

FlexVPN HA Dual Hub Configuration Example

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components used](#)

[Background Information](#)

[Configure](#)

[Network Diagram](#)

[Regular Operational Scenario](#)

[Spoke-to-Spoke \(Shortcut\)](#)

[Routing Tables and Outputs for Regular Operational Scenario](#)

[HUB1 Failure Scenario](#)

[Configurations](#)

[R1-HUB Configuration](#)

[R2-HUB2 Configuration](#)

[R3-SPOKE1 Configuration](#)

[R4-SPOKE2 Configuration](#)

[R5-AGGR1 Configuration](#)

[R6-AGGR2 Configuration](#)

[R7-HOST Configuration \(simulation of HOST in that network\)](#)

[Important Configuration Notes](#)

[Verify](#)

[Troubleshoot](#)

[Related Information](#)

Introduction

This document describes how to configure a full redundancy design for Remote Offices that connect to a Data Center via IPSec-based VPN over an insecure network medium, such as the Internet.

Prerequisites

Requirements

There are no specific requirements for this document.

Components used

The information in this document is based on these technology components:

- [Border Gateway Protocol](#) (BGP) as the routing protocol within the Data Center and between spokes and hubs in the VPN overlay.
- [Bidirectional Forwarding Detection](#) (BFD) as a mechanism that detects down links (router down) that run inside of the Data Center only (not over the overlay tunnels).
- [Cisco IOS[®] FlexVPN](#) between the hubs and spokes, with spoke-to-spoke capabilities enabled via short-cut switching.
- [Generic Routing Encapsulation \(GRE\) tunneling](#) between two hubs in order to enable spoke-to-spoke communication, even when the spokes are connected to different hubs.
- [Enhanced Object Tracking](#) and static routes tied to the tracked objects.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Background Information

When you design remote access solutions for the Data Center, High Availability (HA) is often a key requirement for mission-critical user applications.

The solution that is presented in this document allows fast detection and recovery from failure scenarios in which one of the VPN-terminating hubs goes down due to a reload, upgrade, or power problems. All of the Remote Offices routers (spokes) then use the other operational hub immediately upon detection of such failure.

Here are the advantages of this design:

- Fast network recovery from a VPN hub-down scenario
- No complicated stateful synchronizations (such as IPsec Security Associations (SAs), Internet Security Association and Key Management Protocol (ISAKMP) SAs, and Crypto-routing) between the VPN hubs
- No anti-replay problems due to delays in the Encapsulating Security Payload (ESP) sequence number synchronization with IPsec Stateful HA
- VPN hubs can use different Cisco IOS/IOS-XE based hardware or software
- Flexible load-balancing implementation choices with BGP as the routing protocol that runs in the VPN overlay
- Clear and readable routing on all devices with no hidden mechanisms that run in the background
- Direct spoke-to-spoke connectivity

- All of the [FlexVPN](#) advantages, to include Authentication, Authorization, and Accounting (AAA) integration and per-tunnel Quality of Service (QoS)

