IPv6 Routing

Foreword

- With the advent of the Internet of Everything (IoE), IPv4 address space is running out. It is imperative that we replace IPv4 with IPv6. How to enable reachability between nodes on an IPv6 network? Similar to an IPv4 network, an IPv6 network also supports static routes and dynamic routing protocols.
- IPv6 static routes and IPv4 static routes are configured in a similar way. To support IPv6 networks, the Internet Engineering Task Force (IETF) defines Open Shortest Path First OSPF version 3 (OSPFv3) and extends Intermediate System to Intermediate System (IS-IS) and Border Gateway Protocol (BGP).
- This course describes the concepts and configurations of IPv6 static routes and common IPv6 dynamic routing protocols, including OSPFv3, IS-IS (IPv6), and BGP4+.

Objectives

- On completion of this course, you will be able to:
  - Configure IPv6 static routes.
  - Analyze the differences between OSPFv3 and OSPFv2.
  - Configure basic OSPFv3 functions.
  - Describe the extension of IS-IS to IPv6.
  - Configure basic IS-IS (IPv6) functions.
  - Describe the extension of BGP to IPv6.
  - Configure basic BGP4+ functions.

HUAWEI

Contents

IPv6 Static Routes

- Similar to IPv4 static routes, IPv6 static routes are configured by the administrator and are applicable to IPv6 networks with simple structures.
- When creating an IPv6 static route, you can specify both the outbound interface and next hop, or specify only the outbound interface or next hop.
  - Point-to-point (P2P) interface: specify the outbound interface.
  - Broadcast interface: specify the next hop.
- Load balancing and backup:
  - Setting the same preference value for IPv6 static routes to the same destination implements load balancing among these routes, whereas setting different preference values for such routes implements route backup.

- IPv6 static routes and IPv4 static routes differ mainly in destination and next-hop IP addresses. IPv6 static routes use IPv6 addresses, whereas IPv4 static routes use IPv4 addresses.

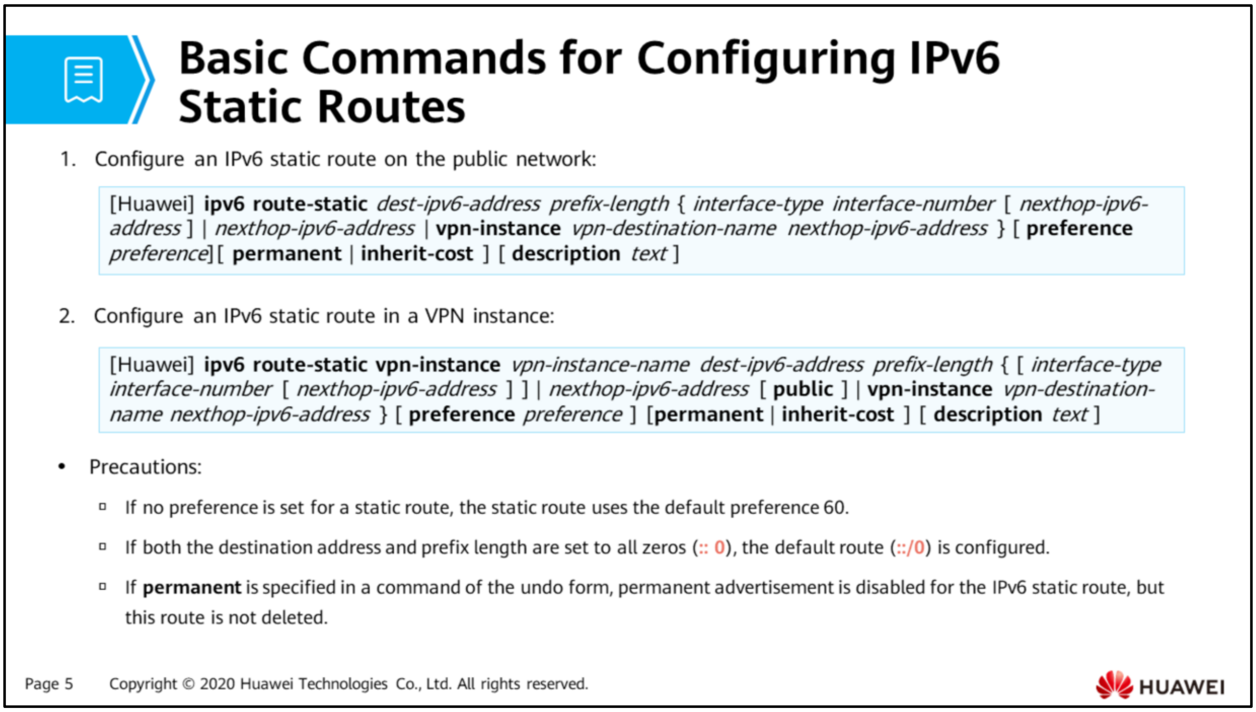

- [Huawei] **ipv6 route-static** *dest-ipv6-address prefix-length* { *interface-type interface-number* [*nexthop-ipv6-address*] | *nexthop-ipv6-address* | **vpn-instance** *vpn-destination-name nexthop-ipv6-address* } [**preference** *preference*][**permanent** | **inherit-cost**] [**description** *text*]
  - **preference** *preference*: specifies a preference value for the route. The value is an integer ranging from 1 to 255. The default value is 60.
  - **permanent**: enables the function of permanently advertising the IPv6 static route.
  - **inherit-cost**: enables the static route to inherit the cost of the recursive route.
  - **description** *text*: specifies a description for the static route. The value is a string of 1 to 80 characters and can contain spaces.
- [Huawei] **ipv6 route-static vpn-instance** *vpn-instance-name dest-ipv6-address prefix-length* { [*interface-type interface-number* [*nexthop-ipv6-address*]] | *nexthop-ipv6-address* [**public**] | **vpn-instance** *vpn-destination-name nexthop-ipv6-address* } [**preference** *preference*] [**permanent** | **inherit-cost**] [**description** *text*]
  - **public**: indicates that *nexthop-ipv6-address* is a public network address instead of an address in the source VPN instance.

Example for Configuring IPv6 Static Routes

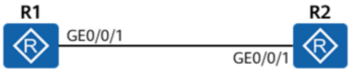

Device	Interface	IPv6 Address
R1	GE0/0/1	2001:DB8:2345:12::1/64
	Loopback0	2001:DB8:2345:1::1/128
R2	GE0/0/1	2001:DB8:2345:12::2/64
	Loopback0	2001:DB8:2345:2::2/128

Scenario description:

- A company has an IPv6 network deployed for service testing. In the initial deployment phase, static routes are used to implement interconnection and communication over the IPv6 network.
- Static routes are also used to enable communication between loopback0 interfaces of R1 and R2.

1. Assign an IPv6 address to each router interface. (The configuration details are not provided here.)
2. Configure IPv6 static routes.

```
[R1] ipv6 route-static 2001:DB8:2345:2::2 128 2001:DB8:2345:12::2
```

```
[R2] ipv6 route-static 2001:DB8:2345:1::1 128 2001:DB8:2345:12::1
```

3. Test the IPv6 network connectivity.

```
[R1] ping ipv6 2001:DB8:2345:2::2
 PING 2001:DB8:2345:2::2 : 56  data bytes, press CTRL_C to break
   Reply from 2001:DB8:2345:2::2
   bytes=56 Sequence=1  hop limit=64  time = 90 ms
   Reply from 2001:DB8:2345:2::2
   bytes=56 Sequence=2  hop limit=64  time = 20 ms
   Reply from 2001:DB8:2345:2::2
   bytes=56 Sequence=3  hop limit=64  time = 20 ms
   Reply from 2001:DB8:2345:2::2
   bytes=56 Sequence=4  hop limit=64  time = 20 ms
   Reply from 2001:DB8:2345:2::2
   bytes=56 Sequence=5  hop limit=64  time = 20 ms
 --- 2001:DB8:2345:2::2 ping statistics ---
   5 packet(s) transmitted
   5 packet(s) received
   0.00% packet loss
   round-trip min/avg/max = 20/34/90 ms
```

Contents

OSPFv3 Overview

- OSPF is a link-state Interior Gateway Protocol (IGP) defined by the IETF. Currently, OSPF version 2 (OSPFv2) is intended for IPv4, and OSPF version 3 (OSPFv3) is intended for IPv6.
- OSPFv3 is mainly developed as a routing protocol independent of any specific network layer. The internal routing information of OSPFv3 has been redesigned to serve this purpose.

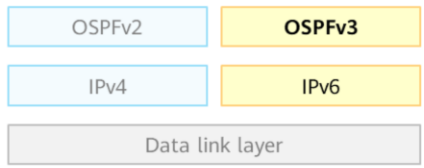

OSPFv2 is an IGP running over IPv4, whereas OSPFv3 is an IGP running over IPv6. They are not compatible.

HUAWEI

Similarities Between OSPFv3 and OSPFv2

OSPF's basic running mechanisms remain unchanged, including:

- Basic concepts:
  - Area division and router types
  - Parameters that affect route calculation: preference and cost
  - Supported network types: broadcast, non-broadcast multiple access (NBMA), P2P, and point-to-multipoint networking (P2MP)
  - Packet types: Hello, database description (DD), link state request (LSR), link state update (LSU), and link state acknowledgment (LSAck)
- Working mechanisms:
  - Establishment of neighbor relationships and transition of the neighbor relationship status
  - Election of the Designated router (DR) and backup DR (BDR)
  - Link state advertisement (LSA) flooding
  - Route calculation process

- Similarities also include support for special areas, support for virtual links, and multi-process support.
- For details, see the "HCIP-Datacom-Core Technology" course.

OSPFv3 Topology and Router Types

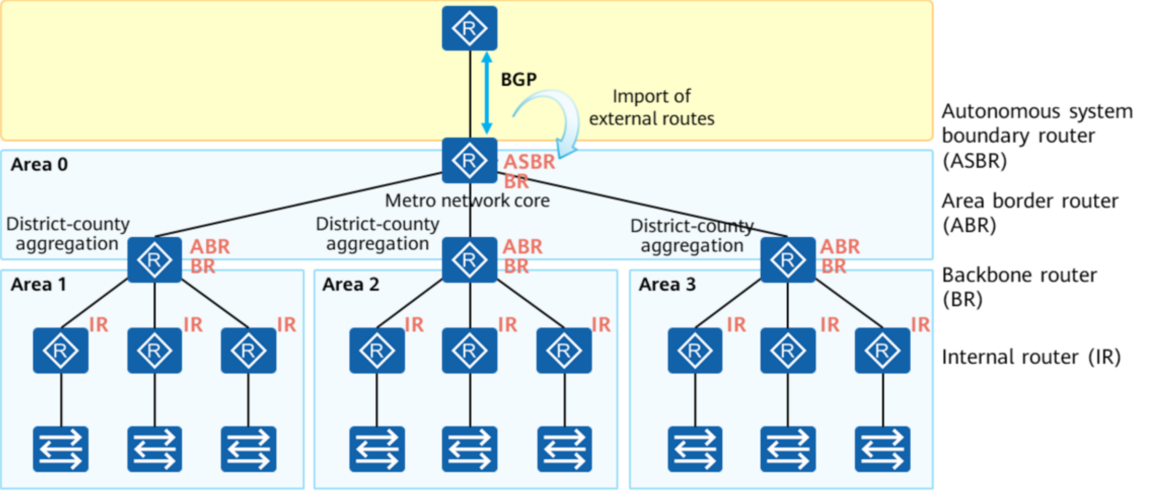

OSPFv3 and OSPFv2 Have Similar Working Mechanisms

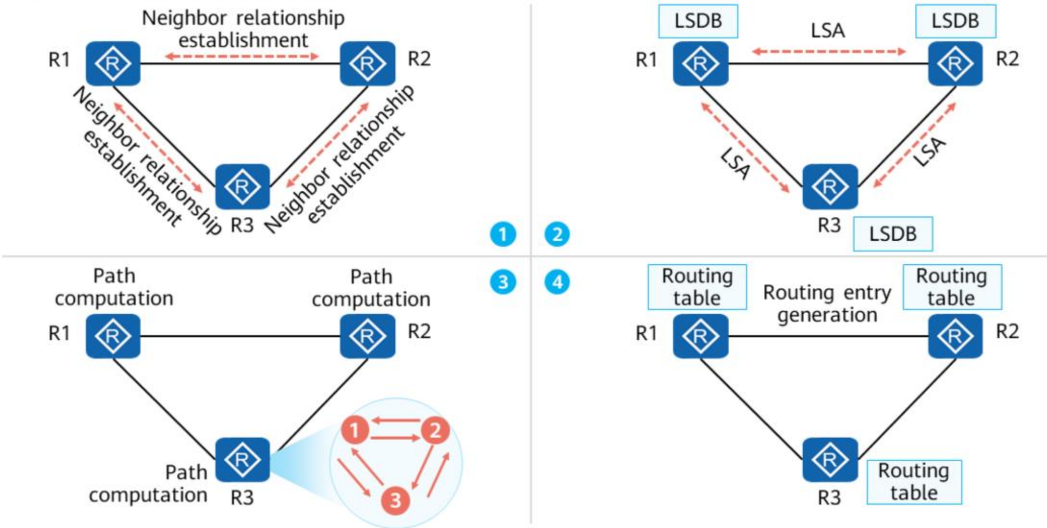

HUAWEI

Differences Between OSPFv3 and OSPFv2

- OSPFv3 runs and performs topology calculation based on links instead of network segments.
- OSPFv3 supports multiple instances per link.
- OSPFv3 packets and LSAs do not contain IP addresses, and the packet and LSA formats are reconstructed.
  - OSPFv3 packets and router-LSAs/network-LSAs do not contain IP addresses.
  - OSPFv3 LSAs define the flooding scopes of LSAs.
  - OSPFv3 has new types of LSAs defined to carry IPv6 addresses and prefixes.
  - OSPFv3 neighbors are identified only by router IDs, not by IP addresses.

Unique Neighbor ID: Router ID

OSPFv3 uses router IDs to identify network devices.

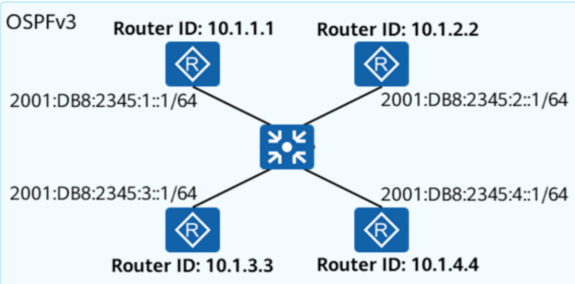

- A router ID uniquely identifies an OSPFv3 device in an AS. If no router ID is configured, an OSPFv3 process cannot run.
- When you configure a router ID for a device, ensure that the router ID in an AS is unique.
- A router ID is a 32-bit local identifier, which is not related to IPv6 addresses and is expressed in dotted decimal notation.

- On broadcast, NBMA, P2P, and P2MP networks, OSPFv2 uses IPv4 interface addresses to identify neighbors. On virtual-link networks, however, OSPFv2 uses router IDs to identify neighbors.

