

Implementa ACI Transit Routing (Multipod)

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Premesse](#)

[Configurazione](#)

[Esempio di rete](#)

[Configurazioni](#)

[Risorse](#)

Introduzione

In questo documento viene descritto come configurare il routing di transito in un ambiente multipod ACI (Application Centric Infrastructure).

Contributo di Matthew Rich, Cisco TAC Engineer

Prerequisiti

Requisiti

Cisco raccomanda la conoscenza dei seguenti argomenti:

1. ACI multi-pod environment
2. Almeno 2 output L3
3. Almeno 1 direttrice in ciascun pod configurato come riflettore di route MP-BGP:
[configurazione di un riflettore di route MP-BGP](#)

Componenti usati

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

1. 2 switch N5K-C5548UP (usati come router esterni)
2. 1 switch foglia N9K-C9332PQ e 1 switch foglia N9K-C93108TC-EX
3. 2 Switch a spine N9K-C9336PQ
4. 1 N9K-C9232C (usato come dispositivo IPN)

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

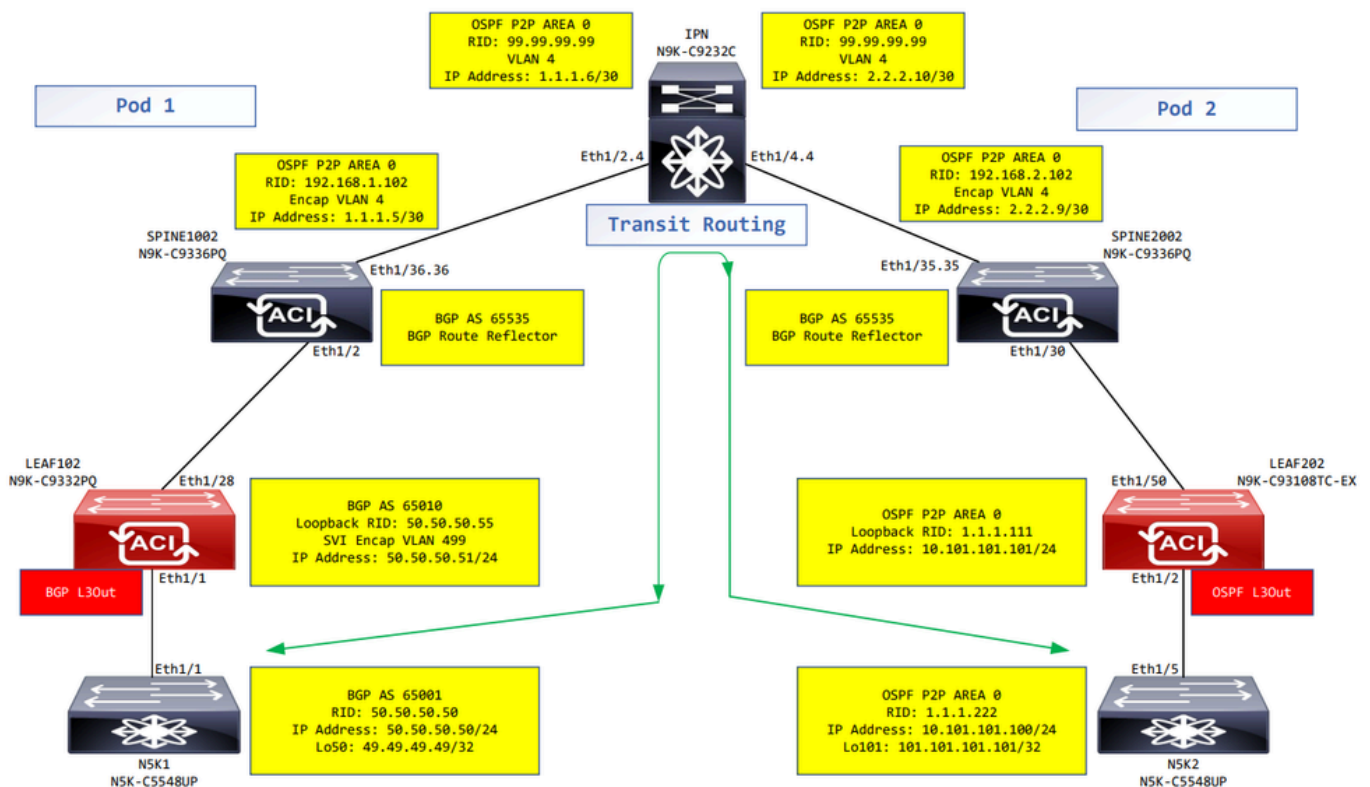
ripristinata ai valori predefiniti. Se la rete è operativa, valutare attentamente eventuali conseguenze derivanti dall'uso dei comandi.

