



Configuring OSPFv2

This chapter contains the following sections:

- [Finding Feature Information](#), on page 1
- [Information About OSPFv2](#), on page 1
- [Prerequisites for OSPFv2](#), on page 12
- [Guidelines and Limitations for OSPFv2](#), on page 12
- [Default Settings for OSPFv2](#), on page 14
- [Configuring Basic OSPFv2](#), on page 15
- [Configuring Advanced OSPFv2](#), on page 24
- [Verifying the OSPFv2 Configuration](#), on page 41
- [Monitoring OSPFv2](#), on page 42
- [Configuration Examples for OSPFv2](#), on page 42
- [Related Documents for OSPFv2](#), on page 43
- [Feature History for OSPFv2](#), on page 43

Finding Feature Information

Your software release might not support all the features documented in this module. For the latest caveats and feature information, see the Bug Search Tool at <https://tools.cisco.com/bugsearch/> and the release notes for your software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the "New and Changed Information" chapter or the Feature History table in this chapter.

Information About OSPFv2

OSPFv2 is an IETF link-state protocol for IPv4 networks. An OSPFv2 router sends a special message, called a hello packet, out each OSPF-enabled interface to discover other OSPFv2 neighbor routers. Once a neighbor is discovered, the two routers compare information in the to determine if the routers have compatible configurations. The neighbor routers try to establish , which means that the routers synchronize their link-state databases to ensure that they have identical OSPFv2 routing information. Adjacent routers share (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 OSPFv2 routers eventually have identical link-state databases. When all OSPFv2 routers have identical link-state

databases, the network is converged. Each router then uses Dijkstra's Shortest Path First (SPF) algorithm to build its route table.

You can divide OSPFv2 networks into areas. Routers send most LSAs only within one area, which reduces the CPU and memory requirements for an OSPF-enabled router.

OSPFv2 supports IPv4, while OSPFv3 supports IPv6. For more information, see the "Configuring OSPFv3" chapter.



Note OSPFv2 on Cisco NX-OS supports RFC 2328. This RFC introduced a different method to calculate route summary costs which is not compatible with the calculation used by RFC1583. RFC 2328 also introduced different selection criteria for AS-external paths. It is important to ensure that all routers support the same RFC. Use the **rfc1583compatibility** command if your network includes routers that are only compliant with RFC1583. The default supported RFC standard for OSPFv2 may be different for Cisco NX-OS and Cisco IOS. You must make adjustments to set the values identically. For more information, see the "OSPF RFC Compatibility Mode Example" section.

Hello Packet

OSPFv2 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. OSPFv2 uses Hello packets for the following tasks:

- Neighbor discovery
- Keepalives
- Bidirectional communications
- Designated router election

The Hello packet contains information about the originating OSPFv2 interface and router, including the assigned OSPFv2 cost of the link, the , and optional capabilities of the originating router. An OSPFv2 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.

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.

OSPFv2 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 (usually a multiple of the hello interval), then the neighbor is removed from the local neighbor table.

Neighbors

An OSPFv2 interface must have a compatible configuration with a remote interface before the two can be considered neighbors. The two OSPFv2 interfaces must match the following criteria:

- Hello interval
- Dead interval

- Area ID
- Authentication
- Optional capabilities

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

- Neighbor ID—The router ID of the neighbor.
- Priority—Priority of the neighbor. The priority is used for designated router election.
- 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 the time since the last Hello packet was received from this neighbor.
- IP Address—The IP address of the neighbor.
- Designated Router—Indication of whether the neighbor has been declared as the designated router or as the backup designated router.
- Local interface—The local interface that received the Hello packet for this neighbor.

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.

Adjacency is established using Database Description packets, Link State Request packets, and Link State Update packets in OSPF. The Database Description packet includes just the LSA headers from the link-state database of the neighbor. The local router compares these headers with its own link-state database and determines which LSAs are new or 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 OSPF. If every router floods the network with LSAs, the same link-state information is sent from multiple sources. Depending on the type of network, OSPFv2 might use a single router, the (DR), to control the LSA floods and represent the network to the rest of the OSPFv2 area. If the DR fails, OSPFv2 selects a (BDR). If the DR fails, OSPFv2 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. OSPFv2 routers establish a DR and BDR that controls LSA flooding on the network. OSPFv2 uses the well-known IPv4 multicast addresses 224.0.0.5 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, OSPFv2 chooses the highest router IDs as the DR and BDR.

All other routers establish adjacency with the DR and the BDR and use the IPv4 multicast address 224.0.0.6 to send LSA updates to the DR and BDR. Figure 3-1 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 1: DR in Multi-Access Network




Areas

You can limit the CPU and memory requirements that OSPFv2 puts on the routers by dividing an OSPFv2 network into . An area is a logical division of routers and links within an OSPFv2 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 you can enter as a number or in dotted decimal notation, such as 10.2.3.1.

