MPLS VPN Deployment and Application

Foreword

- BGP/MPLS IP VPN is widely used on WAN transport because it supports address space overlapping, flexible networking, good scalability, and MPLS traffic engineering (TE).
- MPLS VPN deployment modes vary according to customers' service requirements and networking.
- This course describes several common usage scenarios of MPLS VPN and how to deploy MPLS VPN in these scenarios. In addition, this course describes the extended functions of Open Shortest Path First (OSPF) for MPLS VPN.

- Note: Unless otherwise specified, MPLS VPN in this document indicates BGP/MPLS IP VPN.

Objectives

- Upon completion of this course, you will be able to:
  - Understand the usage scenarios and networking types of MPLS VPN.
  - Deploy the intranet solution with MPLS VPN.
  - Deploy the MPLS VPN Hub&Spoke solution.
  - Understand extended functions of OSPF for MPLS VPN.

Contents

Typical MPLS VPN Applications

Currently, MPLS VPN is mainly used for enterprise interconnection and virtual service networks.

- Enterprise interconnection: MPLS VPN connects the geographically different IP networks of branches, employees on business trips, and partners.
- Virtual service network: Multiple services, such as VoIP and IPTV, can run on the same physical network. A VPN is established for each type of service to isolate services.

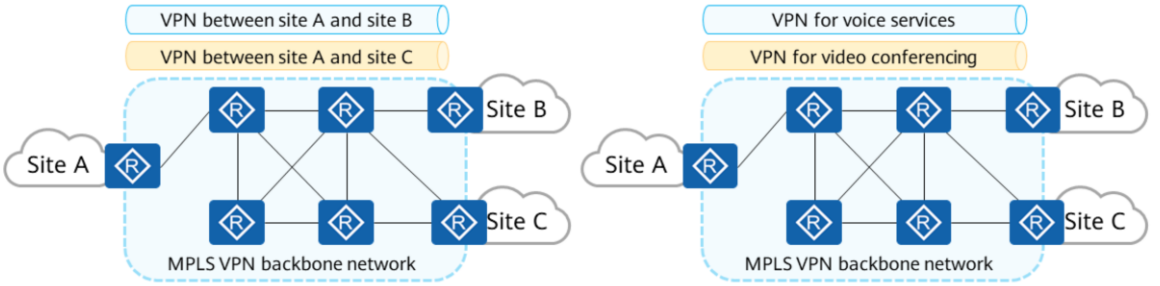

HUAWEI

- The advantages of MPLS VPN include but are not limited to the following:
  - One-hop access and network-wide connectivity can be implemented, and heterogeneous media interconnection is supported. Unlike the traditional leased line, which uses the same medium for a connection between each pair of user devices, the MPLS VPN can provide universal services conveniently.
  - Elastic bandwidth can be implemented. Traffic policing technology is used to ensure the required minimum bandwidth of users and implements best effort scheduling for burst traffic. In addition, the basic bandwidth can be soft expanded. That is, the basic bandwidth can be expanded within a range based on user requirements.
  - The MPLS VPN technology ensures the dedicated bandwidth of each VPN to meet the requirements of different users, traffic models, and QoS of various services.

Basic MPLS VPN Networking — Intranet

When the intranet networking solution is used, all users in a VPN are in a user group isolated from the users of other VPNs and can forward traffic to each other. Users in a VPN cannot communicate with users outside the VPN. The sites of a VPN usually belong to the same organization.

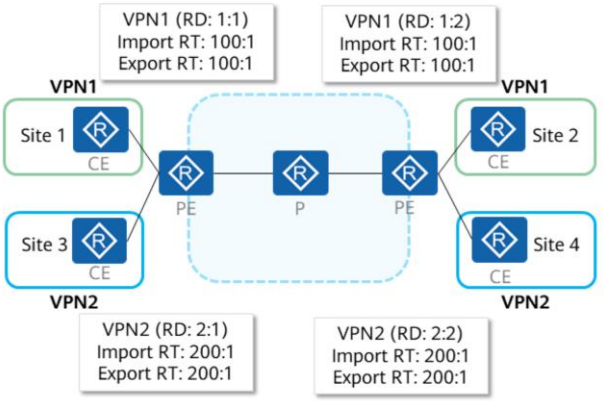

- A PE needs to create a VPN instance for each site and set a unique route distinguisher (RD) for each site.
- Import and export route targets (RTs) are set on PEs to prevent mutual communication between sites in different VPNs.

Basic MPLS VPN Networking — Extranet

When the extranet networking solution is used, VPN users can share network resources in some sites with other VPN users.

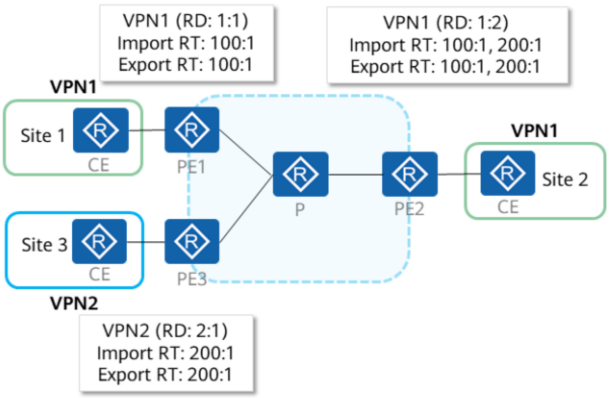

As shown in the figure, site 2 is a shared site that can be accessed by VPN1 and VPN2 users. The following requirements must be met:

- PE2 can receive VPNv4 routes advertised by PE1 and PE3.
- PE1 and PE3 can accept VPNv4 routes advertised by PE2.
- PE2 advertises neither VPNv4 routes received from PE1 to PE3 nor VPNv4 routes received from PE3 to PE1.

Basic MPLS VPN Networking — Hub&Spoke (1)

In the Hub&Spoke solution, one site can be configured as the hub site, and the other sites can be configured as spoke sites. Mutual access between sites must pass through the hub site. Data transmission between sites is centrally managed and controlled by the hub site.

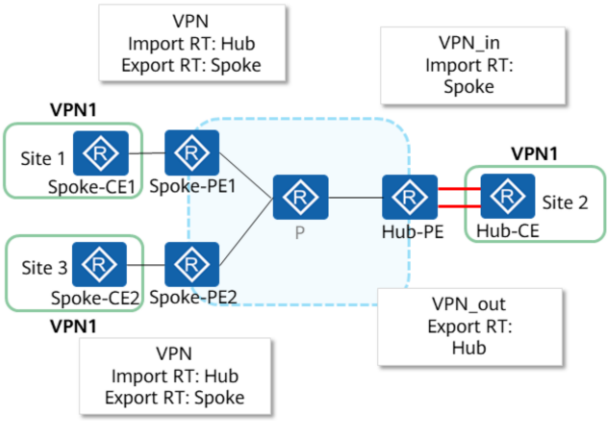

- A spoke site needs to advertise routes to a hub site, and then the hub site advertises the routes to other spoke sites. Spoke sites do not directly exchange routing information.
- For Spoke-PEs, set the export RT to "Spoke" and the import RT to "Hub."
- The Hub-PE needs to use two interfaces or sub-interfaces (bound to two VPN instances that are created). One interface or sub-interface is used to receive the routes from the Spoke-PE, and the import RT of the VPN instance is "Spoke." The other is used to advertise routes to Spoke-PEs, and the export RT of the VPN instance is "Hub."

