Configurazione del router Fusion in SDA

Sommario

Introduzione

Questo documento descrive come configurare i router Fusion in una soluzione Cisco Software-Defined Access (SDA).

Prerequisiti

Requisiti

Nessun requisito specifico previsto per questo documento.

> Nota: è necessario eseguire l'installazione come descritto in Dispositivi supportati, disponibile in Collegamento alle note sulla versione

Componenti usati

Le informazioni di questo documento si basano sulle seguenti versioni hardware:

- DNAC - versione 1.2.1
- Edge and Border - Switch Cat3k
- Fusion - Router Cisco con supporto per perdite tra VRF

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

Nella soluzione Cisco SD-Access, i dispositivi sono gestiti e configurati da Cisco DNA Center. In generale, tutte le parti della struttura SD-Access possono essere configurate e gestite da Cisco DNA Center, come di norma lo sono. Tuttavia, il dispositivo Fusion si trova all'esterno della struttura, pertanto viene configurato manualmente. L'automazione dei confini, di cui si parla più avanti, è una funzionalità di Cisco DNA Center in grado di automatizzare la configurazione dei bordi per il trasferimento di VRF ai dispositivi Fusion.

Talvolta, per ragioni tipicamente legate alla compatibilità con la configurazione corrente, l'automazione dei bordi non è adatta e quindi anche il passaggio dal bordo al dispositivo Fusion può essere configurato manualmente. La comprensione della configurazione utilizzata consente di illustrare importanti dettagli sulla configurazione e sul funzionamento ottimali del sistema complessivo.

Funzionalità di un dispositivo di fusione nella soluzione DNA SD-Access

Un dispositivo Fusion consente il routing e l'inoltro virtuale (VRF) che causa la perdita di dati nei domini dei fabric ad accesso SD e consente la connettività host ai servizi condivisi, ad esempio DHCP, DNS, NTP, ISE, Cisco DNA Center, Wireless LAN Controller (WLC) e simili. Anche se questo ruolo può essere svolto da dispositivi diversi dai router, questo documento si concentra sui router come dispositivi Fusion.

Come accennato in precedenza, i servizi condivisi devono essere resi disponibili a tutte le reti virtuali (VPN) del campus. A tale scopo, è possibile creare peer Border Gateway Protocol (BGP) dai router di confine ai router di fusione. Sul router Fusion, le subnet del VRF del fabric che necessitano di accedere a questi servizi condivisi vengono trapelate nella GRT, o in un VRF dei servizi condivisi, e viceversa. Le route map possono essere utilizzate per contenere tabelle di routing a subnet specifiche di SD-Access Fabric.

> Nota: i nodi di confine ad accesso SD non supportano route di riepilogo che si sovrappongono ai pool IP ad accesso SD. Le route di riepilogo che si sovrappongono ai pool IP devono essere filtrate negli annunci di instradamento dai dispositivi Fusion ai nodi di confine.

Configurazione

I dettagli di configurazione forniti qui si riferiscono alla topologia di rete mostrata di seguito. Questa topologia di rete non è consigliata per le distribuzioni. Viene utilizzato unicamente per facilitare la presentazione degli esempi di configurazione forniti. Per i progetti di installazione consigliati, vedere Design Zone per Cisco Digital Network Architecture.

Esempio di rete

La topologia utilizzata in questo articolo è composta da due router di confine configurati entrambi come bordi esterni e due router di fusione con una connessione a ciascun router di confine rispettivo.

