BGP/MPLS IP VPN Basics

Foreword

- Virtual private network (VPN) is a technology that enables VPN users to communicate with each other over a public network.
- MPLS VPN is a Layer 3 VPN (L3VPN) technology widely used in the industry. Note that in this course, MPLS VPN refers to BGP/MPLS IP VPN.
- This course describes the basic concepts, working process, and typical configuration methods of MPLS VPN.

- Unless otherwise specified, MPLS VPN refers to BGP/MPLS IP VPN.

Objectives

- On completion of this course, you will be able to:
  - Describe the MPLS VPN model.
  - Understand basic concepts of MPLS VPN.
  - Describe the MPLS VPN route transmission and label distribution process.
  - Describe the MPLS VPN data forwarding process.
  - Perform basic MPLS VPN configurations.

Contents

MPLS VPN Definition

- BGP/MPLS IP VPNs are usually constructed by a carrier to provide VPN users with VPN services, and thus implementing route transmission and data forwarding between user networks.
- MPLS VPN uses BGP to advertise VPN routes and uses MPLS to forward VPN packets on the carrier's backbone network (IP network). BGP/MPLS IP VPN, short to MPLS VPN, is a common L3VPN technology.

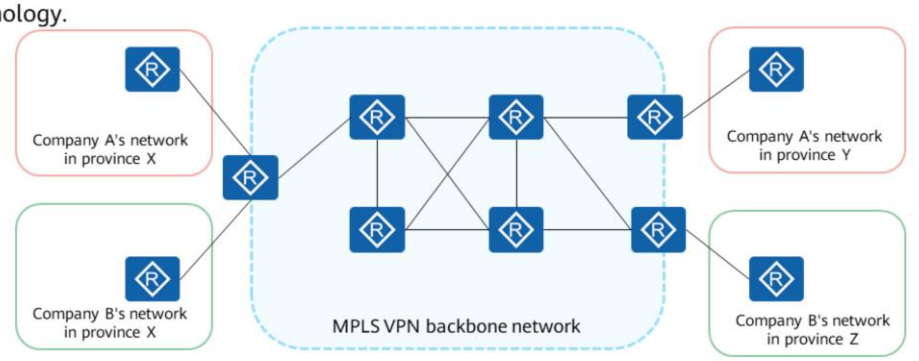

- MPLS VPN backbone networks can also be constructed by enterprises themselves, with technical implementation similarly to that of carriers. This course focuses on the scenario where enterprises purchase MPLS VPN services from carriers.

Network Architecture of MPLS VPN

- A MPLS VPN consists of three parts: customer edge (CE), provider edge (PE), and provider (P). The PE and P are carrier devices, whereas the CE is a MPLS VPN user device.
- A site represents a MPLS VPN user, consisting of CEs and other user devices.

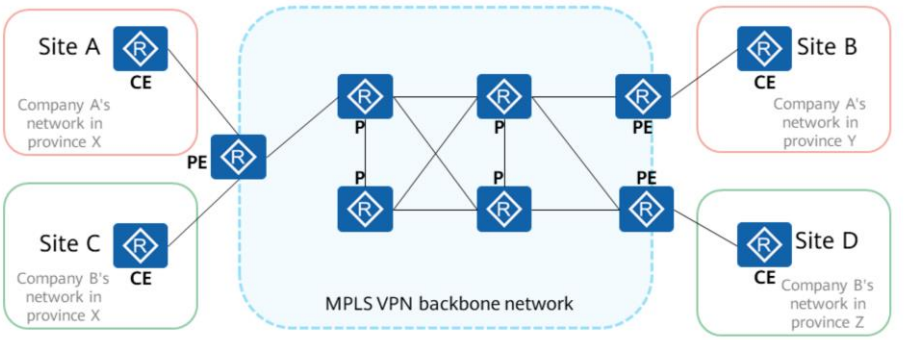

HUAWEI

- CE: an edge device on a user network. A CE provides interfaces that are directly connected to a carrier network. A CE can be a router, switch, or host. CEs are usually unaware VPNs and do not need to support MPLS.
- PE: an edge device on a carrier network and directly connected to a CE. On an MPLS network, PEs process all VPN services, and therefore PEs must have high performance.
- P: a backbone router on a carrier network and not directly connected to a CE. P devices need only to provide basic MPLS forwarding capabilities and do not maintain VPN information.
- The meaning of a site can be understood from the following aspects:
  - A site is a group of IP systems that can communicate without using carrier networks.
  - Sites are classified based on topological relationships between devices rather than the geographical locations of devices. In the preceding figure, the networks in provinces X and Y of company A need to communicate through the backbone network of the carrier. Therefore, the two networks are considered as two sites. If a physical private line is added between the CEs on the networks of provinces X and Y, the two networks can communicate without the need of the carrier network. In this case, the two networks are considered as one site.
- Relationship between sites and VPNs:
  - Sites connected to the same service provider network can be classified into different collections based on configured policies. Only sites that belong to the same collection can access each other, and this collection is defined as a VPN.
  - Devices at a site can belong to multiple VPNs. In other words, a site can belong to more than one VPN.

Technical Architecture of MPLS VPN

MPLS VPN is not a single VPN technology. Rather, it is a comprehensive solution that combines multiple technologies, including:

- MP-BGP: transmits site routing information between PEs.
- LDP: establishes tunnels between PEs.
- VRF: manages VPN users on PEs.
- Static route, IGP, or BGP: exchanges routing information between PEs and CEs.

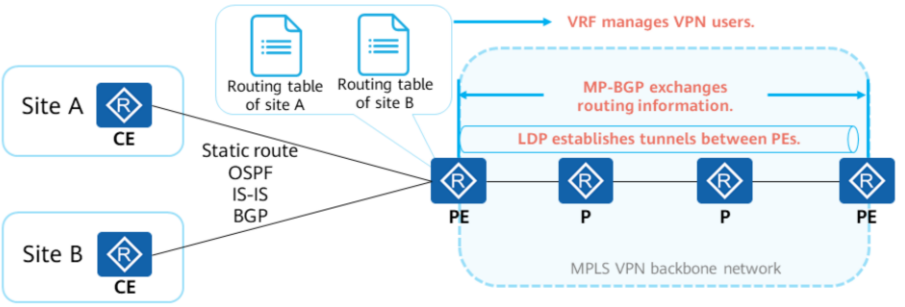

HUAWEI

- Multiprotocol Extensions for BGP (MP-BGP): an extended BGP protocol that supports multiple address families. For details, see related courses.

