

Ejemplo de configuración del Hub dual de FlexVPN HA

Contenido

[Introducción](#)

[prerrequisitos](#)

[Requisitos](#)

[Componentes usados](#)

[Antecedentes](#)

[Configurar](#)

[Diagrama de la red](#)

[Escenario operativo regular](#)

[Spoke al spoke \(acceso directo\)](#)

[Tablas de ruteo y salidas para el escenario operativo regular](#)

[Escenario de falla del HUB1](#)

[Configuraciones](#)

[Configuración R1-HUB](#)

[Configuración R2-HUB2](#)

[Configuración R3-SPOKE1](#)

[Configuración R4-SPOKE2](#)

[Configuración R5-AGGR1](#)

[Configuración R6-AGGR2](#)

[Configuración R7-HOST \(simulación del HOST en esa red\)](#)

[Notas de configuración importantes](#)

[Verificación](#)

[Troubleshooting](#)

[Información Relacionada](#)

Introducción

Este documento describe cómo configurar un diseño de la redundancia completa para las oficinas remotas que conectan con un centro de datos vía el VPN basado en Ipsec sobre un media de la red insegura, tal como Internet.

Prerequisites

Requisitos

No hay requisitos específicos para este documento.

Componentes usados

La información en este documento se basa en estos componentes de la tecnología:

- [Border Gateway Protocol \(BGP\)](#) como el Routing Protocol dentro del centro de datos y entre el spokes y el Hubs en el recubrimiento VPN.
- [Detección bidireccional de la expedición](#) (BFD) como mecanismo que detecta abajo de los links (router abajo) que ejecutado dentro del centro de datos solamente (no sobre los túneles del recubrimiento).
- [FlexVPN del Cisco IOS](#) entre el Hubs y el spokes, con las capacidades del spoke al spoke habilitadas vía la transferencia del atajo.
- [Generic Routing Encapsulation \(GRE\) que hace un túnel](#) entre dos Hubs para habilitar la comunicación del spoke al spoke, incluso cuando el spokes está conectado con diverso Hubs.
- [Rastreo de objetos aumentado](#) y Static rutas atados a los objetos seguidos.

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 la red está funcionando, asegúrese de haber comprendido el impacto que puede tener cualquier comando.

Antecedentes

Cuando usted diseña las soluciones de acceso remoto para el centro de datos, la Alta disponibilidad (HA) es a menudo un requisito dominante para las aplicaciones de usuario de la misión crítica.

La solución que se presenta en este documento permite la detección y la recuperación rápidas de los escenarios de falla en cuál del Hubs VPN-terminal va abajo de debido a una recarga, a una actualización, o a los problemas de alimentación eléctrica. Todo el Routers de las oficinas remotas (spokes) entonces utiliza el otro concentrador operativo inmediatamente al detectar tal error.

Aquí están las ventajas de este diseño:

- Recuperación rápida de la red de un escenario del concentrador-abajo VPN
- Sincronizaciones stateful no complicadas (tales como asociaciones de seguridad IPsec (SA), Internet Security Association and Key Management Protocol (ISAKMP) SA, y Crypto-encaminamiento) entre el Hubs VPN
- Ningunos problemas debido de la anti-respuesta a los retardos en la sincronización del

número de secuencia del Encapsulating Security Payload (ESP) con el IPSec HA stateful

- El Hubs VPN puede utilizar el diverso soporte físico o software basado IOS/IOS-XE de Cisco
- Opciones flexibles de la implementación del balanceo de carga con el BGP como el Routing Protocol que se ejecuta en el recubrimiento VPN
- Encaminamiento clara y legible en todos los dispositivos sin los mecanismos ocultos que se ejecutan en el fondo
- Conectividad directa del spoke al spoke
- Todas las ventajas de [FlexVPN](#), incluir el Calidad de Servicio (QoS) de la integración y del por-túnel del Authentication, Authorization, and Accounting (AAA)

Configurar


Esta sección proporciona los ejemplos de escenario y describe cómo configurar un diseño de la redundancia completa para las oficinas remotas que conectan con el centro de datos vía el VPN basado en Ipsec sobre un media de la red insegura.

Note: Use la [Command Lookup Tool](#) ([clientes registrados solamente](#)) para obtener más información sobre los comandos usados en esta sección.

Diagrama de la red

Ésta es la topología de red que se utiliza en este documento:

Nominal/Default Status – Traffic flowing through HUB1



Note: Todo el Routers que se utiliza en esta topología funciona con la versión deL Cisco IOS 15.2(4)M1, y la nube de Internet utiliza un Esquema de dirección de 172.16.0.0/24.