Premesse

Nel routing di transito, il fabric Cisco ACI annuncia i percorsi che vengono appresi da una connessione di uscita di livello 3 (L3Out) a un'altra connessione L3Out. I domini esterni di layer 3 eseguono il peer con il fabric sugli switch foglia del bordo. La struttura è un dominio MP-BGP (Multiprotocol Border Gateway Protocol) di transito tra i peer.

Configurazione

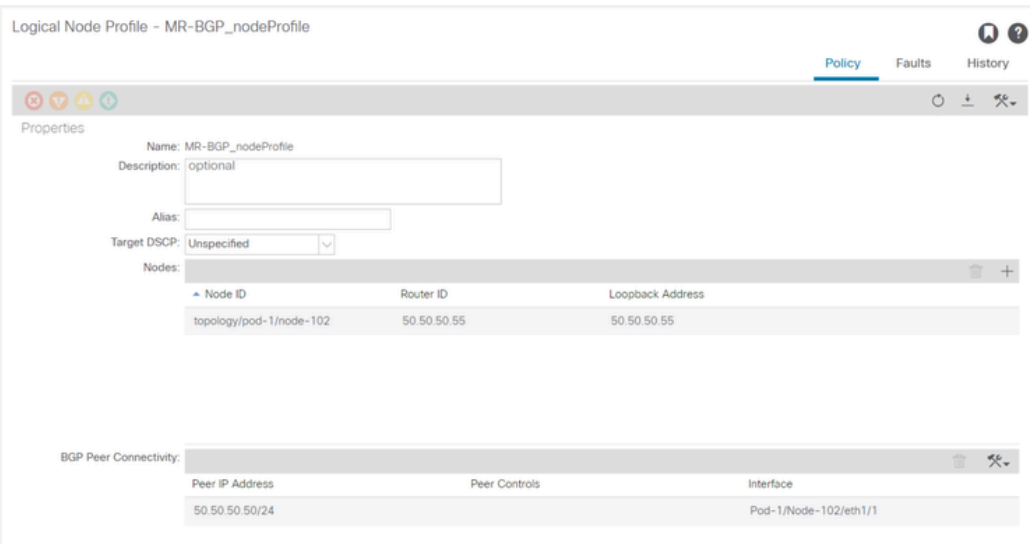
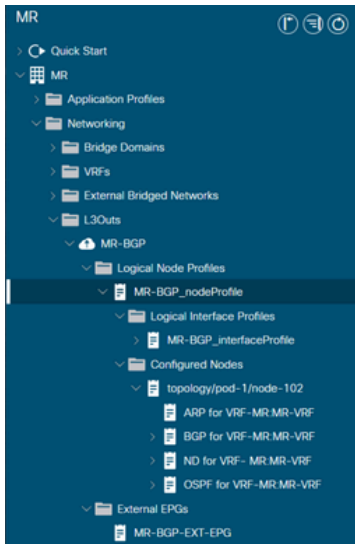
Esempio di rete



