FlexVPN HA双集线器配置示例

目录

简介

先决条件

- 要求
- 使用的组件

背景信息

配置

- 网络图
- 常规操作场景
- 分支到分支（快捷方式）
- 常规运行方案的路由表和输出
- HUB1故障场景
- 配置
  - R1-HUB配置
  - R2-HUB2配置
  - R3-SPOKE1配置
  - R4-SPOKE2配置
  - R5-AGGR1配置
  - R6-AGGR2配置
  - R7-HOST配置（模拟该网络中的HOST）
- 重要配置说明

验证

故障排除

相关信息

简介

本文档介绍如何为通过基于IPSec的VPN在不安全的网络介质（例如Internet）上连接到数据中心的远程办公室配置完全冗余设计。

先决条件

要求

本文档没有任何特定的要求。

使用的组件

本文档中的信息基于以下技术组件：

- 边界网关协议(BGP)作为数据中心内以及VPN重叠中的分支和集线器之间的路由协议。
- 双向转发检测(BFD)是一种机制，可检测仅运行在数据中心内部（而不是通过重叠隧道）的下行链路（路由器关闭）。
- 中心和分支之间的Cisco IOS® FlexVPN，通过快捷交换启用分支到分支功能。
- 在两个中心之间使用通用路由封装(GRE)隧道，以便实现分支到分支的通信，即使分支连接到不同的中心也是如此。
- 增强的对象跟踪和与被跟踪对象关联的静态路由。

本文档中的信息都是基于特定实验室环境中的设备编写的。本文档中使用的所有设备最初均采用原始（默认）配置。如果您使用的是真实网络，请确保您已经了解所有命令的潜在影响。

背景信息

在为数据中心设计远程访问解决方案时，高可用性(HA)通常是任务关键型用户应用程序的关键要求。

本文档中介绍的解决方案允许快速检测和从故障场景中恢复，在这些故障场景中，一个VPN终端集线器因重新加载、升级或电源问题而停机。所有远程办公室路由器（分支）在检测到此类故障后立即使用另一个运行集线器。

此设计的优点如下：

- 从VPN集线器关闭方案快速恢复网络
- VPN中心之间没有复杂的状态同步(例如IPSec安全关联(SA)、互联网安全关联和密钥管理协议(ISAKMP)SA以及加密路由)
- 由于与IPSec有状态HA的封装安全负载(ESP)序列号同步存在延迟，因此没有反重播问题
- VPN集线器可以使用不同的基于Cisco IOS/IOS-XE的硬件或软件
- 灵活的负载均衡实施选择，BGP作为VPN重叠中运行的路由协议
- 所有设备上没有后台运行的隐藏机制的清晰可读的路由
- 直接星型到星型连接
- FlexVPN的所有优势，包括身份验证、授权和记帐(AAA)集成以及每通道服务质量(QoS)

