

Chapter 9: OSPFv3

Instructor Materials

CCNP Enterprise: Advanced Routing



Chapter 9 Content

This chapter covers the following content:

- **OSPFv3 Fundamentals** - This section provides an overview of the OSPFv3 routing protocol, its similarities to OSPFv2, and its configuration.
- **OSPFv3 Configuration** - This section explains and demonstrates how OSPFv3 is used for exchanging IPv6 routes.
- **OSPFv3 LSA Flooding Scope** - This section provides a deeper view of the OSPFv3 link-state advertisement (LSA) structure and the comparison to OSPFv2.

OSPFv3 Fundamentals

- Open Shortest Path First version 3 (OSPFv3) is the latest version of the OSPF protocol.
- The OSPFv3 protocol is not backward compatible with OSPFv2, but the protocol mechanisms are essentially the same.

OSPFv2 and OSPFv3 Differences

The primary differences between OSPFv2 and OSPFv3 protocols are as follows:

- **Support for multiple address families** - OSPFv3 supports IPv4 and IPv6 address families.
- **New LSA types** - New LSA types have been created to carry IPv6 prefixes.
- **Removal of addressing semantics** - The IP prefix information is no longer present in the OSPF packet headers.
- **LSA flooding** - OSPFv3 includes a new link-state type field that is used to determine the flooding scope of LSA, as well as the handling of unknown LSA types.
- **Packet format** - OSPFv3 runs directly over IPv6, and the number of fields in the packet header has been reduced.
- **Router ID** - The router ID is used to identify neighbors, regardless of the network type in OSPFv3.
- **Authentication** - Neighbor authentication has been removed from the OSPF protocol and is now performed through IPsec extension headers in the IPv6 packet.
- **Neighbor adjacencies** - OSPFv3 inter-router communication is handled by IPv6 link-local addressing.
- **Multiple instances** - OSPFv3 packets include an instance ID field that may be used to manipulate which routers on a network segment are allowed to form adjacencies.

OSPFv3 Link-State Advertisement

- The OSPF link-state database (LSDB) information is organized and advertised differently in version 3 than in version 2. OSPFv3 modifies the structure of the router LSA (Type 1), renames the network summary LSA to the inter-area prefix LSA, and renames the autonomous system boundary router (ASBR) summary LSA to inter-area router LSA.
- IP address information is advertised independently by two new LSA types:
 - Intra-area prefix LSA
 - Link-local LSA

Advertising the IP address information using new LSA types eliminates the need for OSPF to perform full shortest path first (SPF) tree calculations every time a new address prefix is added or changed on an interface. The OSPFv3 LSDB creates a shortest path topology tree based on links instead of networks.

OSPFv3 Fundamentals

OSPFv3 LSA Types

- Table 9-2 provides a brief description of each OSPFv3 LSA type.

Table 9-2 OSPFv3 LSA Type

LSA Type	Name	Description
0x2001	Device	Every router generates router LSAs that describe the state and cost of the router's interfaces to the area.
0x2002	Network	A designated router generates network LSAs to announce all the routers attached to the link, including itself.
0x2003	Inter-area prefix	Area border routers generate inter-area prefix LSAs to describe routes to IPv6 address prefixes that belong to other areas.
0x2004	Inter-area router	Area border routers generate inter-area router LSAs to announce the addresses of autonomous system boundary routers in other areas.
0x2005	AS-external	Autonomous system boundary routers advertise AS-external LSAs to announce default routes or routes learned via redistribution from other protocols.
0x2007	NSSA	Autonomous system boundary routers that are located in a not-so-stubby area (NSSA) advertise NSSA LSAs for routes redistributed into the area.
0x0008	Link	The link LSA maps all the global unicast address prefixes associated with an interface to the link-local interface IP address of the router. The link LSA is shared only between neighbors on the same link.
0x2009	Intra-area prefix	The intra-area-prefix LSA is used to advertise one or more IPv6 prefixes that are associated with a router, stub, or transit network segment.