Basic MPLS VPN Networking — Hub&Spoke (2)

The process of advertising routes from site 1 to site 2 is as follows:

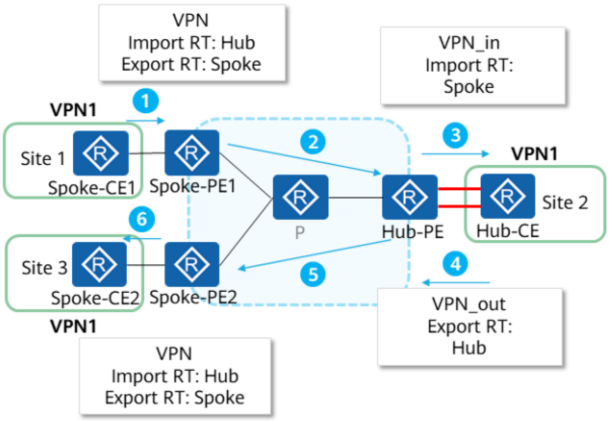

1. Spoke-CE1 advertises the route to Spoke-PE1.
2. Spoke-PE1 advertises the route to the Hub-PE through IBGP.
3. The Hub-PE imports the route into the VPN_in routing table through the import RT attribute of the VPN instance (VPN_in) and advertises the route to the Hub-CE.
4. The Hub-CE learns the route and advertises the route to the VPN instance (VPN_out) of the Hub-PE.
5. The Hub-PE advertises the route to Spoke-PE2 with the export RT attribute of VPN_out.
6. Spoke-PE2 advertises the route to Spoke-CE2.

MCE Networking

- When a private network is divided into VPNs based on services or networks, services of different VPN users must be completely isolated. In this case, configuring a CE for each VPN increases device and maintenance costs.
- A multi-VPN-instance CE (MCE) device can function as a CE for multiple VPN instances on an MPLS VPN network, reducing the investment on network devices.

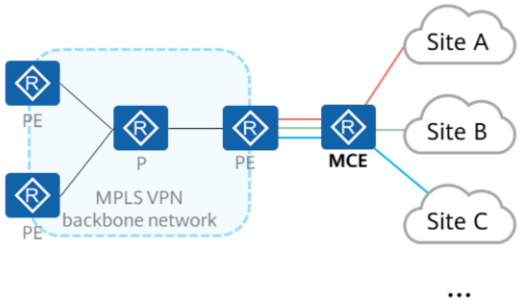

- The MCE device extends some functions of the PE to the CE. By binding different interfaces to VPNs, the MCE creates and maintains an independent routing and forwarding table (multi-VRF) for each VPN.
- The MCE can be connected to the corresponding PE through physical interfaces, sub-interfaces, or logical interfaces. On the PE, these interfaces must be bound to the corresponding VPN instances.

Inter-AS MPLS VPN Networking

- With the wide application of the MPLS VPN solution, the number and geographical scope of terminal users are increasing. The number of sites in an enterprise is increasing. It is common to connect a geographical location to another service provider, for example, between different MANs of a carrier, the backbone networks of the carriers that cooperate with each other may span different autonomous systems (ASs).
- Generally, the MPLS VPN architecture runs in an AS. The routing information of any VPN can only be flooded in one AS as needed. The inter-AS MPLS VPN solution is used to deploy the MPLS VPN between ASs.

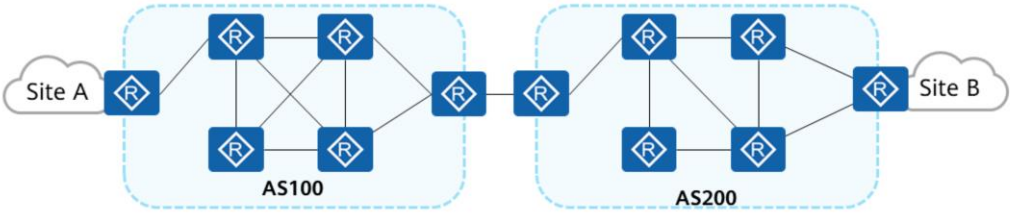

- RFC 2547 defines three inter-AS VPN solutions:
  - Inter-AS VPN Option A (inter-provider backbones Option A) mode: The inter-AS VPN manages its own VPN routes through dedicated interfaces between AS boundary routers (ASBRs). This mode is also called VRF-to-VRF.
  - Inter-AS VPN Option B (inter-provider backbones Option B) mode: ASBRs use MP-EBGP to advertise labeled VPNv4 routes. This mode is also called EBGP redistribution of labeled VPN-IPv4 routes.
  - Inter-AS VPN Option C (inter-provider backbones Option C): PEs use multi-hop MP-EBGP to advertise VPNv4 routes, which are also called multi-hop EBGP redistribution of labeled VPN-IPv4 routes.
- For more information about inter-AS MPLS VPN, see materials of the related HCIE-Datacom courses.

Contents

Deploying MPLS VPN in the Intranet Scenario

- As shown in the figure, customer X and customer Y have two sites each. The two sites need to be interconnected through MPLS VPN, which corresponds to VPNX and VPNY, respectively.
- Interconnection interfaces, AS numbers, and IP addresses are shown in the figure. CEs and PEs exchange routing information using the protocol or method shown in the figure.

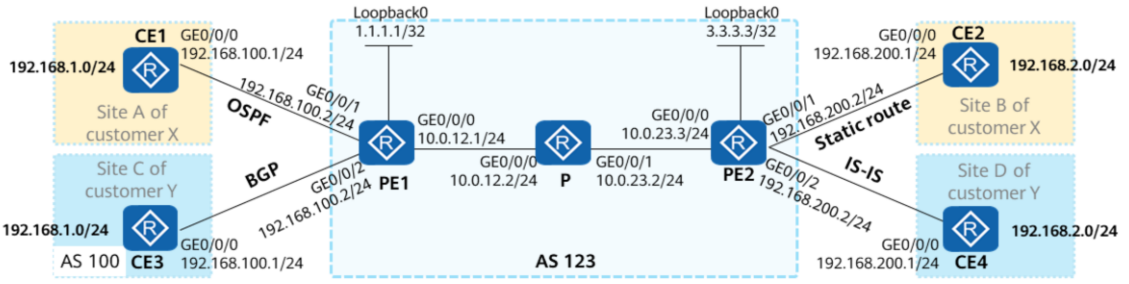

HUAWEI

- Note: This course describes only the non-inter-AS MPLS VPN deployment.

Deployment Roadmap

1. MPLS VPN backbone network configuration
   - 1.1 Configure an IGP to implement IP connectivity on the backbone network.
   - 1.2 Configure MPLS and MPLS Label Distribution Protocol (LDP) and establish an MPLS LSP public network tunnel to transmit VPN data.
   - 1.3 Configure MP-BGP to establish MP-BGP peer relationships for transmitting VPNv4 routes.
2. VPN user access configuration
   - 2.1 Create a VPN instance and set parameters (RTs and RDs).
   - 2.2 Bind an interface to the VPN instance.
   - **2.3 Configure route exchange between PEs and CEs.**

Deploying OSPF Between PEs and CEs (1)

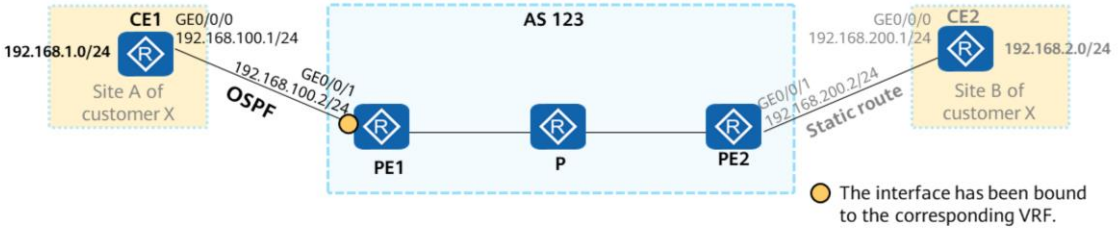


```
[CE1] ospf 1
[CE1-ospf-1] area 0
[CE1-ospf-1-area-0.0.0.0] network 192.168.100.0 0.0.0.255
[CE1-ospf-1-area-0.0.0.0] network 192.168.1.0 0.0.0.255
```

The OSPF configuration on CE1 is similar to the traditional OSPF configuration. CE1 does not need to support VRF.