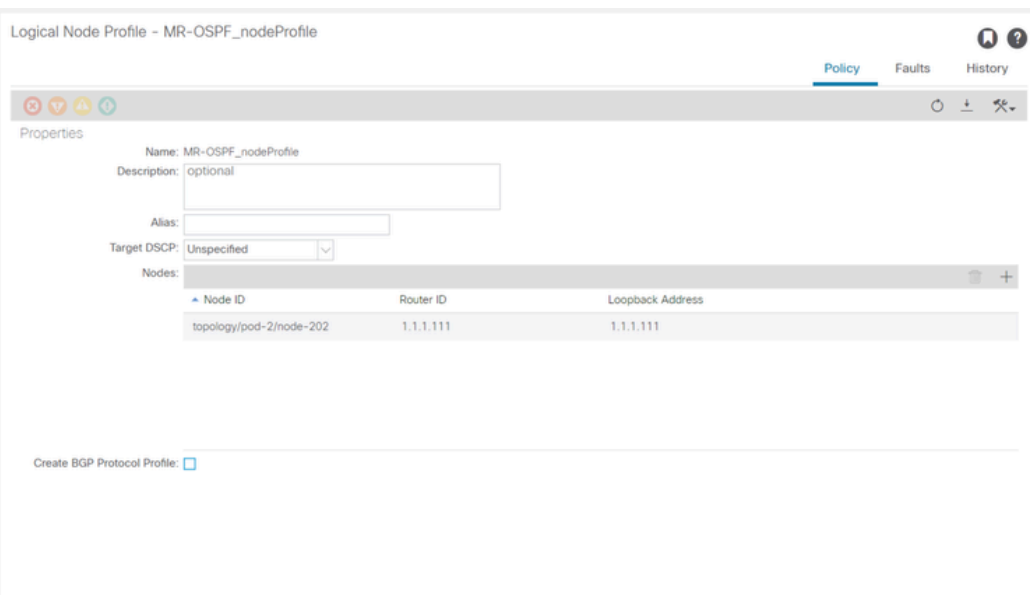
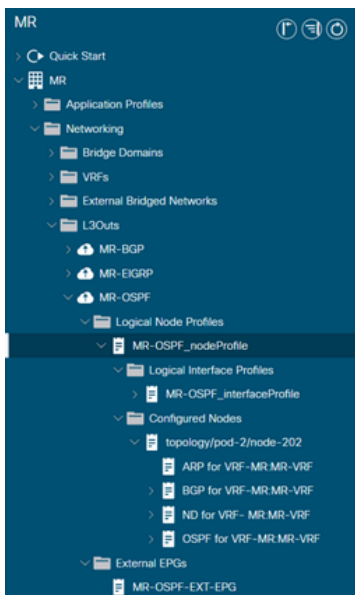
Esempio di rete

Configurazioni

Un profilo di nodo logico viene utilizzato per identificare lo switch foglia connesso alle reti esterne e in grado di distribuire il protocollo di routing o le route statiche verso di esso.



Profilo nodo logico per LEAF102



Profilo nodo logico per LEAF202

Per identificare l'interfaccia L3Out che si connette al dispositivo esterno, viene utilizzato un profilo di interfaccia logica. Come conseguenza di entrambi i profili, vengono visualizzati diversi elementi funzione definiti per il routing e l'inoltro virtuale (VRF): Address Resolution Protocol (ARP), Border Gateway Protocol (BGP), Neighbor Discovery e Open Shortest Path First (OSPF).

Logical Interface Profile - MR-BGP_interfaceProfile

Policy | Faults | History

General | Routed Sub-Interfaces | Routed Interfaces | **SVI** | Floating SVI

Path	Side A IP	Side B IP	Secondary IP Address	IP Address	MAC Address	MTU (bytes)	Encap	Encap Scope
Pod-1/Node-102/eth1/1				50.50.50.51/24	00:22:BD:F8:19:FF	inherit	vlan-499	Local

Profilo interfaccia logica per LEAF102, eth1/1

Logical Interface Profile - MR-OSPF_interfaceProfile

Policy | Faults | History

General | Routed Sub-Interfaces | **Routed Interfaces** | SVI | Floating SVI

Path	IP Address	Secondary IP Address	MAC Address	MTU (bytes)	PTP
Pod-2/Node-202/eth1/2	10.101.101.101/24		00:22:BD:F8:19:FF	9000	Disabled

Profilo interfaccia logica per LEAF202, eth1/2

Un profilo di istanza EPG esterno (External EPG, L3Out EPG) rappresenta un gruppo di subnet esterne con lo stesso comportamento di protezione. Altre subnet possono inoltre essere associate ad altri ambiti, che definiscono il comportamento di routing per la subnet.