Escenario operativo regular

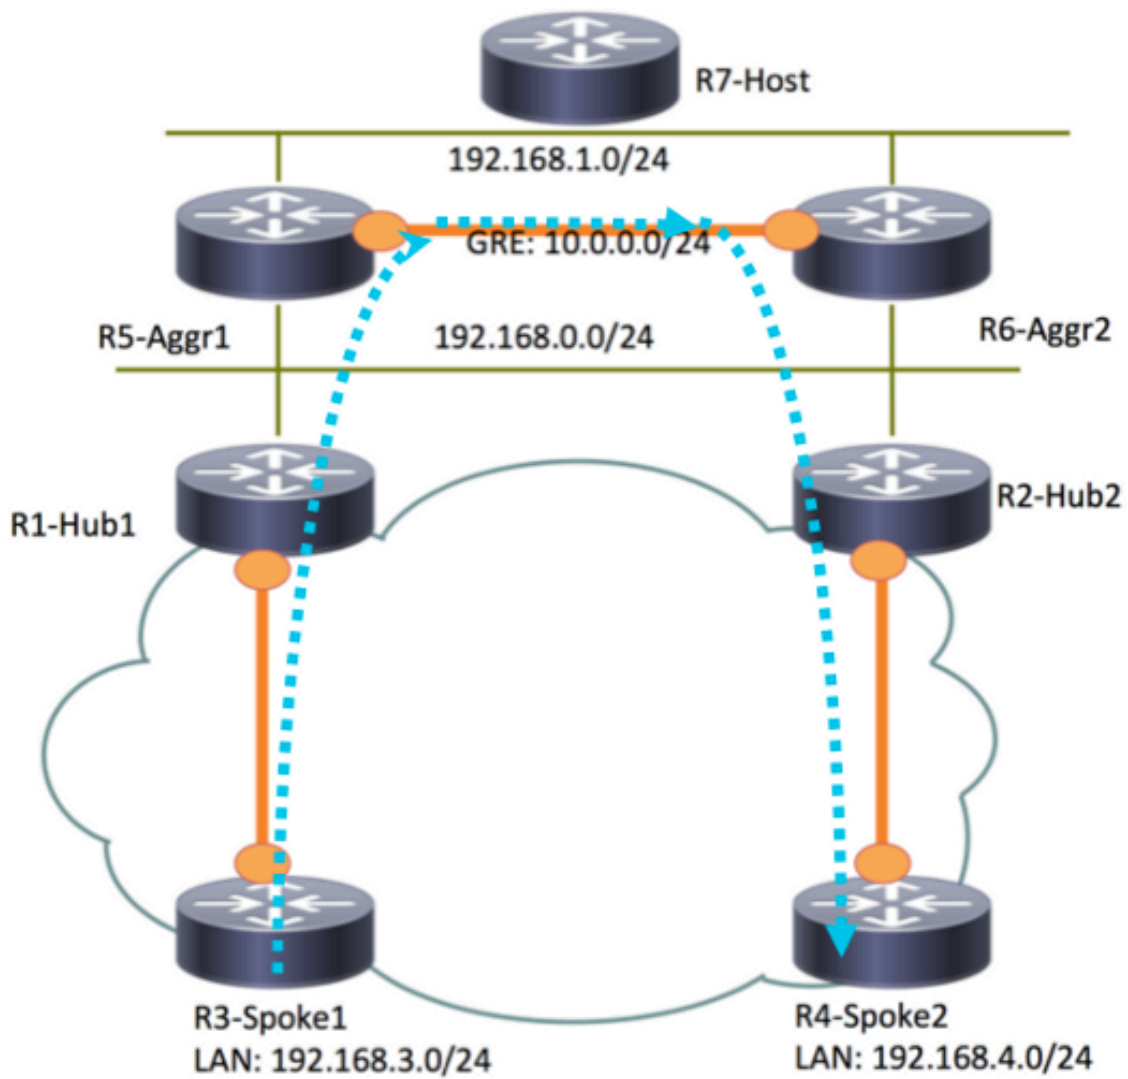
En un escenario operativo normal, cuando todo el Routers es ascendente y operativo, todos los routers radiales rutean todo el tráfico a través del concentrador predeterminado (R1-HUB1). Se alcanza esta preferencia de la encaminamiento cuando la preferencia local del valor por defecto BGP se fija a 200 (refiera a las secciones que siguen para los detalles). Esto se puede ajustar sobre la base de los requisitos del despliegue, tales como balanceo de carga del tráfico.

Spoke al spoke (acceso directo)


Si R3-Spoke1 inicia una conexión a R4-Spoke2, un túnel dinámico del spoke al spoke se crea con la configuración de la transferencia del atajo.

Tip: Para más detalles, refiera al [FlexVPN que configura habló a la guía de configuración radial](#).

Si R3-Spoke1 está conectado solamente con R1-HUB1, y R4-Spoke2 está conectado solamente con R2-HUB2, una conexión directa del spoke al spoke se puede todavía alcanzar con el túnel GRE de punto a punto que se ejecuta entre el Hubs. En este caso, el trayecto del tráfico inicial entre R3-Spoke1 y R4-Spoke2 aparece similares a esto:



Puesto que R1-Hub1 recibe el paquete en la interfaz de acceso virtual, que tiene el mismo ID de la red del Next Hop Resolution Protocol (NHRP) que éste en el túnel GRE, la indicación del tráfico se envía hacia el R3-Spoke1. Esto acciona la creación del túnel dinámico del spoke al spoke:



Tablas de ruteo y salidas para el escenario operativo regular

Aquí está la tabla de ruteo R1-HUB1 en un escenario operativo regular:

```
R1-HUB1# show ip route
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override
```

```
Gateway of last resort is not set
```

```
10.0.0.0/8 is variably subnetted, 11 subnets, 3 masks
S    10.0.0.0/8 is directly connected, Null0
C    10.0.0.0/24 is directly connected, Tunnel0
L    10.0.0.1/32 is directly connected, Tunnel0
C    10.0.1.1/32 is directly connected, Loopback0
```

```

S      10.0.1.2/32 is directly connected, Virtual-Access1
S      10.0.1.3/32 is directly connected, Virtual-Access2
B      10.0.2.1/32 [200/0] via 10.0.0.2, 00:05:40
B      10.0.2.3/32 [200/0] via 10.0.0.2, 00:05:40
B      10.0.2.4/32 [200/0] via 10.0.0.2, 00:05:40
B      10.0.5.1/32 [200/0] via 192.168.0.5, 00:05:40
B      10.0.6.1/32 [200/0] via 192.168.0.6, 00:05:40
      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C      172.16.0.0/24 is directly connected, Ethernet0/0
L      172.16.0.1/32 is directly connected, Ethernet0/0
S      192.168.0.0/16 is directly connected, Null0
      192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.0.0/24 is directly connected, Ethernet0/2
L      192.168.0.1/32 is directly connected, Ethernet0/2
B      192.168.1.0/24 [200/0] via 192.168.0.5, 00:05:40
B      192.168.3.0/24 [200/0] via 10.0.1.4, 00:05:24
B      192.168.4.0/24 [200/0] via 10.0.1.5, 00:05:33

