Risoluzione dei problemi di Firepower Threat Defense Routing

Sommario

Introduzione

Questo documento descrive come Firepower Threat Defense (FTD) inoltra i pacchetti e implementa vari concetti di routing.

Prerequisiti

Requisiti

- Conoscenze base di routing

Componenti usati

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

hardware:

- Cisco Firepower 41xx Threat Defense versione 7.1.x
- Firepower Management Center (FMC) versione 7.1.x

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.

Premesse

Meccanismi di inoltro pacchetti FTD

FTD è un'immagine software unificata costituita da 2 motori principali:

- Datapath engine (LINA)
- Motore Snort

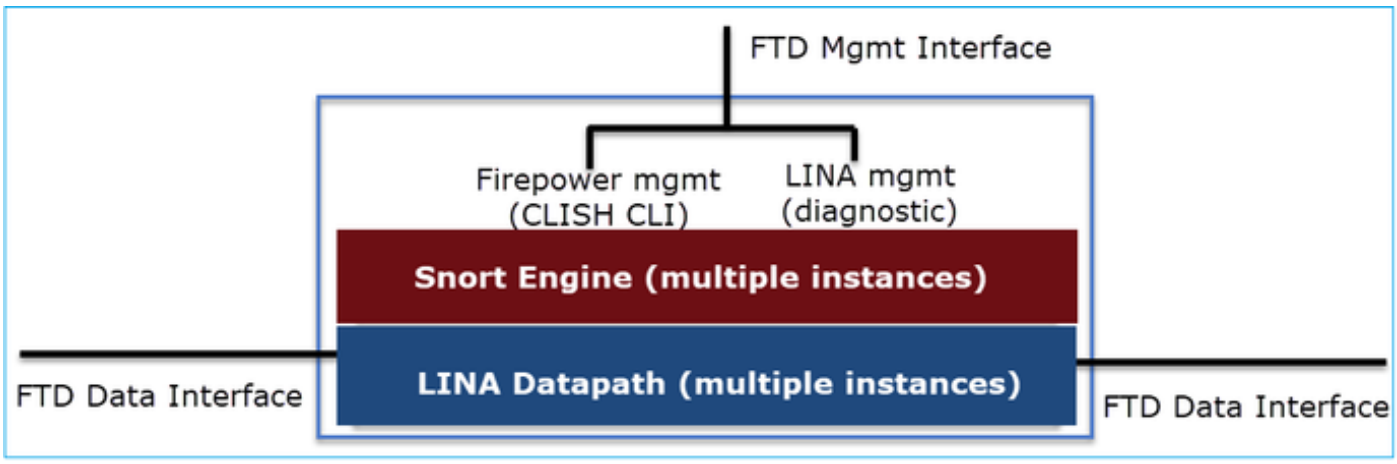

Il datapath e il motore di snort sono le parti principali del piano dati dell'FTD.

Il meccanismo di inoltro del piano dati FTD dipende dalla modalità interfaccia. L'immagine seguente riepiloga le varie modalità di interfaccia e le modalità di distribuzione FTD:

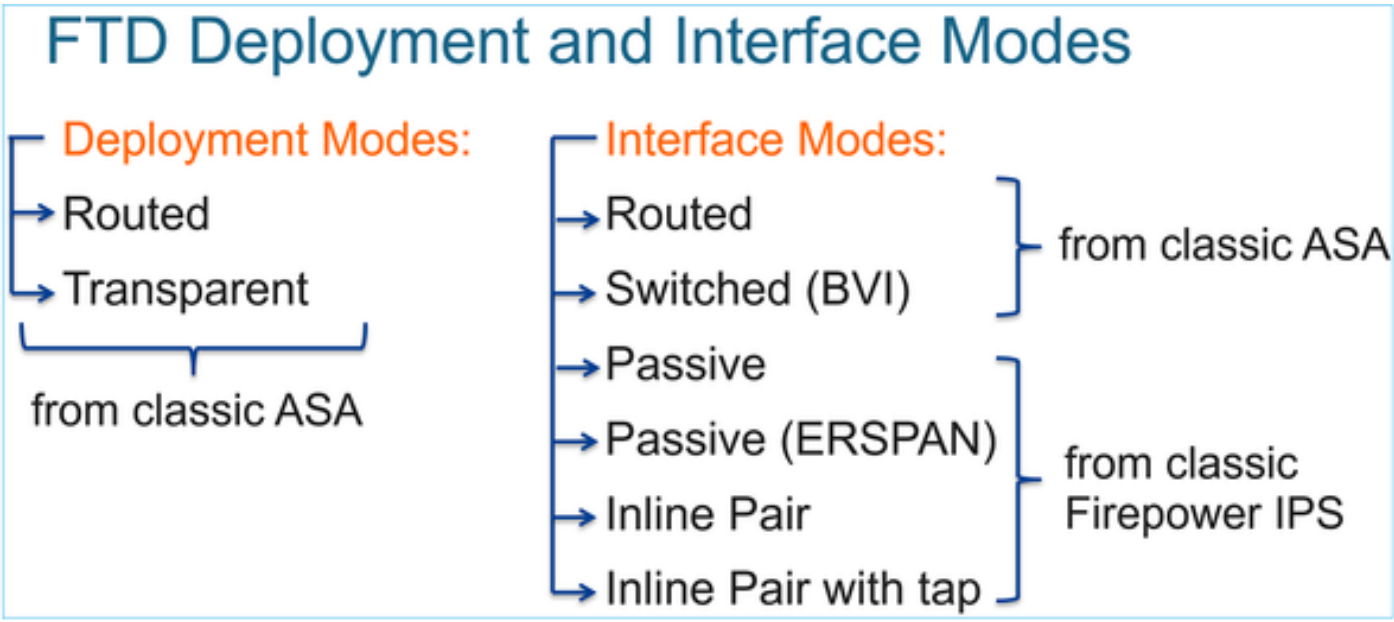

La tabella riepiloga il modo in cui l'FTD inoltra i pacchetti nel piano dati in base alla modalità interfaccia. I meccanismi di inoltro sono elencati in ordine di preferenza:

FTD Deployment mode	FTD Interface mode	Forwarding Mechanism
Routed	Routed	Packet forwarding based on the following order: 1. Connection lookup 2. Nat lookup (xlate) 3. Policy Based Routing (PBR) 4. Global routing table lookup
Routed or Transparent	Switched (BVI)	1. NAT lookup 2. Destination MAC Address L2 Lookup*
Routed or Transparent	Inline Pair	The packet will be forwarded based on the pair configuration.
Routed or Transparent	Inline Pair with Tap	The original packet will be forwarded based on the pair configuration. The copy of the packet will be dropped internally
Routed or Transparent	Passive	The packet is dropped internally
Routed	Passive (ERSPAN)	The packet is dropped internally

* Un FTD in modalità trasparente esegue una ricerca route in alcune situazioni:

MAC Address vs. Route Lookups

For traffic within a bridge group, the outgoing interface of a packet is determined by performing a destination MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for the following situations:

- Traffic originating on the Firepower Threat Defense device—Add a default/static route on the Firepower Threat Defense device for traffic destined for a remote network where a syslog server, for example, is located.
- Voice over IP (VoIP) and TFTP traffic, and the endpoint is at least one hop away—Add a static route on the Firepower Threat Defense device for traffic destined for the remote endpoint so that secondary connections are successful. The Firepower Threat Defense device creates a temporary "pinhole" in the access control policy to allow the secondary connection; and because the connection might use a different set of IP addresses than the primary connection, the Firepower Threat Defense device needs to perform a route lookup to install the pinhole on the correct interface.

  Affected applications include:

  - H.323
  - RTSP
  - SIP
  - Skinny (SCCP)
  - SQL*Net
  - SunRPC
  - TFTP
- Traffic at least one hop away for which the Firepower Threat Defense device performs NAT—Configure a static route on the Firepower Threat Defense device for traffic destined for the remote network. You also need a static route on the upstream router for traffic destined for the mapped addresses to be sent to the Firepower Threat Defense device.

Consultare la guida FMC per ulteriori dettagli.

A partire dalla versione 6.2.x, l'FTD supporta il Routing e Bridging integrato (IRB):

FTD Integrated Routing and Bridging (IRB)

- Available as from 6.2.x
- Allows an FTD in **Routed mode** to have multiple interfaces (up to 64) to be part of the **same VLAN** and perform L2 switching between them
- BVI-to-Routed or BVI-to-BVI Routing is allowed

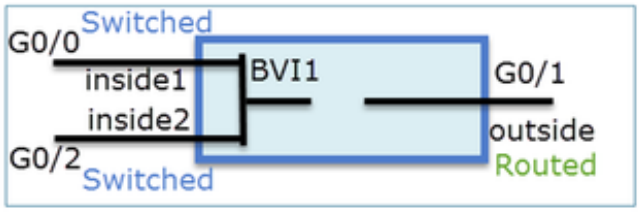

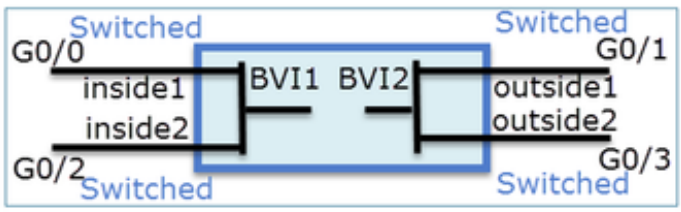

Comandi di verifica BVI:

Verification commands

```
firepower# show bridge-group
```

```
firepower# show ip
Interface                Name                 IP address      Subnet mask     Method
GigabitEthernet0/0       VLAN1576_G0-0        203.0.113.1     255.255.255.0   manual
GigabitEthernet0/1       VLAN1577_G0-1        192.168.1.15    255.255.255.0   manual
GigabitEthernet0/2       VLAN1576_G0-2        203.0.113.1     255.255.255.0   manual
GigabitEthernet0/4.100   SUB1                 203.0.113.1     255.255.255.0   manual
BVI1                     LAN                  203.0.113.1     255.255.255.0   manual
BVI2                     LAN2                 192.168.1.15    255.255.255.0   manual
```

- BVI nameif is used in L3 Routing configuration

```
firepower# show run route
route LAN 1.1.1.0 255.255.255.0 203.0.113.5 1
```

- BVI member nameif is used in policies like NAT configuration

```
firepower# show run nat
nat (VLAN1576_G0-0,VLAN1577_G0-1) source dynamic any interface
nat (VLAN1576_G0-2,VLAN1577_G0-1) source dynamic any interface
```

Punto chiave

Per le interfacce di routing o BVI (IRB), l'inoltro dei pacchetti si basa su questo ordine:

- Ricerca connessione
- Ricerca NAT (destinazione NAT, nota anche come UN-NAT)
- Policy-Based Routing (PBR)
- Ricerca nella tabella di routing globale

E la fonte NAT?

Il NAT di origine viene controllato dopo la ricerca di routing globale.

Nel prosieguo di questo documento viene trattata in modo specifico la modalità dell'interfaccia di routing.

Comportamento di routing del piano dati (LINA)

In modalità di interfaccia indirizzata, FTD LINA inoltra i pacchetti in 2 fasi:

Fase 1 - Determinazione dell'interfaccia in uscita

Fase 2 - Selezione hop successivo

Supponiamo di avere questa topologia:

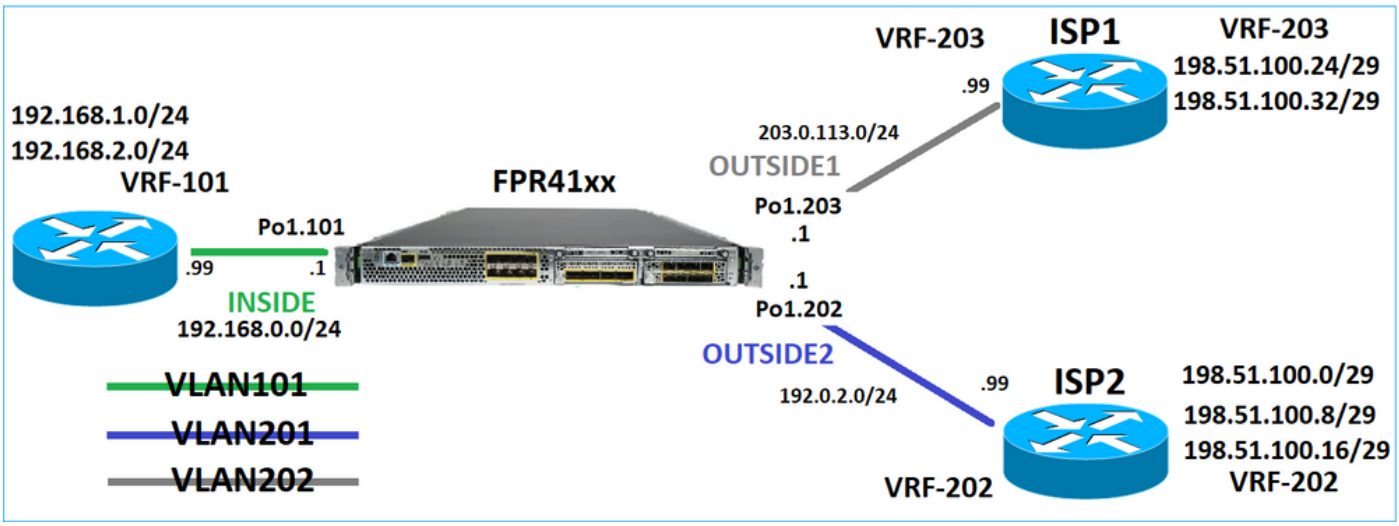

E questo progetto di routing:

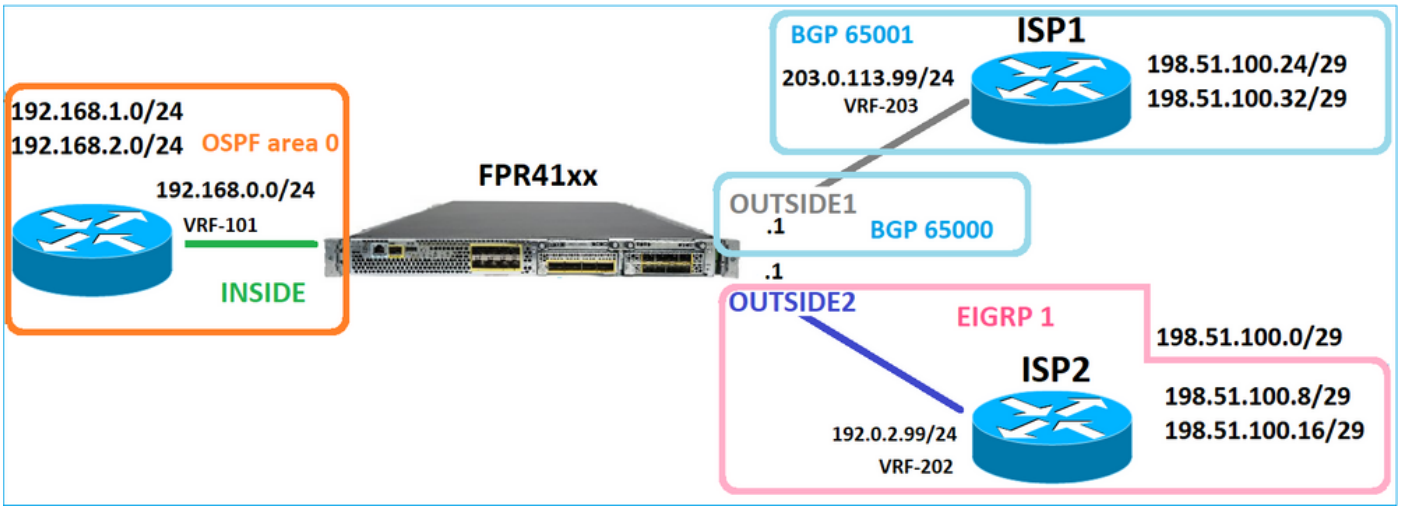

La configurazione del routing FTD:

```
firepower# show run router
router ospf 1
network 192.168.0.0 255.255.255.0 area 0
log-adj-changes
!
router bgp 65000
bgp log-neighbor-changes
bgp router-id vrf auto-assign
address-family ipv4 unicast
neighbor 203.0.113.99 remote-as 65001
neighbor 203.0.113.99 ebgp-multihop 255
neighbor 203.0.113.99 transport path-mtu-discovery disable
neighbor 203.0.113.99 activate
no auto-summary
no synchronization
exit-address-family
!
router eigrp 1
no default-information in
no default-information out
no eigrp log-neighbor-warnings
```

```
no eigrp log-neighbor-changes
network 192.0.2.0 255.255.255.0
!
firepower# show run route
route OUTSIDE2 198.51.100.0 255.255.255.248 192.0.2.99 1
```

Base informazioni ciclo FTD (RIB, Routing Information Base) - Control Plane:

```
firepower# show route | begin Gate
Gateway of last resort is not set

C 192.0.2.0 255.255.255.0 is directly connected, OUTSIDE2
L 192.0.2.1 255.255.255.255 is directly connected, OUTSIDE2
C 192.168.0.0 255.255.255.0 is directly connected, INSIDE
L 192.168.0.1 255.255.255.255 is directly connected, INSIDE
O 192.168.1.1 255.255.255.255
[110/11] via 192.168.0.99, 01:11:25, INSIDE
O 192.168.2.1 255.255.255.255
[110/11] via 192.168.0.99, 01:11:15, INSIDE
S 198.51.100.0 255.255.255.248 [1/0] via 192.0.2.99, OUTSIDE2
D 198.51.100.8 255.255.255.248
[90/130816] via 192.0.2.99, 01:08:11, OUTSIDE2
D 198.51.100.16 255.255.255.248
[90/130816] via 192.0.2.99, 01:08:04, OUTSIDE2
B 198.51.100.24 255.255.255.248 [20/0] via 203.0.113.99, 00:28:29
B 198.51.100.32 255.255.255.248 [20/0] via 203.0.113.99, 00:28:16
C 203.0.113.0 255.255.255.0 is directly connected, OUTSIDE1
L 203.0.113.1 255.255.255.255 is directly connected, OUTSIDE1
```

Tabella di routing ASP (Accelerated Security Path) FTD corrispondente - Piano dati:

```
firepower# show asp table routing
route table timestamp: 91
in 169.254.1.1 255.255.255.255 identity
in 192.168.0.1 255.255.255.255 identity
in 192.0.2.1 255.255.255.255 identity
in 192.168.1.1 255.255.255.255 via 192.168.0.99, INSIDE
in 192.168.2.1 255.255.255.255 via 192.168.0.99, INSIDE
in 203.0.113.1 255.255.255.255 identity
in 169.254.1.0 255.255.255.248 nlp_int_tap
in 198.51.100.0 255.255.255.248 via 192.0.2.99, OUTSIDE2
in 198.51.100.8 255.255.255.248 via 192.0.2.99, OUTSIDE2
in 198.51.100.16 255.255.255.248 via 192.0.2.99, OUTSIDE2
in 198.51.100.24 255.255.255.248 via 203.0.113.99 (unresolved, timestamp: 89)
in 198.51.100.32 255.255.255.248 via 203.0.113.99 (unresolved, timestamp: 90)
in 192.168.0.0 255.255.255.0 INSIDE
in 192.0.2.0 255.255.255.0 OUTSIDE2
in 203.0.113.0 255.255.255.0 OUTSIDE1
in ff02::1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in ff02::1:ff01:3 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in ff02::1:ff00:1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in fe80::200:ff:fe01:3 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in fd00:0:0:1::1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
```

```
in fd00:0:0:1:: ffff:ffff:ffff:ffff:: nlp_int_tap
out 255.255.255.255 255.255.255.255 OUTSIDE1
out 203.0.113.1 255.255.255.255 OUTSIDE1
out 203.0.113.0 255.255.255.0 OUTSIDE1
out 224.0.0.0 240.0.0.0 OUTSIDE1
out 255.255.255.255 255.255.255.255 OUTSIDE2
out 192.0.2.1 255.255.255.255 OUTSIDE2
out 198.51.100.0 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 198.51.100.8 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 198.51.100.16 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 192.0.2.0 255.255.255.0 OUTSIDE2
out 224.0.0.0 240.0.0.0 OUTSIDE2
out 255.255.255.255 255.255.255.255 INSIDE
out 192.168.0.1 255.255.255.255 INSIDE
out 192.168.1.1 255.255.255.255 via 192.168.0.99, INSIDE
out 192.168.2.1 255.255.255.255 via 192.168.0.99, INSIDE
out 192.168.0.0 255.255.255.0 INSIDE
out 224.0.0.0 240.0.0.0 INSIDE
out 255.255.255.255 255.255.255.255 cmi_mgmt_int_tap
out 224.0.0.0 240.0.0.0 cmi_mgmt_int_tap
out 255.255.255.255 255.255.255.255 ha_ctl_nlp_int_tap
out 224.0.0.0 240.0.0.0 ha_ctl_nlp_int_tap
out 255.255.255.255 255.255.255.255 ccl_ha_nlp_int_tap
out 224.0.0.0 240.0.0.0 ccl_ha_nlp_int_tap
out 255.255.255.255 255.255.255.255 nlp_int_tap
out 169.254.1.1 255.255.255.255 nlp_int_tap
out 169.254.1.0 255.255.255.248 nlp_int_tap
out 224.0.0.0 240.0.0.0 nlp_int_tap
out fd00:0:0:1::1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff nlp_int_tap
out fd00:0:0:1:: ffff:ffff:ffff:ffff:: nlp_int_tap
out fe80:: ffc0:: nlp_int_tap
out ff00:: ff00:: nlp_int_tap
out 0.0.0.0 0.0.0.0 via 0.0.0.0, identity
out :: :: via 0.0.0.0, identity
```

Punti chiave

L'FTD (simile all'ASA, Adaptive Security Appliance) determina prima l'interfaccia di uscita (uscita) di un pacchetto (per questo motivo, controlla le voci "in" della tabella di routing ASP). Quindi, per l'interfaccia determinata, tenta di trovare l'hop successivo (a tale scopo, cerca le voci 'out' della tabella di routing ASP). Ad esempio:

```
firepower# show asp table routing | include in.*198.51.100.0
in 198.51.100.0 255.255.255.248 via 192.0.2.99, OUTSIDE2
firepower#
firepower# show asp table routing | include out.*OUTSIDE2
out 255.255.255.255 255.255.255.255 OUTSIDE2
out 192.0.2.1 255.255.255.255 OUTSIDE2
out 198.51.100.0 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 198.51.100.8 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 198.51.100.16 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 192.0.2.0 255.255.255.0 OUTSIDE2
out 224.0.0.0 240.0.0.0 OUTSIDE2
```

Infine, per l'hop successivo risolto, LINA controlla la cache ARP per verificare la presenza di un'adiacenza valida.