The screenshot displays the configuration page for 'External EPG Instance Profile - MR-BGP-EXT-EPG'. The left sidebar shows a tree view with 'MR-OSPF' selected. The main panel has tabs for 'Policy', 'Operational', 'Stats', 'Health', 'Faults', and 'History'. Under the 'Policy' tab, there are sub-tabs for 'General', 'Contracts', and 'Inherited Contracts'. The 'General' sub-tab is active, showing the following properties:

- Name: MR-BGP-EXT-EPG
- Alias: [Empty field]
- Tags: [Dropdown menu]
- Global Alias: [Empty field]
- Description: optional
- pcTag: 49159
- Contract Exception Tag: [Empty field]
- Configured VRF Name: MR-VRF
- Resolved VRF: uni/tn-MR/ctx-MR-VRF
- QoS Class: Unspecified
- Target DSCP: Unspecified
- Configuration Status: applied
- Configuration Issues: [Empty list]
- Preferred Group Member: Exclude Include
- Subnets:

IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
49.49.49.49/32					External Subnets for th...

Profilo istanza EPG esterna per MR-BGP L3Out

The screenshot displays the configuration page for 'External EPG Instance Profile - MR-OSPF-EXT-EPG'. The left sidebar shows a tree view with 'MR-OSPF' selected. The main panel has tabs for 'Policy', 'Operational', 'Stats', 'Health', 'Faults', and 'History'. Under the 'Policy' tab, there are sub-tabs for 'General', 'Contracts', and 'Inherited Contracts'. The 'General' sub-tab is active, showing the following properties:

- Name: MR-OSPF-EXT-EPG
- Alias: [Empty field]
- Tags: [Dropdown menu]
- Global Alias: [Empty field]
- Description: optional
- pcTag: 49156
- Contract Exception Tag: [Empty field]
- Configured VRF Name: MR-VRF
- Resolved VRF: uni/tn-MR/ctx-MR-VRF
- QoS Class: Unspecified
- Target DSCP: Unspecified
- Configuration Status: applied
- Configuration Issues: [Empty list]
- Preferred Group Member: Exclude Include
- Subnets:

IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy
101.101.101.101/32					External Subnets for th...

Profilo istanza EPG esterna per MR-OSPF L3Out

Il contratto MR-PERMIT-ICMP viene applicato sia come contratto fornito che come contratto consumato in entrambi gli EPG esterni.

External EPG Instance Profile - MR-BGP-EXT-EPG

Policy Operational Stats Health Faults History

General Contracts Inherited Contracts

Healthy

Name	Tenant	Tenant Alias	Contract Type	Provided / Consumed	QoS Class	State	Label	Subject Label
Contract Type: Contract								
MR-PERMIT-ICMP	MR		Contract	Provided	Unspecified	formed		
MR-PERMIT-ICMP	MR		Contract	Consumed	Unspecified	formed		

Contratto MR-PERMIT-ICMP applicato a MR-BGP-EXT-EPG

External EPG Instance Profile - MR-OSPF-EXT-EPG

Policy Operational Stats Health Faults History

General Contracts Inherited Contracts

Healthy

Name	Tenant	Tenant Alias	Contract Type	Provided / Consumed	QoS Class	State	Label	Subject Label
Contract Type: Contract								
MR-PERMIT-ICMP	MR		Contract	Provided	Unspecified	formed		
MR-PERMIT-ICMP	MR		Contract	Consumed	Unspecified	formed		

Contratto MR-PERMIT-ICMP applicato a MR-OSPF-EXT-EPG

Su LEAF102, BGP viene stabilito con il router adiacente 50.50.50.50 e riceve la rete esterna 49.49.49.49/32.

