



CHAPTER 6

Configuring OSPFv3

This chapter describes how to configure Open Shortest Path First version 3 (OSPFv3) for IPv6 networks.

This chapter includes the following sections:

- [Information About OSPFv3, page 6-1](#)
- [Licensing Requirements for OSPFv3, page 6-12](#)
- [Prerequisites for OSPFv3, page 6-12](#)
- [Configuration Guidelines and Limitations, page 6-13](#)
- [Configuring Basic OSPFv3, page 6-13](#)
- [Configuring Advanced OSPFv3, page 6-19](#)
- [Verifying OSPFv3 Configuration, page 6-36](#)
- [Displaying OSPFv3 Statistics, page 6-37](#)
- [OSPFv3 Example Configuration, page 6-37](#)
- [Related Topics, page 6-37](#)
- [Default Settings, page 6-38](#)
- [Additional References, page 6-38](#)

Information About OSPFv3

OSPFv3 is an IETF link-state protocol (see “[Overview](#)” section on page 1-1). An OSPFv3 router sends a special message, called a *hello packet*, out each OSPF-enabled interface to discover other OSPFv3 neighbor routers. Once a neighbor is discovered, the two routers compare information in the Hello packet to determine if the routers have compatible configurations. The neighbor routers attempt to establish *adjacency*, which means that the routers synchronize their link-state databases to ensure that they have identical OSPFv3 routing information. Adjacent routers share *link-state advertisements* (LSAs) that include information about the operational state of each link, the cost of the link, and any other neighbor information. The routers then flood these received LSAs out every OSPF-enabled interface so that all OSPFv3 routers eventually have identical link-state databases. When all OSPFv3 routers have identical link-state databases, the network is *converged* (see the “[Convergence](#)” section on page 1-6). Each router then uses Dijkstra’s Shortest Path First (SPF) algorithm to build its route table.

You can divide OSPFv3 networks into areas. Routers send most LSAs only within one area, which reduces the CPU and memory requirements for an OSPF-enabled router.

OSPFv3 supports IPv6. For information about OSPF for IPv4, see [Chapter 5, “Configuring OSPFv2”](#).

Send document comments to nexus7k-docfeedback@cisco.com.

This section includes the following topics:

- [Comparison of OSPFv3 and OSPFv2, page 6-2](#)
- [Hello Packet, page 6-2](#)
- [Neighbors, page 6-3](#)
- [Adjacency, page 6-3](#)
- [Designated Routers, page 6-4](#)
- [Areas, page 6-5](#)
- [Link-State Advertisement, page 6-5](#)
- [OSPFv3 and the IPv6 Unicast RIB, page 6-7](#)
- [Authentication, page 6-8](#)
- [Address Family Support, page 6-8](#)
- [Advanced Features, page 6-8](#)

Comparison of OSPFv3 and OSPFv2

Much of the OSPFv3 protocol is the same as in OSPFv2. OSPFv3 is described in RFC 2740.

The key differences between the OSPFv3 and OSPFv2 protocols are as follows:

- OSPFv3 expands on OSPFv2 to provide support for IPv6 routing prefixes and the larger size of IPv6 addresses.
- LSAs in OSPFv3 are expressed as prefix and prefix length instead of address and mask.
- The router ID and area ID are 32-bit numbers with no relationship to IPv6 addresses.
- OSPFv3 uses link-local IPv6 addresses for neighbor discovery and other features.
- OSPFv3 uses IPv6 for authentication.
- OSPFv3 redefines LSA types.

Hello Packet

OSPFv3 routers periodically send Hello packets on every OSPF-enabled interface. The *hello interval* determines how frequently the router sends these Hello packets and is configured per interface. OSPFv3 uses Hello packets for the following tasks:

- Neighbor discovery
- Keepalives
- Bidirectional communications
- Designated router election (see the [“Designated Routers” section on page 6-4](#))

The Hello packet contains information about the originating OSPFv3 interface and router, including the assigned OSPFv3 cost of the link, the hello interval, and optional capabilities of the originating router. An OSPFv3 interface that receives these Hello packets determines if the settings are compatible with the receiving interface settings. Compatible interfaces are considered neighbors and are added to the neighbor table (see the [“Neighbors” section on page 6-3](#)).

Send document comments to nexus7k-docfeedback@cisco.com.

Hello packets also include a list of router IDs for the routers that the originating interface has communicated with. If the receiving interface sees its own router ID in this list, then bidirectional communication has been established between the two interfaces.

OSPFv3 uses Hello packets as a keepalive message to determine if a neighbor is still communicating. If a router does not receive a Hello packet by the configured *dead interval* (usually a multiple of the hello interval), then the neighbor is removed from the local neighbor table.

Neighbors

An OSPFv3 interface must have a compatible configuration with a remote interface before the two can be considered neighbors. The two OSPFv3 interfaces must match the following criteria:

- Hello interval
- Dead interval
- Area ID (see the “[Areas](#)” section on page 6-5)
- Authentication
- Optional capabilities

If there is a match, the information is entered into the neighbor table:

- Neighbor ID—The router ID of the neighbor router.
- Priority—Priority of the neighbor router. The priority is used for designated router election (see the “[Designated Routers](#)” section on page 6-4).
- State—Indication of whether the neighbor has just been heard from, is in the process of setting up bidirectional communications, is sharing the link-state information, or has achieved full adjacency.
- Dead time—Indication of how long since the last Hello packet was received from this neighbor.
- Link-local IPv6 Address—The link-local IPv6 address of the neighbor.
- Designated Router—Indication of whether the neighbor has been declared the designated router or backup designated router (see the “[Designated Routers](#)” section on page 6-4).
- Local interface—The local interface that received the Hello packet for this neighbor.

When the first Hello packet is received from a new neighbor, the neighbor is entered into the neighbor table in the init state. Once bidirectional communication is established, the neighbor state becomes 2-way. ExStart and exchange states come next, as the two interfaces exchange their link-state database. Once this is all complete, the neighbor moves into the full state, signifying full adjacency. If the neighbor fails to send any Hello packets in the dead interval, then the neighbor is moved to the down state and is no longer considered adjacent.

Adjacency

Not all neighbors establish adjacency. Depending on the network type and designated router establishment, some neighbors become fully adjacent and share LSAs with all their neighbors, while other neighbors do not. For more information, see the “[Designated Routers](#)” section on page 6-4.

Adjacency is established using Database Description packets, Link State Request packets, and Link State Update packets in OSPFv3. The Database Description packet includes just the LSA headers from the link-state database of the neighbor (see the “[Link-State Database](#)” section on page 6-7). The local router compares these headers with its own link-state database and determines which LSAs are new or

Send document comments to nexus7k-docfeedback@cisco.com.

updated. The local router sends a Link State Request packet for each LSA that it needs new or updated information on. The neighbor responds with a Link State Update packet. This exchange continues until both routers have the same link-state information.

Designated Routers

Networks with multiple routers present a unique situation for OSPFv3. If every router floods the network with LSAs, the same link-state information will be sent from multiple sources. Depending on the type of network, OSPFv3 might use a single router, the *designated router (DR)*, to control the LSA floods and represent the network to the rest of the OSPFv3 area (see the “Areas” section on page 6-5). If the DR fails, OSPFv3 selects a *backup designated router (BDR)*. If the DR fails, OSPFv3 uses the BDR.