Cisco NX-OS always displays the area in dotted decimal notation.

If you define more than one area in an OSPFv2 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 (ABRs). An ABR connects to both the backbone area and at least one other defined area.

Figure 2: OSPFv2 Areas



The ABR has a separate link-state database for each area to which it connects. The ABR sends Network Summary (type 3) from one connected area to the backbone area. The backbone area sends summarized information about one area to another area. In the OSPFv2 Areas Figure, Area 0 sends summarized information about Area 5 to Area 3.

OSPFv2 defines one other router type: the autonomous system boundary router (ASBR). This router connects an OSPFv2 area to another autonomous system. An autonomous system is a network controlled by a single technical administration entity. OSPFv2 can redistribute its routing information into another autonomous system or receive redistributed routes from another autonomous system.

Link-State Advertisements

Link-State Advertisements Types

OSPFv2 uses link-state advertisements (LSAs) to build its routing table.

Names	Description
Router LSA	LSA sent by every router. This LSA includes the state and the cost of all links and a list of all OSPFv2 neighbors on the link. Router LSAs trigger an SPF recalculation. Router LSAs are flooded to local OSPFv2 area.
Network LSA	LSA sent by the DR. This LSA lists all routers in the multi-access network. Network LSAs trigger an SPF recalculation.
Network Summary LSA	LSA sent by the area border router to an external area for each destination in the local area. This LSA includes the link cost from the area border router to the local destination.
ASBR Summary LSA	LSA sent by the area border router to an external area. This LSA advertises the link cost to the ASBR only.

Names	Description
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.
NSSA External LSA	LSA generated by the ASBR within a not-so-stubby area (NSSA). This LSA includes the link cost to an external autonomous system destination. NSSA External LSAs are flooded only within the local NSSA.
Opaque LSAs	LSA used to extend OSPF.

Link Cost

Each OSPFv2 interface is assigned a . 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

When an OSPFv2 router receives an LSA, it forwards that LSA out every OSPF-enabled interface, flooding the OSPFv2 area with this information. This LSA flooding guarantees that all routers in the network have identical routing information. LSA flooding depends on the OSPFv2 area configuration. The LSAs are flooded based on the (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 OSPFv2 to pack multiple LSAs into an OSPFv2 Update message.

By default, LSAs with link-state refresh times within 10 seconds 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 OSPFv2 load on your network.

Link-State Database

Each router maintains a link-state database for the OSPFv2 network. This database contains all the collected LSAs, and includes information on all the routes through the network. OSPFv2 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.

Opaque LSAs

Opaque LSAs allow you to extend OSPF functionality. Opaque LSAs consist of a standard LSA header followed by application-specific information. This information might be used by OSPFv2 or by other applications. OSPFv2 uses Opaque LSAs to support OSPFv2 Graceful Restart capability. Three Opaque LSA types are defined as follows:

- LSA type 9—Flooded to the local network.

- LSA type 10—Flooded to the local area.
- LSA type 11—Flooded to the local autonomous system.

OSPFv2 and the Unicast RIB

OSPFv2 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 OSPFv2 route table. When the OSPFv2 network is converged, this route table feeds into the unicast RIB. OSPFv2 communicates with the unicast RIB to do the following:

- Add or remove routes
- Handle route redistribution from other protocols
- Provide convergence updates to remove stale OSPFv2 routes and for stub router advertisements

OSPFv2 also runs a modified Dijkstra algorithm for fast recalculation for summary and external (type 3, 4, 5, and 7) LSA changes.

Authentication

You can configure authentication on OSPFv2 messages to prevent unauthorized or invalid routing updates in your network. Cisco NX-OS supports two authentication methods:

- Simple password authentication
- MD5 authentication digest

You can configure the OSPFv2 authentication for an OSPFv2 area or per interface.

Simple Password Authentication

Simple password authentication uses a simple clear-text password that is sent as part of the OSPFv2 message. The receiving OSPFv2 router must be configured with the same clear-text password to accept the OSPFv2 message as a valid route update. Because the password is in clear text, anyone who can watch traffic on the network can learn the password.

MD5 Authentication

You should use MD5 authentication to authenticate OSPFv2 messages. You configure a password that is shared at the local router and all remote OSPFv2 neighbors. For each OSPFv2 message, Cisco NX-OS creates an MD5 one-way message digest based on the message itself and the encrypted password. The interface sends this digest with the OSPFv2 message. The receiving OSPFv2 neighbor validates the digest using the same encrypted password. If the message has not changed, the digest calculation is identical and the OSPFv2 message is considered valid.

MD5 authentication includes a sequence number with each OSPFv2 message to ensure that no message is replayed in the network.

Advanced Features for OSPFv2

Cisco NX-OS supports advanced OSPFv2 features that enhance the usability and scalability of OSPFv2 in the network.


