

Implementazione dell'handoff da VxLAN/EVPN a MPLS/SR in Nexus 9000

Sommario

[Introduzione](#)

[Prerequisiti](#)

[Requisiti](#)

[Componenti usati](#)

[Premesse](#)

[Vantaggi della SR](#)

[Linee guida e limitazioni per configurare la perfetta integrazione di EVPN con L3VPN \(MPLS SR\)](#)

[Configurazione](#)

[Esempio di rete](#)

[Configurazione](#)

[Verifica](#)

[Informazioni correlate](#)

Introduzione

In questo documento viene descritto come configurare VxLAN VPN over Segment Routing MPLS handoff sugli switch Cisco Nexus serie 9000.

Prerequisiti

Requisiti

Cisco raccomanda la conoscenza dei seguenti argomenti:

- VPN MPLS layer 3
- MP-BGP
- Ciclo segmento

Componenti usati

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

- BL // N9K-C93240YC-FX2 // 9.3(3)
- ASR1K / 16.12.02
- Dorso-1 // N9K-C92160YC-X // 9.2(3)
- PE // N9K-C93240YC-FX2 // 9.3(3)

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

Le installazioni di centri dati (DC) hanno adottato VXLAN EVPN per i suoi vantaggi, quali apprendimento del control-plane EVPN, multitenancy, mobilità perfetta, ridondanza e aggiunte POD più semplici. 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 routing dei segmenti è un metodo di inoltro dei pacchetti sulla rete basato sul paradigma del routing di origine. L'origine sceglie un percorso e lo codifica nell'intestazione del pacchetto come un elenco ordinato di segmenti. I segmenti sono un identificatore per qualsiasi tipo di istruzione.

Vantaggi della SR

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

- Piani di controllo MPLS IGP unificati.
- Pronto per SDN: Il routing dei segmenti è stato creato per SDN ed è alla base di AER (Application Engineered Routing). SR prepara le reti per i modelli aziendali, dove le applicazioni possono indirizzare il comportamento della rete. SR offre il giusto equilibrio tra intelligenza distribuita e ottimizzazione e programmazione centralizzate.
- Configurazione minima: Il routing dei segmenti per TE richiede una configurazione minima sul router di origine.
- Bilanciamento del carico: A differenza di RSVP-TE, il bilanciamento del carico per il routing dei segmenti può essere eseguito in presenza di percorsi multipli (ECMP, Equal Cost Multiple Paths).
- Supporta Fast Reroute (FRR): Fast Reroute consente l'attivazione di un percorso di backup preconfigurato entro 50 millisecondi dall'errore di un percorso.
- Distribuzione plug-and-play: I criteri di routing dei segmenti sono interoperabili con i piani di controllo e dati MPLS esistenti e possono essere implementati in una distribuzione esistente.

Linee guida e limitazioni per configurare la perfetta integrazione di EVPN con L3VPN (MPLS SR)

| Funzionalità | Cisco Nexus 9300-FX2 Switch | Switch Cisco Nexus 9504 e 9508 con schede di linea -R | Commenti |
|-----------------------|-----------------------------|---|---|
| VXLAN VPN su SR-L3VPN | Sì | Sì | Estendere la connettività di layer 3 tra diverse pod DC Sottolineare IGP/BGP con estensioni SR. |