Network types are as follows:

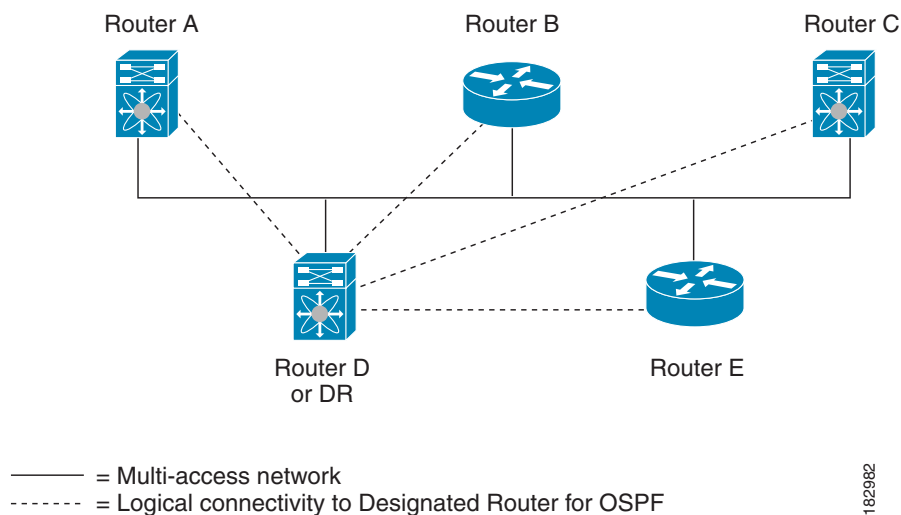
- Point-to-point—A network that exists only between two routers. All neighbors on a point-to-point network establish adjacency and there is no DR.
- Broadcast—A network with multiple routers that can communicate over a shared medium that allows broadcast traffic, such as Ethernet. OSPFv3 routers establish a DR and BDR that controls LSA flooding on the network. OSPFv3 uses the well-known IPv6 multicast addresses, FF02::5, and a MAC address of 0100.5300.0005 to communicate with neighbors.

The DR and BDR are selected based on the information in the Hello packet. When an interface sends a Hello packet, it sets the priority field and the DR and BDR field if it knows who the DR and BDR are. The routers follow an election procedure based on which routers declare themselves in the DR and BDR fields and the priority field in the Hello packet. As a final tie breaker, OSPFv3 chooses the highest router IDs as the DR and BDR.

All other routers establish adjacency with the DR and the BDR and use the IPv6 multicast address FF02::6 to send LSA updates to the DR and BDR. Figure 6-1 shows this adjacency relationship between all routers and the DR.

DRs are based on a router interface. A router might be the DR for one network and not for another network on a different interface.

Figure 6-1 DR in Multi-Access Network



182982


Send document comments to nexus7k-docfeedback@cisco.com.

Areas

You can limit the CPU and memory requirements that OSPFv3 puts on the routers by dividing an OSPFv3 network into *areas*. An area is a logical division of routers and links within an OSPFv3 domain that creates separate subdomains. LSA flooding is contained within an area, and the link-state database is limited to links within the area. You can assign an area ID to the interfaces within the defined area. The Area ID is a 32-bit value that can be expressed as a number or in dotted decimal notation, such as 10.2.3.1.

If you define more than one area in an OSPFv3 network, you must also define the backbone area, which has the reserved area ID of 0. If you have more than one area, then one or more routers become area border routers (ABRs). An ABR connects to both the backbone area and at least one other defined area (see [Figure 6-2](#)).

Figure 6-2 OSPFv3 Areas



182983

The ABR has a separate link-state database for each area which it connects to. The ABR sends Inter-Area Prefix (type 3) LSAs (see the [“Route Summarization”](#) section on page 6-10) from one connected area to the backbone area. The backbone area sends summarized information about one area to another area. In [Figure 6-2](#), Area 0 sends summarized information about Area 5 to Area 3.

OSPFv3 defines one other router type: the autonomous system boundary router (ASBR). This router connects an OSPFv3 area to another autonomous system. An autonomous system is a network controlled by a single technical administration entity. OSPFv3 can redistribute its routing information into another autonomous system or receive redistributed routes from another autonomous system. For more information, see [“Advanced Features”](#) section on page 6-8.

Link-State Advertisement

OSPFv3 uses link-state advertisements (LSAs) to build its routing table.

This section includes the following topics:

- [LSA Types, page 6-6](#)
- [Link Cost, page 6-6](#)

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

- [Flooding and LSA Group Pacing](#), page 6-6
- [Link-State Database](#), page 6-7

LSA Types

Table 6-1 shows the LSA types supported by Cisco NX-OS.

Table 6-1 LSA Types

Type	Name	Description
1	Router LSA	LSA sent by every router. This LSA includes state and cost of all links. Does not include prefix information. Router LSAs trigger an SPF recalculation. Router LSAs are flooded to the local OSPFv3 area.
2	Network LSA	LSA sent by the DR. Lists all routers in the multi-access network. This LSA does not include prefix information. Network LSAs trigger an SPF recalculation. See the “Designated Routers” section on page 6-4.
3	Inter-Area Prefix LSA	LSA sent by the area border router to an external area for each destination in local area. This LSA includes the link cost from area the border router to the local destination. See the “Areas” section on page 6-5.
4	Inter-Area Router LSA	LSA sent by the area border router to an external area. This LSA advertises the link cost to the ASBR only. See the “Areas” section on page 6-5.
5	AS External LSA	LSA generated by the ASBR. This LSA includes the link cost to an external autonomous system destination. AS External LSAs are flooded throughout the autonomous system. See the “Areas” section on page 6-5.
7	Type-7 LSA	LSA generated by the ASBR within an NSSA. This LSA includes the link cost to an external autonomous system destination. Type-7 LSAs are flooded only within the local NSSA. See the “Areas” section on page 6-5.
8	Link LSA	LSA sent by every router, using a link-local flooding scope (see the “Flooding and LSA Group Pacing” section on page 6-6). This LSA includes the link-local address and IPv6 prefixes for this link.
9	Intra-Area Prefix LSA	LSA sent by every router. This LSA includes any prefix or link state changes. Intra-Area Prefix LSAs are flooded to the local OSPFv3 area. This LSA does not trigger an SPF recalculation.
11	Grace LSAs	LSA sent by a restarting router, using a link-local flooding scope. This LSA is used for a graceful restart of OSPFv3. See the “High Availability and Graceful Restart” section on page 6-11.

Link Cost

Each OSPFv3 interface is assigned a *link cost*. The cost is an arbitrary number. By default, Cisco NX-OS assigns a cost that is the configured reference bandwidth divided by the interface bandwidth. By default, the reference bandwidth is 40 Gb/s. The link cost is carried in the LSA updates for each link.

Flooding and LSA Group Pacing

OSPFv3 floods LSA updates to different sections of the network, depending on the LSA type. OSPFv3 uses the following flooding scopes:

Send document comments to nexus7k-docfeedback@cisco.com.

- Link-local—LSA is flooded only on the local link, and no further. Used for Link LSAs and Grace LSAs.
- Area-local—LSA is flooded throughout a single OSPF area only. Used for Router LSAs, Network LSAs, Inter-Area-Prefix LSAs, Inter-Area-Router LSAs, and Intra-Area-Prefix LSAs.
- AS scope—LSA is flooded throughout the routing domain. Used for AS External LSAs.

LSA flooding guarantees that all routers in the network have identical routing information. LSA flooding depends on the OSPFv3 area configuration (see the “[Areas](#)” section on page 6-5). The LSAs are flooded based on the *link-state refresh* time (every 30 minutes by default). Each LSA has its own link-state refresh time.

You can control the flooding rate of LSA updates in your network by using the LSA group pacing feature. LSA group pacing can reduce high CPU or buffer utilization. This feature groups LSAs with similar link-state refresh times to allow OSPFv3 to pack multiple LSAs into an OSPFv3 Update message.

By default, LSAs with link-state refresh times within four minutes of each other are grouped together. You should lower this value for large link-state databases or raise it for smaller databases to optimize the OSPFv3 load on your network.

Link-State Database

Each router maintains a link-state database for the OSPFv3 network. This database contains all the collected LSAs, and includes information on all the routes through the network. OSPFv3 uses this information to calculate the best path to each destination and populates the routing table with these best paths.