配置

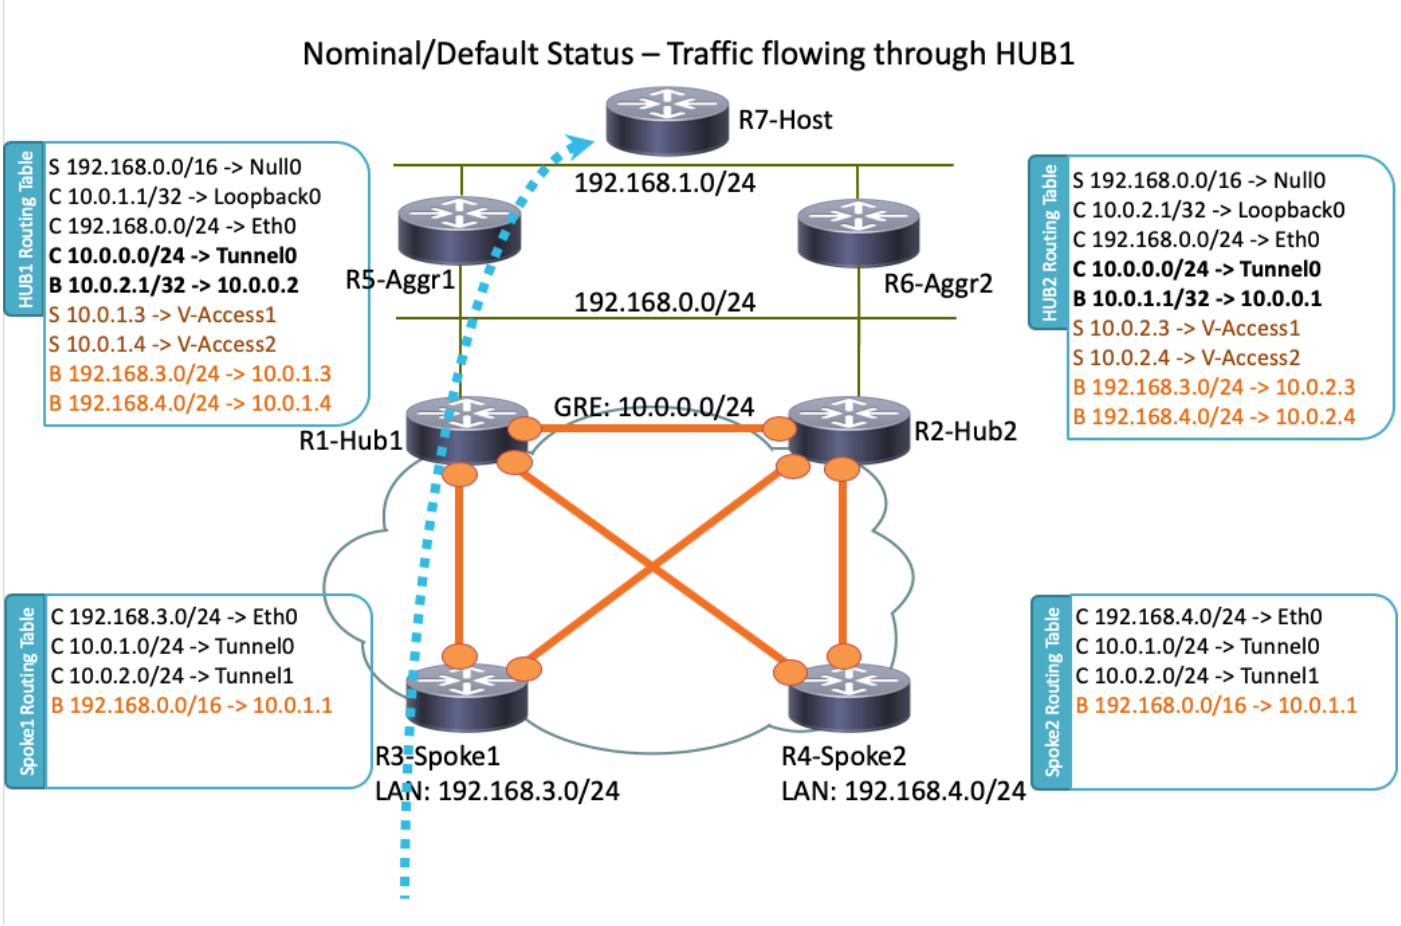
本节提供示例场景并介绍如何为通过不安全网络介质上基于IPSec的VPN连接到数据中心的远程办