BGP Peer Entry - 50.50.50.50

General Address Health Faults History

Properties

- Vrf Name: MR-MR-VRF
- BGP Version: BGP Version 4
- Remote Router Id: 50.50.50.50
- BGP State: Established
- Up For: 2022-07-27T17:17:22.493+00:00
- Remote As: 65001
- Update Source: vlan14
- Restart Time Advertised By Peer: Default
- Hold Time: 180
- Keepalive Interval: 60
- Neighbor: 50.50.50.50
- Link: eBGP
- Peer Index: 1
- Shutdown Reason: Unspecified
- State Reason: none
- Directly Attached Interface: vlan14
- Tcp Md5 Authentication: disabled
- Connection Established: 1
- Connection Dropped: 0
- Connection Attempts: na

Message Statistics

	Sent	Rcvd
Opens	1	1
Notifications	0	0
Updates	8	2
Keepalives	1692	1689
Route Refresh	0	0
Capability	1	1
Total	1702	1693
Total bytes	32485	32186
Bytes in queue	0	0

Next Hop

Address:	Resolved Using:
Refcount	

Voce peer BGP su LEAF102

```
LEAF102# show ip bgp summary vrf MR:MR-VRF
BGP summary information for VRF MR:MR-VRF, address family IPv4 Unicast
BGP router identifier 50.50.50.55, local AS number 65535
BGP table version is 37, IPv4 Unicast config peers 4, capable peers 2
14 network entries and 16 paths using 1952 bytes of memory
BGP attribute entries [12/1776], BGP AS path entries [0/0]
BGP community entries [0/0], BGP clusterlist entries [5/28]

Neighbor      V   AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
50.50.50.50   4 65001   1691    1700     37   0    0    1d04h 1
```

Riepilogo BGP per VRF MR:MR-VRF su LEAF102

```
LEAF102# show ip route bgp vrf MR:MR-VRF
IP Route Table for VRF "MR:MR-VRF"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

49.49.49.49/32, ubest/mbest: 1/0
 *via 50.50.50.50%MR:MR-VRF, [20/0], 1d04h, bgp-65535, external, tag 65010
```

Route BGP per VRF MR:MR-VRF su LEAF102

Su LEAF202, OSPF viene stabilito con il router adiacente 1.1.1.222 e riceve la rete esterna 101.101.101.101/32.

OSPF - MR:MR-VRF

General Health Faults History

PROPERTIES

Name: MR:MR-VRF
Route ID: 1.1.1.111
Distance: 110
Max ECMP: 8
Bandwidth Reference (Mbps): 40000
Operational State: Up

STATS

Interface Count: 2
Activeareact: 1
Active Nssa Areaact: 0
Active Stub Areaact: 0
Active Ext Areaact: 1
Extareaact: 1
Nssa Areaact: 0
Stubareaact: 0
Areaact: 1
Ext Lsact: 1
Opaqas Lsact: 0

Neighbors

Neighbor Id	State	Peer Ip	Interface
1.1.1.222	Full	10.101.101.100	eth1/2

Page 1 of 1 Objects Per Page: 15 Displaying Objects 1 - 1 Of 1

Inter Protocol Route Leak Into OSPF

Name	Redistribution Protocol	Route Map	Scope	Asn
MR:MR-VRF	BGP	exp-ctx-PROTO-2555906	Inter protocol leak	65535
MR:MR-VRF	COOP	exp-ctx-st-2555906	Inter protocol leak	1
MR:MR-VRF	Direct	exp-ctx-st-2555906	Inter protocol leak	1
MR:MR-VRF	EIGRP	exp-ctx-PROTO-2555906	Inter protocol leak	1
MR:MR-VRF	Static	exp-ctx-st-2555906	Inter protocol leak	1

Voce router adiacente OSPF su LEAF202

```
LEAF202# show ip ospf neighbors vrf MR:MR-VRF
OSPF Process ID default VRF MR:MR-VRF
Total number of neighbors: 1
Neighbor ID      Pri State           Up Time   Address          Interface
1.1.1.222        1 FULL/ -         2d04h    10.101.101.100  Eth1/2
```