LSAs are removed from the link-state database if no LSA update has been received within a set interval, called the MaxAge. Routers flood a repeat of the LSA every 30 minutes to prevent accurate link-state information from being aged out. Cisco NX-OS supports the LSA grouping feature to prevent all LSAs from refreshing at the same time. For more information, see the “[Flooding and LSA Group Pacing](#)” section on page 6-6.

OSPFv3 and the IPv6 Unicast RIB

OSPFv3 runs the Dijkstra shortest path first algorithm on the link-state database. This algorithm selects the best path to each destination based on the sum of all the link costs for each link in the path. The resultant shortest path for each destination is then put in the OSPFv3 route table. When the OSPFv3 network is converged, this route table feeds into the IPv6 unicast RIB. OSPFv3 communicates with the IPv6 unicast RIB to do the following:

- Add or remove routes
- Handle route redistribution from other protocols
- Provide convergence updates to remove stale OSPFv3 routes and for stub router advertisements (see the “[Multiple OSPFv3 Instances](#)” section on page 6-11)

OSPFv3 also runs a modified Dijkstra algorithm for fast recalculation for Inter-Area Prefix, Inter-Area Router, AS-External, Type-7, and Intra-Area Prefix (type 3, 4, 5, 7, 8) LSA changes.

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

Authentication

OSPFv3 relies on the IPv6 Authentication Header and the IPv6 Encapsulating Security Payload to ensure integrity, authentication, and confidentiality of OSPFv3 packets. (see [Chapter 3, “Configuring IPv6”](#)).

Address Family Support

Cisco NX-OS supports multiple address families, such as unicast IPv6 and multicast IPv6. OSPFv3 features that are specific to an *address family* are as follows:

- Default routes
- Route summarization
- Route redistribution
- Filter lists for border routers
- SPF optimization

Use the **address-family ipv6 unicast** command to enter the IPv6 unicast address family configuration mode when configuring these features.

Advanced Features

Cisco NX-OS supports a number of advanced OSPFv3 features that enhance the usability and scalability of OSPFv3 in the network.

This section includes the following topics:

- [Stub Area, page 6-8](#)
- [Not-So-Stubby Area, page 6-9](#)
- [Virtual Links, page 6-9](#)
- [Route Redistribution, page 6-10](#)
- [Route redistribution uses route maps to control which external routes are redistributed. See Chapter 15, “Configuring Route Policy Manager,” for details on configuring route maps. You can use route maps to modify parameters in the AS External \(type 5\) and Type-7 LSAs before these external routes are advertised in the local OSPFv3 autonomous system., page 6-10](#)
- [Route Summarization, page 6-10](#)
- [High Availability and Graceful Restart, page 6-11](#)
- [Multiple OSPFv3 Instances, page 6-11](#)
- [SPF Optimization, page 6-12](#)
- [Virtualization Support, page 6-12](#)

Stub Area


You can limit the amount of external routing information that floods an area by making it a *stub area*. A stub area is an area that does not allow AS External (type 5) LSAs (see the [“Link-State Advertisement” section on page 6-5](#)). These LSAs are usually flooded throughout the local autonomous system to propagate external route information. Stub areas have the following requirements:

Send document comments to nexus7k-docfeedback@cisco.com.

- All routers in the stub area are stub routers. See the “Stub Routing” section on page 1-7.
- No ASBR routers exist in the stub area.
- You cannot configure virtual links in the stub area.

Figure 6-3 shows an example an OSPFv3 autonomous system where all routers in Area 10 have to go through the ABR to reach external autonomous systems. Area 10 can be configured as a stub area.

Figure 6-3 Stub Area



Stub areas use a default route for all traffic that needs to go through the backbone area to the external autonomous system. The default route is an Inter-Area-Prefix LSA with prefix length set to 0 for IPv6.

Not-So-Stubby Area

A Not-So-Stubby Area (*NSSA*) is similar to the stub area, except that an NSSA allows you to import autonomous system external routes within an NSSA using redistribution. The NSSA ASBR redistributes these routes and generates Type-7 LSAs that it floods throughout the NSSA. You can optionally configure the ABR that connects the NSSA to other areas to translate this Type-7 LSA to AS External (type 5) LSAs. The ABR then floods these AS External LSAs throughout the OSPFv3 autonomous system. Summarization and filtering are supported during the translation. See the “Link-State Advertisement” section on page 6-5 for details on Type-7 LSAs.

You can, for example, use NSSA to simplify administration if you are connecting a central site using OSPFv3 to a remote site that is using a different routing protocol. Before NSSA, the connection between the corporate site border router and a remote router could not be run as an OSPFv3 stub area because routes for the remote site could not be redistributed into a stub area. You needed to maintain two routing protocols. With NSSA, you can extend OSPFv3 to cover the remote connection by defining the area between the corporate router and remote router as an NSSA (see the “Configuring NSSA” section on page 6-23).


The backbone Area 0 cannot be an NSSA.

Virtual Links

Virtual links allow you to connect an OSPFv3 area ABR to a backbone area ABR when a direct physical connection is not available. Figure 6-4 shows a virtual link that connects Area 3 to the backbone area through Area 5.

Send document comments to nexus7k-docfeedback@cisco.com.

Figure 6-4 Virtual Links



182985

You can also use virtual links to temporarily recover from a partitioned area, which occurs when a link within the area fails, isolating part of the area from reaching the designated ABR to the backbone area.

Route Redistribution

OSPFv3 can learn routes from other routing protocols by using route redistribution. See the [“Route Redistribution” section on page 1-6](#). You configure OSPFv3 to assign a link cost for these redistributed routes or a default link cost for all redistributed routes.

Route redistribution uses route maps to control which external routes are redistributed. See [Chapter 15, “Configuring Route Policy Manager,”](#) for details on configuring route maps. You can use route maps to modify parameters in the AS External (type 5) and Type-7 LSAs before these external routes are advertised in the local OSPFv3 autonomous system.

Route Summarization

Because OSPFv3 shares all learned routes with every OSPF-enabled router, you might want to use route summarization to reduce the number of unique routes that are flooded to every OSPF-enabled router. Route summarization simplifies route tables by replacing more-specific addresses with an address that represents all the specific addresses. For example, you can replace 2010:11:22:0:1000::1 and 2010:11:22:0:2000:679:1 with one summary address, 2010:11:22::/32.

Typically, you would summarize at the boundaries of area border routers (ABRs). Although you could configure summarization between any two areas, it is better to summarize in the direction of the backbone so that the backbone receives all the aggregate addresses and injects them, already summarized, into other areas. The two types of summarization are as follows:

- Inter-area route summarization
- External route summarization

You configure inter-area route summarization on ABRs, summarizing routes between areas in the autonomous system. To take advantage of summarization, you should assign network numbers in areas in a contiguous way to be able to lump these addresses into one range.

Send document comments to nexus7k-docfeedback@cisco.com.

External route summarization is specific to external routes that are injected into OSPFv3 using route redistribution. You should make sure that external ranges that are being summarized are contiguous. Summarizing overlapping ranges from two different routers could cause packets to be sent to the wrong destination. Configure external route summarization on ASBRs that are redistributing routes into OSPF.

When you configure a summary address, Cisco NX-OS automatically configures a discard route for the summary address to prevent routing black holes and route loops.

High Availability and Graceful Restart

Cisco NX-OS supports high-availability. If a Cisco NX-OS system experiences a cold reboot, the network stops forwarding traffic to the system and removes the system from the network topology. In this scenario, OSPFv3 experiences a stateless restart, and removes all neighbor adjacencies on the local system. Cisco NX-OS applies the startup configuration and OSPFv3 rediscovers the neighbors and establishes the adjacencies again.

OSPFv3 automatically restarts if the process experiences problems. After the restart, OSPFv3 initiates a graceful restart so that the platform is not taken out of the network topology. If you manually restart OSPF, it performs a graceful restart, which is similar to a stateful switchover. The running configuration is applied in both cases.