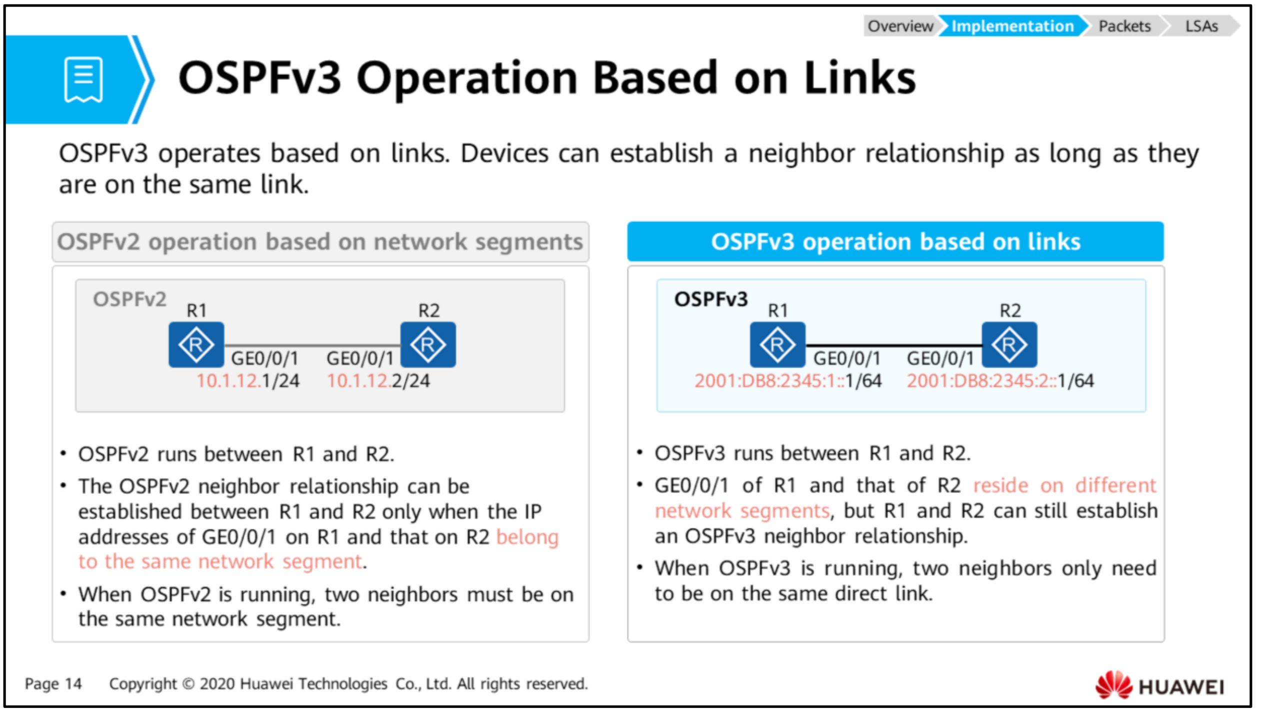

- IPv6 emphasizes the link concept. Multiple IPv6 prefixes that indicate different IP subnets can be allocated to the same link. Different from IPv4, IPv6 allows two nodes on the same link to communicate even if they do not have the same IPv6 prefix. This greatly changes the OSPF behavior.
- In OSPFv3, the concepts "link" and "prefix" are frequently used, which however are independent of each other. The terms "network" and "subnet" used in OSPFv2 should be replaced with the term "link" when OSPFv3 is discussed.

Multi-Instance Support by Links

- An OSPFv3 physical interface can be bound to multiple instances, which are identified by different instance IDs. A single physical link can run multiple OSPFv3 instances.
- These OSPFv3 instances can each establish a neighbor relationship with the device at the remote end of the link, send packets to the remote device, and do not interfere with each other, thereby fully sharing resources on the same link.

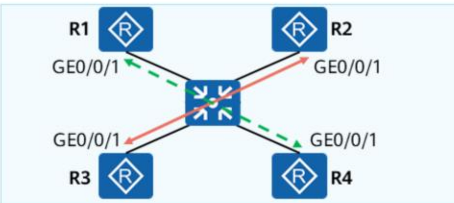

By using instance IDs, R2 and R3 can establish an OSPFv3 neighbor relationship, and R1 and R4 can also establish an OSPFv3 neighbor relationship.

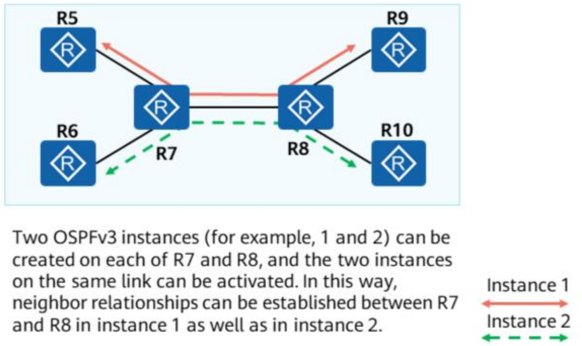

Two OSPFv3 instances (for example, 1 and 2) can be created on each of R7 and R8, and the two instances on the same link can be activated. In this way, neighbor relationships can be established between R7 and R8 in instance 1 as well as in instance 2.

- In multi-instance, each instance is differentiated by adding a specific instance ID to the OSPFv3 packet header. If an instance is assigned a specific instance ID, the OSPFv3 packets that do not match the instance ID are discarded.

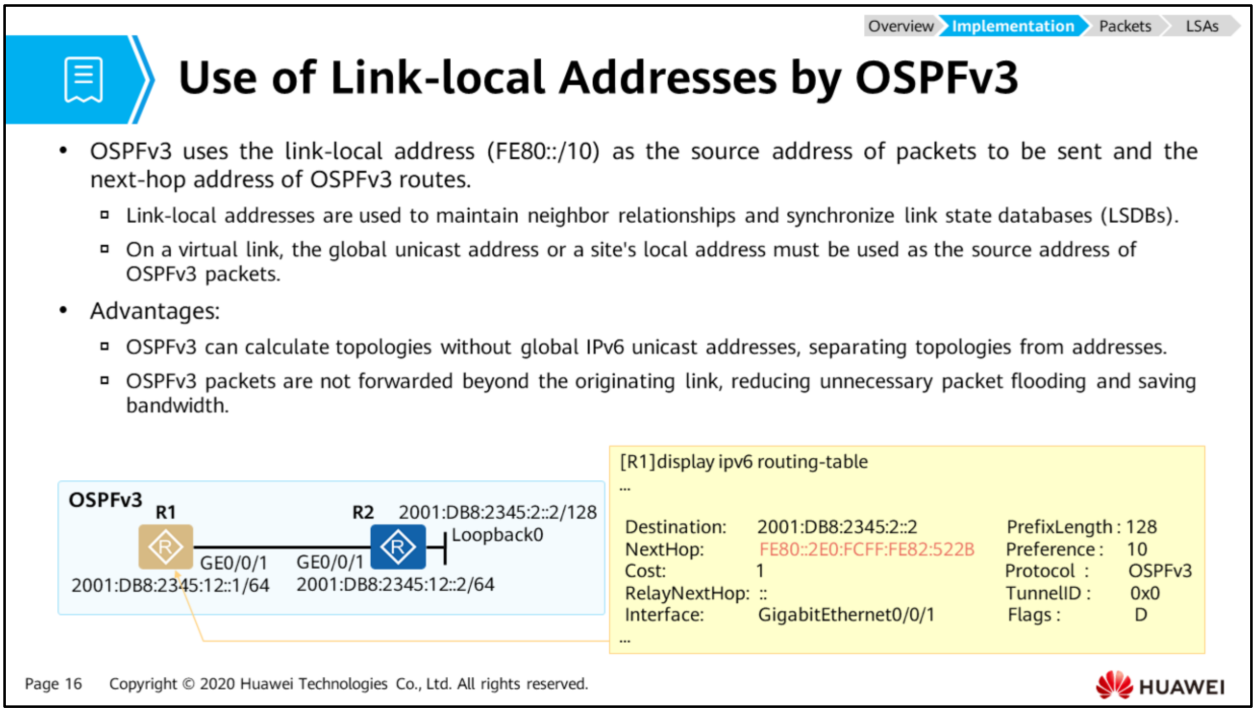

- IPv6 implements neighbor discovery and automatic configuration using link-local addresses. Routers running IPv6 do not forward IPv6 packets whose destination addresses are link-local addresses. Such packets are valid only on the local link.
- OSPFv3 is a routing protocol running on IPv6 and uses link-local addresses to send OSPFv3 packets.
    - OSPFv3 assumes that each router has been assigned a link-local address on each link. All OSPFv3 interfaces except virtual-link interfaces use the associated link-local addresses as the source addresses to send OSPFv3 packets.
    - A router learns the link-local addresses of all the other routers attached to the same link and uses these addresses as the next-hop addresses to forward packets.
    - Note: Description of link-local addresses is only contained in link-LSAs (new type of LSA supported in OSPFv3).

OSPFv3 Packets

- OSPFv3 and OSPFv2 have the same types of packets:
  - Hello, DD, LSR, LSU, and LSAck packets
- OSPFv3 and OSPFv2 use the same protocol number 89.
  - OSPFv2: The Protocol in the IPv4 packet header is 89.
  - OSPFv3: The Next Header in the IPv6 packet header is 89.
- Similar to OSPFv2, OSPFv3 uses a multicast address as the destination address of OSPF packets.
  - OSPFv2 uses IPv4 multicast addresses:
    - OSPF IGP routers use 224.0.0.5, whereas the OSPF IGP DR uses 224.0.0.6.
  - OSPFv3 uses IPv6 multicast addresses:
    - OSPF IGP routers use FF02::5, whereas the OSPF IGP DR uses FF02::6.

- OSPFv3 packets have the following functions:
  - Hello packet: Hello packets are sent periodically to discover, establish, and maintain OSPFv3 neighbor relationships.
  - DD packet: A DD packet describes the summary of a local LSDB and is used for LSDB synchronization between two devices.
  - LSR packet: An LSR packet is used to request the required LSAs from a neighbor. An OSPFv3 device sends LSR packets to its neighbor only after DD packets have been successfully exchanged between them.
  - LSU packet: An LSU packet is sent to a neighbor to provide required LSAs.
  - LSAck packet: An LSAck packet is used to acknowledge the received LSAs.

OSPFv3 Packet Header

- OSPFv2 packet header:

Version=2	Type	Packet Length
Router ID		
Area ID		
Checksum		Auth Type
Authentication		

- OSPFv3 packet header:

Version=3	Type	Packet Length	
Router ID			
Area ID			
Checksum		**Instance ID**	0

- Similar to OSPFv2, OSPFv3 defines the same packet header for its five types of packets. However, some fields in the OSPFv3 packet header differ from those in the OSPFv2 packet header.
- Changes:
  - **Instance ID**: 1 byte. The default value is 0. This field allows multiple OSPFv3 instances to be run on a link. Each instance has a unique instance ID. The instance ID is valid only on the local link.
  - **All authentication fields are removed from the OSPFv3 packet header**: OSPFv3 authentication can be implemented using IPv6's authentication and security processing mechanism. In addition, OSPFv3 packets can be authenticated using OSPFv3's own mechanism.

- Version: indicates the OSPF version, and occupies 1 byte. For OSPFv3, the value is 3.
- Type: indicates the type of an OSPFv3 packet and occupies 1 byte. The following types are available:
  - 1: Hello packet
  - 2: DD packet
  - 3: LSR packet
  - 4: LSU packet
  - 5: LSAck packet
- Packet length: indicates the total length of an OSPFv3 packet, including the packet header. The field occupies 2 bytes.
- Router ID: indicates the router ID of the router that originates the packet, and occupies 4 bytes.
- Area ID: indicates the area in which the packet is sent, and occupies 4 bytes.
- Checksum: uses the standard 16-bit IPv6 checksum and occupies 2 bytes.
- 0: Occupying 1 byte, this field is reserved and must be set to 0.

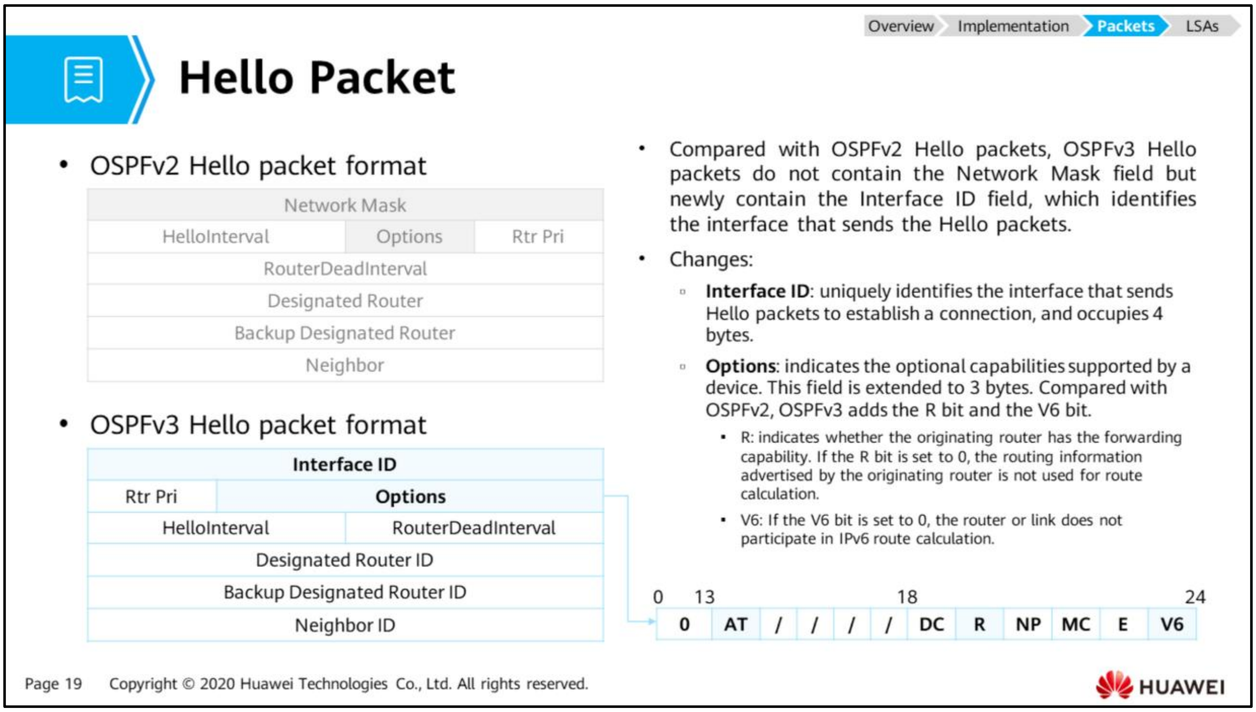

- Rot Pri: indicates the router's router priority, which is used for DR election. This field occupies 1 byte, and the default value is 1. If it is set to 0, the router cannot participate in DR or BDR election.
- Options: indicates the optional capabilities supported by the router and occupies 3 bytes.
  - AT: indicates whether OSPFv3 authentication is supported. This option occupies 1 bit. If the AT bit is 1, an authentication tail field containing authentication information is added to the OSPFv3 packet.
  - DC: indicates whether the capability of processing demand circuits is supported. This option occupies 1 bit.
  - R: indicates whether the originator is a valid router. This option occupies 1 bit.
  - NP: indicates whether the area to which the originating router interface belongs is a not-so-stubby area (NSSA). This option occupies 1 bit.
  - MC: indicates whether multicast data packets can be forwarded. This option occupies 1 bit.
  - E: indicates whether external routes are supported. This option occupies 1 bit.
  - V6: indicates whether the router or link can participate in route calculation. This option occupies 1 bit. If it is set to 0, the router or link does not participate in IPv6 route calculation.
- HelloInterval: indicates the interval at which Hello packets are sent. This field occupies 2 bytes.