Lo strumento di traccia dei pacchetti FTD conferma questo processo:

```
firepower# packet-tracer input INSIDE icmp 192.168.1.1 8 0 198.51.100.1

Phase: 1
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 7582 ns
Config:
Implicit Rule
Additional Information:
MAC Access list

Phase: 2
Type: INPUT-ROUTE-LOOKUP
Subtype: Resolve Egress Interface
Result: ALLOW
Elapsed time: 8474 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

Phase: 3
Type: ACCESS-LIST
Subtype: log
Result: ALLOW
Elapsed time: 5017 ns
Config:
access-group CSM_FW_ACL_ global
access-list CSM_FW_ACL_ advanced permit ip any any rule-id 268434433
access-list CSM_FW_ACL_ remark rule-id 268434433: ACCESS POLICY: mzafeiro_empty - Default
access-list CSM_FW_ACL_ remark rule-id 268434433: L4 RULE: DEFAULT ACTION RULE
Additional Information:
This packet will be sent to snort for additional processing where a verdict will be reached

Phase: 4
Type: CONN-SETTINGS
Subtype:
Result: ALLOW
Elapsed time: 5017 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: 5
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 5017 ns
Config:
```

```
Additional Information:

Phase: 6
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 5017 ns
Config:
Additional Information:

Phase: 7
Type: INSPECT
Subtype: np-inspect
Result: ALLOW
Elapsed time: 57534 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: 8
Type: INSPECT
Subtype: np-inspect
Result: ALLOW
Elapsed time: 3122 ns
Config:
Additional Information:

Phase: 9
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 29882 ns
Config:
Additional Information:

Phase: 10
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 446 ns
Config:
Additional Information:

Phase: 11
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Elapsed time: 20962 ns
Config:
Additional Information:
New flow created with id 178, packet dispatched to next module

Phase: 12
Type: EXTERNAL-INSPECT
Subtype:
Result: ALLOW
Elapsed time: 20070 ns
```

```
Config:
Additional Information:
Application: 'SNORT Inspect'

Phase: 13
Type: SNORT
Subtype:
Result: ALLOW
Elapsed time: 870592 ns
Config:
Additional Information:
Snort Trace:
Packet: ICMP
Session: new snort session
Snort id 1, NAP id 1, IPS id 0, Verdict PASS
Snort Verdict: (pass-packet) allow this packet

Phase: 14
Type: INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Resolve Preferred Egress interface
Result: ALLOW
Elapsed time: 6244 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

Phase: 15
Type: ADJACENCY-LOOKUP
Subtype: Resolve Nexthop IP address to MAC
Result: ALLOW
Elapsed time: 1784 ns
Config:
Additional Information:
Found adjacency entry for Next-hop 192.0.2.99 on interface OUTSIDE2
Adjacency :Active
MAC address 4c4e.35fc.fcd8 hits 5 reference 1

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE2(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 1046760 ns
```

La tabella ARP FTD come viene visualizzata nel Piano di controllo:

```
firepower# show arp
OUTSIDE1 203.0.113.99 4c4e.35fc.fcd8 3051
OUTSIDE2 192.0.2.99 4c4e.35fc.fcd8 5171
```

Per forzare la risoluzione ARP:

```
firepower# ping 192.168.0.99
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.0.99, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
firepower# show arp
INSIDE 192.168.0.99 4c4e.35fc.fcd8 45
OUTSIDE1 203.0.113.99 4c4e.35fc.fcd8 32
OUTSIDE2 192.0.2.99 4c4e.35fc.fcd8 1
```

Tabella ARP FTD come viene visualizzata nel piano dati:

```
firepower# show asp table arp

Context: single_vf, Interface: OUTSIDE1
203.0.113.99 Active 4c4e.35fc.fcd8 hits 2 reference 1

Context: single_vf, Interface: OUTSIDE2
192.0.2.99 Active 4c4e.35fc.fcd8 hits 5 reference 0

Context: single_vf, Interface: INSIDE
192.168.0.99 Active 4c4e.35fc.fcd8 hits 5 reference 0

Context: single_vf, Interface: identity
:: Active 0000.0000.0000 hits 0 reference 0
0.0.0.0 Active 0000.0000.0000 hits 848 reference 0

Last clearing of hits counters: Never
```

Ordine delle operazioni FTD

Nell'immagine è illustrato l'ordine delle operazioni e la posizione in cui vengono eseguiti i controlli di instradamento ASP di input e output:

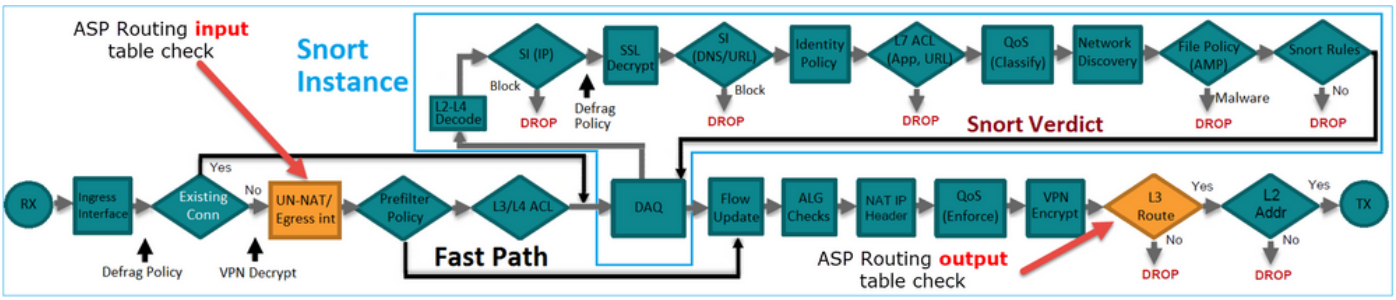

Configurazione

Caso 1 - Inoltro basato sulla ricerca della connessione

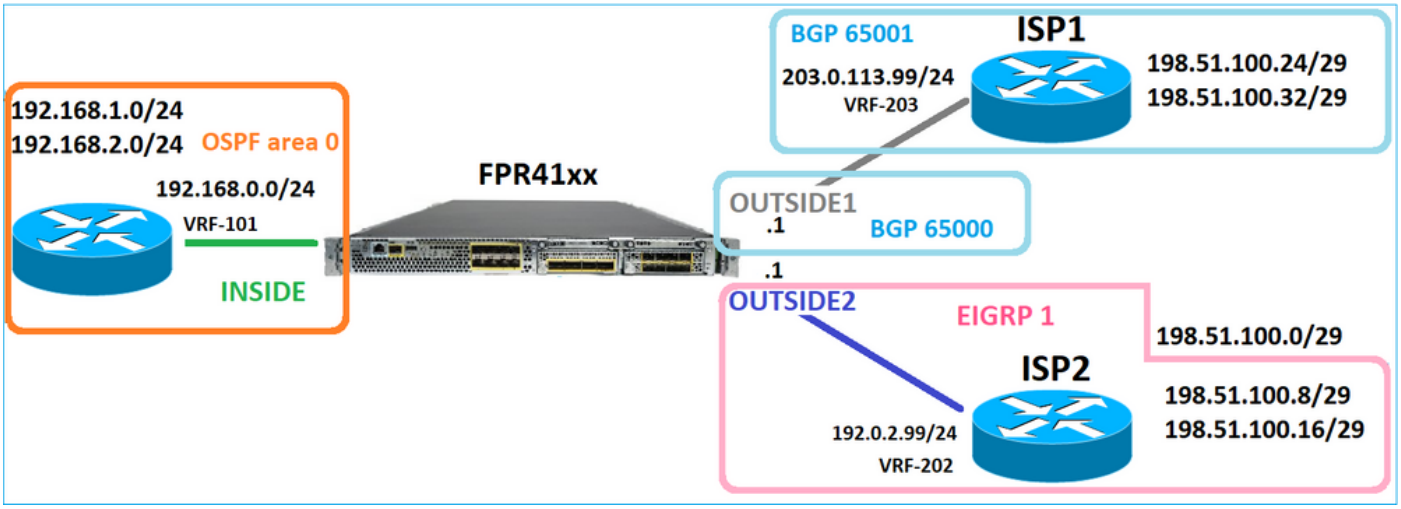

Come già accennato, il componente principale del motore LINA FTD è il processo Datapath (istanze multiple basate sul numero di core del dispositivo). Inoltre, il datapath (noto anche come Accelerated Security Path - ASP) è costituito da 2 percorsi:

1. Percorso lento = responsabile della creazione della nuova connessione (popola il Percorso rapido).
2. Percorso rapido = Gestisce i pacchetti che appartengono a connessioni stabilite.

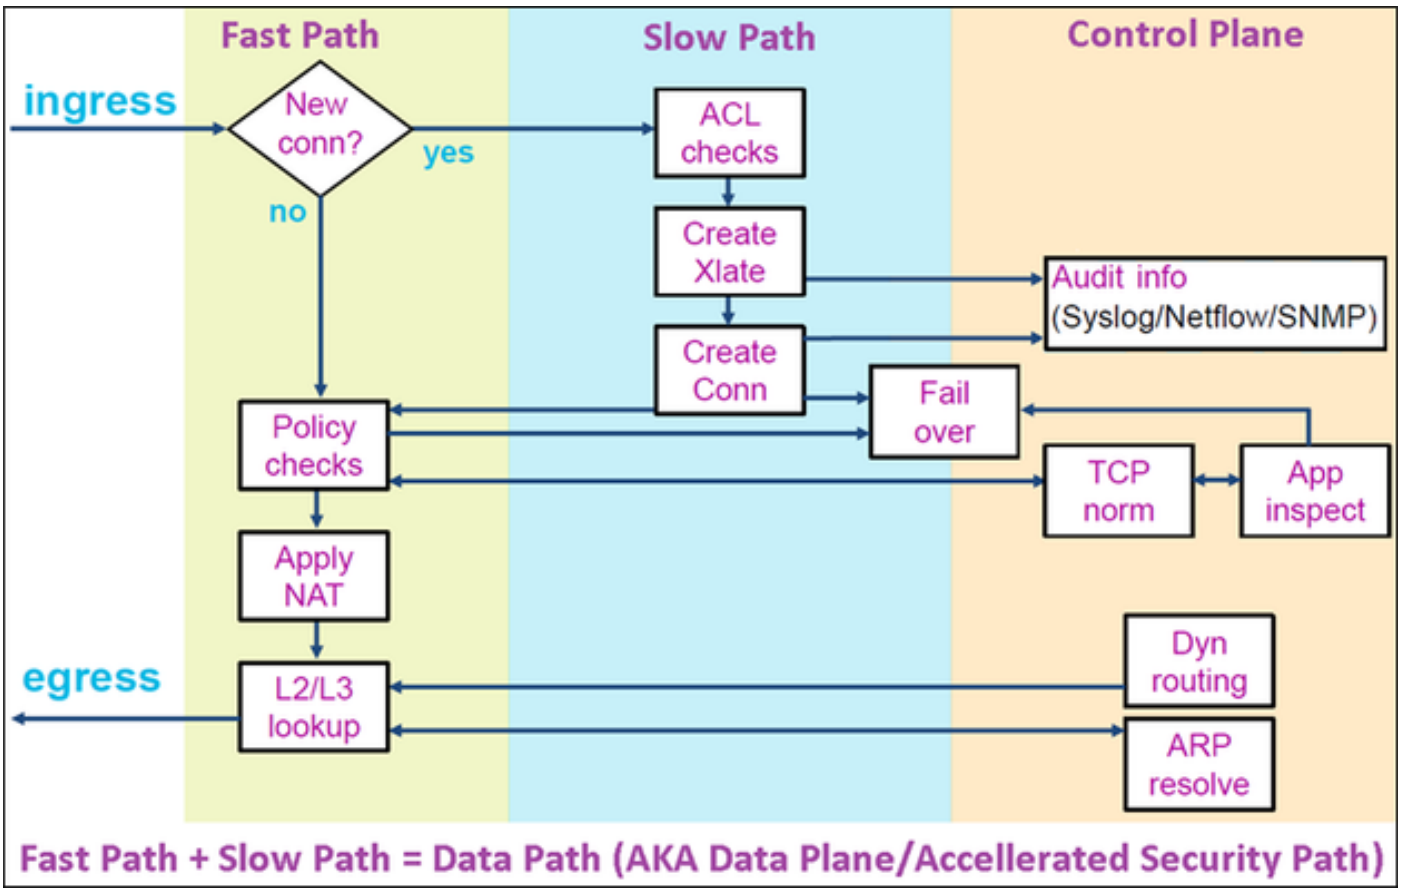

- Comandi quali show route e show arp mostrano il contenuto del Control Plane.
- D'altra parte, comandi quali show asp table routing e show asp table arp mostrano il contenuto di ASP (Datapath) che è ciò che viene effettivamente applicato.

Abilita acquisizione con traccia sull'interfaccia FTD INSIDE:

```
firepower# capture CAPI trace detail interface INSIDE match ip host 192.168.1.1 host 198.51.100.1
```

Aprire una sessione Telnet con l'FTD:

```
Router1# telnet 198.51.100.1 /vrf VRF-101 /source-interface lo1
Trying 198.51.100.1 ... Open
```

Le clip FTD mostrano i pacchetti dall'inizio della connessione (viene acquisito l'handshake a 3 vie TCP):

```
firepower# show capture CAPI

26 packets captured

1: 10:50:38.407190 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: S 1306692135:1306692135(0) w
2: 10:50:38.408929 802.1Q vlan#101 P0 198.51.100.1.23 > 192.168.1.1.57734: S 1412677784:1412677784(0) a
3: 10:50:38.409265 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: . ack 1412677785 win 4128
4: 10:50:38.409433 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: P 1306692136:1306692154(18)
5: 10:50:38.409845 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: . ack 1412677785 win 4128
6: 10:50:38.410135 802.1Q vlan#101 P0 198.51.100.1.23 > 192.168.1.1.57734: . ack 1306692154 win 4110
7: 10:50:38.411355 802.1Q vlan#101 P0 198.51.100.1.23 > 192.168.1.1.57734: P 1412677785:1412677797(12)
8: 10:50:38.413049 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: P 1306692154:1306692157(3) a
9: 10:50:38.413140 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: P 1306692157:1306692166(9) a
10: 10:50:38.414071 802.1Q vlan#101 P0 198.51.100.1.23 > 192.168.1.1.57734: . 1412677797:1412678322(525
...
```

Traccia il primo pacchetto (TCP SYN). Questo pacchetto passa attraverso il percorso lente LINA FTD e viene eseguita una ricerca di routing globale nel caso seguente:

```
firepower# show capture CAPI packet-number 1 trace

26 packets captured

   1: 10:50:38.407190 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: S 1306692135:1306692135(0
Phase: 1
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 4683 ns
Config:
Additional Information:
Forward Flow based lookup yields rule:
in id=0x1505f1d17940, priority=13, domain=capture, deny=false
hits=1783, user_data=0x1505f2096910, cs_id=0x0, l3_type=0x0
```

```
src mac=0000.0000.0000, mask=0000.0000.0000
dst mac=0000.0000.0000, mask=0000.0000.0000
input_ifc=INSIDE, output_ifc=any

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 4683 ns
Config:
Implicit Rule
Additional Information:
Forward Flow based lookup yields rule:
in id=0x1502a7ba4d40, priority=1, domain=permit, deny=false
hits=28, user_data=0x0, cs_id=0x0, l3_type=0x8
src mac=0000.0000.0000, mask=0000.0000.0000
dst mac=0000.0000.0000, mask=0100.0000.0000
input_ifc=INSIDE, output_ifc=any

Phase: 3
Type: INPUT-ROUTE-LOOKUP
Subtype: Resolve Egress Interface
Result: ALLOW
Elapsed time: 5798 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

Phase: 4
Type: ACCESS-LIST
Subtype: log
Result: ALLOW
Elapsed time: 3010 ns
Config:
access-group CSM_FW_ACL_ global
access-list CSM_FW_ACL_ advanced permit ip any any rule-id 268434433
access-list CSM_FW_ACL_ remark rule-id 268434433: ACCESS POLICY: mzafeiro_empty - Default
access-list CSM_FW_ACL_ remark rule-id 268434433: L4 RULE: DEFAULT ACTION RULE
Additional Information:
This packet will be sent to snort for additional processing where a verdict will be reached
Forward Flow based lookup yields rule:
in id=0x1505f1e2e980, priority=12, domain=permit, deny=false
hits=4, user_data=0x15024a56b940, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, ifc=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, ifc=any,, dscp=0x0, nsg_id=none
input_ifc=any, output_ifc=any

Phase: 5
Type: CONN-SETTINGS
Subtype:
Result: ALLOW
Elapsed time: 3010 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:
Forward Flow based lookup yields rule:
in id=0x1505f1f18bc0, priority=7, domain=conn-set, deny=false
```

```
hits=4, user_data=0x1505f1f13f70, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, dscp=0x0, nsg_id=none
input_ifc=INSIDE(vrfid:0), output_ifc=any

Phase: 6
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 3010 ns
Config:
Additional Information:
Forward Flow based lookup yields rule:
in id=0x15052e96b150, priority=0, domain=nat-per-session, deny=false
hits=125, user_data=0x0, cs_id=0x0, reverse, use_real_addr, flags=0x0, protocol=6
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, dscp=0x0, nsg_id=none
input_ifc=any, output_ifc=any

Phase: 7
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 3010 ns
Config:
Additional Information:
Forward Flow based lookup yields rule:
in id=0x1502a7bacde0, priority=0, domain=inspect-ip-options, deny=true
hits=19, user_data=0x0, cs_id=0x0, reverse, flags=0x0, protocol=0
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, dscp=0x0, nsg_id=none
input_ifc=INSIDE(vrfid:0), output_ifc=any

Phase: 8
Type: NAT
Subtype: per-session
Result: ALLOW
Elapsed time: 52182 ns
Config:
Additional Information:
Reverse Flow based lookup yields rule:
in id=0x15052e96b150, priority=0, domain=nat-per-session, deny=false
hits=127, user_data=0x0, cs_id=0x0, reverse, use_real_addr, flags=0x0, protocol=6
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, dscp=0x0, nsg_id=none
input_ifc=any, output_ifc=any

Phase: 9
Type: IP-OPTIONS
Subtype:
Result: ALLOW
Elapsed time: 892 ns
Config:
Additional Information:
Reverse Flow based lookup yields rule:
in id=0x1502a7f9b460, priority=0, domain=inspect-ip-options, deny=true
hits=38, user_data=0x0, cs_id=0x0, reverse, flags=0x0, protocol=0
src ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any
dst ip/id=0.0.0.0, mask=0.0.0.0, port=0, tag=any, dscp=0x0, nsg_id=none
input_ifc=OUTSIDE2(vrfid:0), output_ifc=any

Phase: 10
```

```
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Elapsed time: 25422 ns
Config:
Additional Information:
New flow created with id 244, packet dispatched to next module
Module information for forward flow ...
snp_fp_inspect_ip_options
snp_fp_tcp_normalizer
snp_fp_tcp_proxy
snp_fp_snort
snp_fp_tcp_proxy
snp_fp_translate
snp_fp_tcp_normalizer
snp_fp_adjacency
snp_fp_fragment
snp_ifc_stat

Module information for reverse flow ...
snp_fp_inspect_ip_options
snp_fp_tcp_normalizer
snp_fp_translate
snp_fp_tcp_proxy
snp_fp_snort
snp_fp_tcp_proxy
snp_fp_tcp_normalizer
snp_fp_adjacency
snp_fp_fragment
snp_ifc_stat

Phase: 11
Type: EXTERNAL-INSPECT
Subtype:
Result: ALLOW
Elapsed time: 36126 ns
Config:
Additional Information:
Application: 'SNORT Inspect'

Phase: 12
Type: SNORT
Subtype:
Result: ALLOW
Elapsed time: 564636 ns
Config:
Additional Information:
Snort Trace:
Packet: TCP, SYN, seq 182318660
Session: new snort session
AppID: service unknown (0), application unknown (0)
Snort id 28, NAP id 1, IPS id 0, Verdict PASS
Snort Verdict: (pass-packet) allow this packet

Phase: 13
Type: INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Resolve Preferred Egress interface
Result: ALLOW
Elapsed time: 7136 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)
```

```
Phase: 14
Type: ADJACENCY-LOOKUP
Subtype: Resolve Nexthop IP address to MAC
Result: ALLOW
Elapsed time: 2230 ns
Config:
Additional Information:
Found adjacency entry for Next-hop 192.0.2.99 on interface OUTSIDE2
Adjacency :Active
MAC address 4c4e.35fc.fcd8 hits 10 reference 1

Phase: 15
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 5352 ns
Config:
Additional Information:
Forward Flow based lookup yields rule:
out id=0x150521389870, priority=13, domain=capture, deny=false
hits=1788, user_data=0x1505f1d2b630, cs_id=0x0, l3_type=0x0
src mac=0000.0000.0000, mask=0000.0000.0000
dst mac=0000.0000.0000, mask=0000.0000.0000
input_ifc=OUTSIDE2, output_ifc=any

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE2(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 721180 ns

1 packet shown
firepower#
```

Traccia un altro pacchetto in entrata dallo stesso flusso. Il pacchetto che corrisponde a una connessione attiva:

```
firepower# show capture CAPI packet-number 3 trace

33 packets captured

3: 10:50:38.409265 802.1Q vlan#101 P0 192.168.1.1.57734 > 198.51.100.1.23: . ack 1412677785 win 4128
Phase: 1
Type: CAPTURE
Subtype:
Result: ALLOW
Elapsed time: 2676 ns
Config:
Additional Information:
Forward Flow based lookup yields rule:
in id=0x1505f1d17940, priority=13, domain=capture, deny=false
```

```
hits=105083, user_data=0x1505f2096910, cs_id=0x0, l3_type=0x0
src mac=0000.0000.0000, mask=0000.0000.0000
dst mac=0000.0000.0000, mask=0000.0000.0000
input_ifc=INSIDE, output_ifc=any

Phase: 2
Type: ACCESS-LIST
Subtype:
Result: ALLOW
Elapsed time: 2676 ns
Config:
Implicit Rule
Additional Information:
Forward Flow based lookup yields rule:
in id=0x1502a7ba4d40, priority=1, domain=permit, deny=false
hits=45, user_data=0x0, cs_id=0x0, l3_type=0x8
src mac=0000.0000.0000, mask=0000.0000.0000
dst mac=0000.0000.0000, mask=0100.0000.0000
input_ifc=INSIDE, output_ifc=any

Phase: 3
Type: FLOW-LOOKUP
Subtype:
Result: ALLOW
Elapsed time: 1338 ns
Config:
Additional Information:
Found flow with id 2552, using existing flow
Module information for forward flow ...
snp_fp_inspect_ip_options
snp_fp_tcp_normalizer
snp_fp_snort
snp_fp_translate
snp_fp_tcp_normalizer
snp_fp_adjacency
snp_fp_fragment
snp_ifc_stat

Module information for reverse flow ...
snp_fp_inspect_ip_options
snp_fp_tcp_normalizer
snp_fp_translate
snp_fp_snort
snp_fp_tcp_normalizer
snp_fp_adjacency
snp_fp_fragment
snp_ifc_stat

Phase: 4
Type: EXTERNAL-INSPECT
Subtype:
Result: ALLOW
Elapsed time: 16502 ns
Config:
Additional Information:
Application: 'SNORT Inspect'

Phase: 5
Type: SNORT
Subtype:
Result: ALLOW
Elapsed time: 12934 ns
```

```
Config:
Additional Information:
Snort Trace:
Packet: TCP, ACK, seq 1306692136, ack 1412677785
AppID: service unknown (0), application unknown (0)
Snort id 19, NAP id 1, IPS id 0, Verdict PASS
Snort Verdict: (pass-packet) allow this packet

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
Action: allow
Time Taken: 36126 ns

1 packet shown
firepower#
```

Timeout mobile

Il problema

L'instabilità temporanea del percorso può causare connessioni UDP di lunga durata (elefanti) attraverso l'FTD da stabilire attraverso interfacce FTD diverse da quelle desiderate.