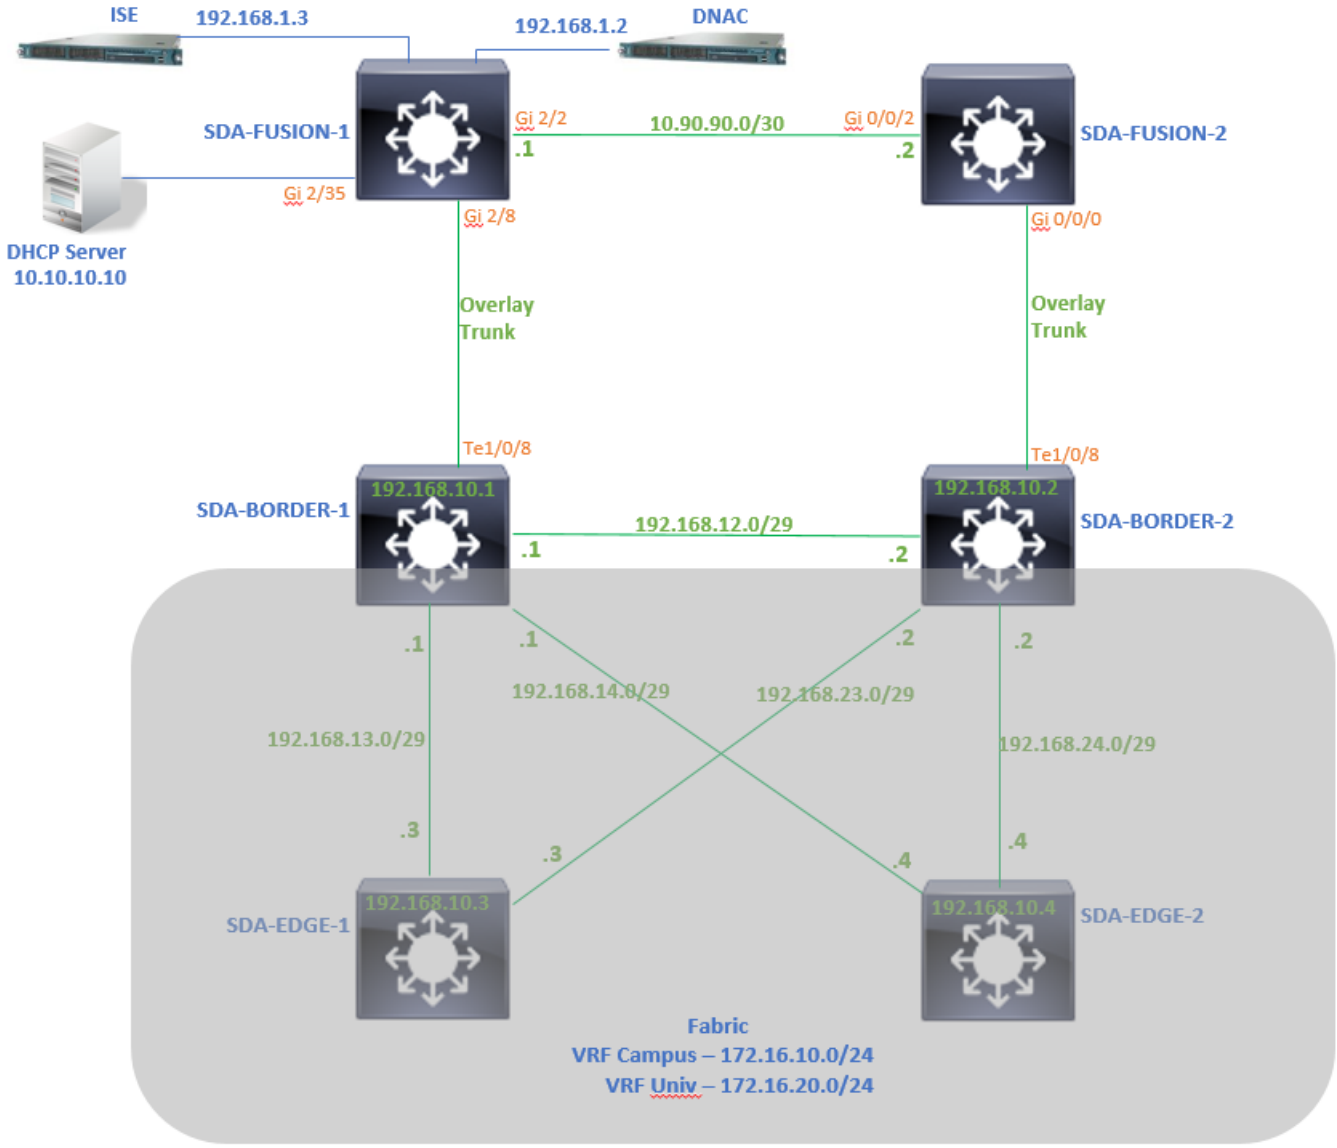

Configurazioni

Passaggio 1. Configurare il collegamento di handoff da DNAC

Durante la procedura di assegnazione dei dispositivi al ruolo di Border Router durante l'aggiunta di quest'ultimo al fabric, è possibile creare un collegamento handoff. Al layer 2 si tratta di un

collegamento trunk collegato al router Fusion. Sono necessarie le seguenti misure:

1. Configurare il numero AS locale per BGP. Questo numero Autonomous System (AS) viene usato per configurare il processo BGP sui router di confine.

2. Aggiungere l'interfaccia sotto Transit. Questa interfaccia è la connessione diretta tra Border e Fusion Router. (in questo esempio, 1/0/8 su Bordo).

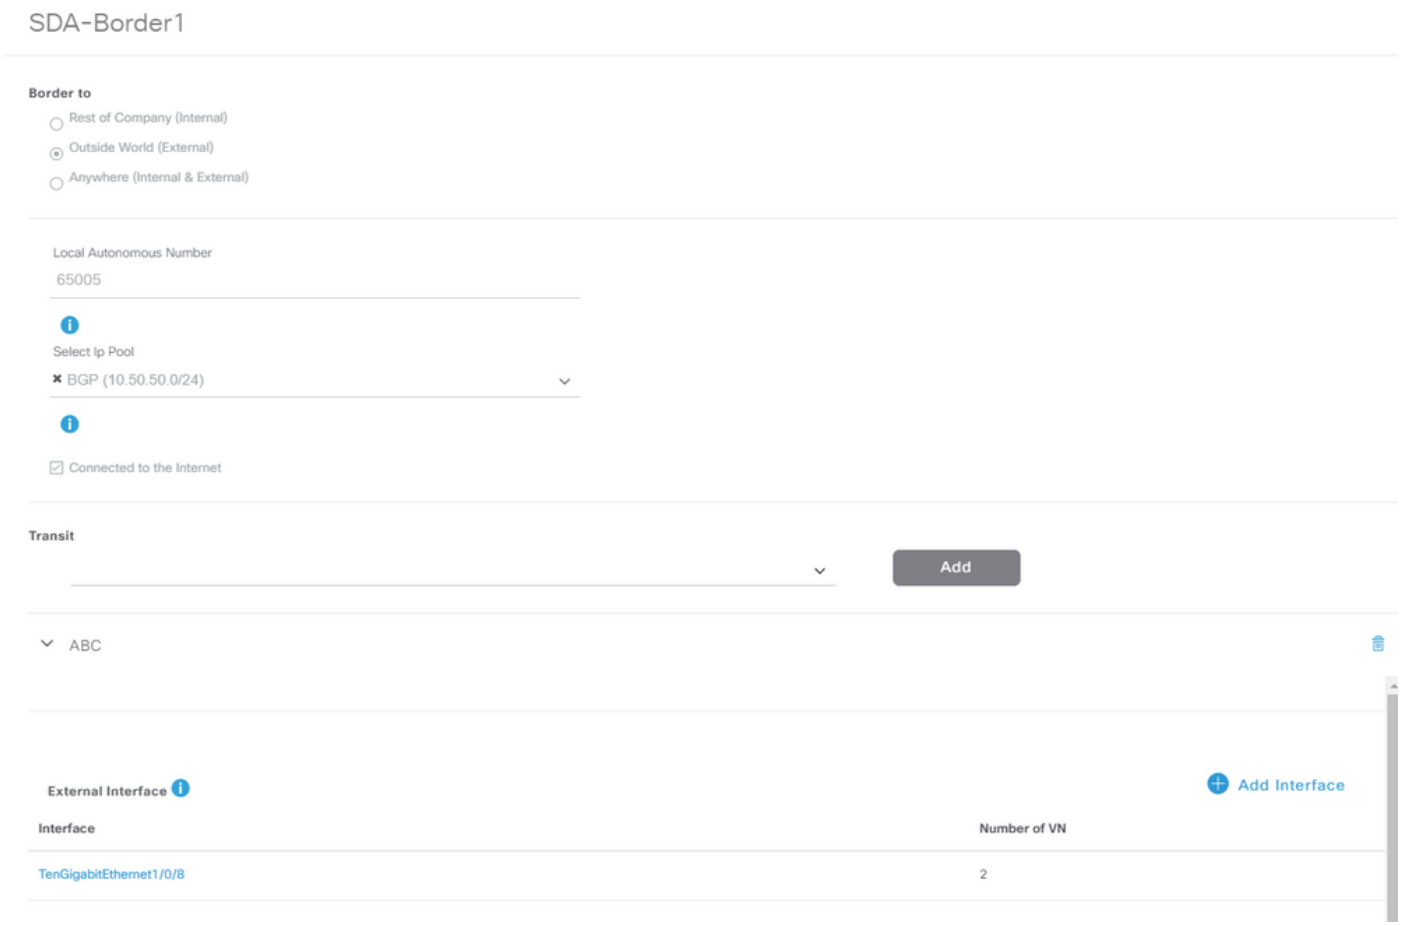

3. Configurare il numero AS remoto. Questo numero AS viene usato sui router di confine per le istruzioni dei router adiacenti verso il router Fusion per configurare i peer BGP (eBGP) esterni.

