

Configuración del router de fusión en SDA

Contenido

[Introducción](#)

[Prerequisites](#)

[Requirements](#)

[Componentes Utilizados](#)

[Antecedentes](#)

[Funcionalidad de un dispositivo Fusion en la solución Cisco DNA SD-Access](#)

[Configurar](#)

[Diagrama de la red](#)

[Configuraciones](#)

[Paso 1. Configuración del enlace de entrega](#)

[Paso 2. Verificar las configuraciones introducidas en los routers de borde](#)

[Paso 3. Configuración de Allowas-In en Routers de Borde](#)

[Paso 4. Configurar routers de fusión](#)

[Paso 5. Configuración de la fuga de VRF en el router Fusion](#)

[Verificación](#)

[Paso 1. Verificación del Peering eBGP entre los Routers de Fusión y de Borde](#)

[Paso 2. Verifique el Peering iBGP entre ambos Routers Fusion](#)

[Paso 3. Verificar Prefijos en Tabla BGP y Tabla de Ruteo](#)

[Configuración manual para redundancia de bordes](#)

[SDA-Border-1](#)

[SDA-Border-2](#)

[Simplificación de la configuración de fusión con el uso de plantillas](#)

[Definición de variable](#)

[Ejemplo de plantilla](#)

[Fusión 1](#)

[Fusión 2](#)

Introducción

Este documento describe cómo configurar los routers Fusion en una solución Cisco Software-Defined Access (SDA).

Prerequisites

Requirements

No hay requisitos específicos para este documento.



Nota: La configuración es obligatoria según los dispositivos admitidos, que se pueden

Componentes Utilizados

La información que contiene este documento se basa en las siguientes versiones de hardware:

- Cisco Digital Network Architecture Controller, versión 1.2.1
- Perímetro y frontera - Switch Cisco Cat3k
- Fusion - Router de Cisco compatible con fugas Inter-VRF

La información que contiene este documento se creó a partir de los dispositivos en un ambiente de laboratorio específico. Todos los dispositivos que se utilizan en este documento se pusieron en funcionamiento con una configuración verificada (predeterminada). Si tiene una red en vivo, asegúrese de entender el posible impacto de cualquier comando.

Antecedentes


En la solución Cisco SD-Access, Cisco Catalyst Center administra y configura los dispositivos. En general, todas las partes del fabric de SD-Access se pueden configurar y gestionar, y normalmente se gestionan, mediante Cisco Catalyst Center. Sin embargo, el dispositivo Fusion está fuera del fabric, por lo que se configura manualmente. La automatización de los bordes, que se trata a continuación, es una función de Cisco Catalyst Center que puede automatizar la configuración de los bordes para la transferencia de VRF a los dispositivos Fusion.

En ocasiones, por razones típicamente relacionadas con la compatibilidad con la configuración actual, la automatización de bordes no es adecuada, y por lo tanto la transferencia de la frontera al dispositivo Fusion también se puede configurar a mano. La comprensión de la configuración que se utiliza ayuda a ilustrar detalles importantes sobre la configuración y el funcionamiento óptimos del sistema en general.

Funcionalidad de un dispositivo Fusion en la solución Cisco DNA SD-Access

Un dispositivo Fusion permite la fuga de Virtual Routing and Forwarding (VRF) a través de dominios de fabric de acceso SD y permite la conectividad de host a servicios compartidos, como DHCP, DNS, NTP, ISE, Cisco Catalyst Center, Wireless LAN Controllers (WLC) y similares. Aunque esta función la pueden realizar otros dispositivos distintos de los routers, este documento se centra en los routers como dispositivos Fusion.

Como se ha mencionado anteriormente, los servicios compartidos deben estar disponibles para todas las redes virtuales (VPN) del campus. Esto se logra con la creación de pares BGP (Border Gateway Protocol) desde los routers de borde a los routers de fusión. En el router Fusion, las subredes del VRF de fabric que necesitan acceso a estos servicios compartidos se filtran en el GRT, o un VRF de servicios compartidos y viceversa. Los mapas de ruta se pueden utilizar para ayudar a contener las tablas de ruteo a las subredes específicas de la estructura de acceso SD.

