

Implemente o roteamento de trânsito da ACI (multipod)

Contents

[Introdução](#)

[Pré-requisitos](#)

[Requisitos](#)

[Componentes Utilizados](#)

[Informações de Apoio](#)

[Configurar](#)

[Diagrama de Rede](#)

[Configurações](#)

[Recursos](#)

Introdução

Este documento descreve como configurar o roteamento de trânsito em um ambiente multipod da Application Centric Infrastructure (ACI).

Contribuição de Matthew Rich, engenheiro do Cisco TAC

Pré-requisitos

Requisitos

A Cisco recomenda que você tenha conhecimento destes tópicos:

1. Ambiente de vários pods da ACI
2. Pelo menos 2 L3Outs
3. Pelo menos 1 coluna em cada pod configurado como um refletor de rota MP-BGP:
[Configurando um refletor de rota MP-BGP](#)

Componentes Utilizados

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

1. 2 switches N5K-C548UP (usados como roteadores externos)
2. 1 switch leaf N9K-C9332PQ e 1 switch leaf N9K-C93108TC-EX
3. 2 switches spine N9K-C9336PQ
4. 1 N9K-C9232C (usado como dispositivo NPI)

As informações neste documento foram criadas a partir de dispositivos em um ambiente de

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

Informações de Apoio

No roteamento em trânsito, a estrutura da Cisco ACI anuncia as rotas aprendidas de uma conexão L3Out (L3Out) de Camada 3 para outra conexão L3Out. Os domínios externos de Camada 3 correspondem à estrutura nos switches leaf de borda. A estrutura é um domínio de protocolo de gateway de borda multiprotocolo (MP-BGP) de trânsito entre os correspondentes.

