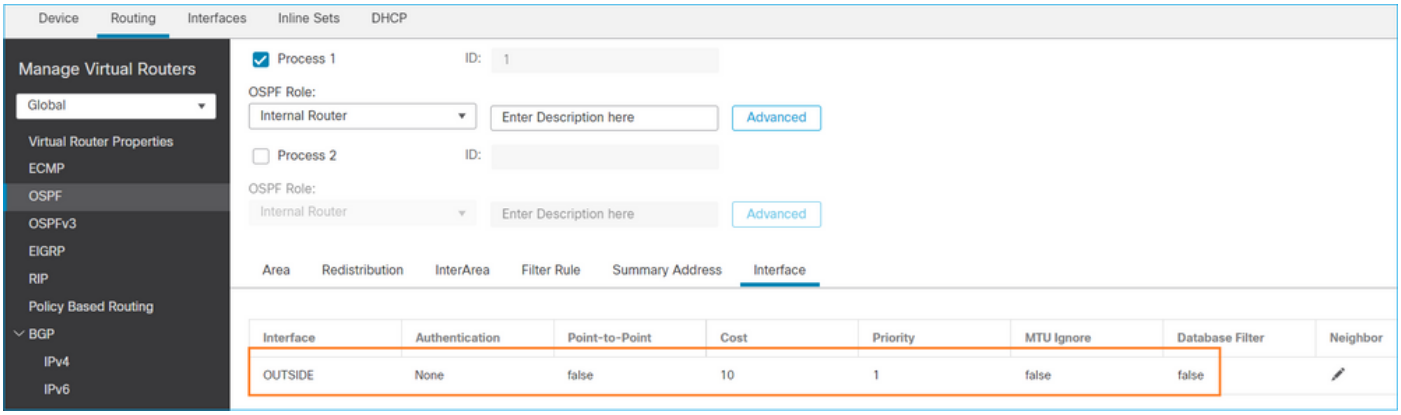

Analogamente, per OSPFv3

Configurazione su CLI 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
```

La configurazione crea queste voci nelle tabelle di autorizzazione ASP (Accelerated Security Path)

FTD in modo che il traffico multicast in entrata non venga bloccato:

<#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
```

Per 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
...
```

Le adiacenze OSPFv2 e OSPFv3 sono attive:

<#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
```

Le sessioni OSPF multicast terminate nella casella sono le seguenti:

<#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
```

Come prova, abilitare l'acquisizione per IPv4 e cancellare le connessioni al dispositivo:

<#root>

```
firepower#

capture CAP interface OUTSIDE trace

firepower#

clear conn all

12 connection(s) deleted.
firepower#

clear capture CAP

firepower# !
```

Avviso: si è verificata un'interruzione dell'alimentazione. L'esempio viene mostrato solo a scopo dimostrativo.

I pacchetti OSPF acquisiti:

<#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]
```

Di seguito viene riportata la modalità di gestione del pacchetto multicast OSPFv2 da parte del 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
```

In questo modo il pacchetto multicast OSPFv3 viene gestito dal 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
```

Attività 2 - Configurazione del multicast di base

Topologia

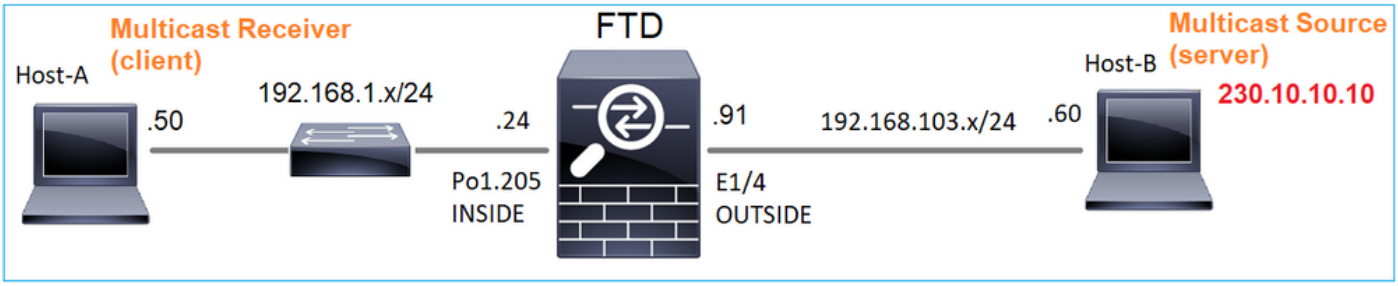

Requisito

Configurare il firewall in modo che il traffico multicast proveniente dal server venga inviato al client multicast su IP 230.10.10.10

Soluzione

Dal punto di vista del firewall, la configurazione minima è abilitare il routing multicast a livello globale. Ciò consente di abilitare in background IGMP e PIM su tutte le interfacce del firewall.

Nell'interfaccia utente del CCP:

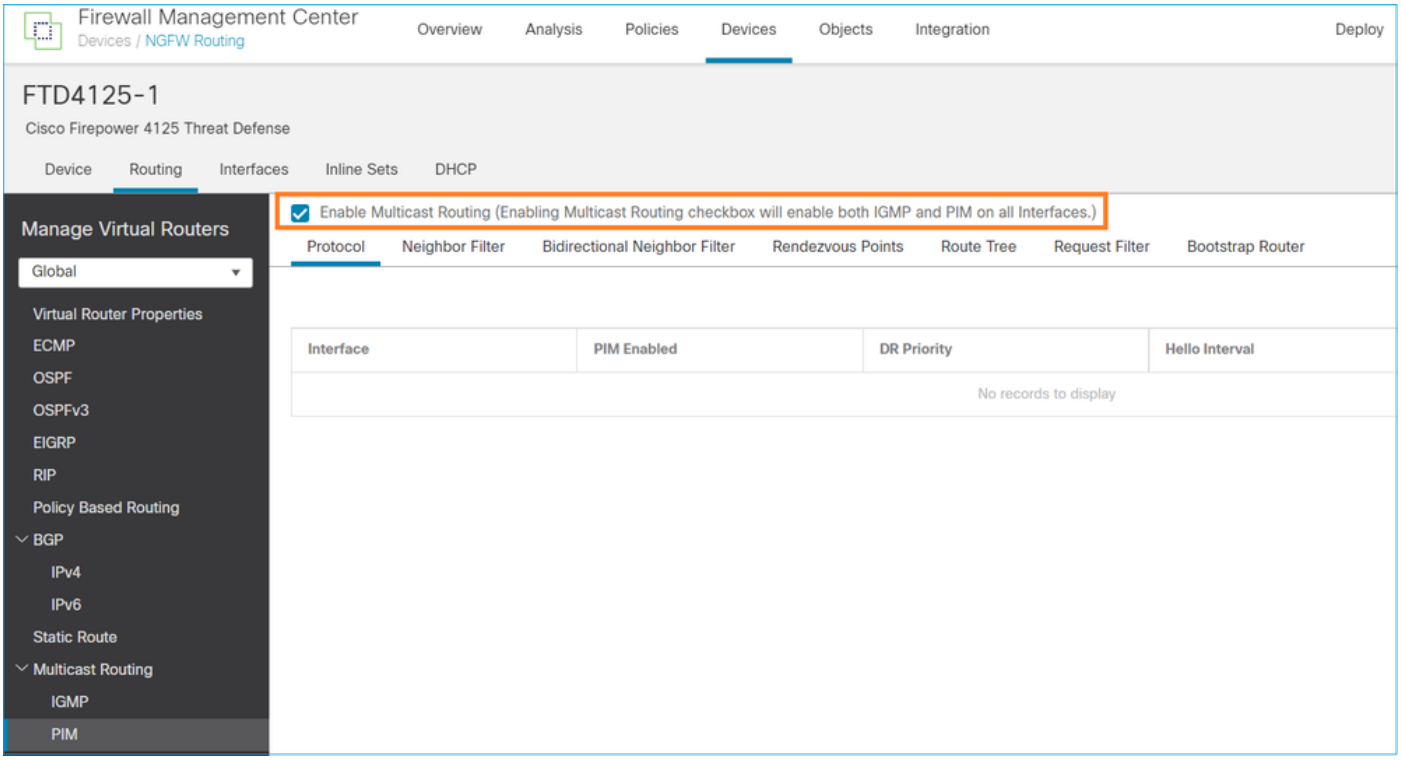

Dalla CLI del firewall, questa è la configurazione push:

<#root>

```
firepower#
show run multicast-routing
multicast-routing
<-- Multicast routing is enabled
```

Verifica 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
```

Verifica 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
```

