

Distribuire l'EVPN di layer 3 su SR MPLS [Ospf / iBGP] [PE-CE is OSPF] in Nexus 9300

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Premesse](#)

[Riepilogo MPLS L3VPN](#)

[Panoramica di EVPN con L3VPN \(MPLS SR\)](#)

[Esempio di rete](#)

[Configurazione](#)

[Verifica](#)

[Informazioni correlate](#)

Introduzione

In questo documento viene descritto come implementare/configurare MPLS (Multiprotocol Label Switching) di EVPN di layer 3 sui prodotti Nexus 9300 con protocollo PE-CE come Open Shortest Path First (OSPF).

Prerequisiti

Requisiti

Cisco raccomanda la conoscenza dei seguenti argomenti:

- Border Gateway Protocol (BGP)
- OSPF (Open Shortest Path First)
- L3VPN
- EVPN
- Routing segmento (SR)

Componenti usati

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

- SPINE Hardware - 93360YC-FX2 in esecuzione con 9.3(3)
- Hardware FOGLIA - 93240YC-FX2 in esecuzione con 9.3(3)
- CLIENT - 93216TC-FX2 (host-1), Catalyst-3750 (host-2)

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

Riepilogo MPLS L3VPN

Una VPN è:

- Rete IP che fornisce servizi di rete privati su un'infrastruttura pubblica.
- Insieme di siti a cui è consentito comunicare tra loro in privato tramite Internet o altre reti pubbliche o private.

Le VPN convenzionali vengono create dalla configurazione di una rete completa di tunnel o di circuiti virtuali permanenti (PVC) per tutti i siti di una VPN. Questo tipo di VPN non è facile da mantenere o espandere, in quanto l'aggiunta di un nuovo sito richiede la modifica di ogni dispositivo periferico nella VPN.

Le VPN basate su MPLS vengono create nel layer 3 e si basano sul modello peer. Il modello peer consente al provider di servizi e al cliente di scambiare informazioni di routing di layer 3. Il provider di servizi inoltra i dati tra le sedi del cliente senza il coinvolgimento del cliente.

Le VPN MPLS sono più facili da gestire ed espandere rispetto alle VPN convenzionali. Quando si aggiunge un nuovo sito a una VPN MPLS, è necessario aggiornare solo il router perimetrale del provider di servizi che fornisce servizi alla sede del cliente.

Questi sono i componenti della VPN MPLS:

- Provider (P) router- Router nel nucleo della rete del provider. I router PE eseguono la commutazione MPLS e non collegano etichette VPN ai pacchetti indirizzati. Le etichette VPN vengono utilizzate per indirizzare i pacchetti di dati alla rete privata corretta o al router perimetrale del cliente.
- PE router: router che collega l'etichetta VPN ai pacchetti in arrivo in base all'interfaccia o alla sottointerfaccia su cui vengono ricevuti e allega anche le etichette di base MPLS. Un router PE si collega direttamente a un router CE.
- Cliente (C) router-router nella rete ISP (Internet Service Provider) o aziendale.
- Router Customer Edge (CE): router perimetrale sulla rete dell'ISP che si connette al router PE sulla rete. Un router CE deve interfacciarsi con un router PE.

Panoramica di EVPN con L3VPN (MPLS SR)

Le implementazioni dei data center (DC) hanno adottato VXLAN EVPN (o MPLS EVPN) per i suoi vantaggi, quali l'apprendimento del control-plane EVPN, la multitenancy, la mobilità perfetta, la ridondanza e le aggiunte più semplici di POD. Analogamente, il CORE è una rete MPLS L3VPN basata sul protocollo LDP (Label Distribution Protocol) o in transizione dal tradizionale sublay

basato su LDP di MPLS L3VPN a una soluzione più sofisticata come il Segment Routing (SR).