La soluzione

Per risolvere questo problema, impostare il timeout floating-conn su un valore diverso da quello predefinito disabilitato:

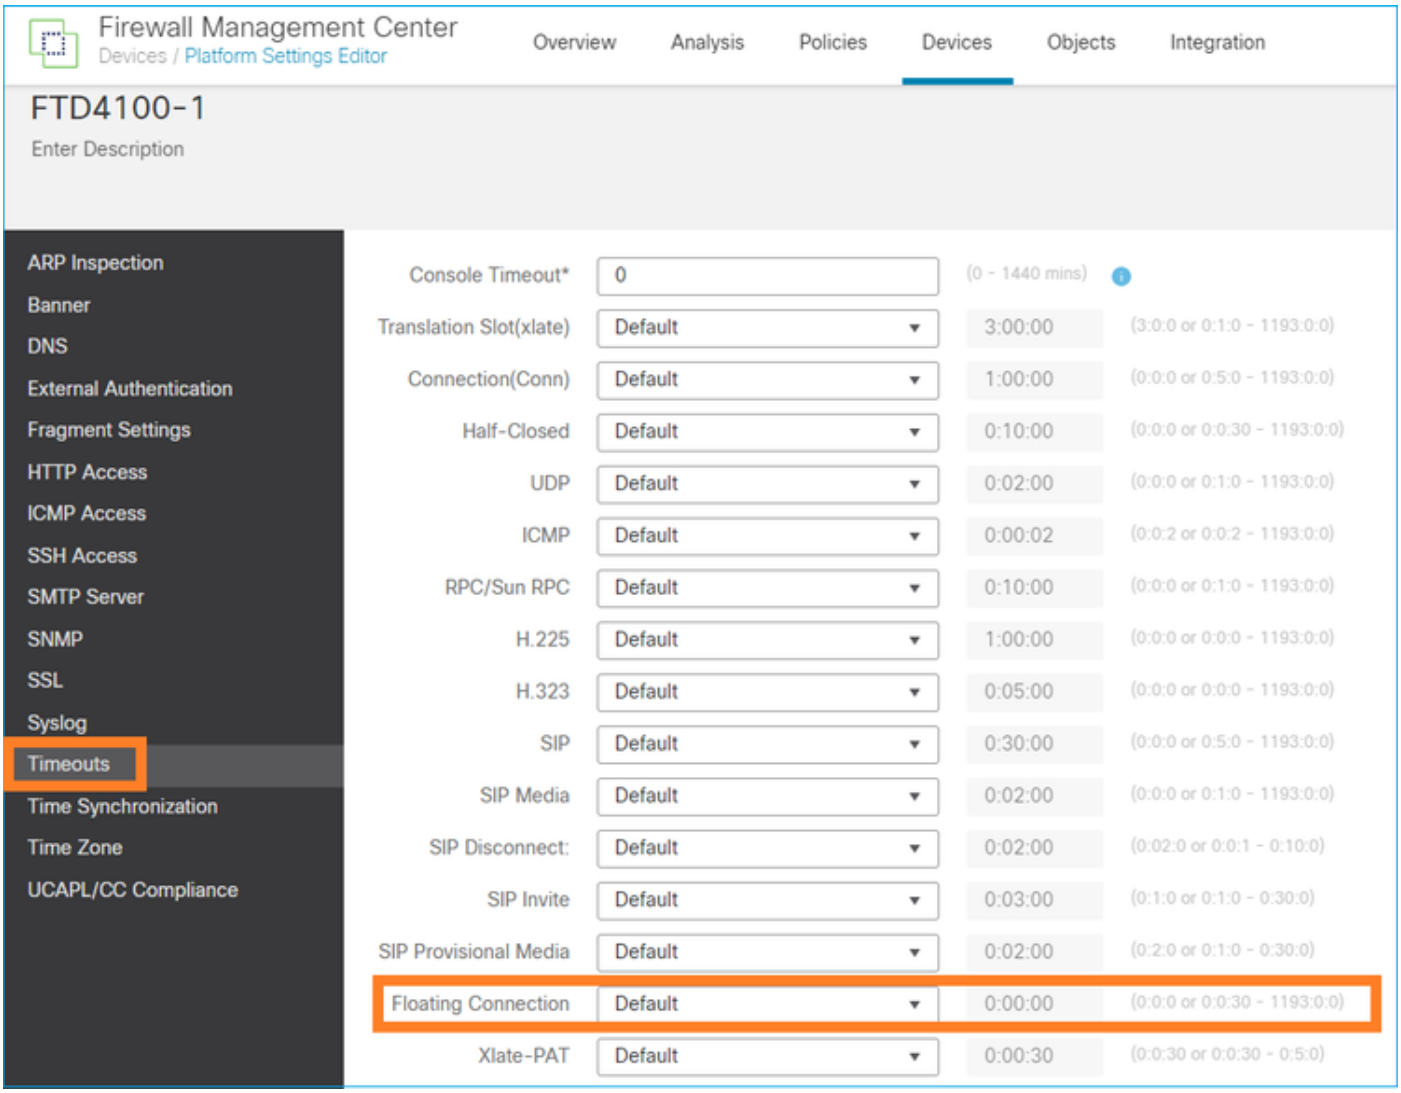

Dalla guida di riferimento per i comandi:

floating-conn	When multiple routes exist to a network with different metrics, the ASA uses the one with the best metric at the time of connection creation. If a better route becomes available, then this timeout lets connections be closed so a connection can be reestablished to use the better route. The default is 0 (the connection never times out). To make it possible to use better routes, set the timeout to a value between 0:0:30 and 1193:0:0.

Per ulteriori informazioni, vedere Case study: UDP Connections Fail After Reload from the Cisco Live BRKSEC-3020 session:

Floating Connection Timeout

- The "bad" connection never times out since the UDP traffic is constantly flowing
  - TCP is stateful, so the connection would terminate and re-establish on its own
  - ASA needs to tear the original connection down when the corresponding route changes
  - ASA 8.4(2)+ introduces **timeout floating-conn** to accomplish this goal

```
asa# show run timeout
timeout xlate 9:00:00
timeout pat-xlate 0:00:30
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 9:00:00 absolute uauth 0:01:00 inactivity
timeout tcp-proxy-reassembly 0:01:00
timeout floating-conn 0:00:00
asa#
asa# configure terminal
asa(config)# timeout floating-conn 0:01:00
```

Schedule the conn entry for termination in **1 minute** if a matching packet yields a different egress interface on route lookup

Timeout Conn-Holddown

Il problema

Un percorso diventa inattivo (viene rimosso), ma il traffico corrisponde a una connessione stabilita.

La soluzione

La funzione di controllo del timeout è stata aggiunta su ASA 9.6.2. La funzionalità è attivata per impostazione predefinita, ma attualmente (7.1.x) non è supportata dall'interfaccia utente di FMC o da FlexConfig. Miglioramenti correlati: [ENH: timeout conn-holddown non disponibile per la configurazione in FMC]

Dalla guida ASA CLI:

conn-holddown	How long the system should maintain a connection when the route used by the connection no longer exists or is inactive. If the route does not become active within this holddown period, the connection is freed. The purpose of the connection holddown timer is to reduce the effect of route flapping, where routes might come up and go down quickly. You can reduce the holddown timer to make route convergence happen more quickly. The default is 15 seconds, the range is 00:00:00 to 00:00:15.

```
firepower# show run all timeout
timeout xlate 3:00:00
timeout pat-xlate 0:00:30
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 sctp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout sip-provisional-media 0:02:00 uauth 0:05:00 absolute
timeout tcp-proxy-reassembly 0:00:30
timeout floating-conn 0:00:00
timeout conn-holddown 0:00:15
timeout igp stale-route 0:01:10
```

Caso 2 - Inoltro basato su ricerca NAT

Requisito

Configura questa regola NAT:

- Tipo: statico
- Interfaccia di origine: INSIDE
- Interfaccia di destinazione: OUTSIDE1
- Fonte originale: 192.168.1.1
- Destinazione originale: 198.51.100.1
- Fonte tradotta: 192.168.1.1
- Destinazione tradotta: 198.51.100.1

Soluzione

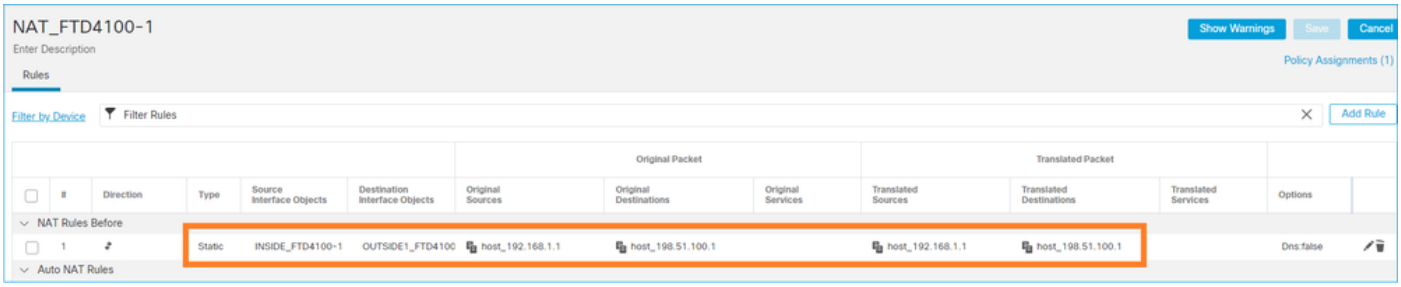

La regola NAT distribuita nella CLI FTD:

```
firepower# show run nat
nat (INSIDE,OUTSIDE1) source static host_192.168.1.1 host_192.168.1.1 destination static host_198.51.10
firepower# show nat
Manual NAT Policies (Section 1)
1 (INSIDE) to (OUTSIDE1) source static host_192.168.1.1 host_192.168.1.1 destination static host_198.51
translate_hits = 0, untranslate_hits = 0
```

Configura 3 clip:

```
firepower# capture CAPI trace detail interface INSIDE match ip host 192.168.1.1 host 198.51.100.1
firepower# capture CAPO1 interface OUTSIDE1 match ip host 192.168.1.1 any
firepower# capture CAPO2 interface OUTSIDE2 match ip host 192.168.1.1 any
firepower# show capture
capture CAPI type raw-data trace detail interface INSIDE [Capturing - 0 bytes]
match ip host 192.168.1.1 host 198.51.100.1
capture CAPO1 type raw-data interface OUTSIDE1 [Capturing - 0 bytes]
```

```
match ip host 192.168.1.1 any
capture CAPO2 type raw-data interface OUTSIDE2 [Capturing - 0 bytes]
match ip host 192.168.1.1 any
```

Avviare una sessione telnet da 192.168.1.1 a 198.51.100.1:

```
Router1# telnet 198.51.100.1 /vrf VRF-101 /source-interface lo1
Trying 198.51.100.1 ...
% Connection timed out; remote host not responding
```

I pacchetti arrivano su FTD, ma niente lascia le interfacce OUTSIDE1 o OUTSIDE2:

```
firepower# show capture
capture CAPI type raw-data trace detail interface INSIDE [Capturing - 156 bytes]
match ip host 192.168.1.1 host 198.51.100.1
capture CAPO1 type raw-data interface OUTSIDE1 [Capturing - 0 bytes]
match ip host 192.168.1.1 any
capture CAPO2 type raw-data interface OUTSIDE2 [Capturing - 0 bytes]
match ip host 192.168.1.1 any
```

Tracciare il pacchetto TCP SYN. La fase 3 (UN-NAT) mostra che NAT (UN-NAT in particolare) ha deviato il pacchetto all'interfaccia OUTSIDE1 per la ricerca dell'hop successivo:

```
firepower# show capture CAPI
2 packets captured
1: 11:22:59.179678 802.1Q vlan#101 P0 192.168.1.1.38790 > 198.51.100.1.23: S 1174675193:1174675193(0) w
2: 11:23:01.179632 802.1Q vlan#101 P0 192.168.1.1.38790 > 198.51.100.1.23: S 1174675193:1174675193(0) w
2 packets shown
firepower#
```

```
firepower# show capture CAPI packet-number 1 trace detail

2 packets captured

1: 11:22:59.179678 4c4e.35fc.fcd8 00be.75f6.1dae 0x8100 Length: 62
802.1Q vlan#101 P0 192.168.1.1.38790 > 198.51.100.1.23: S [tcp sum ok] 1174675193:1174675193(0) win 412
...

Phase: 3
Type: UN-NAT
Subtype: static
Result: ALLOW
```

```
Elapsed time: 6244 ns
Config:
nat (INSIDE,OUTSIDE1) source static host_192.168.1.1 host_192.168.1.1 destination static host_198.51.10
Additional Information:
NAT divert to egress interface OUTSIDE1(vrfid:0)
Untranslate 198.51.100.1/23 to 198.51.100.1/23

...
Phase: 12
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Elapsed time: 25422 ns
Config:
Additional Information:
New flow created with id 2614, packet dispatched to next module
Module information for forward flow ...
snp_fp_inspect_ip_options
snp_fp_tcp_normalizer
snp_fp_tcp_proxy
snp_fp_snort
snp_fp_tcp_proxy
snp_fp_translate
snp_fp_tcp_normalizer
snp_fp_adjacency
snp_fp_fragment
snp_ifc_stat


Phase: 15
Type: INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Resolve Preferred Egress interface
Result: ALLOW
Elapsed time: 8028 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

Phase: 16
Type: SUBOPTIMAL-LOOKUP
Subtype: suboptimal next-hop
Result: ALLOW
Elapsed time: 446 ns
Config:
Additional Information:
Input route lookup returned ifc OUTSIDE2 is not same as existing ifc OUTSIDE1

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE1(vrfid:0)
output-status: up
output-line-status: up
Action: drop
Time Taken: 777375 ns
Drop-reason: (no-adjacency) No valid adjacency, Drop-location: frame 0x00005577204a7287 flow (NA)/NA


1 packet shown
```

In questo caso, SUBOPTIMAL-LOOKUP indica che l'interfaccia in uscita determinata dal processo NAT (OUTSIDE1) è diversa dall'interfaccia in uscita specificata nella tabella di input ASP:

```
firepower# show asp table routing | include 198.51.100.0
in 198.51.100.0 255.255.255.248 via 192.0.2.99, OUTSIDE2
out 198.51.100.0 255.255.255.248 via 192.0.2.99, OUTSIDE2
```

Per ovviare al problema, aggiungere una route statica mobile sull'interfaccia OUTSIDE1:

```
firepower# show run route
route OUTSIDE2 198.51.100.0 255.255.255.248 192.0.2.99 1
route OUTSIDE1 198.51.100.0 255.255.255.248 203.0.113.99 200
```

Nota: se si tenta di aggiungere una route statica con la stessa metrica di quella già esistente, viene visualizzato questo errore:

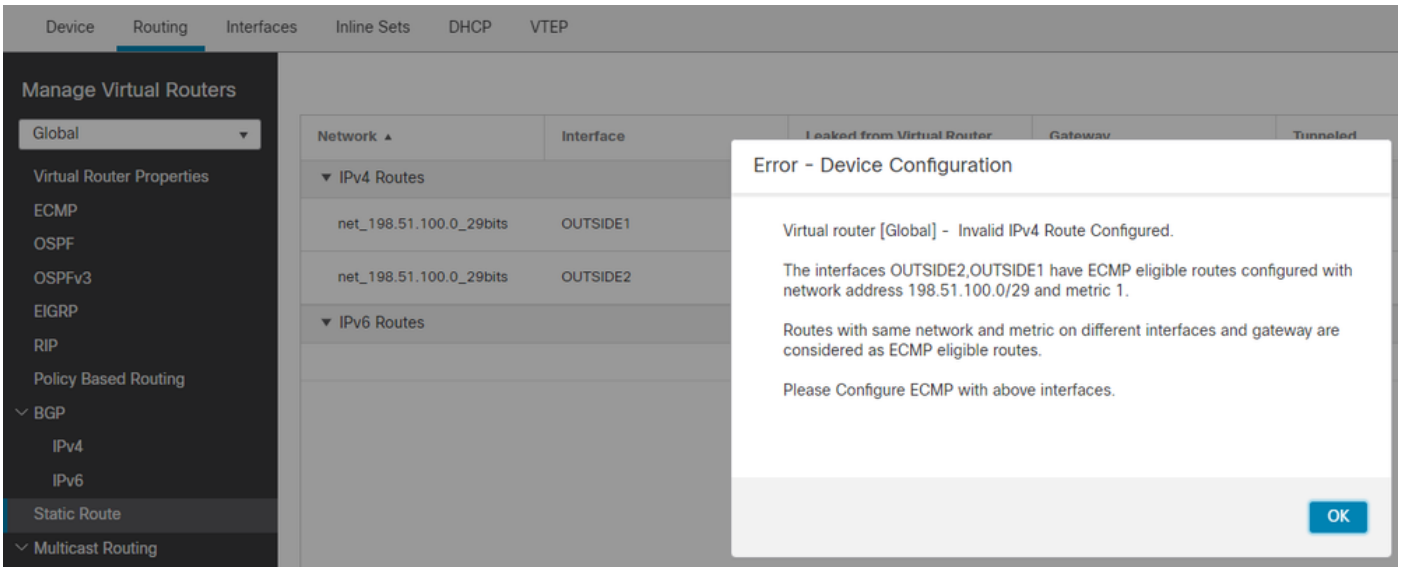

Nota: il percorso mobile con una metrica di distanza pari a 255 non è installato nella tabella di routing.

Provare a connettersi in modalità Telnet per verificare che vi siano pacchetti inviati tramite l'FTD:

```
Router1# telnet 198.51.100.1 /vrf VRF-101 /source-interface lo1
Trying 198.51.100.1 ...
% Connection timed out; remote host not responding
```

```
firepower# show capture
capture CAPI type raw-data trace detail interface INSIDE [Capturing - 156 bytes]
match ip host 192.168.1.1 host 198.51.100.1
capture CAPO1 type raw-data interface OUTSIDE1 [Capturing - 312 bytes]
match ip host 192.168.1.1 any
capture CAPO2 type raw-data interface OUTSIDE2 [Capturing - 386 bytes]
match ip host 192.168.1.1 any
```

Nell'analisi del pacchetto viene mostrato che i pacchetti vengono inoltrati all'interfaccia ISP1 (OUTSIDE1) anziché all'ISP2 a causa di una ricerca NAT:

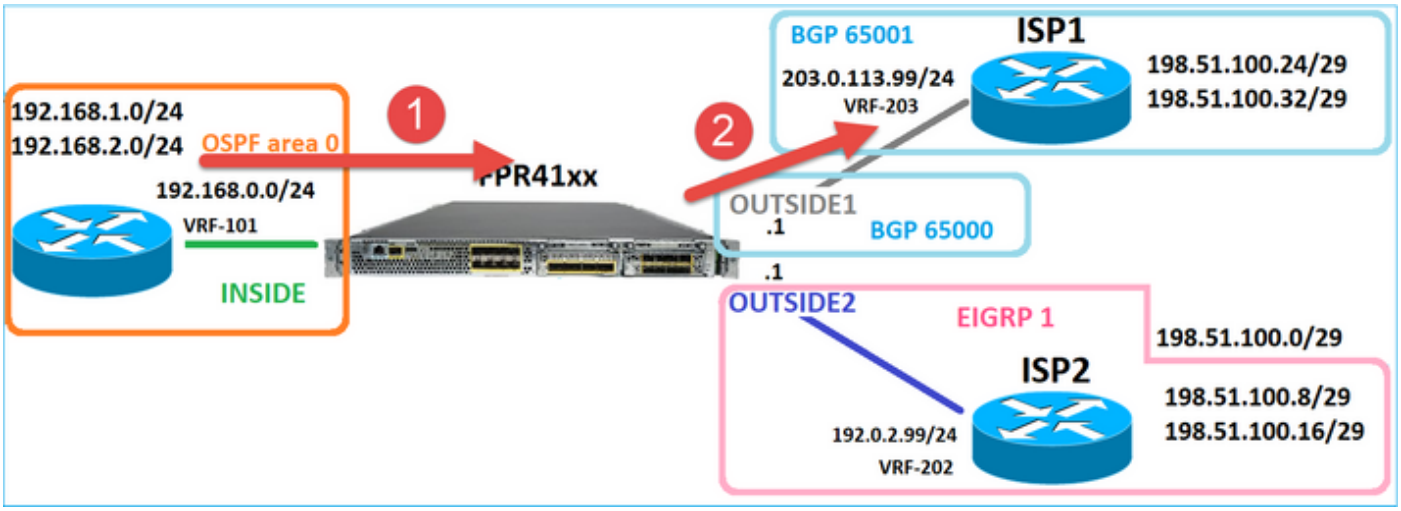


```
firepower# show capture CAPI packet-number 1 trace

2 packets captured

1: 09:03:02.773962 802.1Q vlan#101 P0 192.168.1.1.16774 > 198.51.100.1.23: S 2910053251:2910053251(0) w
...

Phase: 3
Type: UN-NAT
Subtype: static
Result: ALLOW
Elapsed time: 4460 ns
Config:
nat (INSIDE,OUTSIDE1) source static host_192.168.1.1 host_192.168.1.1 destination static host_198.51.10
Additional Information:
NAT divert to egress interface OUTSIDE1(vrfid:0)
Untranslate 198.51.100.1/23 to 198.51.100.1/23

...

Phase: 12
Type: FLOW-CREATION
Subtype:
Result: ALLOW
Elapsed time: 29436 ns
Config:
Additional Information:
New flow created with id 2658, packet dispatched to next module
Module information for forward flow ...
```

```
snp_fp_inspect_ip_options
snp_fp_tcp_normalizer
snp_fp_snort
snp_fp_translate
snp_fp_tcp_normalizer
snp_fp_adjacency
snp_fp_fragment
snp_ifc_stat

Phase: 15
Type: INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Resolve Preferred Egress interface
Result: ALLOW
Elapsed time: 5798 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

Phase: 16
Type: SUBOPTIMAL-LOOKUP
Subtype: suboptimal next-hop
Result: ALLOW
Elapsed time: 446 ns
Config:
Additional Information:
Input route lookup returned ifc OUTSIDE2 is not same as existing ifc OUTSIDE1

Phase: 17
Type: NEXTHOP-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Lookup Nexthop on interface
Result: ALLOW
Elapsed time: 1784 ns
Config:
Additional Information:
Found next-hop 203.0.113.99 using egress ifc OUTSIDE1(vrfid:0)

Phase: 18
Type: ADJACENCY-LOOKUP
Subtype: Resolve Nexthop IP address to MAC
Result: ALLOW
Elapsed time: 1338 ns
Config:
Additional Information:
Found adjacency entry for Next-hop 203.0.113.99 on interface OUTSIDE1
Adjacency :Active
MAC address 4c4e.35fc.fcd8 hits 106 reference 2
...

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE1(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 723409 ns


1 packet shown
firepower#
```