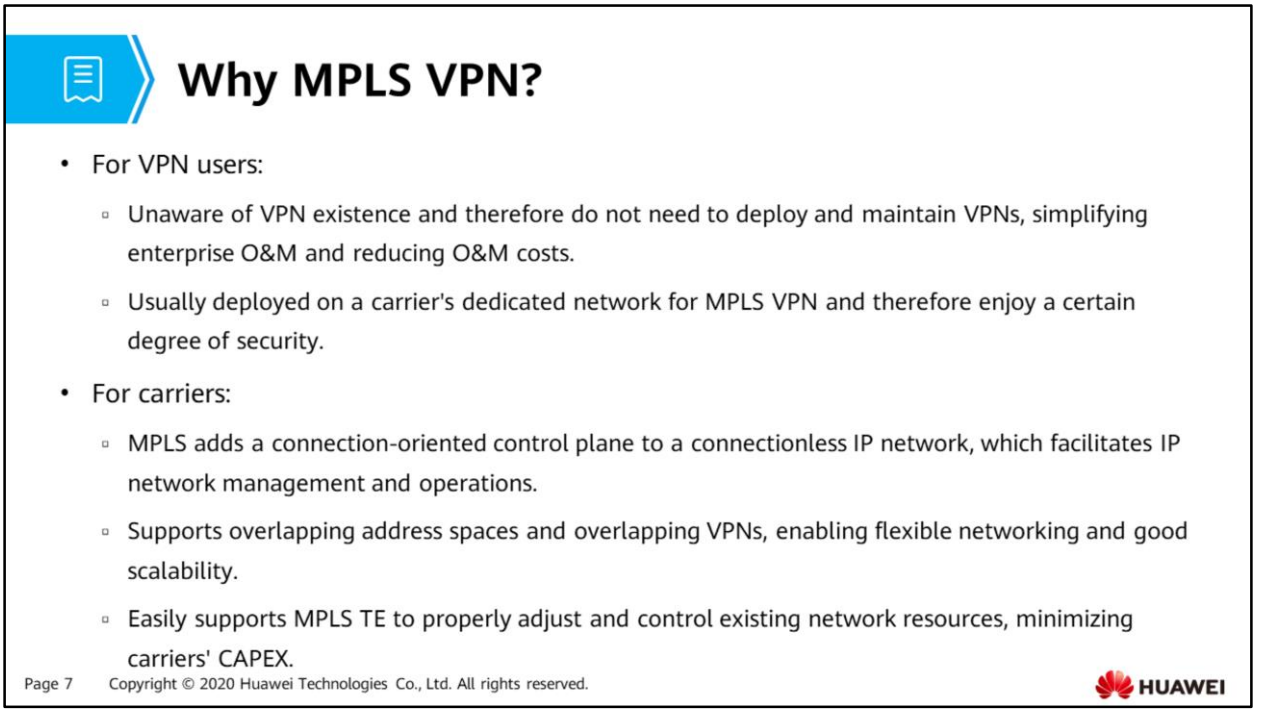

- MPLS traffic engineering (MPLS TE): steers traffic to constrained LSPs for forwarding so that the traffic is transmitted along specified paths. MPLS TE fully uses network resources and provides bandwidth and QoS guarantee without the need for hardware upgrades. It minimizes network costs.

Common MPLS VPN Networking Schemes

The following common networking schemes are available to meet different requirements of VPN users:

- Intranet: All users on a VPN are in a group isolated from users of other VPNs. Users at the same VPN site can communicate with each other, but users in different VPN sites cannot.
- Extranet: applies to the scenario where a VPN user wants to provide some resources of the local site for users of other VPNs.
- Hub&Spoke: can be used if a carrier wants to use a central access control device on a VPN to control the mutual access of users.

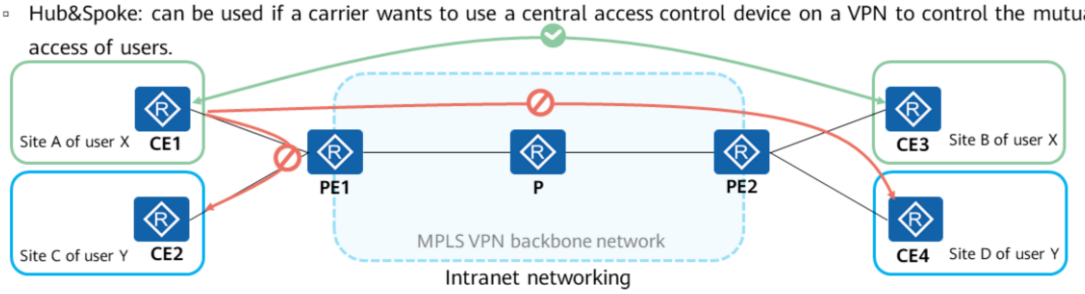

Intranet networking

HUAWEI

- Intranet networking is the simplest and most typical MPLS VPN networking scheme. The following technical implementation of MPLS VPN will be described based on this networking scheme.

Contents

MPLS VPN Route Advertisement

Route exchange must be completed between different sites to allow different sites on the same VPN to communicate with each other. On a basic MPLS VPN, only CEs and PEs are involved in VPN route advertisement. P routers only need to maintain the routes of the backbone network, without the need to know any VPN routes. VPN route advertisement involves the following three phases:

- From the local CE to the ingress PE
- From the ingress PE to the egress PE
- From the egress PE to the remote CE

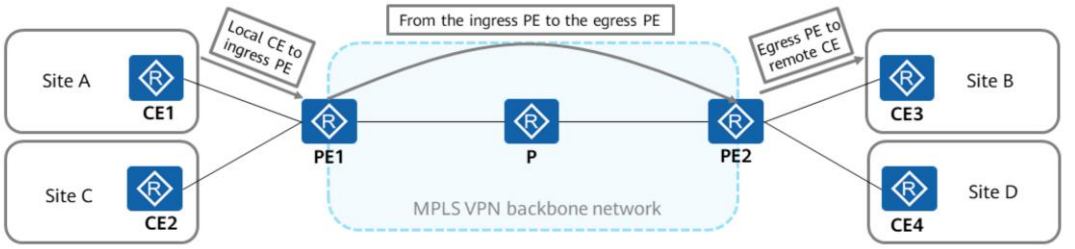

Route Exchange Between CEs and PEs

- As shown in the figure, customers X and Y belong to different VPNs, and each of them has two sites. Routes need to be exchanged between the two sites (belong to the same VPN).
- The CEs and PEs can use static routes, OSPF, IS-IS, or BGP to exchange routes. No matter which routing protocol is used, CEs and PEs exchange standard IPv4 routes.
- The local CEs exchange routes with the ingress PE in the same way as the egress PE exchange routes with the remote CEs.

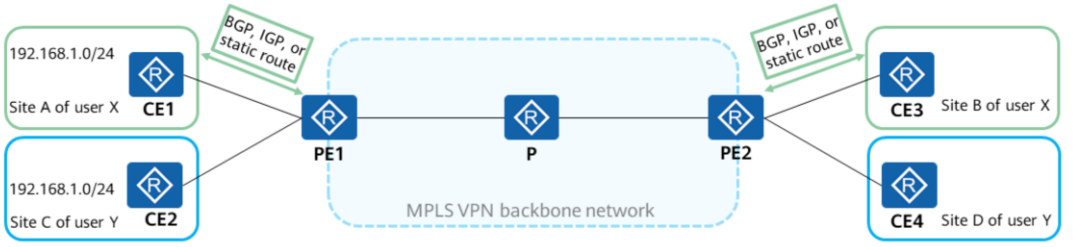

HUAWEI

Route Distribution from the Ingress PE to the Egress PE (1)

After receiving routes from CEs, PEs need to independently store routes of different VPNs and solve the problem where different customers use overlapping IP address spaces.

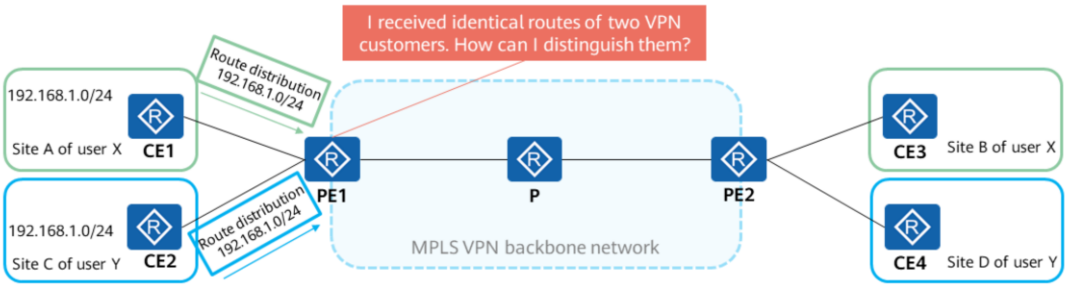

HUAWEI

- A VPN is a private network. Different VPNs independently manage their own address ranges, which are also called address spaces. Address spaces of different VPNs may overlap. For example, in the preceding figure, both user X and user Y use 192.168.1.0/24, indicating that the address spaces overlap. VPNs can use overlapping address spaces in the following situations:
  - They do not share the same site.
  - They share a same site, but devices at the site do not communicate with devices using overlapping address spaces at the other sites of the VPNs.

VRF

Virtual routing and forwarding (VRF), also called VPN instance, is a key technology in the MPLS VPN architecture. Each VRF uses independent routing and forwarding entries to logically isolate VPNs.