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. These LSAs are usually flooded throughout the local autonomous system to propagate external route information. Stub areas have the following requirements:

- All routers in the stub area are stub routers.
- No ASBR routers exist in the stub area.
- You cannot configure virtual links in the stub area.

The following figure shows an example of an OSPFv2 autonomous system where all routers in area 0.0.0.10 have to go through the ABR to reach external autonomous systems. Area 0.0.0.10 can be configured as a stub area.

Figure 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 0.0.0.0 for IPv4.

Not-So-Stubby Area

A Not-so-Stubby Area (NSSA) is similar to a 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 NSSA External (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 NSSA External LSA to AS External (type 5) LSAs. The ABR then floods these AS External LSAs throughout the OSPFv2 autonomous system. Summarization and filtering are supported during the translation.

You can, for example, use NSSA to simplify administration if you are connecting a central site using OSPFv2 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 OSPFv2 stub area because routes for the remote site could not be redistributed into a stub area. With NSSA, you can extend OSPFv2 to cover the remote connection by defining the area between the corporate router and remote router as an NSSA.

The backbone Area 0 cannot be an NSSA.



Note OSPF became compliant with RFC 3101 section 2.5(3). When an Area Border Router attached to a Not-so-Stubby Area receives a default route LSA with P-bit clear, it should be ignored. OSPF had been previously adding the default route under these conditions.


If you have already designed your networks with RFC non-compliant behavior and expect a default route to be added on NSSA ABR, you will see a change in behavior when you upgrade.

If you decide to continue with the old behavior, you have the option to enable it with the **default-route nssa-abr pbit-clear** command.

Virtual Links

Virtual links allow you to connect an OSPFv2 area ABR to a backbone area ABR when a direct physical connection is not available. The figure shows a virtual link that connects Area 3 to the backbone area through Area 5.

Figure 4: Virtual Links



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

OSPFv2 can learn routes from other routing protocols by using route redistribution. You configure OSPFv2 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. You must configure a route map with the redistribution to control which routes are passed into OSPFv2. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. You can use route maps to modify parameters in the AS External (type 5) and NSSA External (type 7) LSAs before these external routes are advertised in the local OSPFv2 autonomous system.

OSPFv2 sets the type-5 LSA's forwarding address as described below:

- If the next-hop for the route is an attached-route then the forwarding address is the next-hop address for that route.
- If the next-hop for the route is a recursive route and next-hop's next-hop is an attached route then the forwarding address is the next-hop's next-hop address.

Route Summarization

Because OSPFv2 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 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one summary address, 10.1.0.0/16.

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.

External route summarization is specific to external routes that are injected into OSPFv2 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 provides a multilevel high-availability architecture. OSPFv2 supports stateful restart, which is also referred to as non-stop routing (NSR). If OSPFv2 experiences problems, it attempts to restart from its previous run-time state. The neighbors do not register any neighbor event in this case. If the first restart is not successful and another problem occurs, OSPFv2 attempts a graceful restart.

A graceful restart, or nonstop forwarding (NSF), allows OSPFv2 to remain in the data forwarding path through a process restart. When OSPFv2 needs to perform a graceful restart, it sends a link-local opaque (type 9) LSA, called a grace LSA. This restarting OSPFv2 platform is called NSF capable.

The grace LSA includes a grace period, which is a specified time that the neighbor OSPFv2 interfaces hold onto the LSAs from the restarting OSPFv2 interface. (Typically, OSPFv2 tears down the adjacency and discards all LSAs from a down or restarting OSPFv2 interface.) The participating neighbors, which are called NSF helpers, keep all LSAs that originate from the restarting OSPFv2 interface as if the interface was still adjacent.

When the restarting OSPFv2 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 the graceful restart has finished.

Stateful restart is used in the following scenarios:

- First recovery attempt after the process experiences problems
- ISSU
- User-initiated switchover using the **system switchover** command
- Active supervisor reload using the **reload module active-sup** command

Graceful restart is used in the following scenarios:

- Second recovery attempt after the process experiences problems within a 4-minute interval
- Manual restart of the process using the **restart ospf** command
- Active supervisor removal



Note The Cisco Nexus 7000 series devices support the Internet Engineering Task Force (IETF) version only. As a result, NSF IETF must be explicitly configured under the routing protocols in the Virtual Switching System (VSS). Use the **nsf ietf** command in router configuration mode for NSF IETF configuration. No additional configuration is required on the Cisco Nexus 7000 pairs because they run NSF IETF graceful-restart by default. However, each neighbor device that will become Layer 3 adjacent must have NSF configured and the same mode of NSF must be enabled to successfully operate a graceful failover.

OSPFv2 Stub Router Advertisements

You can configure an OSPFv2 interface to act as a stub router using the OSPFv2 Stub Router Advertisements feature. Use this feature when you want to limit the OSPFv2 traffic through this router, such as when you want to introduce a new router to the network in a controlled manner or limit the load on a router that is already overloaded. You might also want to use this feature for various administrative or traffic engineering reasons.