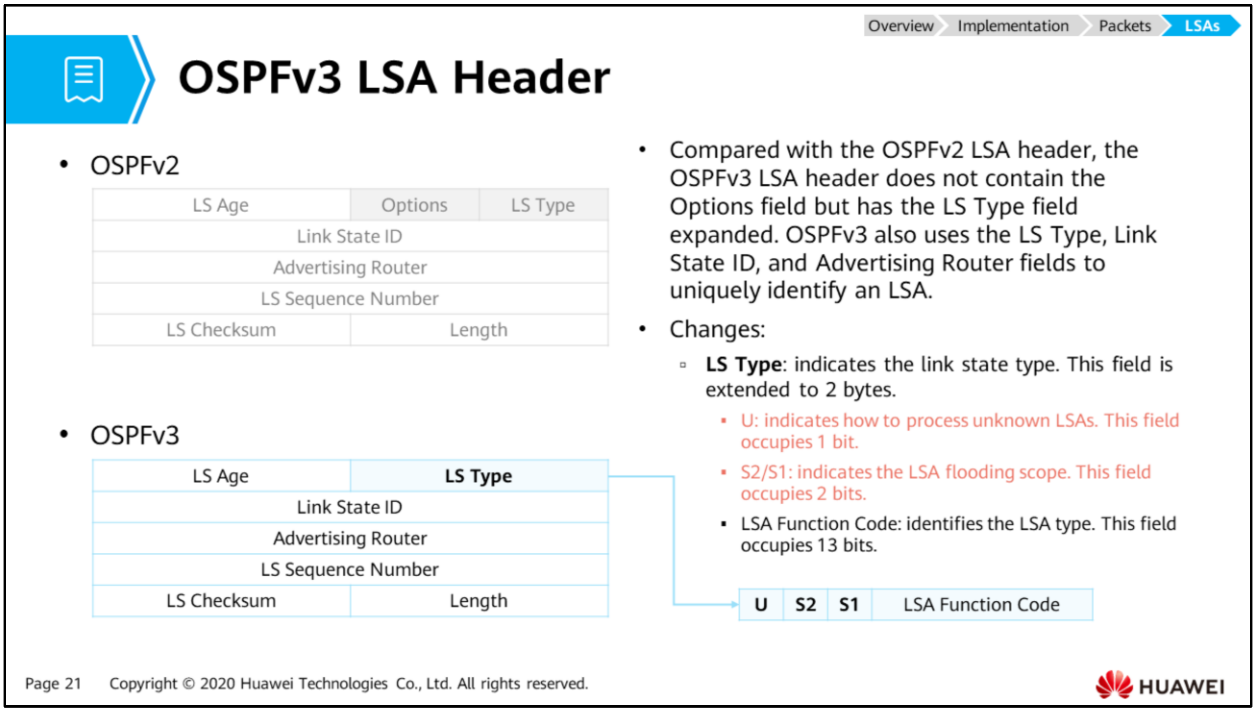

- LS Age: indicates the time elapsed since the LSA was generated, in seconds. This field occupies 2 bytes. The value of this field continually increases regardless of whether the LSA is transmitted over a link or saved in an LSDB.
- LS Type: indicates the LSA type. This field occupies 2 bytes. The high-order three bits of this field identify generic properties of the LSA, whereas the remaining bits identify the LSA's specific function.
  - The U-bit indicates how to process an unknown LSA, that is, how a router that does not recognize an LSA's function code should process this LSA.
    - 0: The LSA is treated as if it had the link-local flooding scope.
    - 1: The LSA is stored and flooded as if its type had been understood.
  - The S2 and S1 bits indicate the flooding scope of the LSA.
    - S2 S1 = 0 0: link-local flooding scope. The LSA is flooded only on the originating link.
    - S2 S1 = 0 1: area flooding scope. The LSA is flooded to all routers in the originating area.
    - S2 S1 = 1 0: AS flooding scope. The LSA is flooded to all routers in the local AS.
    - S2 S1 = 1 1: reserved.
- Link State ID: indicates a local 32-bit identifier, which is irrelevant to an IPv6 address, and occupies 4 bytes. This field, together with the LS Type and Advertising Router fields, describes an LSA in the OSPFv3 domain. Compared with OSPFv2, OSPFv3 does not contain address information in the Link State ID field.

Overview > Implementation > Packets > LSAs

OSPFv3 LSA Types

- Compared with OSPFv2, OSPFv3 LSAs have similar names but slightly different functions.
- OSPFv3 supports two new types of LSAs: link-LSA and intra-area-prefix-LSA.

OSPFv2 LSAs		OSPFv3 LSAs		Similarities and Differences
Type	Name	Type	Name	
1	Router-LSA	0x2001	Router-LSA	The functions are similar, but the OSPFv3 LSAs no longer describe address information. Instead, they only describe topology structures.
2	Network-LSA	0x2002	Network-LSA	
3	Network-summary-LSA	0x2003	Inter-area-prefix-LSA	The functions are similar, but the names are different.
4	ASBR-summary-LSA	0x2004	Inter-area-router-LSA	
5	AS-external-LSA	0x4005	AS-external-LSA	The functions and names are the same.
7	NSSA-LSA	0x2007	NSSA-LSA	
		0x0008	Link-LSA	New in OSPFv3
		0x2009	Intra-area-prefix-LSA	

HUAWEI

- As shown in the figure, the U-bit in the LS Type field of the OSPFv3 LSA header is 0 by default. Except the Type 5 and Type 8 LSAs, the other types of LSAs all have the area flooding scope (S2 S1 = 0 1).
  - Link-local flooding scope: LSAs, including link-LSAs, are flooded only on the local link.
  - Area flooding scope: The following types of LSAs are flooded in a single OSPF area: router-LSA, network-LSA, inter-area-prefix-LSA, inter-area-router-LSA, NSSA-LSA, and intra-area-prefix-LSA.
  - AS flooding scope: LSAs, including AS-external-LSAs, are flooded in an entire routing domain (AS).

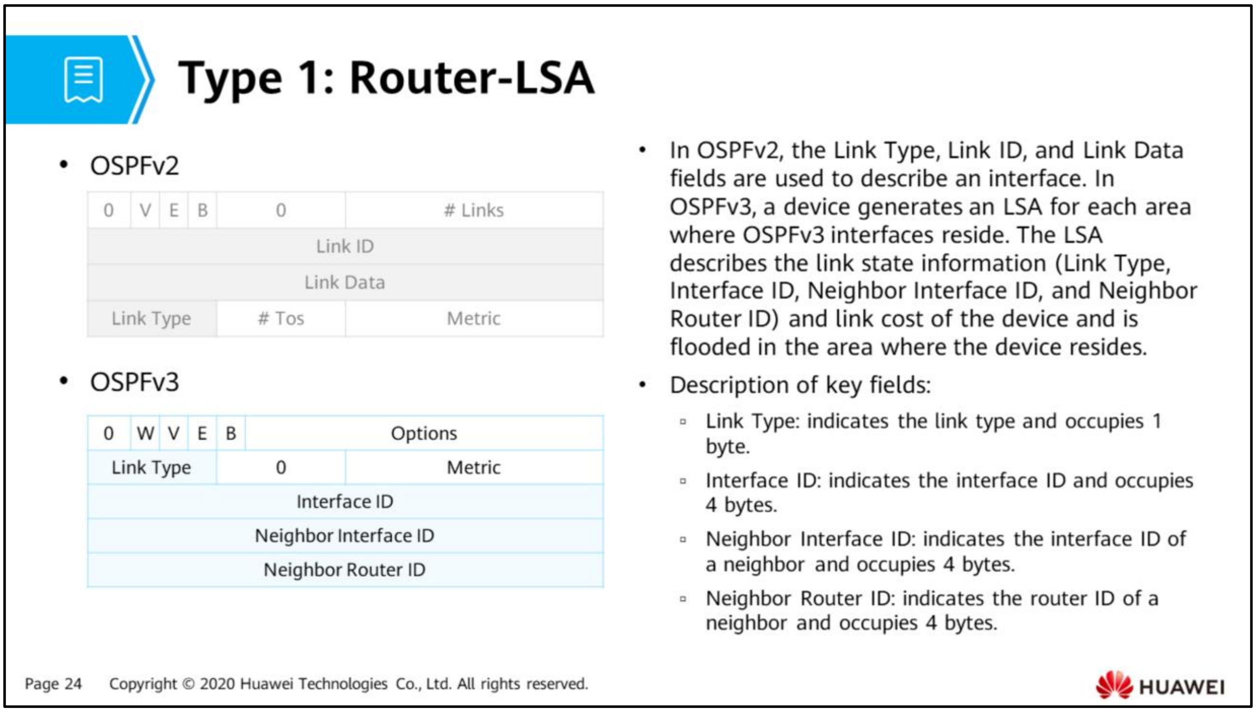

- The fields in an OSPFv3 router-LSA are described as follows:
  - W: wildcard receiver. The value 1 indicates that the router supports multicast routes.
  - V: virtual link. The value 1 indicates that the router that generates the LSA is at one end of the virtual link.
  - E: external. The value 1 indicates that the router that generates the LSA is an ASBR.
  - B: border. The value 1 indicates that the router that generates the LSA is an ABR.
  - Options: indicates the optional capabilities supported by the router and occupies 3 bytes.
    - DC: indicates whether the capability of processing demand circuits is supported. This option occupies 1 bit.
    - R: indicates whether the originator is a valid router. This option occupies 1 bit.
    - NP: indicates whether the area to which the originating router interface belongs is a not-so-stubby area (NSSA). This option occupies 1 bit.

Type 2: Network-LSA

- OSPFv2

Network Mask
Attached Router
Attached Router
...

- OSPFv3

0	Options
Attached Router	
Attached Router	
...	

- Compared with OSPFv2, OSPFv3 network-LSAs do not contain the Network Mask field. Instead, OSPFv3 network-LSAs describe the link state of the local network segment only by using the router IDs of the attached routers. Network-LSAs are generated by the DR and flooded in the area to which the DR belongs.
- Description of key fields:
  - Attached Router: indicates the router ID of each router connected to the same network segment, including the router ID of the DR, and occupies 4 bytes.

- The fields in an OSPFv3 network-LSA are described as follows:
  - Options: same as the Options field in a router-LSA.

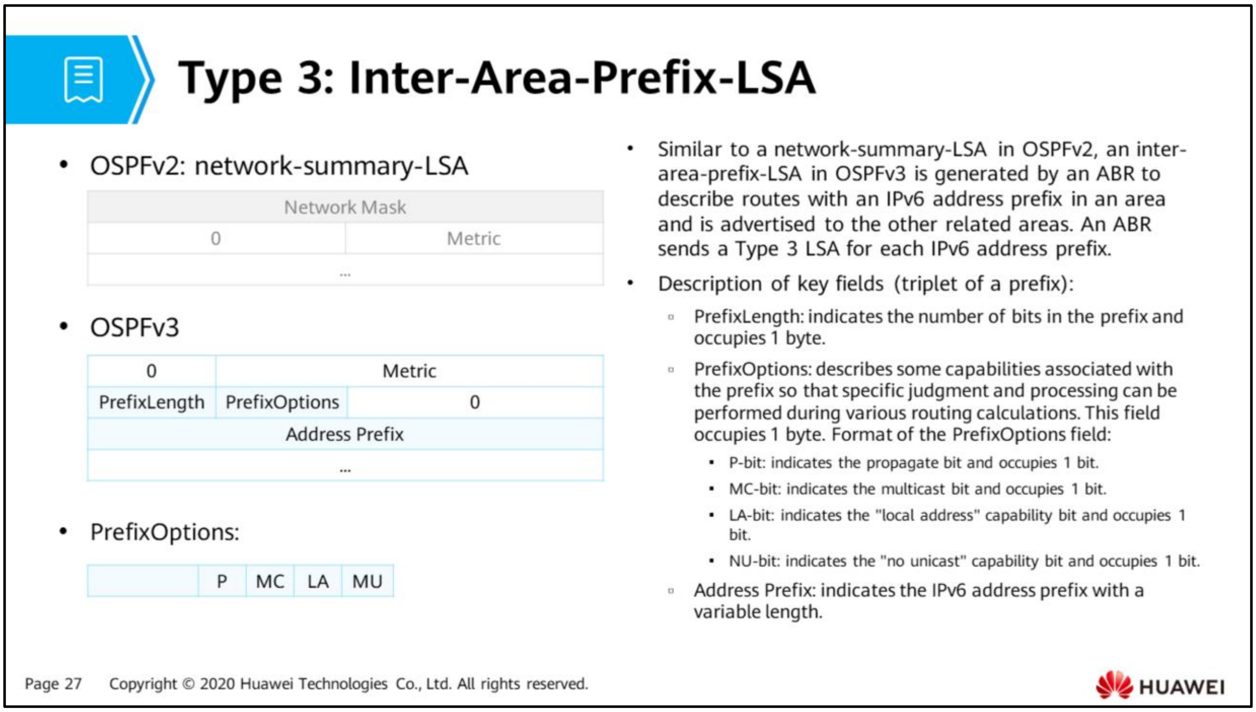

- The fields in an OSPFv3 inter-area-prefix-LSA are described as follows:
  - Metric: indicates the cost of the route to the destination address and occupies 3 bytes.
  - PrefixOptions: Each prefix advertised by an LSA has its own PrefixOptions field.
    - P-bit: propagate bit. This bit needs to be set to 1 if the prefix of an NSSA needs to be advertised by an ABR.
    - MC-bit: multicast bit. If this bit is set to 1, the prefix is used for multicast route calculation. Otherwise, the prefix is not used for multicast route calculation.
    - LA-bit: local address capability bit. If this bit is set to 1, the prefix is an interface address of the router.
    - NU-bit: no unicast capability bit. If this bit is set to 1, the prefix is not used for IPv6 unicast route calculation.
- Note: The prefix length of the default route is 0. An ABR can also originate an inter-area Type 3 LSA to advertise a default route to a stub area.

Type 4: Inter-Area-Router-LSA

- OSPFv2: ASBR-summary-LSA

Network Mask	
0	Metric
...	

- OSPFv3

0	Options
0	Metric
Destination Router ID	

- Similar to a Type 4 LSA in OSPFv2, a Type 4 LSA in OSPFv3 is generated by an ABR, describes the route to an ASBR, and is advertised to areas excluding that to which the ASBR belongs. An ABR sends a Type 4 LSA to describe each ASBR.
- Description of key fields:
  - Destination Router ID: indicates the router ID of the destination router being described by the LSA, that is, the router ID of an ASBR. This field occupies 4 bytes.