OSPFv3 Communication

OSPFv3 packets use protocol ID 89, and routers communicate with each other using the local interface's IPv6 link-local address as the source. Depending on the packet type, the destination address is either a unicast link-local address or the multicast link-local scoped address:

- FF02::05: OSPFv3 AllSPFRouters
- FF02::06: OSPFv3 AllDRouters designated router (DR) router

Table 9-3 OSPFv3 Packet Types

Type	Packet Name	Source	Destination	Purpose
1	Hello	Link-local address	FF02::5 (all routers)	Discover and maintain neighbors
		Link-local address	Link-local address	Initial adjacency forming, immediate hello
2	Database description	Link-local address	Link-local address	Summarize database contents
3	Link-state request	Link-local address	Link-local address	Database information request
4	Link-state update	Link-local address	Link-local address	Initial adjacency forming, in response to link-state request
		Link-local address (from DR)	FF02::5 (all routers)	Database update
		Link-local address (from non-DR)	FF02::6 (DR/BDR)	Database update
5	Link-state acknowledgment	Link-local address	Link-local address	Initial adjacency forming, in response to link-state update
		Link-local address (from DR)	FF02::5 (all routers)	Flooding acknowledgment
		Link-local address (from non-DR)	FF02::6 (DR/BDR)	Flooding acknowledgment

OSPFv3 Configuration

The following section explains the process for configuring and verifying OSPFv3.

OSPFv3 Configuration

The process for configuring OSPFv3 involves the following steps:

Step 1. Initialize the routing process by enabling ipv6 unicast-routing on the router and then configuring OSPFv3 with the command `router ospfv3 [process-id]`.

Step 2. Define the router ID (RID) by using the command `router-id`. The router ID is a 32-bit value that does not need to match an IPv4 address. It may be any number, as long as the value is unique within the OSPF domain. OSPFv3 uses the same algorithm as OSPFv2 for dynamically locating the RID. If there are not any IPv4 interfaces available, the RID is set to 0.0.0.0 and does not allow adjacencies to form.

Step 3. Initialize the address family within the routing process by using the optional command `address-family {ipv6 | ipv4} unicast`. The appropriate address family is enabled automatically when OSPFv3 is enabled on an interface.

Step 4. Use the interface command `ospfv3 process-id ipv6 area area-id` to enable the protocol and assign the interface to an area.

OSPFv3 Configuration

OSPFv3 Topology

Figure 9-1 shows a simple four-router topology to demonstrate OSPFv3 configuration. Area 0 consists of R1, R2, and R3, and Area 34 contains R3 and R4. R3 is the area border router (ABR).




Figure 9-1 OSPFv3 Topology

Example 9-1 provides the OSPFv3 and IPv6 address configurations for R1.

Example 9-1 IPv6 Addressing and OSPFv3 Configuration

```
R1
interface Loopback0
  ipv6 address 2001:DB8::1/128
  ospfv3 1 ipv6 area 0
!
interface GigabitEthernet0/1
  ipv6 address FE80::1 link-local
  ipv6 address 2001:DB8:0:1::1/64
  ospfv3 1 ipv6 area 0
!
interface GigabitEthernet0/2
  ipv6 address FE80::1 link-local
  ipv6 address 2001:DB8:0:12::1/64
  ospfv3 1 ipv6 area 0
!
router ospfv3 1
  router-id 192.168.1.1
```

OSPFv3 Configuration

OSPFv3 Verification

The commands for viewing OSPFv3 settings and statuses are very similar to those used in OSPFv2. In essence, you replace `ip ospf` with `ospfv3`. For example, to view neighbor adjacencies use the command **show ip ospfv3 neighbor**. To view an OSPFv3-enabled interface status use the command **show ospfv3 interface [interface-id]**