Verifica 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 seco
(*,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
```

Traffico multicast attraverso il firewall

In questo caso, l'applicazione VLC media player viene utilizzata come server multicast e client per testare il traffico multicast:

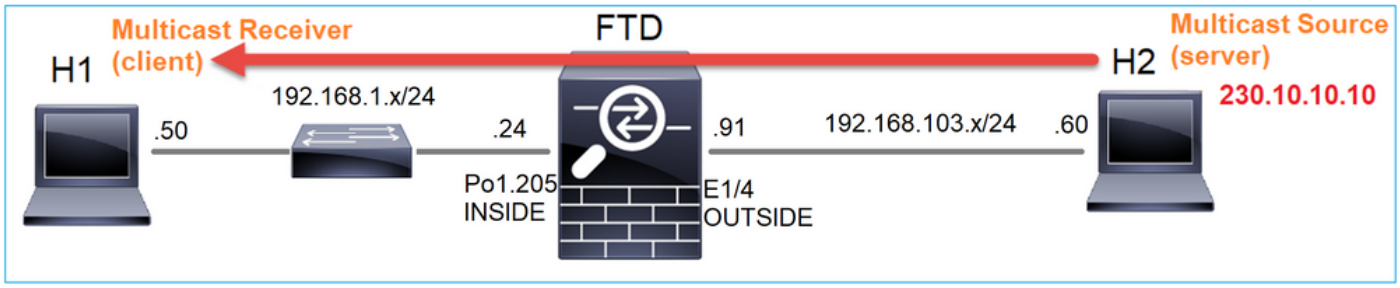

Configurazione server multicast VLC:

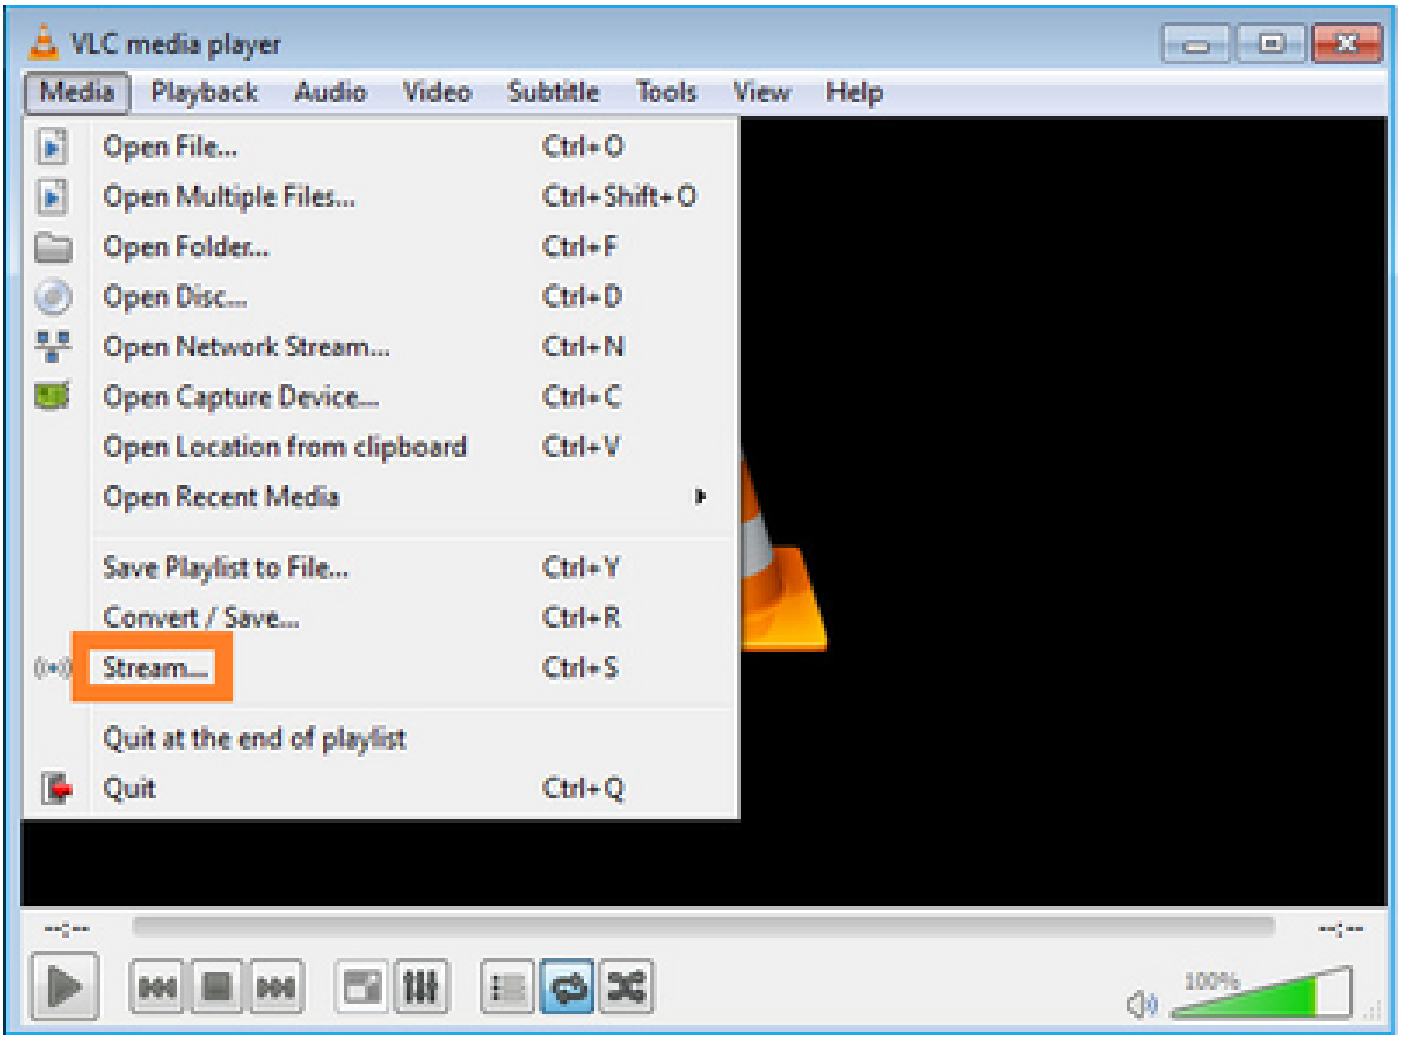
VLC media player
Media
Playback
Audio
Video
Subtitle
Tools
View
Help
Open File... Ctrl+O
Open Multiple Files... Ctrl+Shift+O
Open Folder... Ctrl+F
Open Disc... Ctrl+D
Open Network Stream... Ctrl+N
Open Capture Device... Ctrl+C
Open Location from clipboard Ctrl+V
Open Recent Media
Save Playlist to File... Ctrl+Y
Convert / Save... Ctrl+R
Stream... Ctrl+S
Quit at the end of playlist
Quit Ctrl+Q
100%

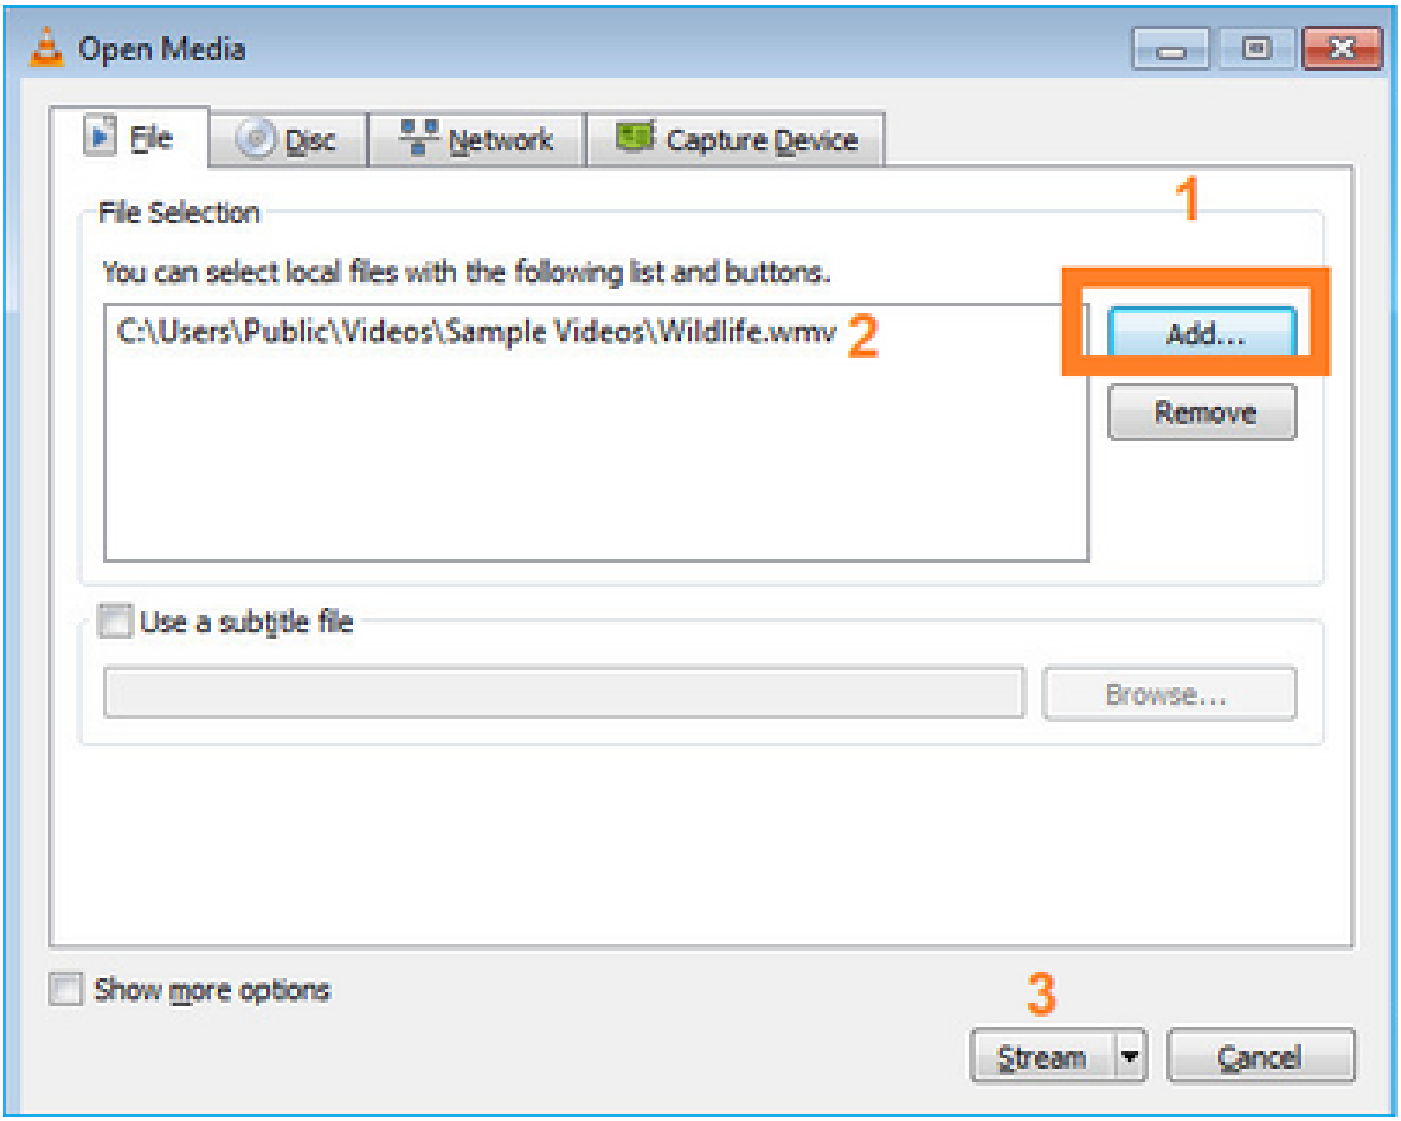

Nella schermata successiva selezionare Avanti.

Selezionare il formato:

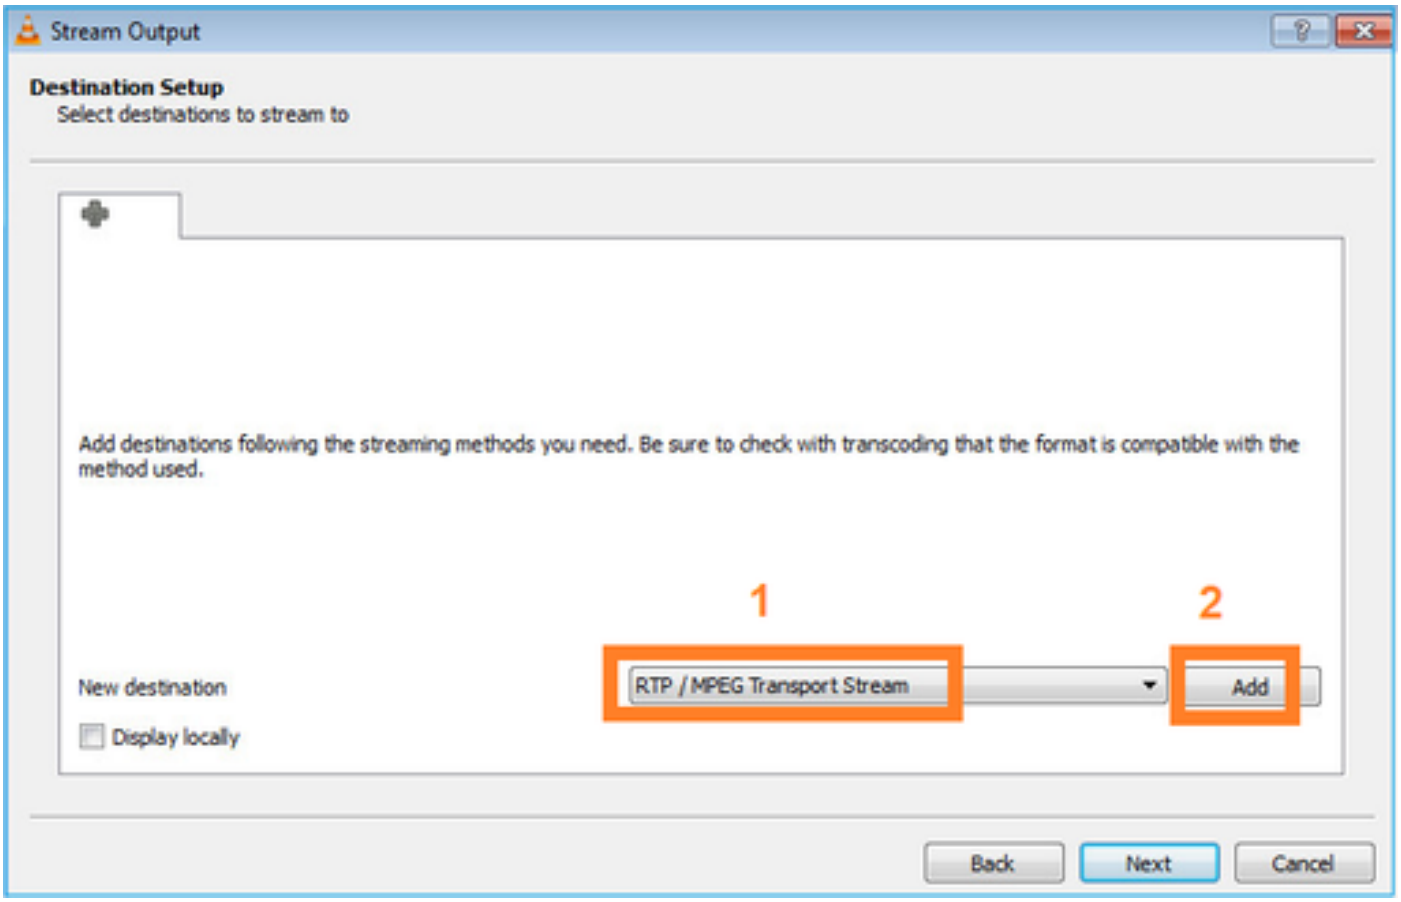

Specificare l'indirizzo IP e la porta multicast:

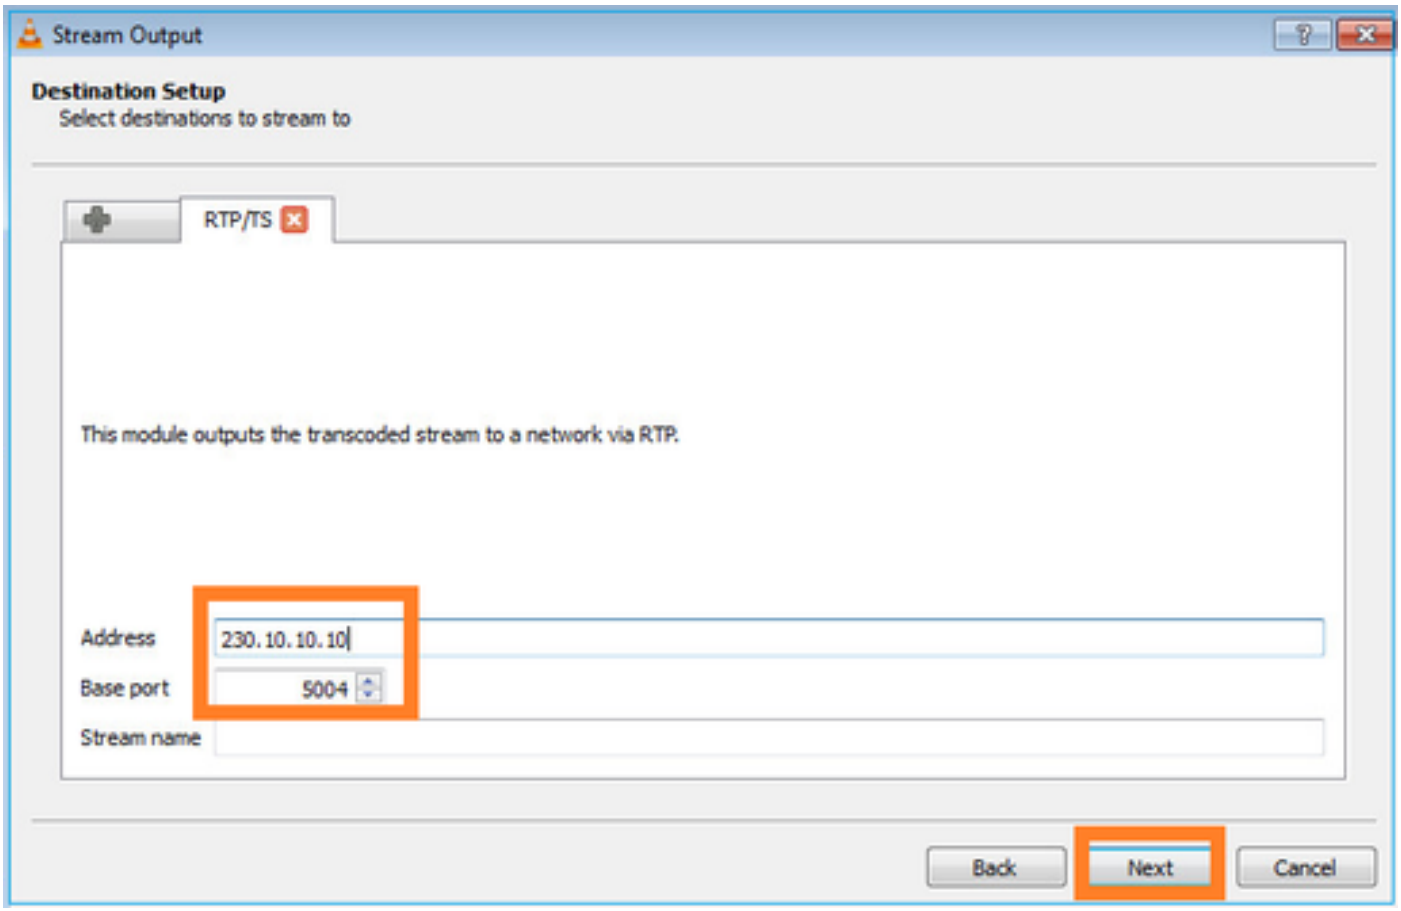

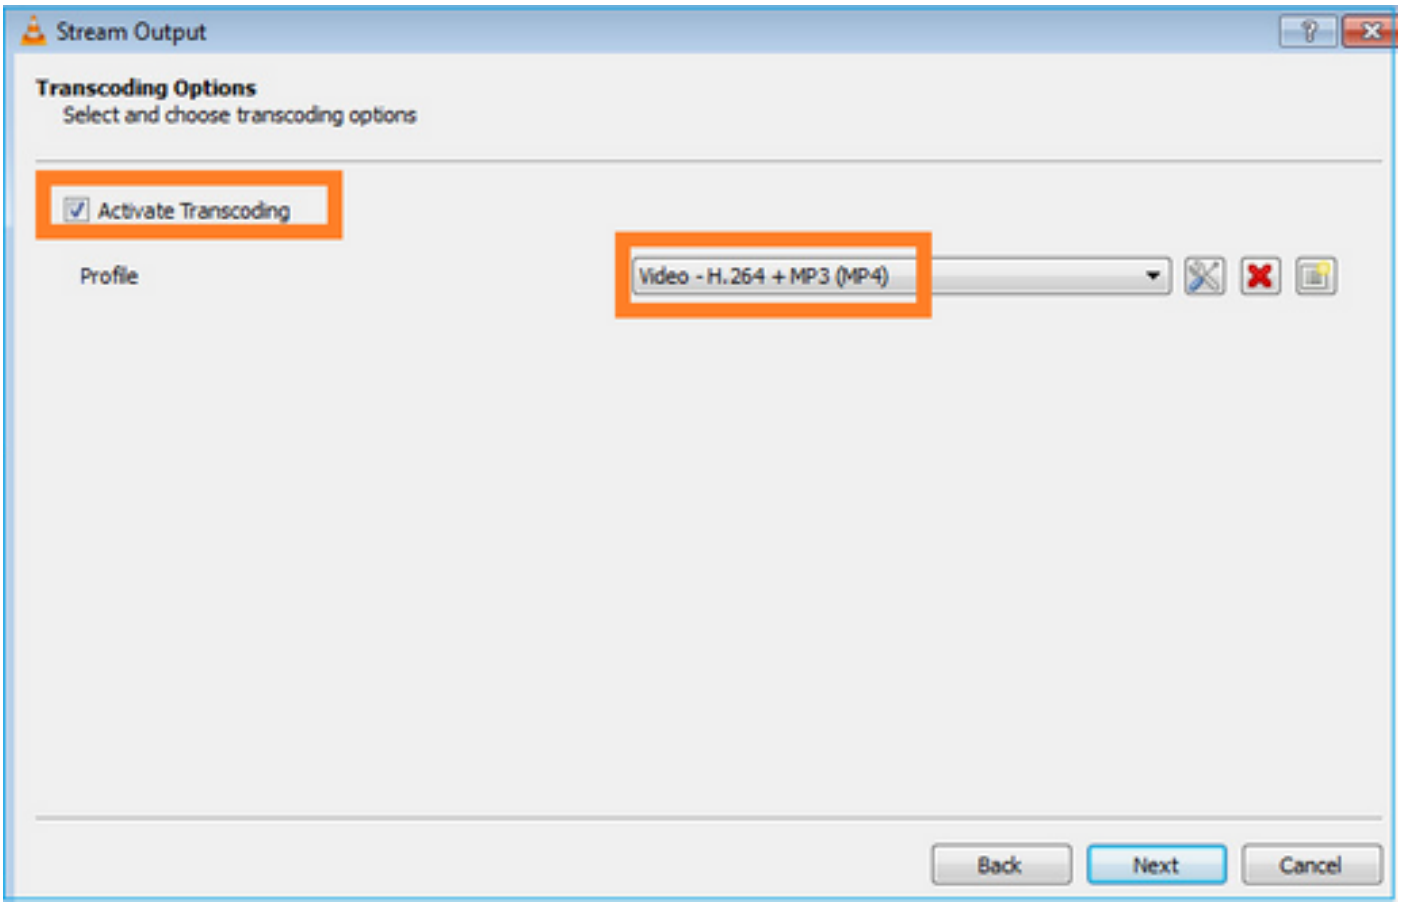

Abilita le clip LINA sul firewall FTD:

<#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
```