公室配置完全冗余设计。

注意：使用命令查找工具（仅限注册用户）可获取有关本部分所使用命令的详细信息。

网络图

以下是本文档中使用的网络拓扑：



注意：此拓扑中使用的所有路由器都运行Cisco IOS版本15.2(4)M1，而Internet云使用的地址方案为172.16.0.0/24。

常规操作场景

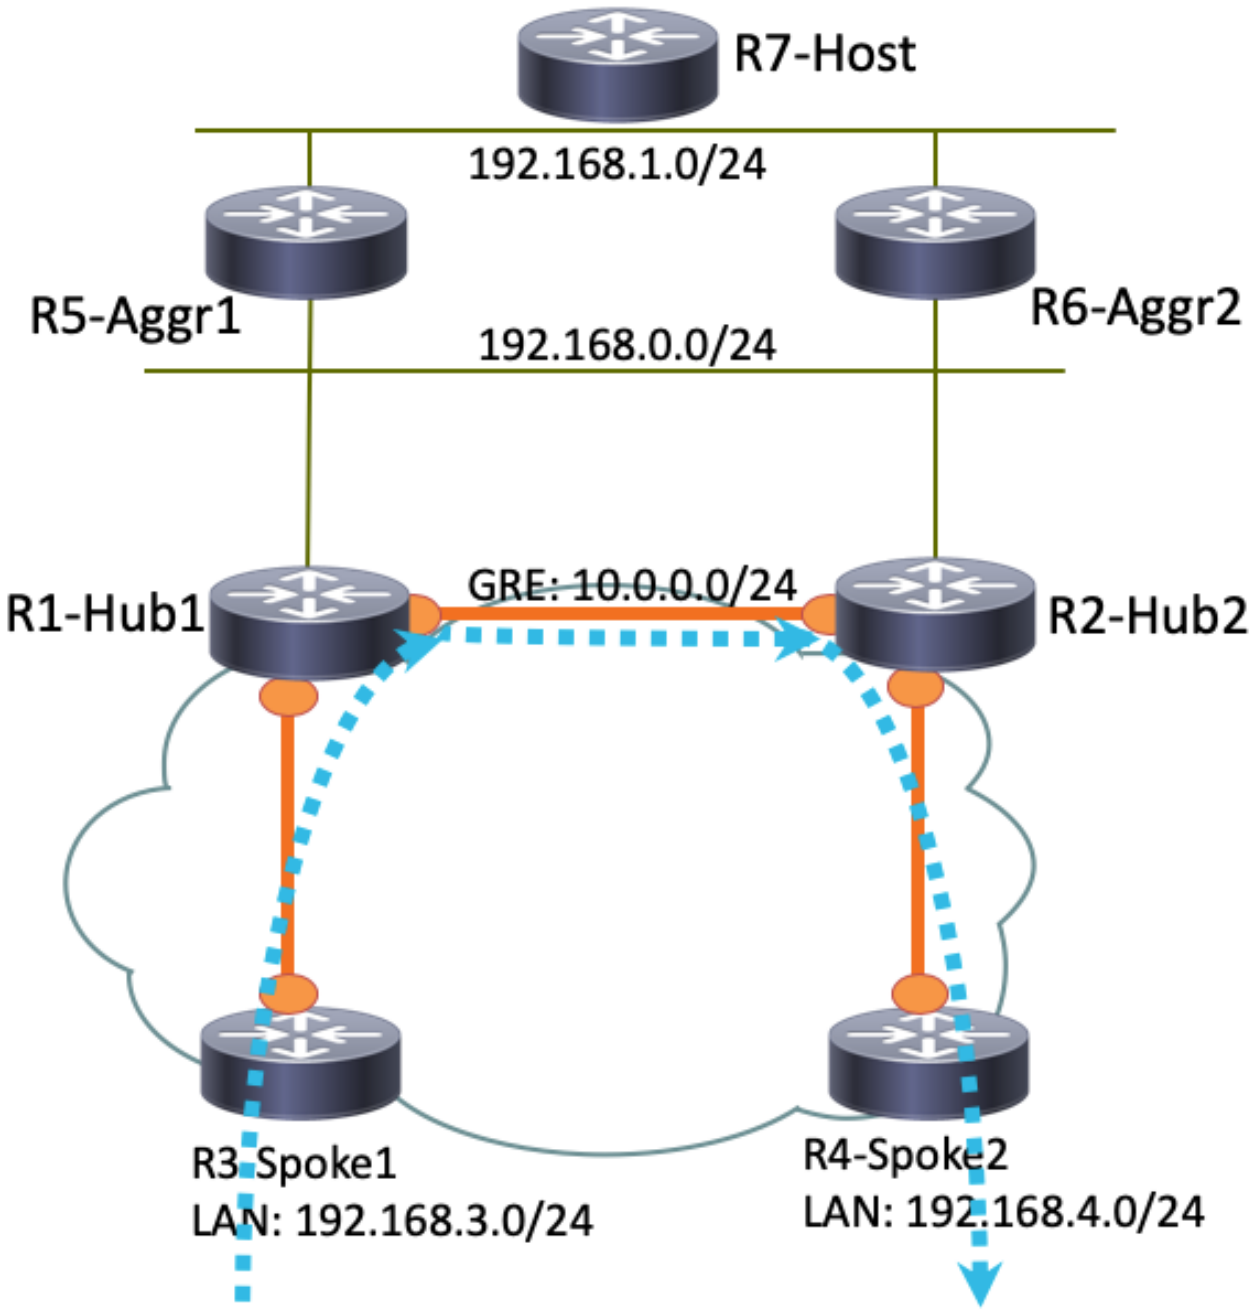
在正常运行情况下，当所有路由器都启动并运行正常时，所有分支路由器都会通过默认集线器(R1-HUB1)路由所有流量。 当默认BGP本地优先级设置为200时（请参阅后续部分了解详细信息），即可实现此路由首选项。这可以根据部署要求（如流量负载均衡）进行调整。