VXLAN VPN su SR-L3VPN Sì Sì

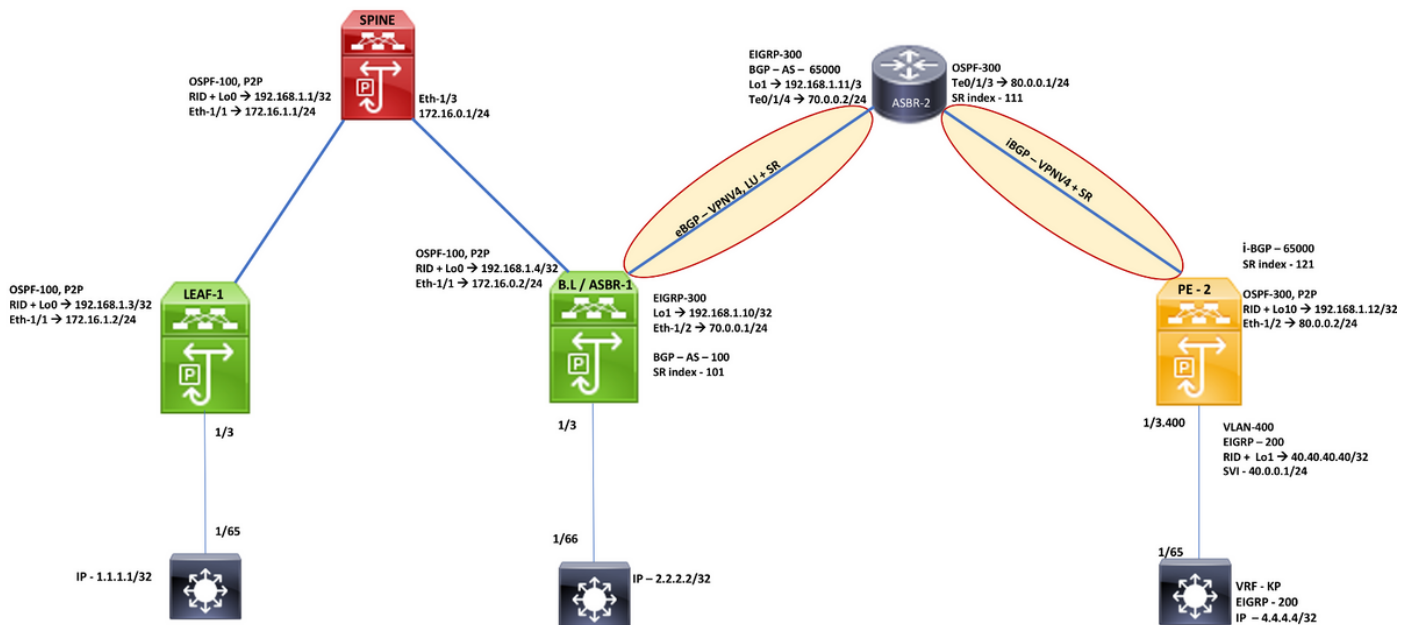
VXLAN VPN su MPLS L3VPN No Sì
(LDP)

Estendere la connettività di layer 3 tra il DC con VXLAN e qualsiasi dominio (DC o CORE) con SR.

Il sottostrato è LDP.