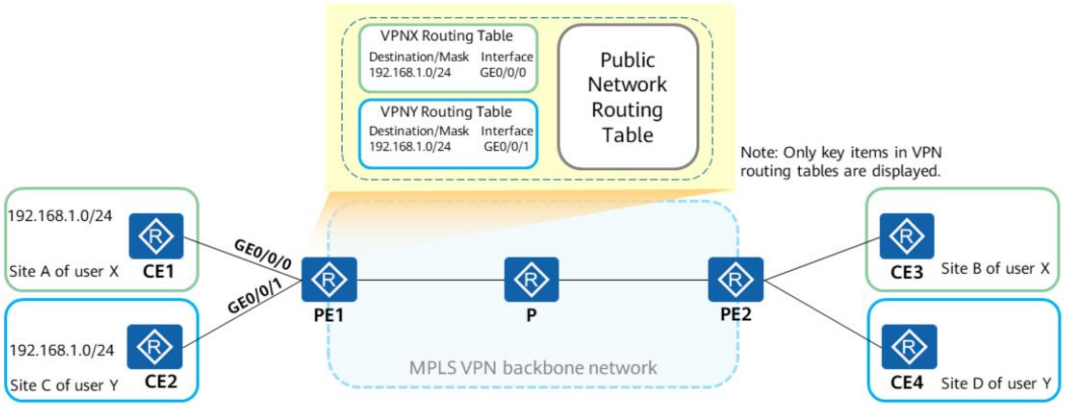

HUAWEI

- For details about VRFs, see the related HCIP-Datacom-Core course.

RD

After a PE receives IPv4 address prefixes from CEs in different VPNs, the PE differentiates these address prefixes based on VPN instance configurations. However, VPN instances are only locally significant. A PE cannot transmit VPN instance information to the peer PE. The route distinguisher (RD) is introduced to resolve this issue.

- An RD is 8 bytes long and used to distinguish IPv4 prefixes with the same address space.
- After a PE receives an IPv4 route from a CE, the PE adds an RD to the IPv4 prefix of the route and converts the route to a globally unique VPN-IPv4 route.

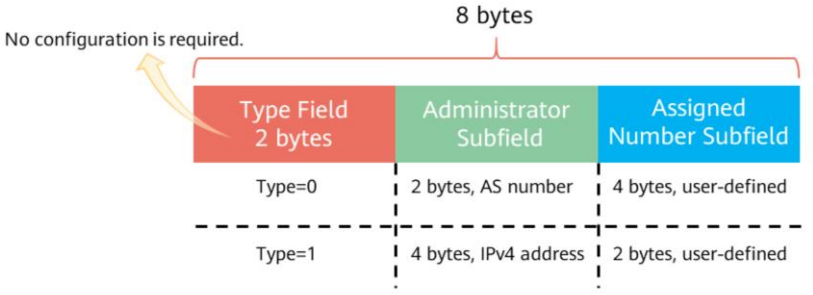

- When configuring an RD, you need to specify only the Administrator and Assigned Number subfields in the RD.
- Four types of RD configuration formats are available. The following two types are commonly used:
  - 16-bit AS number:32-bit user-defined number (for example 100:1)
  - 32-bit IPv4 address:16-bit user-defined number (for example, 172.1.1.1:1)
- The RD structure enables each carrier to allocate RDs independently. In some application scenarios, however, RDs must be globally unique to ensure normal routing.

VPN-IPv4 Address

A VPN-IPv4 address — also called a VPNv4 address — consists of 12 bytes, including an 8-byte RD and a 4-byte IPv4 address prefix.

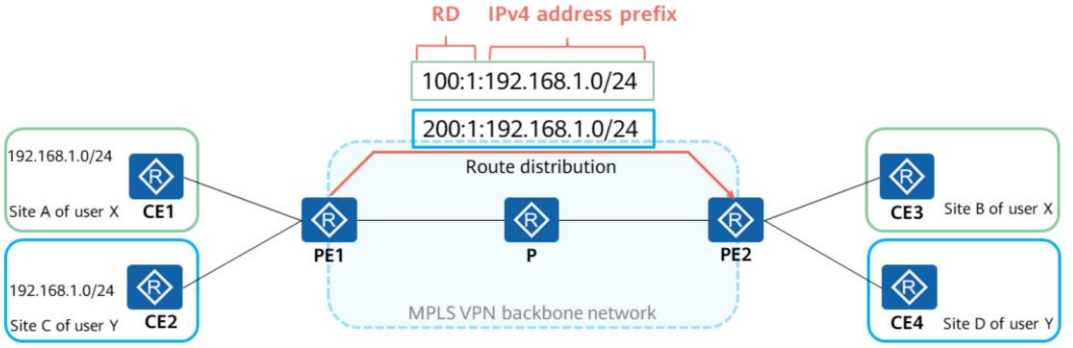

Route Distribution from the Ingress PE to the Egress PE (2)

- PEs establish BGP peer relationships and use BGP to exchange routes with each other. Why is BGP used?
  - BGP uses TCP as its transport layer protocol to improve reliability. It enables PEs that are not directly connected to directly exchange routes.
  - BGP is easy to expand, facilitating VPN route distribution between PEs.
  - The number of routes to be exchanged between PEs may be large. BGP sends only updated routes, enabling more routes to be exchanged while conserving link bandwidth resources.
- Traditional BGP-4 cannot process VPNv4 routes.

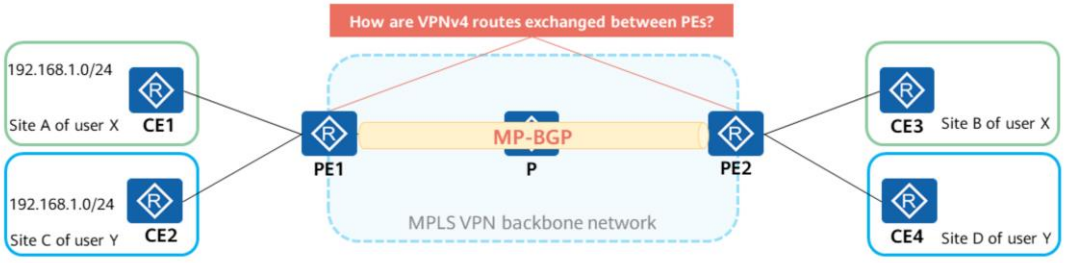

MP-BGP

- MPLS VPN uses MP-BGP defined in RFC 2858 (*Multiprotocol Extensions for BGP-4*) to correctly process VPN routes.
- MP-BGP uses address families to differentiate network layer protocols. An address family can be a traditional IPv4 address family or any other address family, such as a VPN-IPv4 address family or an IPv6 address family.
- The following two path attributes are added for MP-BGP:
  - MP_REACH_NLRI: Multiprotocol Reachable NLRI, used to advertise reachable routes and next hop information.
  - MP_UNREACH_NLRI: Multiprotocol Unreachable NLRI used to withdraw unreachable routes.

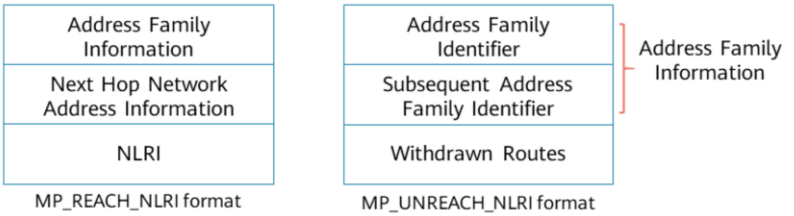

- NLRI: Network Layer Reachability Information
- For the values of address families, see RFC 3232 (*Assigned Numbers*).
- MP_REACH_NLRI is used to advertise reachable routes and next hop information. It consists of one or more 3-tuples <Address Family Information, Next Hop Network Address Information, NLRI>.
  - Address Family Information: consists of a 2-byte Address Family Identifier (AFI) and a 1-byte Subsequent Address Family Identifier (SAFI).
    - The AFI identifies the network layer protocol, corresponding to the address family value defined by "Address Family Number" in RFC 3232. For example, 1 indicates IPv4 and 2 indicates IPv6.
    - The SAFI indicates the NLRI type. If the AFI value is 1 and the SAFI value is 128, the address in the NLRI is an MPLS-labeled VPN-IPv4 address.
  - Next Hop Network Address Information: consists of the 1-byte length of the next hop network address and the variable-length next hop network address.
  - NLRI: consists of one or more 3-tuples <length, label, prefix>. This part will be described in detail in the following slides.
- MP_UNREACH_NLRI is used to instruct a peer to delete unreachable routes. The format of this attribute is as follows:
  - AFI: same as that in the MP_REACH_NLRI attribute
  - SAFI: NLRI type, same as that in the MP_REACH_NLRI attribute
  - Withdrawn Routes: an unreachable route list, consisting of one or more NLRI fields A BGP speaker can withdraw a route by adding the NLRI same as that in a previously advertised reachable route to the Withdrawn Routes field.