Il ciclo di segmenti viene adottato per i relativi vantaggi, quali:

- Piani di controllo IGP e MPLS unificati
- Metodi di progettazione del traffico più semplici
- Configurazione più semplice
- Adozione SDN
- EVPN (RFC 7432) è una soluzione basata su BGP MPLS che è stata utilizzata per i servizi Ethernet di nuova generazione in una rete di centri dati virtualizzata.
- EVPN utilizza diversi elementi di base, quali RD, RT e VRF, dalle tecnologie MPLS esistenti.
- L3 EVPN over SR, introdotto in NXOS 7.0(3)I6(1), utilizza la route EVPN Type-5 con incapsulamento MPLS.
- L3 EVPN over SR offre multi-tenant, scalabilità e prestazioni elevate per i servizi evoluti del centro dati.

Nota: In DC, il piano dati può essere VXLAN o MPLS.

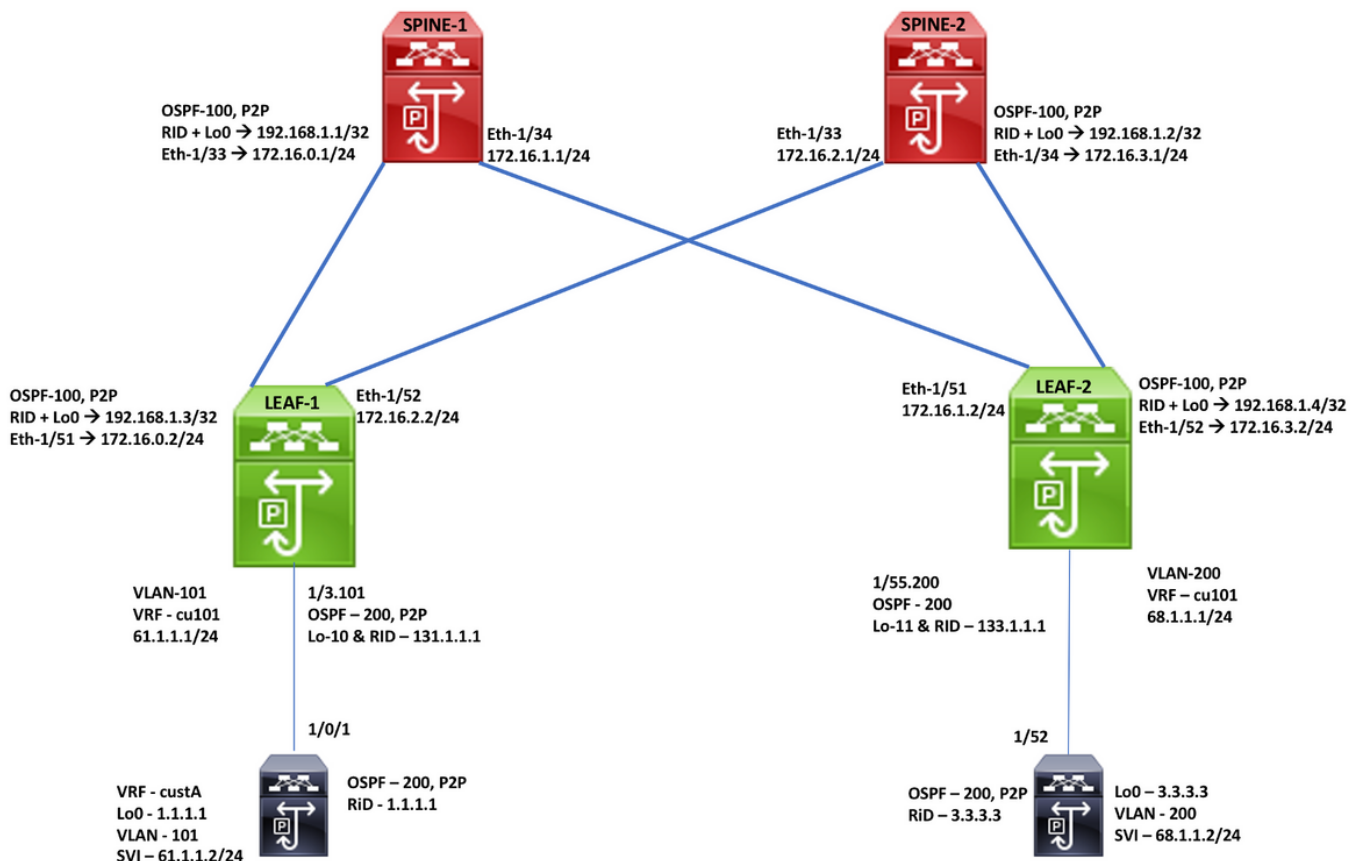
VPN MPLS L3 tradizionale

Blocchi predefiniti principali: RD, RT e VRF
Livello inferiore per trasporto: IGP, LDP e RSVP-TE
Livello overlay per servizio: VPNv4 e VPNv6

MPLS L3 VPN over SR

Blocchi predefiniti principali: RD, RT e VRF
Livello inferiore per trasporto: IGP/BGP-LU e SR-TE
Livello overlay per servizio: EVPN

Esempio di rete



Configurazione

SPINE-1 Configuration		
Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
feature-set mpls feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam mpls label range 5000 450000 segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.1/32 index 211 route-map label-index-spine1 permit 10 set label-index 211	interface Ethernet1/33 ip address 172.16.0.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown interface Ethernet1/34 ip address 172.16.1.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown interface loopback0 ip address 192.168.1.1/32 ip router ospf 100 area 0.0.0.0 router ospf 100 segment-routing mpls router-id 192.168.1.1	router bgp 65001 router-id 192.168.1.1 address-family ipv4 unicast network 192.168.1.1/32 route-map label-index-spine1 allocate-label all address-family ipv4 labeled-unicast address-family l2vpn evpn template