```

Aquí está la tabla de ruteo R3-SPOKE1 en un escenario operativo regular después de que el túnel del spoke al spoke con R4-SPOKE2 se cree:

```
R3-SPOKE1# show ip route
```

```

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

```

```
Gateway of last resort is not set
```

```

      10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
B      10.0.0.0/8 [200/0] via 10.0.1.1, 00:06:27
H      10.0.0.1/32 is directly connected, 00:06:38, Tunnell
S %    10.0.1.1/32 is directly connected, Tunnel0
C      10.0.1.3/32 is directly connected, Tunnel0
H      10.0.1.4/32 is directly connected, 00:01:30, Virtual-Access1
S      10.0.2.1/32 is directly connected, Tunnell
C      10.0.2.3/32 is directly connected, Tunnell
H      10.0.2.4/32 [250/1] via 10.0.2.3, 00:01:30, Virtual-Access1
      172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C      172.16.0.0/24 is directly connected, Ethernet0/0
L      172.16.0.3/32 is directly connected, Ethernet0/0
B      192.168.0.0/16 [200/0] via 10.0.1.1, 00:06:27
      192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C      192.168.3.0/24 is directly connected, Ethernet0/1
L      192.168.3.3/32 is directly connected, Ethernet0/1
      192.168.4.0/32 is subnetted, 1 subnets
H      192.168.4.4 [250/1] via 10.0.1.3, 00:01:30, Virtual-Access1

```

En R3-Spoke1, la tabla BGP tiene dos entradas para la red 192.168.0.0/16 con diversas preferencias locales (se prefiere R1-Hub1):

```
R3-SPOKE1#show ip bgp 192.168.0.0/16
```

```

BGP routing table entry for 192.168.0.0/16, version 8
Paths: (2 available, best #2, table default)
      Not advertised to any peer
      Refresh Epoch 1
      Local

```

10.0.2.1 from 10.0.2.1 (10.0.2.1)

Origin incomplete, metric 0, localpref 100, valid, internal

rx pathid: 0, tx pathid: 0

Refresh Epoch 1

Local

10.0.1.1 from 10.0.1.1 (10.0.1.1)

Origin incomplete, metric 0, localpref 200, valid, internal, best

rx pathid: 0, tx pathid: 0x0

Aquí está la tabla de ruteo R5-AGGR1 en un escenario operativo regular:

R5-LAN1#show ip route

```
10.0.0.0/8 is variably subnetted, 10 subnets, 3 masks
B    10.0.0.0/8 [200/0] via 192.168.0.1, 00:07:22
B    10.0.0.0/24 [200/0] via 192.168.0.1, 00:07:22
B    10.0.1.1/32 [200/0] via 192.168.0.1, 00:07:22
B    10.0.1.3/32 [200/0] via 192.168.0.1, 00:07:17
B    10.0.1.4/32 [200/0] via 192.168.0.1, 00:07:16
B    10.0.2.1/32 [200/0] via 192.168.0.2, 15:44:13
B    10.0.2.3/32 [200/0] via 192.168.0.2, 15:44:13
B    10.0.2.4/32 [200/0] via 192.168.0.2, 15:44:13
C    10.0.5.1/32 is directly connected, Loopback0
B    10.0.6.1/32 [200/0] via 192.168.0.6, 00:07:22
172.16.0.0/24 is subnetted, 1 subnets
B    172.16.0.0 [200/0] via 192.168.0.1, 00:07:22
B    192.168.0.0/16 [200/0] via 192.168.0.1, 00:07:22
192.168.0.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.0.0/24 is directly connected, Ethernet0/0
L    192.168.0.5/32 is directly connected, Ethernet0/0
192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/24 is directly connected, Ethernet0/1
L    192.168.1.5/32 is directly connected, Ethernet0/1
B    192.168.3.0/24 [200/0] via 10.0.1.3, 00:07:06
B    192.168.4.0/24 [200/0] via 10.0.1.4, 00:07:15
```

Aquí está la tabla de ruteo R7-HOST en un escenario operativo regular:

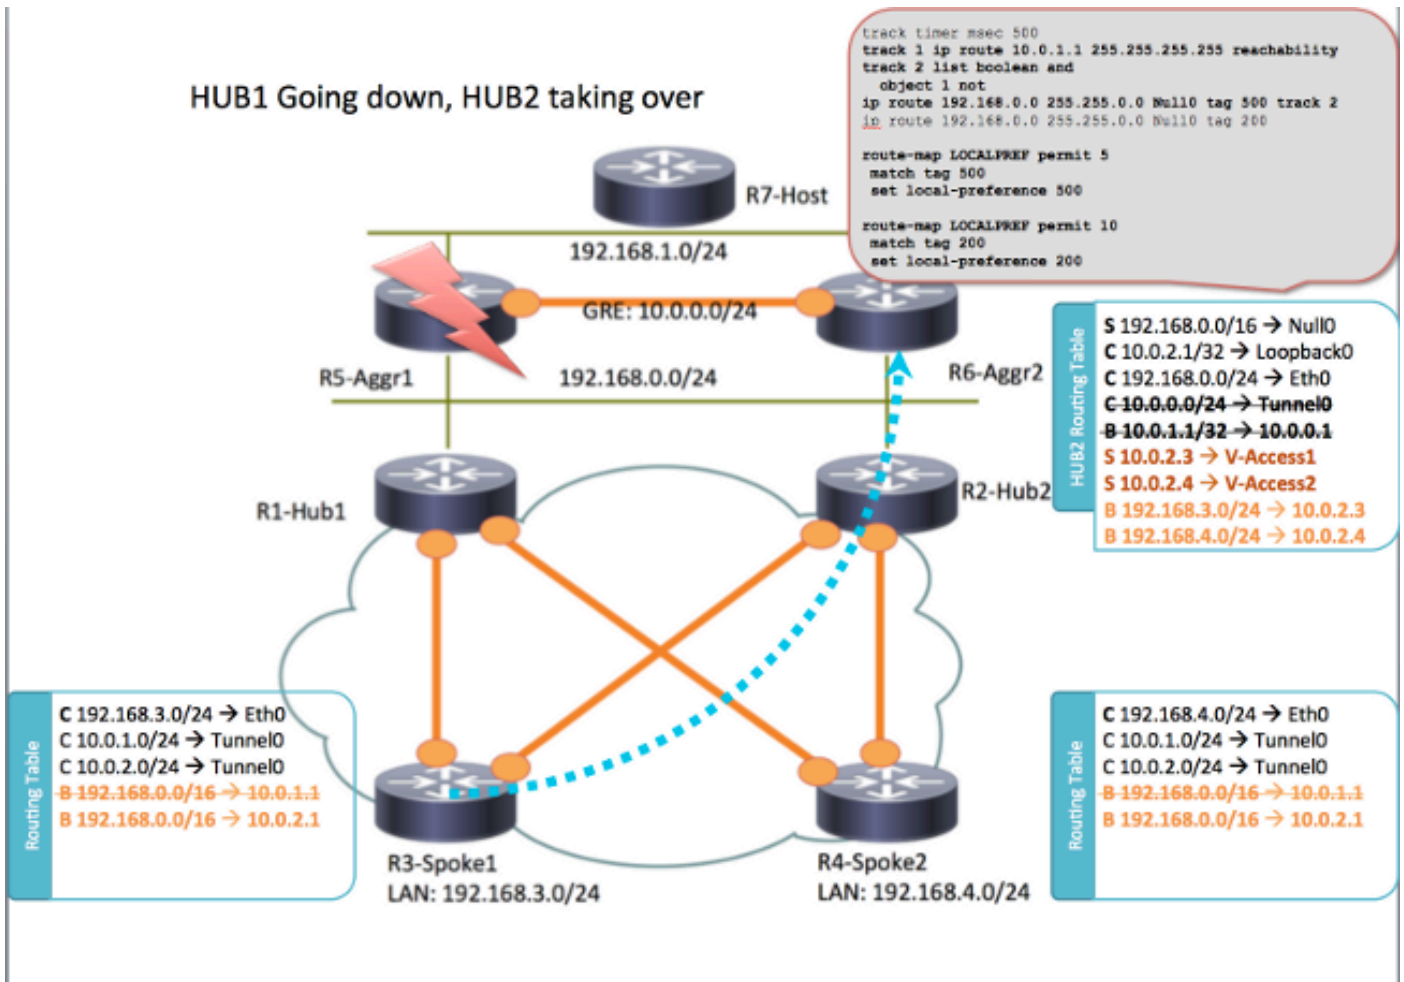
R7-HOST#show ip route

```
S*   0.0.0.0/0 [1/0] via 192.168.1.254
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/24 is directly connected, Ethernet0/0
L    192.168.1.7/32 is directly connected, Ethernet0/0
```

Escenario de falla del HUB1

Aquí está abajo un escenario R1-HUB1 (debido a las acciones tales como interrupciones de la alimentación eléctrica o una actualización):

HUB1 Going down, HUB2 taking over



En este escenario, esta Secuencia de eventos ocurre:

1. El BFD en R2-HUB2 y en el Routers R5-AGGR1 y R6-AGGR2 del agregado LAN detecta el estado de inactividad de R1-HUB1. Como consecuencia, la vecindad BGP va inmediatamente abajo.
2. La detección del objeto de la pista para R2-HUB2 que detecte la presencia del loopback R1-HUB1 va abajo (la pista 1 en el ejemplo de configuración).
3. Esto trágicamente siguió el objeto acciona otra pista para subir (lógico NO). En este ejemplo, la pista 2 sube siempre que vaya la pista 1 abajo.
4. Esto acciona IP estático una entrada de ruteo que se agregará a la tabla de ruteo debido a un valor que sea más bajo que la distancia administrativa predeterminedada. Aquí está la configuración pertinente:

```

! Routes added when second HUB is down
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
    
```

```

! Default static routes are with Tag 200 and admin distance of 150
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
    
```

5. R2-HUB2 redistribuye estas Static rutas con una preferencia local BGP que sea más grande que el valor que se fija para R1-HUB1. En este ejemplo, una preferencia local de **500** se utiliza en el escenario de falla, en vez de los **200** que es fijado por R1-HUB1:

```

route-map LOCALPREF permit 5
  match tag 500
  set local-preference 500
!
route-map LOCALPREF permit 10
  match tag 200
  set local-preference 200
!

```

En R3-Spoke1, usted puede ver esto en las salidas BGP. Observe que todavía existe la entrada al r1, pero no se utiliza:

```

R3-SPOKE1#show ip bgp 192.168.0.0/16
BGP routing table entry for 192.168.0.0/16, version 10
Paths: (2 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.0.2.1 from 10.0.2.1 (10.0.2.1)
      Origin incomplete, metric 0, localpref 500, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
  Refresh Epoch 1
  Local
    10.0.1.1 from 10.0.1.1 (10.0.1.1)
      Origin incomplete, metric 0, localpref 200, valid, internal
      rx pathid: 0, tx pathid: 0

```

6. En este momento, ambos spokes (R3-Spoke1 y R4-Spoke2) comienza a enviar el tráfico a R2-HUB2. Todos estos pasos deben ocurrir dentro del segundo. Aquí está la tabla de ruteo en el spoke 3:

```

R3-SPOKE1#show ip route
  10.0.0.0/8 is variably subnetted, 5 subnets, 2 masks
B   10.0.0.0/8 [200/0] via 10.0.2.1, 00:00:01
S   10.0.1.1/32 is directly connected, Tunnel0
C   10.0.1.3/32 is directly connected, Tunnel0
S   10.0.2.1/32 is directly connected, Tunnel1
C   10.0.2.3/32 is directly connected, Tunnel1
  172.16.0.0/16 is variably subnetted, 2 subnets, 2 masks
C   172.16.0.0/24 is directly connected, Ethernet0/0
L   172.16.0.3/32 is directly connected, Ethernet0/0
B   192.168.0.0/16 [200/0] via 10.0.2.1, 00:00:01
  192.168.3.0/24 is variably subnetted, 2 subnets, 2 masks
C   192.168.3.0/24 is directly connected, Ethernet0/1
L   192.168.3.3/32 is directly connected, Ethernet0/1