Deploying OSPF Between PEs and CEs (2)

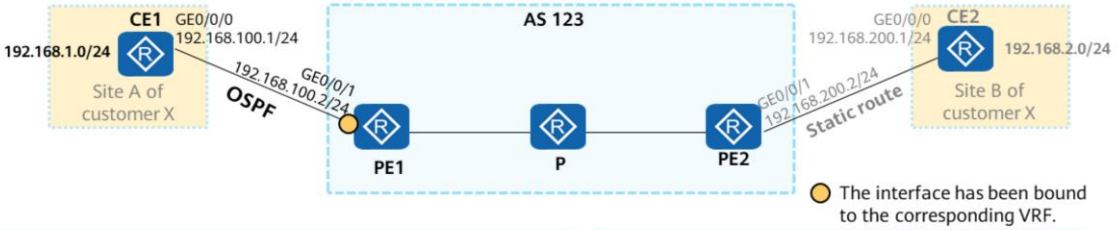


```
[PE1] ospf 1 vpn-instance VPNX
[PE1-ospf-1] area 0
[PE1-ospf-1-area-0.0.0.0] network 192.168.100.0 0.0.0.255
[PE1-ospf-1-area-0.0.0.0] quit
[PE1-ospf-1] import bgp
```

The OSPF process used by PE1 to interconnect with CE1 must be bound to the corresponding VPN instance.
Import the BGP routes in the routing table of VPNX on PE1 (mainly the customer routes learned by PE1 through BGP and destined for site B) into OSPF so that these routes can be advertised to CE1 through OSPF.

```
[PE1] bgp 123
[PE1-bgp] ipv4-family vpn-instance VPNX
[PE1-bgp] import-route ospf 1
```

Import the OSPF routes learned by OSPF process 1 in the routing table of VPNX on PE1 to BGP. Then, convert the customer routes destined for site A to BGP VPNv4 routes and advertise them to PE2.

Deploying Static Routes Between PEs and CEs

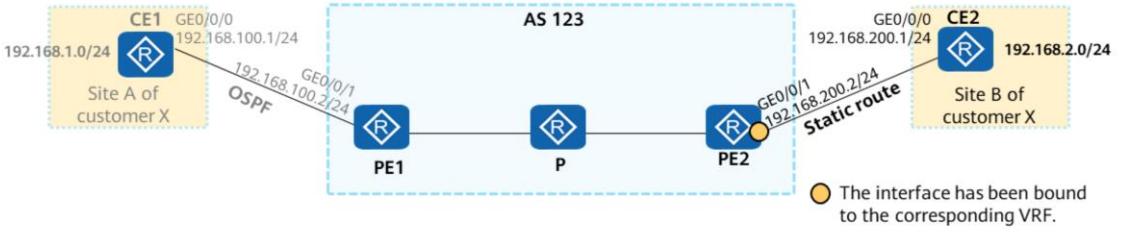

[CE2] **ip route-static** 192.168.1.0 24 192.168.200.2
[CE2] **ip route-static** 192.168.100.0 24 192.168.200.2

A static route to each network segment at site A needs to be configured on CE2.

[PE2] **ip route-static vpn-instance** VPNX 192.168.2.0 24 192.168.200.1

A static route to each network segment at site B needs to be configured on PE2.

[PE2] **bgp** 123
[PE2-bgp] **ipv4-family vpn-instance** VPNX
[PE2-bgp] **import-route static**

Import the static route in the routing table of VPNX on PE2 to BGP so that the static route can be converted into a BGP VPNv4 route and advertised to PE1.

Deploying EBGP Between PEs and CEs

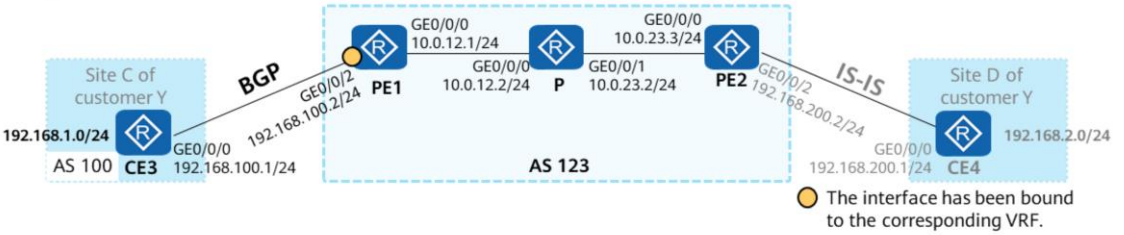


```
[CE3] bgp 100
[CE3-bgp] peer 192.168.100.2 as-number 123
[CE3-bgp] network 192.168.1.0 24
```

CE3 only needs to perform common BGP configurations and does not need to support VRF.

```
[PE1] bgp 123
[PE1-bgp] ipv4-family vpn-instance VPNY
[PE1-bgp-VPNY] peer 192.168.100.1 as-number 100
```

When a PE and a CE use BGP to exchange customer routes, you do not need to manually import routes on the PE. In this example, after PE1 learns a customer route from CE3 using BGP, PE1 automatically converts the route to a VPNv4 route and advertises the route to PE2. After PE1 learns the route to site D from PE2 using BGP, PE1 automatically converts the route to an IPv4 route and advertises the IPv4 route to CE3.

BGP Configuration in Special Scenarios — AS Number Replacement

In an MPLS VPN scenario, if EBGP runs between a PE and a CE to exchange routing information, the AS numbers of the two sites may be the same.

- If CE1 sends a VPN route to PE1 through EBGP and PE2 forwards the route to CE2, CE2 will discard the route due to repetitive AS numbers. As a result, site 1 and site 2 that belong to the same VPN cannot communicate with each other.
- You can run the **peer substitute-as** command on each PE to enable the AS number replacement function. That is, the PE replaces the AS number of the VPN site where the CE resides in the received private network route with the local AS number. The peer CE then does not discard the route with the repetitive AS number.

```
[PE1] bgp 123
[PE1-bgp] ipv4-family vpn-instance vpn1
[PE1-bgp-vpn1] peer 192.168.100.1 substitute-as
```

When sending a BGP route to CE1, PE1 replaces 65001 with the local AS number 123 if the AS_Path attribute contains 65001. Therefore, if a route is transmitted from CE2 to PE2 and then from PE2 to PE1, the AS_Path attribute of the BGP route is {123,123} when PE1 transmits the route to CE1.

BGP Configuration in Special Scenarios — SoO

In a CE multi-homing scenario, if BGP AS number replacement is enabled, routing loops may occur. Therefore, the site of origin (SoO) feature is required to prevent the routing loops.

- Both CE1 and CE3 belong to site 1. CE2 belongs to site 2. The AS numbers of sites 1 and 2 are both 65001. EBGP runs between PEs and CEs. To ensure that the PEs and CEs learn routes from each other, configure AS number replacement on PE1 and PE2.
- CE1 transmits an intra-site route to PE1, and PE1 transmits the route to CE3. Because AS number replacement is configured, CE3 receives the route, which may cause a routing loop.

```
[PE1] bgp 123
[PE1-bgp] ipv4-family vpn-instance vpn1
[PE1-bgp-vpn1] peer 192.168.100.1 soo 200:1
[PE1-bgp-vpn1] peer 192.168.200.1 soo 200:1
```

After the SoO attribute is configured for the BGP peer:

- When a BGP route is received from a peer, the SoO attribute is carried in the path attribute and advertised to other BGP peers.
- Before advertising a BGP route to a peer, the device checks whether the SoO attribute of the route is the same as the configured SoO value. If the SoO attribute of the route is the same as the configured SoO value, the device does not advertise the route to prevent loops.