A graceful restart, or nonstop forwarding (NSF), allows OSPFv3 to remain in the data forwarding path through a process restart. When OSPFv3 needs to restart, it first sends a link-local Grace (type 11) LSA. This restarting OSPFv3 platform is called NSF capable.

The Grace LSA includes a grace period, which is a specified time that the neighbor OSPFv3 interfaces hold onto the LSAs from the restarting OSPFv3 interface. (Typically, OSPFv3 tears down the adjacency and discards all LSAs from a down or restarting OSPFv3 interface.) The participating neighbors, which are called NSF helpers, keep all LSAs that originate from the restarting OSPFv3 interface as if the interface were still adjacent.

When the restarting OSPFv3 interface is operational again, it rediscovers its neighbors, establishes adjacency, and starts sending its LSA updates again. At this point, the NSF helpers recognize that graceful restart has finished.

**Note**

If the restarting OSPFv3 interface does not come back up before the end of the grace period, or if the network experiences a topology change, the OSPFv3 neighbors tear down adjacency with the restarting OSPFv3 and treat it as a normal OSPFv3 restart.

**Note**

You must enable graceful restart to support in-service software upgrades (ISSU) for OSPFv3. If you disable graceful restart, Cisco NX-OS issues a warning that ISSU cannot be supported with this configuration.

Multiple OSPFv3 Instances

Cisco NX-OS supports multiple instances of the OSPFv3 protocol. By default, every instance uses the same system router ID. You must manually configure the router ID for each instance if the instances are in the same OSPFv3 autonomous system.

The OSPFv3 header includes an instance ID field to identify that OSPFv3 packet for a particular OSPFv3 instance. You can assign the OSPFv3 instance. The interface drops all OSPFv3 packets that do not have a matching OSPFv3 instance ID in the packet header.

Send document comments to nexus7k-docfeedback@cisco.com.

Cisco NX-OS allows only one OSPFv3 instance on an interface.

SPF Optimization

Cisco NX-OS optimizes the SPF algorithm in the following ways:

- Partial SPF for Network (type 2) LSAs, Inter-Area Prefix (type 3) LSAs, and AS External (type 5) LSAs—When there is a change on any of these LSAs, Cisco NX-OS performs a faster partial calculation rather than running the whole SPF calculation.
- SPF timers—You can configure different timers for controlling SPF calculations. These timers include exponential backoff for subsequent SPF calculations. The exponential backoff limits the CPU load of multiple SPF calculations.

Virtualization Support

OSPFv3 supports Virtual Routing and Forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0* and see [Chapter 14, “Configuring Layer 3 Virtualization.”](#)

Licensing Requirements for OSPFv3

The following table shows the licensing requirements for this feature:

Product	License Requirement
NX-OS	OSPFv3 requires an Enterprise Services license. For a complete explanation of the NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco Nexus 7000 Series NX-OS Licensing Guide, Release 4.0</i> .

Prerequisites for OSPFv3

OSPFv3 has the following prerequisites:

- You must be familiar with routing fundamentals to configure OSPFv3.
- You must be logged on to the switch.
- You have configured at least one interface for IPv6 that is capable of communicating with a remote OSPFv3 neighbor.
- You have installed the Enterprise Services license.
- You have completed the OSPFv3 network strategy and planning for your network. For example, you must decide whether multiple areas are required.
- You have enabled the OSPF feature (see the [“Enabling the OSPFv3 Feature”](#) section on page 6-13).
- You have installed the Advanced Services license and entered the desired VDC (see to the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0*) if you are configuring VDCs.

Send document comments to nexus7k-docfeedback@cisco.com.

- You are familiar with IPv6 addressing and basic configuration. See [Chapter 3, “Configuring IPv6”](#) for information on IPv6 routing and addressing.

Configuration Guidelines and Limitations

OSPFv3 has the following configuration guidelines and limitations:

- You can have up to four instances of OSPFv3 in a VDC.



Note

If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Configuring Basic OSPFv3

Configure OSPFv3 after you have designed your OSPFv3 network.

This section includes the following topics:

- [Enabling the OSPFv3 Feature, page 6-13](#)
- [Creating an OSPFv3 Instance, page 6-14](#)
- [Configuring Networks in OSPFv3, page 6-16](#)

Enabling the OSPFv3 Feature

You must enable the OSPFv3 feature before you can configure OSPFv3.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **feature ospfv3**
3. **copy running-config startup-config**

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	<code>feature ospfv3</code> Example: switch(config)# feature ospfv3	Enables the OSPFv3 feature.
Step 3	<code>copy running-config startup-config</code> Example: switch(config)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **no feature ospfv3** command to disable the OSPFv3 feature and remove all associated configuration.

Command	Purpose
<code>no feature ospfv3</code> Example: switch(config)# no feature ospfv3	Disables the OSPFv3 feature and removes all associated configuration.

Creating an OSPFv3 Instance

The first step in configuring OSPFv3 is to create an instance or OSPFv3 instance. You assign a unique instance tag for this OSPFv3 instance. The instance tag can be any string. For each OSPFv3 instance, you can also configure the following optional parameters:

- Router ID—Configures the router ID for this OSPFv3 instance. If you do not use this parameter, the router ID selection algorithm is used. See the [“Router IDs” section on page 1-5](#).
- Administrative distance—Rates the trustworthiness of a routing information source. See the [“Administrative Distance” section on page 1-6](#).
- Log adjacency changes—Creates a system message whenever an OSPFv3 neighbor changes its state.
- Maximum paths—Sets the maximum number of equal paths that OSPFv3 installs in the route table for a particular destination. Use this parameter for load balancing between multiple paths.
- Reference bandwidth—Controls the calculated OSPFv3 cost metric for a network. The calculated cost is the reference bandwidth divided by the interface bandwidth. You can override the calculated cost by assigning a link cost when a network is added to the OSPFv3 instance. See the [“Configuring Networks in OSPFv3” section on page 6-16](#).

For more information about OSPFv3 instance parameters, see the [“Configuring Advanced OSPFv3” section on page 6-19](#).

Send document comments to nexus7k-docfeedback@cisco.com.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPFv3 feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that the OSPFv3 instance tag that you plan on using is not already in use on this router.

Use the **show ospfv3 instance-tag** command to verify that the instance tag is not in use.

OSPFv3 must be able to obtain a router identifier (for example, a configured loopback address) or you must configure the router ID option.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **router ospfv3 instance-tag**
3. **router-id ip-address**
4. **show ipv6 ospfv3 instance-tag**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router ospfv3 instance-tag Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	router-id id Example: switch(config-router)# router-id 209.0.2.1	(Optional) Configures the OSPFv3 router ID. This ID uses the dotted decimal notation and identifies this OSPFv3 instance and must exist on a configured interface in the system.
Step 4	show ipv6 ospfv3 instance-tag Example: switch(config-router)# show ipv6 ospfv3 201	(Optional) Displays OSPFv3 information.
Step 5	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **no router ospfv3** command to remove the OSPFv3 instance and all associated configuration.

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
no router ospfv3 <i>instance-tag</i> Example: switch(config)# no router ospfv3 201	Deletes the OSPFv3 instance and all associated configuration.

**Note**

This command does not remove OSPF configuration in interface mode. You must manually remove any OSPFv3 commands configured in interface mode.

You can configure the following optional parameters for OSPFv3 in router configuration mode:

Command	Purpose
log-adjacency-changes [detail] Example: switch(config-router)# log-adjacency-changes	Generates a system message whenever a neighbor changes state.

Use the **address-family ipv6 unicast** command to configure the following optional parameters for OSPFv3 in address family configuration mode:

Command	Purpose
distance <i>number</i> Example: switch(config-router-af)# distance 25	Configures the administrative distance for this OSPFv3 instance. The range is from 1 to 255. The default is 110.
maximum-paths <i>paths</i> Example: switch(config-router-af)# maximum-paths 4	Configures the maximum number of equal OSPFv3 paths to a destination in the route table. The range is from 1 to 16. The default is 8. Used for load balancing.

