

Configurazione del multicast inter-VRF senza perdite di route unicast inter-VRF

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Esempio di rete](#)

[Configurazioni](#)

[Verifica](#)

[Risoluzione dei problemi](#)

Introduzione

In questo documento viene descritto come inoltrare un flusso multicast che proviene da un'origine in un'origine VRF (Virtual Routing and Forwarding) e va verso gli utenti nel ricevitore VRF senza causare perdite di percorso unicast tra VRF.

Prerequisiti

Requisiti

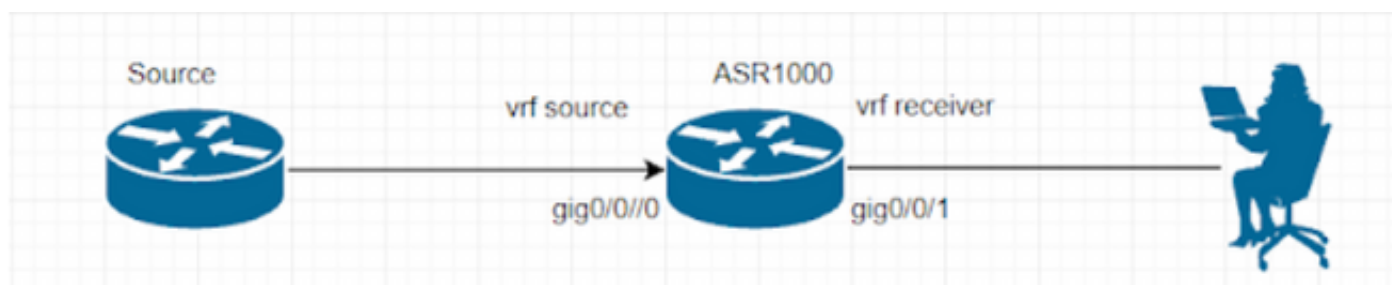
Cisco raccomanda la conoscenza delle funzionalità PIM e ASM.

Componenti usati

Il riferimento delle informazioni contenute in questo documento è ASR1000.

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.

Esempio di rete



Configurazioni

Genera flusso multicast UDP sull'origine.

```
Source#  
  
ip sla 1  
  
udp-echo 239.1.1.1 2000 source-ip 10.1.1.1  
  
ip sla schedule 1 life forever start-time now
```

Il ricevitore è configurato con IGMP e invia i rapporti di appartenenza a IGMP.

Il router ASR1000 sta eseguendo l'inoltro del traffico multicast tra vrf. Il loopback 1 in 'origine' VRF è l'RP per entrambi i VRF.

```
ASR1000#  
  
ip vrf source  
  
rd 1:1  
  
!  
  
ip vrf receiver  
  
rd 2:2  
  
!  
  
ip multicast-routing vrf source distributed  
  
ip multicast-routing vrf receiver distributed  
  
  
ip pim vrf source rp-address 1.1.1.1  
  
ip pim vrf receiver rp-address 1.1.1.1  
  
  
interface Loopback1  
  
ip vrf forwarding source  
  
ip address 1.1.1.1 255.255.255.255  
  
ip pim sparse-mode  
  
  
interface GigabitEthernet0/0/0  
  
ip vrf forwarding source  
  
ip address 10.1.1.2 255.255.255.0  
  
ip pim sparse-mode
```

```
interface GigabitEthernet0/0/1
ip vrf forwarding receiver
ip address 20.1.1.2 255.255.255.0
ip pim sparse-mode
```

(S, G) voce non formata in vrf 'receive' per il momento.

```
ASR1002-1#show ip mroute vrf source
```

```
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
```

```
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
```

```
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
```

```
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
```

```
U - URD, I - Received Source Specific Host Report,
```

```
Z - Multicast Tunnel, z - MDT-data group sender,
```

```
Y - Joined MDT-data group, y - Sending to MDT-data group,
```

```
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
```

```
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
```

```
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
```

```
V - RD & Vector, v - Vector, p - PIM Joins on route,
```

```
x - VxLAN group, c - PFP-SA cache created entry,
```

```
* - determined by Assert
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
```

```
Timers: Uptime/Expires
```

```
Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(* , 239.1.1.1), 00:05:01/stopped, RP 1.1.1.1, flags: SPF
```

```
Incoming interface: Null, RPF nbr 0.0.0.0
```

```
Outgoing interface list: Null
```

```
(10.1.1.1, 239.1.1.1), 00:05:01/00:02:26, flags: PFT
```

```
Incoming interface: GigabitEthernet0/0/0, RPF nbr 0.0.0.0
```

```
Outgoing interface list: Null
```

```
(* , 224.0.1.40), 00:07:03/00:02:59, RP 1.1.1.1, flags: SJCL
```

```
Incoming interface: Null, RPF nbr 0.0.0.0
```