```

7. Sesiones de BGP posteriores entre el spokes y el R1-HUB1 van abajo, y el Dead Peer Detection (DPD) quita los túneles IPsec que se terminan en R1-HUB1. Sin embargo, esto no afecta el reenvío de tráfico, puesto que R2-HUB2 se utiliza ya como el gateway túnel-terminal principal:

```

R3-SPOKE1#show ip bgp 192.168.0.0/16
BGP routing table entry for 192.168.0.0/16, version 10
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.0.2.1 from 10.0.2.1 (10.0.2.1)
      Origin incomplete, metric 0, localpref 500, valid, internal, best
      rx pathid: 0, tx pathid: 0x0

```

Configuraciones

Esta sección proporciona las configuraciones de muestra para el Hubs y el spokes que se utilizan en esta topología.

Configuración R1-HUB

```
version 15.4
!
hostname R1-HUB1
!
aaa new-model
!
aaa authorization network default local
!
aaa session-id common
!
! setting track timers to the lowest possible (the lower this value is
! the faster router will react
track timer ip route msec 500
!
! Monitoring of HUB2's loopback present in routing table
! If it is present it will mean that HUB2 is alive
track 1 ip route 10.0.2.1 255.255.255.255 reachability
!
! Monitoring of loopback of R5-AGGR-1
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
! Monitoring of loopback of R6-AGGR-2
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
! Track 2 should be UP only when HUB2 is not available and both AGGRE routers are up
!
track 2 list boolean and
  object 1 not
  object 3
  object 4
!

! IKEv2 Config Exchange configuration (IP addresses for spokes are assigned from pool)
crypto ikev2 authorization policy default
  pool SPOKES
  route set interface
  route accept any tag 20
!
!
! IKEv2 profile for Spokes - Smart Defaults used
crypto ikev2 profile default
  match identity remote any
  authentication remote pre-share key cisco
  authentication local pre-share key cisco
  aaa authorization group psk list default default
  virtual-template 1
!
interface Loopback0
  ip address 10.0.1.1 255.255.255.255
!
! GRE Tunnel configured to second HUB. It is required for spoke-to-spoke connectivity
! to work in all possible circumstances
```

```

! no BFD echo configuration is required to avoid Traffic Indication sent by remote HUB
! (BFD echo is having the same source and destination IP address)
!
interface Tunnel0
 ip address 10.0.0.1 255.255.255.0
 ip nhrp network-id 1
 ip nhrp redirect
bfd interval 50 min_rx 50 multiplier 3
no bfd echo
 tunnel source Ethernet0/2
 tunnel destination 192.168.0.2
!
interface Ethernet0/0
 ip address 172.16.0.1 255.255.255.0
!
interface Ethernet0/2
 ip address 192.168.0.1 255.255.255.0
bfd interval 50 min_rx 50 multiplier 5
!
interface Virtual-Templatel type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
!
! BGP Configuration
router bgp 1
 bgp log-neighbor-changes
! dynamic peer-groups are used for AGGR routers and SPOKES
 bgp listen range 192.168.0.0/24 peer-group DC
 bgp listen range 10.0.1.0/24 peer-group SPOKES
! BGP timers configured
 timers bgp 15 30
 neighbor SPOKES peer-group
 neighbor SPOKES remote-as 1
 neighbor DC peer-group
 neighbor DC remote-as 1
! Within DC BFD is used to determine neighbour status
 neighbor DC fall-over bfd
 neighbor 10.0.0.2 remote-as 1
! BFD is used to detect HUB2 status
 neighbor 10.0.0.2 fall-over bfd
!
 address-family ipv4
 redistribute connected
! route-map which determines what should be the local-pref
 redistribute static route-map LOCALPREF
 neighbor SPOKES activate
! to spokes only Aggregate/Summary routes are sent
 neighbor SPOKES route-map AGGR out
 neighbor DC activate
 neighbor DC route-reflector-client
 neighbor 10.0.0.2 activate
 neighbor 10.0.0.2 route-reflector-client
 exit-address-family
!
ip local pool SPOKES 10.0.1.2 10.0.1.254
!
! When HUB2 goes down Static Routes with Tag 500 are added and admin distance of 1
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
! Default static routes are with Tag 200 and admin distance of 150
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200

```

```

!
!
ip prefix-list AGGR seq 5 permit 192.168.0.0/16
ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
route-map AGGR permit 10
  match ip address prefix-list AGGR
!
route-map LOCALPREF permit 5
  match tag 500
  set local-preference 500
!
route-map LOCALPREF permit 10
  match tag 200
  set local-preference 200
!
route-map LOCALPREF permit 15
  match tag 20

```

Configuración R2-HUB2

```

hostname R2-HUB2
!
aaa new-model
!
aaa authorization network default local
!
track timer ip route msec 500
!
track 1 ip route 10.0.1.1 255.255.255.255 reachability
!
track 2 list boolean and
  object 1 not
  object 3
  object 4
!
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
!
crypto ikev2 authorization policy default
  pool SPOKES
  route set interface
  route accept any tag 20
!
!
crypto ikev2 profile default
  match identity remote any
  authentication remote pre-share key cisco
  authentication local pre-share key cisco
  aaa authorization group psk list default default
  virtual-template 1
!
!
interface Loopback0
  ip address 10.0.2.1 255.255.255.255
!
interface Tunnel0
  ip address 10.0.0.2 255.255.255.0
  ip nhrp network-id 1