Vicino OSPF per VRF MR:MR-VRF su LEAF202

```
LEAF202# show ip route ospf vrf MR:MR-VRF
IP Route Table for VRF "MR:MR-VRF"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

101.101.101.101/32, ubest/mbest: 1/0
 *via 10.101.101.100, eth1/2, [110/41], 1d00h, ospf-default, intra
```

Route OSPF per VRF MR:MR-VRF su LEAF202

Sia su LEAF102 che su LEAF202, la tabella MP-BGP per il VRF mostra la rete BGP esterna, 49.49.49.49/32, ma appare come esterna su LEAF102 e interna su LEAF202. La rete esterna OSPF 101.101.101.101/32 appare anche nelle tabelle BGP su entrambi i lati, su LEAF202 come ridistribuita da OSPF e su LEAF102 come interna.

```
LEAF102# show bgp vpnv4 unicast vrf MR:MR-VRF
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
BGP table version is 119, local router ID is 10.0.232.68
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

Network          Next Hop          Metric   LocPrf   Weight Path
Route Distinguisher: 102:2555906 (VRF MR:MR-VRF)
*>e49.49.49.49/32  50.50.50.50      0        100     0 65010 65001 i
*>i101.101.101.101/32 20.0.248.0      41        100     0 ?
```


Tabella MP-BGP per VRF MR:MR-VRF su LEAF102

```
LEAF202# show bgp vpnv4 unicast vrf MR:MR-VRF
BGP routing table information for VRF overlay-1, address family VPNv4 Unicast
BGP table version is 95, local router ID is 20.0.248.0
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

  Network          Next Hop          Metric      LocPrf      Weight Path
Route Distinguisher: 202:2555906 (VRF MR:MR-VRF)
*>i49.49.49.49/32  10.0.232.68      100         100         0 65010 65001 i
*>r101.101.101.101/32 0.0.0.0          41          100         32768 ?
```

Tabella MP-BGP per VRF MR:MR-VRF su LEAF202

La tabella BGP IPv4 contiene informazioni equivalenti.

```
LEAF102# show bgp ipv4 unicast vrf MR:MR-VRF
BGP routing table information for VRF MR:MR-VRF, address family IPv4 Unicast
BGP table version is 37, local router ID is 50.50.50.55
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

  Network          Next Hop          Metric      LocPrf      Weight Path
*>e49.49.49.49/32  50.50.50.50      100         100         0 65010 65001 i
*>i101.101.101.101/32 20.0.248.0       41          100         0 ?
```

Tabella BGP IPv4 per VRF MR:MR-VRF su LEAF102

```
LEAF202# show bgp ipv4 unicast vrf MR:MR-VRF
BGP routing table information for VRF MR:MR-VRF, address family IPv4 Unicast
BGP table version is 31, local router ID is 1.1.1.111
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist, I-injected
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup

  Network          Next Hop          Metric      LocPrf      Weight Path
*>i49.49.49.49/32  10.0.232.68      100         100         0 65010 65001 i
*>r101.101.101.101/32 0.0.0.0          41          100         32768 ?
```

Tabella BGP IPv4 per VRF MR:MR-VRF su LEAF202

Tuttavia, la rete esterna OSPF 101.101.101.101/32 non è inclusa nella tabella di routing di N5K1.

```
N5K1# show ip route vrf MR-BGP
IP Route Table for VRF "MR-BGP"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

49.49.49.49/32, ubest/mbest: 2/0, attached
  *via 49.49.49.49, Lo50, [0/0], 1d07h, local
  *via 49.49.49.49, Lo50, [0/0], 1d07h, direct
50.50.50.0/24, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d07h, direct
50.50.50.50/32, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d07h, local
```