È interessante notare che in questo caso, i pacchetti sono visualizzati su INSIDE ed entrambe le interfacce in uscita:

```
firepower# show capture CAPI

2 packets captured

1: 09:03:02.773962 802.1Q vlan#101 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3031010184:3031010184(0) w
2: 09:03:05.176565 802.1Q vlan#101 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3031010184:3031010184(0) w
2 packets shown
firepower# show capture CAPO1

4 packets captured

1: 09:03:02.774358 802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3249840142:3249840142(0) w
2: 09:03:02.774557 802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3249840142:3249840142(0) w
3: 09:03:05.176702 802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3249840142:3249840142(0) w
4: 09:03:05.176870 802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3249840142:3249840142(0) w
4 packets shown
firepower# show capture CAPO2

5 packets captured

1: 09:03:02.774679 802.1Q vlan#202 P0 192.168.1.1.32134 > 198.51.100.1.23: S 194652172:194652172(0) win
2: 09:03:02.775457 802.1Q vlan#202 P0 198.51.100.1.23 > 192.168.1.1.32134: S 4075003210:4075003210(0) a
3: 09:03:05.176931 802.1Q vlan#202 P0 192.168.1.1.32134 > 198.51.100.1.23: S 194652172:194652172(0) win
4: 09:03:05.177282 802.1Q vlan#202 P0 198.51.100.1.23 > 192.168.1.1.32134: . ack 194652173 win 4128
5: 09:03:05.180517 802.1Q vlan#202 P0 198.51.100.1.23 > 192.168.1.1.32134: S 4075003210:4075003210(0) a
```

I dettagli del pacchetto includono le informazioni sull'indirizzo MAC, e una traccia dei pacchetti sulle interfacce OUTSIDE1 e OUTSIDE2 rivela il percorso dei pacchetti:

```
firepower# show capture CAPO1 detail

4 packets captured

1: 09:03:02.774358 00be.75f6.1dae 4c4e.35fc.fcd8 0x8100 Length: 62
802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S [tcp sum ok] 3249840142:3249840142(0) win 412
2: 09:03:02.774557 4c4e.35fc.fcd8 00be.75f6.1dae 0x8100 Length: 62
802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S [tcp sum ok] 3249840142:3249840142(0) win 412
3: 09:03:05.176702 00be.75f6.1dae 4c4e.35fc.fcd8 0x8100 Length: 62
802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S [tcp sum ok] 3249840142:3249840142(0) win 412
4: 09:03:05.176870 4c4e.35fc.fcd8 00be.75f6.1dae 0x8100 Length: 62
802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S [tcp sum ok] 3249840142:3249840142(0) win 412
4 packets shown
```

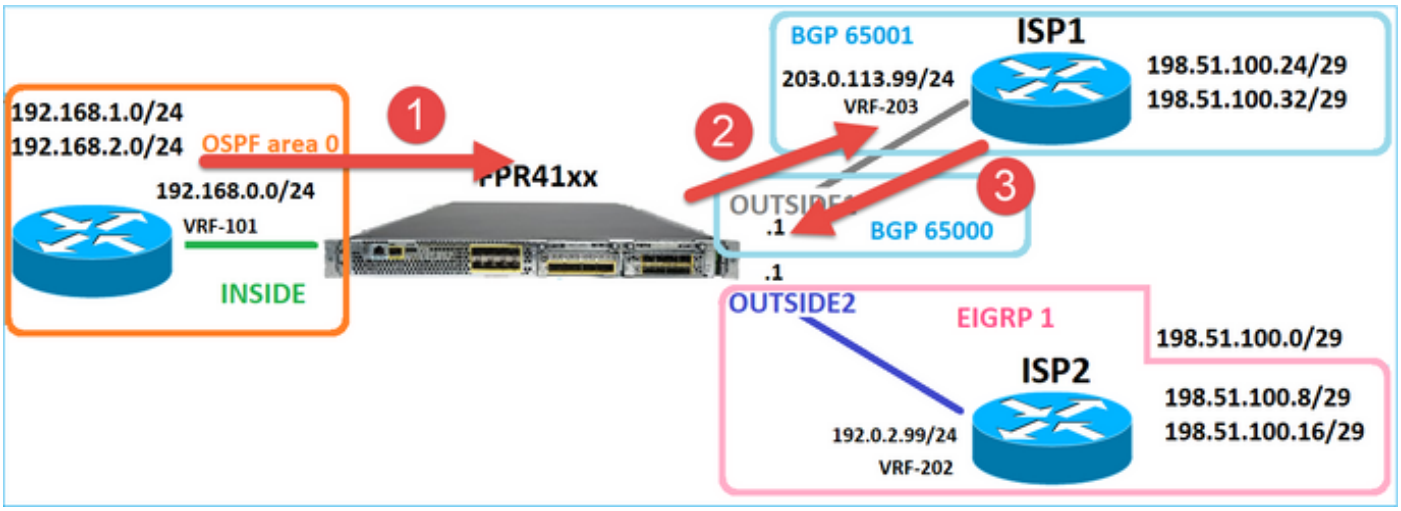

La traccia del pacchetto restituito mostra il reindirizzamento all'interfaccia OUTSIDE2 dovuto alla ricerca della tabella di routing globale:

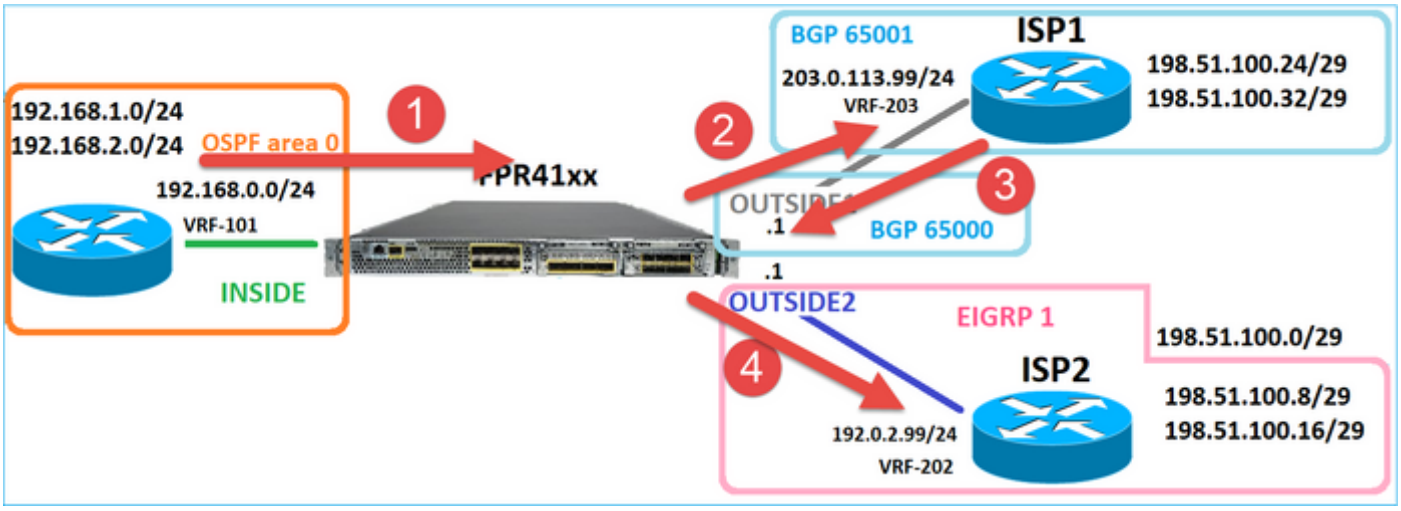


```
firepower# show capture CAPO1 packet-number 2 trace

4 packets captured

2: 09:03:02.774557 802.1Q vlan#203 P0 192.168.1.1.32134 > 198.51.100.1.23: S 3249840142:3249840142(0) w
...

Phase: 3
Type: INPUT-ROUTE-LOOKUP
Subtype: Resolve Egress Interface
Result: ALLOW
Elapsed time: 7136 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

...

Phase: 10
Type: FLOW-CREATION
Subtype:
Result: ALLOW
```

```
Elapsed time: 12488 ns
Config:
Additional Information:
New flow created with id 13156, packet dispatched to next module

...

Phase: 13
Type: INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP
Subtype: Resolve Preferred Egress interface
Result: ALLOW
Elapsed time: 3568 ns
Config:
Additional Information:
Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0)

Phase: 14
Type: ADJACENCY-LOOKUP
Subtype: Resolve Nexthop IP address to MAC
Result: ALLOW
Elapsed time: 1338 ns
Config:
Additional Information:
Found adjacency entry for Next-hop 192.0.2.99 on interface OUTSIDE2
Adjacency :Active
MAC address 4c4e.35fc.fcd8 hits 0 reference 1

...

Result:
input-interface: OUTSIDE1(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE2(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 111946 ns


1 packet shown
firepower#
```

Il router ISP2 invia la risposta (SYN/ACK), ma questo pacchetto viene reindirizzato all'ISP1 perché corrisponde alla connessione stabilita. Il pacchetto viene scartato dall'FTD a causa dell'assenza di adiacenze L2 nella tabella ASP in uscita:

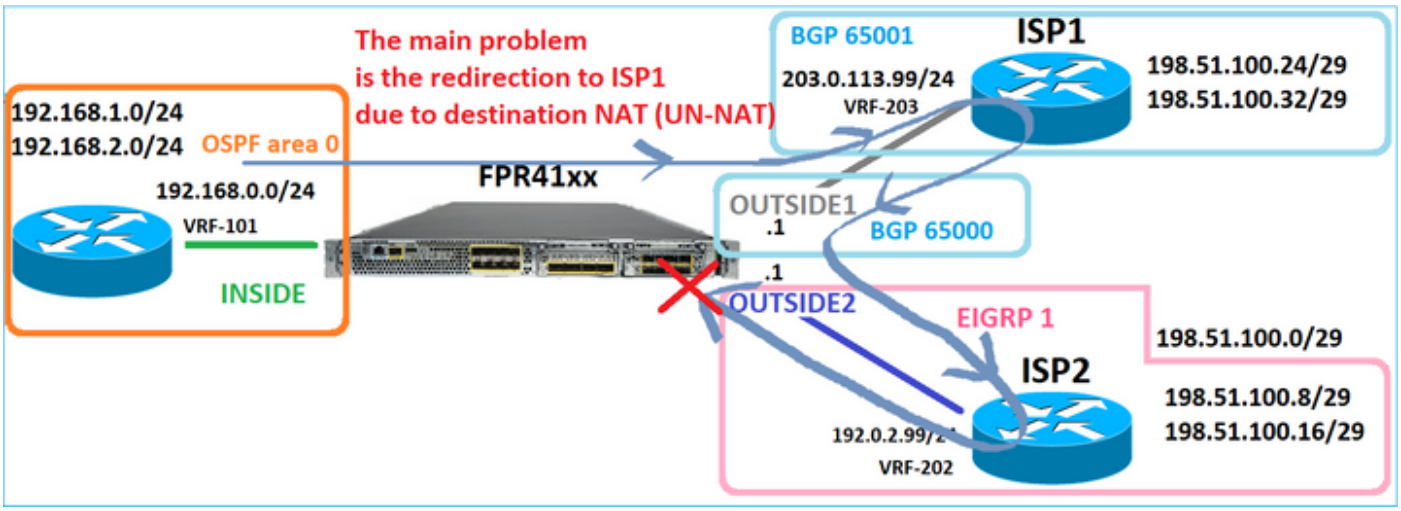


```
firepower# show capture CAPO2 packet-number 2 trace

5 packets captured

2: 09:03:02.775457 802.1Q vlan#202 P0 198.51.100.1.23 > 192.168.1.1.32134: S 4075003210:4075003210(0) a
...

Phase: 3
Type: FLOW-LOOKUP
Subtype:
Result: ALLOW
Elapsed time: 2230 ns
Config:
Additional Information:
Found flow with id 13156, using existing flow

...

Phase: 7
Type: SUBOPTIMAL-LOOKUP
Subtype: suboptimal next-hop
Result: ALLOW
Elapsed time: 0 ns
Config:
Additional Information:
Input route lookup returned ifc INSIDE is not same as existing ifc OUTSIDE1

Result:
input-interface: OUTSIDE2(vrfid:0)
input-status: up
input-line-status: up
output-interface: INSIDE(vrfid:0)
output-status: up
output-line-status: up
Action: drop
Time Taken: 52628 ns
Drop-reason: (no-adjacency) No valid adjacency, Drop-location: frame 0x00005577204a7287 flow (NA)/NA
```

Caso 3 - Inoltro basato su PBR (Policy Based Routing)

Dopo la ricerca del flusso di connessione e la ricerca NAT di destinazione, PBR è l'elemento successivo che può influenzare la determinazione dell'interfaccia in uscita. Il PBR è documentato in: Policy Based Routing

Per la configurazione PBR sul CCP, è importante conoscere le seguenti linee guida:
FlexConfig è stato utilizzato per configurare PBR in FMC per le versioni FTD precedenti alla 7.1. È comunque possibile utilizzare FlexConfig per configurare PBR in tutte le versioni. Tuttavia, per un'interfaccia in entrata, non è possibile configurare PBR utilizzando sia la pagina FlexConfig che la pagina Policy Based Routing di FMC.

In questo caso di studio, l'FTD ha un percorso verso 198.51.100.0/24 che punta verso ISP2:

```
firepower# show route | begin Gate
Gateway of last resort is not set

C 192.0.2.0 255.255.255.0 is directly connected, OUTSIDE2
L 192.0.2.1 255.255.255.255 is directly connected, OUTSIDE2
C 192.168.0.0 255.255.255.0 is directly connected, INSIDE
L 192.168.0.1 255.255.255.255 is directly connected, INSIDE
O 192.168.1.1 255.255.255.255 [110/11] via 192.168.0.99, 5d01h, INSIDE
O 192.168.2.1 255.255.255.255 [110/11] via 192.168.0.99, 5d01h, INSIDE
S 198.51.100.0 255.255.255.248 [1/0] via 192.0.2.99, OUTSIDE2
D 198.51.100.8 255.255.255.248
[90/130816] via 192.0.2.99, 5d01h, OUTSIDE2
D 198.51.100.16 255.255.255.248
[90/130816] via 192.0.2.99, 5d01h, OUTSIDE2
B 198.51.100.24 255.255.255.248 [20/0] via 203.0.113.99, 5d00h
B 198.51.100.32 255.255.255.248 [20/0] via 203.0.113.99, 5d00h
C 203.0.113.0 255.255.255.0 is directly connected, OUTSIDE1
L 203.0.113.1 255.255.255.255 is directly connected, OUTSIDE1
```

Requisito

Configurare un criterio PBR con le caratteristiche seguenti:

- Il traffico proveniente da IP 192.168.2.0/24 e destinato a 198.51.100.5 deve essere inviato all'ISP1 (next-hop 203.0.113.99), mentre le altre origini devono usare l'interfaccia OUTSIDE2.

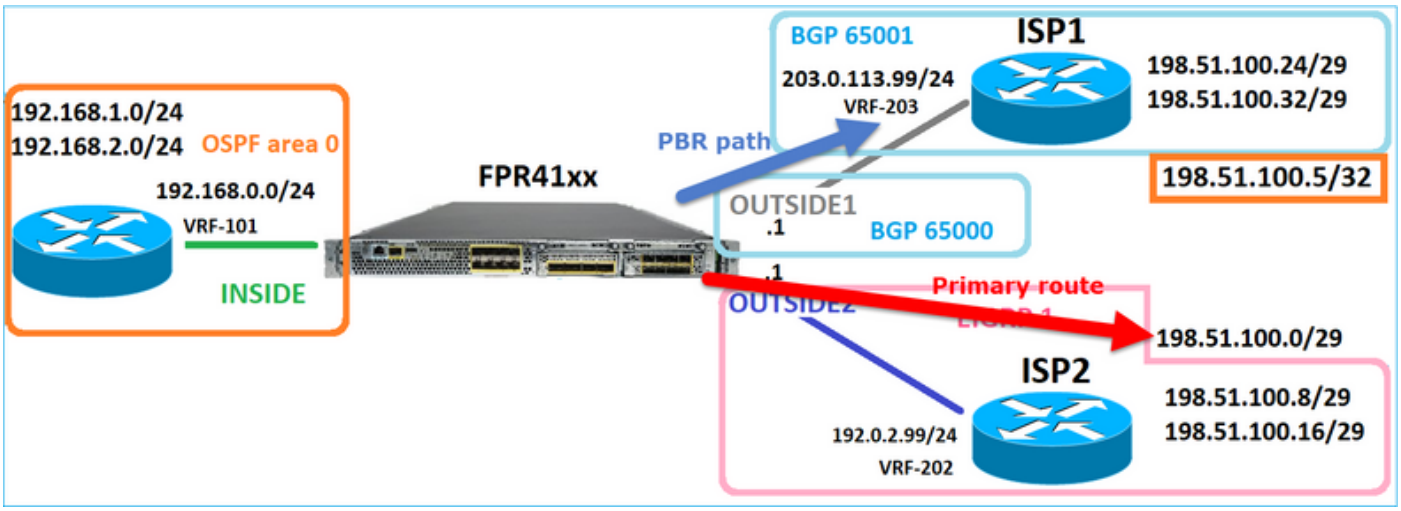

Soluzione

Nelle versioni precedenti alla 7.1, per configurare il PBR:
1. Creare un ACL esteso che corrisponda al traffico interessato (ad esempio, PBR_ACL).
2. Creare una route-map che corrisponda all'ACL creato nel passaggio 1 e impostare l'hop successivo desiderato.
3. Creare un oggetto FlexConfig che abiliti PBR sull'interfaccia in entrata utilizzando la mappa route creata nel passo 2.

Nelle versioni successive alla 7.1, è possibile configurare PBR utilizzando la modalità precedente alla 7.1 oppure utilizzare la nuova opzione Policy Based Routing nella sezione Device > Routing:
1. Creare un ACL esteso che corrisponda al traffico interessato (ad esempio, PBR_ACL).
2. Aggiungere un criterio PBR e specificare:
a. Il traffico corrispondente
b. L'interfaccia in entrata
c. L'hop successivo

Configura PBR (nuovo modo)

Passaggio 1 - Definire un elenco degli accessi per il traffico corrispondente.

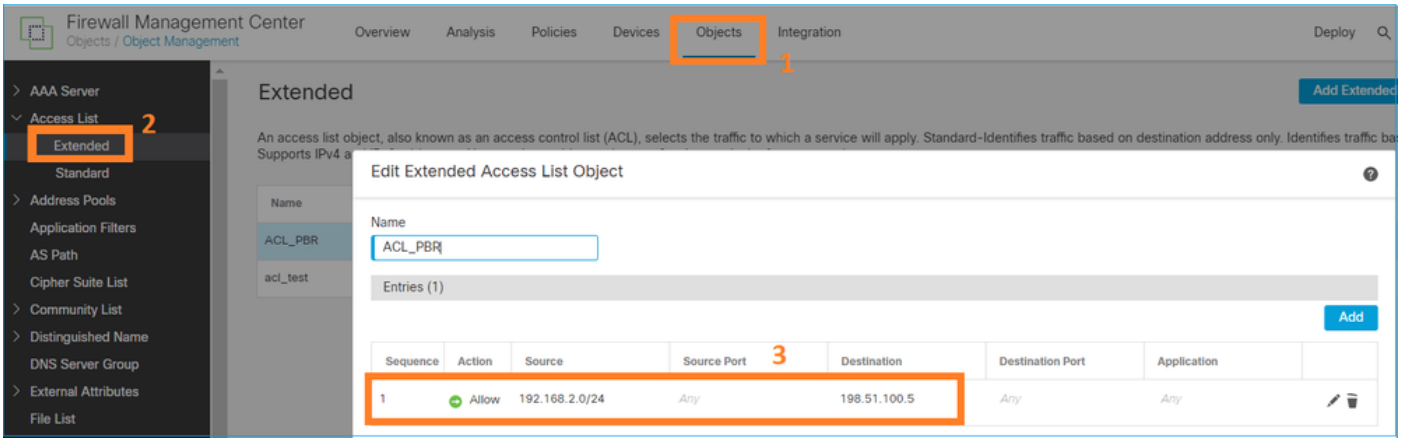

Passaggio 2 - Aggiungere un criterio PBR

Selezionare Dispositivi > Gestione dispositivi e modificare il dispositivo FTD. Scegliere

Instradamento > Instradamento basato su criteri e nella pagina Instradamento basato su criteri selezionare Aggiungi.

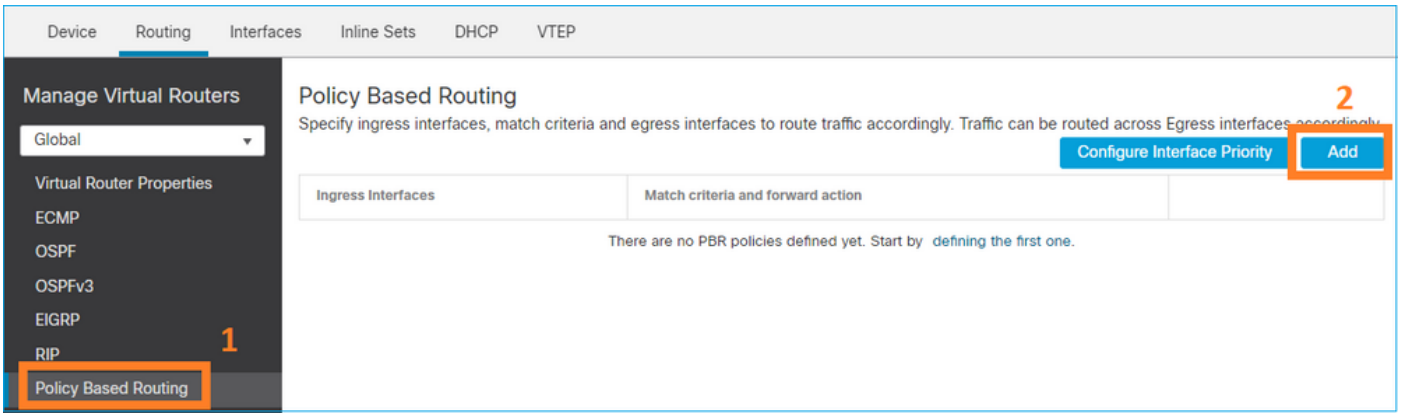

Specificare l'interfaccia in entrata:

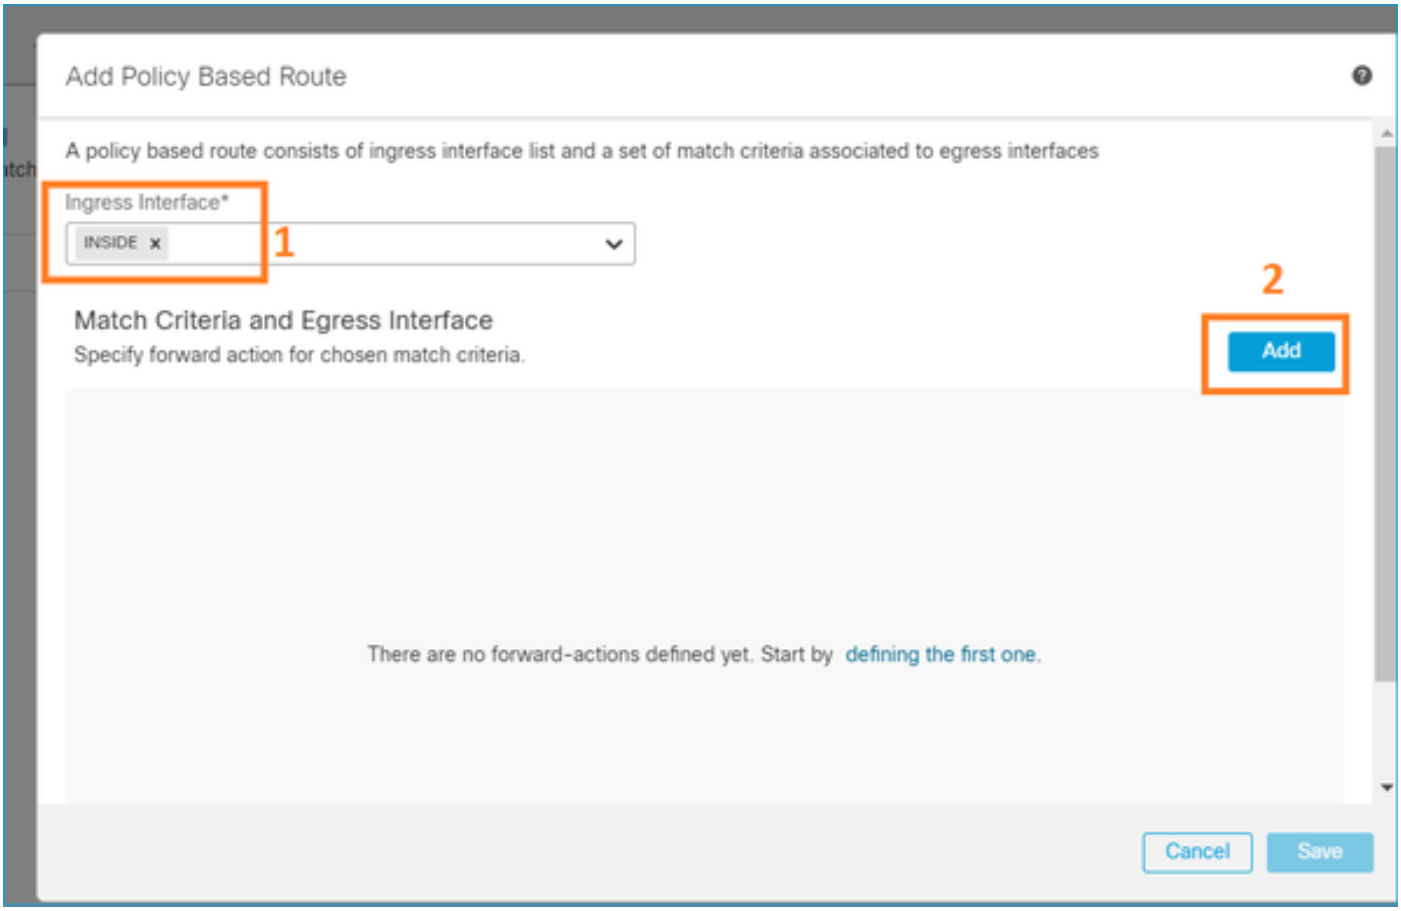

Specificare le azioni di inoltro:

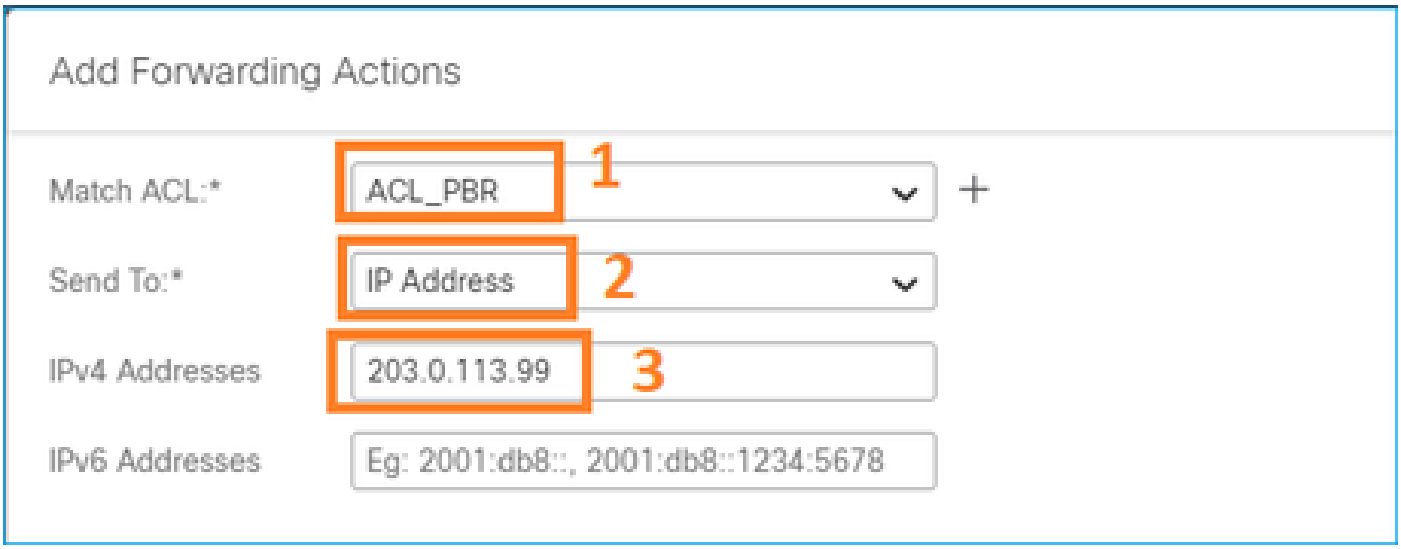

Salva e distribuisci

Nota: se si desidera configurare più interfacce in uscita, è necessario impostare nel campo 'Invia a' l'opzione 'Interfacce in uscita' (disponibile dalla versione 7.0+). Per ulteriori informazioni, vedere: Esempio di configurazione per il routing basato su criteri

---

Configura PBR (modalità legacy)

Passaggio 1 - Definire un elenco degli accessi per il traffico corrispondente.

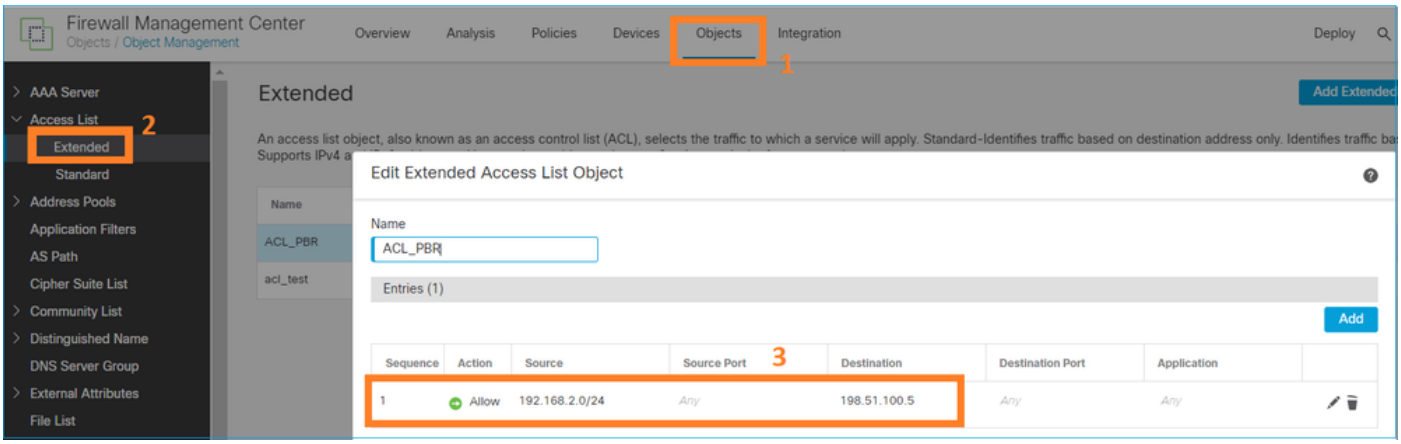

2. Definire una route-map che corrisponda all'ACL e imposti l'hop successivo.

Definire innanzitutto la clausola di corrispondenza:

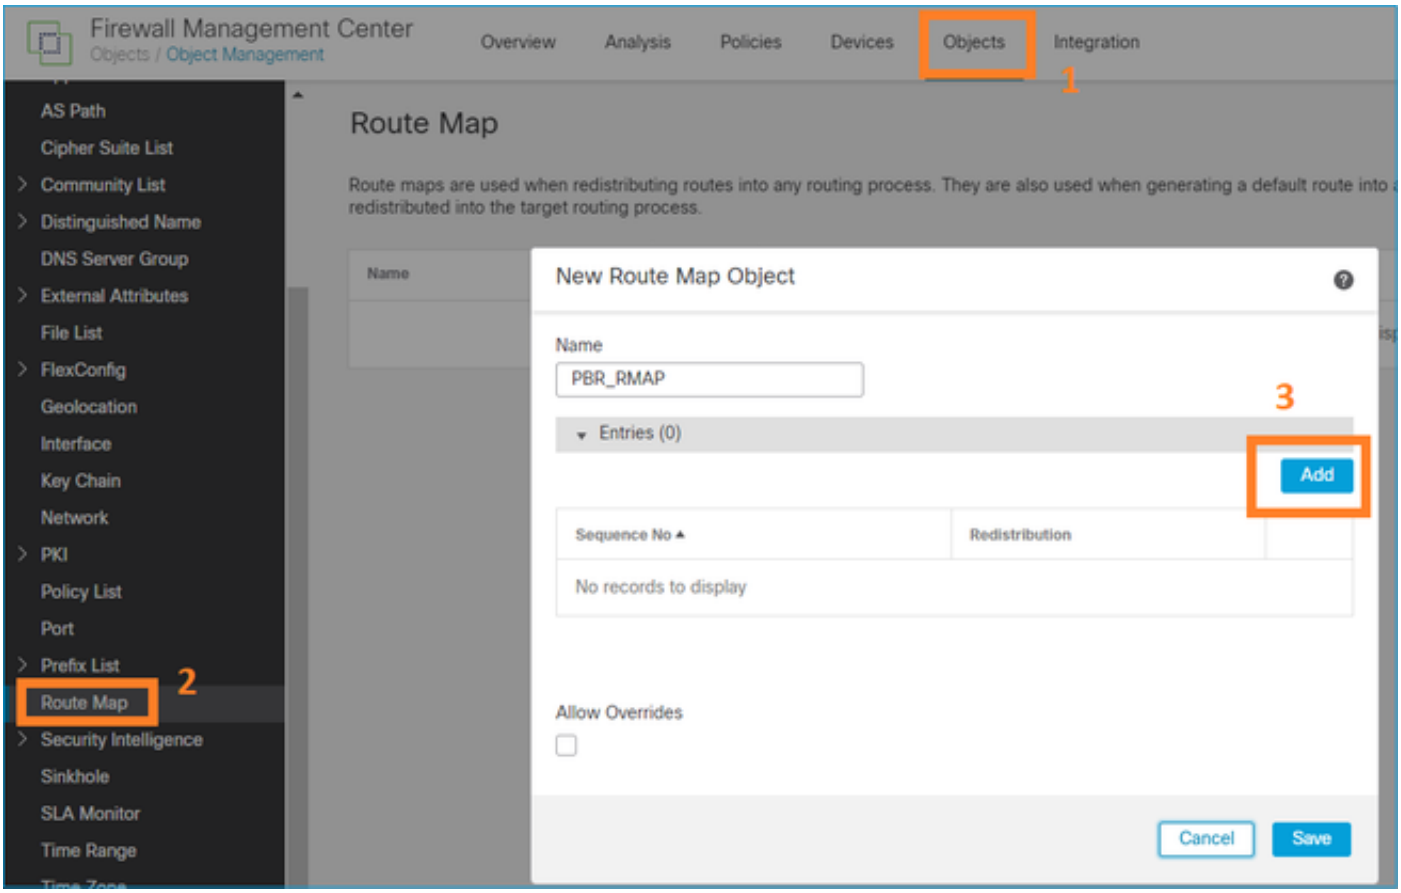

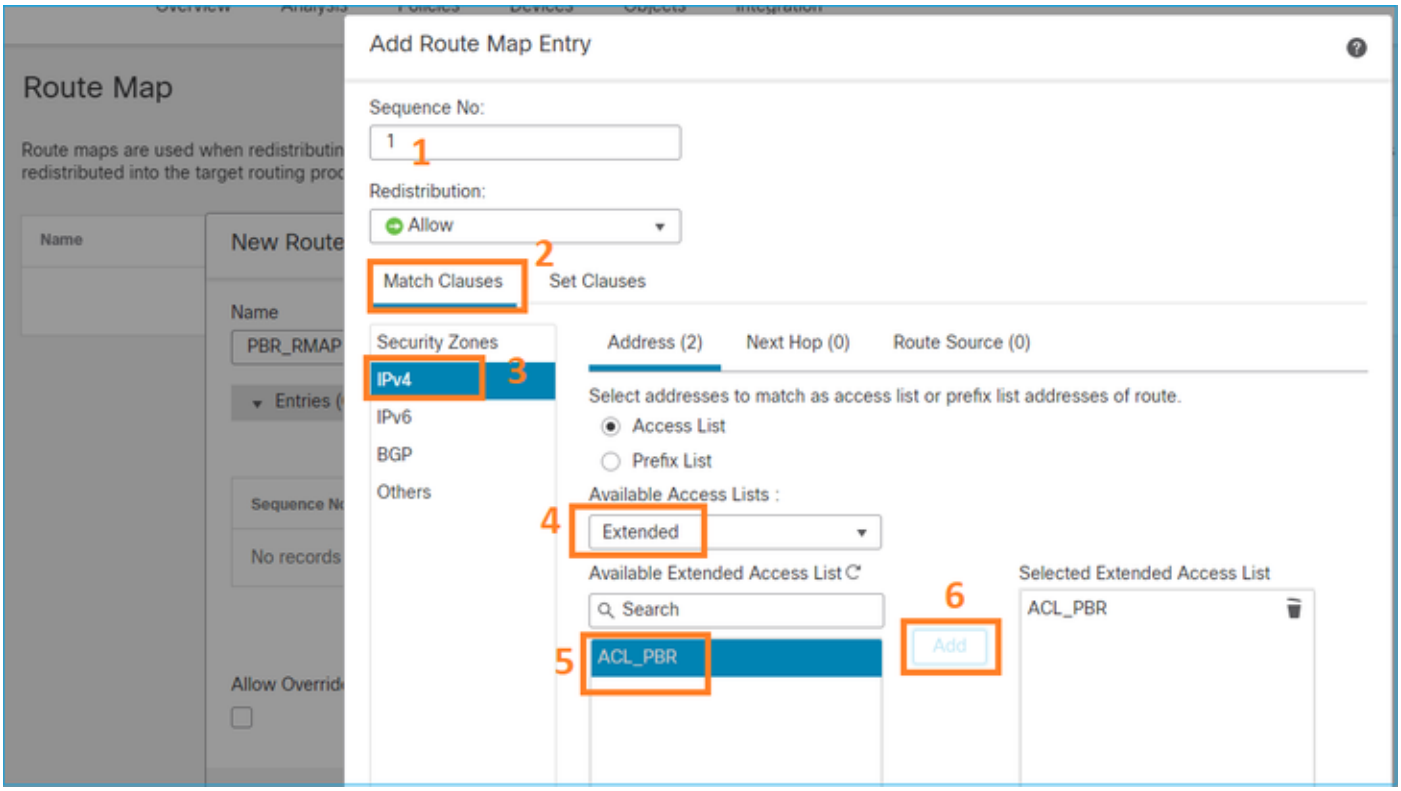

Definire la clausola Set:

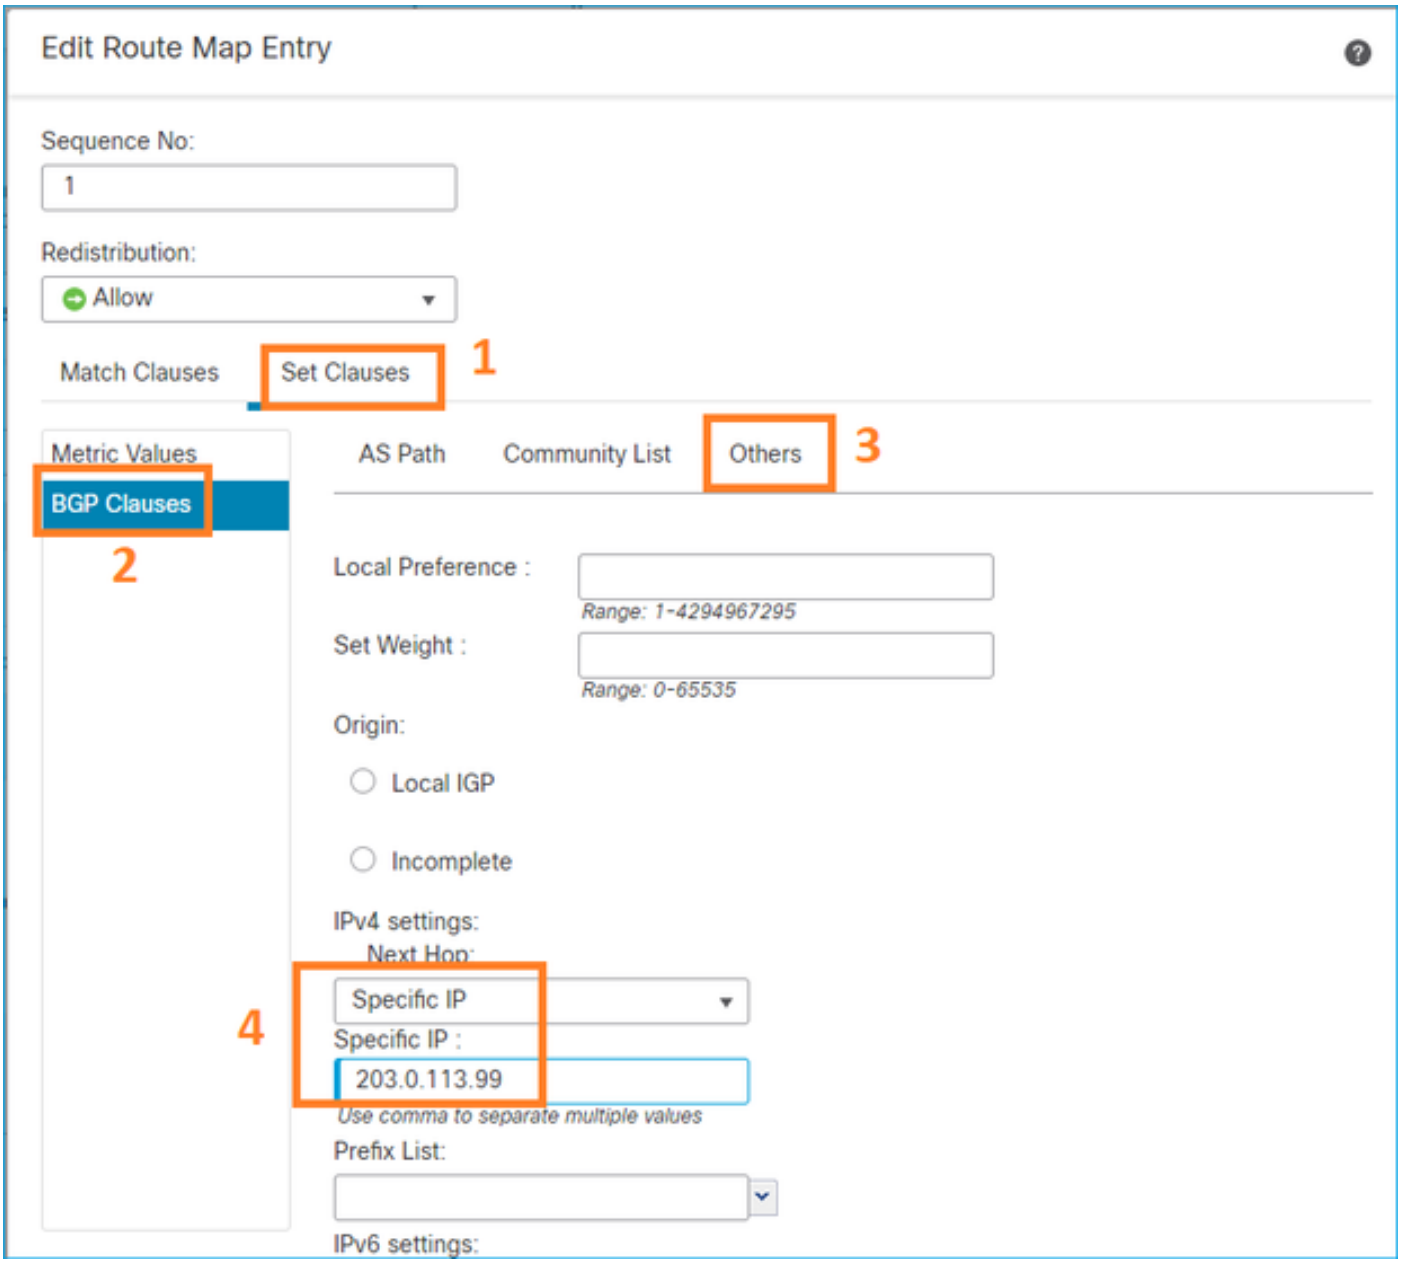

Aggiungi e salva.

Passaggio 3 - Configurare l'oggetto PBR FlexConfig.