Example 9-2 Identifying R3's OSPFv3 Neighbors

```
R3# show ospfv3 ipv6 neighbor
```

```
OSPFv3 1 address-family ipv6 (router-id 192.168.3.3)
```

Neighbor ID	Pri	State	Dead Time	Interface ID	Interface
192.168.2.2	1	FULL/DR	00:00:32	5	GigabitEthernet0/2
192.168.4.4	1	FULL/BDR	00:00:33	5	GigabitEthernet0/4

Example 9-3 Viewing the OSPFv3 Interface Configuration

```
R1# show ospfv3 interface GigabitEthernet0/2
GigabitEthernet0/2 is up, line protocol is up
  Link Local Address FE80::1, Interface ID 3
  Area 0, Process ID 1, Instance ID 0, Router ID 192.168.1.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 192.168.1.1, local address FE80::1
  Backup Designated router (ID) 192.168.2.2, local address FE80::2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Graceful restart helper support enabled
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 192.168.2.2 (Backup Designated Router)
  Suppress hello for 0 neighbor(s)
```

OSPFv3 Configuration

OSPFv3 Verification (Cont.)

The command **show ospfv3 interface brief** shows a brief version of the OSPFv3 interface settings. To view the OSPFv3 IPv6 routing table use the command **show ipv6 route ospf**.

Example 9-4 Viewing the Brief Iteration of OSPFv3 Interfaces

```
R3# show ospfv3 interface brief
```

Interface	PID	Area	AF	Cost	State	Nbrs	F/C
Lo0	1	0	ipv6	1	LOOP	0/0	
Gi0/2	1	0	ipv6	1	BDR	1/1	
Gi0/4	1	34	ipv6	1	DR	1/1	

Example 9-5 Viewing the OSPFv3 Routes in the IPv6 Routing Table

```
R1# show ipv6 route ospf
! Output omitted for brevity

IPv6 Routing Table - default - 11 entries
    RL - RPL, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
    OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
..
O   2001:DB8::2/128 [110/1]
    via FE80::2, GigabitEthernet0/2
O   2001:DB8::3/128 [110/2]
    via FE80::2, GigabitEthernet0/2
OI  2001:DB8::4/128 [110/3]
    via FE80::2, GigabitEthernet0/2
OI  2001:DB8:0:4::/64 [110/4]
    via FE80::2, GigabitEthernet0/2
O   2001:DB8:0:23::/64 [110/2]
    via FE80::2, GigabitEthernet0/2
OI  2001:DB8:0:34::/64 [110/3]
    via FE80::2, GigabitEthernet0/2
```

OSPFv3 Configuration

Passive Interface

An interface is marked as being passive with the command **passive-interface** *interface-id* or globally with the **passive-interface default**; an interface is marked as active with the command **no passive-interface** *interface-id*.

- The command is placed under the OSPFv3 process or under the specific address family.
- Placing the command under the global process cascades the setting to both address families.
- Example 9-6 demonstrates making the LAN interface on R1 explicitly passive and making all interfaces passive on R4 while marking the Gi0/3 interface as active.

The active/passive state of an interface is verified by examining the OSPFv3 interface status using the command **show ospfv3 interface** [*interface-id*] and searching for the passive keyword.

Example 9-6 Configuring OSPFv3 Passive Interfaces

```
R1(config)# router ospfv3 1
R1(config-router)# passive-interface GigabitEthernet0/1
R4(config)# router ospfv3 1
R4(config-router)# passive-interface default
22:10:46.838: %OSPFv3-5-ADJCHG: Process 1, IPv6, Nbr 192.168.3.3 on
GigabitEthernet0/3 from FULL to DOWN, Neighbor Down: Interface down or detached
R4(config-router)# no passive-interface GigabitEthernet 0/3
```

```
R1# show ospfv3 interface gigabitEthernet 0/1 | include Passive
No Hellos (Passive interface)
```

OSPFv3 Configuration

IPv6 Route Summarization

Summarization of internal OSPFv3 routes follows the same rules as for OSPFv2 and must occur on ABRs by using the command **area area-id range prefix/prefix-length**.