RIB per VRF MR-BGP su N5K1

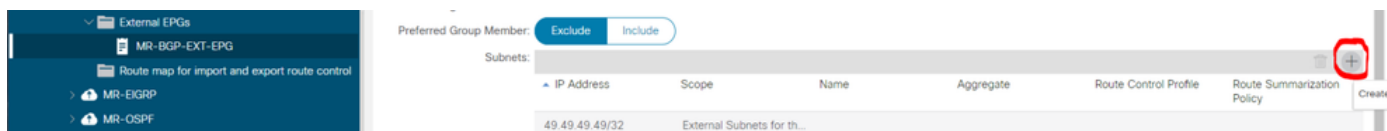
Analogamente, la rete esterna BGP 49.49.49.49/32 non si trova nella scheda RIB di N5K2.

```
N5K2# show ip route vrf MR-OSPF
IP Route Table for VRF "MR-OSPF"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

1.1.1.111/32, ubest/mbest: 1/0
  *via 10.101.101.101, Eth1/5, [110/41], 2d05h, ospf-1, intra
10.101.101.0/24, ubest/mbest: 1/0, attached
  *via 10.101.101.100, Eth1/5, [0/0], 6d22h, direct
10.101.101.100/32, ubest/mbest: 1/0, attached
  *via 10.101.101.100, Eth1/5, [0/0], 6d22h, local
101.101.101.101/32, ubest/mbest: 2/0, attached
  *via 101.101.101.101, Lo101, [0/0], 2d04h, local
  *via 101.101.101.101, Lo101, [0/0], 2d04h, direct
```

RIB per VRF MR-OSPF su N5K2

In BGP L3Out, immettere l'indirizzo IP della subnet esterna ricevuta da OSPF L3Out, 101.101.101.101/32, scegliere Esporta subnet di controllo route nella sezione Controllo route, quindi cancellare le subnet esterne per la classificazione EPG esterno. Fare clic su Invia. L'opzione Esporta subnet di controllo del percorso consente di esportare (annunciare) una rete nel peer esterno.



Crea nuova subnet

Create Subnet dialog box with fields for IP Address (101.101.101.101/32), Name, Route Control options (Export Route Control Subnet checked), Aggregate options, Route Summarization Policy (BGP Route Summarization Policy), Route Control Profile table, External EPG classification options, and Cancel/Submit buttons.

Configurare le opzioni corrette per la nuova subnet

Sul modello N5K1, la rete esterna OSPF 101.101.101.101/32 viene ora ricevuta tramite BGP.

```
N5K1# show ip route vrf MR-BGP
IP Route Table for VRF "MR-BGP"
'*' denotes best ucast next-hop
 '**' denotes best mcast next-hop
 '[x/y]' denotes [preference/metric]
 '%<string>' in via output denotes VRF <string>

49.49.49.32, ubest/mbest: 2/0, attached
  *via 49.49.49.49, Lo50, [0/0], 1d08h, local
  *via 49.49.49.49, Lo50, [0/0], 1d08h, direct
50.50.50.0/24, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d08h, direct
50.50.50.50/32, ubest/mbest: 1/0, attached
  *via 50.50.50.50, Vlan499, [0/0], 1d08h, local
101.101.101.101/32, ubest/mbest: 1/0
  *via 50.50.50.51, [20/0], 00:00:03, bgp-65001, external, tag 65010,
```

RIB per VRF MR-BGP su N5K1

In Uscita OSPF L3immettere l'indirizzo IP della subnet esterna ricevuta da BGP L3Out, 49.49.49.49/32, scegliere Esporta subnet di controllo route nella sezione Controllo route, quindi deselezionare Subnet esterne per la classificazione EPG esterno. Fare clic su Invia.

Preferred Group Member: Exclude Include

Subnets:

IP Address	Scope	Name	Aggregate	Route Control Profile	Route Summarization Policy	
101.101.101.101/32	External Subnets for th...					<input type="checkbox"/> Create

Crea nuova subnet

Create Subnet



IP Address:
address/mask

Name:

Route Control:

- Export Route Control Subnet
- Import Route Control Subnet
- Shared Route Control Subnet

Aggregate

- Aggregate Export
- Aggregate Import
- Aggregate Shared Routes

Route Summarization Policy

Route Control Profile:

Name	Direction
------	-----------

Route control is used for filtering external routes advertised out of the fabric, allowed into the fabric, or leaked to other VRFs within the fabric.