peer EVPN remote-as 65001 update-source loopback0 address-family l2vpn evpn send-community extended route-reflector-client encapsulation mpls template peer Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended route-reflector-client next-hop-self soft-reconfiguration inbound always neighbor 172.16.0.2 inherit peer Labeled-unicast neighbor 172.16.1.2 inherit peer Labeled-unicast neighbor 192.168.1.3 inherit peer EVPN neighbor 192.168.1.4 inherit peer EVPN
SPINE-2 Configuration		
Enabling Features, Label-Range, Route-map, Label-Index	OSPF Configuration	BGP/EVPN Configuration
feature-set mpls feature ospf feature bgp feature mpls segment-routing feature mpls evpn feature interface-vlan feature mpls oam mpls label range 5000 450000 segment-routing mpls global-block 16000 25000 connected-prefix-sid-map address-family ipv4 192.168.1.2/32 index 221 route-map label-index-spine2 permit 10 set label-index 221	interface Ethernet1/33 ip address 172.16.2.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown interface Ethernet1/34 ip address 172.16.3.1/24 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0 mpls ip forwarding no shutdown interface loopback0 ip address 192.168.1.2/32 ip router ospf 100 area 0.0.0.0 router ospf 100 segment-routing mpls router-id 192.168.1.2	router bgp 65001 router-id 192.168.1.2 address-family ipv4 unicast network 192.168.1.2/32 route-map label-index-spine2 allocate-label all address-family ipv4 labeled-unicast address-family l2vpn evpn template peer EVPN remote-as 65001 update-source loopback0 address-family l2vpn evpn send-community extended route-reflector-client encapsulation mpls template peer Labeled-unicast remote-as 65001 address-family ipv4 labeled-unicast send-community extended route-reflector-client next-hop-self soft-reconfiguration inbound always neighbor 172.16.2.2 inherit peer Labeled-unicast neighbor 172.16.3.2 inherit peer Labeled-unicast neighbor 192.168.1.3 inherit peer EVPN neighbor 192.168.1.4 inherit peer EVPN