分支到分支（快捷方式）

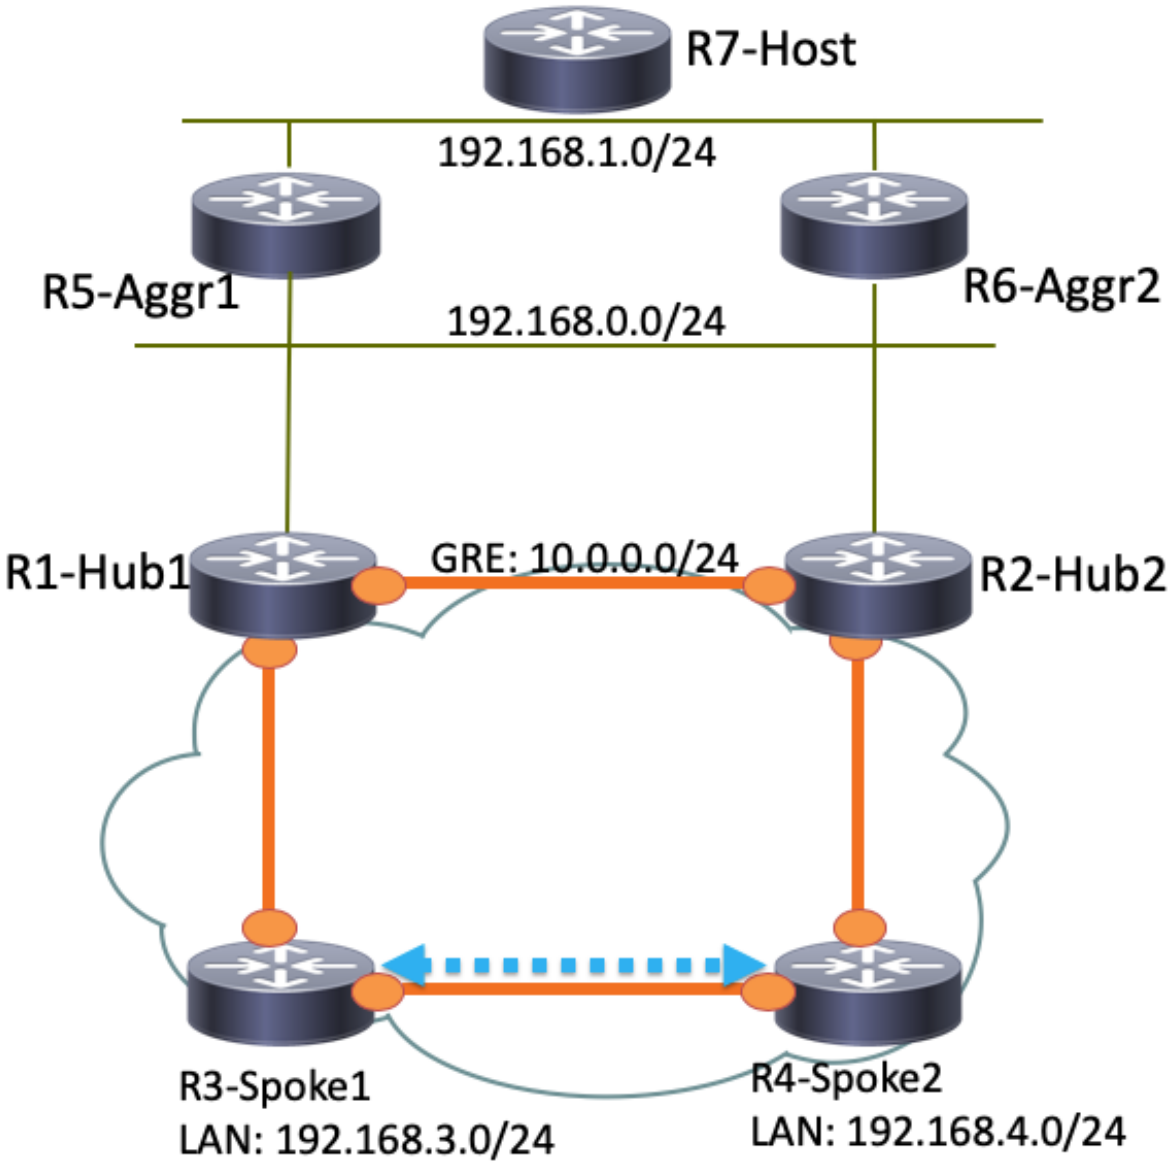
如果R3-Spoke1发起到R4-Spoke2的连接，则会使用快捷交换配置创建动态分支到分支隧道。

