Konfigurieren des Fusion-Routers in SDA

Inhalt

Einleitung

In diesem Dokument wird die Konfiguration von Fusion Routern in einer Cisco SDA-Lösung (Software-Defined Access) beschrieben.

Voraussetzungen

Anforderungen

Es gibt keine spezifischen Anforderungen für dieses Dokument.

> Hinweis: Die Einrichtung ist gemäß den unterstützten Geräten erforderlich. Link zu den Versionshinweisen

Verwendete Komponenten

Die Informationen in diesem Dokument basieren auf folgenden Software- und Hardware-Versionen:

- DNAC - Version 1.2.1
- Rand und Rand - Cat3k Cisco Switch
- Fusion - Cisco Router mit Unterstützung für Inter-VRF Leaking

Die Informationen in diesem Dokument beziehen sich auf Geräte in einer speziell eingerichteten Testumgebung. Alle Geräte, die in diesem Dokument benutzt wurden, begannen mit einer gelöschten (Nichterfüllungs) Konfiguration. Wenn Ihr Netzwerk in Betrieb ist, stellen Sie sicher, dass Sie die möglichen Auswirkungen aller Befehle verstehen.

Hintergrundinformationen

Bei der Cisco SD-Access-Lösung werden die Geräte vom Cisco DNA Center verwaltet und konfiguriert. Im Allgemeinen können alle Komponenten der SD-Access-Fabric vom Cisco DNA Center konfiguriert und verwaltet werden. Das Fusion-Gerät befindet sich jedoch außerhalb des Fabric und wird daher manuell konfiguriert. Border Automation (Grenzautomatisierung) ist eine Funktion im Cisco DNA Center, die die Grenzkonfiguration für die Übergabe von VRFs an die Fusion-Geräte automatisieren kann.

Gelegentlich ist Border Automation aus Gründen, die typischerweise mit der aktuellen Konfiguration zusammenhängen, nicht geeignet, sodass die Übergabe vom Border an das Fusion-Gerät auch von Hand konfiguriert werden kann. Ein Verständnis der verwendeten Konfiguration hilft dabei, wichtige Details über die optimale Konfiguration und den Betrieb des Gesamtsystems zu illustrieren.

Funktionalität eines Fusionsgerätes in einer DNA SD-Zugangslösung

Ein Fusion-Gerät ermöglicht Virtual Routing and Forwarding (VRF)-Lecks über SD-Access Fabric-Domänen und Host-Verbindungen zu gemeinsam genutzten Services wie DHCP, DNS, NTP, ISE, Cisco DNA Center, Wireless LAN Controllern (WLC) und dergleichen. Diese Rolle kann zwar von anderen Geräten als Routern übernommen werden, doch liegt der Schwerpunkt dieses Dokuments auf Routern als Fusion-Geräten.

Wie bereits erwähnt, müssen die Shared Services allen virtuellen Netzwerken (VPNs) auf dem Campus zur Verfügung gestellt werden. Dies wird durch die Erstellung von Border Gateway Protocol (BGP)-Peerings von den Border Routern zu den Fusion Routern erreicht. Auf dem Fusion Router werden die Subnetze der Fabric-VRFs, die Zugriff auf diese gemeinsam genutzten Services benötigen, in die GRT (Shared Services VRF) geleakt, und umgekehrt. Routenzuordnungen können dabei helfen, Routing-Tabellen für Subnetze zu enthalten, die spezifisch für die SD-Access-Fabric sind.

> Hinweis: SD-Access-Grenzknoten unterstützen keine Übersichtsrouten, die sich mit SD-Access-IP-Pools überschneiden. Zusammenfassende Routen, die sich mit IP-Pools überschneiden, müssen beim Routing von Meldungen von Fusion-Geräten zu Grenzknoten gefiltert werden.

Konfigurieren

Die hier angegebenen Konfigurationsdetails beziehen sich auf die als Nächstes gezeigte Netzwerktopologie. Diese Netzwerktopologie ist keine empfohlene Topologie für Bereitstellungen. Es dient hier lediglich dazu, die Darstellung der bereitgestellten Konfigurationsbeispiele zu erleichtern. Die empfohlenen Bereitstellungsdesigns finden Sie in der Design Zone für die Cisco Digital Network Architecture.