Route Distribution from the Ingress PE to the Egress PE (3)

- After MP-BGP distributes VPNv4 routes to the remote PE, the remote PE needs to import the VPNv4 routes to the correct VPN instances.
- MPLS VPN uses a BGP extended community attribute — VPN target, or route target — to control the export and import of VPN routes.
- The local PE adds RTs to its VPNv4 routes before advertising the routes. After receiving the VPNv4 routes, the remote PE imports the routes to the corresponding VPN instances according to the RTs.

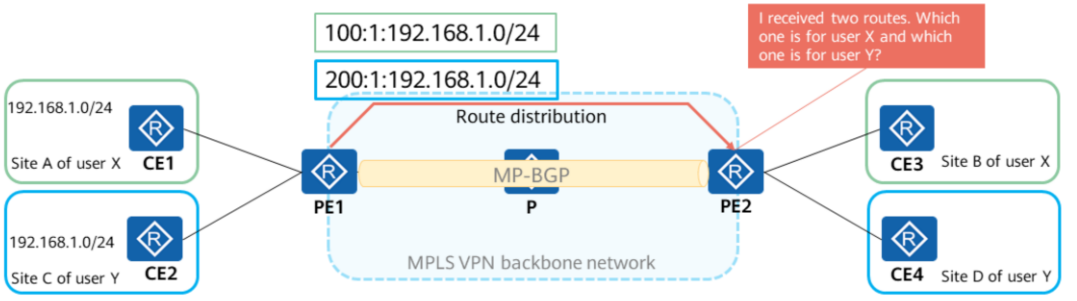

HUAWEI

RT

On a PE, each VPN instance is bound to one or more VPN targets. There are two types of VPN targets:

- Export target (ERT): After learning IPv4 routes from directly connected sites, the local PE converts the routes to VPN-IPv4 routes and adds the ERT attribute to these routes. The ERT attribute, as an extended community attribute of BGP, is distributed with these routes.
- Import target (IRT): After receiving a VPN-IPv4 route distributed by another PE, the PE checks the ERT attribute of the route. If the ERT is identical with the IRT of a VPN instance on the PE, the PE adds the route to the routing table of the VPN instance.

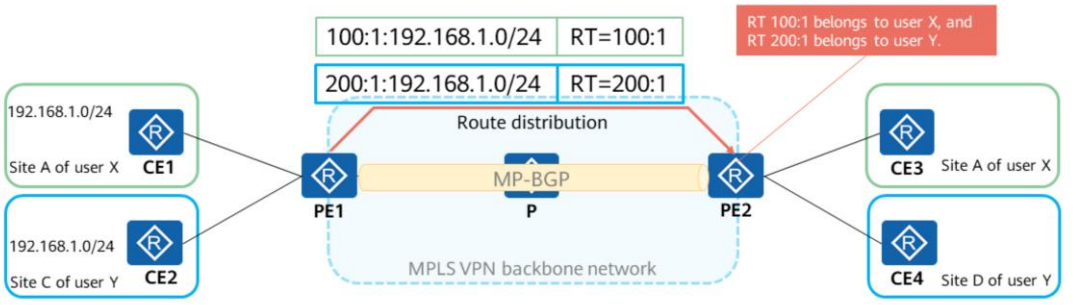

- Similar to an RD, an RT consists of three fields: Type, Administrator, and Assigned Number. The length of an RT is also 8 bytes.
- When configuring a VPN target, you need to specify only the Administrator and Assigned Number subfields in the VPN target. VPN targets have the same configuration formats as RDs.

Route Distribution from the Ingress PE to the Egress PE (4)

- After the egress PE imports a VPNv4 route into a correct VPN instance based on the RT carried in the VPNv4 route, the PE removes the RD from the VPNv4 route and distributes the **IPv4 route** to the corresponding CE.
- The CEs at site B and site D can learn the routes destined for their respective remote sites. Different sites of the same user (in the same VPN) can communicate with each other through similar operations.

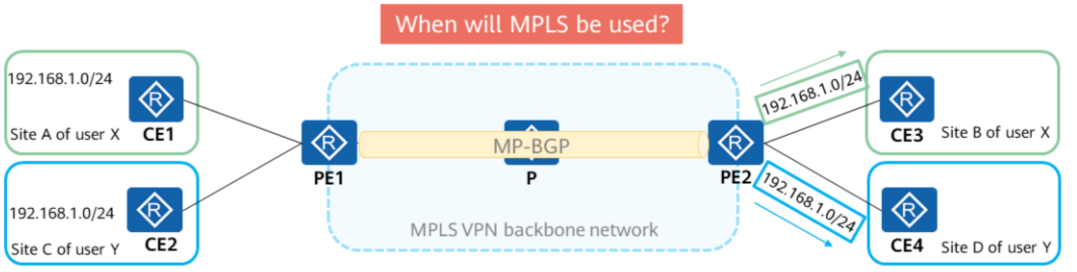

HUAWEI

Problems Encountered During Data Forwarding

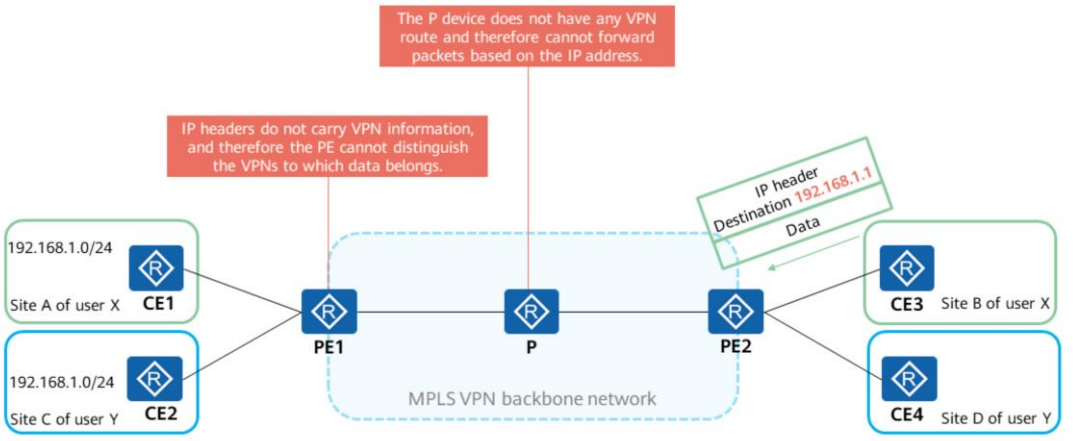

HUAWEI

Solving Problems Through Labels

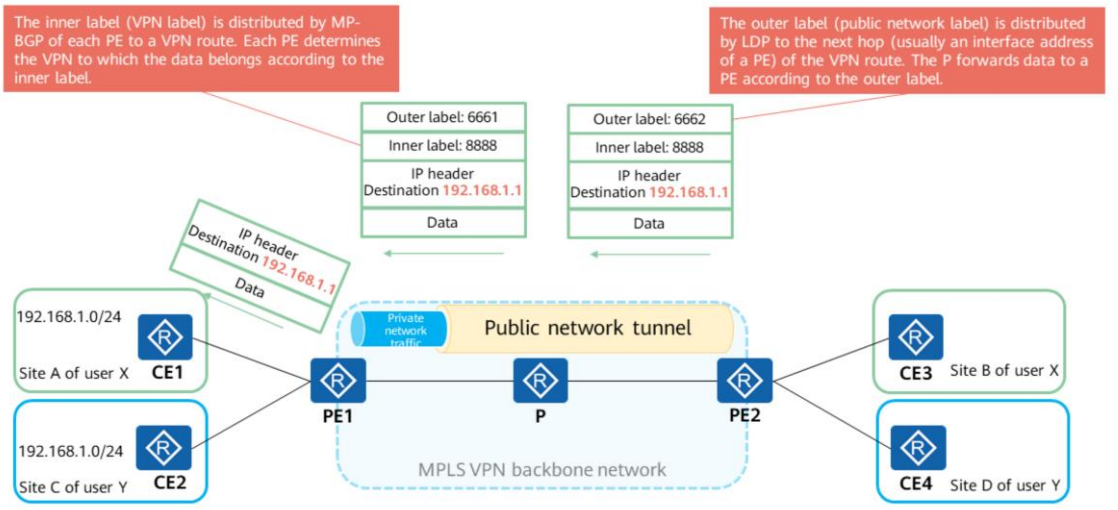

HUAWEI

Route Distribution from the Ingress PE to the Egress PE (5)

- The PEs and P distribute public network labels through LDP to establish LSPs (public network tunnels) between PEs.
- When transmitting VPNv4 routes through MP-BGP, the ingress PE carries VPN labels to differentiate data of different VPNs.
- After receiving VPNv4 routes, the egress PE performs VPN route leaking and tunnel recursion to select routes.