Configurar

Diagrama de Rede

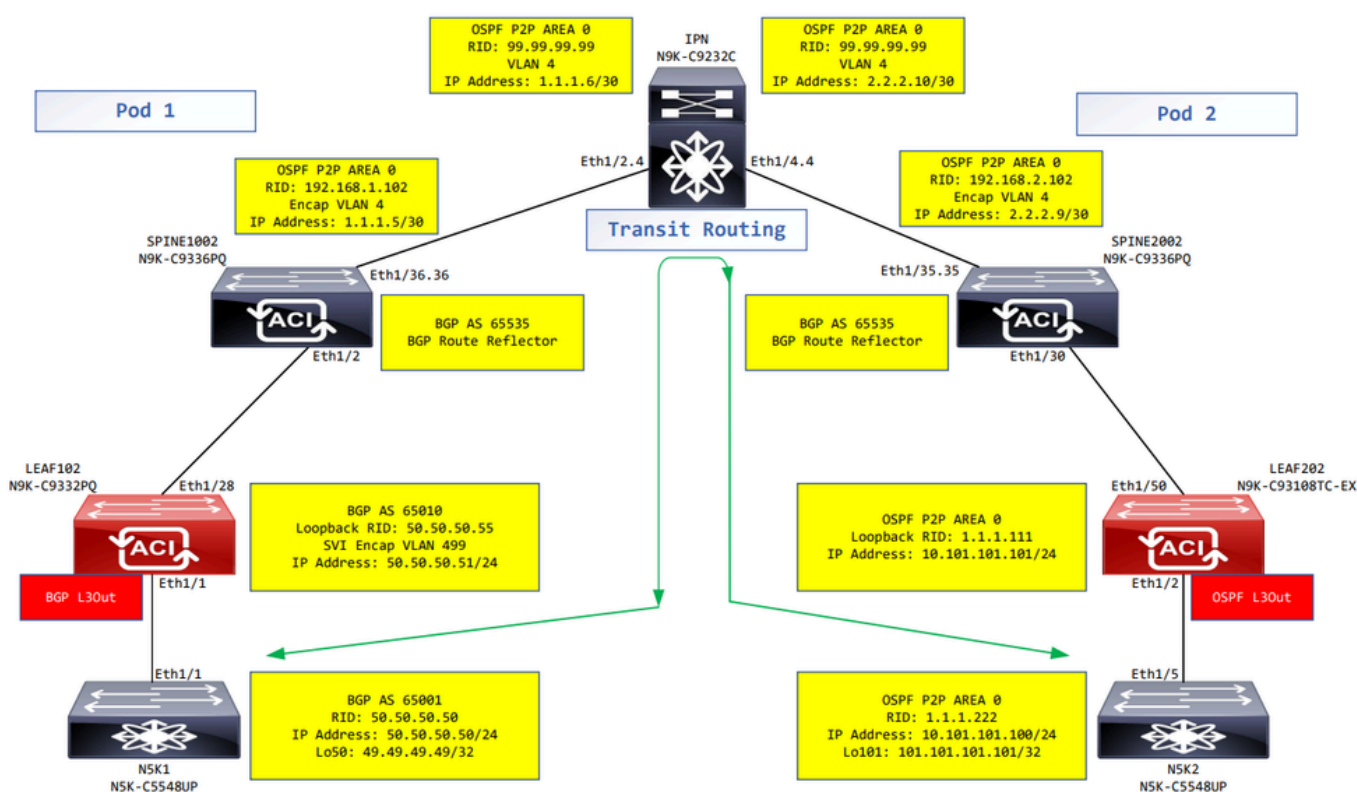
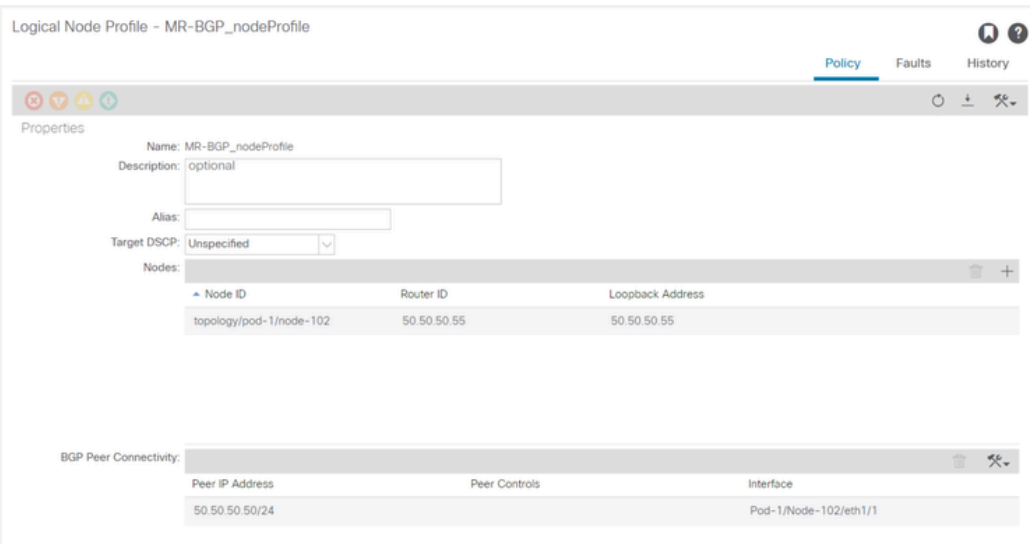
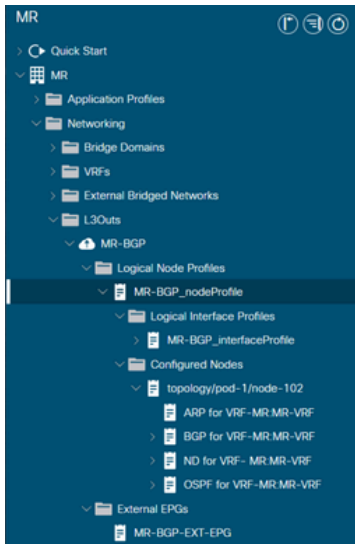