Example 9-8 shows R3's configuration for summarizing these prefixes.

Example 9-9 shows R4's IPv6 routing table after R3 is configured to summarize the Area 0 loopback interfaces. The summary route is highlighted.

Example 9-8 IPv6 Summarization

```
R3# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)# router ospfv3 1
R3(config-router)# address-family ipv6 unicast
R3(config-router-af)# area 0 range 2001:db8:0:0::/65
```

Example 9-9 R4's IPv6 Routing Table After Summarization

```
R4# show ipv6 route ospf | begin Application
1A - LISP away, a - Application
OI 2001:DB8::/65 [110/4]
    via FE80::3, GigabitEthernet0/3
OI 2001:DB8:0:1::/64 [110/4]
    via FE80::3, GigabitEthernet0/3
OI 2001:DB8:0:12::/64 [110/3]
    via FE80::3, GigabitEthernet0/3
OI 2001:DB8:0:23::/64 [110/2]
    via FE80::3, GigabitEthernet0/3
```

OSPFv3 Configuration

Network Type

OSPFv3 supports the same OSPF network types as OSPFv2. Example 9-10 shows how to view the OSPFv3 network type.

Example 9-11 demonstrates changing the OSPFv3 network type by using the interface parameter command **ospfv3 network {point-to-point | point-to-multipoint broadcast | nonbroadcast}**.

Example 9-12 shows how to verify the new settings. The network is now a point-to-point link, and the interface state is indicated as P2P as confirmation.

Example 9-10 Viewing the Dynamic Configured OSPFv3 Network Type

```
R2# show ospfv3 interface GigabitEthernet 0/3 | include Network
      Network Type BROADCAST, Cost: 1
R2# show ospfv3 interface brief
```

Interface	PID	Area	AF	Cost	State	Nbrs	F/C
Lo0	1	0	ipv6	1	LOOP	0/0	
Gi0/3	1	0	ipv6	1	DR	1/1	
Gi0/1	1	0	ipv6	1	BDR	1/1	

Example 9-11 Changing the OSPFv3 Network Type

```
R2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# interface GigabitEthernet 0/3
R2(config-if)# ospfv3 network point-to-point
R3(config)# interface GigabitEthernet 0/2
R3(config-if)# ospfv3 network point-to-point
```

Example 9-12 Viewing the Statically Configured OSPFv3 Network Type

```
R2# show ospfv3 interface GigabitEthernet 0/3 | include Network
      Network Type POINT_TO_POINT, Cost: 1
R2# show ospfv3 interface brief
```

Interface	PID	Area	AF	Cost	State	Nbrs	F/C
Lo0	1	0	ipv6	1	LOOP	0/0	
Gi0/3	1	0	ipv6	1	P2P	1/1	
Gi0/1	1	0	ipv6	1	BDR	1/1	

OSPFv3 Configuration

OSPFv3 Authentication

OSPFv3 does not support neighbor authentication within the protocol itself. Instead, the routing protocol utilizes IP Security (IPsec) to provide authentication. IPv6 Authentication Header (AH) or Encapsulating Security Payload (ESP) extension headers may be added to the OSPF packets to provide authentication, integrity, and confidentiality:

- Authentication Header (AH): Provides authentication
- Encapsulating Security Payload (ESP): Provides authentication and encryption




Figure 9-2 *IPv6 IPsec Packet Format*

OSPFv3 Authentication (Cont.)

OSPFv3 authentication supports IPsec AH authentication using the command **ospfv3 authentication** or ESP authentication and encryption with the command **ospfv3 encryption**.

OSPFv3 neighbor authentication does not perform Internet Key Exchange (IKE) to negotiate the IPsec security association (SA) values. Therefore, the IPsec Security Parameter Index (SPI) hash algorithm and key must be manually defined when configuring OSPFv3 authentication.