4. Selezionare tutte le reti virtuali (VRF) per le quali è richiesta una perdita VRF sul router Fusion.

5. Distribuire la configurazione da DNAC ai dispositivi.

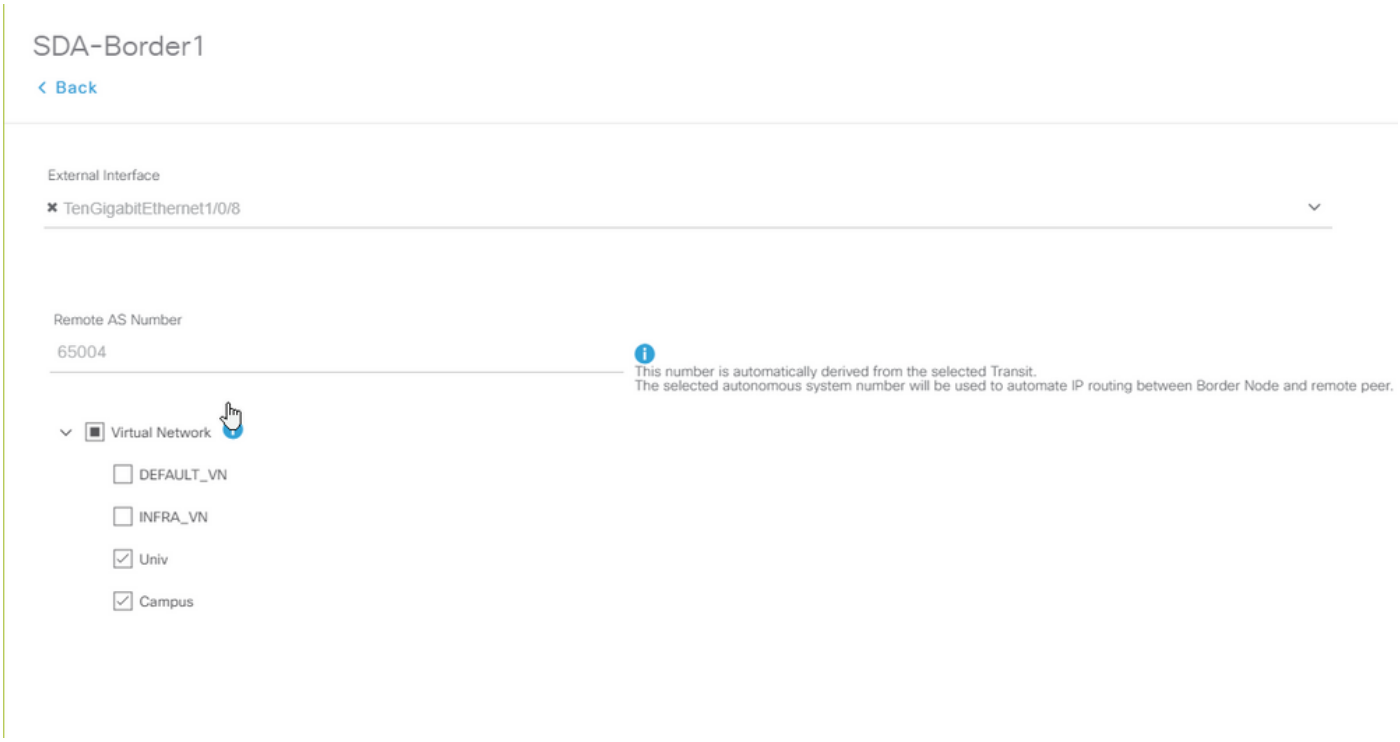

Eseguire la stessa procedura per il dispositivo SDA-Border-2.

Passaggio 2. Verifica delle configurazioni sottoposte a push da DNAC sui router di confine

In questa sezione viene illustrata la verifica della configurazione sui router di confine relativi al protocollo BGP.

SDA-Bordo-1

```
SDA-Border1#show run interface loopback 0
!
interface Loopback0
ip address 192.168.10.1 255.255.255.255
ip router isis
end


SDA-Border1#show run interface tenGigabitEthernet 1/0/8
!
interface TenGigabitEthernet1/0/8
switchport mode trunk
end


SDA-Border1#show run interface loopback 1021

interface Loopback1021
description Loopback Border
vrf forwarding Campus
ip address 172.16.10.1 255.255.255.255
end


SDA-Border1#show run interface loopback 1022
```

```
interface Loopback1022
description Loopback Border
vrf forwarding Univ
ip address 172.16.20.1 255.255.255.255
end


SDA-Border1#show run | section vrf definition Campus
vrf definition Campus
rd 1:4099
!
address-family ipv4
route-target export 1:4099
route-target import 1:4099
exit-address-family


SDA-Border1#show run | section vrf definition Univ
vrf definition Univ
rd 1:4100
!
address-family ipv4
route-target export 1:4100
route-target import 1:4100
exit-address-family
SDA-Border1#


SDA-Border1#show run interface vlan 3007
!
interface Vlan3007 <<< SVI created for BGP Peering under VRF Campus
description vrf interface to External router
vrf forwarding Campus
ip address 10.50.50.25 255.255.255.252
no ip redirects
ip route-cache same-interface
end

SDA-Border1#show run interface vlan 3006
!
interface Vlan3006 <<< SVI created for BGP Peering under VRF Univ
description vrf interface to External router
vrf forwarding Univ
ip address 10.50.50.21 255.255.255.252
no ip redirects
ip route-cache same-interface
end


SDA-Border1#show run | section bgp
router bgp 65005 <<< Local AS Number from DNAC
bgp router-id interface Loopback0
bgp log-neighbor-changes
bgp graceful-restart
!
address-family ipv4
network 192.168.10.1 mask 255.255.255.255
redistribute lisp metric 10
exit-address-family
!
address-family ipv4 vrf Campus
bgp aggregate-timer 0
network 172.16.10.1 mask 255.255.255.255 <<< Anycast IP for Pool in VRF Campus
aggregate-address 172.16.10.0 255.255.255.0 summary-only <<< Only Summary is Advertised
```

```
redistribute lisp metric 10
neighbor 10.50.50.26 remote-as 65004 <<< Peer IP to be used on Fusion for VRF Campus and Remote
AS Number from DNAC
neighbor 10.50.50.26 update-source Vlan3007
neighbor 10.50.50.26 activate
neighbor 10.50.50.26 weight 65535 <<< Weight needed for Fusion peering to make sure locally
originated path from LISP is never preferred
exit-address-family
!
address-family ipv4 vrf Univ
bgp aggregate-timer 0
network 172.16.20.1 mask 255.255.255.255 <<< Anycast IP for Pool in VRF Univ
aggregate-address 172.16.20.0 255.255.255.0 summary-only
redistribute lisp metric 10
neighbor 10.50.50.22 remote-as 65004
neighbor 10.50.50.22 update-source Vlan3006
neighbor 10.50.50.22 activate
neighbor 10.50.50.22 weight 65535
exit-address-family
```