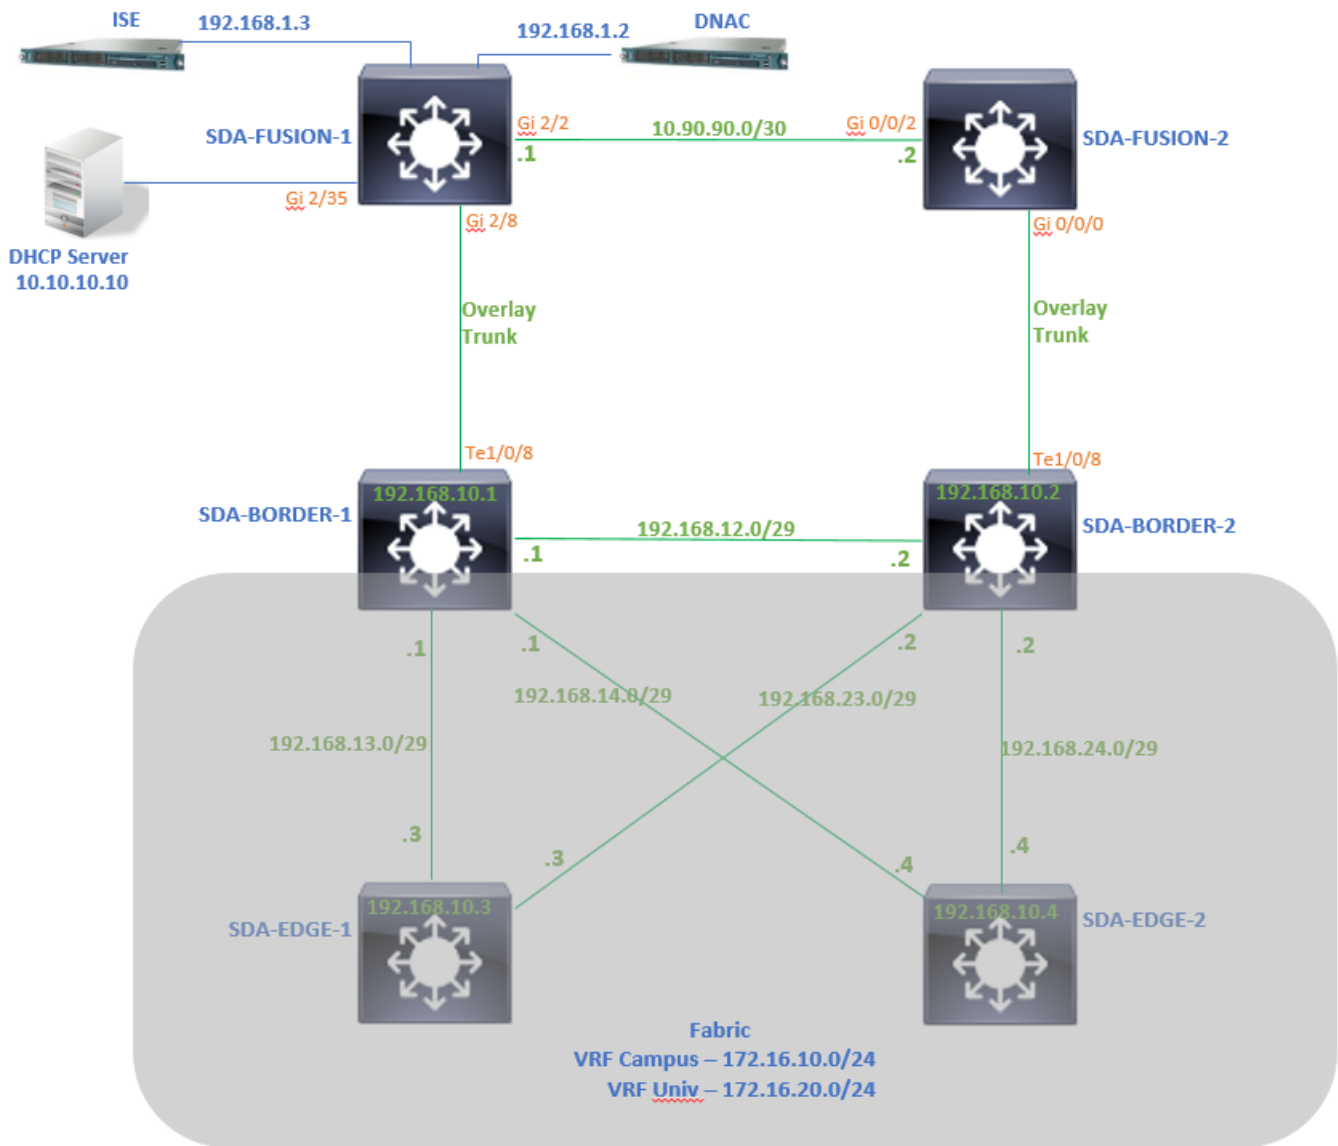
 Nota: Los nodos de borde de acceso SD no admiten rutas de resumen que se superponen con los grupos IP de acceso SD. Las rutas de resumen que se superponen con los grupos de IP se deben filtrar en los anuncios de routing de los dispositivos Fusion a los nodos de borde.