Outgoing interface list:

Loopback1, Forward/Sparse, 00:07:01/00:02:59

```
ASR1002-X-1#sh ip mroute vrf receiver
```

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,

L - Local, P - Pruned, R - RP-bit set, F - Register flag,

T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,

X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,

U - URD, I - Received Source Specific Host Report,

Z - Multicast Tunnel, z - MDT-data group sender,

Y - Joined MDT-data group, y - Sending to MDT-data group,

G - Received BGP C-Mroute, g - Sent BGP C-Mroute,

N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,

Q - Received BGP S-A Route, q - Sent BGP S-A Route,

V - RD & Vector, v - Vector, p - PIM Joins on route,

x - VxLAN group, c - PFP-SA cache created entry,

* - determined by Assert, # - iif-starg configured on rpf intf,

e - encap-helper tunnel flag

Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

```
(*, 239.1.1.1), 00:03:23/00:02:44, RP 1.1.1.1, flags: SJC
```

Incoming interface: Null, RPF nbr 0.0.0.0

Outgoing interface list:

GigabitEthernet0/0/1, Forward/Sparse, 00:03:23/00:02:44

Per risolvere questo problema, è necessario selezionare un VRF in modo da superare il controllo RPF.

```
ip access-list standard 1
```

```
permit 239.1.1.1 log
```

```
exit
```

```
ip multicast vrf receiver rpf select vrf source group-list 1
```

Verifica

Per verificare che la configurazione funzioni correttamente, consultare questa sezione.

Dopo aver eseguito i passaggi descritti in precedenza, è possibile visualizzare l'avanzamento Multicast sul ricevitore.

```
Receiver#show flow monitor test cache format table
```

Cache type: Normal (Platform cache)

Cache size: 200000

Current entries: 1

High Watermark: 3

Flows added: 50

Flows aged: 49

- Inactive timeout (15 secs) 49

IPV4 SRC ADDR IPV4 DST ADDR TRNS SRC PORT TRNS DST PORT INTF INPUT FLOW SAMPLER ID IP TOS IP
PROT ip src as ip dst as ipv4 next hop addr ipv4 src mask ipv4 dst mask tcp flags intf output
bytes pkts time first time last

=====
=====

10.1.1.1 239.1.1.1 57314 1967 Gi0/0/3 0 0x00 17 0 0 0.0.0.0 /0 /0 0x00 Null 80 1 11:09:08.082
11:09:08.082

Packet Capture sul ricevitore.

```
> Frame 1: 94 bytes on wire (752 bits), 94 bytes captured (752 bits)
> Ethernet II, Src: Cisco_93:70:01 (00:a6:ca:93:70:01), Dst: IPv4mcast_01:01:01 (01:00:5e:01:01:01)
v Internet Protocol Version 4, Src: 10.1.1.1, Dst: 239.1.1.1
  0100 .... = Version: 4
  .... 0101 = Header Length: 20 bytes (5)
  > Differentiated Services Field: 0x00 (DSCP: CS0, ECN: Not-ECT)
  Total Length: 80
  Identification: 0x0000 (0)
  > Flags: 0x00
  Fragment Offset: 0
  Time to Live: 254
  Protocol: UDP (17)
  Header Checksum: 0xc198 [validation disabled]
  [Header checksum status: Unverified]
  Source Address: 10.1.1.1
  Destination Address: 239.1.1.1
  > User Datagram Protocol, Src Port: 54527, Dst Port: 1967
  > Data (52 bytes)
```

(S, G) nel ricevitore VRF.

```
ASR1002-1#show ip mroute vrf receiver 239.1.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected, L - Local,
P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT,
M - MSDP created entry, E - Extranet, X - Proxy Join Timer Running,
A - Candidate for MSDP Advertisement, U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender, Y - Joined MDT-data group,
y - Sending to MDT-data group, G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route, V - RD & Vector, v - Vector,
p - PIM Joins on route, x - VxLAN group, c - PFP-SA cache created entry,
* - determined by Assert Outgoing interface flags: H - Hardware switched,
A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state:
Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 00:21:36/stopped, RP 1.1.1.1, flags: SJC
Incoming interface: Loopback1, RPF nbr 1.1.1.1, using vrf source
Outgoing interface list:
GigabitEthernet0/0/1, Forward/Sparse, 00:21:36/00:02:43
```