The following example shows how to create an OSPFv3 instance:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# copy running-config startup-config
```

Configuring Networks in OSPFv3

You can configure a network to OSPFv3 by associating it through the interface that the router uses to connect to that network (see the “Neighbors” section on page 6-3). You can add all networks to the default backbone area (Area 0), or you can create new areas using any decimal number or an IP address.

**Note**

All areas must connect to the backbone area either directly or through a virtual link.

Send document comments to nexus7k-docfeedback@cisco.com.

**Note**

OSPFv3 is not enabled on an interface until you configure a valid IPv6 address for that interface.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPFv3 feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

SUMMARY STEPS

1. `config t`
2. `interface interface-type slot/port`
3. `ipv6 address ipv6-prefix/length`
4. `ipv6 router ospfv3 instance-tag area area-id [secondaries none]`
5. `show ipv6 ospfv3 instance-tag interface interface-type slot/port`
6. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	<code>interface interface-type slot/port</code> Example: switch(config)# interface ethernet 1/2 switch(config-if)#	Enters interface configuration mode.
Step 3	<code>ipv6 address ipv6-prefix/length</code> Example: switch(config-if)# ipv6 address 2001:0DB8::1/48	Assigns an IPv6 address to this interface.
Step 4	<code>ipv6 router ospfv3 instance-tag area area-id [secondaries none]</code> Example: switch(config-if)# ipv6 router ospfv3 201 area 0	Adds the interface to the OSPFv3 instance and area.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 5	<pre>show ipv6 ospfv3 instance-tag interface interface-type slot/port</pre> <p>Example: switch(config-if)# show ipv6 ospfv3 201 interface ethernet 1/2 </p>	(Optional) Displays OSPFv3 information.
Step 6	<pre>copy running-config startup-config</pre> <p>Example: switch(config)# copy running-config startup-config </p>	(Optional) Saves this configuration change.

You can configure the following optional parameters for OSPFv3 in interface configuration mode:

Command	Purpose
<pre>ospfv3 cost number</pre> <p>Example: switch(config-if)# ospfv3 cost 25 </p>	Configures the OSPFv3 cost metric for this interface. The default is to calculate cost metric, based on reference bandwidth and interface bandwidth. The range is from 1 to 65535.
<pre>ospfv3 dead-interval seconds</pre> <p>Example: switch(config-if)# ospfv3 dead-interval 50 </p>	Configures the OSPFv3 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
<pre>ospfv3 hello-interval seconds</pre> <p>Example: switch(config-if)# ospfv3 hello-interval 25 </p>	Configures the OSPFv3 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
<pre>ospfv3 instance instance</pre> <p>Example: switch(config-if)# ospfv3 instance 25 </p>	Configures the OSPFv3 instance ID. The range is from 0 to 255. The default is 0. The instance ID is link-local in scope.
<pre>ospfv3 mtu-ignore</pre> <p>Example: switch(config-if)# ospfv3 mtu-ignore </p>	Configures OSPFv3 to ignore any IP maximum transmission unit (MTU) mismatch with a neighbor. The default is to not establish adjacency if the neighbor MTU does not match the local interface MTU.
<pre>ospfv3 network {broadcast point-point}</pre> <p>Example: switch(config-if)# ospfv3 network broadcast </p>	Sets the OSPFv3 network type.
<pre>ospfv3 passive-interface</pre> <p>Example: switch(config-if)# ospfv3 passive-interface </p>	Suppresses routing updates on the interface.

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
ospfv3 priority <i>number</i> Example: switch(config-if)# ospfv3 priority 25	Configures the OSPFv3 priority, used to determine the DR for an area. The range is from 0 to 255. The default is 1. See the “Designated Routers” section on page 6-4.
ospfv3 shutdown Example: switch(config-if)# ospfv3 shutdown	Shuts down the OSPFv3 instance on this interface.

The following example shows how to add a network area 10 in OSPFv3 instance 201:

```
switch# config t
switch(config)# interface ethernet 1/2
switch(config-if)# ipv6 address 2001:0DB8::1/48
switch(config-if)# ipv6 ospfv3 201 area 10
switch(config-if)# copy running-config startup-config
```

Configuring Advanced OSPFv3

Configure OSPFv3 after you have designed your OSPFv3 network.

This section includes the following topics:

- [Configuring Filter Lists for Border Routers, page 6-19](#)
- [Configuring Stub Areas, page 6-21](#)
- [Configuring a Totally Stubby Area, page 6-22](#)
- [Configuring NSSA, page 6-23](#)
- [Configuring Virtual Links, page 6-25](#)
- [Configuring Redistribution, page 6-27](#)
- [Configuring Route Summarization, page 6-29](#)
- [Modifying the Default Timers, page 6-30](#)
- [Configuring Graceful Restart, page 6-32](#)
- [Restarting an OSPFv3 Instance, page 6-34](#)
- [Configuring OSPFv3 with Virtualization, page 6-34](#)

Configuring Filter Lists for Border Routers

You can separate your OSPFv3 domain into a series of areas that contain related networks. All areas must connect to the backbone area through an area border router (ABR). OSPFv3 domains can connect to external domains as well, through an autonomous system border router (ASBR). See the “Areas” section on page 6-5.

ABRs have the following optional configuration parameters:

- Area range—Configures route summarization between areas. See the “Configuring Route Summarization” section on page 6-29.

Send document comments to nexus7k-docfeedback@cisco.com.

- Filter list—Filters the Inter-Area Prefix (type 3) LSAs on an ABR that are allowed in from an external area.

ASBRs also support filter lists.

BEFORE YOU BEGIN

Create the route map that the filter list uses to filter ip prefixes in incoming or outgoing Inter-Area Prefix (type 3) LSAs. See [Chapter 15, “Configuring Route Policy Manager.”](#)

Ensure that you have enabled the OSPFv3 feature (see the [“Enabling the OSPFv3 Feature”](#) section on [page 6-13](#)).

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

SUMMARY STEPS

1. `config t`
2. `router ospfv3 instance-tag`
3. `address-family ipv6 unicast`
4. `area area-id filter-list route-map map-name {in | out}`
5. `show ipv6 ospfv3 policy statistics area id filter-list {in | out}`
6. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	<code>router ospfv3 instance-tag</code> Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	<code>address-family ipv6 unicast</code> Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	Enters IPv6 unicast address family mode.
Step 4	<code>area area-id filter-list route-map map-name {in out}</code> Example: switch(config-router-af)# area 10 filter-list route-map FilterLSAs in	Filters incoming or outgoing Inter-Area Prefix (type 3) LSAs on an ABR.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 5	<pre>show ipv6 ospfv3 policy statistics area id filter-list {in out}</pre> <p>Example: switch(config-if)# show ipv6 ospfv3 policy statistics area 10 filter-list in</p>	(Optional) Displays OSPFv3 policy information.
Step 6	<pre>copy running-config startup-config</pre> <p>Example: switch(config-router)# copy running-config startup-config</p>	(Optional) Saves this configuration change.

The following example shows how to enable graceful restart if it has been disabled:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router-af)# area 10 filter-list route-map FilterLSAs in
switch(config-router-af)# copy running-config startup-config
```

Configuring Stub Areas

You can configure a stub area for part of an OSPFv3 domain where external traffic is not necessary. Stub areas block AS External (type 5) LSAs, limiting unnecessary routing to and from selected networks. See the “[Stub Area](#)” section on page 6-8. You can optionally block all summary routes from going into the stub area.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPF feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that there are no virtual links or ASBRs in the proposed stub area.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

SUMMARY STEPS

1. `config t`
2. `router ospfv3 instance-tag`
3. `area area-id stub`
4. `address-family ipv6 unicast`
5. `area area-id default-cost cost`
6. `copy running-config startup-config`