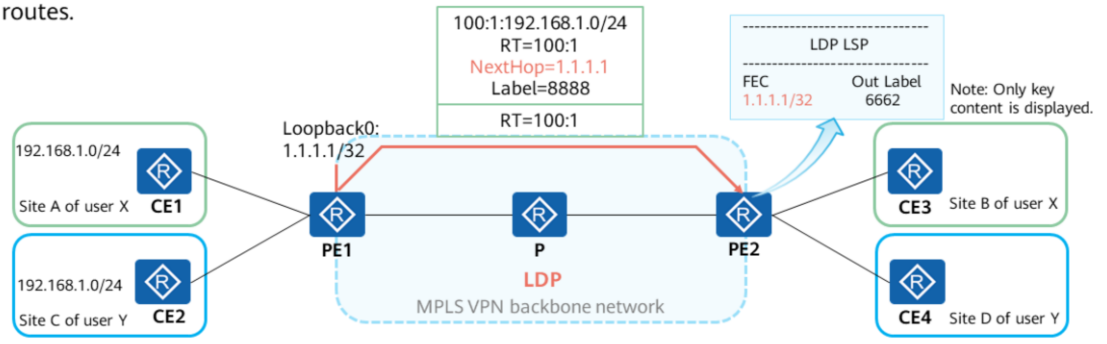

HUAWEI

- A PE device distributes MPLS labels in either of the following ways:
  - One label per route: Each route in a VRF is assigned one label. When many routes exist on the network, the Incoming Label Map (ILM) maintains these entries, requiring high router capacity.
  - One label per instance: Each VPN instance is assigned one label. All the routes of a VPN instance share the same label, reducing the number of labels required.
- VPN route leaking: a process of matching VPNv4 routes against the VPN targets of local VPN instances. After a PE receives a VPNv4 route, the PE directly matches the route against the VPN targets of local VPN instances, without selecting the optimal route or checking whether a desired tunnel exists.
- Tunnel recursion: A public network tunnel is required to transmit VPN traffic from one PE to the other PE over the public network. After VPN route leaking, the route must be successfully recursed to an LSP based on the destination IPv4 prefix before the route is added to the routing table of the corresponding VPN instance. This means that the next hop of the IPv4 route must match an LSP.

Route Exchange Process on a MPLS VPN

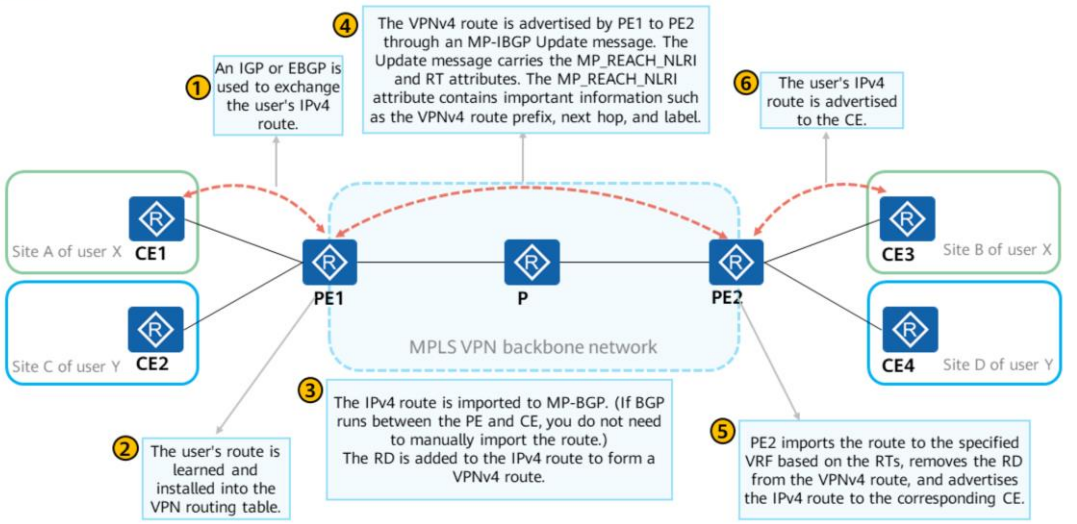

HUAWEI

Section Summary

On a MPLS VPN, VPN routes need to be transmitted between PEs and CEs and between PEs.

- PEs and CEs can use BGP, IGP, or static routes to exchange IPv4 routes.
- PEs use MP-BGP to exchange VPNv4 routing information, including:
  - RD: forms a VPNv4 prefix together with an IPv4 prefix.
  - RT: controls route import and export between PEs.
  - Label: inner label (VPN label), used to differentiate data of different VPNs on a PE during data forwarding.

Contents

Packet Forwarding Process (1)

The following example assumes that site B of user X wants to access 192.168.1.0/24 of site A. The packet forwarding process is as follows:

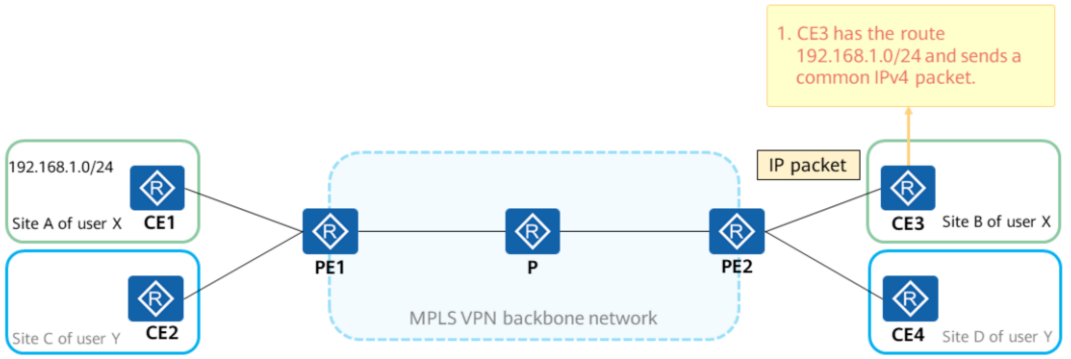

Packet Forwarding Process (2)

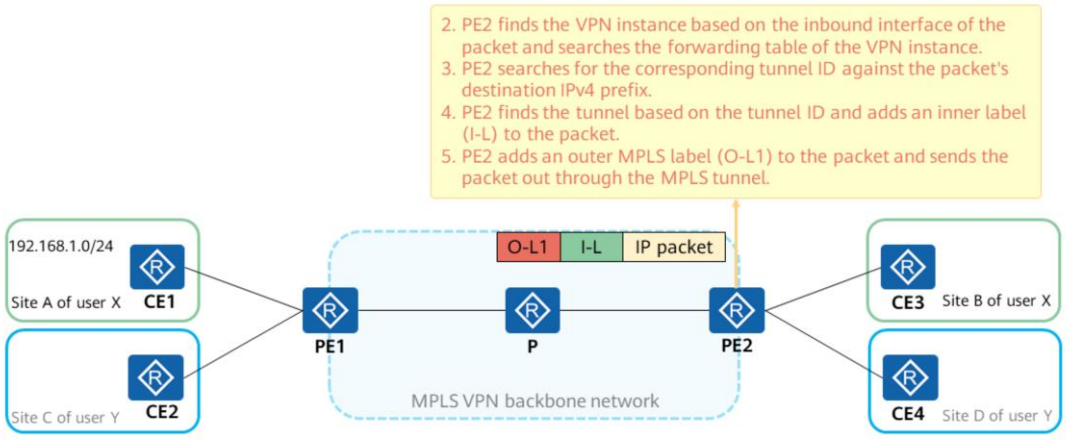

Packet Forwarding Process (3)