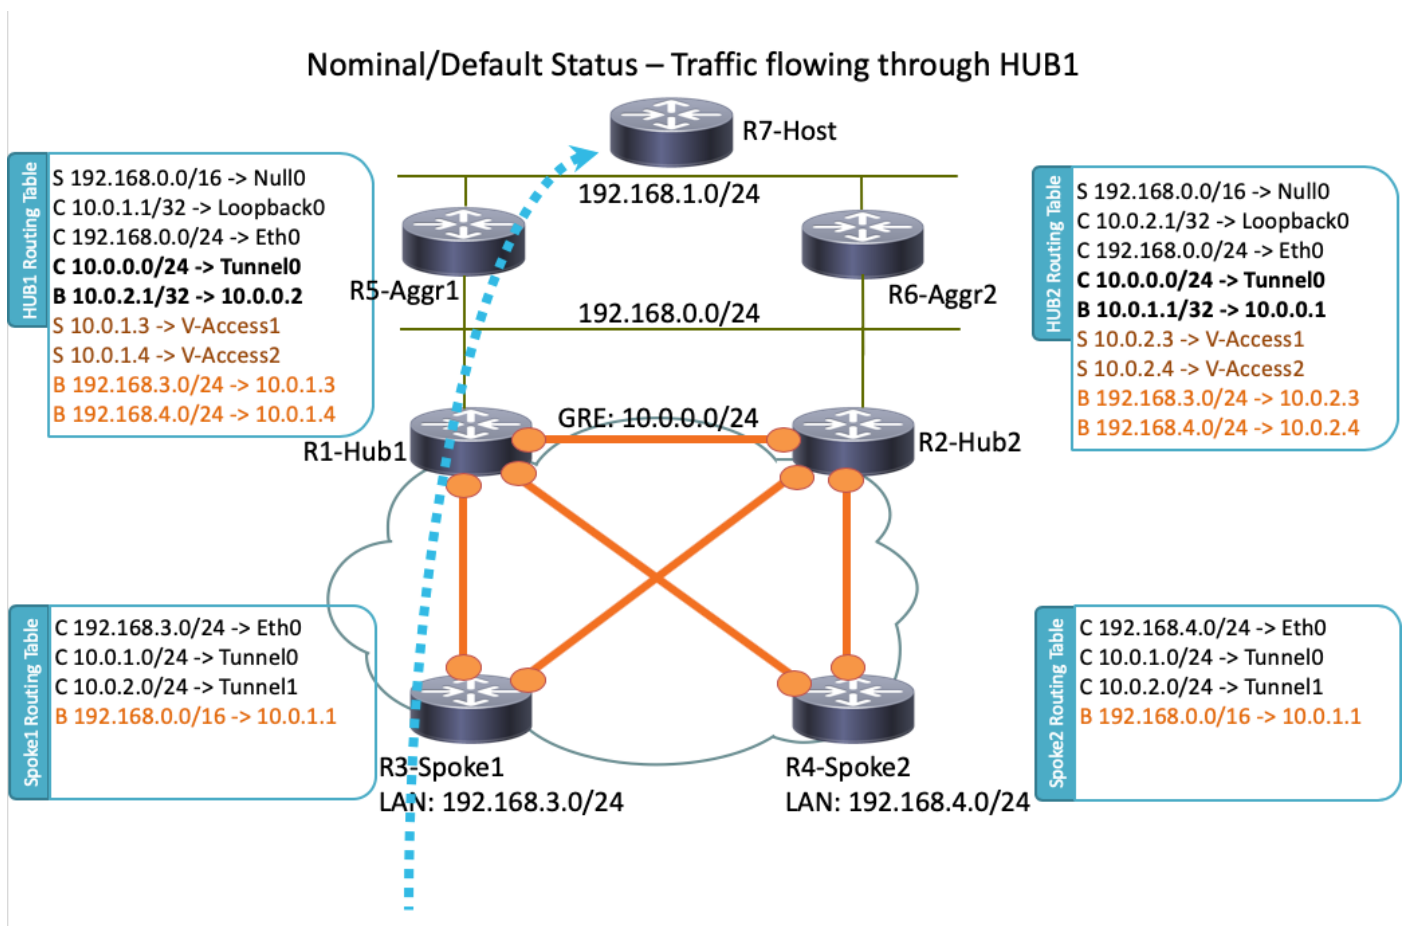
Configure

This section provides example scenarios and describes how to configure a full redundancy design for Remote Offices that connect to the Data Center via IPsec-based VPN over an insecure network medium.

Note: Use the [Command Lookup Tool](#) ([registered](#) customers only) in order to obtain more information on the commands used in this section.

Network Diagram

This is the network topology that is used in this document:



Note: All of the routers that are used in this topology run Cisco IOS Version 15.2(4)M1, and the Internet Cloud uses an address scheme of 172.16.0.0/24.

Regular Operational Scenario

In a normal operational scenario, when all of the routers are up and operational, all spoke routers route all of the traffic through the default hub (R1-HUB1). This routing preference is achieved when the default BGP local preference is set to 200 (refer to the sections that follow for


details). This can be adjusted based on the deployment requirements, such as traffic load-balancing.

Spoke-to-Spoke (Shortcut)


If R3-Spoke1 initiates a connection to R4-Spoke2, a dynamic spoke-to-spoke tunnel is created with the short-cut switching configuration.

Tip: For more details, refer to the [Configuring FlexVPN Spoke to Spoke](#) configuration guide.