OSPFv2 stub router advertisements do not remove the OSPFv2 router from the network topology, but they do prevent other OSPFv2 routers from using this router to route traffic to other parts of the network. Only the traffic that is destined for this router or directly connected to this router is sent.

OSPFv2 stub router advertisements mark all stub links (directly connected to the local router) to the cost of the local OSPFv2 interface. All remote links are marked with the maximum cost (0xFFFF).

Multiple OSPFv2 Instances

Cisco NX-OS supports multiple instances of the OSPFv2 protocol that run on the same node. You cannot configure multiple instances over the same interface. 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 OSPFv2 autonomous system.

SPF Optimization

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

- Partial SPF for Network (type 2) LSAs, Network Summary (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.

BFD

This feature supports bidirectional forwarding detection (BFD). BFD is a detection protocol that provides fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules. See the *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide* for more information.

Virtualization Support for OSPFv2

OSPFv2 supports virtual routing and forwarding (VRF) instances. 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.

Cisco NX-OS Release 6.1 or later supports more than four process instances for OSPFv2 per VDC. However, only the first four configured OSPFv2 instances are supported with MPLS LDP and MPLS TE. Each OSPFv2 instance can support multiple VRFs, up to the system limit. For more information, see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* and the *Cisco Nexus 7000 Series NX-OS Verified Scalability Guide*.

Prerequisites for OSPFv2

OSPFv2 has the following prerequisites:

- You must be familiar with routing fundamentals to configure OSPF.
- You are logged on to the switch.
- You have configured at least one interface for IPv4 that can communicate with a remote OSPFv2 neighbor.
- You have completed the OSPFv2 network strategy and planning for your network. For example, you must decide whether multiple areas are required.
- You have enabled the OSPF feature.
- You have installed the appropriate license and entered the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide* for configuration information and the *Cisco NX-OS Licensing Guide* for licensing information) if you are configuring VDCs.

Guidelines and Limitations for OSPFv2

OSPFv2 has the following configuration guidelines and limitations:

- CE devices install type 3 LSAs with DN-bit or Type 5 LSAs with DN-bit and VPN Route TAG in the RIB (non-default VRF). This behaviour is applicable prior to Cisco NX-OS Release 8.3(2).

- The default-information originate command must be configured so that the MPLS default route is advertised to the CE-VRF. When using default-information originate command, the DN-bit in type 3 5 LSAs options and Route TAGs in Type 5 LSAs are not set for the default route only.
- The Cisco Nexus 7000 supports the Internet Engineering Task Force (IETF) version only. As a result, NSF IETF must be explicitly configured under the routing protocols in the Virtual Switching System (VSS). No additional configuration is required on the Cisco Nexus 7000 pairs because they run NSF IETF graceful-restart by default. However, each neighbor device that will become Layer 3 adjacent must have NSF configured and the same mode of NSF must be enabled to successfully operate a graceful failover.
- Cisco NX-OS displays areas in dotted decimal notation regardless of whether you enter the area in decimal or dotted decimal notation.
- All OSPFv2 routers must operate in the same RFC compatibility mode. OSPFv2 for Cisco NX-OS complies with RFC 2328. Use the **rfc1583compatibility** command in router configuration mode if your network includes routers that support only RFC 1583.
- In scaled scenarios, when the number of interfaces and link-state advertisements in an OSPF process is large, the snmp-walk on OSPF MIB objects is expected to time out with a small-values timeout at the SNMP agent. If you observe a timeout on the querying SNMP agent while polling OSPF MIB objects, increase the timeout value on the polling SNMP agent.
- MTU configured at interface level works in either the data plane or in the control plane but not at both planes at the same time.

When you configure MTU with a size lower than the supported size in data and control planes a few features that have minimum MTU requirements may not work in both the planes.

For example, MPLS VPN is supported in the data plane since this plane supports the MTU of 1500 bytes that the MPLS VPN requires. But control plane does not support MPLS VPN because this plane cannot handle the 1500-byte packets.

To make the configured MTU work in control plane for MPLS VPN, you need to manually configure the OSPF packet size (by using the **packet-size** *size* command) so that OSPF works on the control plane. This is applicable from Cisco NX-OS Release 8.3(2) onwards.

The **packet-size** *size* command is supported on the Ethernet, SVI, and GRE tunnel interfaces.