Configurar

Los detalles de configuración que se proporcionan aquí se relacionan con la topología de red que se muestra a continuación. Esta topología de red no es una topología recomendada para implementaciones. Se utiliza aquí únicamente para facilitar la presentación de los ejemplos de configuración proporcionados. Para ver los diseños de implementación recomendados, consulte la [Zona de diseño para la arquitectura de red digital de Cisco](#).

Diagrama de la red

La topología utilizada para este artículo consiste en dos routers de borde configurados como fronteras externas y dos routers de fusión con una conexión a cada router de borde respectivo.



Configuraciones

Paso 1. Configuración del enlace de entrega

En el paso de asignar a los dispositivos una función de router de borde mientras se agregan al fabric, se puede crear un enlace de transferencia. En la capa 2 es un enlace troncal conectado al router Fusion. Los siguientes pasos son necesarios:

1. Configure el número AS local para BGP. Este número de sistema autónomo (AS) se utiliza para configurar el proceso BGP en los routers de borde.
2. Agregue la interfaz en Tránsito. Esta interfaz es la conexión directa entre el borde y el router de fusión. (El 1/0/8 en Border en este ejemplo.)

SDA-Border1

Border to

- Rest of Company (Internal)
- Outside World (External)
- Anywhere (Internal & External)

Local Autonomous Number

65005



Select Ip Pool

x BGP (10.50.50.0/24)



Connected to the Internet

Transit

Add

ABC

+ Add Interface

External Interface

Interface

Number of VN

TenGigabitEthernet1/0/8

2

3. Configure el número de AS remoto. Este número AS se utiliza en los routers de borde para que las sentencias de vecino hacia el router de fusión configuren pares BGP externos (eBGP).
4. Seleccione todas las redes virtuales (VRF) para las que se requiere la fuga de VRF en el router de fusión.
5. Implemente la configuración de Cisco Catalyst Center en los dispositivos.

SDA-Border1

[< Back](#)

External Interface

* TenGigabitEthernet1/0/8

Remote AS Number

65004



This number is automatically derived from the selected Transit.
The selected autonomous system number will be used to automate IP routing between Border Node and remote peer.

Virtual Network

DEFAULT_VN

INFRA_VN

Univ

Campus

Siga los mismos pasos para el dispositivo SDA-Border-2.

Paso 2. Verificar las configuraciones introducidas en los routers de borde

Esta sección cubre la verificación de la configuración en los Routers de Borde relacionados con el protocolo BGP.