SDA-Bordo-2

```
SDA-Border2#show run interface loopback 0
!
interface Loopback0
 ip address 192.168.10.2 255.255.255.255
 ip router isis
end

SDA-Border2#show run interface tenGigabitEthernet 1/0/8
!
interface TenGigabitEthernet1/0/8
 switchport mode trunk
end


SDA-Border2#show run interface loopback 1021
!
interface Loopback1021
description Loopback Border
vrf forwarding Campus
ip address 172.16.10.1 255.255.255.255
end

SDA-Border2#show run interface loopback 1022
!
interface Loopback1022
description Loopback Border
vrf forwarding Univ
ip address 172.16.20.1 255.255.255.255
end
SDA-Border2#show run | section vrf definition Campus vrf definition Campus rd 1:4099 ! address-
family ipv4 route-target export 1:4099 route-target import 1:4099 exit-address-family SDA-
Border2#show run | section vrf definition Univ vrf definition Univ rd 1:4100 ! address-family
ipv4 route-target export 1:4100 route-target import 1:4100 exit-address-family SDA-Border2#show
run interface vlan 3001 ! interface Vlan3001 description vrf interface to External router vrf
forwarding Campus ip address 10.50.50.1 255.255.255.252 no ip redirects ip route-cache same-
interface end SDA-Border2#show run interface vlan 3003 ! interface Vlan3003 description vrf
interface to External router vrf forwarding Univ ip address 10.50.50.9 255.255.255.252 no ip
redirects ip route-cache same-interface end SDA-Border2#show run | section bgp router bgp 65005
bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart ! address-family
```

```
ipv4 network 192.168.10.2 mask 255.255.255.255 redistribute lisp metric 10 exit-address-family !
address-family ipv4 vrf Campus bgp aggregate-timer 0 network 172.16.10.1 mask 255.255.255.255
aggregate-address 172.16.10.0 255.255.255.0 summary-only redistribute lisp metric 10 neighbor
10.50.50.2 remote-as 65004 neighbor 10.50.50.2 update-source Vlan3001 neighbor 10.50.50.2
activate neighbor 10.50.50.2 weight 65535 exit-address-family ! address-family ipv4 vrf Univ bgp
aggregate-timer 0 network 172.16.20.1 mask 255.255.255.255 aggregate-address 172.16.20.0
255.255.255.0 summary-only redistribute lisp metric 10 neighbor 10.50.50.10 remote-as 65004
neighbor 10.50.50.10 update-source Vlan3003 neighbor 10.50.50.10 activate neighbor 10.50.50.10
weight 65535 exit-address-family
```

Passaggio 3. Configurazione di allowas-in su router di confine

A causa della perdita di VRF sul router di fusione, la famiglia di indirizzi ipv4 per il campus VRF apprende la route originata da VRF Univ (172.16.20.0/24). Tuttavia, sia il router di origine che quello di apprendimento hanno lo stesso numero BGP AS (65005). Per superare i meccanismi di prevenzione del loop BGP e accettare/installare le route sui router di confine, è necessario configurare **allowas-in** per i peer con il router di fusione:

```
SDA-Border1

SDA-Border1(config)#router bgp 65005
SDA-Border1(config-router)#address-family ipv4 vrf Campus
SDA-Border1(config-router-af)#neighbor 10.50.50.26 allowas-in
SDA-Border1(config-router-af)#exit-address-family
SDA-Border1(config-router)#
SDA-Border1(config-router)#address-family ipv4 vrf Univ
SDA-Border1(config-router-af)#neighbor 10.50.50.22 allowas-in
SDA-Border1(config-router-af)#exit-address-family
SDA-Border1(config-router)#


SDA-Border2

SDA-Border2(config)#router bgp 65005
SDA-Border2(config-router)#address-family ipv4 vrf Campus
SDA-Border2(config-router-af)#neighbor 10.50.50.2 allowas-in
SDA-Border2(config-router-af)#exit-address-family
SDA-Border2(config-router)#
SDA-Border2(config-router)#address-family ipv4 vrf Univ
SDA-Border2(config-router-af)#neighbor 10.50.50.10 allowas-in
SDA-Border2(config-router-af)#exit-address-family
SDA-Border2(config-router)#
```

> **Nota**: il comando **allowas-in** deve essere usato per precauzione perché può causare loop. Quando si utilizza un solo dispositivo Fusion con cui entrambi i Bordi si intersecano, è necessario filtrare per assicurarsi che le route originate localmente non vengano accettate di nuovo nel SA dal peer Fusion - all'interno della stessa VPN. In questo caso, il percorso eBGP è preferito al percorso originato localmente a causa del peso massimo dei percorsi eBGP.

Passaggio 4. Configurazione router Fusion

In questa sezione viene illustrata la configurazione manuale per i router Fusion.

SDA-Fusion-1

Configurare il collegamento al router di confine come trunk in modo che corrisponda alla configurazione vlan sul bordo 1:

```
interface GigabitEthernet2/8
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3006, 3007
 switchport mode trunk
end
```

Configurare i VRF richiesti:

```
vrf definition Campus
 rd 1:4099
 !
 address-family ipv4
  route-target export 1:4099
  route-target import 1:4099
 exit-address-family
!

vrf definition Univ
 rd 1:4100
 !
 address-family ipv4
  route-target export 1:4100
  route-target import 1:4100
 exit-address-family
```

Configurare le interfacce SVI:

```
interface Vlan3007
 vrf forwarding Campus
 ip address 10.50.50.26 255.255.255.252
end

interface Vlan3006
 vrf forwarding Univ
 ip address 10.50.50.22 255.255.255.252
end
```

Configurare il peer BGP (eBGP) esterno con SDA-Border-1:

```
router bgp 65004                                    <<< Remote AS from DNAC
 bgp log-neighbor-changes
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv4 vrf Campus
  neighbor 10.50.50.25 remote-as 65005
  neighbor 10.50.50.25 update-source Vlan3007
  neighbor 10.50.50.25 activate
 exit-address-family
 !
```

```
 address-family ipv4 vrf Univ
  neighbor 10.50.50.21 remote-as 65005
  neighbor 10.50.50.21 update-source Vlan3006
  neighbor 10.50.50.21 activate
 exit-address-family
```

Configurare il peer BGP (iBGP) interno con SDA-Fusion-2:

```
interface GigabitEthernet2/2
 description SDA-Fusion1--->SDA-Fusion2
 ip address 10.90.90.1 255.255.255.252
end

router bgp 65004
 neighbor 10.90.90.2 remote-as 65004
 !
 address-family ipv4
  neighbor 10.90.90.2 activate
 exit-address-family
 !
```

Annunciare la subnet del server DHCP nella famiglia di indirizzi globale in cui l'indirizzo IP del server DHCP è 10.10.10.10:

```
interface GigabitEthernet2/35
 description connection to DHCP server
 ip address 10.10.10.9 255.255.255.252
end


router bgp 65004
 !
 address-family ipv4
  network 10.10.10.8 mask 255.255.255.252
 exit-address-family
 !
```

SDA-Fusion-2

Configurare il collegamento verso Border Router. Se un'interfaccia su Fusion è L3 anziché trunk, configurare le sottointerfacce:

```
interface GigabitEthernet0/0/0.3001
 encapsulation dot1Q 3001
 vrf forwarding Campus
 ip address 10.50.50.2 255.255.255.252
end

interface GigabitEthernet0/0/0.3003
 encapsulation dot1Q 3003
 vrf forwarding Univ
 ip address 10.50.50.10 255.255.255.252
end
```

Configurare i VRF corrispondenti:

```
vrf definition Campus
 rd 1:4099
 !
 address-family ipv4
  route-target export 1:4099
  route-target import 1:4099
 exit-address-family
!
!
vrf definition Univ
 rd 1:4100
 !
 address-family ipv4
  route-target export 1:4100
  route-target import 1:4100
 exit-address-family
!
```