If R3-Spoke1 is connected only to R1-HUB1, and R4-Spoke2 is connected only to R2-HUB2, a direct spoke-to-spoke connection can still be achieved with the point-to-point GRE tunnel that runs between the hubs. In this case, the initial traffic path between R3-Spoke1 and R4-Spoke2 appears similar to this:



Since R1-Hub1 receives the packet on the virtual-access interface, which has the same Next Hop Resolution Protocol (NHRP) network ID as that on the GRE tunnel, the Traffic Indication is sent toward the R3-Spoke1. This triggers spoke-to-spoke dynamic tunnel creation:



Routing Tables and Outputs for Regular Operational Scenario

Here is the R1-HUB1 routing table in a regular operational scenario:

```
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

```

Here is the R3-SPOKE1 routing table in a regular operational scenario after the spoke-to-spoke tunnel with R4-SPOKE2 is created:

```
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, Tunnel1
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, Tunnel1
C      10.0.2.3/32 is directly connected, Tunnel1
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

```

On R3-Spoke1, the BGP table has two entries for the **192.168.0.0/16** network with different local-preferences (R1-Hub1 is preferred):

```
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

```

Here is the R5-AGGR1 routing table in a regular operational scenario:

```

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

```

Here is the R7-HOST routing table in a regular operational scenario:

```

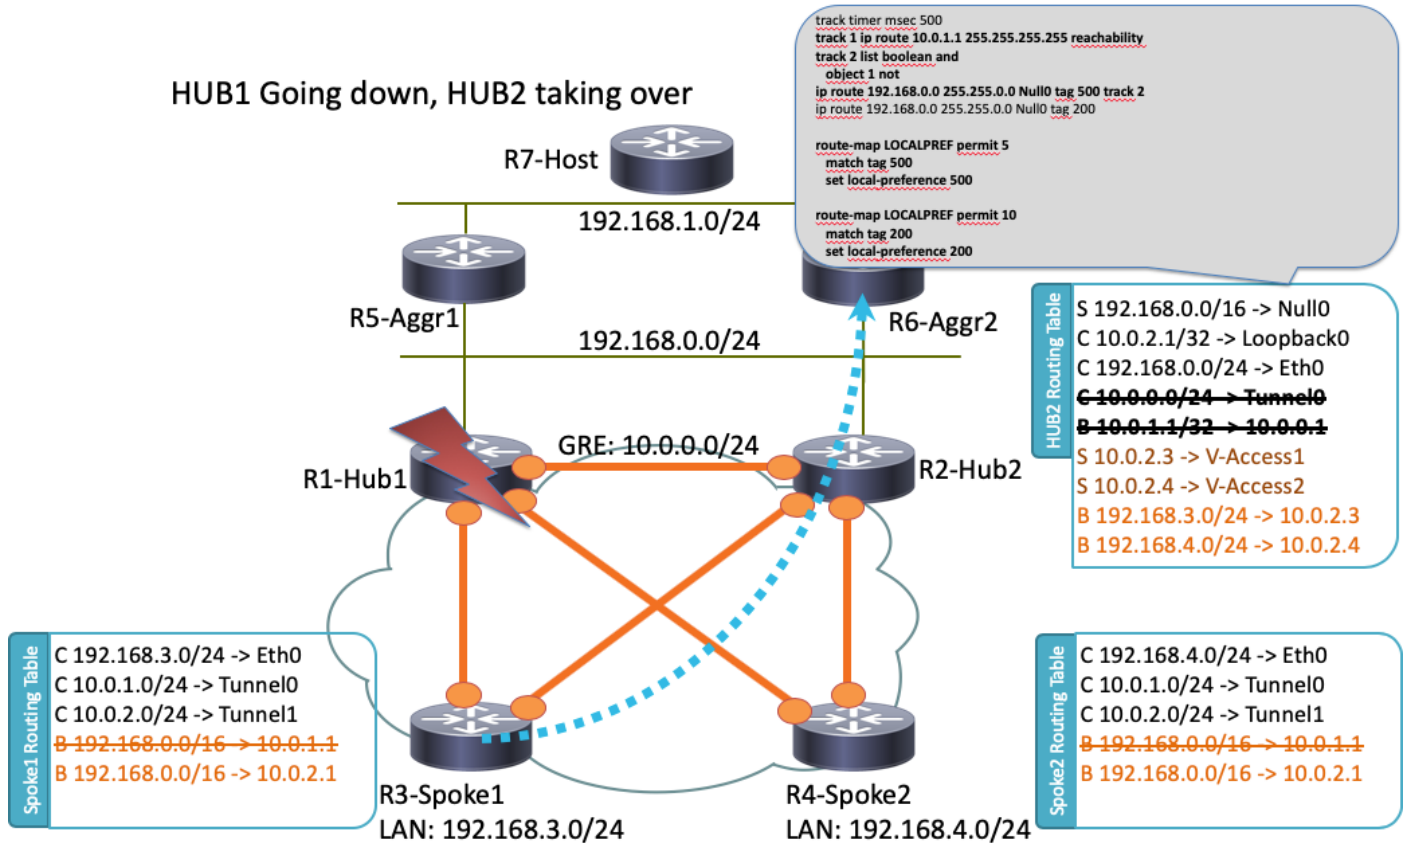
R7-HOST#show ip route
S*   0.0.0.0/0 [1/0] via 192.168.1.254
     192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C    192.168.1.0/24 is directly connected, Ethernet0/0
L    192.168.1.7/32 is directly connected, Ethernet0/0