```

```

ip nhrp redirect
bfd interval 50 min_rx 50 multiplier 3
no bfd echo
tunnel source Ethernet0/2
tunnel destination 192.168.0.1
!
interface Ethernet0/0
 ip address 172.16.0.2 255.255.255.0
!
interface Ethernet0/2
 ip address 192.168.0.2 255.255.255.0
 bfd interval 50 min_rx 50 multiplier 5
!
interface Virtual-Templatel type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
!
router bgp 1
 bgp log-neighbor-changes
 bgp listen range 192.168.0.0/24 peer-group DC
 bgp listen range 10.0.2.0/24 peer-group SPOKES
 timers bgp 15 30
 neighbor SPOKES peer-group
 neighbor SPOKES remote-as 1
 neighbor DC peer-group
 neighbor DC remote-as 1
 neighbor DC fall-over bfd
 neighbor 10.0.0.1 remote-as 1
 neighbor 10.0.0.1 fall-over bfd
!
 address-family ipv4
 redistribute connected
 redistribute static route-map LOCALPREF
 neighbor SPOKES activate
 neighbor SPOKES route-map AGGR out
 neighbor DC activate
 neighbor DC route-reflector-client
 neighbor 10.0.0.1 activate
 neighbor 10.0.0.1 route-reflector-client
 exit-address-family
!
ip local pool SPOKES 10.0.2.2 10.0.2.254
ip forward-protocol nd
!
!
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
!
!
ip prefix-list AGGR seq 5 permit 192.168.0.0/16
ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
route-map AGGR permit 10
 match ip address prefix-list AGGR
!
route-map LOCALPREF permit 5
 match tag 500
 set local-preference 500
!
route-map LOCALPREF permit 10

```

```
match tag 200
set local-preference 100
!
route-map LOCALPREF permit 15
match tag 20
```

Configuración R3-SPOKE1

```
hostname R2-HUB2
!
aaa new-model
!
aaa authorization network default local
!
track timer ip route msec 500
!
track 1 ip route 10.0.1.1 255.255.255.255 reachability
!
track 2 list boolean and
object 1 not
object 3
object 4
!
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
!
crypto ikev2 authorization policy default
pool SPOKES
route set interface
route accept any tag 20
!
!
crypto ikev2 profile default
match identity remote any
authentication remote pre-share key cisco
authentication local pre-share key cisco
aaa authorization group psk list default default
virtual-template 1
!
!
interface Loopback0
ip address 10.0.2.1 255.255.255.255
!
interface Tunnel0
ip address 10.0.0.2 255.255.255.0
ip nhrp network-id 1
ip nhrp redirect
bfd interval 50 min_rx 50 multiplier 3
no bfd echo
tunnel source Ethernet0/2
tunnel destination 192.168.0.1
!
interface Ethernet0/0
ip address 172.16.0.2 255.255.255.0
!
interface Ethernet0/2
ip address 192.168.0.2 255.255.255.0
bfd interval 50 min_rx 50 multiplier 5
!
```

```

interface Virtual-Templatel type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
!
router bgp 1
 bgp log-neighbor-changes
 bgp listen range 192.168.0.0/24 peer-group DC
 bgp listen range 10.0.2.0/24 peer-group SPOKES
 timers bgp 15 30
 neighbor SPOKES peer-group
 neighbor SPOKES remote-as 1
 neighbor DC peer-group
 neighbor DC remote-as 1
 neighbor DC fall-over bfd
 neighbor 10.0.0.1 remote-as 1
 neighbor 10.0.0.1 fall-over bfd
!
 address-family ipv4
  redistribute connected
  redistribute static route-map LOCALPREF
  neighbor SPOKES activate
  neighbor SPOKES route-map AGGR out
  neighbor DC activate
  neighbor DC route-reflector-client
  neighbor 10.0.0.1 activate
  neighbor 10.0.0.1 route-reflector-client
  exit-address-family
!
 ip local pool SPOKES 10.0.2.2 10.0.2.254
 ip forward-protocol nd
!
!
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
!
!
 ip prefix-list AGGR seq 5 permit 192.168.0.0/16
 ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
 route-map AGGR permit 10
  match ip address prefix-list AGGR
!
 route-map LOCALPREF permit 5
  match tag 500
  set local-preference 500
!
 route-map LOCALPREF permit 10
  match tag 200
  set local-preference 100
!
 route-map LOCALPREF permit 15
  match tag 20

```

Configuración R4-SPOKE2

```

hostname R2-HUB2
!

```



```
aaa new-model
!
aaa authorization network default local
!
track timer ip route msec 500
!
track 1 ip route 10.0.1.1 255.255.255.255 reachability
!
track 2 list boolean and
  object 1 not
  object 3
  object 4
!
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
!
crypto ikev2 authorization policy default
  pool SPOKES
  route set interface
  route accept any tag 20
!
!
crypto ikev2 profile default
  match identity remote any
  authentication remote pre-share key cisco
  authentication local pre-share key cisco
  aaa authorization group psk list default default
  virtual-template 1
!
!
interface Loopback0
  ip address 10.0.2.1 255.255.255.255
!
interface Tunnel0
  ip address 10.0.0.2 255.255.255.0
  ip nhrp network-id 1
  ip nhrp redirect
  bfd interval 50 min_rx 50 multiplier 3
  no bfd echo
  tunnel source Ethernet0/2
  tunnel destination 192.168.0.1
!
interface Ethernet0/0
  ip address 172.16.0.2 255.255.255.0
!
interface Ethernet0/2
  ip address 192.168.0.2 255.255.255.0
  bfd interval 50 min_rx 50 multiplier 5
!
interface Virtual-Templatel type tunnel
  ip unnumbered Loopback0
  ip nhrp network-id 1
  ip nhrp redirect
  tunnel protection ipsec profile default
!
router bgp 1
  bgp log-neighbor-changes
  bgp listen range 192.168.0.0/24 peer-group DC
  bgp listen range 10.0.2.0/24 peer-group SPOKES
  timers bgp 15 30
  neighbor SPOKES peer-group
  neighbor SPOKES remote-as 1
```

```

neighbor DC peer-group
neighbor DC remote-as 1
neighbor DC fall-over bfd
neighbor 10.0.0.1 remote-as 1
neighbor 10.0.0.1 fall-over bfd
!
address-family ipv4
redistribute connected
redistribute static route-map LOCALPREF
neighbor SPOKES activate
neighbor SPOKES route-map AGGR out
neighbor DC activate
neighbor DC route-reflector-client
neighbor 10.0.0.1 activate
neighbor 10.0.0.1 route-reflector-client
exit-address-family
!
ip local pool SPOKES 10.0.2.2 10.0.2.254
ip forward-protocol nd
!
!
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
!
!
ip prefix-list AGGR seq 5 permit 192.168.0.0/16
ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
route-map AGGR permit 10
  match ip address prefix-list AGGR
!
route-map LOCALPREF permit 5
  match tag 500
  set local-preference 500
!
route-map LOCALPREF permit 10
  match tag 200
  set local-preference 100
!
route-map LOCALPREF permit 15
  match tag 20