Copiare innanzitutto (duplicare) l'oggetto PBR esistente:

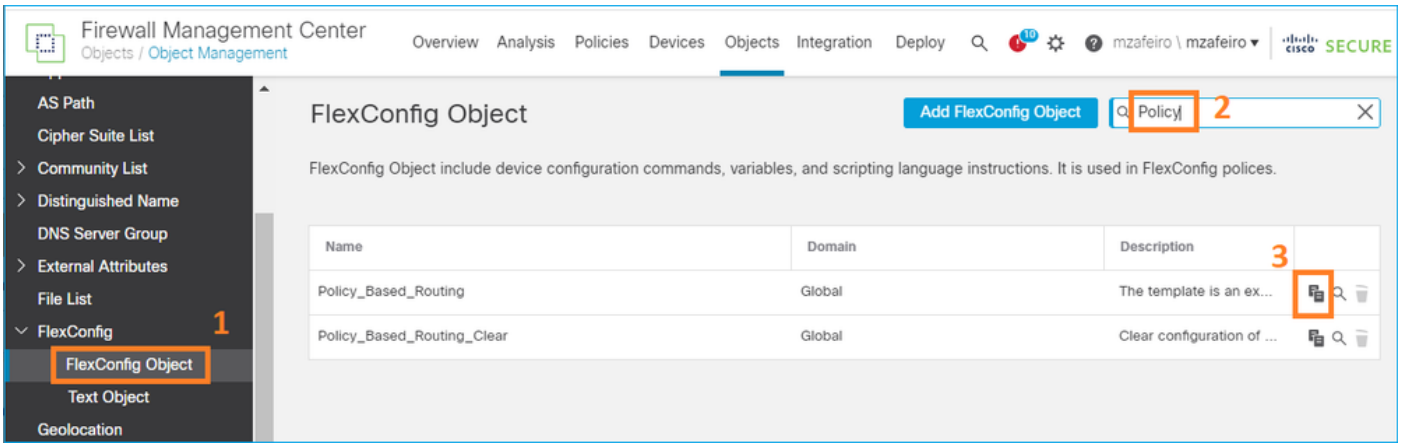

Specificare il nome dell'oggetto e rimuovere l'oggetto route-map predefinito:

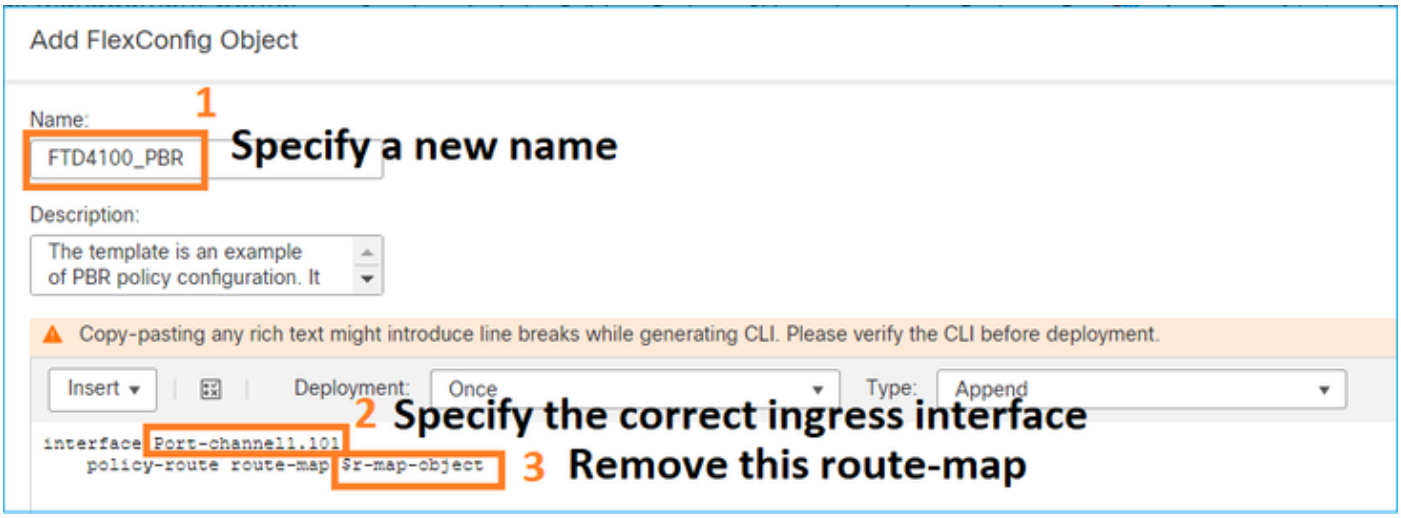

Specificare la nuova route-map:

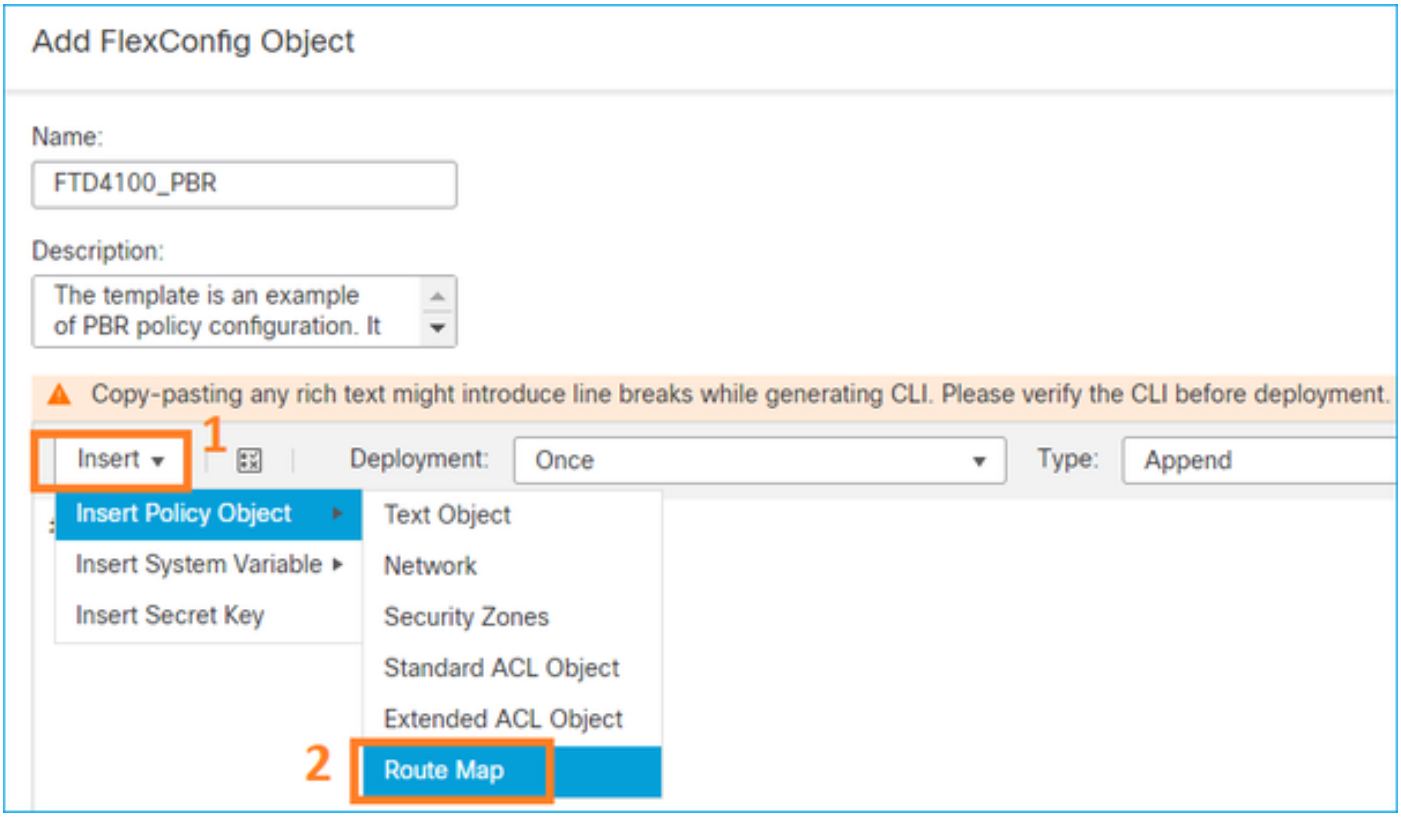

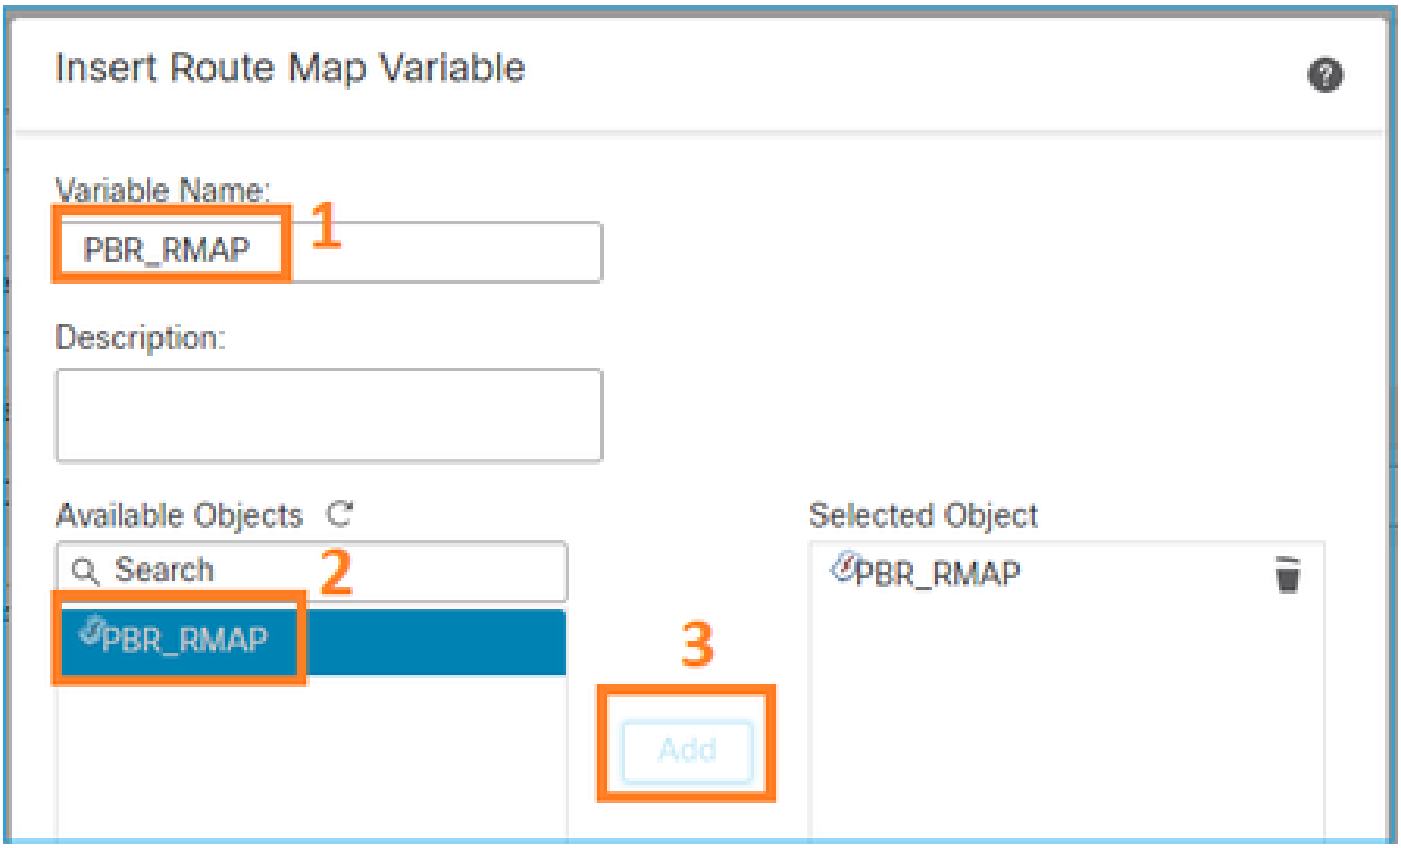

Questo è il risultato finale:

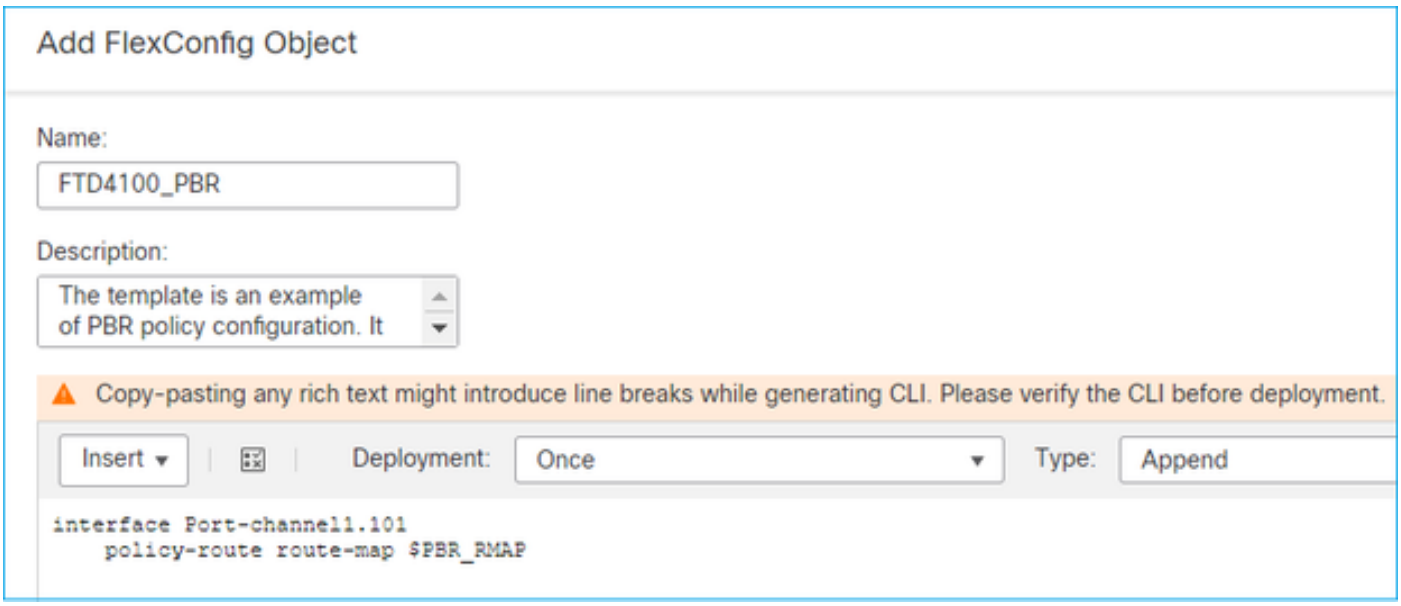

Passaggio 4 - Aggiungere l'oggetto PBR al criterio FTD FlexConfig.

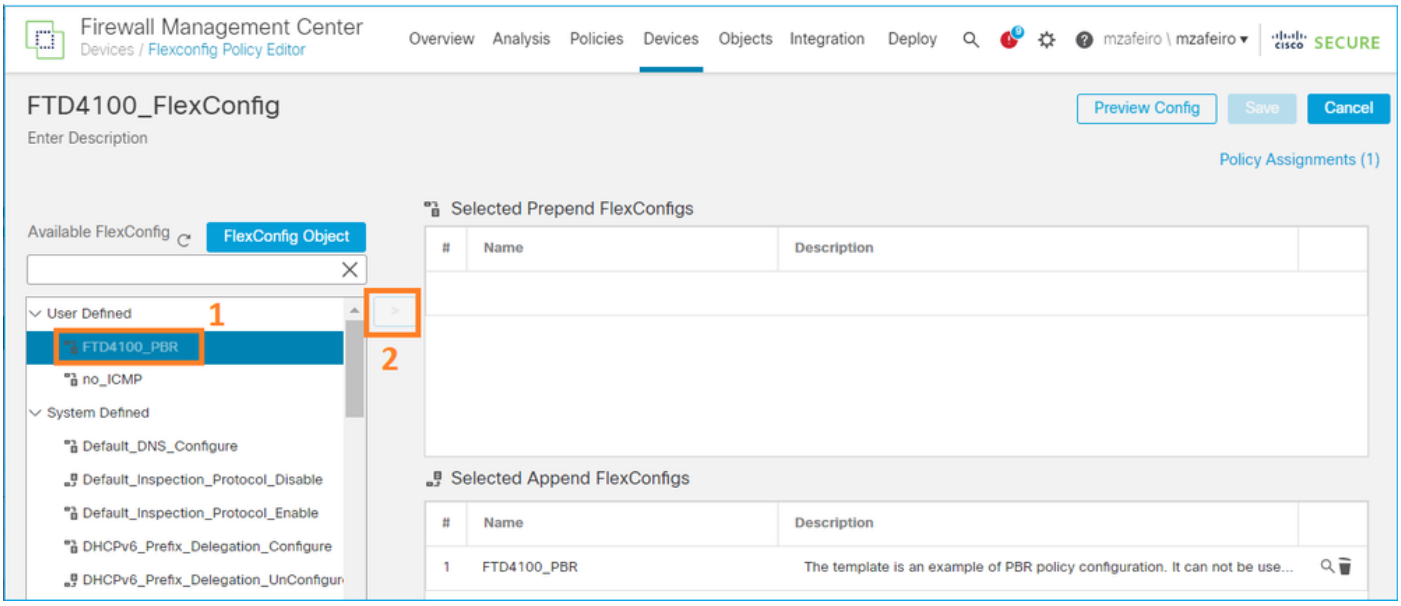

Salvare e selezionare Anteprima configurazione:

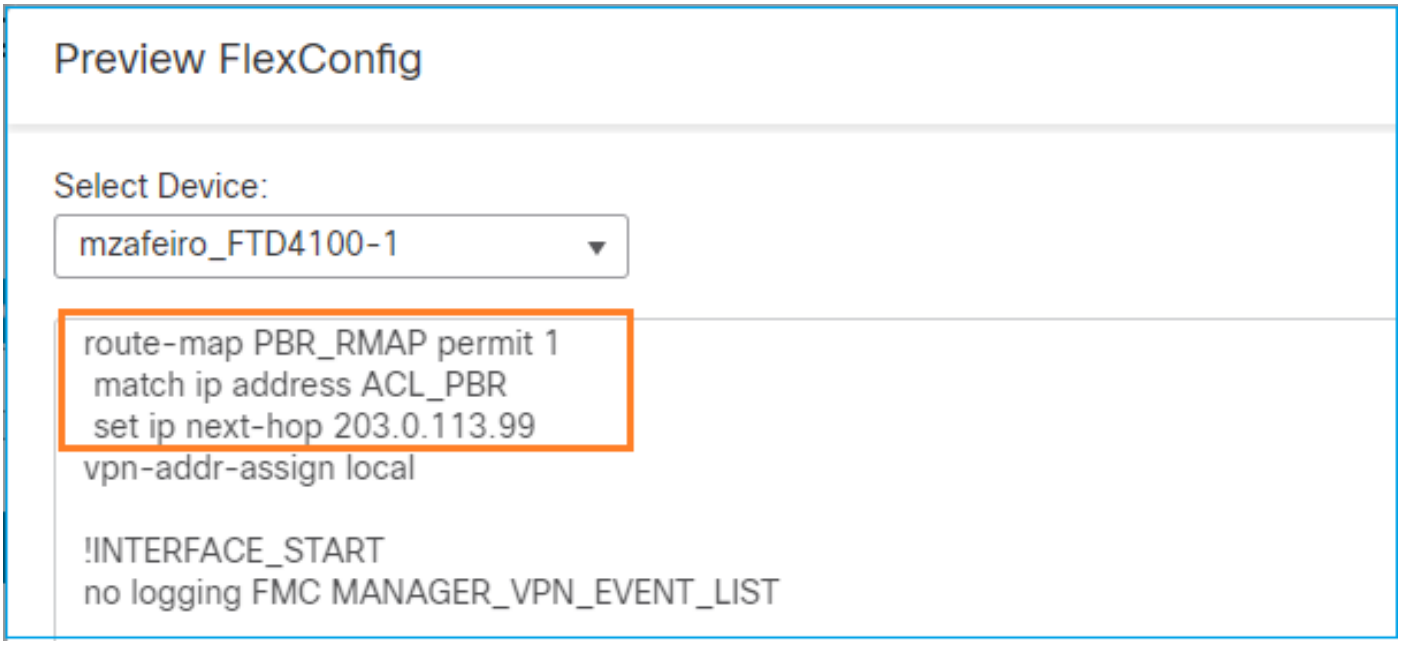

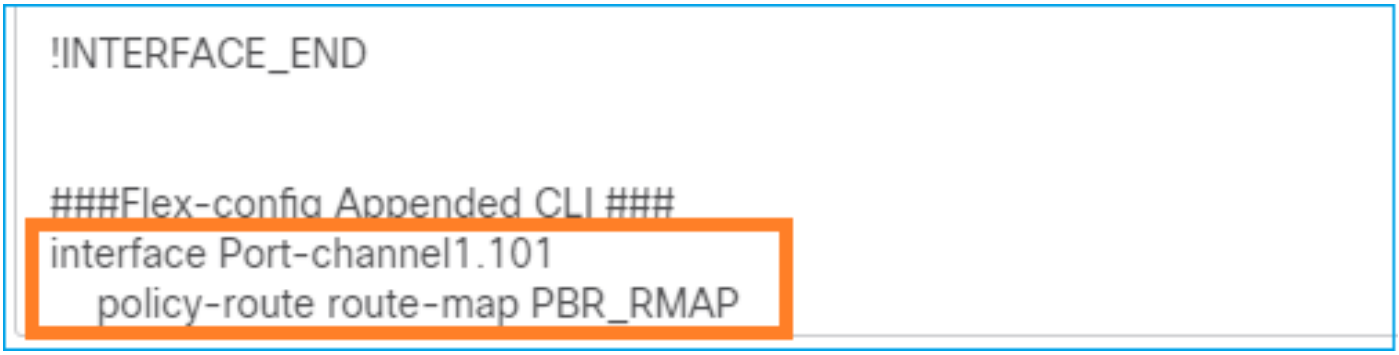

Distribuire infine il criterio.

Nota: non è possibile configurare PBR utilizzando FlexConfig e l'interfaccia utente di FMC per la stessa interfaccia in entrata.