Configurare il peer eBGP con SDA-Border-2:

```
router bgp 65004
 bgp log-neighbor-changes
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv4 vrf Campus
  neighbor 10.50.50.1 remote-as 65005
  neighbor 10.50.50.1 update-source GigabitEthernet0/0/0.3001
  neighbor 10.50.50.1 activate
 exit-address-family
 !
 address-family ipv4 vrf Univ
  neighbor 10.50.50.9 remote-as 65005
  neighbor 10.50.50.9 update-source GigabitEthernet0/0/0.3003
  neighbor 10.50.50.9 activate
 exit-address-family
```

Configurare il peer iBGP con SDA-Fusion-1:

```
interface GigabitEthernet0/0/2
ip address 10.90.90.2 255.255.255.252
negotiation auto
end

router bgp 65004 neighbor 10.90.90.1 remote-as 65004 ! address-family ipv4 neighbor 10.90.90.1
activate exit-address-family
```

Passaggio 5. Configurazione della perdita VRF sul router Fusion

La configurazione per le perdite VRF è identica per entrambi i router Fusion SDA-Fusion-1 e SDA-Fusion-2.

In primo luogo, configurare le perdite VRF tra i due VRF (Campus e Univ), utilizzare l'**importazione route-target**:

```
vrf definition Campus
 !
 address-family ipv4
route-target export 1:4099 route-target import 1:4099
route-target import 1:4100 <<< Import VRF Univ prefixes in VRF Campus
exit-address-family
!
vrf definition Univ
!
address-family ipv4
route-target export 1:4100 route-target import 1:4100
route-target import 1:4099 <<< Import VRF Campus prefixes in VRF Univ
exit-address-family
!
```

Quindi, configurare la perdita di route tra la tabella di routing globale (GRT) e i VRF e tra i VRF e la GRT, utilizzare **import ... map** and **export ... map**:

```
ip prefix-list Campus_Prefix seq 5 permit 172.16.10.0/24    <<< Include Prefixes belonging to
VRF Campus
ip prefix-list Global_Prefix seq 5 permit 10.10.10.8/30     <<< Include Prefixes belonging to
Global (eq DHCP Server Subnet)
ip prefix-list Univ_Prefix seq 5 permit 172.16.20.0/24      <<< Include Prefixes belonging to
VRF Univ

route-map Univ_Map permit 10
 match ip address prefix-list Univ_Prefix
route-map Global_Map permit 10
 match ip address prefix-list Global_Prefix
route-map Campus_Map permit 10
 match ip address prefix-list Campus_Prefix


vrf definition Campus
 !
 address-family ipv4
  import ipv4 unicast map Global_Map     <<< Injecting Global into VRF Campus matching route-map
Global_Map
export ipv4 unicast map Campus_Map <<< Injecting VRF Campus into Global matching route-map
Campus_Map
exit-address-family
!
vrf definition Univ
!
address-family ipv4
import ipv4 unicast map Global_Map <<< Injecting Global into VRF Univ matching route-map
Global_Map
export ipv4 unicast map Univ_Map <<< Injecting VRF Univ into Global matching route-map Univ_Map
exit-address-family
!
```

Verifica

In questa sezione viene descritto come verificare che la configurazione precedente sia stata eseguita correttamente.

Passaggio 1. Verifica del peer eBGP tra router di fusione e router di confine

SDA-Border-1 —Peering—SDA-Fusion-1

```
SDA-Border1#show ip bgp vpnv4 vrf Campus summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.26     4        65004    1294    1295       32    0    0 19:32:22        2



SDA-Border1#show ip bgp vpnv4 vrf Univ summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.22     4        65004    1294    1292       32    0    0 19:32:57        2


-----------------------------------------------


SDA-Fusion1#show ip bgp vpnv4 vrf Campus summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.25     4        65005    1305    1305       31    0    0 19:41:58        1


SDA-Fusion1#show ip bgp vpnv4 vrf Univ summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.21     4        65005    1303    1305       31    0    0 19:42:14        1
```

SDA-Border-2 —Peering—SDA-Fusion-2

```
SDA-Border2#show ip bgp vpnv4 vrf Campus summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.2      4        65004       6       6       61    0    0 00:01:37        2


SDA-Border2#show ip bgp vpnv4 vrf Univ summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.10     4        65004       6       6       61    0    0 00:01:39        2

---------------------------------------------------------------------------

SDA-Fusion2#show ip bgp vpnv4 vrf Campus summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.1      4        65005      17      17        9    0    0 00:11:16        1


SDA-Fusion2#show ip bgp vpnv4 vrf Univ summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.50.50.9      4        65005      17      17        9    0    0 00:11:33        1
```

Passaggio 2. Verifica del peer iBGP tra entrambi i router Fusion

SDA-Fusion-1 —Peering—SDA-Fusion-2

```
SDA-Fusion1#show ip bgp summary
```

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.90.90.2      4        65004      10      12       12    0    0 00:04:57        2


------------------------------------------------------------

SDA-Fusion2#show ip bgp summary

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.90.90.1      4        65004      19      17        4    0    0 00:11:35        3
```

Passaggio 3. Verifica prefissi nella tabella BGP e nella tabella di routing

SDA-Bordo-1

```
SDA-Border1#show ip bgp vpnv4 vrf Campus

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:4099 (default for vrf Campus)
 *>   10.10.10.8/30    10.50.50.26                            65535 65004 i         <<< Prefix
leaked from Global Routing Table on Fusion
 *>   172.16.10.0/24   0.0.0.0                                32768 i               <<< VRF Campus
originated prefix
 *>   172.16.20.0/24   10.50.50.26                            65535 65004 65005 i   <<< Prefix
originated in VRF Univ, leaked on Fusion to VRF Campus


SDA-Border1#show ip route vrf Campus bgp Routing Table: Campus B 10.10.10.8/30 [20/0] via
10.50.50.26, 20:30:30 <<< RIB entry for DHCP Server pool prefix B 172.16.10.0/24 [200/0],
20:32:45, Null0 <<< Null entry created by "aggregate-address" BGP configuration B 172.16.20.0/24
[20/0] via 10.50.50.26, 20:32:45 <<< RIB entry for VRF Univ prefix ----------------------------
------------------------------------------- SDA-Border1#show ip bgp vpnv4 vrf Univ Network
Next Hop Metric LocPrf Weight Path Route Distinguisher: 1:4100 (default for vrf Univ) *>
10.10.10.8/30 10.50.50.22 65535 65004 i <<< Prefix leaked from Global Routing Table on Fusion *>
172.16.10.0/24 10.50.50.22 65535 65004 65005 i <<< Prefix originated in VRF Campus, leaked on
Fusion to VRF Univ *> 172.16.20.0/24 0.0.0.0 32768 i <<< VRF Univ originated prefix SDA-
Border1#show ip route vrf Univ bgp Routing Table: Univ B 10.10.10.8/30 [20/0] via 10.50.50.22,
20:31:06 <<< RIB entry for DHCP Server pool prefix B 172.16.10.0/24 [20/0] via 10.50.50.22,
20:33:21 <<< RIB entry for VRF Campus prefix B 172.16.20.0/24 [200/0], 20:33:21, Null0 <<< Null
entry created by "aggregate-address" BGP configuration
```

SDA-Bordo-2