- Cisco NX-OS Release 6.1 or later supports more than four process instances for OSPFv2 per VDC. However, only the first four configured OSPFv2 instances are supported with MPLS LDP and MPLS TE.
- The **default-information-originate always** command advertises the OSPF default route from Cisco NX-OS Release 7.3(5)D1(1) and later releases and from Cisco NX-OS Release 8.0(1) and later releases in 8.x release train.
- The following guidelines and limitations apply to the administrative distance feature, which is supported beginning with Cisco NX-OS Release 6.1:
 - When an OSPF route has two or more equal cost paths, configuring the administrative distance is non-deterministic for the **match ip route-source** command.
 - Configuring the administrative distance is supported only for the **match route-type**, **match ip address prefix-list**, and **match ip route-source prefix-list** commands. The other match statements are ignored.

- There is no preference among the **match route-type**, **match ip address**, and **match ip route-source** commands for setting the administrative distance of OSPF routes. In this way, the behavior of the table map for setting the administrative distance in Cisco NX-OS OSPF is different from that in Cisco IOS OSPF.
- The discard route is always assigned an administrative distance of 220. No configuration in the table map applies to OSPF discard routes.
- In Cisco NX-OS Release 6.2(6a) and later releases, you can filter next-hop paths for an OSPF route to prevent the path from being added to the RIB. Before Cisco NX-OS Release 6.2(6a), filtering on a specific path was ignored and the entire route was not added to the RIB.



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.

Default Settings for OSPFv2

Table 1: Default OSPFv2 Parameters

Parameters	Default
Administrative distance	110
Hello interval	10 seconds
Dead interval	40 seconds
Discard routes	Enabled
Graceful restart grace period	60 seconds
OSPFv2 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	10 seconds
SPF calculation initial delay time	200 milliseconds
SPF minimum hold time	5000 milliseconds
SPF calculation initial delay time	1000 milliseconds

Configuring Basic OSPFv2

Enabling OSPFv2

You must enable the OSPFv2 feature before you can configure OSPFv2.

Before you begin

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] feature ospf	Enables the OSPFv2 feature. Note Use the no form of this command to disable the OSPFv2 feature and remove all associated configuration.
Step 3	(Optional) switch(config)# show feature	Displays enabled and disabled features.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Creating an OSPFv2 Instance

The first step in configuring OSPFv2 is to create an OSPFv2 instance. You assign a unique instance tag for this OSPFv2 instance. The instance tag can be any string.



Note The OSPF router ID changes without a restart on a Cisco Nexus 7000 switch when you have not configured a manual router ID in the following cases:

- Configuring an SVI or physical interface with a higher IP address than the current router ID on a setup without any configured loopback interfaces.
- Configuring a loopback interface with any given IP address on a setup without any previously configured loopback interfaces.
- Configuring a loopback interface with a higher IP address than the IP address of an existing configured loopback interface.

When a router ID changes, OSPF has to re-advertise all LSAs with the new router ID. To avoid this issue, you can configure a manual OSPF router ID.

Before you begin

Ensure that you have enabled the OSPF feature.

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

OSPFv2 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).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag. Note Use the no form of this command in global configuration mode to remove the OSPFv2 instance and all associated configurations. Using the no form of this command in the interface configuration mode does not remove the OSPF configuration. You must manually remove any OSPFv2 commands configured in interface mode.
Step 3	(Optional) switch(config-router)# router-id ip-address	Configures the OSPFv2 router ID. This IP address identifies this OSPFv2 instance and must exist on a configured interface in the system. This command restarts the OSPF process automatically and changes the router id after it is configured.
Step 4	(Optional) switch(config-router)# show ip ospf instance-tag	Displays OSPF information.
Step 5	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring OSPF Packet Size

MTU configured at interface level works in either the data plane or in the control plane but not at both planes at the same time.

When you configure MTU with a size lower than the supported size in data and control planes a few features that have minimum MTU requirements may not work in both the planes.

For example, MPLS VPN is supported in the data plane since this plane supports the MTU of 1500 bytes that the MPLS VPN requires. But control plane does not support MPLS VPN because this plane cannot handle the 1500-byte packets.

To make the configured MTU work in control plane for MPLS VPN, you need to manually configure the OSPF packet size so that OSPF works on the control plane. This is applicable from Cisco NX-OS Release 8.3(2) onwards.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# [no] router ospf <i>instance-tag</i>	Creates a new OSPF instance with the configured instance tag. Note Use the no form of this command in global configuration mode to remove the OSPFv2 instance and all associated configurations. Using the no form of this command in the interface configuration mode does not remove the OSPF configuration. You must manually remove any OSPFv2 commands configured in interface mode.
Step 3	switch(config-router)# router-id <i>ip-address</i>	Configures the OSPFv2 router ID. This IP address identifies this OSPFv2 instance and must exist on a configured interface in the system. This command restarts the OSPF process automatically and changes the router id after it is configured.
Step 4	switch(config-router)# packet-size <i>size</i>	<ul style="list-style-type: none"> • Configures the OSPFv2 packet size. The size range is from 572 to 9212 bytes. • You can configure the packet-size in the interface configuration mode also. • You can configure the packet-size <i>size</i> command even if their ospf mtu-ignore command is already configured in the network.
Step 5	(Optional) switch(config-router)# show ip ospf interface <i>interface-number</i>	Displays OSPF information.