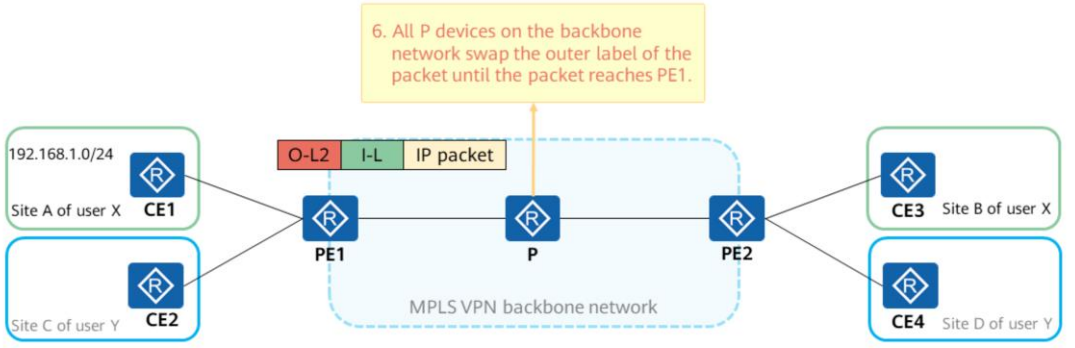

Packet Forwarding Process (4)

7. After receiving the packet, PE1 sends the packet to the MPLS module for processing. The MPLS module removes the outer label.
8. PE1 continues to process the inner label. Specifically, it determines the next hop based on the inner label, removes the inner label, and sends the native IPv4 packet to CE1.

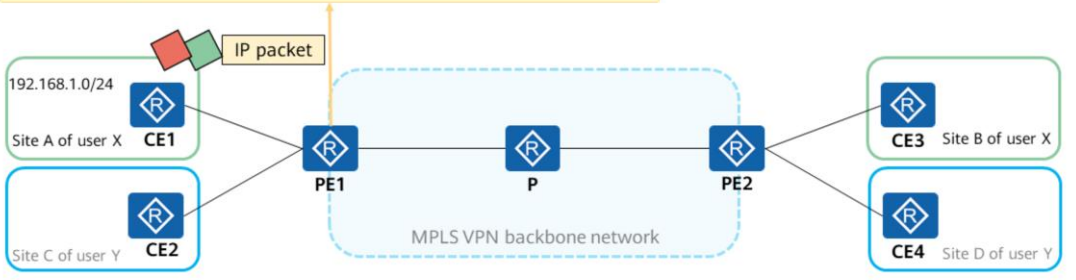

Packet Forwarding Process (5)

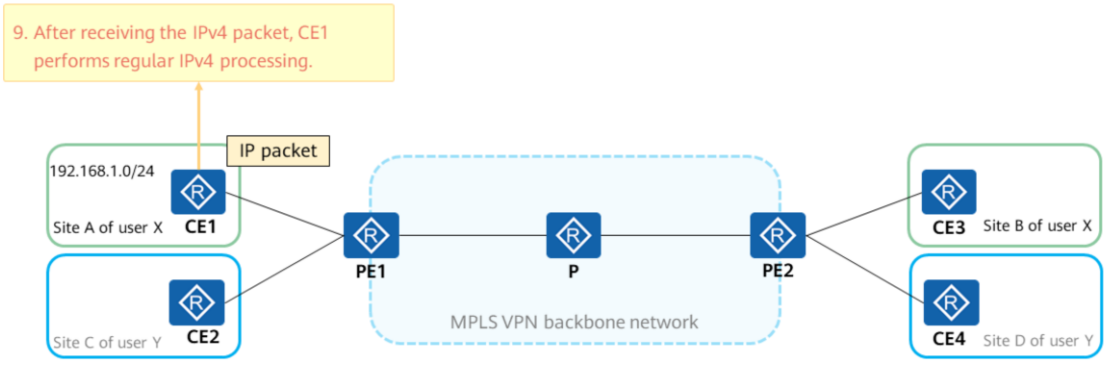

HUAWEI

Contents

Configuration Commands — VPN Instance Configuration (1)

1. Create a VPN instance or enter the VPN instance view.

   [PE] **ip vpn-instance** *vpn-instance-name*

   By default, no VPN instance is configured.
2. Enable the VPN instance IPv4 address family or enter the VPN instance IPv4 address family view.

   [PE-vpn-instance-InstanceName] **ipv4-family**

   By default, the VPN instance IPv4 address family is disabled.
3. Configure an RD for the VPN instance address family.

   [PE-vpn-instance-InstanceName] **route-distinguisher** *route-distinguisher*

   The common formats of RDs are as follows:
   - 2-byte AS number:4-byte user-defined number, for example, 100:1.
   - IPv4 address:2-byte user-defined number, for example, 192.168.122.15:1.

   The RD value must be globally unique regardless of the format.

   An RD configured for a VPN instance address family cannot be modified or deleted. To modify the RD, you need to first disable the corresponding VPN instance address family, or delete the VPN instance and then reconfigure one.

Configuration Commands — VPN Instance Configuration (2)

4. Configure VPN targets for the VPN instance.

   [PE-vpn-instance-InstanceName] **vpn-target** *vpn-target*&<1-8> [**both** | **export-extcommunity** | **import-extcommunity**]

   The **vpn-target** command configures export or import VPN targets for the VPN instance address family.

   - The format of a VPN target is the same as that of an RD.
   - A maximum of eight VPN targets can be configured at a time using the **vpn-target** command. To configure more VPN targets, run the **vpn-target** command multiple times.

5. Bind an interface to the VPN instance.

   [PE-GigabitEthernet0/0/0]**ip binding vpn-instance** *vpn-instance-name*

   The **ip binding vpn-instance** command binds an interface on a PE to a VPN instance. By default, an interface belongs to the root instance and is not bound to any VPN instance. Binding an interface to a VPN instance or deleting the binding will result in the deletion of the IP address of the interface, Layer 3 features, and IP-related routing protocols. These features must be re-configured if needed.

Configuration Commands — MP-BGP Configuration

1. Configure basic BGP functions.

   [PE] **bgp** { *as-number-plain*| *as-number-dot*}

   [PE-bgp] **peer** *ipv4-address* **as-number** *as-number*

   [PE-bgp] **peer** *ipv4-address* **connect-interface loopback** *interface-number*

   A PE must use a loopback interface address with a 32-bit mask to set up an MP-IBGP peer relationship with the peer PE so that VPN routes can be recursed to tunnels.

2. Enable the PE to exchange VPNv4 routes with a specific MP-BGP peer.

   [PE-bgp] **ipv4-family vpnv4** [**unicast**]

   [PE-bgp-af-vpnv4] **peer** *ipv4-address* **enable**

   By default, only the peer relationships in the BGP IPv4 unicast address family view are automatically enabled. In other words, after the **peer as-number** command is run in the BGP view, the system automatically configures the **peer enable** command. In other address family views, however, peering must be enabled manually.

3. Configure the PE to filter VPNv4 routes

   [PE-bgp-af-vpnv4] **policy vpn-target**

   The **policy vpn-target** command enables the PE to filter received VPN routes based on VPN targets. By default, this function is enabled. In some specific networking scenarios, you need to manually disable the filtering function.

Configuration Commands — Route Configuration Between a PE and CE

1. Use EBGP to transmit routes between a PE and a CE.

   [PE-bgp] **ipv4-family vpn-instance** *vpn-instance-name*

   [PE-bgp-InstanceName] **peer** *ipv4-address* **as-number** *as-number*

   On the PE, you need to enter the VPN instance IPv4 address family view and configure the CE as a BGP peer.
   On the CE, common EBGP is configured, and VPN routes are imported to BGP using the import or network mode.

2. Use an IGP to transmit routes between the PE and CE. (OSPF is used as an example).

   [PE] **ospf** *process-id* [**router-id** *router-id*] **vpn-instance** *vpn-instance-name*

   [PE-ospf-processid] **import-route bgp** [**permit-ibgp**] [**cost** *cost* | **route-policy** *route-policy-name* | **tag** *tag* | **type** *type*] *

   [PE-bgp] **ipv4-family vpn-instance** *vpn-instance-name*

   [PE-bgp] **import-route ospf** *process-id* [**med** *med* | **route-policy** *route-policy-name*] *

   An OSPF process to be bound to the VPN instance needs to be created on the PE, and an OSPF neighbor relationship needs to be established between the PE and its connected CE. In addition, OSPF and BGP need to import routes from each other.
   On the CE, common OSPF is configured.