```
SDA-Border2#show ip bgp vpnv4 vrf Campus

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:4099 (default for vrf Campus)
 *>   10.10.10.8/30    10.50.50.2                             65535 65004 i         <<< Prefix
leaked from Global Routing Table on Fusion
 *>   172.16.10.0/24   0.0.0.0                                32768 i               <<< VRF Campus
originated prefix
 *>   172.16.20.0/24   10.50.50.2                             65535 65004 65005 i   <<< Prefix
originated in VRF Univ, leaked on Fusion to VRF Campus


SDA-Border2#show ip route vrf Campus bgp

B        10.10.10.8/30 [20/0] via 10.50.50.2, 01:02:19          <<< RIB entry for DHCP Server
pool prefix
B        172.16.10.0/24 [200/0], 1w6d, Null0                    <<< Null entry created by
"aggregate-address" BGP configuration
```

```
B        172.16.20.0/24 [20/0] via 10.50.50.2, 01:02:27         <<< RIB entry for VRF Univ
Prefix

------------------------------------------------------------------

SDA-Border2#show ip bgp vpnv4 vrf Univ

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:4100 (default for vrf Univ)
 *>   10.10.10.8/30    10.50.50.10                          65535 65004 i         <<< Prefix
leaked from Global Routing Table on Fusion
 *>   172.16.10.0/24   10.50.50.10                          65535 65004 65005 i   <<< Prefix
originated in VRF Campus, leaked on Fusion to VRF Univ
 *>   172.16.20.0/24   0.0.0.0                              32768 i               <<< VRF Univ
originated prefix


SDA-Border2#show ip route vrf Univ bgp

B        10.10.10.8/30 [20/0] via 10.50.50.10, 01:02:29          <<< RIB entry for DHCP Server
pool prefix
B        172.16.10.0/24 [20/0] via 10.50.50.10, 01:02:34         <<< RIB entry for VRF Campus
prefix
B        172.16.20.0/24 [200/0], 1w6d, Null0                     <<< Null entry created by
"aggregate-address" BGP configuration
```

SDA-Fusion-1

```
SDA-Fusion1#show ip bgp

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.10.10.8/30    0.0.0.0                  0         32768 i            <<< Locally
originated Global prefix
 * i 172.16.10.0/24   10.50.50.1               0    100      0 65005 i      <<< Prefix imported
from VRF Campus
 *>                   10.50.50.25              0             0 65005 i
 * i 172.16.20.0/24   10.50.50.9               0    100      0 65005 i      <<< Prefix imported
from VRF Univ
 *>                   10.50.50.21              0             0 65005 i


SDA-Fusion1#show ip route

C        10.10.10.8/30 is directly connected, GigabitEthernet2/35            <<< Prefix for DHCP
Server
B        172.16.10.0 [20/0] via 10.50.50.25 (Campus), 20:50:21               <<< Prefix imported
from VRF Campus
B        172.16.20.0 [20/0] via 10.50.50.21 (Univ), 20:50:21          <<< Prefix imported from
VRF Univ

------------------------------------------------------------------------


SDA-Fusion1#show ip bgp vpnv4 vrf Campus

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:4099 (default for vrf Campus)
Import Map: Global_Map, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
Export Map: Campus_Map, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
 *>  10.10.10.8/30    0.0.0.0                  0         32768 i         <<< Prefix imported
```

```
from Global Routing
 *>  172.16.10.0/24   10.50.50.25              0             0 65005 i <<< Prefix learnt from
Border1 in VRF Campus
 *>  172.16.20.0/24   10.50.50.21              0             0 65005 i <<< Prefix imported fron
VRF Univ


SDA-Fusion1#show ip bgp vpnv4 vrf Campus 172.16.20.0/24
BGP routing table entry for 1:4099:172.16.20.0/24, version 27
Paths: (1 available, best #1, table Campus)
Advertised to update-groups:
5
Refresh Epoch 1
65005, (aggregated by 65005 192.168.10.1), imported path from 1:4100:172.16.20.0/24 (Univ)
10.50.50.21 (via vrf Univ) (via Univ) from 10.50.50.21 (192.168.10.1)
Origin IGP, metric 0, localpref 100, valid, external, atomic-aggregate, best
Extended Community: RT:1:4100
rx pathid: 0, tx pathid: 0x0


SDA-Fusion1#show ip route vrf Campus bgp B 10.10.10.8/30 is directly connected, 20:46:51,
GigabitEthernet2/35 B 172.16.10.0 [20/0] via 10.50.50.25, 20:50:07 B 172.16.20.0 [20/0] via
10.50.50.21 (Univ), 20:50:07 ----------------------------------------------------------------
--------- SDA-Fusion1#show ip bgp vpnv4 vrf Univ Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 1:4100 (default for vrf Univ) Import Map: Global_Map, Address-Family: IPv4
Unicast, Pfx Count/Limit: 1/1000 Export Map: Univ_Map, Address-Family: IPv4 Unicast, Pfx
Count/Limit: 1/1000 *> 10.10.10.8/30 0.0.0.0 0 32768 i <<< Prefix imported from Global Routing
*> 172.16.10.0/24 10.50.50.25 0 0 65005 i <<< Prefix imported fron VRF Campus *> 172.16.20.0/24
10.50.50.21 0 0 65005 i <<< Prefix learnt from Border1 in VRF Univ

SDA-Fusion1#show ip bgp vpnv4 vrf Univ 172.16.10.0/24
BGP routing table entry for 1:4100:172.16.10.0/24, version 25
Paths: (1 available, best #1, table Univ)
Advertised to update-groups:
4
Refresh Epoch 1
65005, (aggregated by 65005 192.168.10.1), imported path from 1:4099:172.16.10.0/24 (Campus)
10.50.50.25 (via vrf Campus) (via Campus) from 10.50.50.25 (192.168.10.1)
Origin IGP, metric 0, localpref 100, valid, external, atomic-aggregate, best
Extended Community: RT:1:4099
rx pathid: 0, tx pathid: 0x0


SDA-Fusion1#show ip route vrf Univ bgp B 10.10.10.8/30 is directly connected, 20:47:01,
GigabitEthernet2/35 B 172.16.10.0 [20/0] via 10.50.50.25 (Campus), 20:50:17 B 172.16.20.0 [20/0]
via 10.50.50.21, 20:50:17
```

SDA-Fusion-2

```
SDA-Fusion2#show ip bgp

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 10.10.10.8/30    10.90.90.1               0    100      0 i
 *>  172.16.10.0/24   10.50.50.1               0             0 65005 i
 * i                  10.50.50.25              0    100      0 65005 i
 *>  172.16.20.0/24   10.50.50.9               0             0 65005 i
 * i                  10.50.50.21              0    100      0 65005 i


SDA-Fusion2#show ip route

B        10.10.10.8/30 [200/0] via 10.90.90.1, 01:25:56
B        172.16.10.0 [20/0] via 10.50.50.1 (Campus), 01:25:56
```