Example

This example shows how to configure the OSPF packet-size:

```
router ospf 1
  router-id 3.3.3.3
  [no] packet-size 2000
```

This example shows the display of the OSPF packet-size:

```
Switch (config-router)# show ip ospf interface ethernet 1/25
Ethernet1/25 is up, line protocol is up
  IP address 1.0.0.1/24
  ----- snip -----
  Number of opaque link LSAs: 0, checksum sum 0
  Max Packet Size: 2000
```

Configuring Optional Parameters on an OSPFv2 Instance

You can configure optional parameters for OSPF. The following commands are available in the router configuration mode.

For more information about OSPFv2 instance parameters, see the “Configuring Advanced OSPFv2” section

Before you begin

Ensure that you have enabled the OSPF feature.

OSPFv2 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).

Procedure

	Command or Action	Purpose
Step 1	switch(config-router)# distance <i>number</i>	Configures the administrative distance for this OSPFv2 instance. The range is from 1 to 255. The default is 110.
Step 2	switch(config-router)# log-adjacency-changes [detail]	Generates a system message whenever a neighbor changes state.
Step 3	switch(config-router)# maximum-paths <i>path-number</i>	Configures the maximum number of equal OSPFv2 paths to a destination in the route table. This command is used for load balancing. The range is from 1 to 16. The default is 8.
Step 4	switch(config-router)# [no]name-lookup <i>path-number</i>	Enables the translation of OSPF router IDs to host names, either by looking up the local hosts database or querying DNS names in IPv6. This command makes it easier to identify a device because it displays the device by name rather than by its router ID or neighbor ID.

	Command or Action	Purpose
		Note To stop displaying OSPF router IDs as DNS names, use the no form of this command.
Step 5	switch(config-router)# passive-interface default	Suppresses routing updates on all interfaces. This command is overridden by the VRF or interface command mode configuration.

Example

This example shows how to create an OSPFv2 instance:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# copy running-config startup-config
```

Configuring Networks in OSPFv2

You can configure a network to OSPFv2 by associating it through the interface that the router uses to connect to that network. 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.



Note OSPF is not enabled on an interface until you configure a valid IP address for that interface.

Before you begin

Ensure that you have enabled the OSPF feature

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ip address <i>ip-prefix/length</i>	Assigns an IP address and subnet mask to this interface.
Step 4	switch(config-if)# ip router ospf <i>instance-tag area area-id</i> [secondaries none]	Adds the interface to the OSPFv2 instance and area.

	Command or Action	Purpose
Step 5	(Optional) switch(config-if)# show ip ospf instance-tag interface interface-type slot/port	Displays OSPF information.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.
Step 7	(Optional) switch(config)# ip ospf cost number	Configures the OSPFv2 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.
Step 8	(Optional) switch(config)# ip ospf dead-interval seconds	Configures the OSPFv2 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
Step 9	(Optional) switch(config)# ip ospf hello-interval seconds	Configures the OSPFv2 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
Step 10	(Optional) switch(config)# ip ospf mtu-ignore	Configures OSPFv2 to ignore any IP MTU mismatch with a neighbor. The default is to not establish adjacency if the neighbor MTU does not match the local interface MTU.
Step 11	(Optional) switch(config)# [default no] ip ospf passive-interface	Suppresses routing updates on the interface. This command overrides the router or VRF command mode configuration. The default option removes this interface mode command and reverts to the router or VRF configuration, if present.
Step 12	(Optional) switch(config)# ip ospf priority number	Configures the OSPFv2 priority, used to determine the DR for an area. The range is from 0 to 255. The default is 1.
Step 13	(Optional) switch(config)# ip ospf shutdown	Shuts down the OSPFv2 instance on this interface.

Example

This example shows how to add a network area 0.0.0.10 in OSPFv2 instance 201:

```
switch# configure terminal
switch(config)# interface ethernet 1/2

switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router ospf 201 area 0.0.0.10
switch(config-if)# copy running-config startup-config
```

Use the **show ip ospf interface** command to verify the interface configuration. Use the **show ip ospf neighbor** command to see the neighbors for this interface.

Configuring Authentication for an Area

You can configure authentication for all networks in an area or for individual interfaces in the area. Interface authentication configuration overrides area authentication.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that all neighbors on an interface share the same authentication configuration, including the shared authentication key.

Create the key chain for this authentication configuration. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

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



Note For OSPFv2, the key identifier in the **key key-id** command supports values from 0 to 255 only.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# area area-id authentication [message-digest]	Configures the authentication mode for an area.
Step 4	switch(config-router)# interface interface-type slot/port	Enters interface configuration mode.
Step 5	(Optional) Configure one of the following commands: <ul style="list-style-type: none"> • ip ospf authentication-key [0 3] key • ip ospf message-digest-key key-id md5 [0 3] key 	The first command configures simple password authentication for this interface. Use this command if the authentication is not set to key-chain or message-digest. The 0 keyword configures the password in clear text. The 3 keyword configures the password as 3DES encrypted. The second command configures message digest authentication for this interface. Use this command if the authentication is set to message-digest. The key-id range is from 1 to 255. The MD5 option 0 configures the password in clear text and 3 configures the pass key as 3DES encrypted.

	Command or Action	Purpose
Step 6	(Optional) switch(config)# show ip ospf <i>instance-tag interface interface-type slot/port</i>	Displays OSPF information.
Step 7	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configuring Authentication for an Interface

You can configure authentication for all networks in an area or for individual interfaces in the area. Interface authentication configuration overrides area authentication.

Before you begin

Ensure that you have enabled the OSPF feature.

Ensure that all neighbors on an interface share the same authentication configuration, including the shared authentication key.

Create the key chain for this authentication configuration. See the *Cisco Nexus 7000 Series NX-OS Security Configuration Guide*.

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



Note For OSPFv2, the key identifier in the `key key-id` command supports values from 0 to 255 only.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# interface <i>interface-type slot/port</i>	Enters interface configuration mode.
Step 3	switch(config-if)# ip ospf authentication [message-digest]	Enables interface authentication mode for OSPFv2 for either cleartext or message-digest type. Overrides area-based authentication for this interface. All neighbors must share this authentication type.
Step 4	(Optional) switch(config-if)# ip ospf authentication key-chain <i>key-name</i>	Configures interface authentication to use key chains for OSPFv2. For details on key chains, see the <i>Cisco Nexus 7000 Series NX-OS Security Configuration Guide</i> .
Step 5	(Optional) switch(config-if)# ip ospf authentication-key [0 3 7] <i>key</i>	Configures simple password authentication for this interface. Use this command if the authentication is not set to key-chain or message-digest.

	Command or Action	Purpose
		<p>The options are as follows:</p> <ul style="list-style-type: none"> • 0—configures the password in clear text. • 3—configures the pass key as 3DES encrypted. • 7—configures the key as Cisco type 7 encrypted.
Step 6	(Optional) switch(config-if)# ip ospf message-digest-key <i>key-id</i> md5 [0 3 7] <i>key</i>	<p>Configures message digest authentication for this interface. Use this command if the authentication is set to message-digest. The key-id range is from 1 to 255. The MD5 options are as follows:</p> <ul style="list-style-type: none"> • 0—configures the password in clear text. • 3—configures the pass key as 3DES encrypted. • 7—configures the key as Cisco type 7 encrypted.
Step 7	(Optional) switch(config-if)# show ip ospf instance-tag interface <i>interface-type slot/port</i>	Displays OSPF information.
Step 8	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to set an interface for simple, unencrypted passwords and set the password for Ethernet interface 1/2:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# exit
switch(config)# interface ethernet 1/2

switch(config-if)# ip router ospf 201 area 0.0.0.10
switch(config-if)# ip ospf authentication
switch(config-if)# ip ospf authentication-key 0 mypass
switch(config-if)# copy running-config startup-config
```

Configuring Advanced OSPFv2

Configuring Filter Lists for Border Routers

You can separate your OSPFv2 domain into a series of areas that contain related networks. All areas must connect to the backbone area through an area border router (ABR). OSPFv2 domains can connect to external domains as well, through an autonomous system border router (ASBR).

ABRs have the following optional configuration parameters:

- Area range—Configures route summarization between areas.
- Filter list—Filters the Network Summary (type 3) LSAs that are allowed in from an external area.

ASBRs also support filter lists.

Before you begin

Ensure that you have enabled the OSPF feature.

Create the route map that the filter list uses to filter IP prefixes in incoming or outgoing Network Summary (type 3) LSAs.