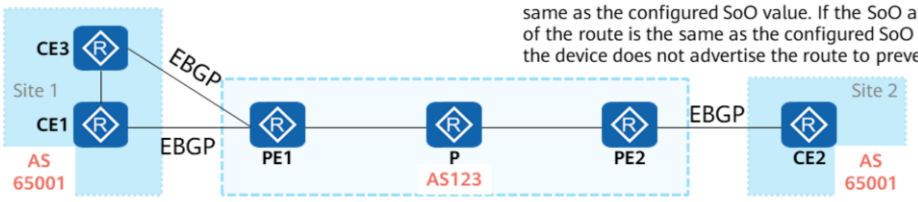

- Note: 192.168.100.1 and 192.168.200.1 are the IP addresses of the interfaces on CE1 and CE3, respectively, that are used to set up the BGP peer relationship with PE1.

Deploying IS-IS Between PEs and CEs

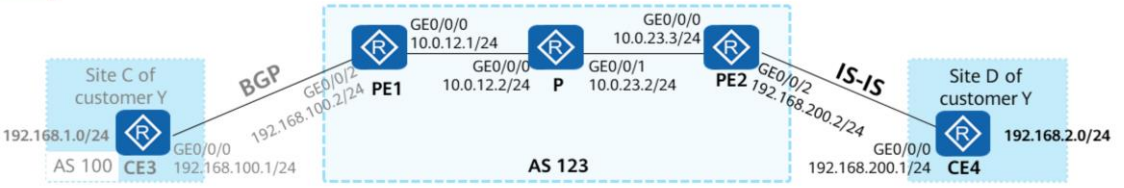

[CE4] **isis** 1
[CE4-isis-1] **network-entity 49.0001.0000.0000.1111.00**
[CE4-isis-1] **is-level level-2**
[CE4-isis-1] **quit**
[CE4] **interface GigabitEthernet** 0/0/0
[CE4-GigabitEthernet0/0/0] **isis enable** 1
[CE4-GigabitEthernet0/0/0]quit
[CE4] **interface GigabitEthernet** 0/0/1
[CE4-GigabitEthernet0/0/1] **isis enable** 1
GE0/0/1 is the interface connected to network segment 192.168.2.0/24.

[PE2] **isis** 1 **vpn-instance** VPNY
[PE2-isis-1] **network-entity 49.0002.0000.0000.2222.00**
[PE2-isis-1] **is-level level-2**
[PE2-isis-1] **import-route bgp level-2**
[PE2-isis-1] **quit**
[PE2] interface GigabitEthernet 0/0/2
[PE2-GigabitEthernet0/0/2] **isis enable** 1
[PE2] **bgp** 123
[PE2-bgp] **ipv4-family vpn-instance** VPNY
[PE2-bgp] **import-route isis 1**

Contents

Deploying MPLS VPN in the Hub&Spoke Scenario

- Hub&Spoke networking solutions are as follows:
  - Method 1: EBGP runs between the Hub-CE and Hub-PE, and between the Spoke-PE and Spoke-CE.
  - Method 2: An IGP runs between the Hub-CE and Hub-PE, and between the Spoke-PE and Spoke-CE.
  - Method 3: EBGP runs between the Hub-CE and Hub-PE, and an IGP runs between the Spoke-PE and Spoke-CE.
- The Hub-CE and Hub-PE cannot use an IGP when the Spoke-PE and Spoke-CE using EBGP to deploy the MPLS VPN in the Hub&Spoke networking.

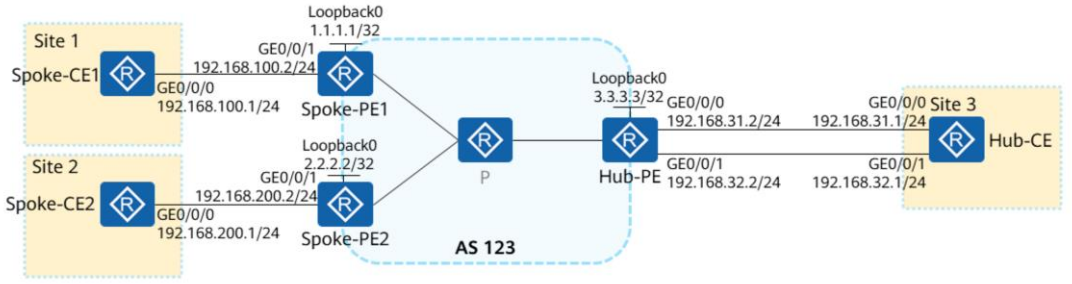

VRF Configuration

- Create a VPN instance on the Spoke-PE. The RT configuration is shown in the figure.
- Create VPN_in and VPN_out on the Hub-PE to import private network routes and export private network routes to the Spoke-PE, respectively. The RT configuration is shown in the figure.

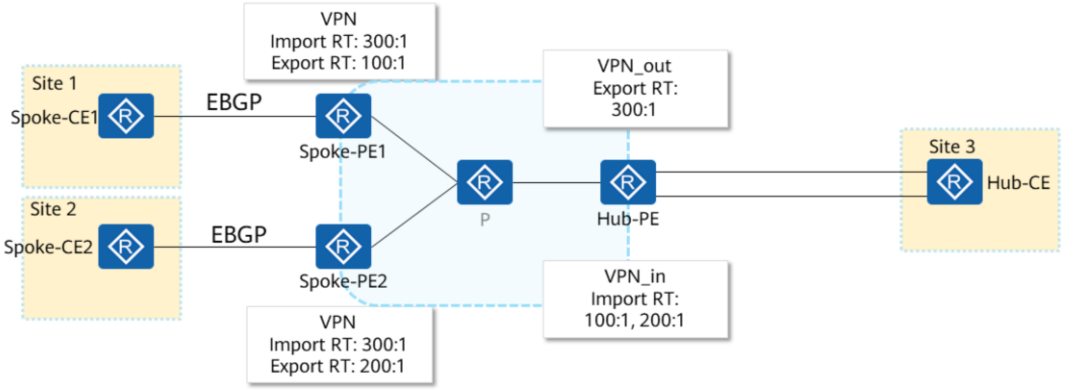

HUAWEI

Deployment Method 1 — Route Advertisement Process

- Spoke-CEs and Spoke-PEs exchange routing information through EBGP. After an EBGP connection is set up, Spoke-CEs and Spoke-PEs advertise routes to BGP.
- Two EBGP connections are set up between the Hub-PE and Hub-CE to separately advertise and accept private network routes.

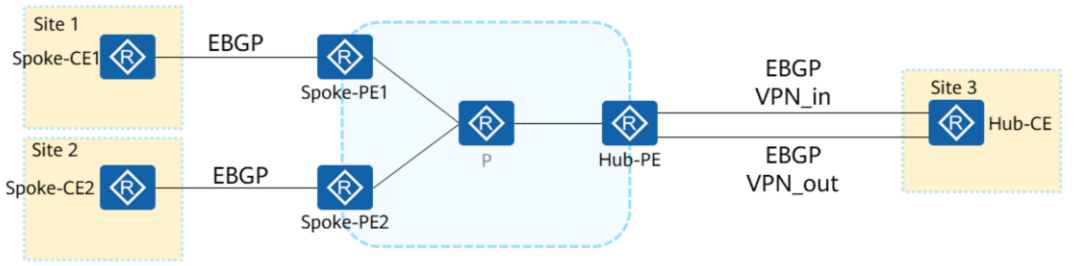

HUAWEI

- The process of advertising routes from Spoke-CE1 to Spoke-CE2 is as follows:
  1. Spoke-CE1 advertises a route to Spoke-PE1 through EBGP.
  2. Spoke-PE1 advertises the route to the Hub-PE through IBGP.
  3. The Hub-PE imports the route into the VPN_in routing table through the import RT attribute of the VPN instance (VPN_in), and then advertises the route to the Hub-CE through EBGP.
  4. The Hub-CE learns the route through the EBGP connection and advertises the route to the VPN instance (VPN_out) of the Hub-PE through another EBGP connection.
  5. The Hub-PE advertises the route with the export RT attribute of VPN_out to all Spoke-PEs.
  6. Spoke-PE2 advertises the route to Spoke-CE2 through EBGP.