LEAF-1 Configuration

Enabling Features, Label-Range, Route-map, Label-Index

```

install feature-set mpls
feature-set mpls
nv overlay evpn
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam
feature nv overlay

mpls label range 5000 450000
segment-routing
mpls
  global-block 16000 25000
  connected-prefix-sid-map
  address-family ipv4
    192.168.1.3/32 index 311

ip prefix-list test1 seq 5 permit 61.1.1.0/24
ip prefix-list test1 seq 10 permit 131.1.1.1/32

ip prefix-list test3 seq 5 permit 1.1.1.1/32

route-map bgp65001 permit 10
  match route-type internal
route-map direct1 permit 10
  match ip address prefix-list test1
  set community 65001:10
route-map label-index-leaf-1 permit 10
  set label-index 311
route-map ospf200 permit 10
  match ip address prefix-list test3

vrf context cu101
  rd auto
  address-family ipv4 unicast
  route-target import 1:101
  route-target import 1:101 evpn
  route-target export 1:101
  route-target export 1:101 evpn
  
```

OSPF Configuration

```

interface Ethernet1/51
  ip address 172.16.0.2/24
  ip ospf network point-to-point
  ip router ospf 100 area 0.0.0.0
  mpls ip forwarding
  no shutdown

interface Ethernet1/52
  ip address 172.16.2.2/24
  ip ospf network point-to-point
  ip router ospf 100 area 0.0.0.0
  mpls ip forwarding
  no shutdown

interface loopback0
  ip address 192.168.1.3/32
  ip router ospf 100 area 0.0.0.0

router ospf 100
  segment-routing mpls
  router-id 192.168.1.3

PE-CE
vrf cu101
  address-family ipv4 unicast

interface Ethernet1/3
  no shutdown
interface Ethernet1/3.101
  encapsulation dot1q 101
  vrf member cu101
  ip address 61.1.1.1/24
  ip ospf network point-to-point
  ip router ospf 200 area 0.0.0.0
  no shutdown

interface loopback10
  vrf member cu101
  ip address 131.1.1.1/32
  ip router ospf 200 area 0.0.0.0

router ospf 200
  vrf cu101
  router-id 131.1.1.1
  redistribute bgp 65001 route-map bgp65001
  
```

BGP/EVPN Configuration

```

router bgp 65001
  router-id 192.168.1.3
  address-family ipv4 unicast
    network 192.168.1.3/32 route-map label-index-leaf-1
  allocate-label all
  address-family ipv4 labeled-unicast
  address-family l2vpn evpn

template peer EVPN
  remote-as 65001
  update-source loopback0
  address-family l2vpn evpn
  send-community extended
  encapsulation mpls

template peer Labeled-unicast
  remote-as 65001
  address-family ipv4 labeled-unicast
  send-community extended
  soft-reconfiguration inbound always

neighbor 172.16.0.1
  inherit peer Labeled-unicast
neighbor 172.16.2.1
  inherit peer Labeled-unicast
neighbor 192.168.1.1
  inherit peer EVPN
neighbor 192.168.1.2
  inherit peer EVPN

vrf cu101
  router-id 131.1.1.1
  address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map direct1
  redistribute ospf 200 route-map ospf200
  
```

LEAF-2 Configuration

Enabling Features, Label-Range, Route-map, Label-Index

```

install feature-set mpls
feature-set mpls
nv overlay evpn
feature ospf
feature bgp
feature mpls segment-routing
feature mpls evpn
feature interface-vlan
feature mpls oam
feature nv overlay

mpls label range 5000 450000
segment-routing
mpls
  global-block 16000 25000
  connected-prefix-sid-map
  address-family ipv4
    192.168.1.4/32 index 321

ip prefix-list new seq 5 permit 68.1.1.0/24
ip prefix-list new seq 10 permit 133.1.1.1/32

ip prefix-list new1 seq 5 permit 3.3.3.3/32

ip prefix-list redtoospf seq 5 permit 61.1.1.0/24
ip prefix-list redtoospf seq 10 permit 1.1.1.1/32

route-map bgp65001 permit 10
  match route-type internal
route-map direct1 permit 10
  match ip address prefix-list new
route-map label-index-Leaf2 permit 10
  set label-index 321
route-map ospf200 permit 10
  match ip address prefix-list new1

vrf context cu101
rd auto
address-family ipv4 unicast
route-target import 1:101
route-target import 1:101 evpn
route-target export 1:101
route-target export 1:101 evpn
  