MPLS VPN Configuration Example — Networking Requirements

- The two sites of each user (user X and user Y) need to be interconnected through MPLS VPN. User X and user Y correspond to VPNX and VPNY, respectively.
- The interconnection interfaces, AS numbers, and IP addresses are shown in the following figure.
- Sites of user X use OSPF to exchange route with the PEs, and sites of user Y uses BGP to exchange routes with the PEs.

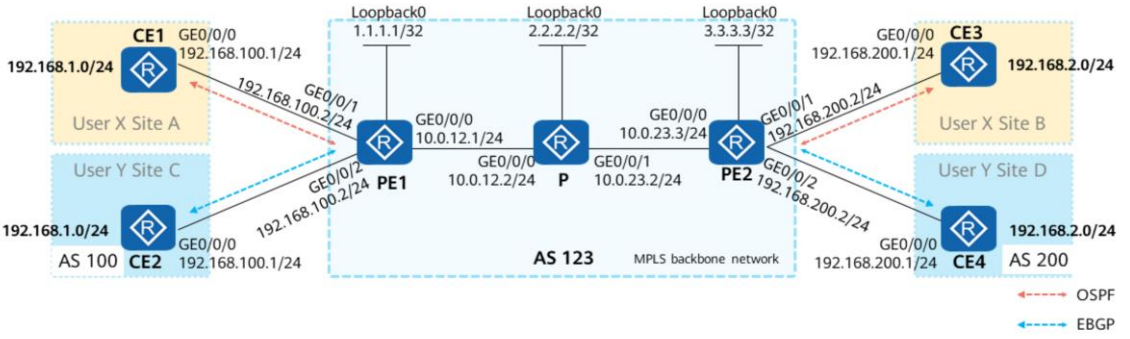

MPLS VPN Configuration Example — Configuration Roadmap

1. Configure the MPLS VPN backbone network.
   1.1 Configure an IGP to implement IP connectivity on the backbone network.
   1.2 Configure MPLS and MPLS LDP to establish MPLS LSPs for transmitting VPN data through the public network.
   1.3 Configure MP-BGP to establish MP-BGP peer relationships for transmitting VPNv4 routes.
2. Configure VPN user access.
   2.1 Create VPN instances and set parameters (RTs and RDs).
   2.2 Bind the interfaces to the VPN instances.
   2.3 Configure route exchange between the PEs and their connected CEs.

MPLS VPN Configuration Example — Data Planning

- The MPLS backbone network uses single-area OSPF to implement route interworking. MPLS LDP is enabled on all interconnection interfaces of PEs and Ps.
- The following table lists the VPN-related configurations on the PEs.

Item	Description			
	PE1		PE2	
VPN name	VPNX	VPNY	VPNX	VPNY
RD	100:1	200:1	100:1	200:1
IRT	100:321	200:234	100:123	200:432
ERT	100:123	200:432	100:321	200:234
Interface	GE0/0/1	GE0/0/2	GE0/0/1	GE0/0/2
MP-BGP	Source interface: Loopback0		Source interface: Loopback0	

MPLS VPN Backbone Network Configuration (1)

1.1 Configure OSPF on the MPLS VPN backbone network to enable route connectivity within the backbone network. The OSPF configuration on PE1 is used as an example.

```
[PE1]ospf 100 router-id 1.1.1.1
[PE1-ospf-100]area 0
[PE1-ospf-100-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[PE1-ospf-100-area-0.0.0.0]network 1.1.1.1 0.0.0.0
```

1.2 Configure MPLS and LDP on PE1, P, and PE2. The configuration on PE1 is used as an example.

```
[PE1]mpls lsr-id 1.1.1.1
[PE1]mpls
Info: Mpls starting, please wait... OK!
[PE1-mpls]mpls ldp
[PE1-mpls-ldp]Interface GigabitEthernet 0/0/0
[PE1-GigabitEthernet0/0/0]mpls
[PE1-GigabitEthernet0/0/0]mpls ldp
```

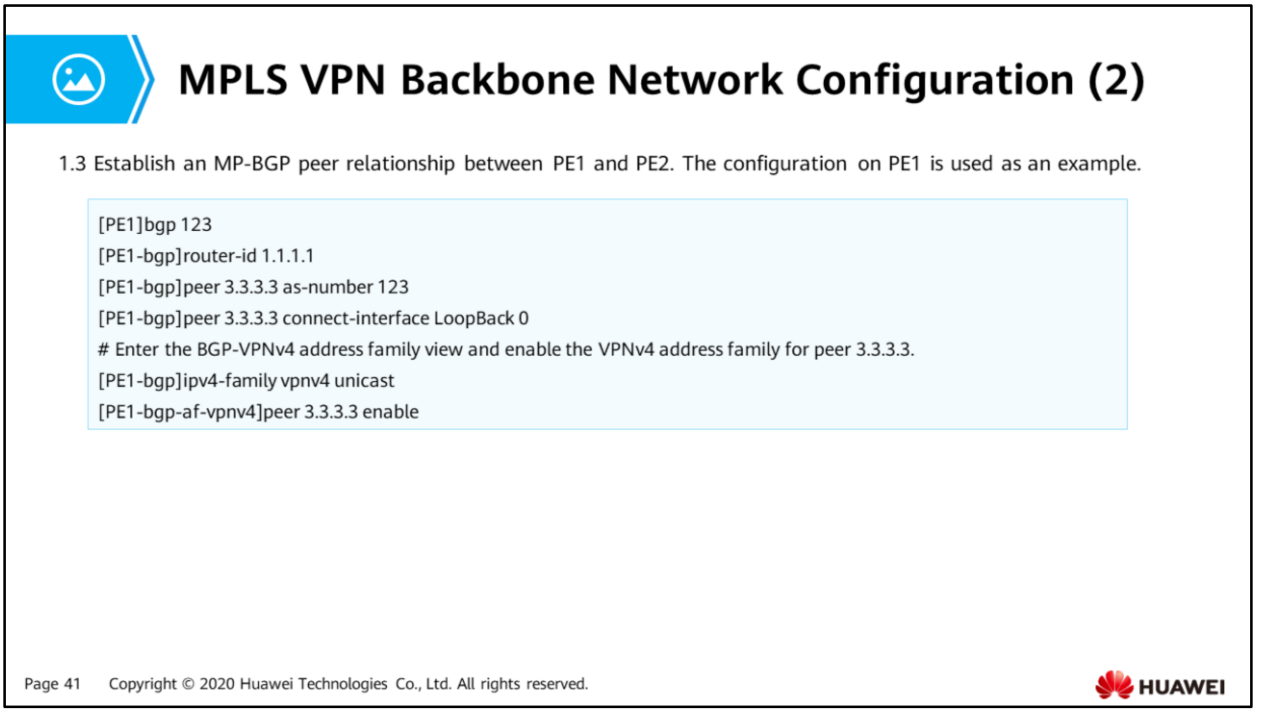

- By default, only the peer relationships in the BGP IPv4 unicast address family view are automatically enabled. In other words, after the **peer as-number** command is run in the BGP view, the system automatically configures the **peer enable** command. In other address family views, however, peering must be enabled manually.