Selezionare il pulsante Stream per il dispositivo per avviare il flusso multicast:

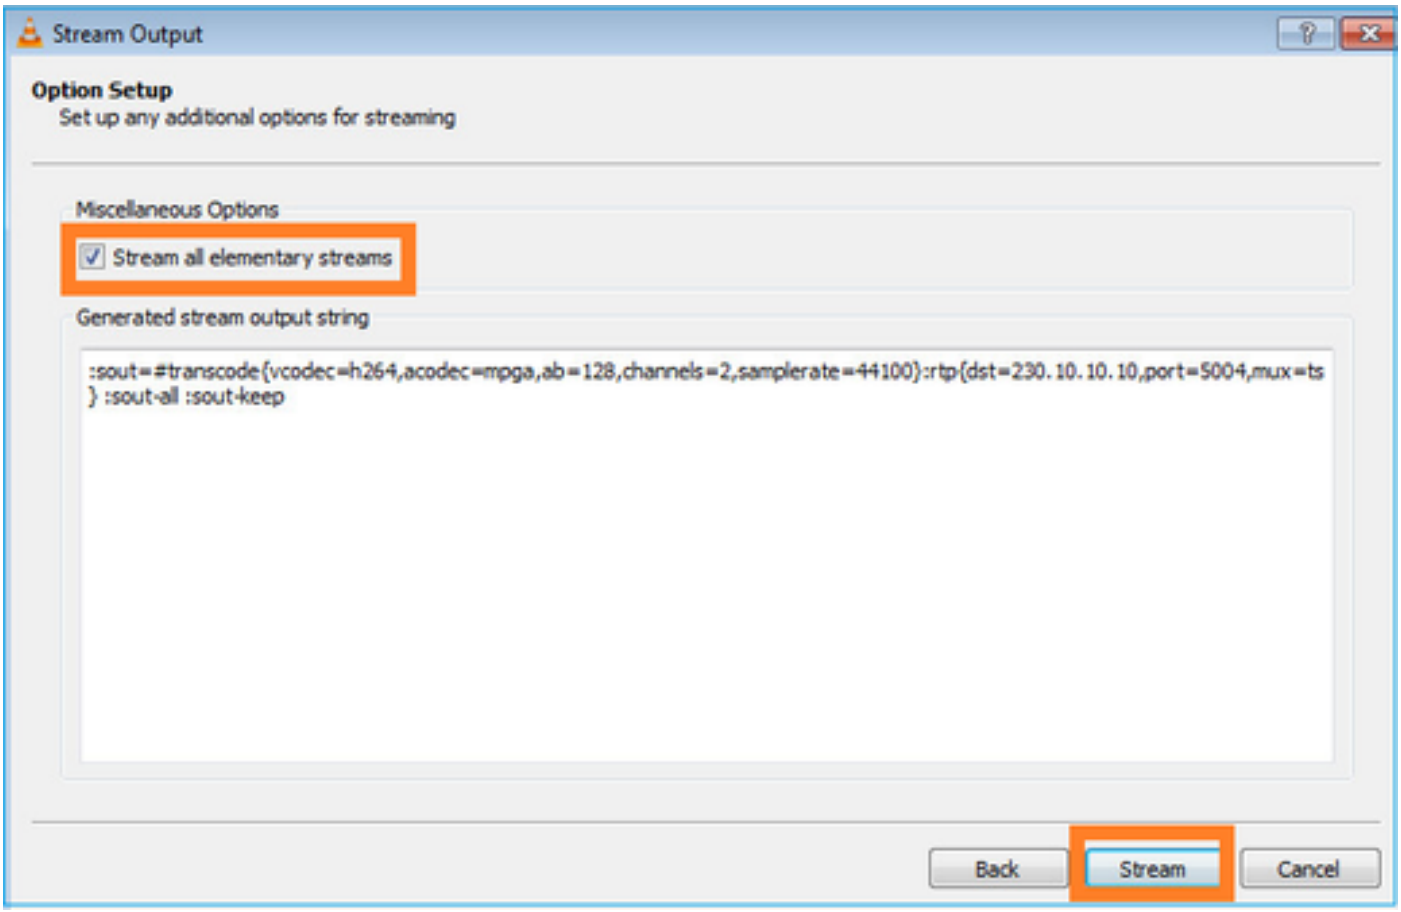

Abilitare l'opzione "loop" in modo che il flusso venga inviato in modo continuo:

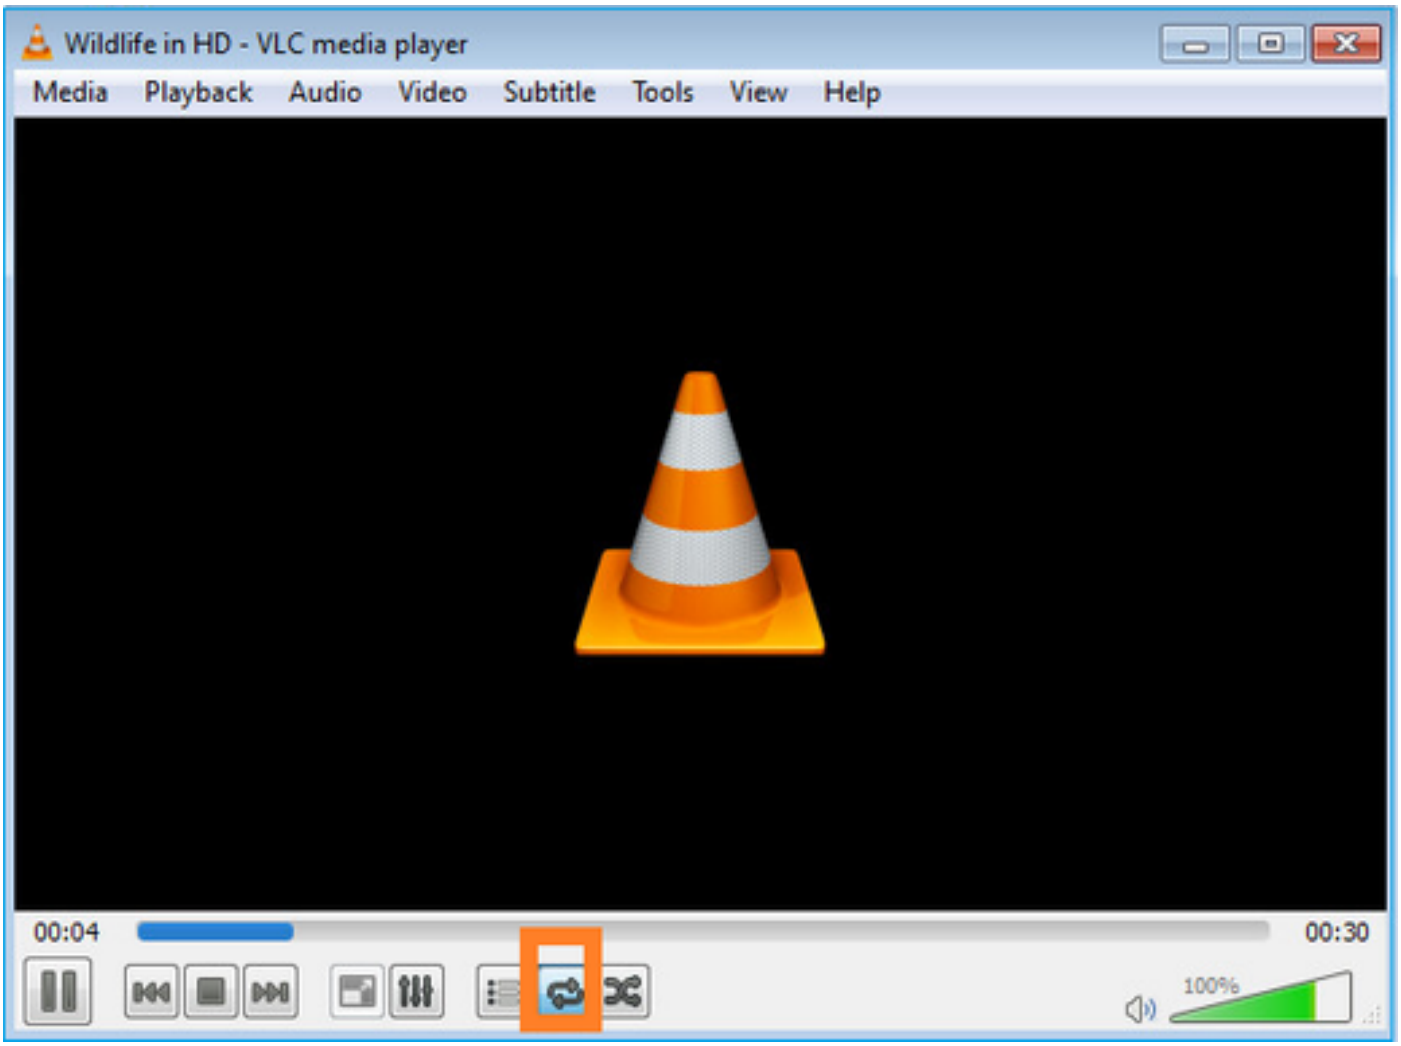

Verifica (scenario non operativo)

Questo scenario è la dimostrazione di uno scenario non operativo. L'obiettivo è dimostrare il comportamento del firewall.

Il dispositivo firewall ottiene il flusso multicast, ma non lo inoltra:

<#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
```

Gocce ASP LINA del firewall visualizzate:

<#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
```

Per tracciare un pacchetto, è necessario acquisire il primo pacchetto del flusso multicast. Per questo motivo, cancellare i flussi correnti:

<#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
...
```

L'opzione ‘detail’ rivela l'indirizzo 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)
...
```

La traccia di un pacchetto reale mostra che il pacchetto è autorizzato, ma non è questo ciò che accade realmente:

<#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
```

In base ai contatori mroute e mfib, i pacchetti vengono scartati perché l'elenco delle interfacce in uscita (OIL) è vuoto:

<#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
```

I contatori MFIB mostrano i guasti di RPF, che in questo caso non è quello che succede veramente:

<#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
```

Errori RPF simili nell'output '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
```

Configurare il ricevitore multicast VLC:

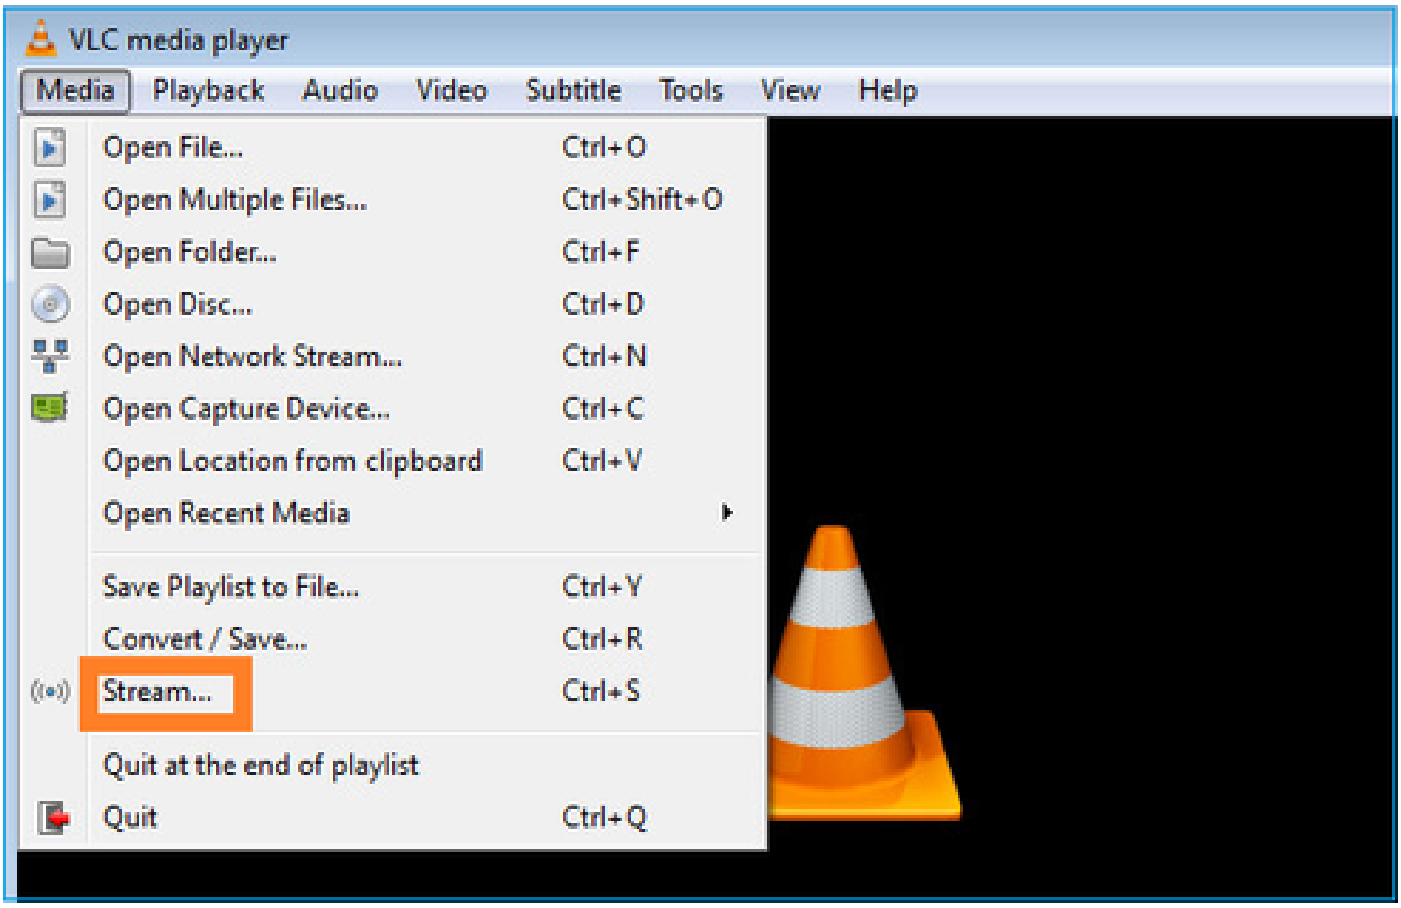

Specificare l'indirizzo IP di origine multicast e selezionare Play:

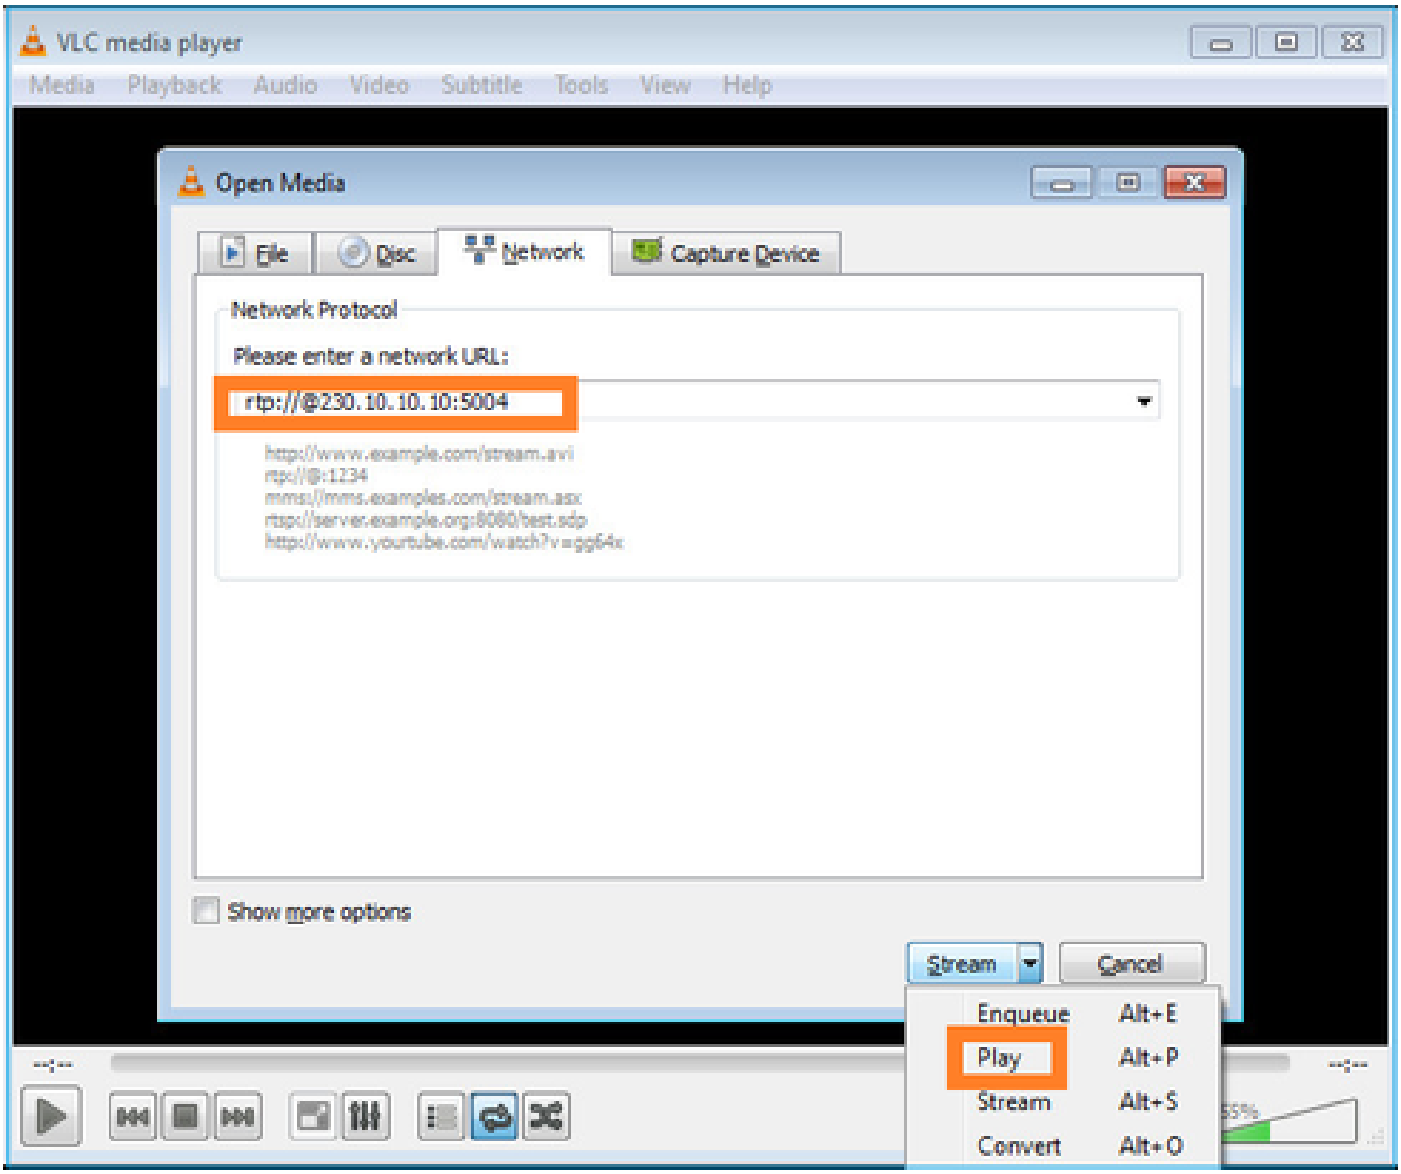

Nel back-end, non appena si seleziona Play, l'host annuncia la propria disponibilità a unirsi al gruppo multicast specifico e invia un messaggio IGMP Report:

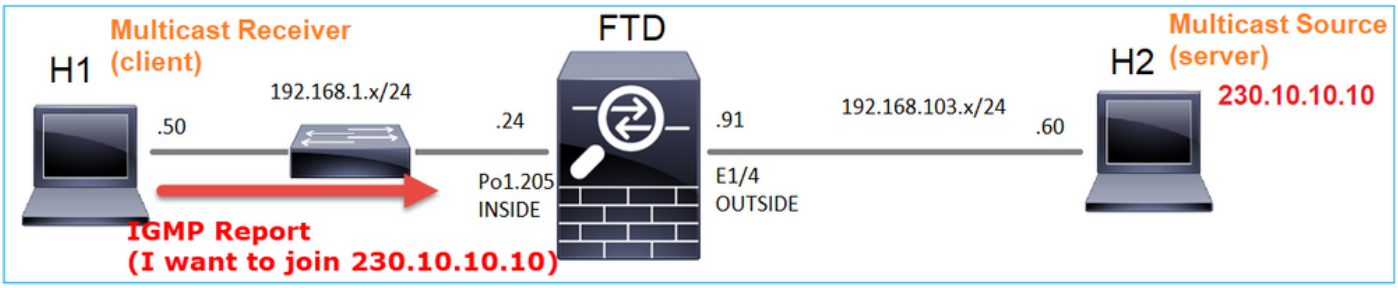

Se si abilita un debug, è possibile visualizzare i messaggi del report IGMP:

<#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
```

Verrà avviato il flusso:

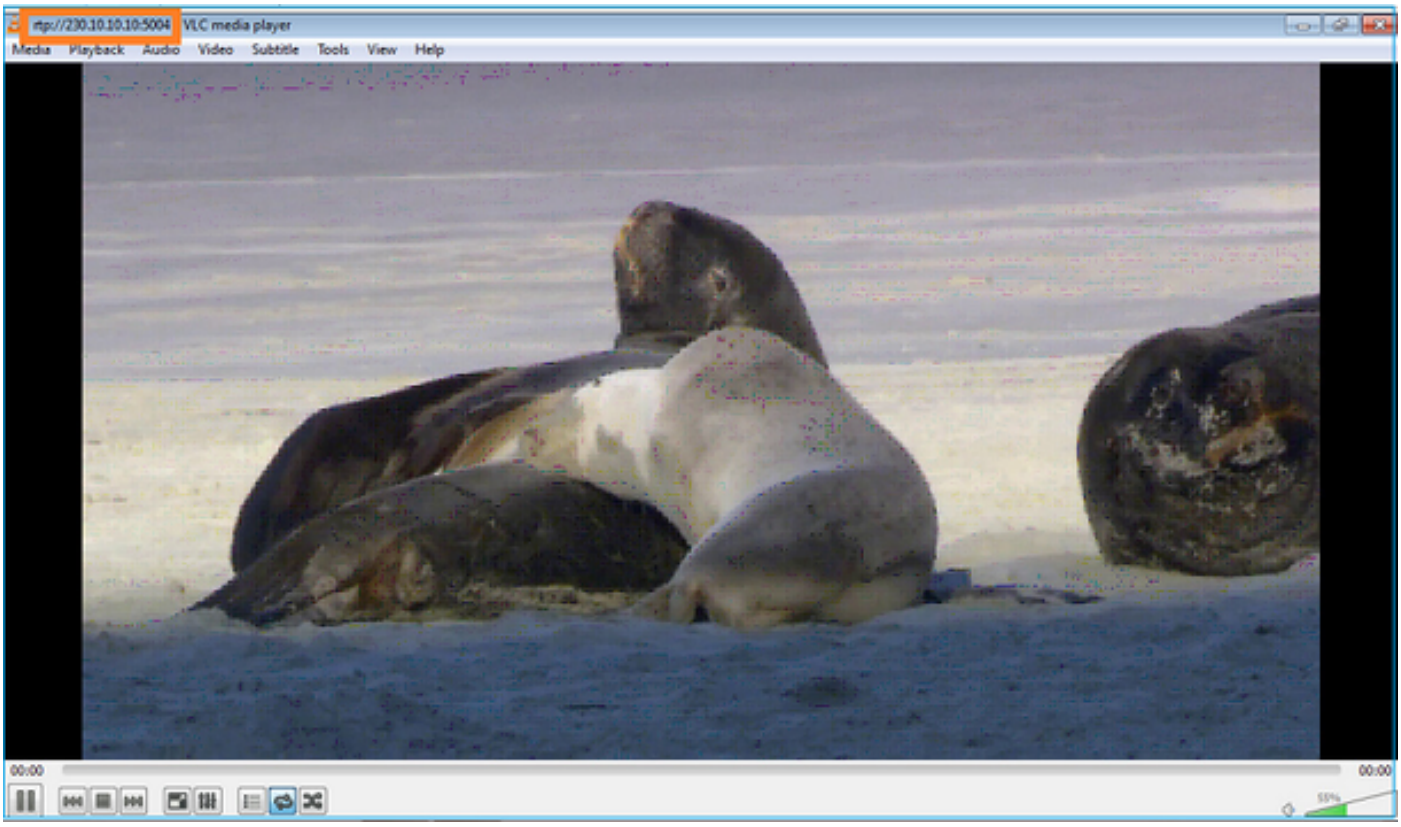

Verifica (scenario operativo)

<#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
```

Tabella di route del 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
```

contatori 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
```

Snooping IGMP

- Lo snooping IGMP è un meccanismo utilizzato sugli switch per prevenire il trasferimento del multicast.
- Lo switch esegue il monitoraggio dei report IGMP per determinare dove si trovano gli host (ricevitori).
- Lo switch esegue il monitoraggio delle query IGMP per determinare la posizione dei router/firewall (mittenti).
- Lo snooping IGMP è abilitato per impostazione predefinita sulla maggior parte degli switch Cisco. Per ulteriori informazioni, consultare le guide di commutazione correlate. Di seguito viene riportato un esempio di output per uno 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
```

Attività 3 - Gruppo statico IGMP e join-group IGMP

Panoramica

	ip igmp static-group	join-group ip igmp
Applicato sull'interfaccia FTD?	Sì	Sì
L'FTD attrae un flusso multicast?	Sì, viene inviato un join PIM verso il dispositivo a monte. verso l'origine o verso il punto di rendering (RP). Questo si verifica solo se l'FTD con questo comando è il PIM Designated Router (DR) su quell'interfaccia.	Sì, viene inviato un join PIM verso il dispositivo a monte. verso l'origine o verso il punto di rendering (RP). Questo si verifica solo se l'FTD con questo comando è il PIM Designated Router (DR) su quell'interfaccia.
L'FTD inoltra il traffico multicast dall'interfaccia?	Sì	Sì
L'FTD utilizza e risponde al traffico multicast?	No	Sì, l'FTD reindirizza il flusso multicast alla CPU, lo consuma e risponde all'origine.
Impatto CPU	Minimo perché il pacchetto non è indirizzato alla CPU.	Può influire sulla CPU FTD in quanto ogni pacchetto multicast appartenente al gruppo viene indirizzato alla CPU FTD.

Attività richiesta

Supponiamo di avere questa topologia:

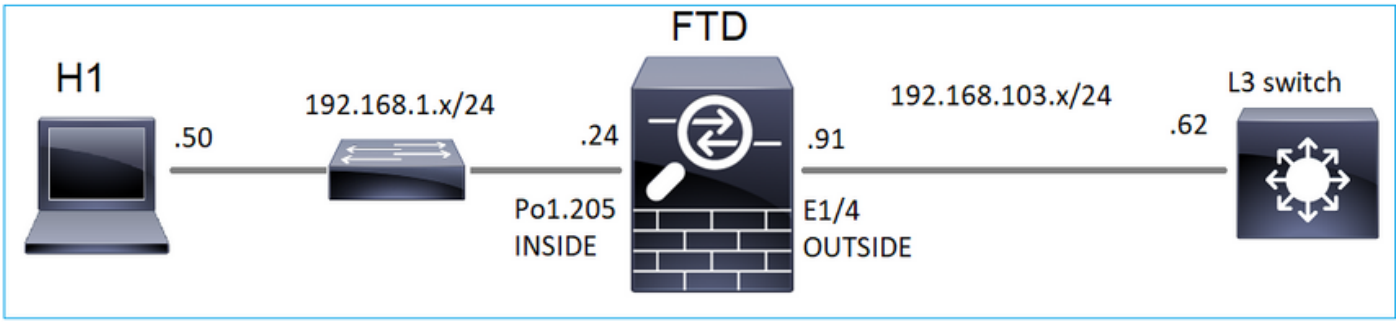

Sul firewall abilitare queste clip:

<#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. Usare il comando ping ICMP dallo switch L3 per inviare il traffico multicast a IP 230.11.11.11 e controllare come viene gestito dal firewall.
2. Abilitare il comando igmp static-group sull'interfaccia INSIDE del firewall e controllare come il flusso multicast (IP 230.11.11.11) viene gestito dal firewall.
3. Abilitare il comando igmp static-group sull'interfaccia INSIDE del firewall e controllare come il flusso multicast (IP 230.11.11.11) viene gestito dal firewall.

Soluzione

Il firewall non ha alcun percorso per 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
```