```

Configuración R5-AGGR1

```

hostname R2-HUB2
!
aaa new-model
!
aaa authorization network default local
!
track timer ip route msec 500
!
track 1 ip route 10.0.1.1 255.255.255.255 reachability
!
track 2 list boolean and
  object 1 not
  object 3
  object 4
!

```

```
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
!
crypto ikev2 authorization policy default
  pool SPOKES
  route set interface
  route accept any tag 20
!
!
crypto ikev2 profile default
  match identity remote any
  authentication remote pre-share key cisco
  authentication local pre-share key cisco
  aaa authorization group psk list default default
  virtual-template 1
!
!
interface Loopback0
  ip address 10.0.2.1 255.255.255.255
!
interface Tunnel0
  ip address 10.0.0.2 255.255.255.0
  ip nhrp network-id 1
  ip nhrp redirect
  bfd interval 50 min_rx 50 multiplier 3
  no bfd echo
  tunnel source Ethernet0/2
  tunnel destination 192.168.0.1
!
interface Ethernet0/0
  ip address 172.16.0.2 255.255.255.0
!
interface Ethernet0/2
  ip address 192.168.0.2 255.255.255.0
  bfd interval 50 min_rx 50 multiplier 5
!
interface Virtual-Templatel type tunnel
  ip unnumbered Loopback0
  ip nhrp network-id 1
  ip nhrp redirect
  tunnel protection ipsec profile default
!
router bgp 1
  bgp log-neighbor-changes
  bgp listen range 192.168.0.0/24 peer-group DC
  bgp listen range 10.0.2.0/24 peer-group SPOKES
  timers bgp 15 30
  neighbor SPOKES peer-group
  neighbor SPOKES remote-as 1
  neighbor DC peer-group
  neighbor DC remote-as 1
  neighbor DC fall-over bfd
  neighbor 10.0.0.1 remote-as 1
  neighbor 10.0.0.1 fall-over bfd
!
address-family ipv4
  redistribute connected
  redistribute static route-map LOCALPREF
  neighbor SPOKES activate
  neighbor SPOKES route-map AGGR out
  neighbor DC activate
  neighbor DC route-reflector-client
```

```

neighbor 10.0.0.1 activate
neighbor 10.0.0.1 route-reflector-client
exit-address-family
!
ip local pool SPOKES 10.0.2.2 10.0.2.254
ip forward-protocol nd
!
!
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
!
!
ip prefix-list AGGR seq 5 permit 192.168.0.0/16
ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
route-map AGGR permit 10
  match ip address prefix-list AGGR
!
route-map LOCALPREF permit 5
  match tag 500
  set local-preference 500
!
route-map LOCALPREF permit 10
  match tag 200
  set local-preference 100
!
route-map LOCALPREF permit 15
  match tag 20

```

Configuración R6-AGGR2

```

hostname R2-HUB2
!
aaa new-model
!
aaa authorization network default local
!
track timer ip route msec 500
!
track 1 ip route 10.0.1.1 255.255.255.255 reachability
!
track 2 list boolean and
  object 1 not
  object 3
  object 4
!
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
!
crypto ikev2 authorization policy default
  pool SPOKES
  route set interface
  route accept any tag 20
!
!
crypto ikev2 profile default
  match identity remote any

```

```
authentication remote pre-share key cisco
authentication local pre-share key cisco
aaa authorization group psk list default default
virtual-template 1
!
!
interface Loopback0
 ip address 10.0.2.1 255.255.255.255
!
interface Tunnel0
 ip address 10.0.0.2 255.255.255.0
 ip nhrp network-id 1
 ip nhrp redirect
 bfd interval 50 min_rx 50 multiplier 3
 no bfd echo
 tunnel source Ethernet0/2
 tunnel destination 192.168.0.1
!
interface Ethernet0/0
 ip address 172.16.0.2 255.255.255.0
!
interface Ethernet0/2
 ip address 192.168.0.2 255.255.255.0
 bfd interval 50 min_rx 50 multiplier 5
!
interface Virtual-Templatel type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
!
router bgp 1
 bgp log-neighbor-changes
 bgp listen range 192.168.0.0/24 peer-group DC
 bgp listen range 10.0.2.0/24 peer-group SPOKES
 timers bgp 15 30
 neighbor SPOKES peer-group
 neighbor SPOKES remote-as 1
 neighbor DC peer-group
 neighbor DC remote-as 1
 neighbor DC fall-over bfd
 neighbor 10.0.0.1 remote-as 1
 neighbor 10.0.0.1 fall-over bfd
!
 address-family ipv4
 redistribute connected
 redistribute static route-map LOCALPREF
 neighbor SPOKES activate
 neighbor SPOKES route-map AGGR out
 neighbor DC activate
 neighbor DC route-reflector-client
 neighbor 10.0.0.1 activate
 neighbor 10.0.0.1 route-reflector-client
 exit-address-family
!
ip local pool SPOKES 10.0.2.2 10.0.2.254
ip forward-protocol nd
!
!
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
!
```

```

!
ip prefix-list AGGR seq 5 permit 192.168.0.0/16
ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
route-map AGGR permit 10
  match ip address prefix-list AGGR
!
route-map LOCALPREF permit 5
  match tag 500
  set local-preference 500
!
route-map LOCALPREF permit 10
  match tag 200
  set local-preference 100
!
route-map LOCALPREF permit 15
  match tag 20