Configurazione

Esempio di rete



Configurazione

Border Leaf - ASBR1 Configuration

```
hostname BL
install feature-set mpls
  allow feature-set mpls
feature-set mpls
nv overlay evpn
feature ospf
feature bgp
feature pim
feature eigrp
feature mpls l3vpn
feature mpls segment-routing
feature interface-vlan
feature vn-segment-vlan-based
feature lldp
feature mpls oam
feature nv overlay

ip pim rp-address 192.168.1.1 group-list 224.0.0.0/4
ip pim ssm range 232.0.0.0/8
mpls label range 5000 450000
```

```
segment-routing
mpls
  global-block 16000 25000
  connected-prefix-sid-map
  address-family ipv4
  192.168.1.10/32 index 100
```

```
ip prefix-list adveig seq 5 permit 2.2.2.2/32
ip prefix-list localsub seq 5 permit 30.0.0.0/24
ip prefix-list localsub seq 10 permit 30.30.30.30/32
ip prefix-list localsub seq 15 permit 192.168.1.10/32
ip prefix-list localsub seq 20 permit 70.0.0.0/24
```

```
route-map direct permit 10
  match ip address prefix-list localsub
route-map eig permit 10
  match ip address prefix-list adveig
route-map lbgpout permit 10
  match route-type external internal
route-map label-index-Leaf2 permit 10
  set label-index 101
```

```
vlan 1,200,300
vlan 200
  vn-segment 201
vlan 300
  vn-segment 5000
```

```
interface Vlan200
  no shutdown
  vrf member KP
  ip address 30.0.0.1/24
  ip router eigrp 200
```

```
interface Vlan201
  no shutdown
  vrf member KP
  ip forward
```

```
vrf context KP
vni 5000
ip pim rp-address 192.168.1.1 group-list 224.0.0.0/4
rd auto
address-family ipv4 unicast
  route-target import 65000:1
  route-target export 65000:1
  route-target both auto
route-target both auto evpn
```

```
interface nve1
  no shutdown
  host-reachability-protocol bgp
  suppress mac-route
  source-interface loopback0
  member vni 201
  mcast-group 239.0.0.1
  member vni 5000 associate-vrf
```

```
interface loopback0
  ip address 192.168.1.4/32
  ip router ospf 100 area 0.0.0.0
  ip pim sparse-mode
```

```
interface loopback1
  ip address 192.168.1.10/32
  ip router eigrp 300
```

```
interface loopback10
  vrf member KP
  ip address 30.30.30.30/32
  ip router eigrp 200
```

```
interface Ethernet1/1
  ip address 172.16.0.2/24
  ip ospf network point-to-point
  ip router ospf 100 area 0.0.0.0
  ip pim sparse-mode
  no shutdown
```

```
interface Ethernet1/2
  mtu 9216
  ip address 70.0.0.1/24
  ip router eigrp 300
  mpls ip forwarding
  no shutdown
```

```
interface Ethernet1/3
  switchport
  switchport mode trunk
  switchport trunk allowed vlan 200-201,300
  no shutdown
```

```
router eigrp 200
  router-id 30.30.30.30
  vrf KP
  address-family ipv4 unicast
  redistribute bgp 100 route-map lbgpout
router eigrp 300
  router-id 192.168.1.10
  router ospf 100
  router-id 192.168.1.4
```

```
router bgp 100
  router-id 192.168.1.4
  address-family ipv4 unicast
  network 192.168.1.10/32 route-map label-index-Leaf2
  redistribute direct route-map direct
  allocate-label all
  address-family ipv6 unicast
  address-family vpnv4 unicast
  address-family ipv4 labeled-unicast
  address-family l2vpn evpn
  neighbor 70.0.0.2
  remote-as 65000
  ebgp-multihop 10
  address-family ipv4 labeled-unicast
  send-community
  send-community extended
  neighbor 192.168.1.1
  remote-as 100
  update-source loopback0
  address-family ipv4 unicast
  send-community extended
  address-family l2vpn evpn
  send-community extended
  import vpn unicast reoriginate
```

```
neighbor 192.168.1.11
  remote-as 65000
  update-source loopback1
  disable-connected-check
  ebgp-multihop 10
  address-family vpnv4 unicast
  send-community extended
  import l2vpn evpn reoriginate
  address-family l2vpn evpn
  send-community extended
  rewrite-evpn-rt-asn
vrf KP
  address-family ipv4 unicast
  redistribute direct route-map direct
  redistribute eigrp 200 route-map eig
```

ASR1K - ASBR2 Configuration

```
hostname P-Router
interface Loopback0
ip address 192.168.1.11 255.255.255.255
!
!
segment-routing mpls
global-block 16000 25000
!
connected-prefix-sid-map
address-family ipv4
192.168.1.11/32 index 111 range 1
exit-address-family
!
```

```
!
interface TenGigabitEthernet0/1/3
ip address 80.0.0.1 255.255.255.0
ip ospf network point-to-point
cdp enable
mpls ip
!
router ospf 300
router-id 192.168.1.11
segment-routing area 0 mpls
segment-routing mpls
segment-routing prefix-sid-map advertise-local
redistribute connected
redistribute static
redistribute bgp 65000
network 80.0.0.0 0.0.0.255 area 0
network 192.168.1.11 0.0.0.0 area 0
!
!
router eigrp 300
network 70.0.0.0 0.0.0.255
network 192.168.1.11 0.0.0.0
eigrp router-id 192.168.1.11
!
!
interface TenGigabitEthernet0/1/4
mtu 9216
ip address 70.0.0.2 255.255.255.0
ip ospf network point-to-point
cdp enable
mpls ip
mpls bgp forwarding
!
```

```
!
router bgp 65000
bgp router-id 192.168.1.11
no bgp transport path-mtu-discovery
bgp log-neighbor-changes
no bgp default route-target filter
neighbor 70.0.0.1 remote-as 100
neighbor 70.0.0.1 ebgp-multihop 10
neighbor 192.168.1.10 remote-as 100
neighbor 192.168.1.10 ebgp-multihop 10
neighbor 192.168.1.10 disable-connected-check
neighbor 192.168.1.10 update-source Loopback0
neighbor 192.168.1.12 remote-as 65000
neighbor 192.168.1.12 update-source Loopback0
!
address-family ipv4
redistribute connected
redistribute ospf 300
segment-routing mpls
neighbor 70.0.0.1 activate
neighbor 70.0.0.1 send-community both
neighbor 70.0.0.1 send-label
neighbor 192.168.1.10 activate
neighbor 192.168.1.10 send-community both
neighbor 192.168.1.10 soft-reconfiguration inbound
neighbor 192.168.1.10 send-label
no neighbor 192.168.1.12 activate
exit-address-family
!
address-family vpnv4
neighbor 192.168.1.10 activate
neighbor 192.168.1.10 send-community both
neighbor 192.168.1.12 activate
neighbor 192.168.1.12 send-community both
neighbor 192.168.1.12 next-hop-self
exit-address-family
!
address-family l2vpn evpn
exit-address-family
!
```