SDA-Border-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
  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
```

<<< SVI created for BGP Peering under VRF C

```
SDA-Border1#show run interface vlan 3006
```

```
!
interface Vlan3006
  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
```

<<< SVI created for BGP Peering under VRF U

```
SDA-Border1#show run | section bgp
```

```
router bgp 65005
  bgp router-id interface Loopback0
  bgp log-neighbor-changes
  bgp graceful-restart
  !
  address-family ipv4
  network 192.168.10.1 mask 255.255.255.255
  redistribute lisp metric 10
  exit-address-family
```

<<< Local AS Number from Cisco Catalyst Cent

```

!
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 Cam
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
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-Border-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
```

Paso 3. Configuración de Allowas-In en Routers de Borde


Debido a la fuga de VRF en el router Fusion, la familia de direcciones ipv4 para el campus VRF aprende la ruta originada por VRF Univ (172.16.20.0/24). Sin embargo, tanto el router de origen como el de aprendizaje tienen el mismo número AS BGP (65005). Para superar los mecanismos de prevención de loop BGP y aceptar/instalar las rutas en los routers de borde, allowas-in debe configurarse para los pares con el router de fusión:


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: El comando allowas-in debe utilizarse con precaución, ya que puede provocar bucles. Cuando utiliza un solo dispositivo Fusion con el que ambos Bordes se emparejan, el filtrado es necesario para asegurarse de que las rutas originadas localmente no sean aceptadas nuevamente en el AS desde el par Fusion - dentro de la misma VPN. Si esto sucede, se

 prefiere la trayectoria eBGP a la trayectoria originada localmente debido al peso máximo para las trayectorias eBGP.

Paso 4. Configurar routers de fusión

Esta sección ilustra la configuración manual de los routers Fusion.

SDA-Fusion-1

Configure el link hacia el Router de borde como un trunk para que coincida con la configuración de vlan en el Border-1:

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

Configure los VRF requeridos:

```
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
```

Configurar interfaces 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
```

Configure el peering BGP externo (eBGP) con SDA-Border-1:

```
router bgp 65004                                     <<< Remote AS from Cisco Catalyst Center
  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
```

Configure el peering BGP interno (iBGP) 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
  !
```

Anuncie la subred del servidor DHCP en la familia global de direcciones donde la IP del servidor DHCP es 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

Configure el link hacia el Router de borde. Si una interfaz en Fusion es L3 en lugar de trunk - configure subinterfaces:

```
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
```

Configure los VRFs correspondientes:

```
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
!
```

Configuración del Peering 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

```

Configuración del Peering 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

```

Paso 5. Configuración de la fuga de VRF en el router Fusion

La configuración para la fuga de VRF es idéntica para los routers de fusión SDA-Fusion-1 y SDA-Fusion-2.

En primer lugar, configure la fuga de VRF entre los dos VRF (Campus y Univ) y utilice la importación de destino de ruta:

```

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
!
```

A continuación, configure la fuga de ruta entre la tabla de routing global (GRT) a los VRF y desde los VRF a los GRT, utilice import ... map y 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 (e
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
 export ipv4 unicast map Campus_Map <<< Injecting VRF Campus into Global matching route-map Campus
exit-address-family
!
```

```

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

Verificación

Esta sección contiene los pasos de verificación para asegurarse de que la configuración anterior ha surtido efecto correctamente.

Paso 1. Verificación del Peering eBGP entre los Routers de Fusión y de Borde

SDA-Border-1 -----Peering-----SDA-Fusion-1

```
SDA-Border1#show ip bgp vpv4 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 vpv4 vrf Univ summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	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 vpv4 vrf Campus summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	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 vpv4 vrf Univ summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	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 vpv4 vrf Campus summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	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 vpv4 vrf Univ summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	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 vpv4 vrf Campus summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	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 vpv4 vrf Univ summary

Neighbor	V	AS	MsgRcvd	MsgSent	Tb1Ver	InQ	OutQ	Up/Down	State/PfxRcd
10.50.50.9	4	65005	17	17	9	0	0	00:11:33	1

Paso 2. Verifique el Peering iBGP entre ambos Routers 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

Paso 3. Verificar Prefijos en Tabla BGP y Tabla de Ruteo

SDA-Border-1

SDA-Border1#show ip bgp vpv4 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
*> 172.16.10.0/24	0.0.0.0	32768	i		<<< VRF Campus originat
*> 172.16.20.0/24	10.50.50.26	65535	65004	65005 i	<<< Prefix originated i