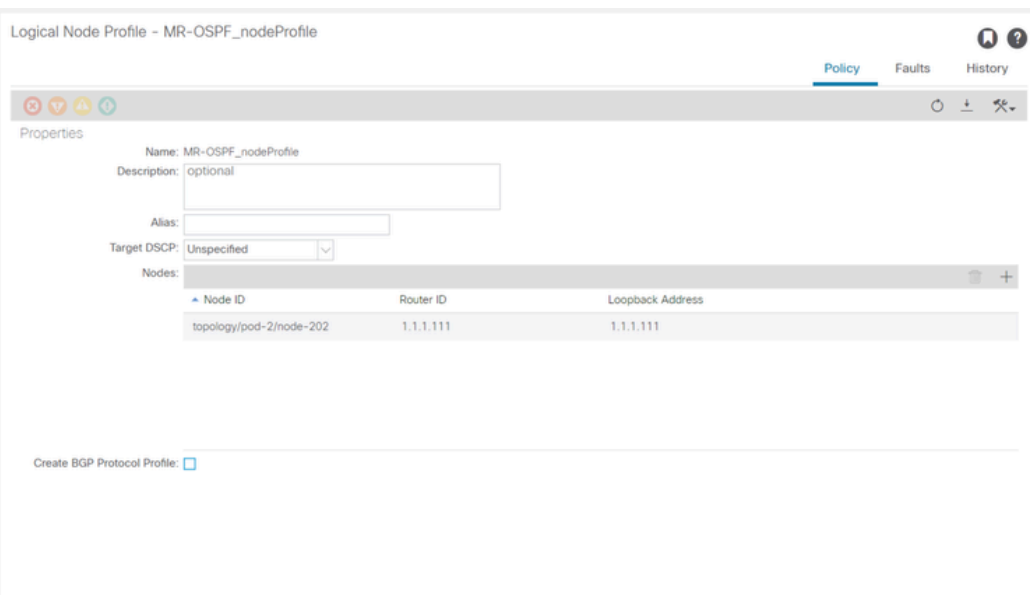
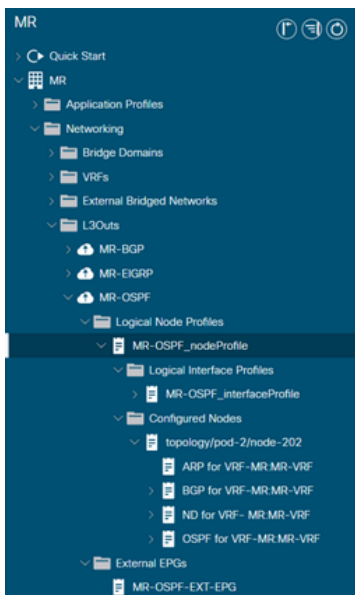
Diagrama de Rede

Configurações

Um perfil de nó lógico é usado para identificar o switch de folha que está conectado a redes externas e que pode implantar o protocolo de roteamento ou rotas estáticas para ele.



Perfil de nó lógico para LEAF102



Perfil de nó lógico para LEAF202

Um perfil de interface lógica é usado para identificar a interface L3Out que se conecta ao dispositivo externo. Você vê vários elementos de função que são definidos para roteamento e encaminhamento virtual (VRF): Address Resolution Protocol (ARP), Border Gateway Protocol (BGP), Neighbor Discovery e Open Shortest Path First (OSPF) como consequência de ambos os perfis.

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

Perfil de interface lógica para 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

Perfil de interface lógica para LEAF202, eth1/2

Um perfil de instância de EPG externo (EPG externo, EPG L3Out) representa um grupo de sub-redes externas que têm o mesmo comportamento de segurança. Outras sub-redes também podem ser associadas a outros escopos, que definem o comportamento de roteamento para essa sub-rede.

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

Policy | Operational | Stats | Health | Faults | History

General | Contracts | Inherited Contracts

Properties

Name: MR-BGP-EXT-EPG
 Alias:
 Tags: (enter tags separated by comma)
 Global Alias:
 Description: optional

pcTag: 49159
 Contract Exception Tag:
 Configured VRF Name: MR-VRF
 Resolved VRF: uni/tn-MR/ctx-MR-VRF
 QoS Class: Unspecified
 Target DSCP: Unspecified

Configuration Status: applied
 Configuration Issues:

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...

Perfil de instância de EPG externo para MR-BGP L3Out

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

Policy | Operational | Stats | Health | Faults | History

General | Contracts | Inherited Contracts

Properties

Name: MR-OSPF-EXT-EPG
 Alias:
 Tags: (enter tags separated by comma)
 Global Alias:
 Description: optional

pcTag: 49156
 Contract Exception Tag:
 Configured VRF Name: MR-VRF
 Resolved VRF: uni/tn-MR/ctx-MR-VRF
 QoS Class: Unspecified
 Target DSCP: Unspecified

Configuration Status: applied
 Configuration Issues:

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...

Perfil de instância de EPG externo para MR-OSPF L3Out

O contrato MR-PERMIT-ICMP é aplicado como um contrato fornecido e consumido em ambos os EPGs externos.

The screenshot shows the configuration page for 'External EPG Instance Profile - MR-BGP-EXT-EPG'. The left sidebar shows the navigation tree with 'MR-BGP' expanded. The main content area has tabs for 'Policy', 'Operational', 'Stats', 'Health', 'Faults', and 'History'. The 'Contracts' tab is active, showing a table of applied contracts.