Far-End PE2 Configuration

```
hostname PE2
install feature-set mpls
allow feature-set mpls
feature-set mpls
feature ospf
feature bgp
feature eigrp
feature mpls l3vpn
feature mpls segment-routing
feature interface-vlan
feature lldp
feature mpls oam

vlan 1400
segment-routing
mpls
global-block 16000 25000
connected-prefix-sid-map
address-family ipv4
192.168.1.12/32 index 121

ip prefix-list adveig seq 5 permit 4.4.4.4/32
ip prefix-list localsub seq 5 permit 40.0.0.0/24
ip prefix-list localsub seq 10 permit 40.40.40.40/32
route-map direct permit 10
match ip address prefix-list localsub
route-map eig permit 10
match ip address prefix-list adveig
route-map ibgpout permit 10
match tag 100
match route-type internal

route-map label-index-Leaf2 permit 10
set label-index 121

vrf context KP
rd auto
address-family ipv4 unicast
route-target import 100:5000
route-target import 65000:1
route-target export 100:5000
route-target export 65000:1
```

```
interface Ethernet1/2
ip address 80.0.0.2/24
ip ospf network point-to-point
ip router ospf 300 area 0.0.0.0
mpls ip forwarding
no shutdown

interface loopback10
ip address 192.168.1.12/32
ip router ospf 300 area 0.0.0.0

router ospf 300
segment-routing mpls
router-id 192.168.1.12
distance 115

interface Ethernet1/3
switchport
switchport mode trunk
switchport trunk allowed vlan 400
no shutdown

interface Vlan400
no shutdown
vrf member KP
ip address 40.0.0.1/24
ip router eigrp 200

interface loopback1
vrf member KP
ip address 40.40.40.40/32
ip router eigrp 200
```

```
router bgp 65000
router-id 192.168.1.12
address-family ipv4 unicast
network 192.168.1.12/32 route-map label-index-Leaf2
redistribute direct route-map direct
redistribute eigrp 200 route-map eig
allocate-label all
address-family vpnv4 unicast
address-family ipv4 labeled-unicast
neighbor 192.168.1.11
remote-as 65000
update-source loopback10
address-family vpnv4 unicast
send-community extended
vrf KP
address-family ipv4 unicast
redistribute direct route-map direct
redistribute eigrp 200 route-map eig
```

Verifica

Host-1 & Leaf-1 Verification

host1# sh ip int br | ex down

```
IP Interface Status for VRF "default"(1)
Interface IP Address Interface Status
Vlan201 20.0.0.2 protocol-up/link-up/admin-up
Lo0 1.1.1.1 protocol-up/link-up/admin-up
```

host1# traceroute 4.4.4.4 source 1.1.1.1

```
traceroute to 4.4.4.4 (4.4.4.4) from 1.1.1.1 (1.1.1.1), 30 hops max, 40 byte packets
 1 20.0.0.1 (20.0.0.1) 0.92 ms 0.531 ms 0.513 ms
 2 30.0.0.1 (30.0.0.1) 1.043 ms 0.819 ms 0.733 ms
 3 40.0.0.1 (40.0.0.1) 0.912 ms 0.673 ms 0.624 ms
 [Label=492287 E=0 TTL=253 S=1]
 4 4.4.4.4 (4.4.4.4) 0.882 ms 0.825 ms 0.561 ms
host1#
```