HUAWEI

- The fields in an OSPFv3 inter-area-router-LSA are described as follows:
  - Options: This field describes the optional capabilities supported by the destination router instead of those supported by the source router. Therefore, the value of this field should equal that of the Options field in the router-LSA generated by the destination router.
  - Metric: indicates the cost of the route to the destination address and occupies 3 bytes.

Type 5: AS-External-LSA

- OSPFv2

Network Mask		
E	0	Metric
Forwarding Address		
External Route Tag		
...		

- OSPFv3

0	E F T	Metric
PrefixLength	PrefixOptions	Referenced LS Type
Address Prefix		
Forwarding Address (Optional)		
External Route Tag (Optional)		
Referenced Link State ID (Optional)		

- Similar to a Type 5 LSA in OSPFv2, a Type 5 LSA in OSPFv3 is generated by an ASBR, describes the route with a prefix outside the local AS, and is advertised to all areas except the stub area and NSSA.
- Description of key fields:
  - Referenced LS Type: indicates whether the LSA needs to reference another LSA and occupies 2 bytes.
    - 0: does not reference any other LSA.
    - 1: references a router-LSA.
    - 2: references a network-LSA.

- The fields in an OSPFv3 AS-external-LSA are described as follows:
  - Bit E: indicates the cost type of an AS external route and occupies 1 bit.
    - The value 1 indicates the cost of a Type 2 external route. This cost does not increase during route transmission.
    - The value 0 indicates the cost of a Type 1 external route. This cost increases during route transmission.
  - Bit F: occupies 1 bit. The value 1 indicates that the Forwarding Address field (optional) is included.
  - Bit T: occupies 1 bit. The value 1 indicates that the External Route Tag field (optional) is included.
  - Metric: indicates the cost of the route to the destination address and occupies 3 bytes.
  - PrefixLength, PrefixOptions, and Address Prefix are triplets that describe a prefix and have the same meanings as those in an inter-area-prefix-LSA.
  - Forwarding Address: is an optional 128-bit IPv6 address and occupies 4 bytes. This field is included if bit F is 1. In this case, a data packet needs to be forwarded to this address before reaching its destination.

Type 8: Link-LSA (Added in OSPFv3)

Link-LSA format

Rtr Pri	Options	
Link-Local Interface Address		
# Prefixes		
PrefixLength	PrefixOptions	0
Address Prefix ...		
...		
PrefixLength	PrefixOptions	0
Address Prefix ...		

- Each device generates a link-LSA for each attached link and floods the link-LSA only on the originating link.
- Link-LSA functions:
  - Advertises the link-local address of the local interface to other routers attached to the originating link.
  - Advertises a list of IPv6 prefixes of the local interface to other routers attached to the originating link.
  - Advertises a collection of Options bits in the network-LSA originated on the local link to other routers attached to the originating link.
- Description of key fields:
  - Link-Local Interface Address: indicates the link-local address of the interface that connects the originating router to the link, and occupies 16 bytes. This address is contained only in the Link-LSA.

- The fields in an OSPFv3 link-LSA are described as follows:
  - Rtr Pri: indicates the router priority of the interface attaching the originating router to the link and occupies 1 byte.
  - Options: indicates a collection of Options bits that the router sets in the network-LSA and occupies 3 bytes.
  - Number of Prefixes: indicates the number of IPv6 address prefixes carried in the LSA, and occupies 4 bytes.
  - PrefixLength, PrefixOptions, and Address Prefix are triplets that describe a prefix and have the same meanings as those in an inter-area-prefix-LSA.

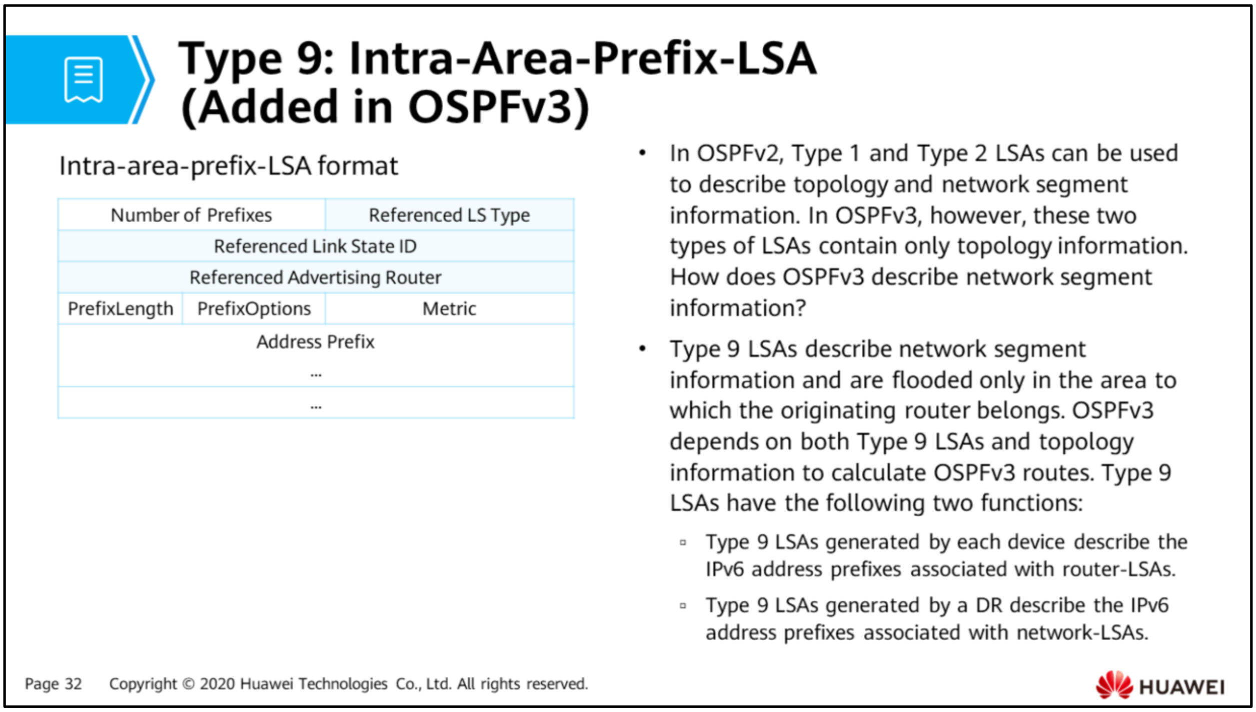

- The fields in an OSPFv3 intra-area-prefix-LSA are described as follows:
  - Number of Prefixes: indicates the number of IPv6 address prefixes carried in the LSA, and occupies 4 bytes. If necessary, prefixes can be carried in multiple intra-area-prefix-LSAs to limit the size of each Type 9 LSA.
  - Referenced LS type: indicates whether the LSA references a router-LSA or a network-LSA, and occupies 4 bytes.
    - 1: A router-LSA is referenced.
    - 2: A network-LSA is referenced.
  - Referenced Link State ID: 4 bytes.
    - If the LSA references a router-LSA, this field is set to 0.
    - If the LSA references a network-LSA, this field is set to the interface ID of the DR on the attached link.
  - Referenced Advertising Router: 4 bytes.
    - If the LSA references a router-LSA, this field is set to the router ID of the associated router.
    - If the LSA references a network-LSA, this field is set to the router ID of the DR on the attached link.

Examples of OSPFv3 LSAs

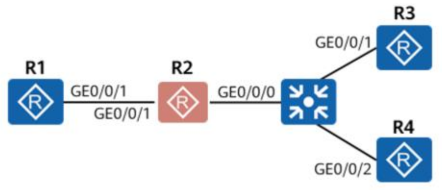

Device	Router ID	Interface	IP Address
R1	10.1.1.1	GE0/0/1	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	2001:DB8:2345:12::2/64
		GE0/0/0	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	2001:DB8:2345:23::3/64
R4	10.1.4.4	GE0/0/2	2001:DB8:2345:23::4/64

```
[R2]display ospfv3 lsdb
    * indicates STALE LSA

        OSPFv3 Router with ID (10.1.2.2) (Process 1)
                                                                Type 8 LSA
            Link-LSA (Interface GigabitEthernet0/0/0)
Link State ID   Origin Router   Age   Seq#          CkSum   Prefix
0.0.0.3         10.1.2.2        0080  0x80000002 0x1c52     1
0.0.0.4         10.1.3.3        0075  0x80000002 0x2a59     1
0.0.0.5         10.1.4.4        0089  0x80000002 0xe2a3     1

            Link-LSA (Interface GigabitEthernet0/0/1)
Link State ID   Origin Router   Age   Seq#          CkSum   Prefix
0.0.0.4         10.1.1.1        0070  0x80000002 0xe971     1
0.0.0.4         10.1.2.2        0069  0x80000002 0x9fdd     1

            Router-LSA (Area 0.0.0.0)
...
            Network-LSA (Area 0.0.0.0)
Link State ID   Origin Router   Age   Seq#          CkSum
0.0.0.4         10.1.2.2        0030  0x80000001 0x9564
0.0.0.5         10.1.4.4        0036  0x80000002 0xc810
                                                                Type 9 LSA
            Intra-Area-Prefix-LSA (Area 0.0.0.0)
Link State ID   Origin Router   Age   Seq#          CkSum   Prefix Reference
0.0.0.1         10.1.2.2        0028  0x80000006 0x1527     1    Network-LSA
0.0.0.1         10.1.4.4        0034  0x80000003 0xfc28     1    Network-LSA
```


- As shown in the figure, R1, R2, R3, and R4 run OSPFv3 and are all deployed in the backbone area.
- After the network becomes stable, check the LSDB of R2. The command output shows information about the following types of LSAs: router-LSA (Type 1), network-LSA (Type 2), Link-LSA (Type 8), and intra-area-prefix-LSA (Type 9).

Link-LSA Example

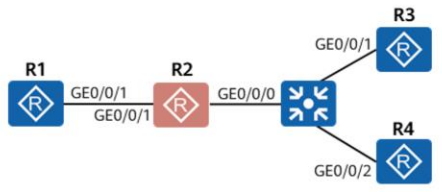

Device	Router ID	Interface	IP Address
R1	10.1.1.1	GE0/0/1	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	2001:DB8:2345:12::2/64
		GE0/0/0	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	2001:DB8:2345:23::3/64
R4	10.1.4.4	GE0/0/2	2001:DB8:2345:23::4/64

Check the link-LSA with the Link State ID of 0.0.0.5 on R2. The command output is as follows:

```
[R2]display ospfv3 lsdb link 0.0.0.5
        OSPFv3 Router with ID (10.1.2.2) (Process 1)
          Link-LSA (Interface GigabitEthernet0/0/0)
LS Age: 446
LS Type: Link-LSA
Link State ID: 0.0.0.5
Originating Router: 10.1.4.4
LS Seq Number: 0x80000002
Retransmit Count: 0
Checksum: 0xE2A3
Length: 56
Priority: 1
Options: 0x000013 (-|R|-|-|E|V6)
Link-Local Address: FE80::2E0:FCFF:FEAC:6C5
Number of Prefixes: 1
  Prefix: 2001:DB8:2345:23::/64
  Prefix Options: 0 (-|-|-|-|-)
```

LSA header

Type 8 LSA

Link-local address

IPv6 address prefix of the network segment where the physical interface resides

HUAWEI

- The command output is described as follows:
  - LS Age: aging time of the LSA.
  - LS Type: LSA type, which can be any of the following:
    - Router-LSA, Network-LSA, Inter-Area-Prefix-LSA, Inter-Area-Router-LSA, AS-external-LSA, NSSA-LSA, Link-LSA, or Intra-Area-Prefix-LSA
  - Link State ID: link state ID in the LSA header.
  - Originating Router: router that generates the LSA.
  - LS Seq Number: sequence number of the LSA. This field is carried in the LSA header.
  - Checksum: checksum of the LSA.
  - Length: length of the LSA.
  - Priority: priority of the interface attached to the link.
  - Options: optional capabilities of the link.
  - Link-Local Address: link-local address.
  - Number of Prefixes: number of IPv6 prefixes contained in the LSA.
  - Prefix: IPv6 prefix.
  - Prefix Options: optional capabilities associated with the prefix.

Intra-Area-Prefix-LSA Example

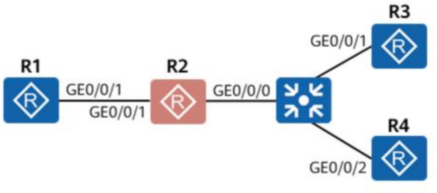

Device	Router ID	Interface	IP Address
R1	10.1.1.1	GE0/0/1	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	2001:DB8:2345:12::2/64
		GE0/0/0	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	2001:DB8:2345:23::3/64
R4	10.1.4.4	GE0/0/2	2001:DB8:2345:23::4/64

Check the intra-area-prefix-LSA on R2. The command output is as follows:

```
[R2]display ospfv3 lsdb intra-prefix

        OSPFv3 Router with ID (10.1.2.2) (Process 1)

            Intra-Area-Prefix-LSA (Area 0.0.0.0)
...

LS Age: 556
LS Type: Intra-Area-Prefix-LSA
Link State ID: 0.0.0.1
Originating Router: 10.1.4.4
LS Seq Number: 0x80000003
Retransmit Count: 0
Checksum: 0xFC28
Length: 44
Number of Prefixes: 1
Referenced LS Type: 0x2002
Referenced Link State ID: 0.0.0.5
Referenced Originating Router: 10.1.4.4

 Prefix: 2001:DB8:2345:23::/64
  Prefix Options: 0 (-|-|-|-|-)
   Metric: 0
```

LSA header

Type 9 LSA

Reference the network-LSA generated by the DR (R4)

IPv6 address prefix of the transit network segment to which the local router is directly connected. This LSA is generated by the DR.

HUAWEI

Contents

HUAWEI

Basic OSPFv3 Configuration Commands (1)

1. Start OSPFv3.

[Huawei] **ospfv3** [*process-id*] [**vpn-instance** *vpn-instance-name*]

An OSPFv3 process is created and runs, and (optional) the process is bound to a VPN instance.

[Huawei-ospfv3-1] **router-id** *router-id*

A router ID is configured for the local device in this OSPFv3 process.
Note: If no router ID is configured, the OSPFv3 process cannot run.

2. Enable OSPFv3 on an interface.

[Huawei-GigabitEthernet0/0/1] **ospfv3** *process-id* **area** *area-id* [**instance** *instance-id*]

The interface is enabled in the specified OSPFv3 process, and the area to which the interface belongs is specified. You can also specify the ID of the instance to which the interface belongs.
Note: Before running this command, you must create an OSPFv3 process and enable the IPv6 function.

HUAWEI

- The configuration commands and methods of OSPFv3 are similar to those of OSPFv2. For details, see the "HCIP-Datacom-Core Technology" course.

Basic OSPFv3 Configuration Commands (2)

3. (Optional) Configure an OSPFv3 network type for the interface.

   [Huawei-GigabitEthernet0/0/1] **ospfv3 network-type** { **broadcast** | **nbma** | **p2mp** [**non-broadcast**] | **p2p** } [**instance** *instance-id*]

   By default, the OSPFv3 network type of a physical interface depends on the data link layer encapsulation of the interface. The default network type of an Ethernet interface is broadcast, whereas that of a serial interface (encapsulated with PPP or HDLC) is P2P.