Deployment Method 1 — Configuration Between Hub-PE and Hub-CE

- The Hub-PE advertises the routes learned from spoke sites to the hub site through the EBGP connection corresponding to VPN_in.
- Hub-CE advertises the routes to spoke sites through EBGP corresponding to VPN_out.

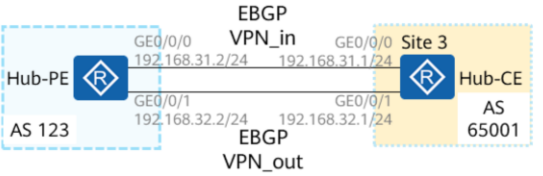

Establish two EBGP connections between the Hub-PE and Hub-CE.
[Hub-PE] bgp 123
[Hub-PE-bgp] ipv4-family vpn-instance VPN_in
[Hub-PE-bgp-VPN_in] peer 192.168.31.1 as-number 65001
[Hub-PE-bgp-VPN_in]quit
[Hub-PE-bgp] ipv4-family vpn-instance VPN_out
[Hub-PE-bgp-VPN_out] peer 192.168.32.1 as-number 65001
[Hub-PE-bgp-VPN_out] **peer 192.168.32.1** allow-as-loop

Set up two EBGP connections between the Hub-CE and Hub-PE.
[Hub-CE] **bgp 65001**
[Hub-CE-bgp] peer 192.168.31.2 as-number 123
[Hub-CE-bgp] peer 192.168.32.2 as-number 123

Because the routes that the Hub-CE sends to the Hub-PE through the EBGP connection corresponding to VPN_out may contain AS 123, these routes will be discarded by the Hub-PE. To prevent such a problem, the Hub-PE must be manually configured to allow repetitive local AS numbers.

Deployment Method 2 — Route Advertisement Process

- The following example uses OSPF as an IGP.
  - Spoke-CEs and Spoke-PEs exchange routing information through OSPF process 100.
  - The Hub-PE uses two OSPF processes to establish OSPF neighbor relationships with the Hub-CE, which separately advertise and accept private network routes.

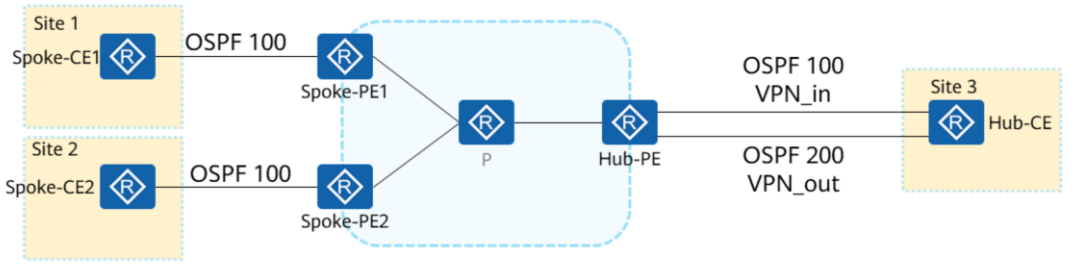

- The process of advertising routes from Spoke-CE1 to Spoke-CE2 is as follows:
  1. Spoke-CE1 advertises a route to Spoke-PE1 through OSPF 100.
  2. Spoke-PE1 advertises the route to the Hub-PE through IBGP.
  3. The Hub-PE imports the route to the VPN_in routing table through the import RT attribute of the VPN instance (VPN_in). After the BGP route is imported into OSPF 100, the route transmitted from Spoke-PE1 is advertised to the Hub-CE.
  4. The Hub-CE receives the route through OSPF 100. After route import is configured, the route is advertised to OSPF 200, and then OSPF 200 advertises the route to the Hub-PE.
  5. The VPN instance (VPN_out) of the Hub-PE imports the route of OSPF 200 multi-instance and advertises the route with the export RT attribute to all Spoke-PEs.
  6. Spoke-PE2 advertises the route to Spoke-CE2 through OSPF 100.

Deployment Method 2 — Configuration Between the Hub-PE and Hub-CE

- The Hub-PE advertises the routes learned from the spoke site to the hub site through OSPF process 100 corresponding to VPN_in.
- The Hub-CE advertises the route to the Hub-PE through OSPF (process 200) corresponding to VPN_out, and then advertises the route to all spoke sites.

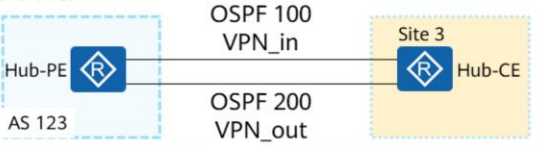


```
# Configure OSPF and BGP to import routes from each other on the Hub-PE.
[Hub-PE] OSPF 100 vpn-instance VPN_in
[Hub-PE-ospf-100]import-route bgp
[Hub-PE-ospf-100]quit
[Hub-PE]bgp 100
[Hub-PE-bgp]ipv4-family vpn-instance VPN_out
[Hub-PE-bgp-VPN_out]import-route ospf 200
```

```
# Import routes from OSPF 200 to OSPF 100 on the Hub-CE.
[Hub-CE]OSPF 200
[Hub-CE-ospf-200]import-route OSPF 100
```

Deployment Method 3 — Route Advertisement Process

- Use OSPF as an example. The Spoke-CEs and Spoke-PEs exchange routing information through an OSPF neighbor relationship (process 100).
- Establish two EBGP connections between the Hub-PE and Hub-CE to accept and advertise private network routes, respectively. The configurations of the Hub-PE and Hub-CE are similar to those in method 1.

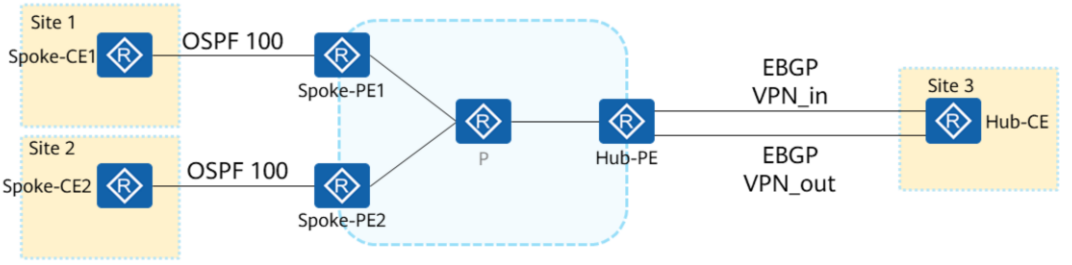

- The process of advertising routes from Spoke-CE1 to Spoke-CE2 is as follows:
  1. Spoke-CE1 advertises a route to Spoke-PE1 through OSPF 100.
  2. Spoke-PE1 advertises the route to the Hub-PE through IBGP.
  3. The Hub-PE imports the route into the VPN_in routing table through the import RT attribute of the VPN instance (VPN_in), and then advertises the route to the Hub-CE through EBGP.
  4. The Hub-CE learns the route through the EBGP connection and advertises the route to the VPN instance (VPN_out) of the Hub-PE through another EBGP connection.
  5. The Hub-PE advertises the route with the export RT attribute of VPN_out to Spoke-PE2.
  6. Spoke-PE2 advertises the route to Spoke-CE2 through OSPF 100.

Why Is There No Method 4?

- The Hub-CE and Hub-PE **cannot** use an IGP, when the Spoke-PE and Spoke-CE **cannot** use EBGP to deploy the MPLS VPN in the Hub&Spoke networking.