Per la configurazione del contratto di servizio PBR, consultare il documento: [Configure PBR with IP SLAs for DUAL ISP on FTD Managed by FMC]

Verifica PBR

Verifica interfaccia in ingresso:

```
firepower# show run interface Po1.101
!
interface Port-channel1.101
vlan 101
nameif INSIDE
cts manual
propagate sgt preserve-untag
policy static sgt disabled trusted
security-level 0
ip address 192.168.0.1 255.255.255.0
policy-route route-map FMC_GENERATED_PBR_1649228271478
ospf authentication null
```

Verifica route-map:

```
firepower# show run route-map
!
route-map FMC_GENERATED_PBR_1649228271478 permit 5
 match ip address ACL_PBR
 set ip next-hop 203.0.113.99
```

```
firepower# show route-map
route-map FMC_GENERATED_PBR_1649228271478, permit, sequence 5
Match clauses:
ip address (access-lists): ACL_PBR

Set clauses:
adaptive-interface cost OUTSIDE1 (0)
```

Verifica route criteri:

```
firepower# show policy-route
Interface Route map
Port-channel1.101 FMC_GENERATED_PBR_1649228271478
```

Packet-Tracer prima e dopo la modifica:

Senza PBR	Con PBR
firepower# packet-tracer input INSIDE tcp 192.168.2.100 1111 198.51.100.5 23 .... Phase: 3 Type: INPUT-ROUTE-LOOKUP Subtype: Resolve Egress Interface Result: ALLOW Elapsed time: 11596 ns Config: Additional Information: Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0) ... Phase: 13 Type: INPUT-ROUTE-LOOKUP-FROM-OUTPUT-ROUTE-LOOKUP Subtype: Resolve Preferred Egress interface Result: ALLOW Elapsed time: 6244 ns Config: Additional Information: Found next-hop 192.0.2.99 using egress ifc OUTSIDE2(vrfid:0) Phase: 14 Type: ADJACENCY-LOOKUP Subtype: Resolve Nexthop IP address to MAC Result: ALLOW Elapsed time: 2230 ns Config: Additional Information: Found adjacency entry for Next-hop 192.0.2.99 on interface OUTSIDE2 Adjacency :Active MAC address 4c4e.35fc.fcd8 hits 0 reference 1 Result: input-interface: INSIDE(vrfid:0) input-status: up input-line-status: up output-interface: OUTSIDE2(vrfid:0) output-status: up output-line-status: up Action: allow Time Taken: 272058 ns	firepower# packet-tracer i ... Phase: 3 Type: SUBOPTIMAL-LOOKUP Subtype: suboptimal next-h Result: ALLOW Elapsed time: 39694 ns Config: Additional Information: Input route lookup returne Phase: 4 Type: ECMP load balancing Subtype: Result: ALLOW Elapsed time: 2230 ns Config: Additional Information: ECMP load balancing Found next-hop 203.0.113.9 Phase: 5 Type: PBR-LOOKUP Subtype: policy-route Result: ALLOW Elapsed time: 446 ns Config: route-map FMC_GENERATED_PB match ip address ACL_PBR set adaptive-interface cos Additional Information: Matched route-map FMC_GENE Found next-hop 203.0.113.9 ... Phase: 15 Type: ADJACENCY-LOOKUP Subtype: Resolve Nexthop I Result: ALLOW Elapsed time: 5352 ns Config: Additional Information: Found adjacency entry for Adjacency :Active MAC address 4c4e.35fc.fcd8 Result: input-interface: INSIDE(vr input-status: up input-line-status: up output-interface: OUTSIDE1 output-status: up output-line-status: up Action: allow Time Taken: 825100 ns

Test con traffico reale

Configurare l'acquisizione dei pacchetti con una traccia:

```
firepower# capture CAPI trace interface INSIDE match ip host 192.168.2.1 host 198.51.100.5
firepower# capture CAPO1 trace interface OUTSIDE1 match ip host 192.168.2.1 host 198.51.100.5
firepower# capture CAPO2 trace interface OUTSIDE2 match ip host 192.168.2.1 host 198.51.100.5
```

```
Router1# telnet 198.51.100.5 /vrf VRF-101 /source-interface lo2
Trying 198.51.100.5 ... Open
```

L'acquisizione mostra:

```
firepower# show capture
capture CAPI type raw-data trace interface INSIDE [Capturing - 4389 bytes]
match ip host 192.168.2.1 host 198.51.100.5
capture CAPO1 type raw-data trace interface OUTSIDE1 [Capturing - 4389 bytes]
match ip host 192.168.2.1 host 198.51.100.5
capture CAPO2 type raw-data trace interface OUTSIDE2 [Capturing - 0 bytes]
match ip host 192.168.2.1 host 198.51.100.5
```

Traccia del pacchetto TCP SYN:

```
firepower# show capture CAPI packet-number 1 trace

44 packets captured

1: 13:26:38.485585 802.1Q vlan#101 P0 192.168.2.1.49032 > 198.51.100.5.23: S 571152066:571152066(0) win
...

Phase: 3
Type: SUBOPTIMAL-LOOKUP
Subtype: suboptimal next-hop
Result: ALLOW
Elapsed time: 13826 ns
Config:
Additional Information:
Input route lookup returned ifc OUTSIDE2 is not same as existing ifc OUTSIDE1

Phase: 4
Type: ECMP load balancing
Subtype:
Result: ALLOW
Elapsed time: 1784 ns
Config:
Additional Information:
ECMP load balancing
```

```
Found next-hop 203.0.113.99 using egress ifc OUTSIDE1(vrfid:0)

Phase: 5
Type: PBR-LOOKUP
Subtype: policy-route
Result: ALLOW
Elapsed time: 446 ns
Config:
route-map FMC_GENERATED_PBR_1649228271478 permit 5
match ip address ACL_PBR
set adaptive-interface cost OUTSIDE1
Additional Information:
Matched route-map FMC_GENERATED_PBR_1649228271478, sequence 5, permit
Found next-hop 203.0.113.99 using egress ifc OUTSIDE1

...

Phase: 15
Type: ADJACENCY-LOOKUP
Subtype: Resolve Nexthop IP address to MAC
Result: ALLOW
Elapsed time: 4906 ns
Config:
Additional Information:
Found adjacency entry for Next-hop 203.0.113.99 on interface OUTSIDE1
Adjacency :Active
MAC address 4c4e.35fc.fcd8 hits 348 reference 2

...

Result:
input-interface: INSIDE(vrfid:0)
input-status: up
input-line-status: up
output-interface: OUTSIDE1(vrfid:0)
output-status: up
output-line-status: up
Action: allow
Time Taken: 222106 ns
```

Nella tabella ASP PBR sono illustrati i conteggi delle visite ai criteri:

```
firepower# show asp table classify domain pbr

Input Table
in id=0x1505f26d3420, priority=2147483642, domain=pbr, deny=false
hits=7, user_data=0x1505f26e7590, cs_id=0x0, use_real_addr, flags=0x0, protocol=0
src ip/id=192.168.2.0, mask=255.255.255.0, port=0, tag=any
dst ip/id=198.51.100.5, mask=255.255.255.255, port=0, tag=any, dscp=0x0, nsg_id=none
input_ifc=INSIDE(vrfid:0), output_ifc=any

Output Table:

L2 - Output Table:

L2 - Input Table:

Last clearing of hits counters: Never
```

Nota: il packet-tracer aumenta anche il contatore visite.

Debug PBR

Avviso: in un ambiente di produzione, il debug può produrre molti messaggi.

Abilita debug:

```
firepower# debug policy-route
debug policy-route enabled at level 1
```

Invia traffico reale:

```
Router1# telnet 198.51.100.5 /vrf VRF-101 /source-interface lo2
Trying 198.51.100.5 ... Open
```

Il debug mostra:

```
firepower#

pbr: policy based route lookup called for 192.168.2.1/37256 to 198.51.100.5/23 proto 6 sub_proto 0 rece
pbr: First matching rule from ACL(2)
pbr: route map FMC_GENERATED_PBR_1649228271478, sequence 5, permit; proceed with policy routing
pbr: policy based routing applied; egress_ifc = OUTSIDE1 : next_hop = 203.0.113.99
```

Nota: Packet-tracer genera anche un output di debug.

Questo diagramma di flusso può essere utilizzato per la risoluzione dei problemi relativi al PBR:

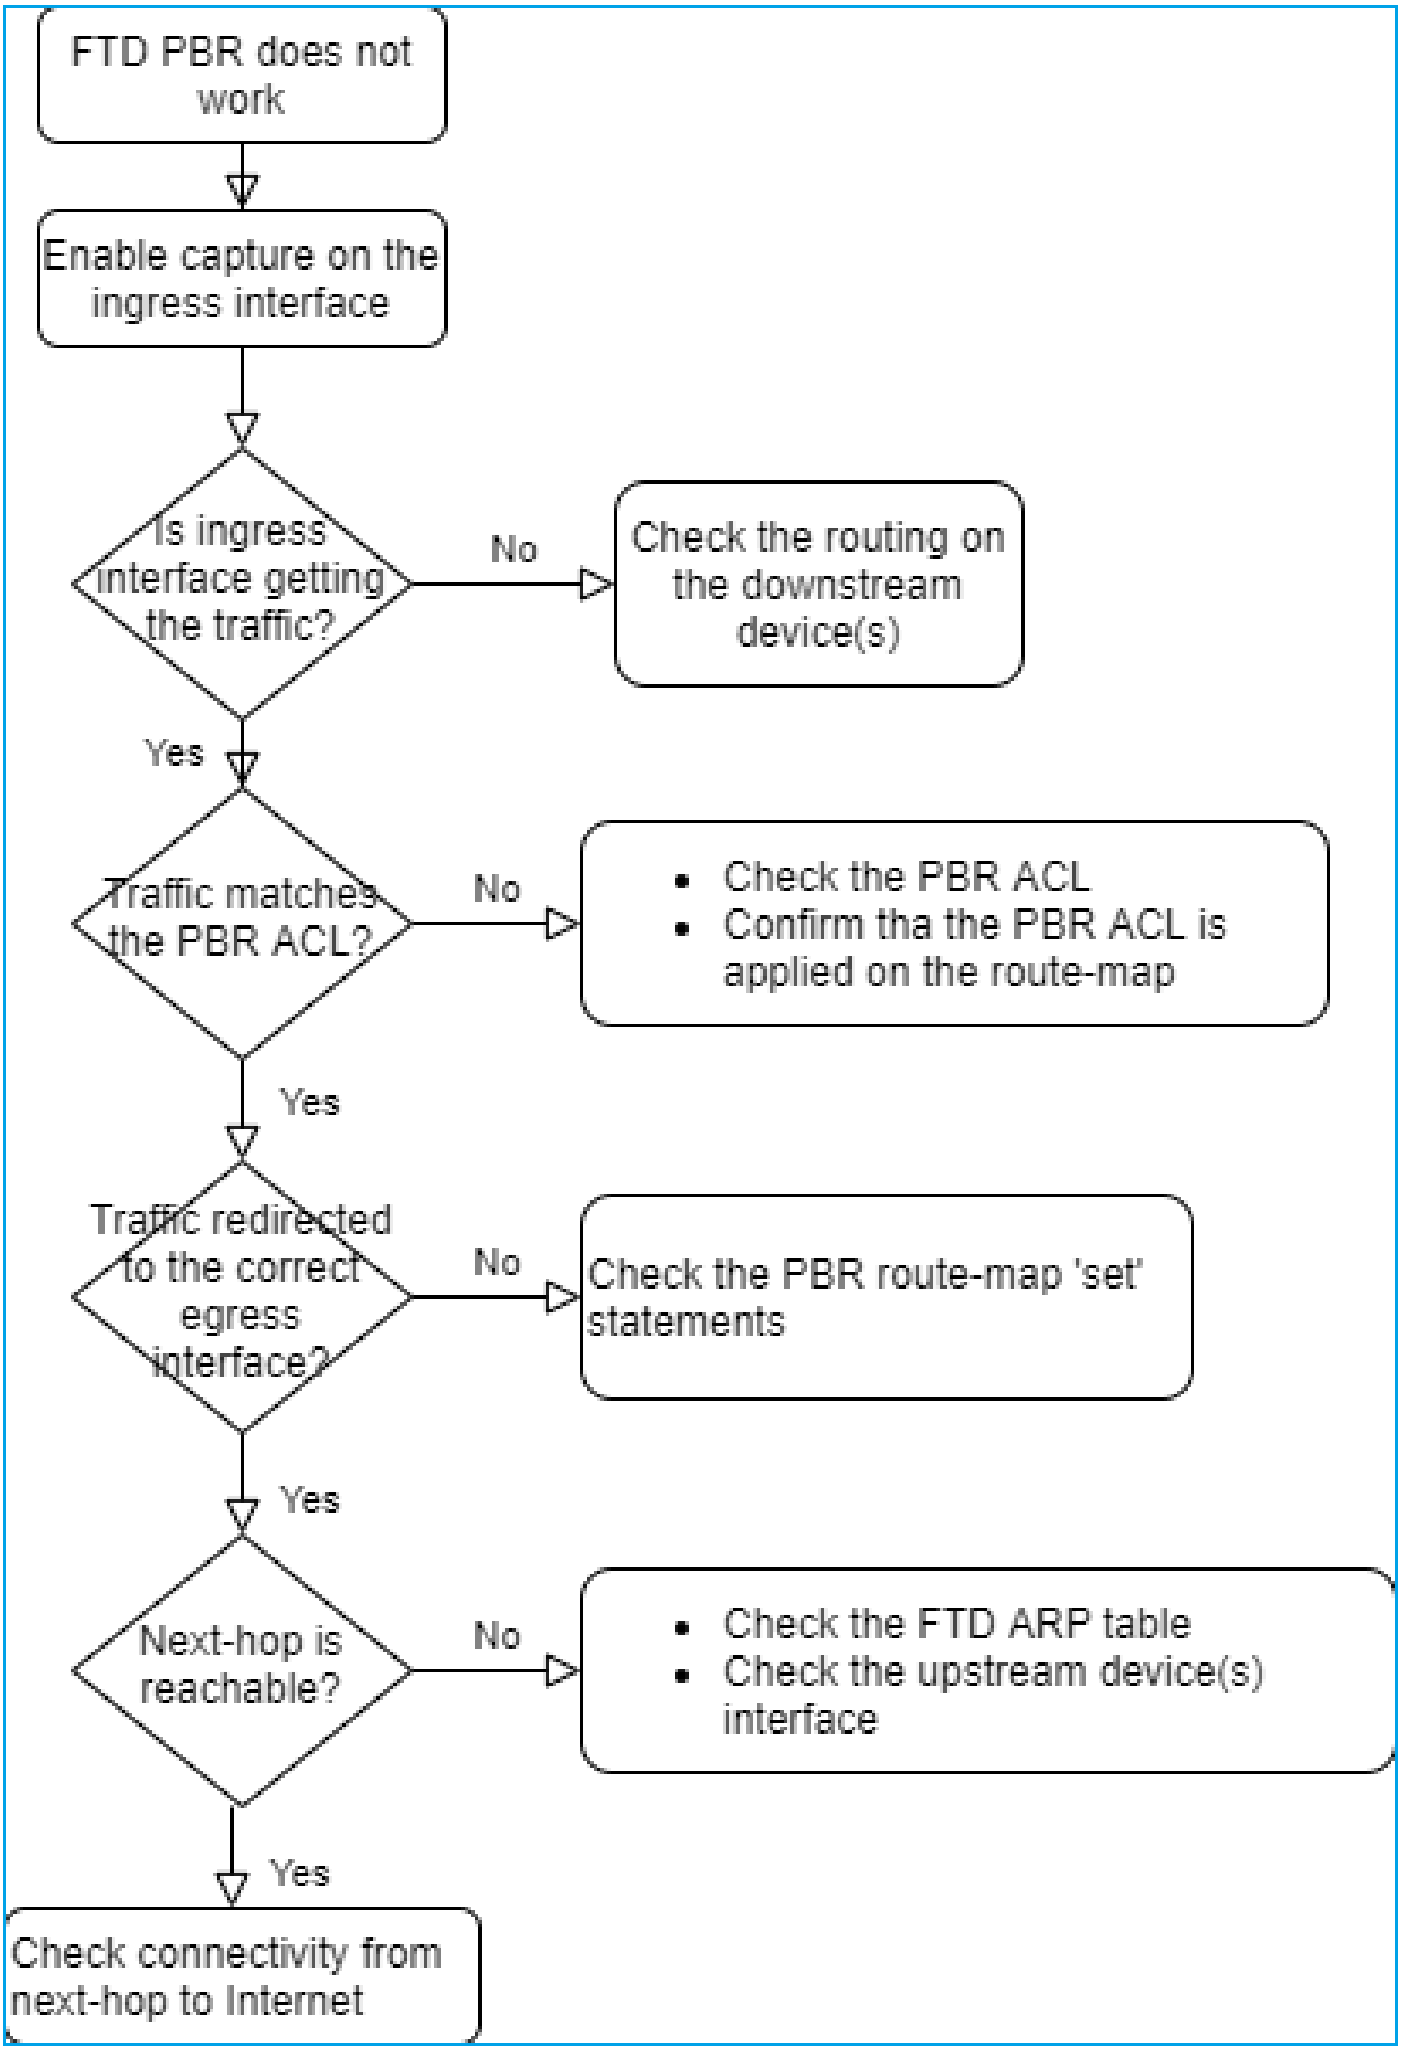

Riepilogo dei comandi PBR

```
show asp drop
```

Caso 4 - Inoltro basato sulla ricerca di routing globale

Dopo la ricerca della connessione, la ricerca NAT e PBR, l'ultimo elemento controllato per determinare l'interfaccia di uscita è la tabella di routing globale.

Verifica tabella di routing

Esaminiamo l'output di una tabella di routing FTD:

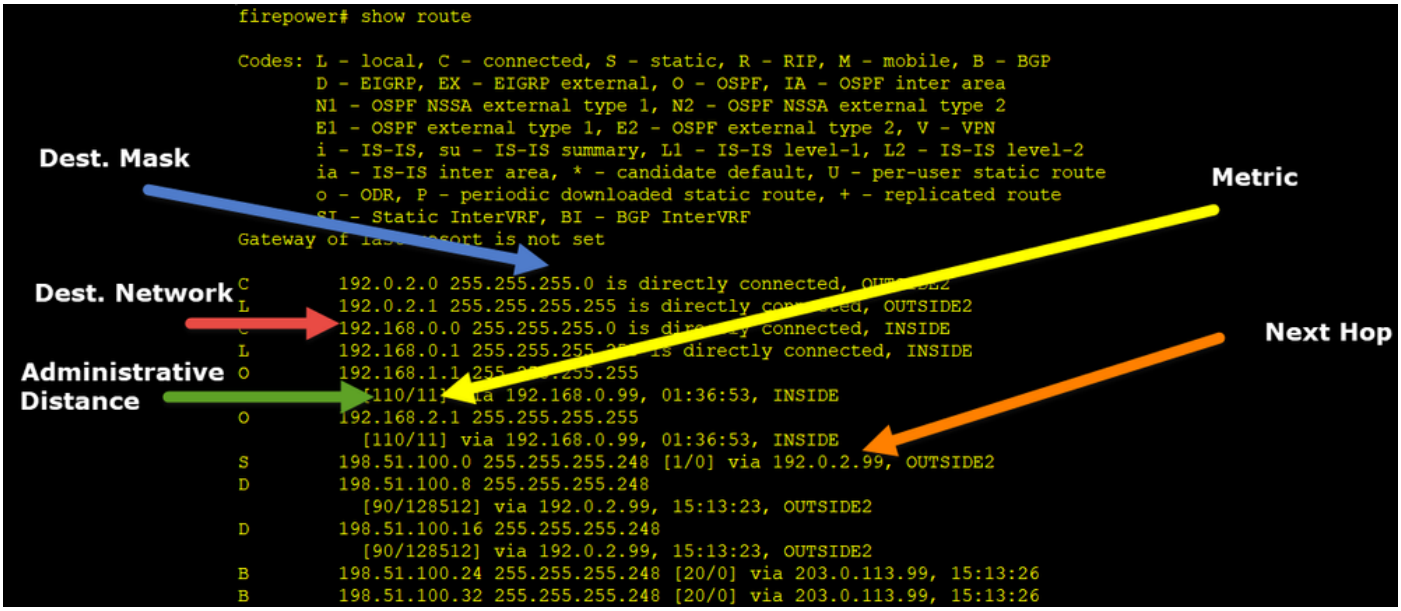

L'obiettivo principale del processo di routing è trovare l'hop successivo. La route selezionata è nell'ordine seguente:

1. Vince partita più lunga
2. AD minimo (tra diverse origini del protocollo di routing)
3. Metrica minima (nel caso in cui i percorsi vengano appresi dalla stessa origine, ovvero dal protocollo di routing)

Modalità di popolamento della tabella di routing:

- IGP (R, D, EX, O, IA, N1, N2, E1, E2, i, su, L1, L2, ia, o)

- BGP (B)

- BGP InterVRF (BI)

- Statico (S)

- InterVRF statica (SI)

- Connesso (C)

- IP locali (L)

- VPN (V)

-Ridistribuzione

-Predefinito

Per visualizzare il riepilogo della tabella di routing, utilizzare questo comando:

<#root>

```
firepower#

show route summary



IP routing table maximum-paths is 8
Route Source   Networks Subnets Replicates Overhead Memory (bytes)
connected      0        8       0          704      2368
static         0        1       0          88       296
ospf 1         0        2       0          176      600
Intra-area: 2 Inter-area: 0 External-1: 0 External-2: 0
NSSA External-1: 0 NSSA External-2: 0
bgp 65000      0        2       0          176      592
External: 2 Internal: 0 Local: 0
eigrp 1        0        2       0          216      592
internal       7                                    3112

Total          7        15      0          1360     7560
```

È possibile tenere traccia degli aggiornamenti della tabella di routing con questo comando:

<#root>

```
firepower#

debug ip routing



IP routing debugging is on
```

Ad esempio, questo è quanto mostra il debug quando la route OSPF 192.168.1.0/24 viene rimossa dalla tabella di routing globale:

<#root>

```
firepower#
```

```
RT: ip_route_delete 192.168.1.0 255.255.255.0 via 192.0.2.99, INSIDE

ha_cluster_synced 0 routetype 0
RT: del 192.168.1.0 via 192.0.2.99, ospf metric [110/11]NP-route: Delete-Output 192.168.1.0/24 hop_coun
RT: delete network route to 192.168.1.0 255.255.255.0NP-route: Delete-Output 192.168.1.0/24 hop_count:1
NP-route: Delete-Input 192.168.1.0/24 hop_count:1 Distance:110 Flags:0X0 , via 0.0.0.0, INSIDE
```

Quando viene aggiunto di nuovo:

<#root>

```
firepower#

RT: NP-route: Add-Output 192.168.1.0/24 hop_count:1 , via 192.0.2.99, INSIDE

NP-route: Add-Input 192.168.1.0/24 hop_count:1 Distance:110 Flags:0X0 , via 192.0.2.99, INSIDE
```

Interfaccia Null0

L'interfaccia Null0 può essere utilizzata per eliminare il traffico indesiderato. Questo rilascio ha un impatto minore sulle prestazioni rispetto al calo del traffico con una regola ACL (Access Control Policy).

Requisito

Configurare una route Null0 per l'host 198.51.100.4/32.

Soluzione

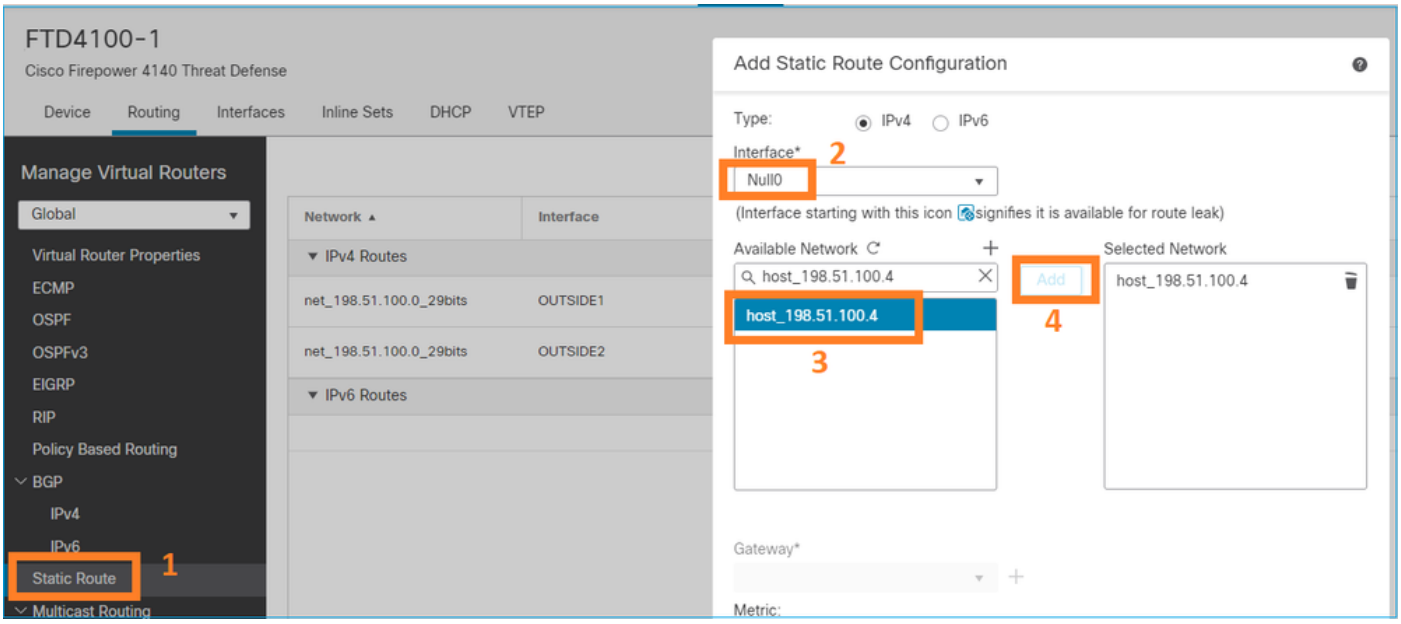

Salva e distribuisci.