Un modo semplice per verificare il multicast è usare lo strumento ping ICMP. In questo caso, eseguire un ping tra l'indirizzo R2 e l'indirizzo 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:
.............................
```

Sul firewall, viene creato dinamicamente un percorso e l'OLIO è vuoto:

<#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
```

L'acquisizione sul 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
```

Il firewall crea le connessioni per ciascun ping, ma scarta automaticamente i pacchetti:

<#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 230.1
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.
May 17 2022 11:05:49: %FTD-6-302021: Teardown ICMP connection for faddr 192.168.103.62/6 gaddr 230.11.1
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 230.1
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.
May 17 2022 11:05:53: %FTD-6-302021: Teardown ICMP connection for faddr 192.168.103.62/6 gaddr 230.11.1
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
```

Nota: l'acquisizione drop ASP LINA non visualizza i pacchetti eliminati

---

L'indicazione principale delle perdite di pacchetti multicast è:

<#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

In FMC configurare un gruppo IGMP statico:

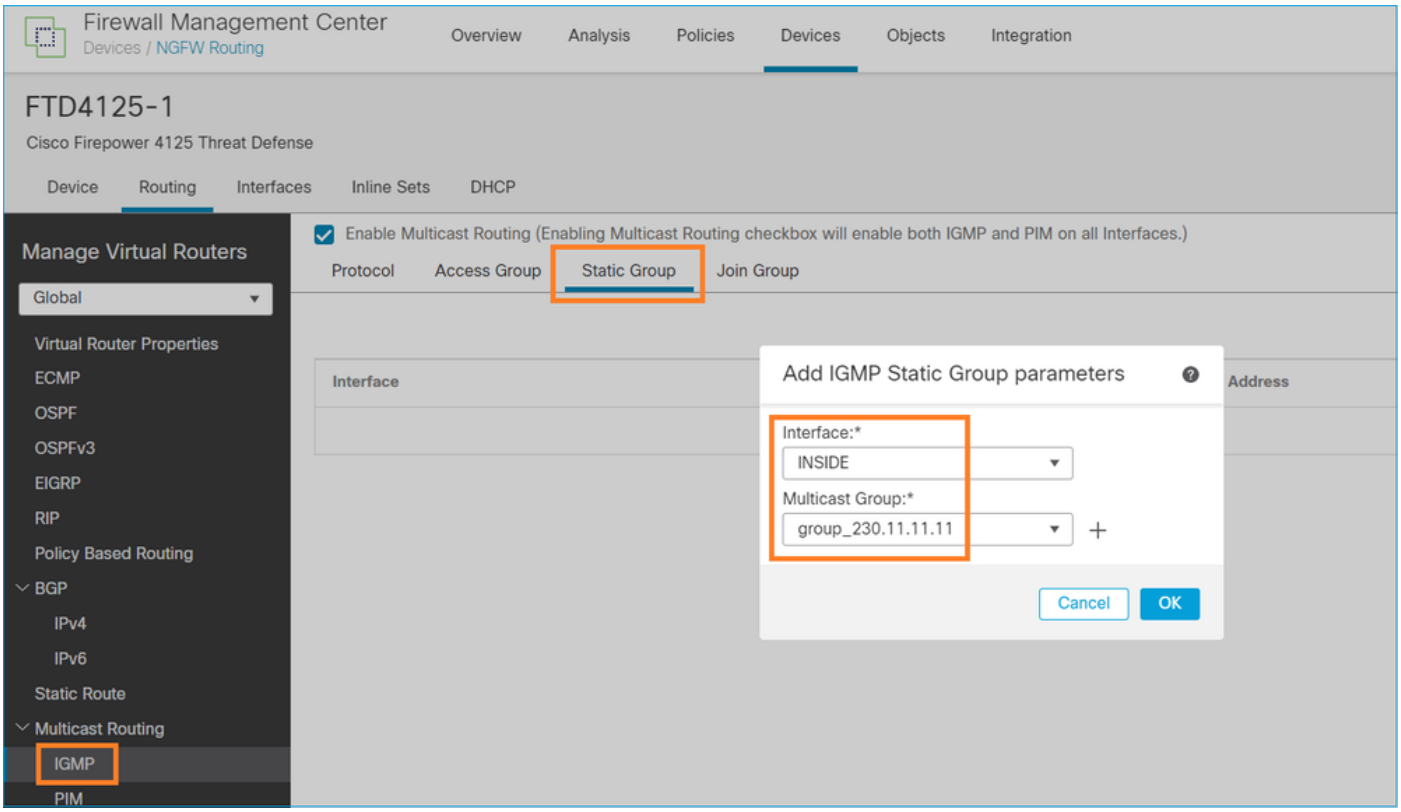

Questo è ciò che viene distribuito in background:

<#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
```

Il ping ha esito negativo, ma il traffico multicast ICMP viene ora inoltrato attraverso il 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
```

Nota: la traccia del pacchetto mostra un output errato (l'interfaccia in entrata è la stessa dell'uscita). Per ulteriori informazioni, consultare l'ID bug 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
```

Suggerimento: è possibile eseguire il ping con il timeout 0 dall'host di origine e controllare i contatori mfib del 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
```

join-group igmp

Sul telecomando di FMC, eseguire la configurazione del gruppo statico precedentemente configurato e configurare un gruppo di join IGMP:

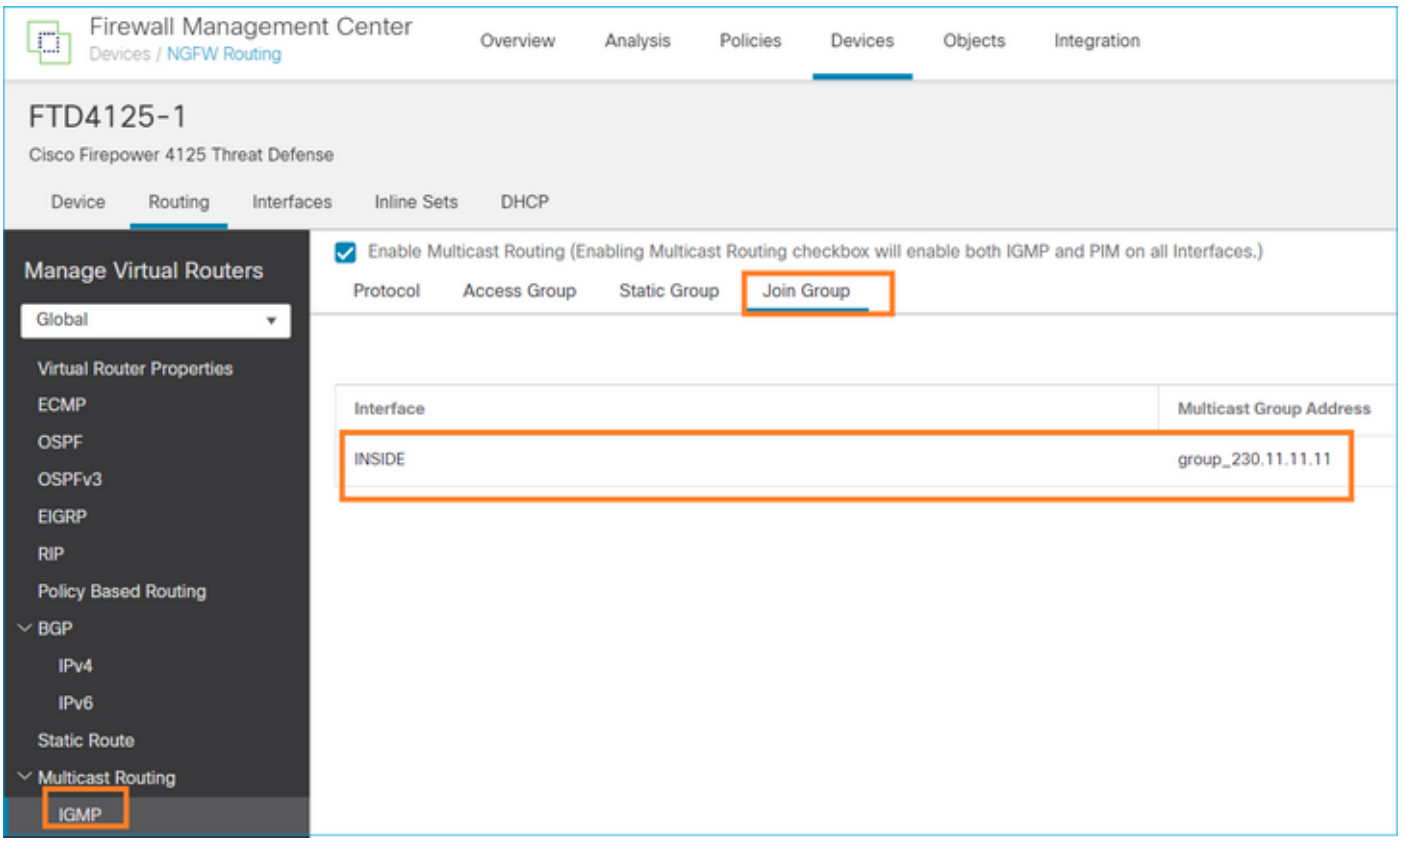

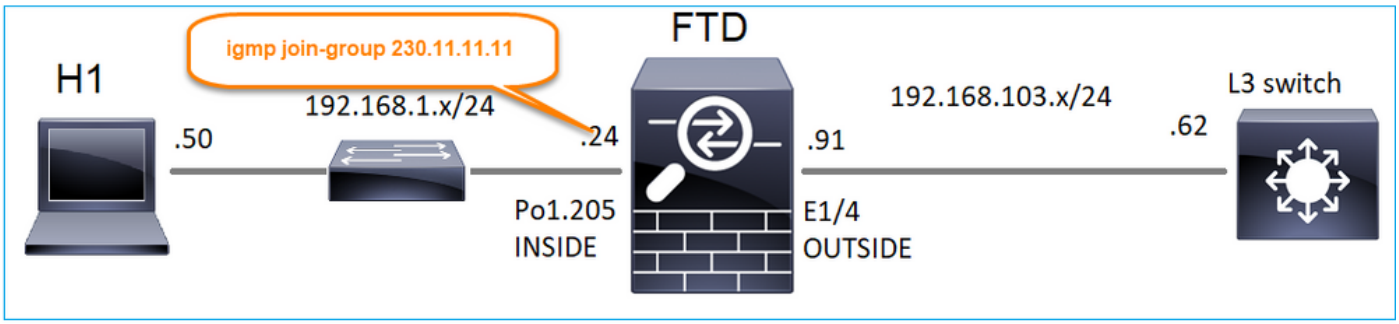

La configurazione distribuita:

<#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
```

Il gruppo 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
```

Dall'host di origine, provare il primo test multicast ICMP verso 230.11.11.11 IP:

<#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
```

Nota: se non si visualizzano tutte le risposte, controllare l'ID bug Cisco CSCvm90069.

Task 4 - Configurazione del routing multicast degli stub IGMP

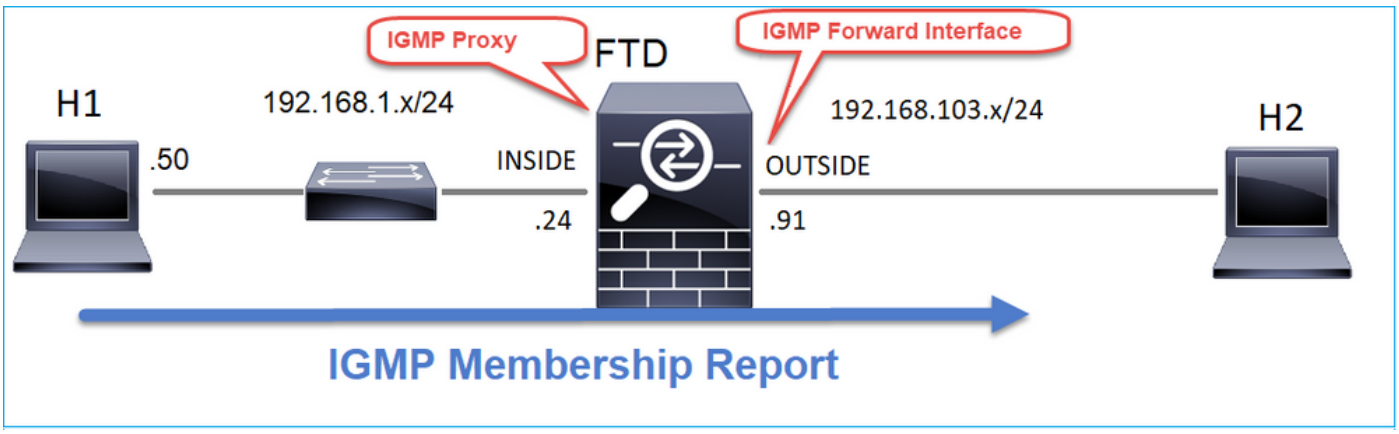

Configurare il routing multicast stub su FTD in modo che i messaggi IGMP Membership Report ricevuti sull'interfaccia INSIDE vengano inoltrati all'interfaccia OUTSIDE.

Soluzione

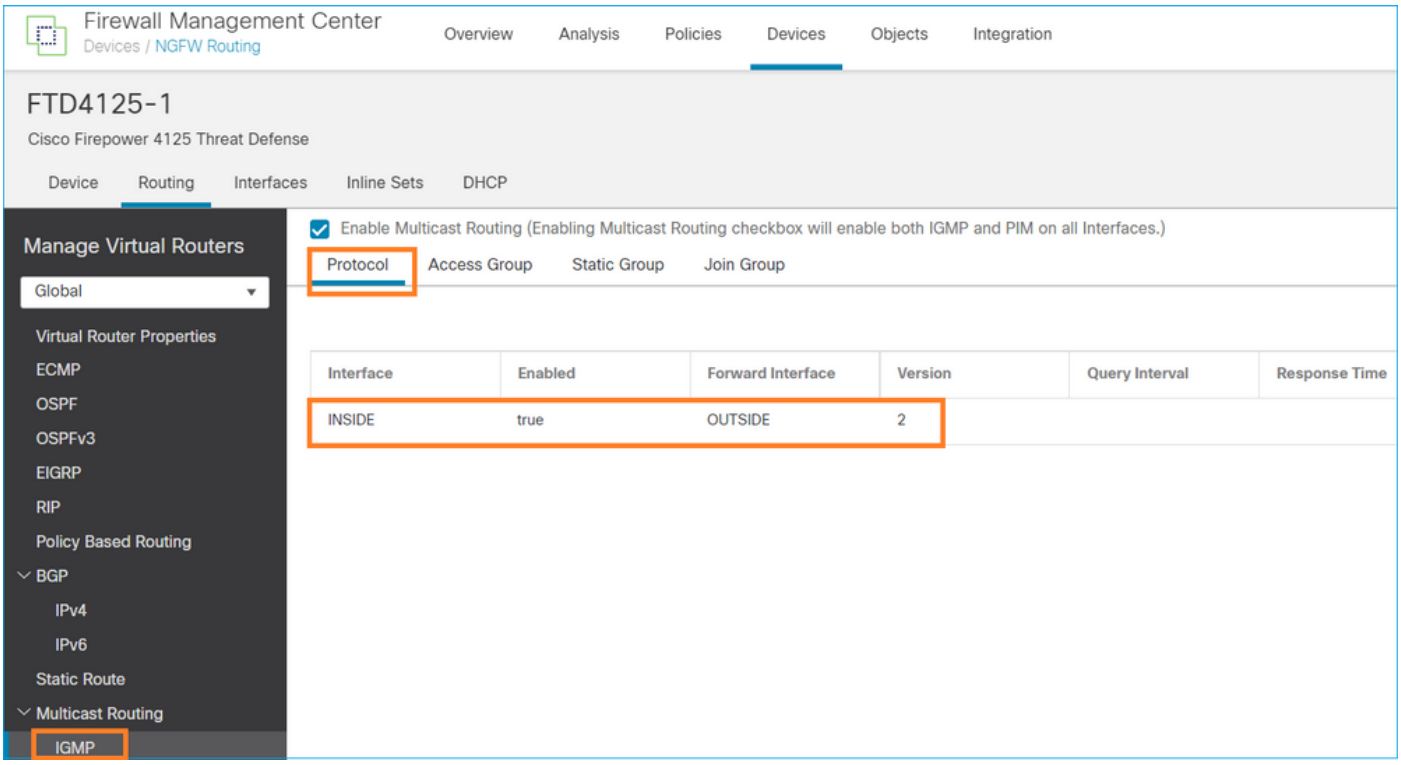

La configurazione distribuita:

<#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
```

Verifica

Abilita acquisizioni su 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
```

Verifica

Per forzare un report di appartenenza IGMP, è possibile utilizzare un'applicazione come VLC:

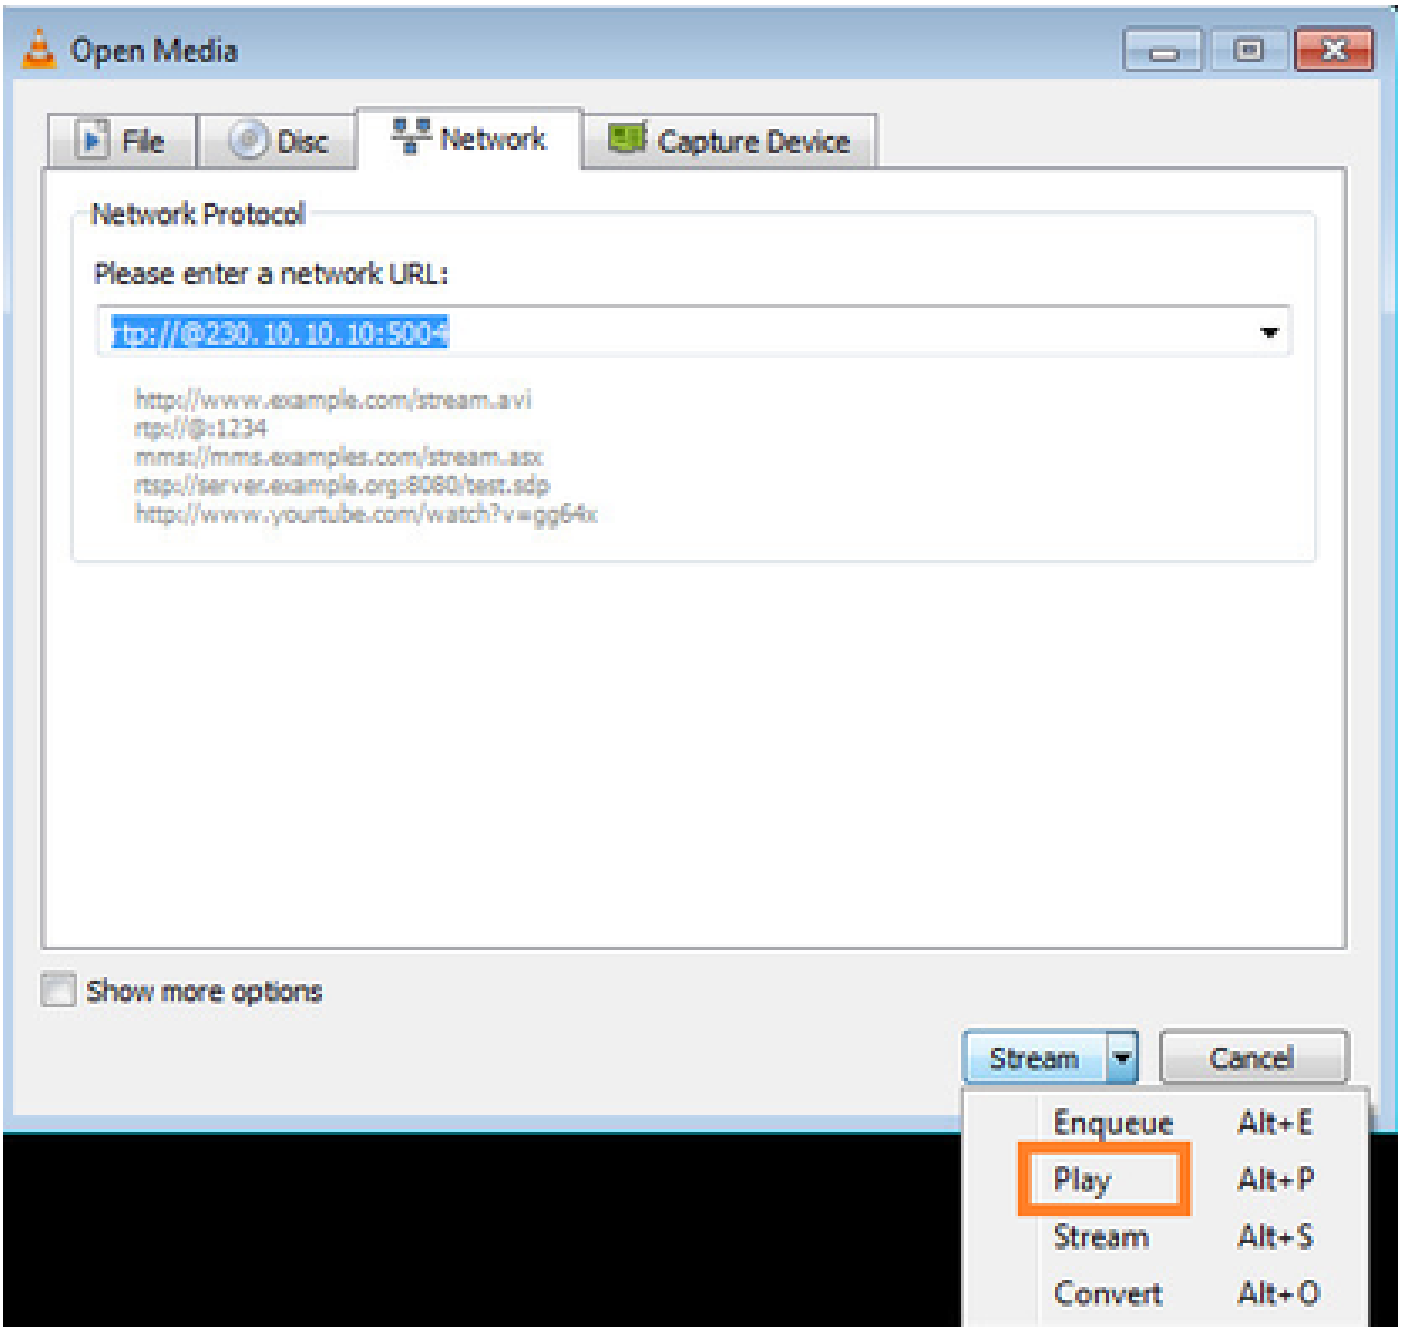

L'FTD proxy i pacchetti IGMP:

<#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
```

L'FTD modifica l'IP di origine:

<#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 si controlla il cappuccio in Wireshark, si osserverà che il pacchetto è stato completamente rigenerato dal firewall (l'identificazione IP cambia).