4. Enter the OSPFv3 area view.

   [Huawei-ospfv3-1] **area** *area-id*

   The specified area ID can be a decimal integer or in the IPv4 address format. Regardless of the specified format, the area ID is displayed as an IPv4 address.

HUAWEI

Verifying the Configuration of Basic OSPFv3 Functions

1. Check the OSPFv3 interface information.

   [Huawei] **display ospfv3** [*process-id*] **interface** [**area** *area-id*] [*interface-type interface-number*]

2. Check the OSPFv3 neighbor information.

   [Huawei] **display ospfv3** [*process-id*] [**area** *area-id*] **peer** [*interface-type interface-number* | *neighbor-id*] [**verbose**]

   neighbor-id: specifies the router ID of a neighbor.

3. Check the OSPFv3 LSDB.

   [Huawei] **display ospfv3** [*process-id*] **lsdb** [**area** *area-id*] [**originate-router** *advertising-router-id* | **self-originate**] [{ **router** | **network** | **inter-router** [**asbr-router** *asbr-router-id*] | { **inter-prefix** | **nssa** } [*ipv6-address prefix-length*] | **link** | **intra-prefix** } [*link-state-id*]]

4. Check the OSPFv3 routing table.

   [Huawei] **display ospfv3** [*process-id*] **routing** [*ipv6-address prefix-length* | **abr-routes** | **asbr-routes** | **intra-routes** | **inter-routes** | **ase-routes** | **nssa-routes** | [**statistics**]]

HUAWEI

- [Huawei] **display ospfv3** [*process-id*] **lsdb** [**area** *area-id*] [**originate-router** *advertising-router-id* | **self-originate**] [{ **router** | **network** | **inter-router** [**asbr-router** *asbr-router-id*] | { **inter-prefix** | **nssa** } [*ipv6-address prefix-length*] | **link** | **intra-prefix** | **grace** } [*link-state-id*]]
  - *process-id*: specifies the ID of an OSPFv3 process. The value is an integer ranging from 1 to 65535.
  - **area** *area-id*: specifies the ID of an area. The area ID can be a decimal integer or in the IPv4 address format. For a decimal integer, the value ranges from 0 to 4294967295. For the IPv4 address format, the value is in dotted decimal notation.
  - **external**: displays information about AS-external LSAs in the LSDB.
  - **inter-prefix**: displays information about inter-area-prefix LSAs in the LSDB.
  - **inter-router**: displays information about inter-area-router LSAs in the LSDB.
  - **intra-prefix**: displays information about intra-area-prefix LSAs in the LSDB.
  - **nssa**: displays information about NSSA LSAs in the LSDB.
  - **link**: displays information about link-LSAs in the LSDB.
  - **network**: displays information about network-LSAs in the LSDB.
  - **router**: displays information about router-LSAs in the LSDB.
  - *link-state-id*: specifies a link state ID. The value is in dotted decimal notation.

Example for Configuring OSPF IPv4/IPv6 Dual Stack

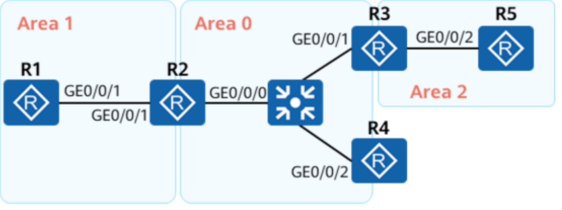

Device	Router ID	Interface	IPv4 Address	IPv6 Address
R1	10.1.1.1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
		GE0/0/0	10.1.23.2/24	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	10.1.23.3/24	2001:DB8:2345:23::3/64
		GE0/0/2	10.1.35.3/24	2001:DB8:2345:35::3/64
R4	10.1.4.4	GE0/0/2	10.1.23.4/24	2001:DB8:2345:23::4/64
R5	10.1.5.5	GE0/0/2	10.1.35.5/24	2001:DB8:2345:35::5/64

Scenario description:

- A company has OSPFv2 deployed over an IPv4 network to implement interconnection and communication. To ensure future service development, the company has an IPv6 network deployed for service testing. OSPFv3 runs on the IPv6 network to implement IPv6 interconnection and communication.
- Each router runs both OSPFv2 and OSPFv3, and the entire AS is divided into three areas. After the configuration is complete, each router should learn the IPv4 and IPv6 routes to all network segments in the AS.

Deploying the IPv4 Network (1)

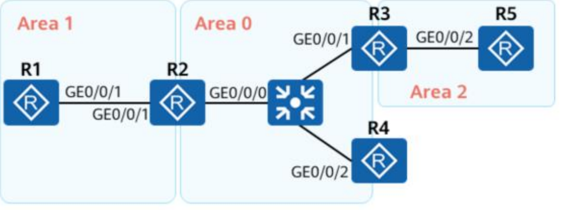

Device	Router ID	Interface	IPv4 Address	IPv6 Address
R1	10.1.1.1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
		GE0/0/0	10.1.23.2/24	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	10.1.23.3/24	2001:DB8:2345:23::3/64
		GE0/0/2	10.1.35.3/24	2001:DB8:2345:35::3/64
R4	10.1.4.4	GE0/0/2	10.1.23.4/24	2001:DB8:2345:23::4/64
R5	10.1.5.5	GE0/0/2	10.1.35.5/24	2001:DB8:2345:35::5/64

1. Assign an IPv4 address to each router interface. (The configuration details are not provided here.)
2. Configure basic OSPF functions.

```
[R1] ospf 1 router-id 10.1.1.1
[R1-ospf-1] area 1
[R1-ospf-1-area-0.0.0.0] network 10.1.12.1 0.0.0.0
```

```
[R2] ospf 1 router-id 10.1.2.2
[R2-ospf-1] area 1
[R2-ospf-1-area-0.0.0.1] network 10.1.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.1] quit
[R2-ospf-1] area 0
[R2-ospf-1-area-0.0.0.0] network 10.1.23.2 0.0.0.0
```

```
[R3] ospf 1 router-id 10.1.3.3
[R3-ospf-1] area 0
[R3-ospf-1-area-0.0.0.0] network 10.1.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0] quit
[R3-ospf-1] area 2
[R3-ospf-1-area-0.0.0.2] network 10.1.35.3 0.0.0.0
```

Deploying the IPv4 Network (2)

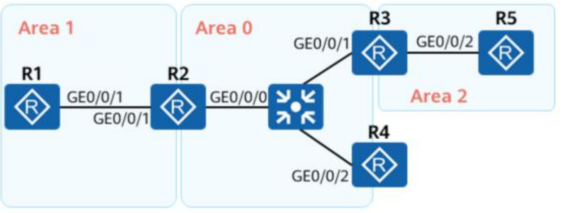


```
[R4] ospf 1 router-id 10.1.4.4
[R4-ospf-1] area 0
[R4-ospf-1-area-0.0.0.0] network 10.1.23.4 0.0.0.0
```

```
[R5] ospf 1 router-id 10.1.5.5
[R5-ospf-1] area 2
[R5-ospf-1-area-0.0.0.2] network 10.1.35.5 0.0.0.0
```

Device	Router ID	Interface	IPv4 Address	IPv6 Address
R1	10.1.1.1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
		GE0/0/0	10.1.23.2/24	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	10.1.23.3/24	2001:DB8:2345:23::3/64
		GE0/0/2	10.1.35.3/24	2001:DB8:2345:35::3/64
R4	10.1.4.4	GE0/0/2	10.1.23.4/24	2001:DB8:2345:23::4/64
R5	10.1.5.5	GE0/0/2	10.1.35.5/24	2001:DB8:2345:35::5/64

HUAWEI

Deploying the IPv6 Network (1)

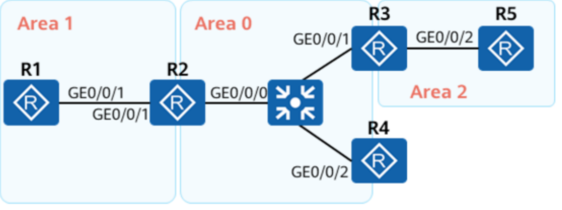

Device	Router ID	Interface	IPv4 Address	IPv6 Address
R1	10.1.1.1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
		GE0/0/0	10.1.23.2/24	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	10.1.23.3/24	2001:DB8:2345:23::3/64
		GE0/0/2	10.1.35.3/24	2001:DB8:2345:35::3/64
R4	10.1.4.4	GE0/0/2	10.1.23.4/24	2001:DB8:2345:23::4/64
R5	10.1.5.5	GE0/0/2	10.1.35.5/24	2001:DB8:2345:35::5/64

1. Enable IPv6 globally and on interfaces, and assign IPv6 addresses to interfaces on each router. (The configuration details are not provided here.)
2. Enable OSPFv3.

```
[R1] ospfv3 1
[R1-ospfv3-1] router-id 10.1.1.1
```

```
[R2] ospfv3 1
[R2-ospfv3-1] router-id 10.1.2.2
```

```
[R3] ospfv3 1
[R3-ospfv3-1] router-id 10.1.3.3
```

```
[R4] ospfv3 1
[R4-ospfv3-1] router-id 10.1.4.4
```

```
[R5] ospfv3 1
[R5-ospfv3-1] router-id 10.1.5.5
```


HUAWEI

Deploying the IPv6 Network (2)

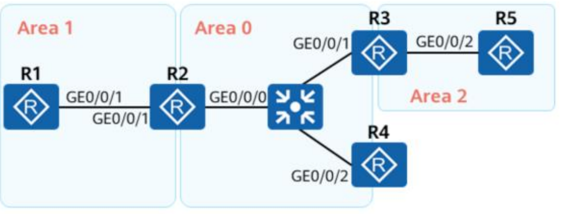

Device	Router ID	Interface	IPv4 Address	IPv6 Address
R1	10.1.1.1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
R2	10.1.2.2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
		GE0/0/0	10.1.23.2/24	2001:DB8:2345:23::2/64
R3	10.1.3.3	GE0/0/1	10.1.23.3/24	2001:DB8:2345:23::3/64
		GE0/0/2	10.1.35.3/24	2001:DB8:2345:35::3/64
R4	10.1.4.4	GE0/0/2	10.1.23.4/24	2001:DB8:2345:23::4/64
R5	10.1.5.5	GE0/0/2	10.1.35.5/24	2001:DB8:2345:35::5/64

3. Enable OSPFv3 on interfaces.

```
[R1] interface gigabitethernet 0/0/1
[R1-GigabitEthernet0/0/1] ospfv3 1 area 1
```

```
[R2] interface gigabitethernet 0/0/1
[R2-GigabitEthernet0/0/1] ospfv3 1 area 1
[R2-GigabitEthernet0/0/1] quit
[R2] interface gigabitethernet 0/0/0
[R2-GigabitEthernet0/0/0] ospfv3 1 area 0
```

```
[R3] interface gigabitethernet 0/0/1
[R3-GigabitEthernet0/0/1] ospfv3 1 area 0
[R3-GigabitEthernet0/0/1] quit
[R3] interface gigabitethernet 0/0/2
[R3-GigabitEthernet0/0/2] ospfv3 1 area 2
```

```
[R4] interface gigabitethernet 0/0/2
[R4-GigabitEthernet0/0/2] ospfv3 1 area 0
```

```
[R5] interface gigabitethernet 0/0/2
[R5-GigabitEthernet0/0/2] ospfv3 1 area 2
```

Checking Neighbor Information on the OSPFv3 Network

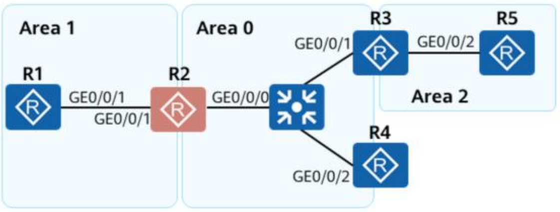


```
[R2]display ospfv3 peer
OSPFv3 Process (1)
OSPFv3 Area (0.0.0.0)
Neighbor ID     Pri  State          Dead Time Interface     Instance ID
10.1.3.3          1  Full/Backup    00:00:31    GE0/0/0       0
10.1.4.4          1  Full/DR        00:00:36    GE0/0/0       0
OSPFv3 Area (0.0.0.1)
Neighbor ID     Pri  State          Dead Time Interface     Instance ID
10.1.1.1          1  Full/Backup    00:00:37    GE0/0/1       0
```

If no instance ID is specified, the default value 0 is used.
The instance IDs of the directly connected interfaces on the two devices must be the same. Otherwise, the OSPFv3 neighbor relationship cannot be established.

HUAWEI

- You can run the **display ospf peer** command to check OSPFv2 neighbor information.
- By comparing OSPFv2 neighbor information and OSPFv3 neighbor information, you will find that the elected DR and BDR are the same, indicating that the DR election modes between OSPFv2 and OSPFv3 are the same.

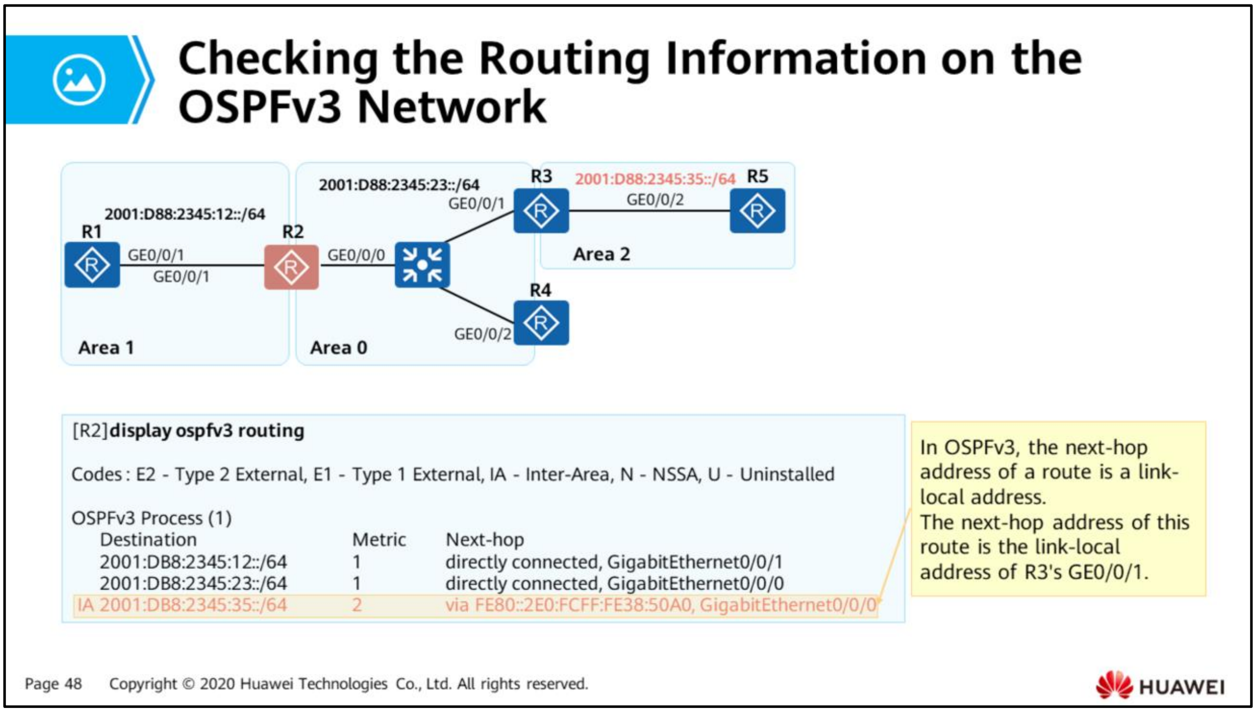

- You can run the **display ospf routing** command to check the routing information on the OSPFv2 network.
- By comparing OSPFv2 routing information and OSPFv3 routing information, you will find that the paths to the same network segment are the same, indicating that the route calculation methods between OSPFv2 and OSPFv3 are the same.