```

OSPF Configuration

```

interface Ethernet1/51
ip address 172.16.1.2/24
ip ospf network point-to-point
ip router ospf 100 area 0.0.0.0
mpls ip forwarding
no shutdown

interface Ethernet1/52
ip address 172.16.3.2/24
ip ospf network point-to-point
ip router ospf 100 area 0.0.0.0
mpls ip forwarding
no shutdown

interface loopback0
ip address 192.168.1.4/32
ip router ospf 100 area 0.0.0.0

router ospf 100
segment-routing mpls
router-id 192.168.1.4

PE-CE
vrf cu101
  address-family ipv4 unicast
interface Ethernet1/55
no shutdown
interface Ethernet1/55.200
encapsulation dot1q 200
vrf member cu101
ip address 68.1.1.1/24
ip ospf network point-to-point
ip router ospf 200 area 0.0.0.0
no shutdown

interface loopback11
vrf member cu101
ip address 133.1.1.1/32
ip router ospf 200 area 0.0.0.0

router ospf 200
vrf cu101
  router-id 133.1.1.1
  redistribute bgp 65001 route-map bgp65001
  
```

BGP/EVPN Configuration

```

router bgp 65001
router-id 192.168.1.4
address-family ipv4 unicast
  network 192.168.1.4/32 route-map label-index-Leaf2
  allocate-label all
address-family ipv4 labeled-unicast
address-family l2vpn evpn

template peer EVPN
remote-as 65001
update-source loopback0
address-family l2vpn evpn
  send-community extended
  encapsulation mpls

template peer Labeled-unicast
remote-as 65001
address-family ipv4 labeled-unicast
  send-community extended
  soft-reconfiguration inbound always

neighbor 172.16.1.1
inherit peer Labeled-unicast
neighbor 172.16.3.1
inherit peer Labeled-unicast
neighbor 192.168.1.1
inherit peer EVPN
neighbor 192.168.1.2
inherit peer EVPN

vrf cu101
router-id 133.1.1.1
address-family ipv4 unicast
  advertise l2vpn evpn
  redistribute direct route-map direct1
  redistribute ospf 200 route-map ospf200
  
```

End-Host Configuration

Host-1 / Cat-3750

```
vrf definition custA
rd 101:1
!
address-family ipv4
exit-address-family
!

interface Loopback0
vrf forwarding custA
ip address 1.1.1.1 255.255.255.255

interface GigabitEthernet1/0/1
switchport trunk allowed vlan 101
switchport trunk encapsulation dot1q
switchport mode trunk
!

interface Vlan101
vrf forwarding custA
ip address 61.1.1.2 255.255.255.0
ip ospf network point-to-point
ip ospf 200 area 0.0.0.0

router ospf 200 vrf custA
router-id 1.1.1.1
network 1.1.1.1 0.0.0.0 area 0.0.0.0
network 61.1.1.0 0.0.0.255 area 0.0.0.0
```

Host-2 / N9K

```
feature ospf
feature interface-vlan

interface Ethernet1/52
switchport
switchport mode trunk
switchport trunk allowed vlan 200
no shutdown

interface Vlan200
no shutdown
ip address 68.1.1.2/24
ip ospf network point-to-point
ip router ospf 200 area 0.0.0.0

interface loopback0
ip address 3.3.3.3/32
ip router ospf 200 area 0.0.0.0

router ospf 200
router-id 3.3.3.3
```

Verifica

Host2# show ip int brief

```
IP Interface Status for VRF "default"(1)
Interface      IP Address      Interface Status
Vlan200        68.1.1.2        protocol-up/link-up/admin-up
Vlan1001       100.0.0.100    protocol-down/link-down/admin-up
Lo0            3.3.3.3         protocol-up/link-up/admin-up
```