Viene creata una voce gruppo nell'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
```

Il firewall FTD crea due connessioni del control plane:

<#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
```

Traccia del primo pacchetto:

<#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
```

Problemi noti

Filtra il traffico multicast nelle zone di destinazione

Non è possibile specificare un'area di sicurezza di destinazione per la regola dei criteri di controllo di accesso corrispondente al traffico multicast:

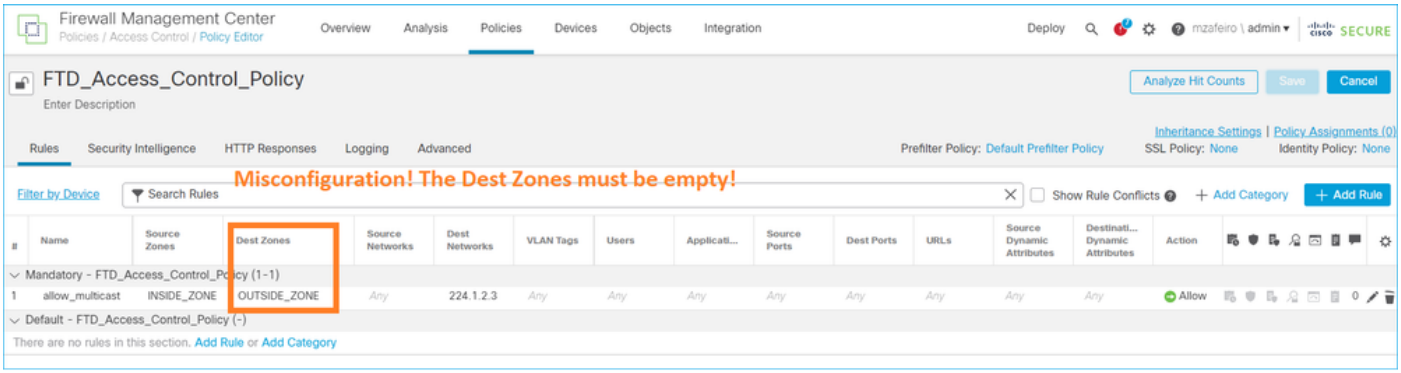

Questo è documentato anche nel manuale per l'utente del CCP:

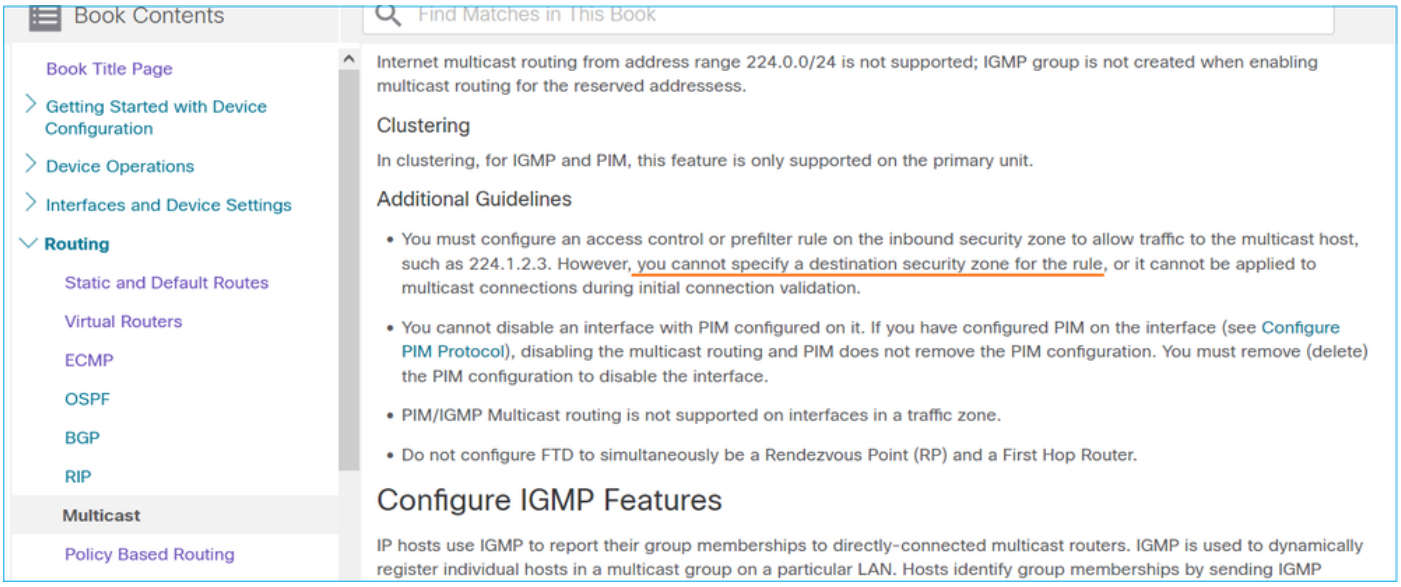

I report IGMP vengono rifiutati dal firewall quando viene superato il limite dell'interfaccia IGMP

Per impostazione predefinita, il firewall consente un massimo di 500 join attivi correnti (report) su un'interfaccia. Se questa soglia viene superata, il firewall ignora i rapporti IGMP aggiuntivi in arrivo provenienti dai ricevitori multicast.

Per controllare il limite IGMP e i join attivi, eseguire il comando show igmp interface namese:

<#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)
```

Il comando IGMP debug igmp visualizza questo output:

<#root>

```
asa#
debug igmp

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

Versioni software con la correzione dell'ID bug Cisco CSCvw60976 consente agli utenti di configurare fino a 5000 gruppi per interfaccia.

Il firewall ignora i report IGMP per l'intervallo di indirizzi 232.x.x.x/8

L'intervallo di indirizzi 232.x.x.x/8 deve essere utilizzato con SSM (Source Specific Multicast). Il firewall non supporta la funzionalità multicast (SSM) specifico dell'origine PIM e la configurazione correlata.

Il comando IGMP debug igmp visualizza questo output:

<#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 bug Cisco CSCsr53916

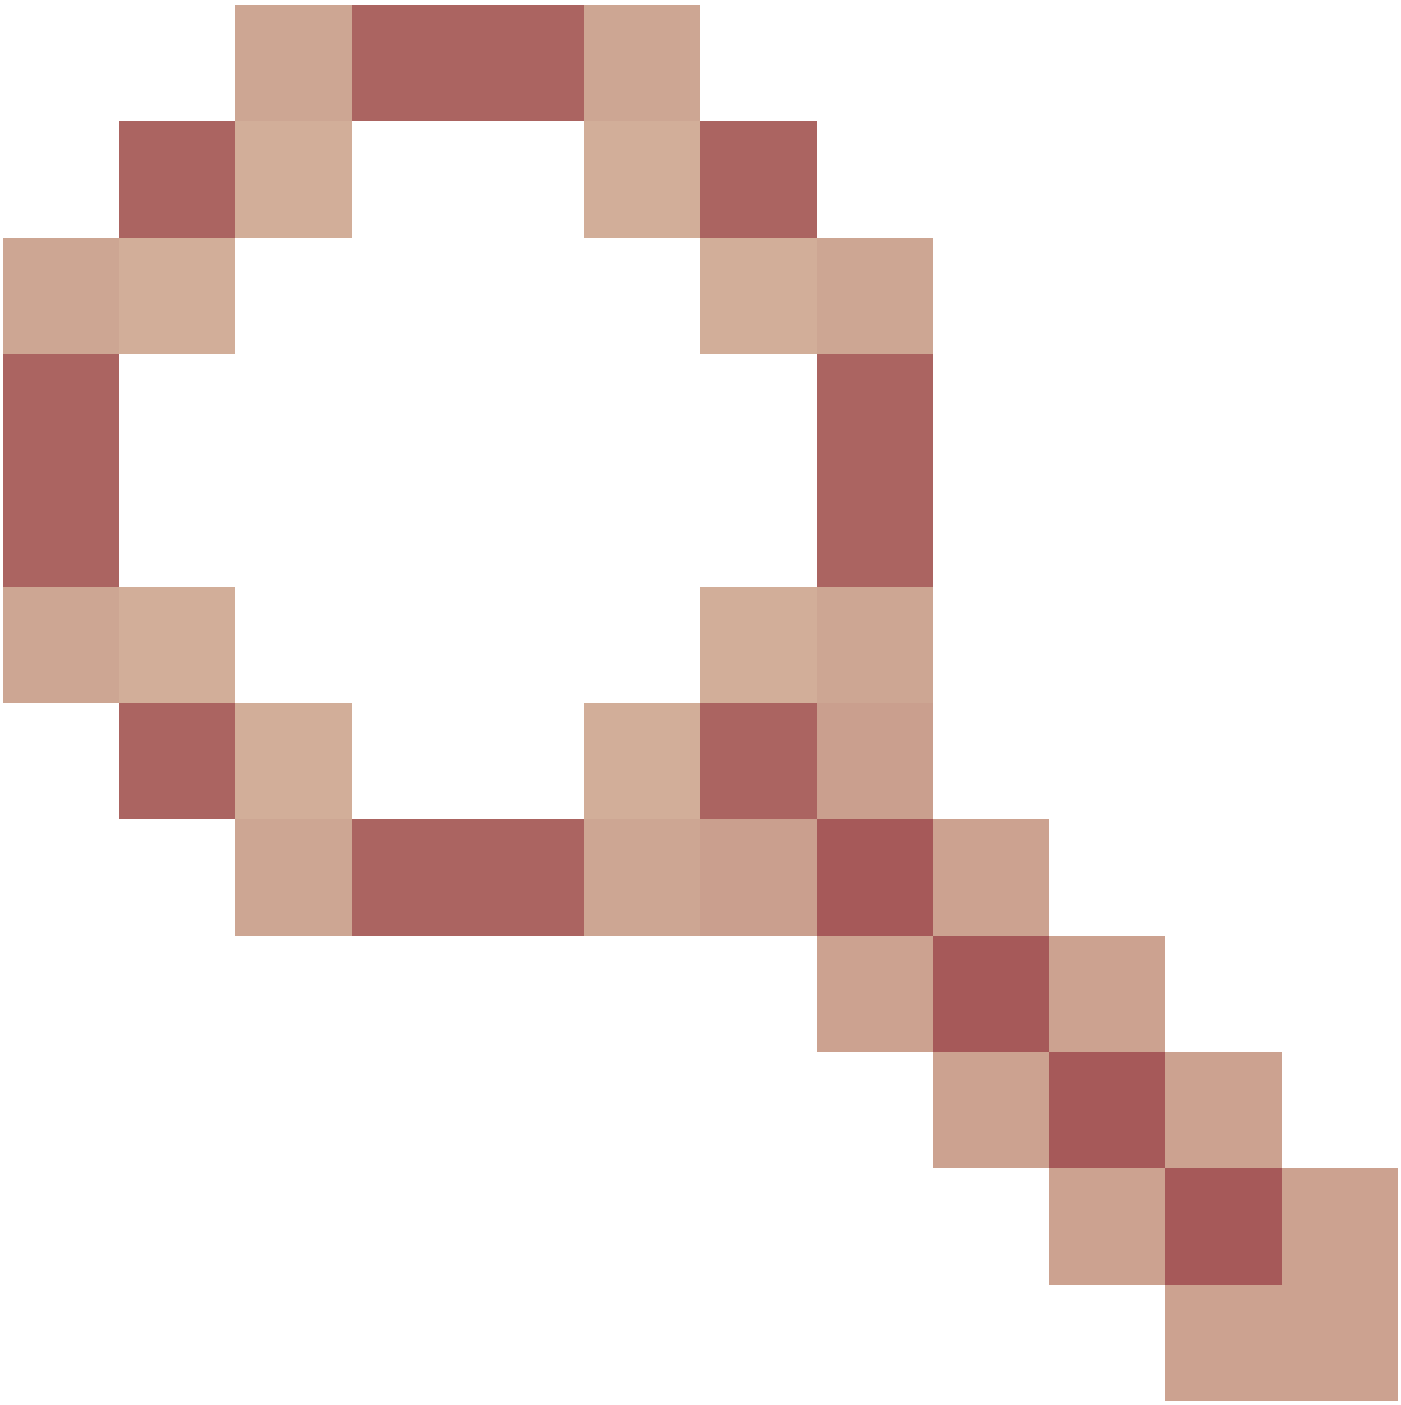

tiene traccia del miglioramento per il supporto dell'intervallo SSM.

Informazioni correlate

- Multicast Routing per Firepower Threat Defense
- Risoluzione dei problemi di Firepower Threat Defense e ASA Multicast PIM