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		

Contrato MR-PERMIT-ICMP aplicado ao MR-BGP-EXT-EPG

The screenshot shows the configuration page for 'External EPG Instance Profile - MR-OSPF-EXT-EPG'. The left sidebar shows the navigation tree with 'MR-OSPF' expanded. The main content area has tabs for 'Policy', 'Operational', 'Stats', 'Health', 'Faults', and 'History'. The 'Contracts' tab is active, showing a table of applied contracts.

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		

Contrato MR-PERMIT-ICMP aplicado ao MR-OSPF-EXT-EPG

Em LEAF102, o BGP é estabelecido com o vizinho 50.50.50.50 e está recebendo a rede externa 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	

Entrada de par BGP em 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
```

Resumo do BGP para VRF MR:MR-VRF no 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
```

Rota BGP para VRF MR:MR-VRF em LEAF102

No LEAF202, o OSPF é estabelecido com o vizinho 1.1.1.222 e recebe a rede externa 101.101.101.101/32.

Entrada de vizinho OSPF em 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
```

Vizinho OSPF para VRF MR:MR-VRF no 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
```

Rota OSPF para VRF MR:MR-VRF em LEAF202

Tanto no LEAF102 como no LEAF202, a tabela MP-BGP para o VRF mostra a rede BGP externa, 49.49.49.49/32, mas aparece como externa no LEAF102 e interna no LEAF202. A rede externa OSPF 101.101.101.101/32 também aparece nas tabelas BGP em ambas as folhas, em LEAF202 como redistribuída do OSPF e em LEAF102 como 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        0        65010 65001 i
*>i101.101.101.101/32 20.0.248.0       41       100      0 ?
```


Tabela MP-BGP para VRF MR:MR-VRF no 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 ?
```

Tabela MP-BGP para VRF MR:MR-VRF no LEAF202

A tabela BGP IPv4 contém informações equivalentes.

```
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 ?
```

Tabela BGP IPv4 para VRF MR:MR-VRF no 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 ?
```

Tabela BGP IPv4 para VRF MR:MR-VRF no LEAF202

No entanto, a rede externa OSPF 101.101.101.101/32 não está na tabela de roteamento de 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 para VRF MR-BGP em N5K1


Da mesma forma, a rede externa BGP 49.49.49.49/32 não está no RIB de 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 para VRF MR-OSPF no N5K2

No BGP L3Out, insira o endereço IP da sub-rede externa recebida do OSPF L3Out, 101.101.101.101/32, escolha Export Route Control Subnet na seção Route Control e limpe a classificação External Subnets for the External EPG. Clique em Submit. A opção Export Route Control Subnet permite que uma rede seja exportada (anunciada) para o peer externo.



Criar nova sub-rede

Create Subnet form with fields for IP Address (101.101.101.101/32), Name, Route Control options (Export Route Control Subnet checked), Aggregate options, Route Summarization Policy, and External EPG classification. The Submit button is circled in red.

Configurar as opções corretas para a nova sub-rede

No N5K1, a rede externa do OSPF 101.101.101.101/32 é agora recebida pelo 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 para VRF MR-BGP em N5K1

No OSPF L3Out, insira o endereço IP da sub-rede externa recebida do BGP L3Out, 49.49.49.49/32, escolha Export Route Control Subnet na seção Route Control e limpe External Subnets para a classificação External EPG. Clique em Submit.

The screenshot shows the configuration page for 'MR-OSPF-EXT-EPG'. On the left is a navigation tree with 'External EPGs' expanded. The main area shows 'Preferred Group Member' with 'Exclude' selected. Below is a 'Subnets' table with columns: IP Address, Scope, Name, Aggregate, Route Control Profile, and Route Summarization Policy. A 'Create' button is circled in red at the top right of the table.

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

Criar nova sub-rede

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

Configurar as opções corretas para a nova sub-rede

Agora no N5K2, a rede externa BGP 49.49.49.49/32 é recebida através do 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 para VRF MR-OSPF no N5K2

O ping funciona entre as duas redes por causa do contrato MR-PERMIT-ICMP que foi aplicado a ambos os EPGs externos anteriormente.

```
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
```

Verificação de comunicação no 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
```

Verificação de comunicação no N5K2

Recursos

[Guia de configuração de rede da camada 3 do Cisco APIC, versão 6.0\(x\)](#)

[Fundamentos da Cisco Application Centric Infrastructure, versão 4.2\(x\)](#)

[Guia de configuração de rede da camada 3 do Cisco APIC, versão 3.x e anterior](#)

Sobre esta tradução

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

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

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