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 pre
B	172.16.10.0/24 [200/0], 20:32:45, Null0	<<< Null entry created by "aggregate-a
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 vpv4 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
*> 172.16.10.0/24	10.50.50.22	65535	65004	65005 i	<<< Prefix originated i
*> 172.16.20.0/24	0.0.0.0	32768	i		<<< VRF Univ originated

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 pre
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-a

SDA-Border-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
*> 172.16.10.0/24	0.0.0.0	32768			i	<<< VRF Campus originat
*> 172.16.20.0/24	10.50.50.2	65535		65004	65005 i	<<< Prefix originated i

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 pref
B	172.16.10.0/24	[200/0],	1w6d,	Null0	<<< Null entry created by "aggregate-ad
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
*> 172.16.10.0/24	10.50.50.10	65535		65004	65005 i	<<< Prefix originated i
*> 172.16.20.0/24	0.0.0.0	32768			i	<<< VRF Univ originated

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 pre
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-a

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 Glob
* i 172.16.10.0/24	10.50.50.1	0	100	0	65005 i	<<< Prefix imported from VR
*>	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 VR
*>	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 V
B	172.16.20.0	[20/0]	via 10.50.50.21 (Univ),	20:50:21	<<< Prefix imported from VRF

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 G
*> 172.16.10.0/24	10.50.50.25	0		0 65005	i <<< Prefix learnt from B
*> 172.16.20.0/24	10.50.50.21	0		0 65005	i <<< Prefix imported from

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 G
*> 172.16.10.0/24	10.50.50.25	0		0 65005	i <<< Prefix imported from
*> 172.16.20.0/24	10.50.50.21	0		0 65005	i <<< Prefix learnt from Bor

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
```

Configuración manual para redundancia de bordes

Para la redundancia entre los PETRs cuando falla un link externo de borde, para los bordes Externo y Externo+Interno, debe construir manualmente sesiones iBGP entre los dos Bordes para cada VNs. Además, en el caso del borde externo+interno donde BGP se importa en LISP y LISP se redistribuye de nuevo en BGP, las etiquetas son necesarias para evitar las importaciones de ruta de iBGP a LISP y, por lo tanto, evitar bucles potenciales.

SDA-Border-1

```
<#root>
```

```
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 Cisco Catalyst Ce
  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-f8894
!
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-f88949a
!

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

route-map DENY-Campus deny 5                                       <<< route-map pushed and used in route-import
  match ip address prefix-list Campus
!
route-map DENY-Campus deny 10
  match ip address prefix-list 13handoff-prefixes
!
route-map DENY-Campus deny 15
  match community 1                                                 <<< match on community-list 1 to deny iBGP prefixes
!
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 13handoff-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
set community 655370                                               <<< setting community/tag to 655370

!

```

SDA-Border-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-f9c6ebda80
!
```

```
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-f9c6ebda80
!
```

```
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 13handoff-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 13handoff-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  
!
```

Simplificación de la configuración de fusión con el uso de plantillas

Esta sección contiene ejemplos de configuración de la plantilla Fusion Template para ayudar a simplificar la configuración.

A continuación se indican las variables que se deben definir en función del diseño de implementación. En este ejemplo, las configuraciones y las VN se basan en la topología anterior que tenía dos VN, Campus y Univ.

Definición de variable

```
interface_Fusion1: GigabitEthernet2/8  
interface_Fusion2: GigabitEthernet0/0/0  
  
Global_prefixes = 10.10.10.8/30  
  
FUSION_BGP_AS = 65004  
BORDER_BGP_AS = 65005
```

Para 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
```


Para 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
Fusion2_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
```

Ejemplo de plantilla

Fusión 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
!

```

Fusión 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
```

Acerca de esta traducción

Cisco ha traducido este documento combinando la traducción automática y los recursos humanos a fin de ofrecer a nuestros usuarios en todo el mundo contenido en su propio idioma.

Tenga en cuenta que incluso la mejor traducción automática podría no ser tan precisa como la proporcionada por un traductor profesional.

Cisco Systems, Inc. no asume ninguna responsabilidad por la precisión de estas traducciones y recomienda remitirse siempre al documento original escrito en inglés (insertar vínculo URL).