(10.1.1.1, 239.1.1.1), 00:03:55/stopped, flags: T
Incoming interface: GigabitEthernet0/0/0, RPF nbr 0.0.0.0, using vrf source
Outgoing interface list:
GigabitEthernet0/0/1, Forward/Sparse, 00:03:55/00:02:43

ASR1002-1#show ip mroute vrf source 239.1.1.1

IP Multicast Routing Table

Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry,
* - determined by Assert

Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join

Timers: Uptime/Expires

Interface state: Interface, Next-Hop or VCD, State/Mode

(* , 239.1.1.1), 00:22:36/stopped, RP 1.1.1.1, flags: SJCFE

Incoming interface: Null, RPF nbr 0.0.0.0

Outgoing interface list: Null

Extranet receivers in vrf receiver:

(* , 239.1.1.1), 00:21:54/stopped, RP 1.1.1.1, OIF count: 1, flags: SJC

(10.1.1.1, 239.1.1.1), 00:04:14/00:02:55, flags: FTE

Incoming interface: GigabitEthernet0/0/0, RPF nbr 0.0.0.0

Outgoing interface list: Null

Extranet receivers in vrf receiver:

(10.1.1.1, 239.1.1.1), 00:04:14/stopped, OIF count: 1, flags: T

ASR1002-1#show ip mfib vrf receiver 239.1.1.1

Entry Flags: C - Directly Connected, S - Signal, IA - Inherit A flag,

ET - Data Rate Exceeds Threshold, K - Keepalive

DDE - Data Driven Event, HW - Hardware Installed

ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB

MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary

MS - MoFRR Entry in Sync, MC - MoFRR entry in MoFRR Client.

I/O Item Flags: IC - Internal Copy, NP - Not platform switched,

NS - Negate Signalling, SP - Signal Present,

A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,

MA - MFIB Accept, A2 - Accept backup,

RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

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

Other counts: Total/RPF failed/Other drops

I/O Item Counts: HW Pkt Count/FS Pkt Count/PS Pkt Count Egress Rate in pps

VRF receiver

(* , 239.1.1.1) Flags: C HW

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

HW Forwarding: 0/0/0/0, Other: 1/1/0

GigabitEthernet0/0/1 Flags: NS

(10.1.1.1, 239.1.1.1) Flags: HW

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

```
HW Forwarding: 15/0/94/0, Other: 0/0/0
GigabitEthernet0/0/1 Flags: NS
```

Inoltre, non vi è alcuna voce unicast nel ricevitore VRF RIB per RP e IP sorgente.

```
ASR1002-1#show ip route vrf receiver 1.1.1.1
Routing Table: receiver % Network not in table
```

```
ASR1002-1#show ip route vrf receiver 10.1.1.1
Routing Table: receiver % Network not in table
```

Pertanto RPF indica come non riuscito, il che è previsto, ma il traffico multicast verrà inoltrato.

```
ASR1002-1#show ip rpf vrf receiver 1.1.1.1
failed, no route exists
ASR1002-1#show ip rpf vrf receiver 10.1.1.1
failed, no route exists
```

Esiste un altro modo per raggiungere questo obiettivo, ovvero:

```
ip mroute vrf receiver 10.1.1.0 255.255.255.0 fallback-lookup vrf source
```

```
ip mroute vrf receiver 1.1.1.1 255.255.255.255 fallback-lookup vrf source
```

In questo caso, il controllo RPF avrà esito positivo.

```
ASR1002-1#show ip rpf vrf receiver 1.1.1.1
RPF information for ? (1.1.1.1)
RPF interface: Loopback1
RPF neighbor: ? (1.1.1.1) - directly connected
RPF route/mask: 1.1.1.1/32
RPF type: multicast (connected)
Doing distance-preferred lookups across tables
Using Extranet RPF Rule: Static Fallback Lookup, RPF VRF: source
RPF topology: ipv4 multicast base
```

```
ASR1002-1#show ip rpf vrf receiver 10.1.1.1
RPF information for ? (10.1.1.1)
RPF interface: GigabitEthernet0/0/0
RPF neighbor: ? (10.1.1.1) - directly connected
RPF route/mask: 10.1.1.0/24
RPF type: multicast (connected)
Doing distance-preferred lookups across tables
Using Extranet RPF Rule: Static Fallback Lookup, RPF VRF: source
RPF topology: ipv4 multicast base
```

Tuttavia, è possibile utilizzare il comando **RPF select** o il comando **fallback lookup**, entrambi in ultima analisi soddisfano il nostro scopo.

Risoluzione dei problemi

Al momento non sono disponibili informazioni specifiche per la risoluzione dei problemi di questa configurazione.