提示：有关详细信息，请参阅配置FlexVPN分支到分支配置指南。

如果R3-Spoke1仅连接到R1-HUB1，而R4-Spoke2仅连接到R2-HUB2，则仍可以通过在集线器之间运行的点对点GRE隧道实现直接辐射到辐射的连接。在这种情况下，R3-Spoke1和R4-Spoke2之间的初始流量路径如下所示：



由于R1-Hub1在虚拟访问接口上收到数据包，该虚拟访问接口与GRE隧道上的虚拟访问接口具有相同的下一跳解析协议(NHRP)网络ID，因此流量指示将发送到R3-Spoke1。这将触发创建分支到分支动态隧道：



常规运行方案的路由表和输出

以下是正常运行场景中的R1-HUB1路由表：

<#root>

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
```

以下是使用R4-SPOKE2创建辐射点到辐射点隧道后常规操作场景中的R3-SPOKE1路由表：

<#root>

```
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
```

在R3-Spoke1上，BGP表有两个用于192.168.0.0/16网络的条目，具有不同的本地首选项（首选R1-Hub1）：

```
<#root>

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
```

以下是正常运行场景中的R5-AGGR1路由表：

```
<#root>

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
```

以下是正常运行场景中的R7-HOST路由表：

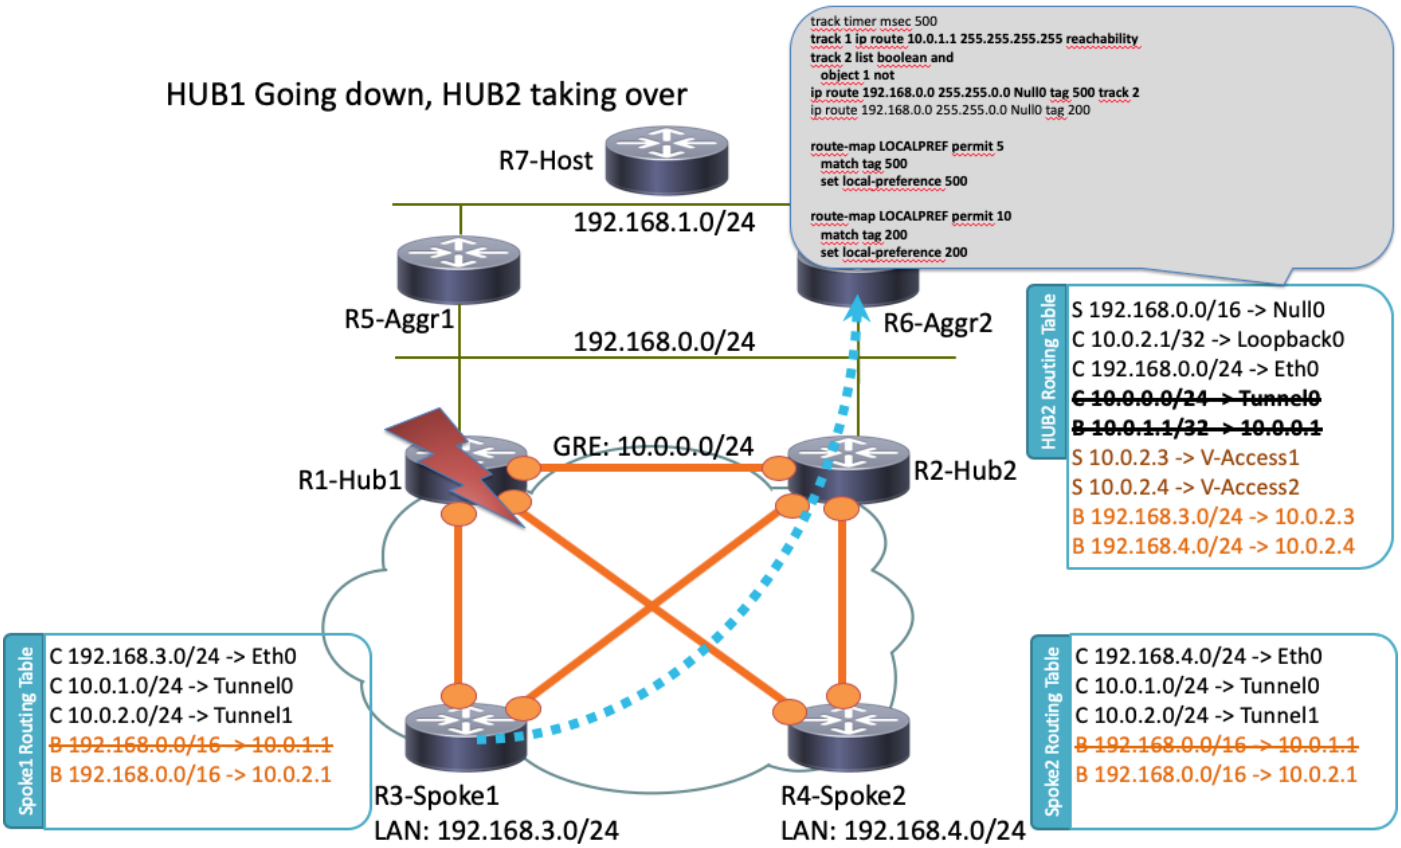
<#root>

```
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故障场景

以下是R1-HUB1关闭场景（由于断电或升级等操作）：



在此场景中，发生以下一系列事件：

1. R2-HUB2和LAN聚合路由器R5-AGGR1和R6-AGGR2上的BFD检测到R1-HUB1的关闭状态。因此，BGP邻居关系立即关闭。

2. 检测R1-HUB1环回是否存在的R2-HUB2的跟踪对象检测会中断（示例配置中的跟踪1）。

3. 此被下载的跟踪对象会触发另一个跟踪进入（逻辑非）。 在本例中，当轨道1关闭时，轨道2将启动。

4. 这会触发静态IP路由条目添加到路由表，因为它的值小于默认管理距离。以下是相关配置：

```
<#root>

! 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使用大于为R1-HUB1设置的值的BGP本地优先级重分布这些静态路由。在本示例中，本地优先级500用于故障方案，而不是由R1-HUB1设置的200:

```
<#root>

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


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


!
```

在R3-Spoke1上，可以在BGP输出中看到此消息。请注意，到R1的条目仍然存在，但未使用：

<#root>

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. 此时，两个分支（R3-Spoke1和R4-Spoke2）开始向R2-HUB2发送流量。所有这些步骤都应在一秒内完成。分支3上的路由表如下：

<#root>

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. 分支和R1-HUB1之间的后续BGP会话关闭，失效对等体检测(DPD)将删除R1-HUB1上终止的IPSec隧道。但是，这不会影响流量转发，因为R2-HUB2已用作主隧道终止网关：

<#root>

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
```

配置

本节提供此拓扑中使用的集线器和辐条的配置示例。

R1-HUB配置

<#root>

```
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-Template1 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配置

<#root>

```
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配置

```
hostname R3-SPOKE1
!
aaa new-model
!
aaa authorization network default local
!
!
crypto ikev2 authorization policy default
 route set interface
!
!
crypto ikev2 profile default
 match identity remote any
 authentication remote pre-share key cisco
 authentication local pre-share key cisco
 dpd 10 2 on-demand
 aaa authorization group psk list default default
!
! Tunnel to the HUB1
!
interface Tunnel0
 ip address negotiated
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 2
 tunnel source Ethernet0/0
 tunnel destination 172.16.0.1
 tunnel protection ipsec profile default
!
! Tunnel to the HUB2
!
interface Tunnel1
 ip address negotiated
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 2
 tunnel source Ethernet0/0
 tunnel destination 172.16.0.2
 tunnel protection ipsec profile default
!
interface Ethernet0/0
 description INTERNET-CLOUD
 ip address 172.16.0.3 255.255.255.0
!
```

```
interface Ethernet0/1
 description LAN
 ip address 192.168.3.3 255.255.255.0
!
interface Virtual-Template2 type tunnel
 ip unnumbered Ethernet0/1
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 2
 tunnel protection ipsec profile default
!
router bgp 1
 bgp log-neighbor-changes
 timers bgp 15 30
 neighbor 10.0.1.1 remote-as 1
 neighbor 10.0.2.1 remote-as 1
 !
 address-family ipv4
  network 192.168.3.0
  neighbor 10.0.1.1 activate
  neighbor 10.0.2.1 activate
 exit-address-family
```

R4-SPOKE2配置

```
hostname R4-SPOKE2
!
aaa new-model
!
aaa authorization network default local
!
!
crypto ikev2 authorization policy default
 route set interface
!
crypto ikev2 profile default
 match identity remote any
 authentication remote pre-share key cisco
 authentication local pre-share key cisco
 dpd 10 2 on-demand
 aaa authorization group psk list default default
!
interface Tunnel0
 ip address negotiated
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 2
 tunnel source Ethernet0/0
 tunnel destination 172.16.0.1
 tunnel protection ipsec profile default
!
interface Tunnel1
 ip address negotiated
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 2
 tunnel source Ethernet0/0
 tunnel destination 172.16.0.2
 tunnel protection ipsec profile default
!
```

```
interface Ethernet0/0
 ip address 172.16.0.4 255.255.255.0
!
interface Ethernet0/1
 ip address 192.168.4.4 255.255.255.0
!
interface Virtual-Template2 type tunnel
 ip unnumbered Ethernet0/1
 ip nhrp network-id 1
 ip nhrp shortcut virtual-template 2
 tunnel protection ipsec profile default
!
router bgp 1
 bgp log-neighbor-changes
 timers bgp 15 30
 neighbor 10.0.1.1 remote-as 1
 neighbor 10.0.2.1 remote-as 1
 !
 address-family ipv4
  network 192.168.4.0
  neighbor 10.0.1.1 activate
  neighbor 10.0.2.1 activate
 exit-address-family
!
```