IPsec peers cannot reuse the same SPI values.

The command **show crypto ipsec sa | include spi** may be used to determine the active IPsec sessions and currently used SPI values. The full interface command **ospfv3 encryption {ipsec spi spi esp encryption-algorithm {key-encryption-type key} authentication-algorithm {key-encryption-type key} | null}** encrypts and authenticates the OSPFv3 packet in IOS using ESP.

OSPFv3 Interface Authentication and Encryption

Example 9-13 demonstrates how to configure encryption and authentication for OSPFv3 packets using ESP. The following fabricated values are included in the configuration to establish the IPsec session:

- Security policy index: = 500
- Encryption algorithm: = 3des
- Encryption key: = 012345678901234567890123456789012345678901234567
- Authentication algorithm: = sha1
- Authentication key: = 0123456789012345678901234567890123456789

Example 9-13 *OSPFv3 Interface Authentication and Encryption*

```
interface GigabitEthernet0/1
ospfv3 encryption ipsec spi 500 esp 3des 01234567890123456789012345678901234567
      8901234567 sha1 0123456789012345678901234567890123456789
! The ospfv3 encryption rolls over to two lines in the example, but it is only
! one single CLI command.
```

OSPFv3 Area Authentication and Encryption (Cont.)

Example 9-14 demonstrates how to configure area authentication and encryption using the same IPsec settings.

Example 9-15 displays the output of the command `show ospfv3 interface [interface id]`. This show command can be used to verify that authentication and encryption are enabled on the interface and that a secure connection has formed with the neighbor.

Example 9-14 OSPFv3 Area Authentication and Encryption

```
router ospfv3 100
 area 0 encryption ipsec spi 500 esp 3des 123456789012345678901234567
      8901234567 sha1 0123456789012345678901234567890123456789
! The ospfv3 encryption rolls over to two lines in the example, but it is
! entered as one long command. The running configuration will display the
! password encrypted
```

Example 9-15 OSPFv3 IPsec Verification

```
R2# show ospfv3 interface
GigabitEthernet0/1 is up, line protocol is up
  Link Local Address FE80::2, Interface ID 3
  Area 0, Process ID 100, Instance ID 0, Router ID 100.0.0.2
  Network Type BROADCAST, Cost: 1
  3DES encryption SHA-1 auth SPI 500, secure socket UP (errors: 0)
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 100.0.0.2, local address FE80::2
  Backup Designated router (ID) 100.0.0.1, local address FE80::1
! Output omitted for brevity
```

OSPFv3 Link-Local Forwarding

Significant changes have occurred in how OSPFv3 builds the area topology. The OSPFv3 LSDB creates a shortest path topology tree based on links instead of networks. This means that transit links only require IPv6 link-local addresses for forwarding traffic.

Example 9-16 demonstrates the removal of the global IPv6 unicast addresses from the transit links on R1, R2, and R3.

Example 9-16 *Removal of Global IPv6 Addresses*

```
R1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# interface g10/2
R1(config-if)# no ipv6 address 2001:DB8:0:12::1/64
```

```
R2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# interface g10/1
R2(config-if)# no ipv6 address 2001:DB8:0:12::2/64
R2(config-if)# interface G10/3
R2(config-if)# no ipv6 address 2001:DB8:0:23::2/64
```

```
R3# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)# interface gigabitEthernet 0/2
R3(config-if)# no ipv6 address 2001:DB8:0:23::3/64
R3(config-if)# interface GigabitEthernet 0/4
R3(config-if)# no ipv6 address 2001:DB8:0:34::3/64
```

OSPFv3 Link-Local Forwarding (Cont.)

Example 9-17 shows the OSPFv3 learned routes from R4's perspective. Notice that the transit networks no longer appear.

R4 still maintains full connectivity to those networks in Example 9-17 because the topology is built using the IPv6 link-local address. As long as the source and destination devices have routes to each other, communication can still exist.