External EPG classification:

- External Subnets for External EPG
- Shared Security Import Subnet

External EPG classification is used to identify the external networks associated with this external EPG for policy enforcement (Contracts).

Cancel

Submit

Configurare le opzioni corrette per la nuova subnet

Ora su N5K2, la rete esterna BGP 49.49.49.49/32 viene ricevuta su OSPF.

```
N5K2# show ip route vrf MR-OSPF
IP Route Table for VRF "MR-OSPF"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

1.1.1.111/32, ubest/mbest: 1/0
  *via 10.101.101.101, Eth1/5, [110/41], 2d05h, ospf-1, intra
10.101.101.0/24, ubest/mbest: 1/0, attached
  *via 10.101.101.100, Eth1/5, [0/0], 6d22h, direct
10.101.101.100/32, ubest/mbest: 1/0, attached
  *via 10.101.101.100, Eth1/5, [0/0], 6d22h, local
49.49.49.49/32, ubest/mbest: 1/0
  *via 10.101.101.101, Eth1/5, [110/1], 00:01:59, ospf-1, type-2, tag 4294967295,
101.101.101.101/32, ubest/mbest: 2/0, attached
  *via 101.101.101.101, Lo101, [0/0], 2d05h, local
  *via 101.101.101.101, Lo101, [0/0], 2d05h, direct
```

RIB per VRF MR-OSPF su N5K2

Il ping funziona tra le due reti a causa del contratto MR-PERMIT-ICMP applicato in precedenza a entrambi gli EPG esterni.

```
N5K1# ping 101.101.101.101 vrf MR-BGP source 49.49.49.49
PING 101.101.101.101 (101.101.101.101) from 49.49.49.49: 56 data bytes
64 bytes from 101.101.101.101: icmp_seq=0 ttl=252 time=3.059 ms
64 bytes from 101.101.101.101: icmp_seq=1 ttl=252 time=2.963 ms
64 bytes from 101.101.101.101: icmp_seq=2 ttl=252 time=7.928 ms
64 bytes from 101.101.101.101: icmp_seq=3 ttl=252 time=2.954 ms
64 bytes from 101.101.101.101: icmp_seq=4 ttl=252 time=2.982 ms

--- 101.101.101.101 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 2.954/3.977/7.928 ms
```

Verifica della comunicazione su N5K1

```
N5K2# ping 49.49.49.49 vrf MR-OSPF source 101.101.101.101
PING 49.49.49.49 (49.49.49.49) from 101.101.101.101: 56 data bytes
64 bytes from 49.49.49.49: icmp_seq=0 ttl=252 time=3.107 ms
64 bytes from 49.49.49.49: icmp_seq=1 ttl=252 time=2.99 ms
64 bytes from 49.49.49.49: icmp_seq=2 ttl=252 time=2.98 ms
64 bytes from 49.49.49.49: icmp_seq=3 ttl=252 time=2.986 ms
64 bytes from 49.49.49.49: icmp_seq=4 ttl=252 time=2.99 ms

--- 49.49.49.49 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 2.98/3.01/3.107 ms
```

Verifica della comunicazione su N5K2

Risorse

[Guida alla configurazione delle reti di layer 3 di Cisco APIC, versione 6.0\(x\)](#)

[Nozioni fondamentali sull'infrastruttura incentrata sulle applicazioni Cisco, versione 4.2\(x\)](#)

[Guida alla configurazione delle reti di layer 3 di Cisco APIC, versione 3.x e precedenti](#)

Informazioni su questa traduzione

Cisco ha tradotto questo documento utilizzando una combinazione di tecnologie automatiche e umane per offrire ai nostri utenti in tutto il mondo contenuti di supporto nella propria lingua. Si noti che anche la migliore traduzione automatica non sarà mai accurata come quella fornita da un traduttore professionista. Cisco Systems, Inc. non si assume alcuna responsabilità per l'accuratezza di queste traduzioni e consiglia di consultare sempre il documento originale in inglese (disponibile al link fornito).