Checking the LSDB on the OSPFv3 Network

```
[R2]display ospfv3 lsdb
* indicates STALE LSA
        OSPFv3 Router with ID (10.1.2.2) (Process 1)
           Link-LSA (Interface GigabitEthernet0/0/0)
Link State ID Origin Router    Age   Seq#           CkSum  Prefix
0.0.0.3       10.1.2.2         0862  0x80000003     0x486d   1
0.0.0.4       10.1.3.3         0854  0x80000003     0xc512   1
0.0.0.5       10.1.4.4         0840  0x80000003     0x4e3c   1
           Link-LSA (Interface GigabitEthernet0/0/1)
Link State ID Origin Router    Age   Seq#           CkSum  Prefix
0.0.0.4       10.1.1.1         0878  0x80000003     0x9d8f   1
0.0.0.4       10.1.2.2         0868  0x80000003     0xcbf8   1
           Router-LSA (Area 0.0.0.0)
Link State ID Origin Router    Age   Seq#           CkSum  Link
0.0.0.0       10.1.2.2         0567  0x8000000f     0xd411   1
0.0.0.0       10.1.3.3         0576  0x8000000e     0xd70c   1
0.0.0.0       10.1.4.4         0577  0x8000000a     0xdd08   1
           Network-LSA (Area 0.0.0.0)
Link State ID Origin Router    Age   Seq#           CkSum
0.0.0.5       10.1.4.4         0577  0x80000006     0xc014
           Inter-Area-Prefix-LSA (Area 0.0.0.0)
Link State ID Origin Router    Age   Seq#           CkSum
0.0.0.1       10.1.2.2         0862  0x80000003     0xceb3
0.0.0.1       10.1.3.3         0782  0x80000004     0x3824
           Intra-Area-Prefix-LSA (Area 0.0.0.0)
Link State ID Origin Router    Age   Seq#           CkSum Prefix Reference
0.0.0.1       10.1.4.4         0576  0x8000000a     0xee2f    1   Network-LSA
```

```
           Router-LSA (Area 0.0.0.1)
Link State ID Origin Router    Age   Seq#       CkSum    Link
0.0.0.0       10.1.1.1         0819  0x80000007 0x9460     1
0.0.0.0       10.1.2.2         0827  0x80000007 0x8a67     1
           Network-LSA (Area 0.0.0.1)
Link State ID Origin Router    Age   Seq#       CkSum
0.0.0.4       10.1.2.2         0828  0x80000003 0x9166
           Inter-Area-Prefix-LSA (Area 0.0.0.1)
Link State ID Origin Router    Age   Seq#       CkSum
0.0.0.1       10.1.2.2         0570  0x80000001 0x510f
0.0.0.2       10.1.2.2         0596  0x80000001 0xfb76
           Intra-Area-Prefix-LSA (Area 0.0.0.1)
Link State ID Origin Router    Age   Seq#       CkSum  Prefix  Reference
0.0.0.1       10.1.2.2         0826  0x80000004 0x1925   1     Network-LSA
```

The LSDB on R2 contains the Type 1, Type 2, Type 3, Type 8, and Type 9 LSAs.

HUAWEI

- You can run the **display ospf lsdb** command to check the OSPFv2 LSDB. You will find that the LSDB contains the Type 1, Type 2, and Type 3 LSAs.

Contents

IS-IS (IPv6) Overview

- IS-IS is a link-state dynamic routing protocol, which was initially designed for the OSI network. Later, to support IPv4 routing, extensions were made to apply IS-IS to IPv4 networks. This is called integrated IS-IS.
- IS-IS packets are classified into Hello protocol data units (PDUs), link state PDUs (LSPs), and sequence number PDUs (SNPs).
  - Packet format

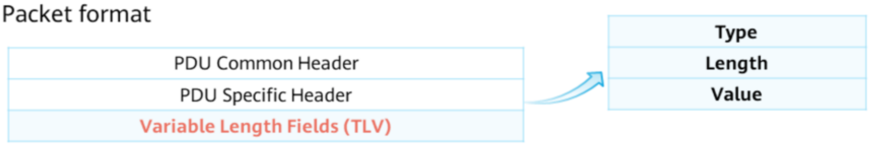

  - The Variable Length Fields (TLV) part consists of multiple TLV triplets. Using the TLV structure to construct packets improves the flexibility and scalability of IS-IS. To support new features, only new TLVs need to be added.
- To support the processing and calculation of IPv6 routes, IS-IS adds two TLVs and one network layer protocol identifier (NLPID).

New TLVs

TLV 232 (IPv6 Interface Address)

It is equivalent to TLV 132 (used to describe IPv4 interface addresses). The only difference is that TLV 232 describes 128-bit IPv6 addresses instead of 32-bit IPv4 addresses.

Type=232	Length	Interface Address 1
Interface Address 1		
Interface Address 1		
Interface Address 1		
Interface Address 1		Interface Address 2
...		

The content of the Interface Address field varies according to different PDUs:

- For Hello PDUs, the Interface Address TLVs can contain only the link-local IPv6 addresses assigned to the interfaces that send the Hello PDUs.
- For LSPs, the Interface Address TLVs can contain only the non-link-local IPv6 addresses assigned to the device interfaces.

TLV 236 (IPv6 Reachability)

It is equivalent to TLVs 128 and 130, and uses the X-bit to identify internal or external reachability information.

Type=236	Length	Metric				
Metric		U	X	S	R	Prefix Length
Prefix						
Sub-TLV Len	Sub-TLV					

The IPv6 Reachability TLV may appear any number of times (including 0) within an LSP. Link-local prefixes are not advertised using this TLV.

- Fields in TLV 232 (IPv6 Interface Address) are described as follows:
  - Type: indicates the TLV type and occupies 8 bits. The value is 232 (0xE8).
  - Length: indicates the length of the Value field in the TLV and occupies 8 bits.
  - Interface Address: indicates a 128-bit IPv6 address.
- Fields in TLV 236 (IPv6 Reachability) are described as follows:
  - Type: indicates the TLV type and occupies 8 bits. The value is 236 (0xEC).
  - Length: indicates the length of the Value field in the TLV and occupies 8 bits.
  - Metric: a 32-bit field indicating the cost.
  - U: up/down bit. This 1-bit field indicates whether the prefix is advertised from a higher level to a lower level.
  - X: external original bit. This 1-bit field indicates whether the prefix was imported into IS-IS from another routing protocol.
  - S: sub-TLV present bit. This 1-bit field is optional.
  - R: This 5-bit field is reserved.
  - Prefix Length: indicate the prefix length and occupies 8 bits.
  - Prefix: indicates an IPv6 address prefix.
  - Sub-TLV Length: indicates the length of a sub-TLV and occupies 8 bits. If the S-bit is set to 1, this field is included.
  - Sub-TLV: If the S-bit is set to 1, this field is included.

New NLPID in TLV 129

- To support the processing and calculation of IPv6 routes, IS-IS adds a new NLPID to TLV 129.
- TLV 129 (Protocol Supported)

Type=129	Length	NLPID	NLPID
...			

Field description:

- Type: indicates the TLV type and occupies 8 bits. The value is 129 (0x81), indicating the Protocol Supported TLV.
- Length: indicates the length of the Value field in the TLV and occupies 8 bits.
- NLPID: indicates the network layer protocol identifier and occupies 8 bits. If IPv4 is supported, the value is 204 (0xCC). If IPv6 is supported, the value is 142 (0x8E).

- If IS-IS supports IPv6, the IPv6 routing information that it advertises must carry the IPv6 NLPID.

IS-IS Multi-Topology Background

By default, on an IS-IS network, a mix of IPv4 and IPv6 topologies is considered as an integrated topology, and IS-IS calculates the same SPT for IPv4 and IPv6.

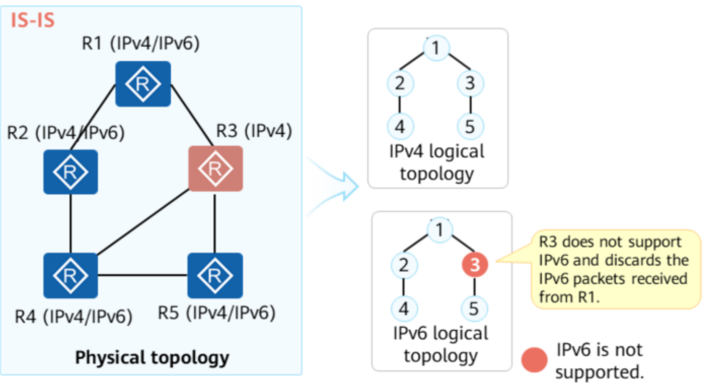

IS-IS single-topology has the following problems:

- In a mixed topology, routers that support the IPv4/IPv6 dual stack cannot detect the routers or links that do not support IPv6 and may still forward IPv6 packets to them. As a result, the IPv6 packets are discarded due to a forwarding failure.
- Similarly, if routers and links that do not support IPv4 exist in the topology, IPv4 packets are discarded.

- IS-IS single-topology has the following disadvantages:
  - Network deployment is not suitable for topology separation.
  - To maintain the same topology, each interface must run both IS-IS (IPv4) and IS-IS (IPv6), which is not flexible.
  - IPv4 areas cannot be used to connect different IPv6 areas. That is, IPv4 networks cannot be used to address IPv6 network isolation.

IS-IS Multi-Topology Overview

- The IS-IS multi-topology (MT) feature indicates that multiple independent IP topologies run in an IS-IS domain. For example, IPv4 and IPv6 topologies are not considered as a single integrated topology. This ensures that IS-IS routes are calculated for both IPv4 and IPv6 networks. Based on the IP versions supported by links, separate SPF calculation is performed in each topology to shield networks from each other.
- IS-IS MT implementation
  - Topology establishment: Neighbor relationships are established through packet exchange to set up multiple topologies.
  - SPF calculation: SPF calculation is performed in each topology.

- The IS-IS MT feature can overcome the disadvantages of IS-IS single topology.

IS-IS MT Implementation

- IS-IS defines new TLVs, which contain MT information, that is, information about the topologies to which interfaces belong. The advertisement of MT information enables the network to perform separate SPF calculations for different topologies, implementing topology separation.
- Multi-Topology TLV 229

Type=229	Length	O	A	R	MT ID
...					

- Key fields are described as follows:
    - Type: indicates the TLV type and occupies 8 bits. The value is 229 (0xE5), indicating support for MT.
    - O: Overload bit, which occupies 1 bit.
    - A: Attach bit, which occupies 1 bit.
    - MT ID: indicates the topology to which the interface belongs, and occupies 12 bits.

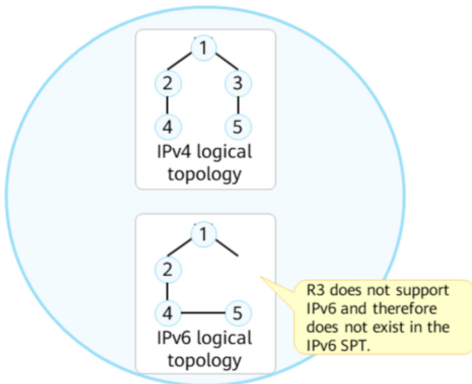

IS-IS calculates a separate topology for each of the IPv4 and IPv6 networks.

HUAWEI

- To support MT, IS-IS defines multiple types of TLVs, including Multi-Topology TLV, MT Intermediate Systems TLV, Multi-Topology Reachable IPv4 Prefixes TLV, and Multi-Topology Reachable IPv6 Prefixes TLV. This course focuses on the Multi-Topology TLV and does not elaborate on the other ones.
- Multi-Topology TLV:
    - This TLV is contained only in IIH PDUs and fragment zero LSPs.
    - Reserved MT IDs:
        - MT ID=0, equivalent to the standard IPv4 topology.
        - MT ID=2, reserved for the IPv6 routing topology.

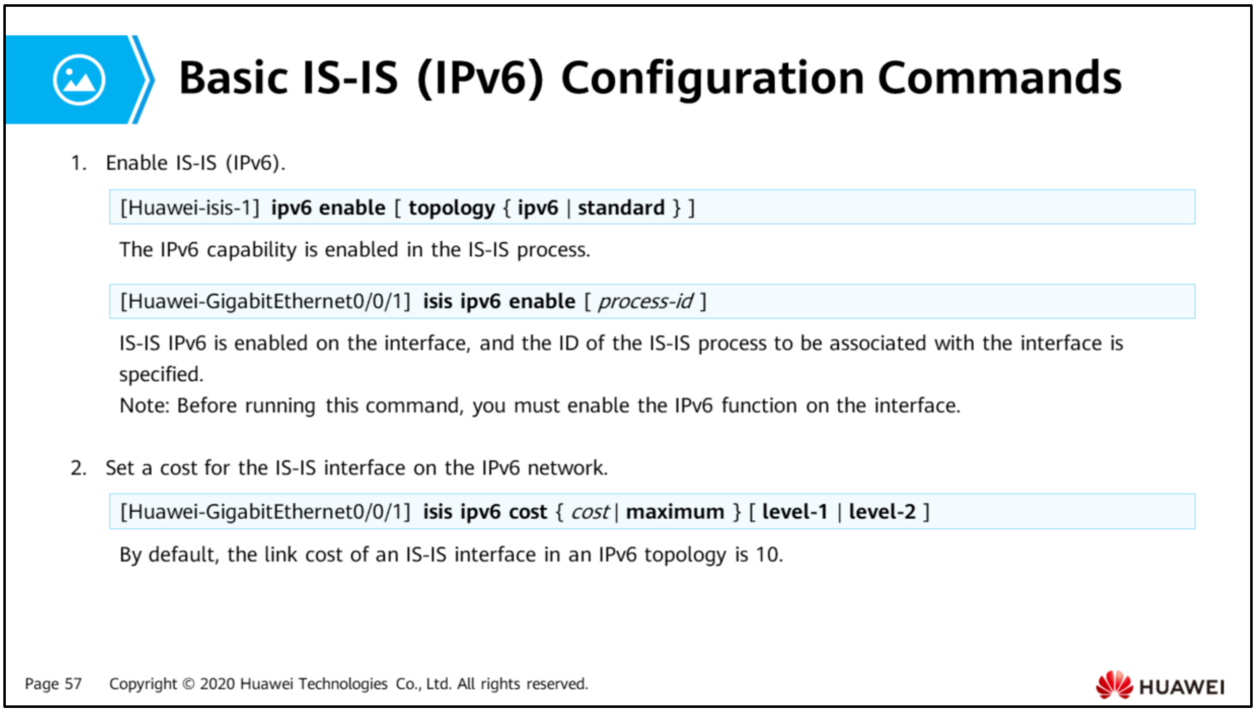

- The basic configuration commands and methods of IS-IS (IPv6) are the same as those of IS-IS (IPv4). For details, see the "HCIP-Datacom-Core Technology" course.
- [Huawei-isis-1] **ipv6 enable** [**topology** { **ipv6** | **standard** }]
  - **topology**: sets a network topology type.
  - **ipv6**: sets the topology type to IPv6. That is, the IPv6 capability for the IS-IS process is enabled in an IPv6 topology. Links on the network can be configured as IPv4 or IPv6 links. SPF calculation is performed separately in IPv4 and IPv6 topologies.
  - **standard**: sets the topology type to standard. That is, the IPv6 capability for the IS-IS process is enabled in an integrated topology. A network administrator must ensure that all links on the network support the same topology type. By default, the **standard** type is used when the IPv6 capability is enabled for an IS-IS process.
- [Huawei-GigabitEthernet0/0/1] **isis ipv6 cost** { *cost* | **maximum** } [**level-1** | **level-2**]
  - *cost*: specifies the link cost of an IPv6 interface. The value is an integer that varies according to the following cost styles:
    - If the cost style is narrow, narrow-compatible, or compatible, the value ranges from 1 to 63.
    - If the cost style is wide or wide-compatible, the value ranges from 1 to 16777214.
    - The default value is 10.