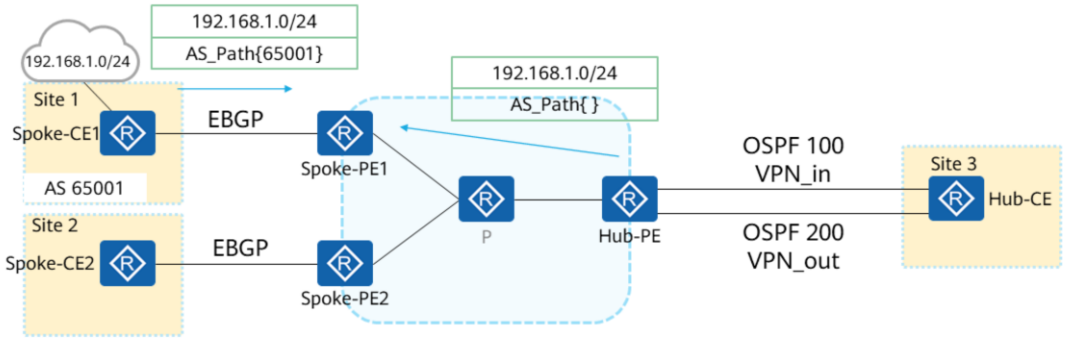

HUAWEI

- The following takes the advertisement of the route destined for 192.168.1.0/24 from Spoke-CE1 to Spoke-CE2 as an example. The process is as follows:
  1. Spoke-CE1 advertises a route to Spoke-PE1 through EBGP.
  2. Spoke-PE1 advertises the route to the Hub-PE through IBGP.
  3. The Hub-PE imports the route to the VPN_in routing table through the import RT attribute of the VPN instance (VPN_in) and advertises the route to the Hub-CE through OSPF 100.
  4. Hub-CE learns the route through OSPF 100 and advertises the route to the Hub-PE through OSPF 200.
  5. The VPN instance (VPN_out) of the Hub-PE imports the route of OSPF 200 and advertises the route with the export RT attribute of VPN_out to all Spoke-PEs.
  6. The VPN instance on Spoke-PE2 imports the route based on the import RT. Spoke-PE2 advertises the route to Spoke-CE2 through EBGP.
- The VPN instance (VPN_out) on the Hub-PE advertises the route to Spoke-PE2 and Spoke-PE1 at the same time with the export RT. The route is imported by the Hub-PE through an IGP (OSPF 200). Because the IGP route does not carry the AS_Path attribute, the AS_Path attribute is null. The AS_Path of the route destined for 192.168.1.0/24 from Spoke-CE1 is not null. Therefore, the route returned by the Hub-PE takes precedence over the route from Spoke-CE1. As a result, route flapping occurs.

Contents

OSPF/BGP in MPLS VPN

- When OSPF is deployed between the PE and CE to exchange routing information, if the standard BGP/OSPF exchange process (BGP/OSPF interoperation for short) is used on the PE to exchange routing information, the remote PE directly generates Type 5 LSAs when importing BGP into the OSPF process of the VPN instance. Different sites consider the routes of other sites as AS external routes (AS_external).
- To solve the problem of OSPF routing information loss caused by standard BGP/OSPF interoperability, BGP and OSPF are extended accordingly.

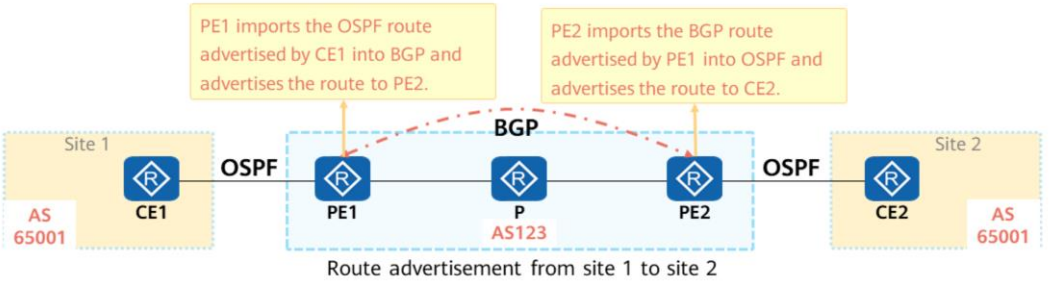

Route advertisement from site 1 to site 2

HUAWEI

- In actual applications, if two sites that need to communicate are in the same AS, each site should consider the route of the other site as an inter-area route rather than an AS external route.

BGP Extended Community Attributes

To retain OSPF routing information, BGP adds some community attributes that can carry OSPF routing information.

- Domain ID: identifies a domain.
- Route Type: contains the area ID and route type of the OSPF route imported to BGP.
  - Area-ID: ID of the VRF OSPF process of the PE that establishes an adjacency with a CE.
  - Route Type: type of the imported OSPF route:
    - 1 or 2: intra-area route, that is, the route calculated by the PE based on Type 1 and Type 2 LSAs
    - 3: inter-area route
    - 5: OSPF external route, that is, the route calculated by the PE through Type 5 LSAs. When the value of the Route Type field is 5, the value of the Area-ID field must be 0.0.0.0.
    - 7: NSSA route, that is, the route calculated by the PE through Type 7 LSAs

Domain ID

- When OSPF routes are imported into BGP on a PE, the PE adds the domain ID attribute to the BGP routes according to the local configuration. The domain ID is transmitted as the extended community attribute of BGP.
- When a PE imports a BGP route to OSPF, if the domain ID carried in the BGP route is the same as the local domain ID, the two sites belong to the same OSPF routing domain. If they are different, they are considered not in the same routing domain.

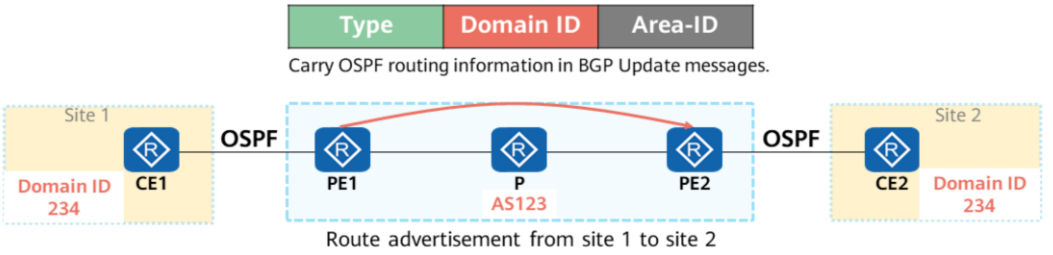

- The domain ID can be configured using the **domain-id** command in the view of the OSPF process bound to the VRF instance.
  - By default, the domain ID is 0 (NULL). If different OSPF domains use NULL as the domain ID, these OSPF domains cannot be distinguished. Consequently, the routes between different OSPF domains are considered as intra-area routes.
  - If an OSPF domain is configured with a non-zero domain ID, NULL is no longer the domain ID of the OSPF domain.
- It is recommended that all OSPF instances related to the same VPN use the same domain ID or the default domain ID.

Domain ID and Route Type

Based on the domain ID and route type in the BGP route, a PE generates different types of OSPF LSAs and advertises them to the OSPF process of the VRF.