Leaf1_N3k# show bgp l2vpn evpn

BGP routing table information for VRF default, address family L2VPN EVPN
 BGP table version is 42, Local Router ID is 192.168.1.3
 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, l-injected
 Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|---|-------------|--------|--------|--------|------|
| Route Distinguisher: 192.168.1.4:3 | | | | | |
| *>i[5]:[0]:[0]:[24]:[30.0.0.0]/224 | 192.168.1.4 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[24]:[40.0.0.0]/224 | 192.168.1.4 | 100 | 0 | 65000 | ? |
| *>i[5]:[0]:[0]:[32]:[2.2.2.2]/224 | 192.168.1.4 | 130816 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[4.4.4.4]/224 | 192.168.1.4 | 100 | 0 | 65000 | ? |
| *>i[5]:[0]:[0]:[32]:[30.30.30.30]/224 | 192.168.1.4 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[40.40.40.40]/224 | 192.168.1.4 | 100 | 0 | 65000 | ? |
| Route Distinguisher: 192.168.1.3:3 (L3VNI 5000) | | | | | |
| *>i[5]:[0]:[0]:[24]:[20.0.0.0]/224 | 192.168.1.3 | 0 | 100 | 32768 | ? |
| *>i[5]:[0]:[0]:[24]:[30.0.0.0]/224 | 192.168.1.4 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[24]:[40.0.0.0]/224 | 192.168.1.4 | 100 | 0 | 65000 | ? |
| *>i[5]:[0]:[0]:[32]:[1.1.1.1]/224 | 192.168.1.3 | 130816 | 100 | 32768 | ? |
| *>i[5]:[0]:[0]:[32]:[2.2.2.2]/224 | 192.168.1.4 | 130816 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[4.4.4.4]/224 | 192.168.1.4 | 100 | 0 | 65000 | ? |
| *>i[5]:[0]:[0]:[32]:[20.20.20.20]/224 | 192.168.1.3 | 0 | 100 | 32768 | ? |
| *>i[5]:[0]:[0]:[32]:[30.30.30.30]/224 | 192.168.1.4 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[40.40.40.40]/224 | 192.168.1.4 | 100 | 0 | 65000 | ? |

Leaf1_N3k#

Border Leaf Verification

BL# sh bgp l2vpn evpn

BGP routing table information for VRF default, address family L2VPN EVPN
 BGP table version is 30, Local Router ID is 192.168.1.4
 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, l-injected
 Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|---|-------------|--------|--------|--------|------|
| Route Distinguisher: 192.168.1.3:3 | | | | | |
| *>i[5]:[0]:[0]:[24]:[20.0.0.0]/224 | 192.168.1.3 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[1.1.1.1]/224 | 192.168.1.3 | 130816 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[20.20.20.20]/224 | 192.168.1.3 | 0 | 100 | 0 | ? |
| Route Distinguisher: 192.168.1.4:3 (L3VNI 5000) | | | | | |
| *>i[5]:[0]:[0]:[24]:[20.0.0.0]/224 | 192.168.1.3 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[24]:[30.0.0.0]/224 | 192.168.1.4 | 0 | 100 | 32768 | ? |
| *>i[5]:[0]:[0]:[24]:[40.0.0.0]/224 | 192.168.1.4 | 0 | 65000 | ? | ? |
| *>i[5]:[0]:[0]:[32]:[1.1.1.1]/224 | 192.168.1.3 | 130816 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[2.2.2.2]/224 | 192.168.1.4 | 130816 | 100 | 32768 | ? |
| *>i[5]:[0]:[0]:[32]:[4.4.4.4]/224 | 192.168.1.4 | 0 | 65000 | ? | ? |
| *>i[5]:[0]:[0]:[32]:[20.20.20.20]/224 | 192.168.1.3 | 0 | 100 | 0 | ? |
| *>i[5]:[0]:[0]:[32]:[30.30.30.30]/224 | 192.168.1.4 | 0 | 100 | 32768 | ? |
| *>i[5]:[0]:[0]:[32]:[40.40.40.40]/224 | 192.168.1.4 | 0 | 65000 | ? | ? |