Example 9-18 demonstrates that R4 still maintains connectivity to R1's LAN interface.

Example 9-17 R4's Routing Table after Removal of Global IPv6 Addresses

```
R4# show ipv6 route ospf
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
        H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
        IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
        ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
        RL - RPL, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
        OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
        la - LISP alt, lr - LISP site-registrations, ld - LISP dyn-eid
        lA - LISP away, a - Application

OI 2001:DB8::1/128 [110/3]
    via FE80::3, GigabitEthernet0/3
OI 2001:DB8::2/128 [110/2]
    via FE80::3, GigabitEthernet0/3
OI 2001:DB8::3/128 [110/1]
    via FE80::3, GigabitEthernet0/3
OI 2001:DB8:0:1::/64 [110/4]
    via FE80::3, GigabitEthernet0/3
```

Example 9-18 Connectivity Test with Link-Local Forwarding

```
R4# ping 2001:DB8:0:1::1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:DB8:0:1::1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/5/6 ms
```

OSPFv3 LSA Flooding Scope

OSPFv3 allows for three flooding scopes: link-local scope, area scope and autonomous system scope.

OSPFv3 LSA Flooding Scope

The OSPFv3 flooding scopes:

- **Link-local scope** - Limited to the local link
- **Area scope** - Contains LSA flooding to the local area
- **Autonomous system scope** - Floods LSAs throughout the entire OSPF routing domain

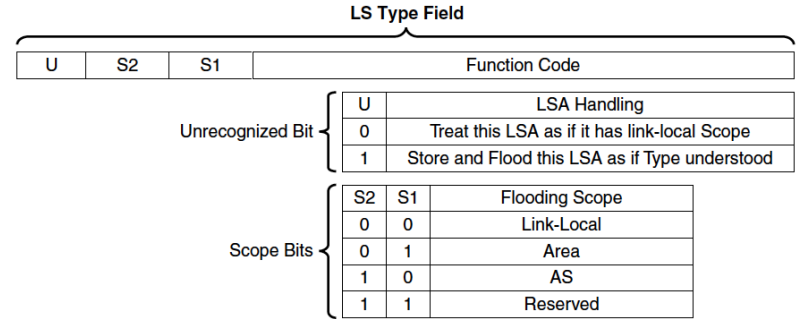


Figure 9-3 LS Type Field

- The LS type field in OSPFv3 has been modified from 8 bits to 16 bits.
- Figure 9-3 shows the new LS Type field format.
- The 3 high-order bits of the new LS Type field allow for the encoding of flood information.
- The first bit, U (unrecognized), indicates how a router should handle an LSA if it is unrecognized.
- The second and third bits, both S (scope) bits, indicate how the LSA should be flooded.
- The remaining bits of the link-state field indicate the function code of the LSA.

OSPFv3 LSA Flooding Scope

OSPFv3 Database

The router LSA describes the router's interface state and cost.

Example 9-19 shows the output of the command **show ospfv3 database router[self-originate | adv-router *RID*]**.

The optional **self-originate** keyword filters the LSAs to those created by the router on which the command is executed. The **adv-router *RID*** keyword allows for selection of the LSAs for a specific router's LSAs that exist in the local router's LSDB.

Example 9-19 *Viewing the Self-Originating LSAs in the OSPFv3 Database*

```
R1# show ospfv3 database router self-originate

      OSPFv3 1 address-family ipv6 (router-id 192.168.1.1)

      Router Link States (Area 0)

LS age: 563
Options: (V6-Bit, E-Bit, R-Bit, DC-Bit)
LS Type: Router Links
Link State ID: 0
Advertising Router: 192.168.1.1
LS Seq Number: 80000012
Checksum: 0x13FB
Length: 40
Number of Links: 1

Link connected to: a Transit Network
Link Metric: 1
Local Interface ID: 4
Neighbor (DR) Interface ID: 4
Neighbor (DR) Router ID: 192.168.1.1
```


OSPFv3 LSA Flooding Scope

OSPFv3 Options

OSPFv3 LSAs include an options bit field that describes the router's capabilities. Table 9-5 describes the various service options.