Whether the domain ID is the same as the local domain ID	Route Type	Types of OSPF LSAs generated by PEs
Yes	1, 2, 3	3
	5, 7	5, 7
No	1, 2, 3, 5, 7	5, 7

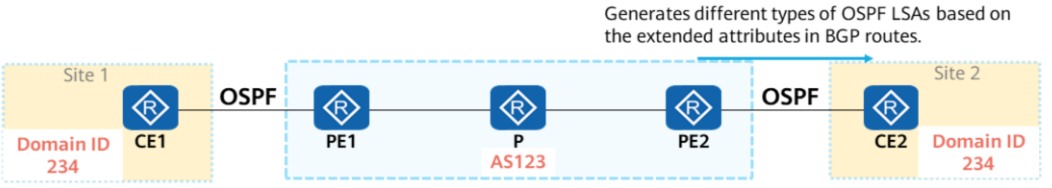

Route advertisement from site 1 to site 2

Contents

Type 3 Routing Loop Prevention — Case

The following figure shows an example of Type 3 LSA routing loops.

- Site 1 and site 2 belong to VPN1.
- Site 1 is connected to PE1 on the backbone network through OSPF area 0.
- Site 2 is connected to PE2 and PE3 on the backbone network through OSPF area 0 (in the dual-homing load balancing scenario).

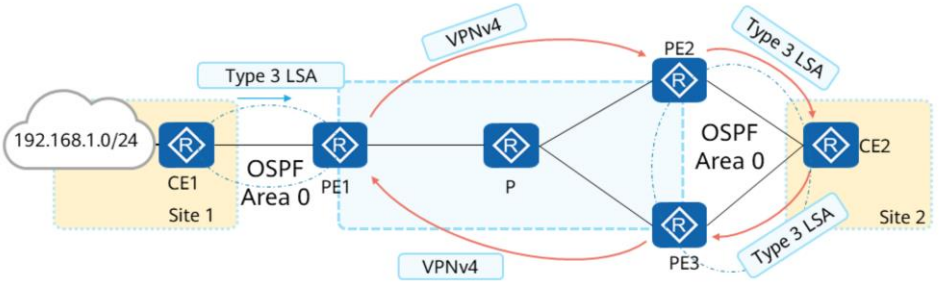

- The loop generation process is as follows:
  1. CE1 at site 1 advertises a route destined for 192.168.1.0/24 to PE1 using a Type 3 LSA.
  2. PE1 imports the route of the OSPF VPN1 process to BGP and advertises the route to PE2 and PE3 through MP-IBGP.
  3. PE2 is configured to import routes from BGP to OSPF. Therefore, PE2 generates Type 3 LSAs and sends them to CE2. CE2 then advertises the Type 3 LSAs received from PE2 to PE3.
  4. PE3 receives two routes destined for 192.168.1.0/24. One is advertised by PE1, and the other is imported by PE2. By default, IGP (OSPF) routes have a higher priority than IBGP routes. Therefore, PE3 selects OSPF routes.
  5. PE3 advertises the optimal route learned from OSPF to PE1 through MP-IBGP.

  - In this case, PE1 has two routes to the destination network segment 192.168.1.0/24. One is learned from CE1 through OSPF, and the other is learned from PE3 through MP-IBGP. The following problems may occur:
    - PE1 withdraws the route to 192.168.1.0/24. If the link between PE1 and PE2 is blocked, BGP Update messages cannot be sent to PE2. As a result, the route sent by PE3 to PE1 still exists. (In normal cases, the route is withdrawn when PE1 sends Update messages to PE2.) The next hop of the route from PE1 to 192.168.1.0/24 is PE3. Routing loops occur.
    - If the priority of the MP-IBGP route on PE1 is higher than that of the OSPF route, PE1 preferentially selects the BGP route advertised by PE3. In this case, PE1 needs to withdraw the BGP route advertised to PE2. As a result, PE3 withdraws the BGP route advertised to PE1, and PE1 preferentially selects the OSPF route again. As a result, route flapping occurs.

Type 3 Routing Loop Prevention — DN Bit

- To prevent Type 3 LSA loops, the OSPF multi-instance process uses an unused bit in the LSA Options field as a flag bit, which is called the DN bit. The DN bit is used to prevent Type 3 LSA loops.
- When performing SPF calculation, the OSPF instance process on a PE ignores Type 3 LSAs with the DN bit being 1.

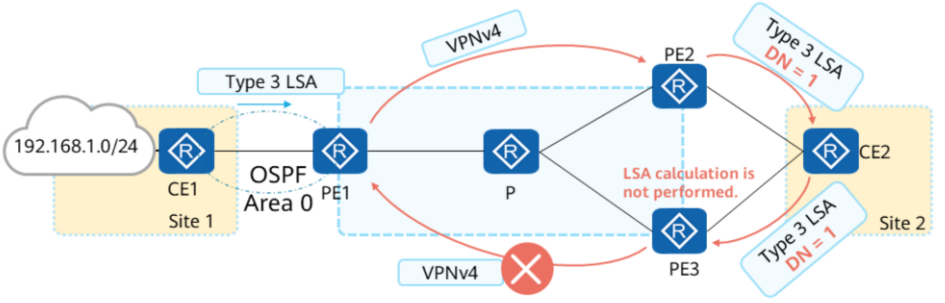

HUAWEI

- By default, the DN bit in LSAs generated by OSPF is set to 1. You can run the **dn-bit-set disable** command to disable OSPF from setting the DN bit in LSAs.

Type 5/7 Routing Loop Prevention — Case

The following figure shows an example of a Type 5 LSA routing loop.

- Site 1 and site 2 belong to VPN1.
- Site 1 is connected to PE1 on the backbone network through EBGP.
- Site 2 is connected to PE2 and PE3 on the backbone network through OSPF.

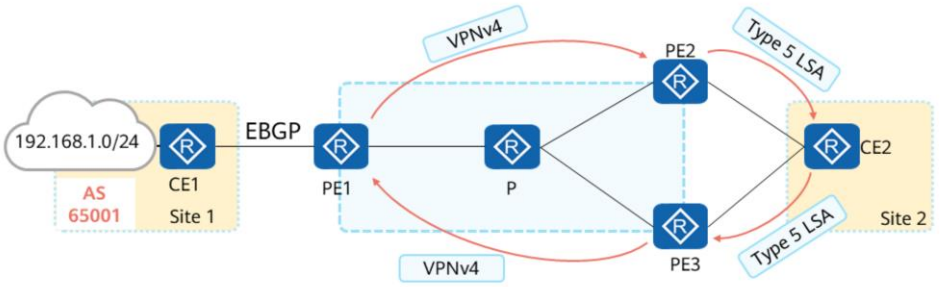

- The loop generation process is as follows:
  1. CE1 advertises a route destined for 192.168.1.0/24 to PE1 through EBGP. The AS_Path of the route is 65001.
  2. PE1 advertises the route to PE2 and PE3 through MP-IBGP.
  3. PE2 imports BGP routes to the OSPF VPN1 process and advertises a Type 5 LSA destined for 192.168.1.0/24 to CE2.
  4. CE2 advertises the Type5 LSA to PE3.
  5. PE3 preferentially selects an OSPF route (the OSPF route has a higher priority than the IBGP route) and advertises an Update message to PE1 through MP-IBGP.
  6. PE1 receives the MP-IBGP Update message from PE3. The MP-IBGP route advertised by PE3 has a higher priority than the EBGP route advertised by CE1 because the MP-IBGP route is an IGP (OSPF) route imported by BGP on PE3 and its AS_Path is null. PE1 prefers the route advertised by PE3.
  7. In this case, a routing loop is formed: PE3 -> CE2 -> PE2 -> PE1 -> PE3.
- Because PE1 does not preferentially select the route learned from CE1, PE1 withdraws the route advertised to PE2. The imported BGP route is also withdrawn in the OSPF VPN instance process on PE2. Then, both CE2 and PE3 withdraw the OSPF routes. The BGP route advertised by PE3 to PE1 is also withdrawn. On PE1, the route learned from CE1 becomes the optimal route. As a result, route flapping occurs.
- The generation and elimination of Type 7 LSA-related loops are similar to those of Type 5 LSA-related loops, and are not described here.