Verifica:

<#root>

```
firepower#

show run route


route OUTSIDE2 198.51.100.0 255.255.255.248 192.0.2.99 1
route OUTSIDE1 198.51.100.0 255.255.255.248 203.0.113.99 200

route Null0 198.51.100.4 255.255.255.255 1
```

<#root>

```
firepower#

show route | include 198.51.100.4


S 198.51.100.4 255.255.255.255 [1/0] is directly connected, Null0
```

Provare ad accedere all'host remoto:

<#root>

```
Router1#

ping vrf VRF-101 198.51.100.4


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 198.51.100.4, timeout is 2 seconds:

.....


Success rate is 0 percent (0/5)
```

I log FTD mostrano:

<#root>

```
firepower#

show log | include 198.51.100.4
```

```
Apr 12 2022 12:35:28:
%FTD-6-110002: Failed to locate egress interface for ICMP from INSIDE:192.168.0.99/0 to 198.51.100.4/0
```

Le interruzioni ASP mostrano:

<#root>

```
firepower#
show asp drop

Frame drop:
No route to host (no-route)                1920
```

Equal Cost Multi-Path (ECMP)

Zone di traffico

- La zona di traffico ECMP consente a un utente di raggruppare le interfacce (nota come zona ECMP).
- Ciò consente il routing ECMP e il bilanciamento del carico del traffico su più interfacce.
- Quando le interfacce sono associate alla zona traffico ECMP, l'utente è in grado di creare percorsi statici pari costo attraverso le interfacce. Le route statiche pari costo sono route verso la stessa rete di destinazione con lo stesso valore di metrica.

Prima della versione 7.1, Firepower Threat Defense supportava il routing ECMP tramite i criteri FlexConfig. A partire dalla versione 7.1, è possibile raggruppare le interfacce in zone di traffico e configurare il routing ECMP in Firepower Management Center.

EMCP è documentato in: ECMP

Nell'esempio, viene rilevato un routing asimmetrico e il traffico di ritorno viene interrotto:

<#root>

```
firepower#
show log

Apr 13 2022 07:20:48: %FTD-6-302013:
B
uilt inbound TCP connection 4046 for INSIDE:192.168.1.1/23943 (192.168.1.1/23943) to OUTSIDE1:198.51.100
```

```
Apr 13 2022 07:20:48: %FTD-6-106015:
Deny TCP (no connection) from 198.51.100.100/23 to 192.168.1.1/23943 flags SYN ACK on interface OUTSIDE2
```

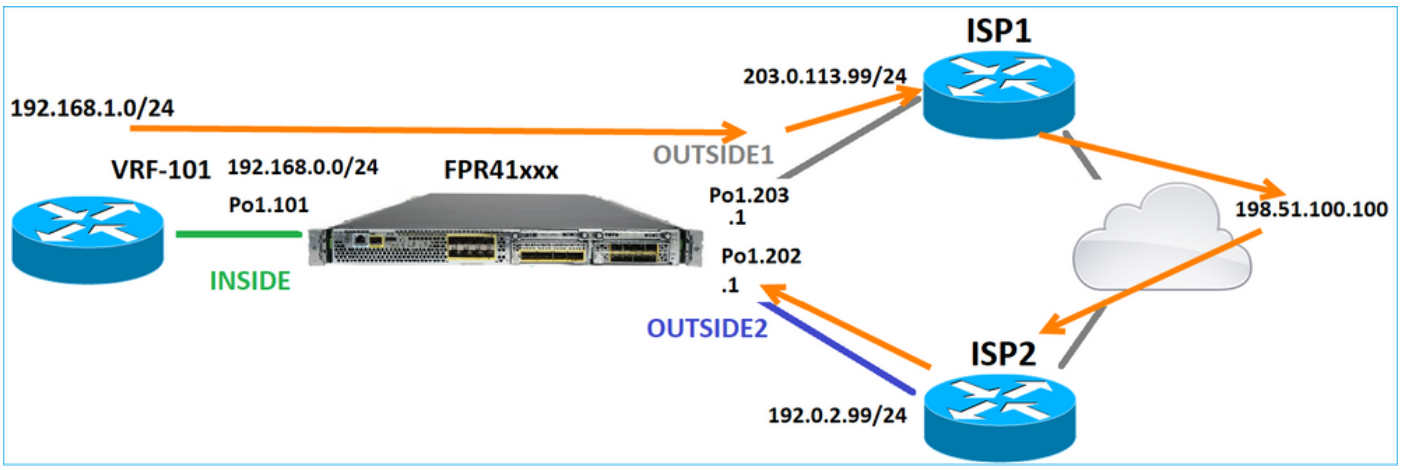

Configurare ECMP dall'interfaccia utente di FMC:

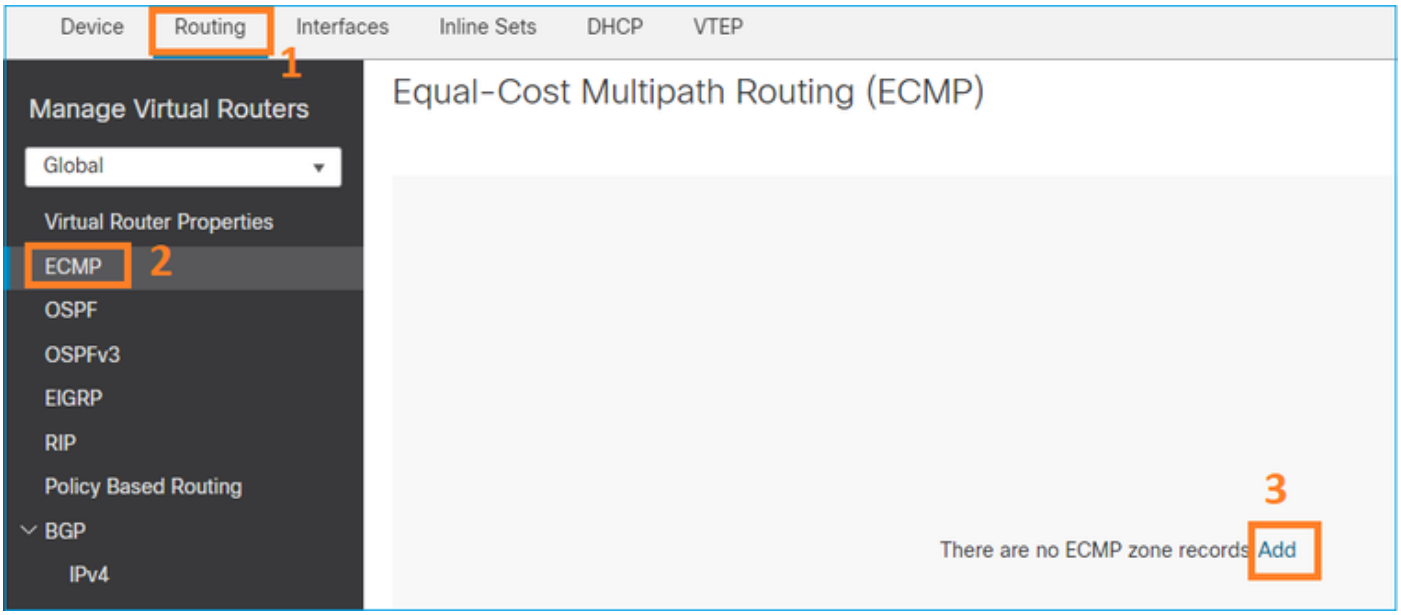

Aggiungere le due interfacce nel gruppo ECMP:

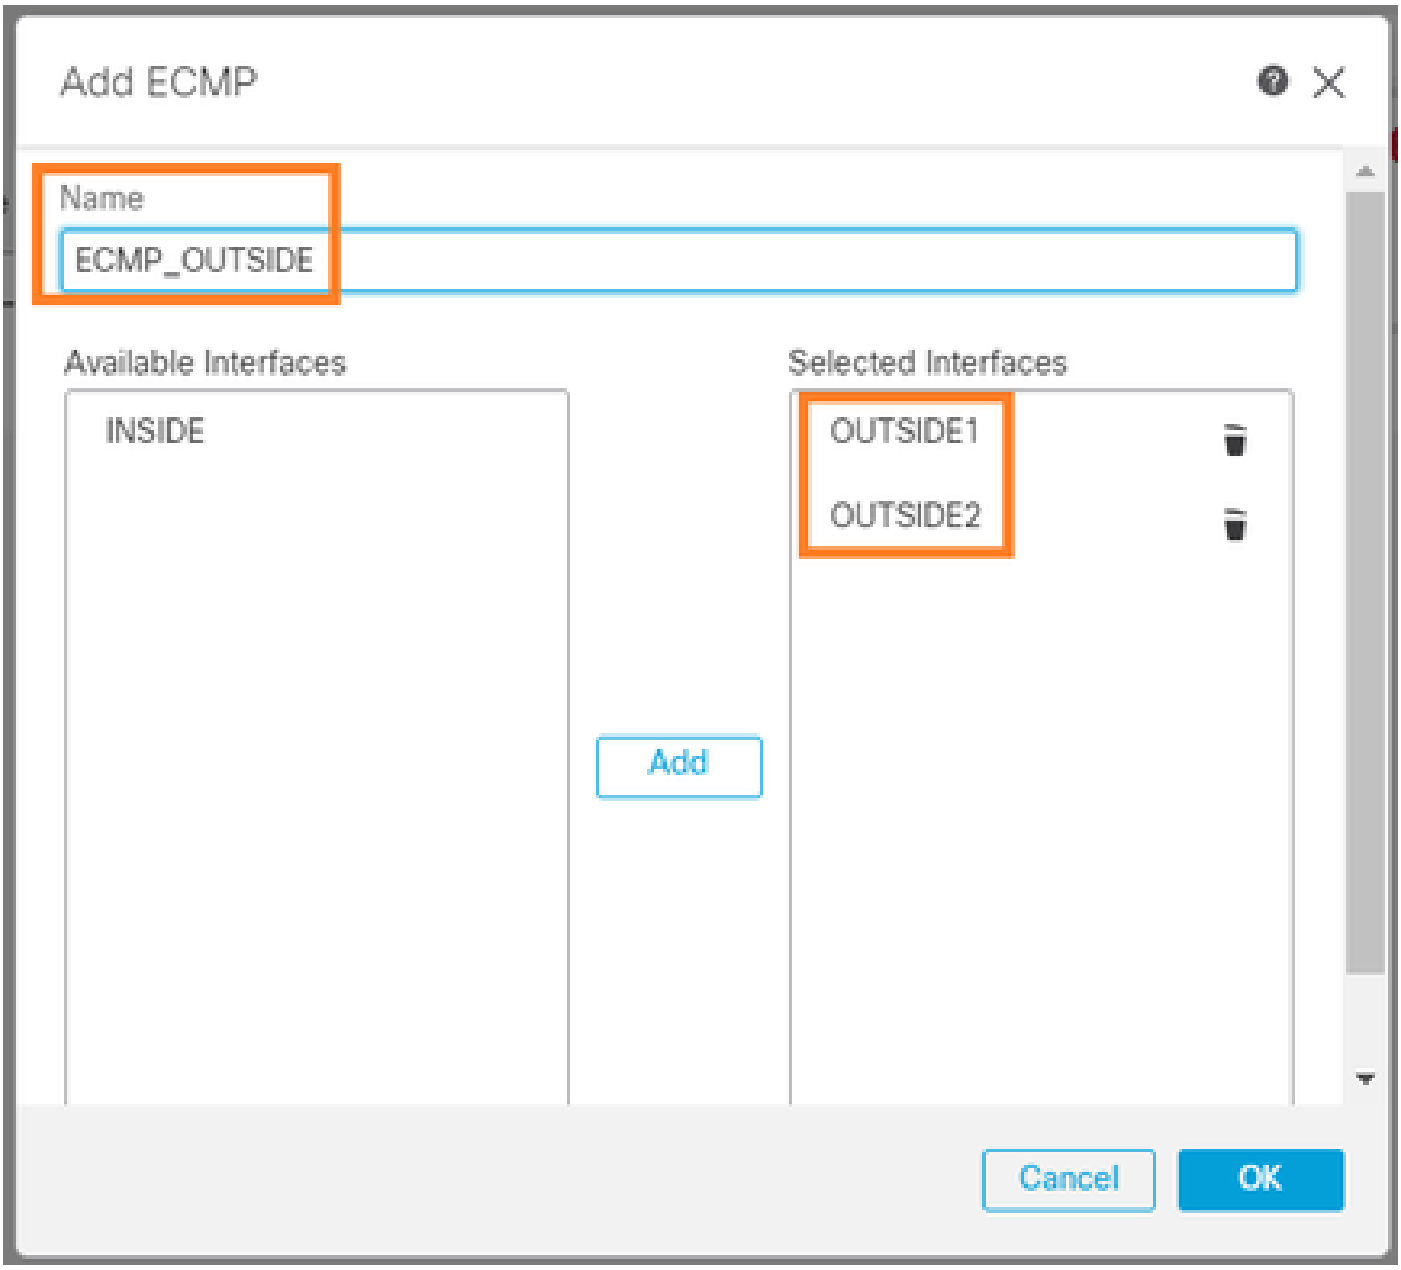

Il risultato:

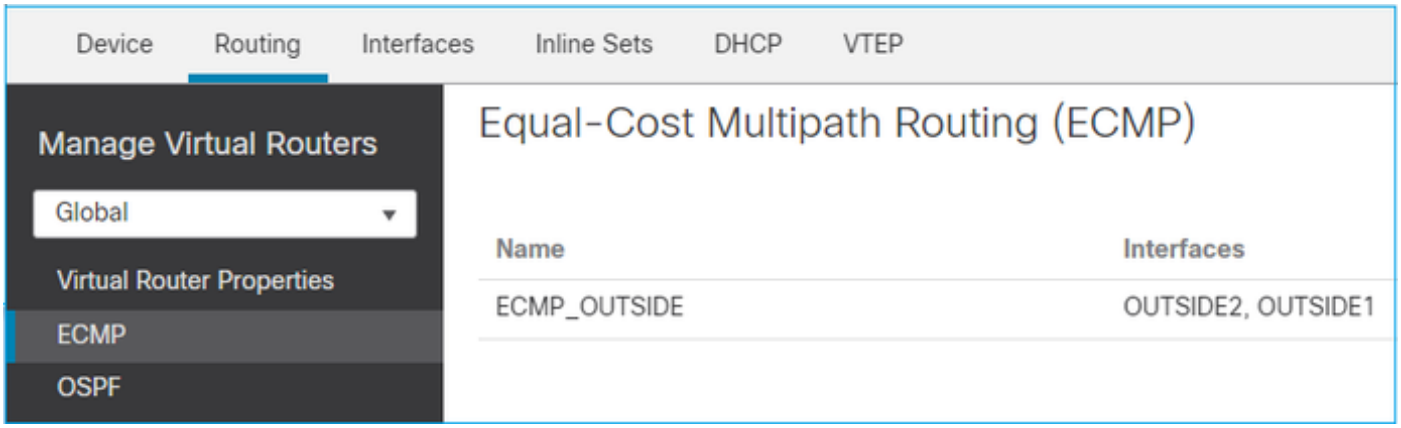

Salva e distribuisci.

Verifica della zona ECMP:

<#root>

```
firepower#
show run zone

zone ECMP_OUTSIDE ecmp

firepower#
show zone

Zone: ECMP_OUTSIDE ecmp

Security-level: 0

Zone member(s): 2

OUTSIDE1 Port-channel1.203

OUTSIDE2 Port-channel1.202
```

Verifica interfaccia:

<#root>

```
firepower#
show run int po1.202

!
interface Port-channel1.202
vlan 202
nameif OUTSIDE2
cts manual
propagate sgt preserve-untag
policy static sgt disabled trusted
security-level 0
zone-member ECMP_OUTSIDE

ip address 192.0.2.1 255.255.255.0

firepower#
show run int po1.203
```

```
!
interface Port-channel1.203
vlan 203
nameif OUTSIDE1
cts manual
propagate sgt preserve-untag
policy static sgt disabled trusted
security-level 0
```

zone-member ECMP_OUTSIDE

```
ip address 203.0.113.1 255.255.255.0
```

A questo punto, il traffico di ritorno è autorizzato e la connessione è attiva:

<#root>

```
Router1#
```

telnet 198.51.100.100 /vrf VRF-101 /source-interface lo1

Trying 198.51.100.100 ... Open

L'opzione Acquisisci su interfaccia ISP1 visualizza il traffico in uscita:

<#root>

```
firepower#
```

show capture CAP1

```
5 packets captured

1: 10:03:52.620115 802.1Q vlan#203 P0 192.168.1.1.56199 > 198.51.100.100.23: S 1782458734:1782458734(0)
2: 10:03:52.621992 802.1Q vlan#203 P0 192.168.1.1.56199 > 198.51.100.100.23: . ack 2000807246 win 4128
3: 10:03:52.622114 802.1Q vlan#203 P0 192.168.1.1.56199 > 198.51.100.100.23: . ack 2000807246 win 4128
4: 10:03:52.622465 802.1Q vlan#203 P0 192.168.1.1.56199 > 198.51.100.100.23: P 1782458735:1782458753(18
5: 10:03:52.622556 802.1Q vlan#203 P0 192.168.1.1.56199 > 198.51.100.100.23: . ack 2000807246 win 4128
```

L'acquisizione sull'interfaccia ISP2 mostra il traffico di ritorno:

<#root>

```
firepower#
```

show capture CAP2

```
6 packets captured

1: 10:03:52.621305 802.1Q vlan#202 P0 198.51.100.100.23 > 192.168.1.1.56199:

S

 2000807245:2000807245(0)

ack

 1782458735 win 64240 <mss 1460>
3: 10:03:52.623808 802.1Q vlan#202 P0 198.51.100.100.23 > 192.168.1.1.56199: . ack 1782458753 win 64222
```

Piano di gestione FTD

L'FTD dispone di 2 piani di gestione:

- Interfaccia Management0 - Fornisce l'accesso al sottosistema Firepower
- Interfaccia diagnostica LINA - Accesso al sottosistema FTD LINA

Per configurare e verificare l'interfaccia Management0, utilizzare rispettivamente i comandi configure network e show network.

D'altra parte, le interfacce LINA forniscono accesso alla LINA stessa. Le voci dell'interfaccia FTD nella RIB FTD possono essere visualizzate come route locali:

<#root>

```
firepower#

show route | include L


Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
L 192.0.2.1 255.255.255.255 is directly connected, OUTSIDE2
L 192.168.0.1 255.255.255.255 is directly connected, INSIDE
L 203.0.113.1 255.255.255.255 is directly connected, OUTSIDE1
```

Analogamente, possono essere viste come voci di identità nella tabella di routing ASP:

<#root>

```
firepower#

show asp table routing | include identity


in 169.254.1.1 255.255.255.255 identity
in

192.0.2.1 255.255.255.255 identity
```

```
in
203.0.113.1 255.255.255.255 identity

in
192.168.0.1 255.255.255.255 identity

in ff02::1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in ff02::1:ff01:3 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in ff02::1:ff00:1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in fe80::200:ff:fe01:3 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
in fd00:0:0:1::1 ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff identity
out 0.0.0.0 0.0.0.0 via 0.0.0.0, identity
out :: :: via 0.0.0.0, identity
```

Punto principale

Quando un pacchetto arriva su un FTD e l'IP di destinazione corrisponde a uno degli IP di identità, l'FTD sa che deve consumare il pacchetto.

Routing interfaccia diagnostica LINA FTD

L'FTD (come le appliance ASA con codice successivo alla 9.5) gestisce una tabella di routing simile al VRF per tutte le interfacce configurate come sola gestione. Un esempio di interfaccia di questo tipo è l'interfaccia diagnostica.

Sebbene FMC non consenta (senza ECMP) di configurare 2 route predefinite su 2 interfacce diverse con la stessa metrica, è possibile configurare 1 route predefinita su un'interfaccia dati FTD e un'altra route predefinita sull'interfaccia diagnostica:

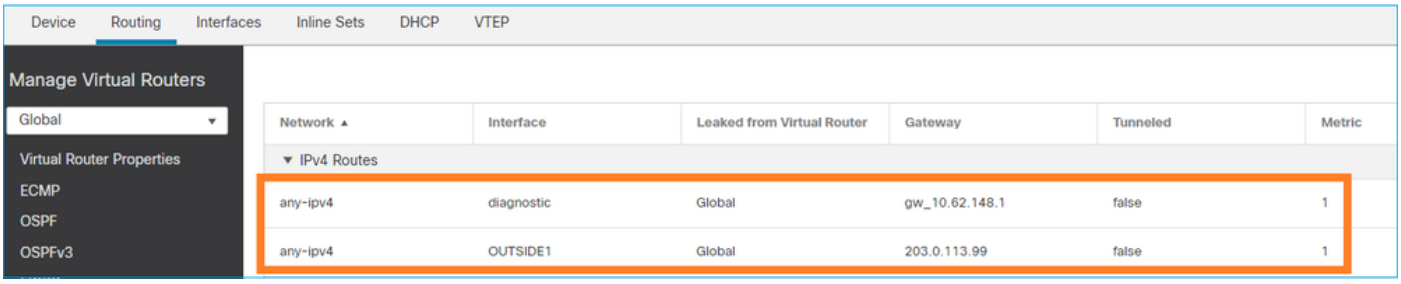

Il traffico del piano dati utilizza il gateway predefinito della tabella globale, mentre il traffico del piano di gestione utilizza il GW predefinito di diagnostica:

<#root>

```
firepower#
show route management-only
```

```
Routing Table: mgmt-only


Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, V - VPN
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, + - replicated route
SI - Static InterVRF, BI - BGP InterVRF

Gateway of last resort is 10.62.148.1 to network 0.0.0.0



S* 0.0.0.0 0.0.0.0 [1/0] via 10.62.148.1, diagnostic
```

Il gateway della tabella di routing globale:

<#root>

```
firepower#

show route | include S\*|Gateway


Gateway of last resort is 203.0.113.99 to network 0.0.0.0


S* 0.0.0.0 0.0.0.0 [1/0] via 203.0.113.99, OUTSIDE1
```

Quando si invia il traffico dall'FTD (traffico preconfigurato), l'interfaccia di uscita viene selezionata in base a:

1. Tabella di routing globale
2. Tabella di routing di sola gestione

È possibile sovrascrivere la selezione dell'interfaccia di uscita se questa viene specificata manualmente.

Provare a eseguire il ping del gateway dell'interfaccia di diagnostica. Se non si specifica l'interfaccia di origine, il ping ha esito negativo perché l'FTD utilizza prima la tabella di routing globale che, in questo caso, contiene un percorso predefinito. Se nella tabella globale non è presente alcuna route, l'FTD esegue una ricerca della route nella tabella di routing di sola gestione:

<#root>

```
firepower#
```

```
ping 10.62.148.1


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.62.148.1, timeout is 2 seconds:

?????


Success rate is 0 percent (0/5)
firepower#

show capture CAP1 | include 10.62.148.1


1: 10:31:22.970607 802.1Q vlan#203 P0

203.0.113.1 > 10.62.148.1 icmp: echo request


2: 10:31:22.971431 802.1Q vlan#203 P0

10.1.1.2 > 203.0.113.1 icmp: host 10.62.148.1 unreachable
```

<#root>

```
firepower#

ping diagnostic 10.62.148.1


Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.62.148.1, timeout is 2 seconds:

!!!!!


Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Lo stesso vale se si tenta di copiare un file da LINA CLI con il comando copy.

Rilevamento inoltro bidirezionale (BFD)

Il supporto BFD è stato aggiunto sulla versione ASA 9.6 classica e solo per il protocollo BGP: Routing di rilevamento inoltro bidirezionale

Su FTD:

- Sono supportati i protocolli BGP IPv4 e BGP IPv6 (software 6.4).
- I protocolli OSPFv2, OSPFv3 e EIGRP non sono supportati.
- BFD per route statiche non supportato.

Router virtuali (VRF)

Il supporto VRF è stato aggiunto nella versione 6.6. Per ulteriori informazioni, consultare il documento: Esempi di configurazione per i router virtuali

Informazioni correlate

- Route statiche e predefinite FTD