Host2# show ip route

```
IP Route Table for VRF "default"
*** denotes best ucast next-hop
**** denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

1.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
3.3.3.3/32, ubest/mbest: 2/0, attached
  *via 3.3.3.3, Lo0, [0/0], 20:16:34, local
  *via 3.3.3.3, Lo0, [0/0], 20:16:34, direct
61.1.1.0/24, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
68.1.1.0/24, ubest/mbest: 1/0, attached
  *via 68.1.1.2, Vlan200, [0/0], 20:20:55, direct
68.1.1.2/32, ubest/mbest: 1/0, attached
  *via 68.1.1.2, Vlan200, [0/0], 20:20:55, local
131.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/1], 00:29:24, ospf-200, type-2, tag 65001
133.1.1.1/32, ubest/mbest: 1/0
  *via 68.1.1.1, Vlan200, [110/41], 20:15:32, ospf-200, intra
```

Host2# traceroute 1.1.1.1

```
traceroute to 1.1.1.1 (1.1.1.1), 30 hops max, 40 byte packets
 1 68.1.1.1 (68.1.1.1)  0.989 ms  0.585 ms  0.407 ms
 2 172.16.3.1 (172.16.3.1)  0.886 ms  172.16.1.1 (172.16.1.1)  0.765 ms  0.731 ms
   [Label=16311 E=0 TTL=1 S=0, Label=492289 E=0 TTL=1 S=1]
   [Label=16311 E=0 TTL=1 S=0, Label=492289 E=0 TTL=1 S=1]
 3 172.16.0.2 (172.16.0.2)  0.717 ms  172.16.2.2 (172.16.2.2)  0.509 ms  172.16.0.2 (172.16.0.2)  0.678 ms
   [Label=492289 E=0 TTL=1 S=1]
   [Label=492289 E=0 TTL=1 S=1]
 4 61.1.1.2 (61.1.1.2)  2.061 ms * 1.315 ms
```

Host2# ping 1.1.1.1 source 3.3.3.3

```
PING 1.1.1.1 (1.1.1.1) from 3.3.3.3: 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=251 time=5.538 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=251 time=1.338 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=251 time=2.201 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=251 time=2.217 ms
64 bytes from 1.1.1.1: icmp_seq=4 ttl=251 time=4.021 ms

--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min/avg/max = 1.338/3.063/5.538 ms
```

```
Leaf1# show ip route 3.3.3.3/32 vrf cu101
```

```
IP Route Table for VRF "cu101"  
** denotes best ucast next-hop  
*** denotes best mcast next-hop  
'[x/y]' denotes [preference/metric]  
'<string>' in via output denotes VRF <string>  
  
3.3.3.3/32, ubest/mbest: 1/0  
  *via 192.168.1.4%default, [200/2], 00:44:27, bgp-65001, internal, tag 65001 (mpls-vpn)
```

```
Leaf1# show forwarding mpls 192.168.1.4/32
```

```
slot 1  
=====
```

Local Label	Prefix Table Id	FEC (Prefix/Tunnel id)	Next-Hop	Interface	Out Label	
16321	0x1	192.168.1.4/32	172.16.0.1	Eth1/51	16321	SWAP
"	0x1	192.168.1.4/32	172.16.2.1	Eth1/52	16321	SWAP

```
Leaf1# show forwarding 3.3.3.3/32 vrf cu101
```

```
slot 1  
=====
```

IPv4 routes for table cu101/base

Prefix	Next-hop	Interface	Labels	Partial Install
*3.3.3.3/32	172.16.0.1	Ethernet1/51	PUSH 16321 492288	
	172.16.2.1	Ethernet1/52	PUSH 16321 492288	

Informazioni correlate

- [Multiprotocollo BGP MPLS VPN](#)
- [White paper sul routing dei segmenti degli switch con piattaforma Cisco Nexus 9500, 9300, 9200, 3200 e 3100](#)
- [Configurazione di EVPN di layer 3 e VPN di layer 3 su MPLS di routing del segmento](#)