```
B        172.16.20.0 [20/0] via 10.50.50.9 (Univ), 01:25:56

------------------------------------------------------------------------------

SDA-Fusion2#show ip bgp vpnv4 vrf Campus

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:4099 (default for vrf Campus)
Import Map: Global_Map, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
Export Map: Campus_Map, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
 *>i 10.10.10.8/30    10.90.90.1               0    100      0 i
 *>  172.16.10.0/24   10.50.50.1               0             0 65005 i
 *>  172.16.20.0/24   10.50.50.9               0             0 65005 i


SDA-Fusion2#show ip route vrf Campus bgp


B        10.10.10.8/30 [200/0] via 10.90.90.1, 01:26:09
B        172.16.10.0 [20/0] via 10.50.50.1, 01:26:13
B        172.16.20.0 [20/0] via 10.50.50.9 (Univ), 01:26:13

------------------------------------------------------------------------------

SDA-Fusion2#show ip bgp vpnv4 vrf Univ

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:4100 (default for vrf Univ)
Import Map: Global_Map, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
Export Map: Univ_Map, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
 *>i 10.10.10.8/30    10.90.90.1               0    100      0 i
 *>  172.16.10.0/24   10.50.50.1               0             0 65005 i
 *>  172.16.20.0/24   10.50.50.9               0             0 65005 i


SDA-Fusion2#show ip route vrf Univ bgp


B        10.10.10.8/30 [200/0] via 10.90.90.1, 01:26:19
B        172.16.10.0 [20/0] via 10.50.50.1 (Campus), 01:26:23
B        172.16.20.0 [20/0] via 10.50.50.9, 01:26:23
```

Configurazione manuale per la ridondanza dei bordi

Per la ridondanza tra i PETR quando un collegamento esterno di confine si interrompe, per le frontiere esterne e esterne+interne, è necessario creare manualmente sessioni iBGP tra i due bordi per ciascuna VN. Inoltre, in caso di confine esterno+interno in cui BGP viene importato in LISP e LISP viene ridistribuito in BGP, sono necessarie delle etichette per impedire l'importazione di route iBGP in LISP ed evitare quindi potenziali loop.

SDA-Bordo-1

```
interface Vlan31
 description vrf interface to SDA-Border-2
 vrf forwarding Campus
 ip address 10.31.1.1 255.255.255.252
!
interface Vlan33
```

```
 description vrf interface to SDA-Border-2
 vrf forwarding Univ
 ip address 10.33.1.1 255.255.255.252
!

router bgp 65005
!
address-family ipv4 vrf Campus
redistribute lisp metric 10 <<< open redistribution pushed by DNAC
neighbor 10.31.1.2 remote-as 65005     <<< iBGP peering with SDA-Border-2
neighbor 10.31.1.2 activate
neighbor 10.31.1.2 send-community     <<< we need to send community/tag to the neighbor
neighbor 10.31.1.2 route-map tag_local_eids out    <<< route-map used to tag prefixes sent out
!
address-family ipv4 vrf Univ
redistribute lisp metric 10
 neighbor 10.33.1.2 remote-as 65005
neighbor 10.33.1.2 activate
neighbor 10.33.1.2 send-community
neighbor 10.33.1.2 route-map tag_local_eids out
!

router lisp
!
instance-id 4099
service ipv4
eid-table vrf Campus
 route-import database bgp 65005 route-map DENY-Campus locator-set rloc_a0602921-91eb-4e27-a294-
f88949a1ca37 <<< pushed by DNAC if Border is (also) Internal
!
 instance-id 4103
 service ipv4
 eid-table vrf Univ
 route-import database bgp 65005 route-map DENY-Univ locator-set rloc_a0602921-91eb-4e27-a294-
f88949a1ca37
!

ip community-list 1 permit 655370 <<< community-list matching tag 655370 - pushed by DNAC
!

route-map DENY-Campus deny 5 <<< route-map pushed by DNAC and used in route-import
match ip address prefix-list Campus
!
route-map DENY-Campus deny 10
match ip address prefix-list l3handoff-prefixes
!
route-map DENY-Campus deny 15
match community 1 <<< match on community-list 1 to deny iBGP prefixes to be imported into LISP
!
route-map DENY-Campus deny 25
match ip address prefix-list deny_0.0.0.0
!
route-map DENY-Campus permit 30
!

route-map DENY-Univ deny 5 <<< similar route-map is pushed for Univ VN
match ip address prefix-list Univ
!
route-map DENY-Univ deny 10
match ip address prefix-list l3handoff-prefixes
!
route-map DENY-Univ deny 15
match community 1
!
```

```
route-map DENY-Univ deny 25
match ip address prefix-list deny_0.0.0.0
!
route-map DENY-Univ permit 30
!

route-map tag_local_eids permit 5 <<< route-map we need to create in order to tag the routes
advertised to the iBGP peer
set community 655370 <<< setting community/tag to 655370
!
```

SDA-Bordo-2

```
interface Vlan31
 description vrf interface to SDA-Border-1
 vrf forwarding Campus
 ip address 10.31.1.2 255.255.255.252
!
interface Vlan33
 description vrf interface to SDA-Border-1
 vrf forwarding Univ
 ip address 10.33.1.2 255.255.255.252
!

router bgp 65005
!
address-family ipv4 vrf Campus
 neighbor 10.31.1.1 remote-as 65005
 neighbor 10.31.1.1 activate
 neighbor 10.31.1.1 send-community
 neighbor 10.31.1.1 route-map tag_local_eids out
!
address-family ipv4 vrf Univ
 neighbor 10.33.1.1 remote-as 65005
 neighbor 10.33.1.1 activate
 neighbor 10.33.1.1 send-community
 neighbor 10.33.1.1 route-map tag_local_eids out
!

router lisp
!
 instance-id 4099
 service ipv4
  eid-table vrf Campus
route-import database bgp 65005 route-map DENY-Campus locator-set rloc_677c0a8a-0802-49f9-99cc-
f9c6ebda80f3     <<< pushed by DNAC
!

 instance-id 4103
  service ipv4
   eid-table vrf Univ
route-import database bgp 65005 route-map DENY-Univ locator-set rloc_677c0a8a-0802-49f9-99cc-
f9c6ebda80f3
!

ip community-list 1 permit 655370
!

route-map DENY-Campus deny 5
match ip address prefix-list Campus
!
route-map DENY-Campus deny 10
match ip address prefix-list l3handoff-prefixes
```

```
!
route-map DENY-Campus deny 15
match community 1
!
route-map DENY-Campus deny 25
match ip address prefix-list deny_0.0.0.0
!
route-map DENY-Campus permit 30
!

route-map DENY-Univ deny 5
match ip address prefix-list Univ
!
route-map DENY-Univ deny 10
match ip address prefix-list l3handoff-prefixes
!
route-map DENY-Univ deny 15
match community 1
!
route-map DENY-Univ deny 25
match ip address prefix-list deny_0.0.0.0
!
route-map DENY-Univ permit 30
!