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router ospfv3 instance-tag Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	area area-id stub Example: switch(config-router)# area 10 stub	Creates this area as a stub area.
Step 4	address-family ipv6 unicast Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	(Optional) Enters IPv6 unicast address family mode.
Step 5	area area-id default-cost cost Example: switch(config-router-af)# area 10 default-cost 25	(Optional) Sets the cost metric for the default summary route sent into this stub area. The range is from 0 to 16777215.
Step 6	copy running-config startup-config Example: switch(config-router)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to create a stub area that blocks all summary route updates:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 stub no-summary
switch(config-router)# copy running-config startup-config
```

Configuring a Totally Stubby Area

You can create a totally stubby area and prevent all summary route updates from going into the stub area.

To create a totally stubby area, use the following command in router configuration mode:

Command	Purpose
area area-id stub no-summary Example: switch(config-router)# area 20 stub no-summary	Creates this area as a totally stubby area.

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring NSSA

You can configure an NSSA for part of an OSPFv3 domain where limited external traffic is required. See the “[Not-So-Stubby Area](#)” section on page 6-9. You can optionally translate this external traffic to an AS External (type 5) LSA and flood the OSPFv3 domain with this routing information. An NSSA can be configured with the following optional parameters:

- **No redistribution**—Redistributed routes bypass the NSSA and are redistributed to other areas in the OSPFv3 autonomous system. Use this option when the NSSA ASBR is also an ABR.
- **Default information originate**—Generates a Type-7 LSA for a default route to the external autonomous system. Use this option on an NSSA ASBR if the ASBR contains the default route in the routing table. This option can be used on an NSSA ABR whether or not the ABR contains the default route in the routing table.
- **Route map**—Filters the external routes so that only those routes you want are flooded throughout the NSSA and other areas.
- **Translate**—Translates Type-7 LSAs to AS External (type 5) LSAs for areas outside the NSSA. Use this command on an NSSA ABR to flood the redistributed routes throughout the OSPFv3 autonomous system. You can optionally suppress the forwarding address in these AS External LSAs.
- **No summary**—Blocks all summary routes from flooding the NSSA. Use this option on the NSSA ABR.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPF feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that there are no virtual links in the proposed NSSA and that it is not the backbone area.

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **router ospfv3 *instance-tag***
3. **area *area-id* nssa [no-redistribution] [default-information-originate [route-map *map-name*]] [no-summary] [translate type7 {always | never}] [suppress-fa]**
4. **address-family ipv6 unicast**
5. **area *area-id* default-cost *cost***
6. **copy running-config startup-config**

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router ospfv3 instance-tag Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	area area-id nssa [no-redistribution] [default-information-originate] [route-map map-name] [no-summary] [translate type7 {always never}] [suppress-fa] Example: switch(config-router)# area 10 nssa	Creates this area as an NSSA.
Step 4	address-family ipv6 unicast Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	(Optional) Enters IPv6 unicast address family mode.
Step 5	area area-id default-cost cost Example: switch(config-router-af)# area 10 default-cost 25	(Optional) Sets the cost metric for the default summary route sent into this NSSA. The range is from 0 to 16777215.
Step 6	copy running-config startup-config Example: switch(config-router)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to create an NSSA that blocks all summary route updates:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 nssa no-summary
switch(config-router)# copy running-config startup-config
```

The following example shows how to create an NSSA that generates a default route;

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 nssa default-info-originate
switch(config-router)# copy running-config startup-config
```

The following example shows how to create an NSSA that filters external routes and blocks all summary route updates:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 nssa route-map ExternalFilter no-summary
switch(config-router)# copy running-config startup-config
```

Send document comments to nexus7k-docfeedback@cisco.com.

The following example shows how to create an NSSA that always translates Type-7 LSAs to AS External (type 5) LSAs:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 nssa translate type 7 always
switch(config-router)# copy running-config startup-config
```

The following example shows how to create an NSSA that blocks all summary route updates:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 nssa no-summary
switch(config-router)# copy running-config startup-config
```

Configuring Virtual Links

A virtual link connects an isolated area to the backbone area through an intermediate area. See the “[Virtual Links](#)” section on page 6-9. You can configure the following optional parameters for a virtual link:

- Authentication—Sets simple password or MD5 message digest authentication and associated keys.
- Dead interval—Sets the time that a neighbor waits for a Hello packet before declaring the local router as dead and tearing down adjacencies.
- Hello interval—Sets the time between successive Hello packets.
- Retransmit interval—Sets the estimated time between successive LSAs.
- Transmit delay—Sets the estimated time to transmit an LSA to a neighbor.



Note

You must configure the virtual link on both routers involved before the link becomes active.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPF feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **router ospfv3 *instance-tag***
3. **area *area-id* virtual-link *router-id***
4. **show ipv6 ospfv3 virtual-link [brief]**
5. **copy running-config startup-config**

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router ospfv3 instance-tag Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	area area-id virtual-link router-id Example: switch(config-router)# area 10 virtual-link 2001:0DB8::1 switch(config-router-vlink)#	Creates one end of a virtual link to a remote router. You must create the virtual link on that remote router to complete the link.
Step 4	show ipv6 ospfv3 virtual-link [brief] Example: switch(config-if)# show ipv6 ospfv3 virtual-link	(Optional) Displays OSPFv3 virtual link information.
Step 5	copy running-config startup-config Example: switch(config-router)# copy running-config startup-config	(Optional) Saves this configuration change.

You can configure the following optional commands in virtual link configuration mode:

Command or Action	Purpose
dead-interval seconds Example: switch(config-router-vlink)# dead-interval 50	(Optional) Configures the OSPFv3 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
hello-interval seconds Example: switch(config-router-vlink)# hello-interval 25	(Optional) Configures the OSPFv3 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
retransmit-interval seconds Example: switch(config-router-vlink)# retransmit-interval 50	(Optional) Configures the OSPFv3 retransmit interval, in seconds. The range is from 1 to 65535. The default is 5.
transmit-delay seconds Example: switch(config-router-vlink)# transmit-delay 2	(Optional) Configures the OSPFv3 transmit-delay, in seconds. The range is from 1 to 450. The default is 1.

Send document comments to nexus7k-docfeedback@cisco.com.

The following example shows how to create a simple virtual link between two ABRs:

Configuration for ABR 1 (router ID 2001:0DB8::1) is as follows:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# area 10 virtual-link 2001:0DB8::10
switch(config-router)# copy running-config startup-config
```

Configuration for ABR 2 (router ID 2001:0DB8::10) is as follows:

```
switch# config t
switch(config)# router ospf 101
switch(config-router)# area 10 virtual-link 2001:0DB8::1
switch(config-router)# copy running-config startup-config
```

Configuring Redistribution

You can redistribute routes learned from other routing protocols into an OSPFv3 autonomous system through the ASBR.

You can configure the following optional parameters for route redistribution in OSPF:

- **Default information originate**—Generates an AS External (type 5) LSA for a default route to the external autonomous system.
- **Default metric**—Sets all redistributed routes to the same cost metric.

BEFORE YOU BEGIN

Create the necessary route maps used for redistribution.

Ensure that you have enabled the OSPF feature (see the [“Enabling the OSPFv3 Feature”](#) section on page 6-13).

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **router ospfv3 *instance-tag***
3. **address-family ipv6 unicast**
4. **redistribute {*bgp id* | *direct* | *isis id* | *rip id* | *static*} route-map *map-name***
5. **default-information originate [*always*] [*route-map map-name*]**
6. **default-metric *cost***
7. **copy running-config startup-config**

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router ospfv3 instance-tag Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	address-family ipv6 unicast Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	Enters IPv6 unicast address family mode.
Step 4	redistribute {bgp id direct isis id rip id static} route-map map-name Example: switch(config-router-af)# redistribute bgp route-map FilterExternalBGP	Redistributes the selected protocol into OSPFv3, through the configured route map.
Step 5	default-information originate [always] [route-map map-name] Example: switch(config-router-af)# default-information-originate route-map DefaultRouteFilter	Creates a default route into this OSPFv3 domain.
Step 6	default-metric cost Example: switch(config-router-af)# default-metric 25	Sets the cost metric for the redistributed routes. The range is from 1 to 16777214.
Step 7	copy running-config startup-config Example: switch(config-router)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to redistribute the Border Gateway Protocol (BGP) into OSPFv3:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# copy running-config startup-config
```