R5-AGGR1配置

```
hostname R5-LAN1
!
no aaa new-model
!
!
interface Loopback0
 ip address 10.0.5.1 255.255.255.255
!
interface Ethernet0/0
 ip address 192.168.0.5 255.255.255.0
 bfd interval 50 min_rx 50 multiplier 5
!
! HSRP configuration on the LAN side
!
interface Ethernet0/1
 ip address 192.168.1.5 255.255.255.0
 standby 1 ip 192.168.1.254
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 192.168.0.1 remote-as 1
 neighbor 192.168.0.1 fall-over bfd
 neighbor 192.168.0.2 remote-as 1
 neighbor 192.168.0.2 fall-over bfd
 !
 address-family ipv4
  redistribute connected
  redistribute static
  neighbor 192.168.0.1 activate
  neighbor 192.168.0.2 activate
```

```
 exit-address-family
```

R6-AGGR2配置

```
hostname R6-LAN2
!
interface Loopback0
 ip address 10.0.6.1 255.255.255.255
!
interface Ethernet0/0
 ip address 192.168.0.6 255.255.255.0
 bfd interval 50 min_rx 50 multiplier 5
!
interface Ethernet0/1
 ip address 192.168.1.6 255.255.255.0
 standby 1 ip 192.168.1.254
 standby 1 priority 200
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 192.168.0.1 remote-as 1
 neighbor 192.168.0.1 fall-over bfd
 neighbor 192.168.0.2 remote-as 1
 neighbor 192.168.0.2 fall-over bfd
 !
 address-family ipv4
  redistribute connected
  redistribute static
  neighbor 192.168.0.1 activate
  neighbor 192.168.0.2 activate
 exit-address-family
!
```

R7-HOST配置（模拟该网络中的HOST）

```
hostname R7-HOST
!
no aaa new-model
!
interface Ethernet0/0
 ip address 192.168.1.7 255.255.255.0
!
ip route 0.0.0.0 0.0.0.0 192.168.1.254
```

重要配置说明

以下是有关前面部分中介绍的配置的一些重要说明：

- 两个中心之间的点对点GRE隧道是星型到辐射型连接在所有场景下正常工作所必需的，尤其是要包括某些分支仅连接到其中一个中心而其它分支连接到另一个中心的场景。

- 为了避免从另一个集线器发出的流量指示，需要在两个集线器之间的GRE隧道接口中配置no bfd echo。BFD回显具有相同的源IP地址和目的IP地址，与发送BFD回显的路由器的IP地址相同。由于这些数据包由响应的路由器路由回，因此会生成NHRP流量指示。

- 在BGP配置中，不需要将网络通告到分支的路由映射过滤，但是它使配置更加优化，因为仅通告聚合/汇总路由：

  <#root>

  neighbor SPOKES route-map AGGR out

- 在集线器上，需要route-map LOCALPREF配置才能设置适当的BGP本地优先级，并且它会将重分发的静态路由过滤为仅汇总路由和IKEv2配置模式路由。

- 此设计不解决远程办公室位置（分支）的冗余问题。 如果分支上的WAN链路断开，VPN也不起作用。为分支路由器添加另一条链路，或在同一位置添加另一台分支路由器以解决此问题。

总之，本文档中介绍的冗余设计可视为有状态切换(SSO)/有状态功能的现代替代方案。它非常灵活，可以进行微调，以满足您的特定部署要求。

验证

当前没有可用于此配置的验证过程。

故障排除

目前没有针对此配置的故障排除信息。

相关信息

- Cisco IOS FlexVPN产品手册
- 配置FlexVPN分支到分支
- 技术支持和文档 - Cisco Systems

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。