route-map tag_local_eids permit 5
set community 655370
!
```

Semplificazione della configurazione di Fusion con l'utilizzo di modelli

In questa sezione vengono forniti esempi di configurazione di Fusion Template per semplificare la configurazione.

Di seguito sono riportate le variabili che devono essere definite in base al progetto di distribuzione. In questo esempio, le configurazioni e le VN si basano sulla topologia precedente che ha due VN, Campus e Univ.

Definizione variabile

```
interface_Fusion1: GigabitEthernet2/8
interface_Fusion2: GigabitEthernet0/0/0

Global_prefixes = 10.10.10.8/30

FUSION_BGP_AS = 65004
BORDER_BGP_AS = 65005
```

Per VN1:

```
VN1 = Campus
Fusion1_VN1_VLAN = 3007
Fusion2_VN1_VLAN = 3001
VN1_prefixes = 172.16.10.0/24

Fusion1_VN1_IP = 10.50.50.26
```

```
Fusion1_VN1_MASK = 255.255.255.252

Fusion2_VN1_IP = 10.50.50.2

Fusion2_VN1_MASK = 255.255.255.252
VN1_RD = 4099
VN1_ border1_neighbor_IP = 10.50.50.25
VN1_border2_neighbor_IP = 10.50.50.1
```

Per VN2:

```
VN2 = Univ
Fusion1_VN2_VLAN = 3006
Fusion2_VN2_VLAN = 3003
VN2_prefixes = 172.16.20.0/24

Fusion1_VN2_IP = 10.50.50.22

Fusion1_VN2_MASK = 255.255.255.252
Fusino2_VN2_IP2 = 10.50.50.10

Fusion2_VN2_MASK = 255.255.255.252
VN2_RD = 4100
VN2_border1_neighbor_IP = 10.50.50.21
VN2_border2_neighbor_IP = 10.50.50.9
```

Esempio di modello

Fusione 1

```
interface $interface_Fusion1
switchport
switchport mode trunk
switchport trunk allowed vlan add $Fusion1_VN1_VLAN, $Fusion1_VN2_VLAN
!
vlan $Fusion1_VN1_VLAN
no shut
!
vlan $Fusion1_VN2_VLAN
no shut
!
vrf definition $VN1
rd 1:$VN1_RD
!
address-family ipv4
route-target export 1:$VN1_RD
route-target import 1:$VN1_RD
route-target import 1:$VN2_RD
exit-address-family
!
vrf definition $VN2
rd 1:$VN2_RD
!
address-family ipv4
route-target export 1:$VN2_RD
route-target import 1:$VN2_RD
route-target import 1:$VN1_RD
exit-address-family
!
interface Vlan $Fusion1_VN1_VLAN
vrf forwarding $VN1
```

```
ip address $Fusion1_VN1_IP $Fusion1_VN1_MASK
!
interface Vlan $Fusion1_VN2_VLAN
vrf forwarding $VN2
ip address $Fusion1_VN2_IP $Fusion1_VN2_MASK
!
router bgp $FUSION_BGP_AS
bgp log-neighbor-changes
!
address-family ipv4
exit-address-family
!
address-family ipv4 vrf $VN1
neighbor $VN1_border1_neighbor_IP remote-as $BORDER_BGP_AS
neighbor $VN1_border1_neighbor_IP update-source Vlan $Fusion1_VN1_VLAN
neighbor $VN1_border1_neighbor_IP activate
exit-address-family
!
address-family ipv4 vrf $VN2
neighbor $VN2_border1_neighbor_IP remote-as $BORDER_BGP_AS
neighbor $VN2_border1_neighbor_IP update-source $Fusion1_VN2_VLAN
neighbor $VN2_border1_neighbor_IP activate
exit-address-family

ip prefix-list ${VN1}_Prefix seq 5 permit $VN1_prefixes
ip prefix-list Global_Prefix seq 5 permit $Global_prefixes
ip prefix-list ${VN2}_Prefix seq 5 permit $VN2_prefixes

route-map ${VN2}_Map permit 10
match ip address prefix-list ${VN2}_Prefix
route-map Global_Map permit 10
match ip address prefix-list Global_Prefix
route-map ${VN1}_Map permit 10
match ip address prefix-list ${VN1}_Prefix

vrf definition $VN1
!
address-family ipv4
import ipv4 unicast map Global_Map
export ipv4 unicast map ${VN1}_Map
exit-address-family
!
vrf definition $VN2
!
address-family ipv4
import ipv4 unicast map Global_Map
export ipv4 unicast map ${VN2}_Map
exit-address-family
!
```

Fusion 2

```
interface $interface_Fusion2.$Fusion2_VN1_VLAN
encapsulation dot1Q $Fusion2_VN1_VLAN
vrf forwarding $VN1
ip address $Fusion2_VN1_IP2 $Fusion2_VN1_MASK
!
interface $interface_Fusion2.$Fusion2_VN2_VLAN
encapsulation dot1Q $Fusion2_VN2_VLAN
vrf forwarding $VN2
ip address $Fusion2_VN2_IP2 $Fusion2_VN2_MASK
!
```

```
vlan $Fusion2_VN1_VLAN
no shut
!
vlan $Fusion2_VN2_VLAN
no shut
!
vrf definition $VN1
rd 1:$VN1_RD
!
address-family ipv4
route-target export 1:$VN1_RD
route-target import 1:$VN1_RD
route-target import 1:$VN2_RD
exit-address-family
!
vrf definition $VN2
rd 1:$VN2_RD
!
address-family ipv4
route-target export 1:$VN2_RD
route-target import 1:$VN2_RD
route-target import 1:$VN1_RD
exit-address-family
!
router bgp $FUSION_BGP_AS
bgp log-neighbor-changes
!
address-family ipv4
exit-address-family
!
address-family ipv4 vrf $VN1
neighbor $VN1_border2_neighbor_IP remote-as $BORDER_BGP_AS
neighbor $VN1_border2_neighbor_IP update-source $interface_Fusion2.$Fusion2_VN1_VLAN
neighbor $VN1_bordre2_neighbor_IP activate
exit-address-family
!
address-family ipv4 vrf $VN2
neighbor $VN2_border2_neighbor_IP remote-as $BORDER_BGP_AS
neighbor $VN2_border2_neighbor_IP update-source $interface_Fusion2.$Fusion2_VN2_VLAN
neighbor $VN2_border2_neighbor_IP activate
exit-address-family

ip prefix-list ${VN1}_Prefix seq 5 permit $VN1_prefixes
ip prefix-list Global_Prefix seq 5 permit $Global_prefixes
ip prefix-list ${VN2}_Prefix seq 5 permit $VN2_prefixes

route-map ${VN2}_Map permit 10
match ip address prefix-list ${VN2}_Prefix
route-map Global_Map permit 10
match ip address prefix-list Global_Prefix
route-map ${VN}_Map permit 10
match ip address prefix-list ${VN1}_Prefix

vrf definition $VN1
!
address-family ipv4
import ipv4 unicast map Global_Map
export ipv4 unicast map ${VN1}_Map
exit-address-family
!
vrf definition $VN2
!
address-family ipv4
import ipv4 unicast map Global_Map
```

```
export ipv4 unicast map ${VN2}_Map
exit-address-family
!
End
```

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).