```

Configuración R7-HOST (simulación del HOST en esa red)

```

hostname R2-HUB2
!
aaa new-model
!
aaa authorization network default local
!
track timer ip route msec 500
!
track 1 ip route 10.0.1.1 255.255.255.255 reachability
!
track 2 list boolean and
  object 1 not
  object 3
  object 4
!
track 3 ip route 10.0.5.1 255.255.255.255 reachability
!
track 4 ip route 10.0.6.1 255.255.255.255 reachability
!
!
crypto ikev2 authorization policy default
  pool SPOKES
  route set interface
  route accept any tag 20
!
!
crypto ikev2 profile default
  match identity remote any
  authentication remote pre-share key cisco
  authentication local pre-share key cisco
  aaa authorization group psk list default default
  virtual-template 1
!
!
interface Loopback0
  ip address 10.0.2.1 255.255.255.255
!
interface Tunnel0
  ip address 10.0.0.2 255.255.255.0
  ip nhrp network-id 1
  ip nhrp redirect

```

```

bfd interval 50 min_rx 50 multiplier 3
no bfd echo
tunnel source Ethernet0/2
tunnel destination 192.168.0.1
!
interface Ethernet0/0
 ip address 172.16.0.2 255.255.255.0
!
interface Ethernet0/2
 ip address 192.168.0.2 255.255.255.0
 bfd interval 50 min_rx 50 multiplier 5
!
interface Virtual-Templatel type tunnel
 ip unnumbered Loopback0
 ip nhrp network-id 1
 ip nhrp redirect
 tunnel protection ipsec profile default
!
router bgp 1
 bgp log-neighbor-changes
 bgp listen range 192.168.0.0/24 peer-group DC
 bgp listen range 10.0.2.0/24 peer-group SPOKES
 timers bgp 15 30
 neighbor SPOKES peer-group
 neighbor SPOKES remote-as 1
 neighbor DC peer-group
 neighbor DC remote-as 1
 neighbor DC fall-over bfd
 neighbor 10.0.0.1 remote-as 1
 neighbor 10.0.0.1 fall-over bfd
!
 address-family ipv4
 redistribute connected
 redistribute static route-map LOCALPREF
 neighbor SPOKES activate
 neighbor SPOKES route-map AGGR out
 neighbor DC activate
 neighbor DC route-reflector-client
 neighbor 10.0.0.1 activate
 neighbor 10.0.0.1 route-reflector-client
 exit-address-family
!
ip local pool SPOKES 10.0.2.2 10.0.2.254
ip forward-protocol nd
!
!
ip route 192.168.0.0 255.255.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 tag 500 track 2
ip route 10.0.0.0 255.0.0.0 Null0 150 tag 200
ip route 192.168.0.0 255.255.0.0 Null0 150 tag 200
!
!
ip prefix-list AGGR seq 5 permit 192.168.0.0/16
ip prefix-list AGGR seq 10 permit 10.0.0.0/8
!
route-map AGGR permit 10
 match ip address prefix-list AGGR
!
route-map LOCALPREF permit 5
 match tag 500
 set local-preference 500
!
route-map LOCALPREF permit 10
 match tag 200

```

```
set local-preference 100
!  
route-map LOCALPREF permit 15  
match tag 20
```

Notas de configuración importantes

Aquí están algunas NOTAS IMPORTANTES sobre las configuraciones que se describen en las secciones anteriores:

- El túnel GRE de punto a punto entre el dos Hubs se requiere para que la Conectividad del spoke al spoke trabaje en los escenarios Allto, incluir específicamente esos escenarios en los cuales algo del spokes esté conectado solamente con uno del Hubs y otros a otro concentrador.
- No se requiere la **ninguna** configuración de la **generación de eco del bfd** en la interfaz de túnel GRE entre el dos Hubs para evitar la indicación del tráfico que se envía de otro concentrador. La generación de eco BFD tiene el mismo IP Address de origen y de destino, que es igual a la dirección IP del router que envía la generación de eco BFD. Puesto que estos paquetes son ruteados detrás por el router que responde, se generan las indicaciones del tráfico NHRP.
- En la configuración BGP, la filtración del route-map que hace publicidad de las redes hacia el spokes no se requiere, solamente él hace las configuraciones más óptimas puesto que solamente global/las rutas de resumen se hacen publicidad:

```
neighbor SPOKES route-map AGGR out
```

- En el Hubs, la configuración del **route-map LOCALPREF** se requiere para configurar la preferencia local apropiada BGP, y filtra las Static rutas redistribuidas a las rutas solamente del resumen y del modo de configuración IKEv2.
- Este diseño no dirige la Redundancia en las ubicaciones de la oficina remota (spoke). Si va el link PÁLIDO en el spoke abajo, el VPN también no trabaja. Agregue un segundo link al router radial o agregue a un segundo router radial dentro de la misma ubicación para abordar este problema.

En resumen, el diseño de la Redundancia que se presenta en este documento se puede tratar como alternativa moderna a la característica del Stateful Switchover (SSO) /Stateful. Es altamente flexible y puede ser ajustado para cumplir sus requisitos específicos del despliegue.

Verificación

Actualmente, no hay un procedimiento de verificación disponible para esta configuración.

Troubleshooting

Actualmente, no hay información específica de troubleshooting disponible para esta configuración.

Información Relacionada

- [Hoja de datos de FlexVPN del Cisco IOS](#)
- [Configurar FlexVPN habló al spoke](#)
- [Soporte Técnico y Documentación - Cisco Systems](#)