MPLS VPN Backbone Network Configuration — Configuration Verification

Check whether LDP LSPs are established.

```
[PE1]display mpls lsp
------------------------------------------------------------------------------
                   LSP Information: LDP LSP
------------------------------------------------------------------------------
FEC             In/Out Label    In/Out IF       Vrf Name
3.3.3.3/32      NULL/1025       -/GE0/0/0
1.1.1.1/32      3/NULL          -/-
```

```
[PE2]display mpls lsp
------------------------------------------------------------------------------
                   LSP Information: LDP LSP
------------------------------------------------------------------------------
FEC             In/Out Label    In/Out IF       Vrf Name
3.3.3.3/32      3/NULL          -/-
1.1.1.1/32      NULL/1024       -/GE0/0/0
```

Check the MP-BGP peer status. The following example uses the command output on PE1.

```
[PE1]display bgp vpnv4 all peer
 BGP local router ID : 1.1.1.1
 Local AS number : 123
 Total number of peers : 1                 Peers in established state : 1

  Peer        V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State Pre              fRcv
  3.3.3.3     4         123  16        18       0     00:14:20      Established            0
```

VPN User Access Configuration (1)

2.1 Create VPN instances and set RDs and RTs as planned. The following uses PE1 as an example.

```
[PE1]ip vpn-instance VPNX
[PE1-vpn-instance-VPNX]route-distinguisher 100:1
[PE1-vpn-instance-VPNX-af-ipv4] vpn-target 100:321 import-extcommunity
 IVT Assignment result:
Info: VPN-Target assignment is successful.
[PE1-vpn-instance-VPNX-af-ipv4] 100:123 export-extcommunity
 EVT Assignment result:
Info: VPN-Target assignment is successful.
[PE1-vpn-instance-VPNX-af-ipv4] quit
[PE1-vpn-instance-VPNX]quit
[PE1]ip vpn-instance VPNY
[PE1-vpn-instance-VPNY]route-distinguisher 200:1
[PE1-vpn-instance-VPNY-af-ipv4]vpn-target 200:234 import-extcommunity
[PE1-vpn-instance-VPNY-af-ipv4]vpn-target 200:432 export-extcommunity
[PE1-vpn-instance-VPNY-af-ipv4]quit
[PE1-vpn-instance-VPNY]quit
```

VPN User Access Configuration (2)

2.2 Bind the interfaces to the VPN instances.

```
[PE1]interface GigabitEthernet 0/0/1
[PE1-GigabitEthernet0/0/1]ip binding vpn-instance VPNX
Info: All IPv4 related configurations on this interface are removed!
Info: All IPv6 related configurations on this interface are removed!
[PE1-GigabitEthernet0/0/1]ip address 192.168.100.2 24
[PE1-GigabitEthernet0/0/1]interface GigabitEthernet 0/0/2
[PE1-GigabitEthernet0/0/2]ip binding vpn-instance VPNY
Info: All IPv4 related configurations on this interface are removed!
Info: All IPv6 related configurations on this interface are removed!
[PE1-GigabitEthernet0/0/2]ip address 192.168.100.2 24
```

VPN User Access Configuration (3)

2.3.1 Configure route exchange between CE1 and PE1 and between CE3 and PE2. The configuration on PE1 is used as an example.

```
# Create an OSPF process and bind it to the instance.
[PE1]ospf 2 vpn-instance VPNX
[PE1-ospf-2]area 0
[PE1-ospf-2-area-0.0.0.0]network 192.168.100.0 0.0.0.255
[PE1-ospf-2-area-0.0.0.0]quit
```

```
# Configure route import between OSPF and MP-BGP.
[PE1]ospf 2 vpn-instance VPNX
[PE1-ospf-2]import-route bgp
[PE1-ospf-2]quit
[PE1]bgp 123
[PE1-bgp]ipv4-family vpn-instance VPNX
[PE1-bgp-VPNX]import-route ospf 2
```

2.3.2 Configure route exchange between CE2 and PE1 and between CE4 and PE2. The configurations on CE2 and PE1 are used as an example.

```
# Configure EBGP on CE2 and import the direct route
192.168.1.0/24.
[CE2]BGP 200
[CE2-bgp]peer 192.168.100.2 as-number 123
[CE2-bgp]network 192.168.1.0 24
[CE2-bgp]quit
```

```
# Configure an EBGP peer for the VPN instance VPNY on PE1.
[PE1]bgp 123
[PE1-bgp]ipv4-family vpn-instance VPNY
[PE1-bgp-VPNY]peer 192.168.100.1 as-number 200
```

Configuration Verification (1)

Check the routes learned on the CEs for user X in VPNX.

[CE1]display ip routing-table
Route Flags: R - relay, D - download to fib

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
192.168.1.0/24	Direct	0	0	D	192.168.1.254	GigabitEthernet0/0/1
192.168.2.0/24	OSPF	10	4	D	192.168.100.2	GigabitEthernet0/0/0
192.168.100.0/24	Direct	0	0	D	192.168.100.1	GigabitEthernet0/0/0
192.168.200.0/24	O_ASE	150	1	D	192.168.100.2	GigabitEthernet0/0/0

[CE3]dis ip routing-table
Route Flags: R - relay, D - download to fib

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
192.168.1.0/24	OSPF	10	4	D	192.168.200.2	GigabitEthernet0/0/0
192.168.2.0/24	Direct	0	0	D	192.168.2.254	GigabitEthernet0/0/1
192.168.100.0/24	O_ASE	150	1	D	192.168.200.2	GigabitEthernet0/0/0
192.168.200.0/24	Direct	0	0	D	192.168.200.1	GigabitEthernet0/0/0

Configuration Verification (2)

Check the routes learned on the CEs for user Y in VPNY.

```
[CE2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
192.168.1.0/24      Direct  0    0           D   192.168.1.254   GigabitEthernet0/0/1
192.168.2.0/24      EBGP    255  0           D   192.168.100.2   GigabitEthernet0/0/0
192.168.100.0/24    Direct  0    0           D   192.168.100.1   GigabitEthernet0/0/0
```

```
[CE4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
192.168.2.0/24      Direct  0    0           D   192.168.1.254   GigabitEthernet0/0/1
192.168.1.0/24      EBGP    255  0           D   192.168.100.2   GigabitEthernet0/0/0
192.168.200.0/24    Direct  0    0           D   192.168.100.1   GigabitEthernet0/0/0
```

Configuration Verification (3)

```
[PE2] display bgp vpnv4 vpn-instance VPNX routing-table 192.168.1.0 24

 BGP local router ID : 3.3.3.3
 Local AS number : 123
 VPN-Instance VPNX, Router ID 3.3.3.3:
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 192.168.1.0/24:
 Label information (Received/Applied): 1026/NULL
 From: 1.1.1.1 (1.1.1.1)
Relay token: 0x1
 Original nexthop: 1.1.1.1
```

```
[PE2]display mpls lsp
------------------------------------------------------------------------------
LSP Information: LDP LSP
------------------------------------------------------------------------------
FEC                 In/Out Label   In/Out IF    Vrf Name
1.1.1.1/32          NULL/1024      -/GE0/0/0
1.1.1.1/32          1024/1024      -/GE0/0/0
```

Use the data from 192.168.2.0/24 to 192.168.1.0/24 as an example. The outer label is 1024, which is allocated by MPLS LDP. The inner label is 1026, which is allocated by MP-BGP.

```
No.     Time        Source           Destination      Protocol  Length  Info
5       12.109000   192.168.2.254    192.168.1.254    ICMP      102     Echo (ping) request
Frame 5: 102 bytes on wire (816 bits), 102 bytes captured (816 bits) on interface 0
Ethernet II, Src: HuaweiTe_b1:15:3e (00:e0:fc:b1:15:3e), Dst: HuaweiTe_49:20:bb (00:e0:fc:49:20:bb)
MultiProtocol Label Switching Header, Label: 1024, Exp: 0, S: 0, TTL: 254
MultiProtocol Label Switching Header, Label: 1026, Exp: 0, S: 1, TTL: 254
Internet Protocol Version 4, Src: 192.168.2.254, Dst: 192.168.1.254
Internet Control Message Protocol
```

Quiz

1. (Single) Which of the following route targets is carried when MP-BGP transmits VPNv4 routes? ( )
   A. Export RT
   B. Implied RT
   C. Import RT
   D. Extended RT
2. (Multiple) In basic MPLS VPN networking, which of the following configurations are required for a VPN instance?
   A. Configure an import RT.
   B. Configure an export RT.
   C. Configure an RD.
   D. Bind an interface to the VPN instance.

1. A
2. ABCD

Summary

- MPLS VPNs are usually constructed by a carrier to provide VPN services for different users. In this manner, routes and data of different users can be transmitted over the MPLS VPN and are isolated from each other.
- MPLS VPN control plane:
  - VRFs and RDs are used to isolate VPN routes of different users and construct unique VPNv4 routes.
  - RTs are to control the import and export of VPNv4 routes.
  - MP-BGP is used to transmit VPNv4 route prefixes, labels, and RTs.
  - MPLS LDP is used to establish LSPs over the public network.
- MPLS forwarding plane: Data is forwarded based on inner and outer labels. The inner labels differentiate VPN data, whereas the outer labels are used for public network transport.


Thank You
www.huawei.com