Table 9-5 OSPFv3 Options Field Bits

Option	Description
V6	The V6 bit indicates that the router participates in IPv6 routing.
E	The E bit indicates that the router is capable of processing external LSAs. A router in a stubby area sets the E bit to clear (0). Neighboring routers do not form adjacencies if they have mismatched E bit settings.
R	The R bit indicates that the router actively participates in forwarding traffic. The R bit set to clear (0) indicates that the router is not to be used as a transit router for forwarding traffic but is still capable of exchanging route information.
DC	The DC bit is set to indicate that the router is capable of suppressing future hellos from being sent over the interface. The interface must be configured as a demand circuit for hello suppression to occur. (Demand circuits are typically used on costly low-bandwidth legacy ISDN BRI circuits, which are beyond the scope of this book.)
MC	The MC bit indicates that the router is capable of multicast extensions for OSPF (MOSPF). This bit is not used and is listed only for reference. In 2008, RFC 5340 deprecated MOSPF along with the Group-Membership-LSA.
N	The N bit indicates that the router supports Type 7 LSAs (NSSA area). Neighboring routers do not form an adjacency if they have mismatched N bit settings.

Example 9-20 shows a portion of R3's router LSA's LSDB. The highlighted bits indicate the functionality the router can perform in each area.

Example 9-20 Viewing R3's LSAs in the OSPFv3 Database

```
R1# show ospfv3 database router adv-router 192.168.3.3

OSPFv3 1 address-family ipv6 (router-id 192.168.1.1)

Router Link States (Area 0)

LSA ignored in SPF calculation
LS age: 136
Options: (V6-Bit, E-Bit, R-Bit, DC-Bit)
LS Type: Router Links
Link State ID: 0
Advertising Router: 192.168.3.3
LS Seq Number: 80000011
Checksum: 0x34D4
Length: 40
Area Border Router
Number of Links: 1

Link connected to: another Router (point-to-point)
Link Metric: 1
Local Interface ID: 4
Neighbor Interface ID: 5
Neighbor Router ID: 192.168.2.2
```

OSPFv3 LSA Flooding Scope

OSPFv3 Database Network

Example 9-21 shows the output of the command **show ospfv3 database network [self-originate]**.

The link LSA is responsible for providing details for the IPv6 prefixes associated with an interface. Example 9-22 shows the output of the command **show ospfv3 database link [self-originate]**.

Example 9-21 *OSPFv3 Database Network*

```
R1# show ospfv3 database network self-originate

OSPFv3 1 address-family ipv6 (router-id 192.168.1.1)

Net Link States (Area 0)

LS age: 1791
Options: (V6-Bit, E-Bit, R-Bit, DC-Bit)
LS Type: Network Links
Link State ID: 4 (Interface ID of Designated Router)
Advertising Router: 192.168.1.1
LS Seq Number: 8000000B
Checksum: 0x9F17
Length: 32
Attached Router: 192.168.1.1
Attached Router: 192.168.2.2
```

Example 9-22 *OSPFv3 Database Link*

```
R1# show ospfv3 database link self-originate

OSPFv3 1 address-family ipv6 (router-id 192.168.1.1)

Link (Type-8) Link States (Area 0)

LS age: 1572
Options: (V6-Bit, E-Bit, R-Bit, DC-Bit)
LS Type: Link-LSA (Interface: GigabitEthernet0/2)
Link State ID: 4 (Interface ID)
Advertising Router: 192.168.1.1
LS Seq Number: 8000000C
Checksum: 0x389C
Length: 56
Router Priority: 1
Link Local Address: FE80::1
Number of Prefixes: 1
Prefix Address: 2001:DB8:0:12::
Prefix Length: 64, Options: None

LS age: 1829
Options: (V6-Bit, E-Bit, R-Bit, DC-Bit)
LS Type: Link-LSA (Interface: GigabitEthernet0/1)
Link State ID: 3 (Interface ID)
Advertising Router: 192.168.1.1
LS Seq Number: 8000000B
Checksum: 0xBB2C
Length: 56
Router Priority: 1
Link Local Address: FE80::1
Number of Prefixes: 1
Prefix Address: 2001:DB8:0:1::
Prefix Length: 64, Options: None
```