Type 5/7 Route Loop Prevention — VPN Route Tag

- A VPN route tag can be used to prevent Type 5 and Type 7 routing loops.
- When a PE generates Type 5 or Type 7 LSAs based on received BGP VPN routes, the LSAs carry VPN route flags. If a PE finds that the VPN route tag in an LSA is the same as the locally configured one, the PE ignores the LSA. This prevents routing loops.

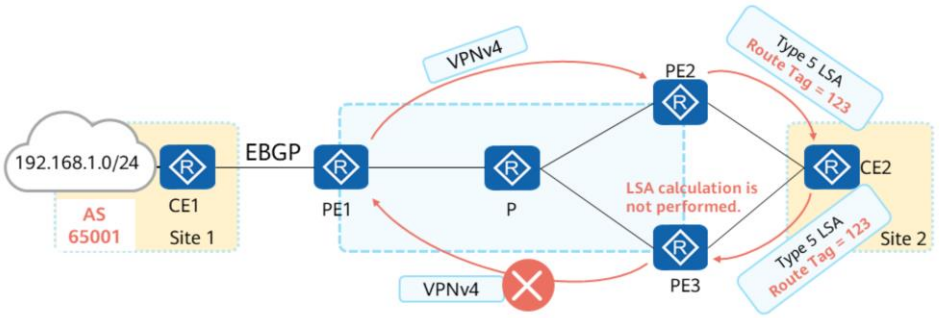

- The VPN route tag is not transmitted in the BGP extended community attribute. The VPN route tag is valid only on the PEs that receive BGP routes and generate OSPF LSAs.
- By default, the VPN route tag is calculated based on the AS number of BGP. If BGP is not configured, the default value is 0.
- You can run the **route-tag** command to set a VPN route tag.

Contents

HUAWEI

Sham Link Usage Scenarios

- Generally, BGP peers use BGP extended community attributes to carry routing information on the MPLS VPN backbone network. OSPF running on the peer PE can use the information to generate Type 3 LSAs from the PE to CE. These Type 3 LSAs are inter-area routes.
- If a backdoor link is added between CE1 and CE2 and OSPF is run to exchange routes, the route learned through the backdoor link is an intra-area route.
- Because intra-area routes take precedence over inter-area routes, the backdoor link is preferentially selected. To allow the backdoor link as a backup link, use the sham link.

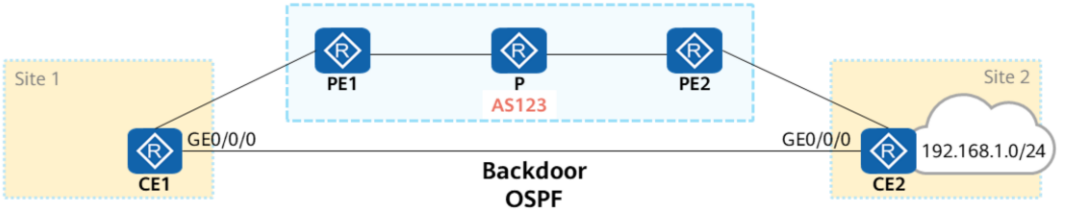

Working Mechanism of Sham Link

- The sham link creates an intra-area link between two PEs. When LSAs are flooded on a sham link, all OSPF route types remain unchanged and are not changed to Type 3 or 5 LSAs.
- A sham link is considered as a link between two VPN instances. The addresses of the two ends of the link are the addresses of the PEs, which are used as the source and destination addresses of the connection. The source and destination addresses of a sham link are loopback interface addresses with 32-bit masks. The loopback interface must be bound to a VPN instance and advertised through BGP.

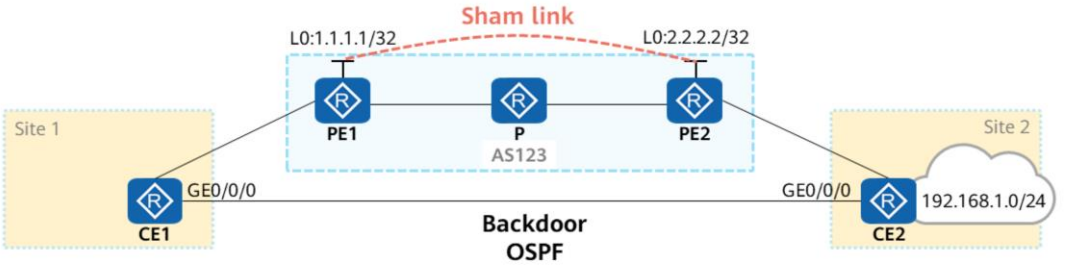

HUAWEI

- Multiple sham links of the same OSPF process can share the same endpoint address, but different OSPF processes cannot have two sham links with the same endpoint address.

Sham Link Configuration Example

1. Create an interface on the PE to set up a sham link. The configuration of PE2 is similar to that of PE1.

```
[PE1]interface LoopBack0
[PE1-LoopBack0] ip binding vpn-instance VPNA
[PE1-LoopBack0]ip address 1.1.1.1 32
# Advertise the routes in the BGP VPN address family.
[PE1-bgp-VPNA]network  1.1.1.1 32
```

3. Configure the sham link on PE1. The configuration of PE2 is similar to that of PE1.

```
[PE1-ospf-1]area 0
[PE1-ospf-1-area-0.0.0.0]sham-link 1.1.1.1 2.2.2.2
```

4. Adjust the cost value to ensure that the cost value of the backdoor link is greater than that of the sham link.

```
[CE1-GigabitEthernet0/0/0]ospf cost  1000
```


- When configuring a sham link, you can specify the route cost of the sham link. The default value is 1.

Verifying the Configuration of the Sham Link

1. For details about common OSPF VPN configurations, see the preceding configuration examples.

```
<PE1>display ospf sham-link  area  0

        OSPF Process 1 with Router ID 1.1.1.1

 Sham-Link: 1.1.1.1 --> 2.2.2.2
 Neighbor ID: 2.2.2.2 ,    State: Full ,    GR status: Normal
 Area: 0.0.0.0
 Cost: 1 , State: P-2-P , Type: Sham
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
```

2. Check OSPF routes on CE1. The command output shows that the peer route has been learned as an intra-area route.

```
<CE1>display  ospf routing

 Routing for Network
 Destination        Cost      Type       NextHop         AdvRouter       Area
 1.1.1.1/32         0         Stub       1.1.1.1         10.1.12.1       0.0.0.0
 10.1.12.0/24       1         Transit    10.1.12.1       10.1.12.1       0.0.0.0
 192.168.1.0/24     3         Stub       10.0.12.1       2.2.2.2         0.0.0.0
 10.1.23.0/24       3         Transit    10.1.12.2       10.1.23.2       0.0.0.0
```

Quiz

1. (Multiple) On an MPLS VPN network, when a PE imports VPN routes learned from other PEs to OSPF, which of the following LSAs may be generated? ( )
   A. Type 1 LSA
   B. Type 3 LSA
   C. Type 5 LSA
   D. Type 7 LSA
2. (TorF) When a CE transmits routes to a PE through BGP, the routes may carry the SoO attribute. ( )
   A. True
   B. False

1. BCD
2. B

Summary

- MPLS VPN has different networking solutions in different scenarios. The common networking solutions are intranet, extranet, and Hub&Spoke. In addition, MPLS VPN networking can be classified into inter-AS and intra-AS networking based on whether the MPLS VPN backbone network is an inter-AS network.
- PEs and CEs can use static, OSPF, IS-IS, or BGP routes to exchange routing information. OSPF provides the following extended features for MPLS VPN:
  - Domain ID: identifies whether the routes imported to a VPN instance belong to the same OSPF domain.
  - DN bit: used to prevent routing loops because of Type 3 LSAs.
  - VPN route tag: is used to prevent routing loops caused by Type 5 or Type 7 LSAs.
  - Sham link: controls OSPF route selection in special scenarios.


Thank You
www.huawei.com