Netzwerkdiagramm

Die für diesen Artikel verwendete Topologie besteht aus zwei Border Routern, die beide als External Borders konfiguriert sind, und zwei Fusion Routern mit einer Verbindung zu jedem jeweiligen Border Router.

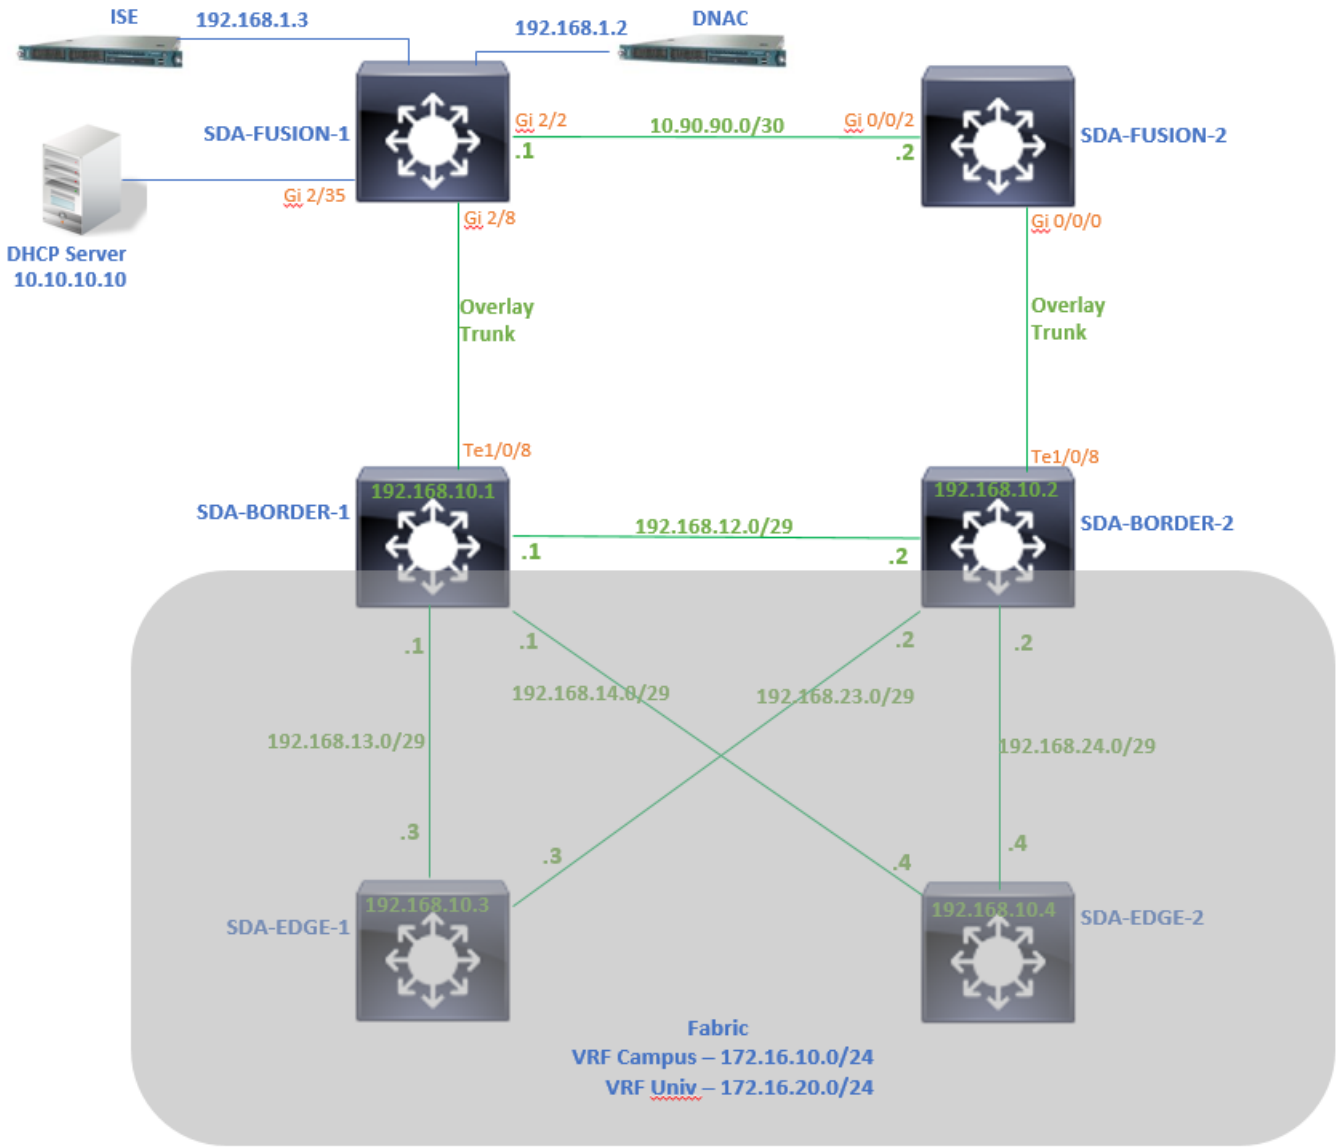

Konfigurationen

Schritt 1: Konfigurieren Sie den Übergabe-Link von DNAC aus.

Wenn Sie Geräten beim Hinzufügen zum Fabric die Rolle eines Border Routers zuweisen, kann

ein Handoff-Link erstellt werden. Auf Layer 2 ist ein Trunk-Link mit dem Fusion-Router verbunden. Die nächsten Schritte sind erforderlich:

1. Konfigurieren der lokalen AS-Nummer für BGP Diese AS-Nummer (Autonomous System) wird zur Konfiguration des BGP-Prozesses auf den Border Routern verwendet.

2. Schnittstelle unter Transit hinzufügen. Diese Schnittstelle ist die direkte Verbindung zwischen Border und Fusion Router. (In diesem Beispiel ist 1/0/8 on Border zu verwenden).