Verifying the Configuration of Basic IS-IS (IPv6) Functions

1. Check information about the interface with IS-IS (IPv6) enabled.

   [Huawei] **display isis interface** *interface-type interface-number* [**verbose**]

2. Check IS-IS (IPv6) neighbor information.

   [Huawei] **display isis** *process-id* **peer** [**verbose**]

3. Check IS-IS (IPv6) routing information.

   [Huawei] **display isis route** [*process-id* | **vpn-instance** *vpn-instance-name*] **ipv6** [**verbose** | [**level-1** | **level-2**] | *ipv6-address* [*prefix-length*]]

Example for Configuring IS-IS IPv4/IPv6 Dual Stack

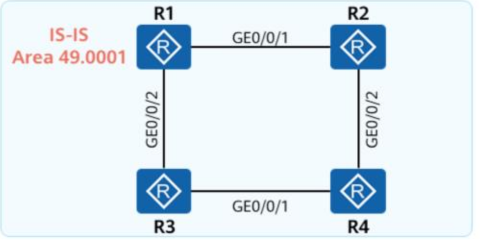

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	GE0/0/2	10.1.13.1/24	2001:DB8:2345:13::1/64
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	GE0/0/2	10.1.24.2/24	2001:DB8:2345:24::2/64
R3	GE0/0/1	10.1.34.3/24	2001:DB8:2345:34::3/64
	GE0/0/2	10.1.13.3/24	2001:DB8:2345:13::3/64
R4	GE0/0/1	10.1.34.4/24	2001:DB8:2345:34::4/64
	GE0/0/2	10.1.24.4/24	2001:DB8:2345:24::4/64

Scenario description:

- A company has IS-IS deployed over an IPv4 network to implement IPv4 network interconnection. To ensure future service development, the company has an IPv6 network deployed for service testing. In this case, the IPv6 function needs to be enabled in the IS-IS process on the network.
- As device performance varies, link cost modification is required to ensure that IPv4 and IPv6 services are routed through different paths, and IS-IS MT is required to ensure that a separate topology is calculated for the IPv6 network.
- All routers run IS-IS and are all Level-2 routers. The entire network is in area 49.0001. After the configuration is complete, each router should learn the IPv4 and IPv6 routes to all network segments in the AS.

HUAWEI

Deploying the IPv4 Network

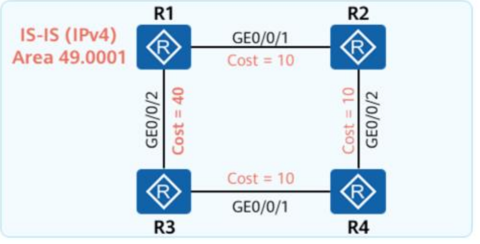

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	GE0/0/2	10.1.13.1/24	2001:DB8:2345:13::1/64
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	GE0/0/2	10.1.24.2/24	2001:DB8:2345:24::2/64
R3	GE0/0/1	10.1.34.3/24	2001:DB8:2345:34::3/64
	GE0/0/2	10.1.13.3/24	2001:DB8:2345:13::3/64
R4	GE0/0/1	10.1.34.4/24	2001:DB8:2345:34::4/64
	GE0/0/2	10.1.24.4/24	2001:DB8:2345:24::4/64

1. Assign an IPv4 address to each router interface. (The configuration details are not provided here.)
2. Configure basic IS-IS functions.

```
[R1] isis 1
[R1-isis-1] is-level level-2
[R1-isis-1] network-entity 49.0001.0000.0000.0001.00
#
[R1] interface gigabitethernet 0/0/1
[R1-GigabitEthernet0/0/1] isis enable 1
[R1-GigabitEthernet0/0/1] quit
[R1] interface gigabitethernet 0/0/2
[R1-GigabitEthernet0/0/2] isis enable 1
[R1-GigabitEthernet0/0/2] isis cost 40 level-2
```

The configurations of R2, R3, and R4 are similar to the configuration of R1, and are not provided here.

The network entity titles (NETs) are as follows:

R2: 49.0001.0000.0000.0002.00

R3: 49.0001.0000.0000.0003.00

R4: 49.0001.0000.0000.0004.00

HUAWEI

Deploying the IPv6 Network

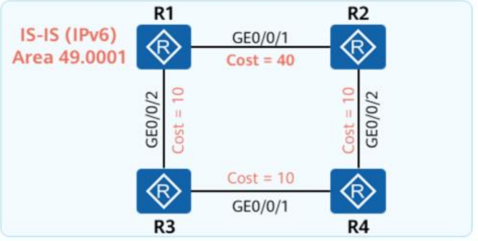

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	GE0/0/2	10.1.13.1/24	2001:DB8:2345:13::1/64
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	GE0/0/2	10.1.24.2/24	2001:DB8:2345:24::2/64
R3	GE0/0/1	10.1.34.3/24	2001:DB8:2345:34::3/64
	GE0/0/2	10.1.13.3/24	2001:DB8:2345:13::3/64
R4	GE0/0/1	10.1.34.4/24	2001:DB8:2345:34::4/64
	GE0/0/2	10.1.24.4/24	2001:DB8:2345:24::4/64

1. Assign an IPv6 address to each router interface. (The configuration details are not provided here.)

2. Enable IS-IS (IPv6).

```
[R1] isis 1
[R1-isis-1] ipv6 enable topology ipv6
#
[R1] interface gigabitethernet 0/0/1
[R1-GigabitEthernet0/0/1] isis ipv6 enable 1
[R1-GigabitEthernet0/0/1] isis ipv6 cost 40 level-2
[R1-GigabitEthernet0/0/1] quit
[R1] interface gigabitethernet 0/0/2
[R1-GigabitEthernet0/0/2] isis ipv6 enable 1
```

The configurations of R2, R3, and R4 are similar to the configuration of R1, and are not provided here.

HUAWEI

Check the Routing Information on the IS-IS (IPv4) Network

The following example uses R1 as the root to calculate the SPT.

The IS-IS (IPv4) logical topology is shown in the right figure.

```
[R1]dis ip routing-table protocol isis
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : ISIS
        Destinations : 2        Routes : 2

ISIS routing table status : <Active>
        Destinations : 2        Routes : 2

Destination/Mask   Proto   Pre  Cost     Flags NextHop      Interface

     10.1.24.0/24  ISIS-L2  15   20        D   10.1.12.2    GigabitEthernet0/0/1
     10.1.34.0/24  ISIS-L2  15   30        D   10.1.12.2    GigabitEthernet0/0/1

ISIS routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

IS-IS (IPv4) logical topology

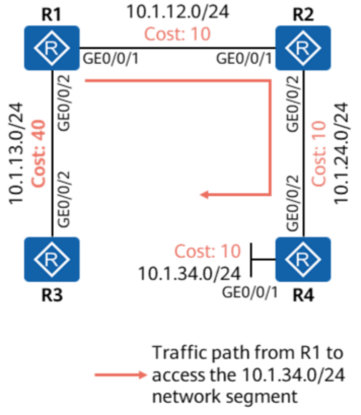

HUAWEI

Check the Routing Information on the IS-IS (IPv6) Network

The following example uses R1 as the root to calculate the SPT.

The IS-IS (IPv6) logical topology is shown in the right figure.

```
[R1]dis ipv6 routing-table protocol isis
Public Routing Table : ISIS
Summary Count : 2

ISIS Routing Table's Status : < Active >
Summary Count : 2

 Destination :          2001:DB8:2345:24::                  PrefixLength : 64
 NextHop :              FE80::2E0:FCFF:FE90:3D3A            Preference   : 15
 Cost :                 30                                  Protocol     : ISIS-L2
 RelayNextHop :         ::                                  TunnelID     : 0x0
 Interface :   GigabitEthernet0/0/2                 Flags         : D

 Destination :          2001:DB8:2345:34::                  PrefixLength : 64
 NextHop :   FE80::2E0:FCFF:FE90:3D3A               Preference    : 15
 Cost :                 20                          Protocol      : ISIS-L2
 RelayNextHop :         ::                                  TunnelID     : 0x0
 Interface :   GigabitEthernet0/0/2                 Flags         : D