OSPFv3 LSA Flooding Scope

OSPFv3 LSDB

Example 9-23 shows R3's database. Notice that R3's router LSA bits are set to B, indicating that it is an ABR router.

Example 9-23 Summary View of an OSPFv3 LSDB

```
R3# show ospfv3 database
```

```
! Output Omitted for brevity
```

```
OSPFv3 1 address-family ipv6 (router-id 192.168.3.3)
```

Router Link States (Area 0)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
192.168.1.1	416	0x80000005	0	1	None
192.168.2.2	375	0x80000007	0	2	None
192.168.3.3	351	0x80000005	0	1	B

Net Link States (Area 0)

ADV Router	Age	Seq#	Link ID	Rtr count
192.168.2.2	375	0x80000002	3	2
192.168.3.3	351	0x80000002	4	2

Router Link States (Area 34)

ADV Router	Age	Seq#	Fragment ID	Link count	Bits
192.168.3.3	351	0x80000004	0	1	B
192.168.4.4	399	0x80000005	0	1	None

Net Link States (Area 34)

ADV Router	Age	Seq#	Link ID	Rtr count
192.168.4.4	399	0x80000002	5	2

Inter Area Prefix Link States (Area 34)

ADV Router	Age	Seq#	Prefix
192.168.3.3	351	0x80000002	2001:DB8:0:23::/64
192.168.3.3	351	0x80000002	2001:DB8:0:12::/64
192.168.3.3	1572	0x80000001	2001:DB8:0:1::/64
192.168.3.3	6	0x80000001	2001:DB8::3/128
192.168.3.3	6	0x80000001	2001:DB8::2/128
192.168.3.3	6	0x80000001	2001:DB8::1/128

Link (Type-8) Link States (Area 34)

ADV Router	Age	Seq#	Link ID	Interface
192.168.3.3	351	0x80000002	6	Gi0/4
192.168.4.4	399	0x80000002	5	Gi0/4

Intra Area Prefix Link States (Area 34)

ADV Router	Age	Seq#	Link ID	Ref-ls type	Ref-LSID
192.168.4.4	399	0x80000002	5120	0x2002	5

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 9

Description
OSPFv3 fundamentals
OSPFv3 link-state advertisement
OSPFv3 communication
OSPFv3 configuration
OSPFv3 verification
IPv6 route summarization
Network type
OSPFv3 authentication
OSPFv3 flooding scope

Prepare for the Exam

Command Reference for Chapter 9

Task	Command Syntax
Configure OSPFv3 on a router and enable it on an interface	router ospfv3 [<i>process-id</i>] interface <i>interface-id</i> ospfv3 <i>process-id</i> { ipv4 ipv6 } area <i>area-id</i>
Configure a specific OSPFv3 interface as passive	passive-interface <i>interface-id</i>
Configure all OSPFv3 interfaces as passive	passive-interface default
Summarize an IPv6 network range on an ABR	area <i>area-id</i> range <i>prefix/prefix-length</i>
Configure an OSPFv3 interface as point-to-point or broadcast network type	ospfv3 network { point-to-point broadcast }
Display OSPFv3 interface settings	show ospfv3 interface [<i>interface-id</i>]
Display OSPFv3 IPv6 neighbors	show ospfv3 ipv6 neighbor
Display OSPFv3 router LSAs	show ospfv3 database router
Display OSPFv3 network LSAs	show ospfv3 database network
Display OSPFv3 link LSAs	show ospfv3 database link