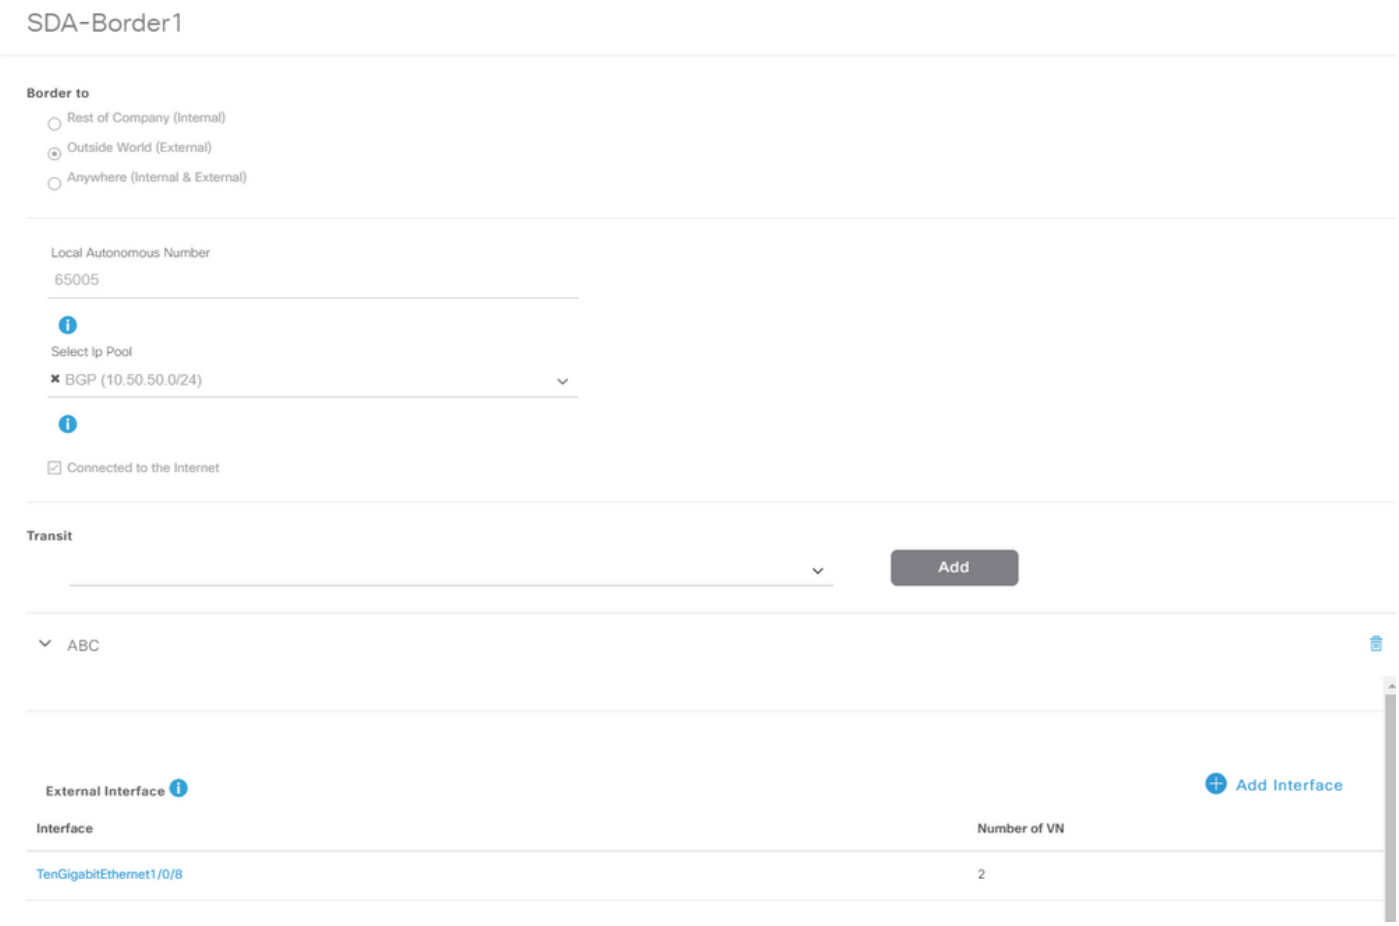

3. Konfigurieren der Remote-AS-Nummer Diese AS-Nummer wird auf Border Routern für Nachbaranweisungen zum Fusion Router verwendet, um eBGP-Peers (External BGP) zu konfigurieren.