ISIS Routing Table's Status : < Inactive >
Summary Count : 0
```

IS-IS (IPv6) logical topology

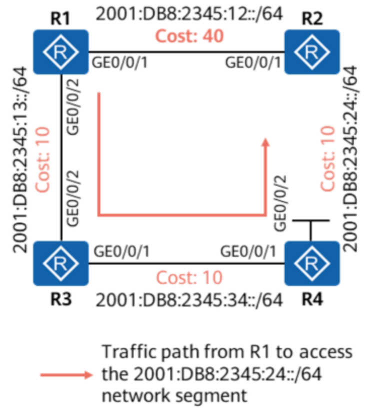

HUAWEI

Contents

BGP4+ Overview

- Traditional BGP-4 can manage only IPv4 unicast routing information. Multiprotocol Extensions for BGP (MP-BGP) can carry routing information for multiple network layer protocols. Currently, MP-BGP uses extended attributes and address families to support IPv6, multicast, and VPN, without changing the messaging and routing mechanisms of BGP-4.
- MP-BGP for IPv6 unicast networks is called BGP4+. BGP4+ establishes an independent topology for an IPv6 unicast network and stores routing information in an independent routing table. This ensures that routing information on an IPv4 unicast network is isolated from that on an IPv6 unicast network.

HUAWEI

- To support IPv6, BGP needs to map IPv6 routing information to the NLRI attributes.

Address Families That MP-BGP Supports

- MP-BGP uses address families to differentiate network layer protocols. To enable BGP peers to exchange different types of routing information, activating the peers and advertising BGP routes need to be performed in the correct address families.
- MP-BGP supports the following address families:
  - IPv4 unicast address family
  - IPv4 multicast address family
  - IPv6 unicast address family
  - VPNv4 address family
  - VPNv6 address family
  - ...

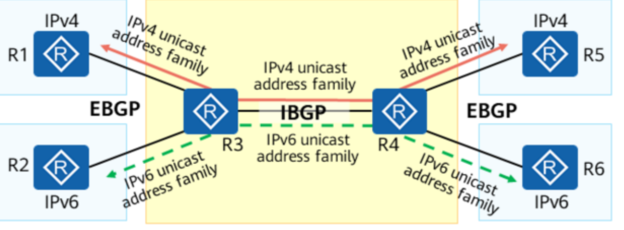

Peers on the IPv4 network must be activated in the IPv4 unicast address family, whereas peers on the IPv6 network must be activated in the IPv6 unicast address family.

HUAWEI

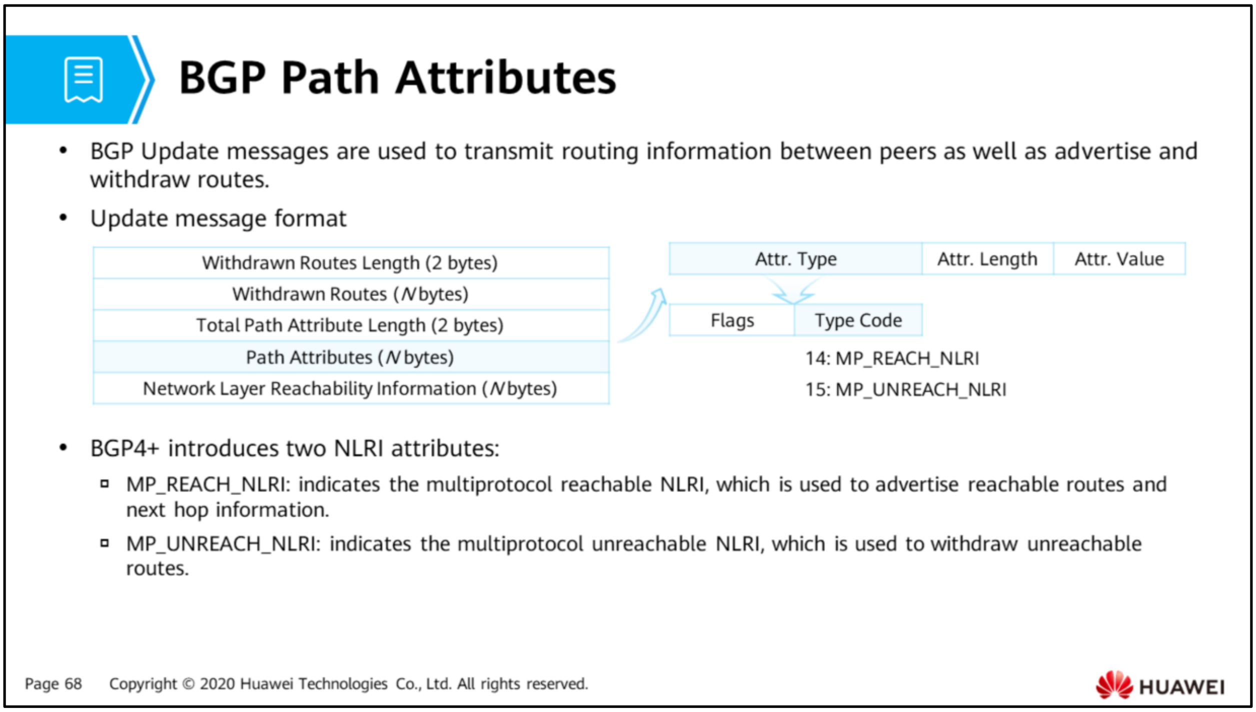

- Update message:
    - An Update message can be used to advertise multiple routes with the same path attribute. These routes are stored in the NLRI attribute. An Update message can also carry multiple unreachable routes, which are stored in the Withdrawn Routes field, to instruct peers to withdraw these routes.

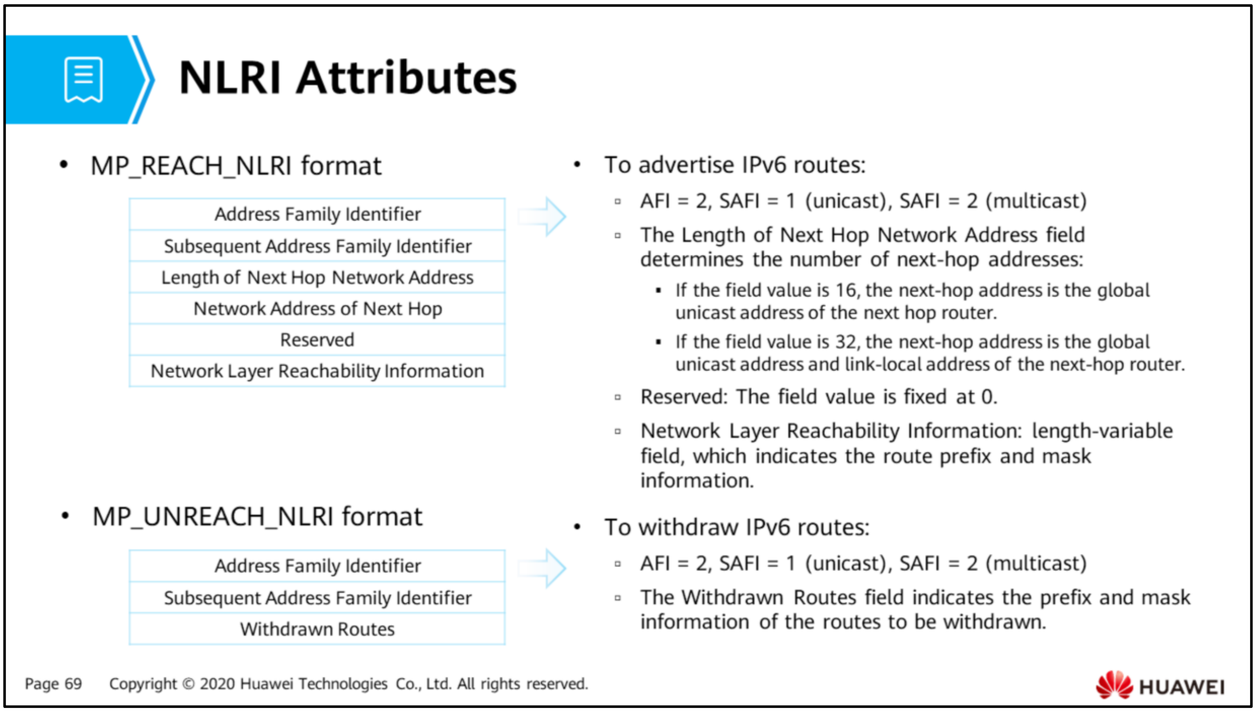

- Fields in the MP_REACH_NLRI attribute are described as follows:
  - Address Family Information: consists of a 2-byte address family identifier (AFI) and a 1-byte subsequent address family identifier (SAFI).
  - Length of Next Hop Network Address: indicates the length of the next-hop address and occupies 1 byte. Generally, the value is 16.
  - Network Address of Next Hop: The length is variable and depends on the preceding field. Generally, the value is a global unicast address.
  - Reserved: 1 byte. The value must be 0.
  - Network Layer Reachability Information: contains information about the routes with the same attributes. The value 0 indicates the default route.
- Fields in the MP_UNREACH_NLRI field are described as follows:
  - Withdrawn Routes: indicates the routes to be withdrawn. The value is in the <mask length, route prefix> format. If the mask length is 0, the associated route is a default route.

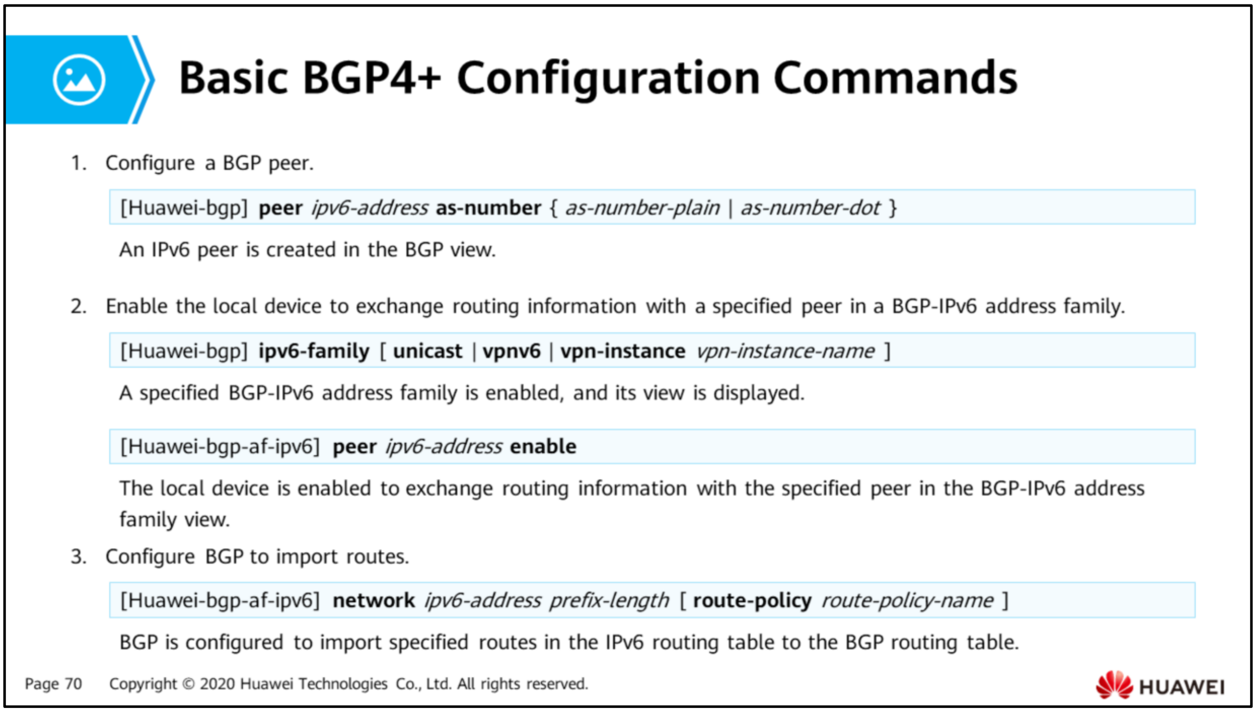

- The basic configuration commands and methods of BGP4+ are the same as those of BGP. For details, see the "HCIP-Datacom-Core Technology" course.
- [Huawei-bgp] **ipv6-family** [**unicast** | **vpnv6** | **vpn-instance** *vpn-instance-name*]
  - **unicast**: enters the IPv6 unicast address family view.
  - **vpnv6**: enters the BGP-VPNv6 address family view.
  - **vpn-instance** *vpn-instance-name*: associates a specified VPN instance with the IPv6 address family and enters the BGP-VPN instance IPv6 address family view. The value is a string of 1 to 31 case-sensitive characters. If spaces are used, the string must start and end with double quotation marks (").
- [Huawei-bgp-af-ipv6] **network** *ipv6-address prefix-length* [**route-policy** *route-policy-name*]
  - *ipv6-address*: specifies the IPv6 address advertised by BGP. The value is a 32-digit hexadecimal number, in the format of *X:X:X:X:X:X:X:X*.
  - *prefix-length*: specifies the prefix length of the IPv6 address advertised by BGP. The value is an integer ranging from 0 to 128.
  - **route-policy** *route-policy-name*: specifies the name of the route-policy applied to route advertisement. The value is a string of 1 to 40 case-sensitive characters. If spaces are used, the string must start and end with double quotation marks (").

Verifying the Configuration of Basic BGP4+ Functions

1. Check BGP4+ peer information.

```
[Huawei] display bgp ipv6 peer ipv6-address [ verbose ]
```

2. Check BGP4+ routing information.

```
[Huawei] display bgp ipv6 routing-table
```


HUAWEI

Example for Configuring BGP IPv4/IPv6 Dual Stack

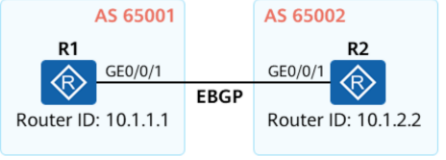

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	Loopback0	10.1.1.1/32	2001:DB8:2345:1::1/128
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	Loopback0	10.1.2.2/32	2001:DB8:2345:2::2/128

Scenario description:

- Two branches of a company communicate with each other through BGP. To ensure future service development, an IPv6 network is deployed in each branch for service testing. In this case, BGP4+ needs to be deployed to implement IPv6 network interworking.

Deploying the IPv4 Network

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	Loopback0	10.1.1.1/32	2001:DB8:2345:1::1/128
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	Loopback0	10.1.2.2/32	2001:DB8:2345:2::2/128

1. Assign an IPv4 address to each router interface. (The configuration details are not provided here.)

2. Configure basic BGP functions: establish an EBGP peer relationship.

```
[R1] bgp 65001
[R1-bgp] router-id 10.1.1.1
[R1-bgp] peer 10.1.12.2 as-number 65002
```

```
[R2] bgp 65002
[R2-bgp] router-id 10.1.2.2
[R2-bgp] peer 10.1.12.1 as-number 65001
```

3. Configure BGP to advertise IPv4 routes.

```
[R1-bgp] iPv4-family unicast
[R1-bgp-af-ipv4] network 10.1.1.1 32
```

```
[R2-bgp] iPv4-family unicast
[R2-bgp-af-ipv4] network 10.1.2.2 32
```

Deploying the IPv6 Network

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	Loopback0	10.1.1.1/32	2001:DB8:2345:1::1/128
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	Loopback0	10.1.2.2/32	2001:DB8:2345:2::2/128

1. Assign an IPv6 address to each router interface. (The configuration details are not provided here.)

2. Configure basic BGP4+ functions: establish an EBGP peer relationship and enable the local device to exchange routing information with the EBGP peer.

```
[R1] bgp 65001
[R1-bgp] peer 2001:DB8:2345:12::2 as-number 65002
[R1-bgp] iPv6-family unicast
[R1-bgp-af-ipv6] peer 2001:DB8:2345:12::2 enable
```

```
[R2] bgp 65002
[R2-bgp] peer 2001:DB8:2345:12::1 as-number 65001
[R2-bgp] iPv6-family unicast
[R2-bgp-af-ipv6] peer 2001:DB8:2345:12::1 enable
```

3. Configure BGP4+ to advertise IPv6 routes.

```
[R1-bgp-af-ipv6] network 2001:DB8:2345:1::1 128
```

```
[R2-bgp-af-ipv6] network 2001:DB8:2345:2::2 128
```


HUAWEI

Checking the BGP Peer Information

- Run the **display bgp [ipv6] peer** command to check information about the peers on the IPv4 and IPv6 networks.
- By comparing the command outputs, you will find that the peer information is basically the same except for the addresses used to establish the peer relationships.

BGP

```
[R1]display bgp peer

 BGP local router ID :  10.1.1.1
 Local AS number:       65001
 Total number of peers:  1                  Peers in established state : 1

  Peer            V     AS     MsgRcvd  MsgSent  OutQ    Up/Down   State        PrefRcv
  10.1.12.2       4     65002  11       12       0       00:06:49  Established    1
```

BGP4+

```
[R1]display bgp ipv6 peer

 BGP local router ID : 10.1.1.1
 Local AS number : 65001
 Total number of peers : 1                  Peers in established state : 1

  Peer               V     AS     MsgRcvd  MsgSent  OutQ    Up/Down   State        PrefRcv
  2001:DB8:2345:12::2  4   65002  14       15       0       00:07:35  Established    1
```

Checking the BGP4+ Routing Information

Device	Interface	IPv4 Address	IPv6 Address
R1	GE0/0/1	10.1.12.1/24	2001:DB8:2345:12::1/64
	Loopback0	10.1.1.1/32	2001:DB8:2345:1::1/128
R2	GE0/0/1	10.1.12.2/24	2001:DB8:2345:12::2/64
	Loopback0	10.1.2.2/32	2001:DB8:2345:2::2/128

```
[R1]display bgp ipv6 routing-table

 BGP Local router ID is 10.1.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 2
*> Network :      2001:DB8:2345:1::1        PrefixLen :   128
   NextHop :      ::                        LocPrf  :
   MED :          0                         PrefVal :     0
   Label  :
   Path/Ogn :     i
*> Network :      2001:DB8:2345:2::2        PrefixLen :   128
   NextHop :      2001:DB8:2345:12::2       LocPrf :
   MED :          0                         PrefVal :     0
   Label :
   Path/Ogn :     65002  i
```

Check the BGP4+ routing table. The command output shows that a total of two valid routes exist.

Checking the NLRI Attributes of BGP4+

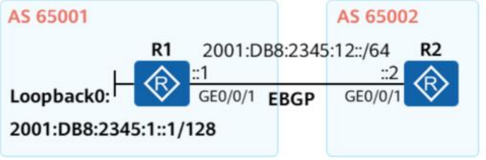

By analyzing the Update message sent by R1, you can check information about the MP_REACH_NLRI attribute. In the attribute:

- Next-hop address of the reachable route: 2001:db8:2345:12::1
- Prefix and prefix length of the reachable route: 2001:db8:2345:1::1/128

UPDATE Message-Border Gateway Protocol
- UPDATE Message
  - Marker: 16 bytes
  - Length: 85 bytes
  - Type: UPDATE Message (2)
  - Unfeasible routes length: 0 bytes
  - Total path attribute length: 62 bytes
  - Path attributes
    - ORIGIN: IGP (4 bytes)
    - AS_PATH: 65001 (9 bytes)
    - MULTI_EXIT_DISC: 0 (7 bytes)
    - **MP_REACH_NLRI** (42 bytes)
      - Flags: 0x90 (Optional, Non-transitive, Complete, Extended Length)
      - Type code: MP_REACH_NLRI (14)
      - Length: 38 bytes
      - Address family: IPv6 (2)
      - Subsequent address family identifier: Unicast (1)
      - Next hop network address (16 bytes)
        - Next hop: 2001:db8:2345:12::1 (16)
      - Subnetwork points of attachment: 0
      - Network layer reachability information (17 bytes)
        - 2001:db8:2345:1::1/128
          - MP Reach NLRI prefix length: 128
          - MP Reach NLRI prefix: 2001:db8:2345:1::1

HUAWEI

Quiz

1. (Multiple) Which of the following types of LSAs belong to OSPFv3? ( )
   A. Router-LSA
   B. Network-LSA
   C. Network-summary-LSA
   D. Link-LSA
2. (TorF) IS-IS can calculate the SPTs for IPv4 and IPv6 networks separately. ( )
   A. True
   B. False

HUAWEI

1. ABD
2. A

Summary

- IPv6 static routes can meet network interconnection requirements. However, as the network scale expands, dynamic routing protocols need to be leveraged. Dynamic routing protocols are required to support IPv6 and carry IPv6 addresses for route advertisement.
- To support IPv6 networks, the IETF enhanced and improved OSPFv2 and finally developed the new protocol, OSPFv3.
- IS-IS features high scalability and supports IPv6 through new TLVs.
- Similarly, BGP supports IPv6 by using its own extended path attributes and multi-protocol address families.


Thank You
www.huawei.com