```

HUB1 Failure Scenario

Here is an R1-HUB1 down scenario (due to actions such as power outages or an upgrade):

HUB1 Going down, HUB2 taking over



In this scenario, this sequence of events occurs:

1. The BFD on R2-HUB2 and on LAN aggregate routers R5-AGGR1 and R6-AGGR2 detect the down status of R1-HUB1. As a result, the BGP neighborhood immediately goes down.
2. The track object detection for R2-HUB2 that detects the presence of the R1-HUB1 loopback goes down (Track 1 in the example configuration).
3. This downed tracked object triggers another track to go up (Logical NOT). In this example, Track 2 goes up whenever Track 1 goes down.
4. This triggers a static IP routing entry to be added to the routing table due to a value that is lower than the default administrative distance. Here is the relevant configuration:

```

! 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 redistributes these static routes with a BGP local-preference that is larger than the value that is set for R1-HUB1. In this example, a local-preference of **500** is used in the failure scenario, instead of the **200** that is set by 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
!

```

On R3-Spoke1, you can see this in the BGP outputs. Note that the entry to R1 still exists, but it is not used:

```

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. At this point, both spokes (R3-Spoke1 and R4-Spoke2) begin to send traffic to R2-HUB2. All of these steps should occur within one second. Here is the routing table on 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. Later BGP sessions between the spokes and R1-HUB1 goes down, and Dead Peer Detection (DPD) removes the IPsec tunnels that are terminated on R1-HUB1. However, this does not impact traffic forwarding, since R2-HUB2 is already used as the main tunnel-terminating gateway:

```

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

Configurations

This section provides sample configurations for the hubs and spokes that are used in this topology.

R1-HUB Configuration

```
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

```

R2-HUB2 Configuration

```

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-Template1 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
```

R3-SPOKE1 Configuration

```
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-Template1 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

```

R4-SPOKE2 Configuration

```

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

```

R5-AGGR1 Configuration

```

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-Template1 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

```

R6-AGGR2 Configuration

```

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
```

R7-HOST Configuration (simulation of HOST in that network)

```
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

```

Important Configuration Notes

Here are some important notes about the configurations that are described in the previous sections:

- The point-to-point GRE tunnel between the two hubs is required for spoke-to-spoke connectivity to work in all scenarios, specifically to include those scenarios in which some of the spokes are connected only to one of the hubs and others to another hub.
- The **no bfd echo** configuration in the GRE tunnel interface between the two hubs is required in order to avoid the Traffic Indication that is sent out from another hub. The BFD Echo has the same source and destination IP address, which is equal to the IP address of the router that sends the BFD Echo. Since these packets are routed back by the router that responds, the NHRP Traffic Indications are generated.
- In the BGP configuration, route-map filtering that advertises the networks towards spokes is not required, but it makes the configurations more optimal since only aggregate/summary routes are advertised:

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

- On the hubs, the **route-map LOCALPREF** configuration is required in order to set up the proper BGP local preference, and it filters the redistributed static routes to only the summary and IKEv2 configuration mode routes.
- This design does not address redundancy at Remote Office locations (spoke). If the WAN link on the spoke goes down, VPN also does not work. Add a second link to the spoke router or add a second spoke router within the same location in order to address this issue.

In summary, the redundancy design that is presented in this document can be treated as a modern alternative to the Stateful Switchover (SSO)/Stateful feature. It is highly flexible and can be fine-tuned in order to meet your specific deployment requirements.

Verify

There is currently no verification procedure available for this configuration.

Troubleshoot

There is currently no specific troubleshooting information available for this configuration.

Related Information

- [Cisco IOS FlexVPN Data Sheet](#)
- [Configuring FlexVPN Spoke to Spoke](#)
- [Technical Support & Documentation - Cisco Systems](#)