4. Wählen Sie alle virtuellen Netzwerke (VRFs) aus, für die VRF-Leaking auf dem Fusion Router erforderlich ist.

5. Konfiguration von DNAC auf Geräte bereitstellen.

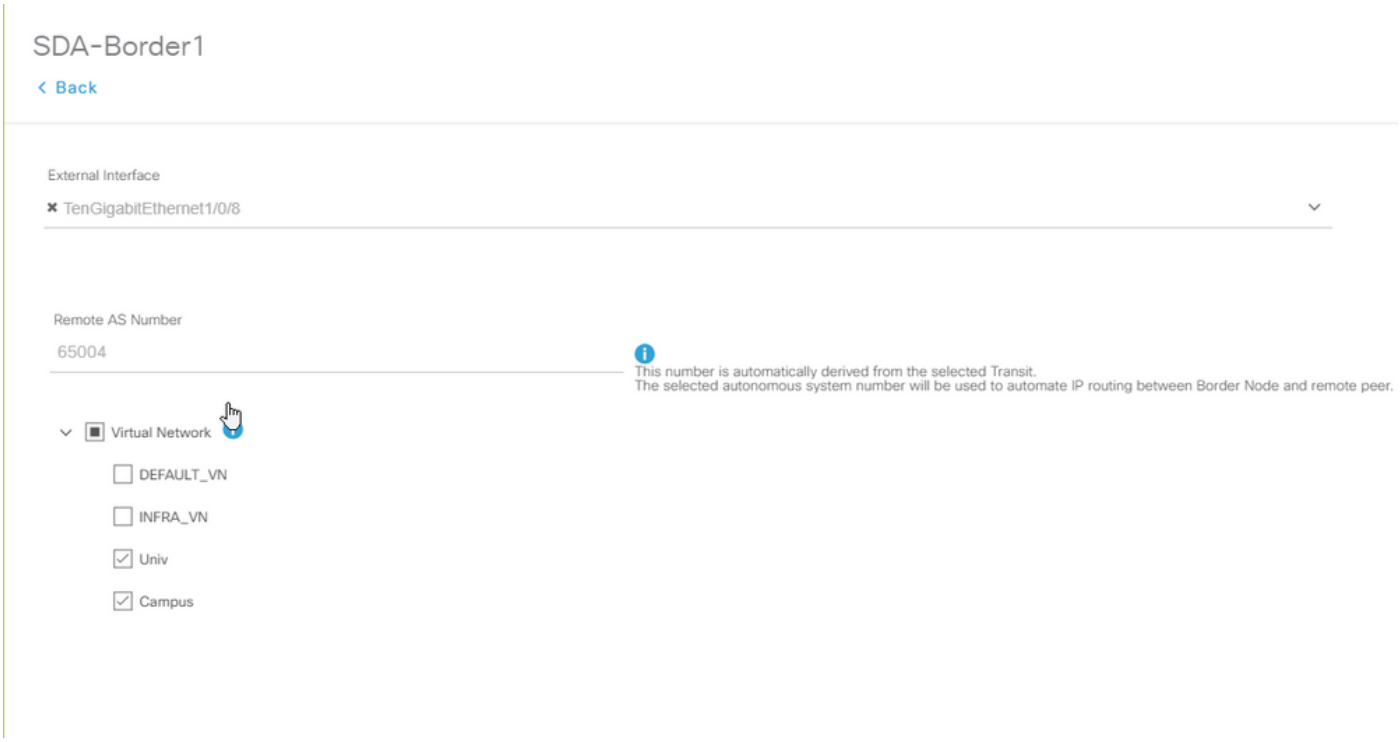

Führen Sie die gleichen Schritte für das SDA-Border-2-Gerät aus.

Schritt 2: Überprüfung der von DNAC an Border Router weitergeleiteten Konfigurationen

In diesem Abschnitt wird die Verifizierung der Konfiguration von Border Routern in Bezug auf das BGP-Protokoll behandelt.

SDA-Rahmen-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-Rahmen-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
```

Schritt 3: Zulassen auf Border Routern konfigurieren

Aufgrund der VRF-Leaking auf dem Fusion Router erkennt die address-family-IPv4 für VRF-Campus die von VRF Univ (172.16.20.0/24) generierte Route. Der Ausgangs- und der lernende Router haben jedoch dieselbe BGP-AS-Nummer (65005). Um die BGP-Schleifenvermeidungsmechanismen zu umgehen und die Routen auf Border Routern zu akzeptieren/zu installieren, muss **allowas-in** für die Peerings mit dem Fusion Router konfiguriert werden:

```
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)#
```

> **Hinweis**: Der Befehl **allowas-in** muss mit Vorsicht verwendet werden, da er Schleifen verursachen kann. Wenn Sie nur ein Fusion-Gerät verwenden, mit dem beide Borders-Peers arbeiten, ist eine Filterung erforderlich, um sicherzustellen, dass lokal generierte Routen nicht vom Fusion-Peer zurück in das AS akzeptiert werden - innerhalb derselben VN. In diesem Fall wird der eBGP-Pfad aufgrund der maximalen Gewichtung der eBGP-Pfade dem lokal erstellten Pfad vorgezogen.

Schritt 4: Fusion-Router konfigurieren

In diesem Abschnitt wird die manuelle Konfiguration für die Fusion-Router erläutert.

SDA-Fusion-1

Konfigurieren Sie den Link zum Border Router als Trunk, der der VLAN-Konfiguration an Border-1 entspricht:

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

Konfigurieren der erforderlichen VRFs:

```
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
```

Konfigurieren von SVI-Schnittstellen:

```
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
```

Konfigurieren Sie externes BGP (eBGP)-Peering mit 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
```

Konfigurieren Sie internes BGP (iBGP)-Peering mit 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
 !
```

Geben Sie das DHCP-Server-Subnetz unter der globalen Adressfamilie an, wobei die IP-Adresse des DHCP-Servers 10.10.10.10 lautet:

```
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

Konfigurieren Sie den Link zu Border Router. Wenn eine Schnittstelle in Fusion L3 anstatt Trunk ist, konfigurieren Sie Subschnittstellen:

```
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
```

Konfigurieren Sie die entsprechenden VRFs:

```
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
!
```

Konfigurieren Sie eBGP-Peering mit 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
```

Konfigurieren Sie iBGP-Peering mit 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
```

Schritt 5: VRF-Leaking auf Fusion Router konfigurieren

Die Konfiguration für VRF-Leaking ist für die Fusion-Router SDA-Fusion-1 und SDA-Fusion-2 identisch.

Konfigurieren Sie zunächst das VRF-Leaking zwischen den beiden VRF-Instanzen (Campus und Univ), und verwenden Sie **route-target import**:

```
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
!
```

Konfigurieren Sie dann das Route Leaking zwischen der Global Routing Table (GRT) zu den VRFs, und verwenden Sie von den VRFs zur GRT **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
!
```

Überprüfung

In diesem Abschnitt werden die erforderlichen Verifizierungsschritte beschrieben, um sicherzustellen, dass die vorherige Konfiguration ordnungsgemäß durchgeführt wurde.

Schritt 1: Überprüfung des eBGP-Peering zwischen Fusion und Border Router

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
```

Schritt 2: Überprüfen des iBGP-Peering zwischen beiden Fusion-Routern

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
```

Schritt 3: Überprüfen von Präfixen in der BGP-Tabelle und Routing-Tabelle

SDA-Rahmen-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-Rahmen-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
```

Manuelle Konfiguration für Grenzredundanz

Für Redundanz zwischen den PETRs bei Ausfall einer externen Grenzverbindung sowie für externe und interne Grenzen müssen Sie für jedes VPNs manuell iBGP-Sitzungen zwischen den beiden Grenzen erstellen. Wenn BGP in LISP importiert und wieder in das BGP zurückverteilt wird, sind darüber hinaus Tags erforderlich, um iBGP- und LISP-Routing-Importe zu verhindern und somit potenzielle Schleifen zu vermeiden.

SDA-Rahmen-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-Rahmen-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
!
```

Vereinfachte Fusion-Konfiguration mithilfe von Vorlagen

Dieser Abschnitt enthält Beispiele für Fusion-Vorlagen zur Vereinfachung der Konfiguration.

Im Folgenden werden die Variablen aufgeführt, die basierend auf dem Bereitstellungsdesign definiert werden müssen. In diesem Beispiel basieren die Konfigurationen und VPNs auf der vorherigen Topologie mit zwei VPNs, "Campus" und "Univ".

Variablendefinition

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

Global_prefixes = 10.10.10.8/30

FUSION_BGP_AS = 65004
BORDER_BGP_AS = 65005
```

Für 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
```

Für 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
```

Vorlagenbeispiel

Kernfusion 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
!
```

Kernfusion 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

Informationen zu dieser Übersetzung

Cisco hat dieses Dokument maschinell übersetzen und von einem menschlichen Übersetzer editieren und korrigieren lassen, um unseren Benutzern auf der ganzen Welt Support-Inhalte in ihrer eigenen Sprache zu bieten. Bitte beachten Sie, dass selbst die beste maschinelle Übersetzung nicht so genau ist wie eine von einem professionellen Übersetzer angefertigte. Cisco Systems, Inc. übernimmt keine Haftung für die Richtigkeit dieser Übersetzungen und empfiehlt, immer das englische Originaldokument (siehe bereitgestellter Link) heranzuziehen.