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring Route Summarization

You can configure route summarization for inter-area routes by configuring an address range that is summarized. You can also configure route summarization for external, redistributed routes by configuring a summary address for those routes on an ASBR. See the “[Route Summarization](#)” section on page 6-10.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPF feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

SUMMARY STEPS

1. `config t`
2. `router ospfv3 instance-tag`
3. `address-family ipv6 unicast`
4. `area area-id range ipv6-prefix/length [no-advertise]`
or
5. `summary-address ipv6-prefix/length [no-advertise] [tag tag]`
6. `show ipv6 ospfv3 summary-address`
7. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	<code>router ospfv3 instance-tag</code> Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	<code>address-family ipv6 unicast</code> Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	Enters IPv6 unicast address family mode.
Step 4	<code>area area-id range ipv6-prefix/length [no-advertise]</code> Example: switch(config-router-af)# area 10 range 2001:0DB8::/48 advertise	Creates a summary address on an ABR for a range of addresses. Optionally advertises this summary address in a Inter-Area Prefix (type 3) LSA.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 5	<pre>summary-address ipv6-prefix/length [no-advertise][tag tag] Example: switch(config-router-af)# summary-address 2001:0DB8::/48 tag 2</pre>	Creates a summary address on an ASBR for a range of addresses and optionally assigns a tag for this summary address that can be used for redistribution with route maps.
Step 6	<pre>show ipv6 ospfv3 summary-address Example: switch(config-router)# show ipv6 ospfv3 summary-address</pre>	(Optional) Displays information about OSPFv3 summary addresses.
Step 7	<pre>copy running-config startup-config Example: switch(config-router)# copy running-config startup-config</pre>	(Optional) Saves this configuration change.

The following example shows how to create summary addresses between areas on an ABR:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# address-family ipv6 unicast
switch(config-router)# area 10 range 2001:0DB8::/48
switch(config-router)# copy running-config startup-config
```

The following example shows how to create summary addresses on an ASBR:

```
switch# config t
switch(config)# router ospf 201
switch(config-router)# address-family ipv6 unicast
switch(config-router)# summary-address 2001:0DB8::/48
switch(config-router)# copy running-config startup-config
```

Modifying the Default Timers

OSPFv3 includes a number of timers that control the behavior of protocol messages and shortest path first (SPF) calculations. OSPFv3 includes the following optional timer parameters:

- LSA arrival time—Sets the minimum interval allowed between LSAs arriving from a neighbor. LSAs that arrive faster than this time are dropped.
- Pacing LSAs—Set the interval at which LSAs are collected into a group and refreshed, checksummed, or aged. This timer controls how frequently LSA updates occur and optimizes how many are sent in an LSA update message (see the [“Flooding and LSA Group Pacing”](#) section on page 6-6).
- Throttle LSAs—Set rate limits for generating LSAs. This timer controls how frequently an LSA is generated if no topology change occurs.
- Throttle SPF calculation—Controls how frequently the SPF calculation is run.

At the interface level, you can also control the following timers:

- Retransmit interval—Sets the estimated time between successive LSAs.
- Transmit delay—Sets the estimated time to transmit an LSA to a neighbor.

See the [“Configuring Networks in OSPFv3”](#) section on page 6-16 for information on the hello interval and dead timer.

Send document comments to nexus7k-docfeedback@cisco.com.

BEFORE YOU BEGIN

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **router ospfv3 *instance-tag***
3. **timers lsa-arrival *msec***
4. **timers lsa-group-pacing *seconds***
5. **timers throttle lsa *hold-interval***
6. **address-family ipv6 unicast**
7. **timers throttle spf *delay-time hold-time***
8. **interface *type slot/port***
9. **ospfv3 retransmit-interval *seconds***
10. **ospfv3 transmit-delay *seconds***
11. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router ospfv3 <i>instance-tag</i> Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	timers lsa-arrival <i>msec</i> Example: switch(config-router)# timers lsa-arrival 2000	Sets the LSA arrival time in milliseconds. The range is from 10 to 600000. The default is 1000 milliseconds.
Step 4	timers lsa-group-pacing <i>seconds</i> Example: switch(config-router)# timers lsa-group-pacing 2000	Sets the interval in seconds for grouping LSAs. The range is from 1 to 1800. The default is 240 seconds.
Step 5	timers throttle lsa <i>ihold-interval</i> Example: switch(config-router)# timers throttle lsa 3000	Sets the rate limit in milliseconds for generating LSAs. The range is from 10 to 600000. The default is 5000 milliseconds.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 6	address-family ipv6 unicast Example: switch(config-router)# address-family ipv6 unicast switch(config-router-af)#	Enters IPv6 unicast address family mode.
Step 7	timers throttle spf delay-time hold-time Example: switch(config-router)# timers throttle spf 3000 2000	Sets the SPF best path schedule initial delay time and the minimum hold time in seconds between SPF best path calculations. The range is from 1 to 600000. The default is no delay time and 5000 millisecond hold time.
Step 8	interface type slot/port Example: switch(config)# interface ethernet 1/2 switch(config-if)#	Enters interface configuration mode.
Step 9	ospfv3 retransmit-interval seconds Example: switch(config-if)# ospfv3 retransmit-interval 30	Sets the estimated time in seconds between LSAs transmitted from this interface. The range is from 1 to 65535. The default is 5.
Step 10	ospfv3 transmit-delay seconds Example: switch(config-if)# ospfv3 transmit-delay 600 switch(config-if)#	Sets the estimated time in seconds to transmit an LSA to a neighbor. The range is from 1 to 450. The default is 1.
Step 11	copy running-config startup-config Example: switch(config-if)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to control LSA flooding with the `lsa-group-pacing` option:

```
switch# config t
switch(config)# router ospf 201
switch(config-router)# timers lsa-group-pacing 300
switch(config-router)# copy running-config startup-config
```

Configuring Graceful Restart

Graceful restart is enabled by default. You can configure the following optional parameters for graceful restart in an OSPFv3 instance:

- Grace period—Configures how long neighbors should wait after a graceful restart has started before tearing down adjacencies.
- Helper mode disabled—Disables helper mode on the local OSPFv3 instance. OSPFv3 will not participate in the graceful restart of a neighbor.
- Planned graceful restart only—Configures OSPFv3 to support graceful restart only in the event of a planned restart.

Send document comments to nexus7k-docfeedback@cisco.com.

BEFORE YOU BEGIN

Ensure that you have enabled the OSPFv3 feature (see the “[Enabling the OSPFv3 Feature](#)” section on page 6-13).

Ensure that all neighbors are configured for graceful restart with matching optional parameters set.

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

SUMMARY STEPS

1. `config t`
2. `router ospfv3 instance-tag`
3. `graceful-restart`
4. `graceful-restart grace-period seconds`
5. `graceful-restart helper-disable`
6. `graceful-restart planned-only`
7. `show ipv6 ospfv3 instance-tag`
8. `copy running-config startup-config`

DETAILED STEPS

	Command	Purpose
Step 1	<code>config t</code> Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	<code>router ospfv3 instance-tag</code> Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 3	<code>graceful-restart</code> Example: switch(config-router)# graceful-restart	Enables graceful restart. A graceful restart is enabled by default.
Step 4	<code>graceful-restart grace-period seconds</code> Example: switch(config-router)# graceful-restart grace-period 120	Sets the grace period, in seconds. The range is from 5 to 1800. The default is 60 seconds.
Step 5	<code>graceful-restart helper-disable</code> Example: switch(config-router)# graceful-restart helper-disable	Disables helper mode. Enabled by default.
Step 6	<code>graceful-restart planned-only</code> Example: switch(config-router)# graceful-restart planned-only	Configures graceful restart for planned restarts only.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 7	<code>show ipv6 ospfv3 instance-tag</code> Example: <code>switch(config-if)# show ipv6 ospfv3 201</code>	(Optional) Displays OSPFv3 information.
Step 8	<code>copy running-config startup-config</code> Example: <code>switch(config)# copy running-config startup-config</code>	(Optional) Saves this configuration change.

The following example shows how to enable graceful restart if it has been disabled, and set the grace period to 120 seconds:

```
switch# config t
switch(config)# router ospfv3 201
switch(config-router)# graceful-restart
switch(config-router)# graceful-restart grace-period 120
switch(config-router)# copy running-config startup-config
```

Restarting an OSPFv3 Instance

You can restart an OSPFv3 instance. This clears all neighbors for the instance.

To restart an OSPFv3 instance and remove all associated neighbors, use the following command:

Command	Purpose
<code>restart ospfv3 instance-tag</code> Example: <code>switch(config)# restart ospfv3 201</code>	Restarts the OSPFv3 instance and removes all neighbors.

Configuring OSPFv3 with Virtualization

You can configure multiple OSPFv3 instances in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple OSPFv3 instances in each VRF. You assign an OSPFv3 interface to a VRF.



Note

Configure all other parameters for an interface after you configure the VRF for an interface. Configuring a VRF for an interface deletes all the configuration for that interface.

BEFORE YOU BEGIN

Create the VDCs.

Ensure that you have enabled the OSPF feature (see the [“Enabling the OSPFv3 Feature”](#) section on page 6-13).

Ensure that you are in the correct VDC (or use the `switchto vdc` command).

Send document comments to nexus7k-docfeedback@cisco.com.

SUMMARY STEPS

1. **config t**
2. **vrf context** *vrf_name*
3. **router ospfv3** *instance-tag*
4. **vrf** *vrf-name*
5. <optional parameters configured>
6. **interface** *type slot/port*
7. **vrf member** *vrf-name*
8. **ipv6 address** *ipv6-prefix/length*
9. **ipv6 ospfv3** *instance-tag area area-id*
10. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	vrf context <i>vrf-name</i> Example: switch(config)# vrf context RemoteOfficeVRF switch(config-vrf)#	Creates a new VRF and enters VRF configuration mode.
Step 3	router ospfv3 <i>instance-tag</i> Example: switch(config)# router ospfv3 201 switch(config-router)#	Creates a new OSPFv3 instance with the configured instance tag.
Step 4	vrf <i>vrf-name</i> Example: switch(config-router)# vrf RemoteOfficeVRF switch(config-router-vrf)#	Enters VRF configuration mode.
Step 5	maximum-paths <i>paths</i> Example: switch(config-router-vrf)# maximum-paths 4	(Optional) Configures the maximum number of equal OSPFv3 paths to a destination in the route table for this VRF. Used for load balancing.
Step 6	interface <i>type slot/port</i> Example: switch(config)# interface ethernet 1/2 switch(config-if)#	Enters interface configuration mode.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 7	vrf member <i>vrf-name</i> Example: switch(config-if)# vrf member RemoteOfficeVRF	Adds this interface to a VRF.
Step 8	ipv6 address <i>ipv6-prefix/length</i> Example: switch(config-if)# ipv6 address 2001:0DB8::1/48	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 9	ipv6 ospfv3 <i>instance-tag area area-id</i> Example: switch(config-if)# ipv6 ospfv3 201 area 0	Assigns this interface to the OSPFv3 instance and area configured.
Step 10	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to create a VRF and add an interface to the VRF:

```
switch# config t
switch(config)# vrf context NewVRF
switch(config-vrf)# exit
switch(config)# router ospfv3 201
switch(config-router)# exit
switch(config)# interface ethernet 1/2
switch(config-if)# vrf member NewVRF
switch(config-if)# ipv6 address 2001:0DB8::1/48
switch(config-if)# ipv6 ospfv3 201 area 0
switch(config-if)# copy running-config startup-config
```

Verifying OSPFv3 Configuration

To verify OSPFv3 configuration, use one of the following commands:

Command	Purpose
show ipv6 ospfv3	Displays the OSPFv3 configuration.
show ipv6 ospfv3 border-routers	Displays the internal OSPF routing table entries to an ABR and ASBR.
show ipv6 ospfv3 database	Displays lists of information related to the OSPFv3 database for a specific router.
show ipv6 ospfv3 interface <i>type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv3 interface configuration.
show ipv6 ospfv3 neighbors	Displays the neighbor information. Use the clear ospfv3 neighbors command to remove adjacency with all neighbors.
show ipv6 ospfv3 request-list	Displays a list of LSAs requested by a router.

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
<code>show ipv6 ospfv3 retransmission-list</code>	Displays a list of LSAs waiting to be retransmitted.
<code>show ipv6 ospfv3 summary-address</code>	Displays a list of all summary address redistribution information configured under an OSPFv3 instance.
<code>show running-configuration ospfv3</code>	Displays the current running OSPFv3 configuration.

Displaying OSPFv3 Statistics

To display OSPFv3 statistics, use the following commands:

Command	Purpose
<code>show ipv6 ospfv3 memory</code>	Displays the OSPFv3 memory usage statistics.
<code>show ipv6 ospfv3 policy statistics area <i>area-id</i> filter-list {in out} [vrf {<i>vrf-name</i> all default management}]</code>	Displays the OSPFv3 route policy statistics for an area.
<code>show ipv6 ospfv3 policy statistics redistribute {<i>bgp id</i> <i>direct</i> <i>isis id</i> <i>rip id</i> <i>static</i>} vrf {<i>vrf-name</i> all default management}</code>	Displays the OSPFv3 route policy statistics.
<code>show ipv6 ospfv3 statistics [vrf {<i>vrf-name</i> all default management}]</code>	Displays the OSPFv3 event counters.
<code>show ipv6 ospfv3 traffic [<i>interface-type number</i>] [vrf {<i>vrf-name</i> all default management}]</code>	Displays the OSPFv3 packet counters.

OSPFv3 Example Configuration

The following example shows how to configure OSPFv3:

```
feature ospfv3
router ospfv3 201
  router-id 290.0.2.1

interface ethernet 1/2
  ipv6 address 2001:0DB8::1/48
  ipv6 ospfv3 201 area 10
```

Related Topics

The following topics can give more information on OSPF:

- [Chapter 5, “Configuring OSPFv2”](#)

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

- Chapter 15, “Configuring Route Policy Manager”

Default Settings

Table 6-2 lists the default settings for OSPFv3 parameters.

Table 6-2 *Default OSPFv3 Parameters*

Parameters	Default
Hello interval	10 seconds
Dead interval	40 seconds
Graceful restart grace period	60 seconds
Graceful restart notify period	15 seconds
OSPFv3 feature	Disabled
Stub router advertisement announce time	600 seconds
Reference bandwidth for link cost calculation	40 Gb/s
LSA minimal arrival time	1000 milliseconds
LSA group pacing	240 seconds
SPF calculation initial delay time	0 milliseconds
SPF calculation hold time	5000 milliseconds
SPF calculation initial delay time	0 milliseconds

Additional References

For additional information related to implementing OSPF, see the following sections:

- [Related Documents, page 6-39](#)
- [MIBs, page 6-39](#)

Send document comments to nexus7k-docfeedback@cisco.com.

Related Documents

Related Topic	Document Title
OSPFv3 CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0</i>
VDCs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0</i>

MIBs

MIBs	MIBs Link
<ul style="list-style-type: none">• OSPF-MIB• OSPF-TRAP-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Send document comments to nexus7k-docfeedback@cisco.com.