BL#

BL# show bgp vpnv4 unicast

BGP routing table information for VRF default, address family VPNv4 Unicast
 BGP table version is 81, Local Router ID is 192.168.1.4
 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, l-injected
 Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath, & - backup, 2 - best2

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|---|--------------|--------|--------|--------|------|
| Route Distinguisher: 192.168.1.4:3 (VRF KP) | | | | | |
| *>i1.1.1.1/32 | 192.168.1.3 | 130816 | 100 | 0 | ? |
| *>2.2.2.2/32 | 0.0.0.0 | 130816 | 100 | 32768 | ? |
| *>e4.4.4.4/32 | 192.168.1.11 | 0 | 65000 | ? | ? |
| *>20.0.0.0/24 | 192.168.1.3 | 0 | 100 | 0 | ? |
| *>30.0.0.0/24 | 192.168.1.3 | 0 | 100 | 0 | ? |
| *>30.0.0.0/24 | 0.0.0.0 | 0 | 100 | 32768 | ? |
| *>30.30.30.30/32 | 0.0.0.0 | 0 | 100 | 32768 | ? |
| *>e40.0.0.0/24 | 192.168.1.11 | 0 | 65000 | ? | ? |
| *>e40.40.40.40/32 | 192.168.1.11 | 0 | 65000 | ? | ? |
| Route Distinguisher: 192.168.1.12:3 | | | | | |
| *>e4.4.4.4/32 | 192.168.1.11 | 0 | 65000 | ? | ? |
| *>e40.0.0.0/24 | 192.168.1.11 | 0 | 65000 | ? | ? |
| *>e40.40.40.40/32 | 192.168.1.11 | 0 | 65000 | ? | ? |

BL#

BL# show segment-routing mpls clients

Segment-Routing Mpls Client Info

Client: bgp-100
 PIB Index: 1 UUID: 0x11b PID: 4611 MTS SAP: 7255
 TIBs registered:
 VRF: default Table: base

Total Clients: 1
 BL#

BL# show segment-routing mpls ipv4 connected-prefix-sid-map
 Segment-Routing Mpls Prefix-SID Mappings
 Prefix-SID mappings for VRF default Table base

| Prefix | SID | Type | Range | SRGB |
|-----------------|-----|------|-------|------|
| 192.168.1.10/32 | 100 | Indx | 1 | Y |

*SRGB - Indicates whether Prefix-Sid is within configured SRGB
 BL#

BL# show segment-routing mpls detail

Segment-Routing Mpls Global Info
 Service Name: segment-routing
 State: Enabled
 Process ID: 2505
 Configured SRGB: 16000 - 25000
 SRGB Allocation status: Alloc-Successful
 Current SRGB: 16000 - 25000
 Cleanup Interval: 60
 Retry Interval: 180
 SRGB ULIB handle: 4096002
 ULIB RegStatus: Done
 ULIB PIB handle: 0x2
 BL#

BL# show mpls switching

Legend:
 (P)=Protected, (F)=FRR active, (*)=more labels in stack.

IPv4:
 In-Label Out-Label FEC name Out-Interface Next-Hop
 VRF default
 5000 Pop Label 80.0.0.0/24 Eth1/2 70.0.0.2
 16111 Pop Label 192.168.1.11/32 Eth1/2 70.0.0.2
 16121 16121 192.168.1.12/32 Eth1/2 70.0.0.2

In-Label VRF
 492288 default
 492287 KP

| Local | Out-Label | Out-Interface | Next-Hop |
|---------------------|----------------|---------------|--------------|
| FEC: 4.4.4.4/32 | 192.168.1.12:3 | 5001 22 | 192.168.1.11 |
| FEC: 40.0.0.0/24 | 192.168.1.12:3 | 5002 23 | 192.168.1.11 |
| FEC: 40.40.40.40/32 | 192.168.1.12:3 | 5003 24 | 192.168.1.11 |

| Block | Label-Range |
|-------|---------------|
| 1 | 16000 - 25000 |

BL#