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

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id filter-list route-map map-name {in out}</code>	Filters incoming or outgoing Network Summary (type 3) LSAs on an ABR.
Step 4	(Optional) <code>switch(config-if)# show ip ospf policy statistics area id filter-list {in out}</code>	Displays OSPF policy information.
Step 5	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure a filter list in area 0.0.0.10:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 filter-list route-map FilterLSAs in
switch(config-router)# copy running-config startup-config
```


Configuring Stub Areas

You can configure a stub area for part of an OSPFv2 domain where external traffic is not necessary. Stub areas block AS External (type 5) LSAs and limit unnecessary routing to and from selected networks. You can optionally block all summary routes from going into the stub area.

Before you begin

Ensure that you have enabled the OSPF feature.

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).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id stub</code>	Creates this area as a stub area.
Step 4	(Optional) <code>switch(config-router)# area area-id default-cost cost</code>	Sets the cost metric for the default summary route sent into this stub area. The range is from 0 to 16777215. The default is 1.
Step 5	(Optional) <code>switch(config-if)# show ip ospf instance-tag</code>	Displays OSPF information.
Step 6	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create a stub area:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 stub
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
<code>router ospf instance-tag</code>	Creates this area as a totally stubby area.

Configuring NSSA

You can configure an NSSA for part of an OSPFv2 domain where limited external traffic is required. You can optionally translate this external traffic to an AS External (type 5) LSA and flood the OSPFv2 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 OSPFv2 autonomous system. Use this option when the NSSA ASBR is also an ABR.
- Default information originate—Generates an NSSA External (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 that you want are flooded throughout the NSSA and other areas.
- Translate—Translates NSSA External LSAs to AS External LSAs for areas outside the NSSA. Use this command on an NSSA ABR to flood the redistributed routes throughout the OSPFv2 autonomous system. You can optionally suppress the forwarding address in these AS External LSAs. If you choose this option, the forwarding address is set to 0.0.0.0.
- 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.

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).

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id nssa [no-redistribution] [default-information-originate]originate [route-map map-name] [no-summary] [translate type7 {always never}] [suppress-fa]</code>	Creates this area as an NSSA.
Step 4	(Optional) <code>switch(config-router)# area area-id default-cost cost</code>	Sets the cost metric for the default summary route sent into this NSSA.
Step 5	(Optional) <code>switch(config-if)# show ip ospf instance-tag</code>	Displays OSPF information.

	Command or Action	Purpose
Step 6	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

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

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa no-summary
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that generates a default route:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa default-info-originate
switch(config-router)# copy running-config startup-config
```

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

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa route-map ExternalFilter no-summary
switch(config-router)# copy running-config startup-config
```

This example shows how to create an NSSA that always translates NSSA External (type 5) LSAs to AS External (type 7) LSAs:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 nssa translate type 7 always
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. You can configure the following optional parameters for a virtual link:

- Authentication—Sets a 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.

You cannot add a virtual link to a stub area.

Before you begin

Ensure that you have enabled the OSPF feature.

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

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# area area-id virtual link router-id</code>	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	(Optional) <code>switch(config-router-vlink)# show ip ospf virtual-link [brief]</code>	Displays OSPF virtual link information.
Step 5	(Optional) <code>switch(config-router-vlink)# authentication [key-chain key-id message-digest null]</code>	Overrides area-based authentication for this virtual link.
Step 6	(Optional) <code>switch(config-router-vlink)# authentication-key [0 3] key</code>	Configures a simple password for this virtual link. Use this command if the authentication is not set to key-chain or message-digest. 0 configures the password in clear text. 3 configures the password as 3DES encrypted.
Step 7	(Optional) <code>switch(config-router-vlink)# dead-interval seconds</code>	Configures the OSPFv2 dead interval, in seconds. The range is from 1 to 65535. The default is four times the hello interval, in seconds.
Step 8	(Optional) <code>switch(config-router-vlink)# hello-interval seconds</code>	Configures the OSPFv2 hello interval, in seconds. The range is from 1 to 65535. The default is 10 seconds.
Step 9	(Optional) <code>switch(config-router-vlink)# message-digest-key key-id md5 [0 3] key</code>	Configures message digest authentication for this virtual link. Use this command if the authentication is set to message-digest. 0 configures the password in cleartext. 3 configures the pass key as 3DES encrypted.
Step 10	(Optional) <code>switch(config-router-vlink)# retransmit-interval seconds</code>	Configures the OSPFv2 retransmit interval, in seconds. The range is from 1 to 65535. The default is 5.

	Command or Action	Purpose
Step 11	(Optional) switch(config-router-vlink)# transmit-delay <i>seconds</i>	Configures the OSPFv2 transmit-delay, in seconds. The range is from 1 to 450. The default is 1.
Step 12	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to create a simple virtual link between two ABRs.

The configuration for ABR 1 (router ID 27.0.0.55) is as follows:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 virtual-link 10.1.2.3
switch(config-router)# copy running-config startup-config
```

The configuration for ABR 2 (Router ID 10.1.2.3) is as follows:

```
switch# configure terminal
switch(config)# router ospf 101
switch(config-router)# area 0.0.0.10 virtual-link 27.0.0.55
switch(config-router)# copy running-config startup-config
```

Configuring Redistribution

You can redistribute routes learned from other routing protocols into an OSPFv2 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.



Note Default information originate ignores match statements in the optional route map.

- **Default metric**—Sets all redistributed routes to the same cost metric.



Note If you redistribute static routes, Cisco NX-OS requires the configuration of **default-information originate** command under the router OSPF process to successfully redistribute or generate the default static route.

Before you begin

Ensure that you have enabled the OSPF feature.

Create the necessary route maps used for redistribution.

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

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# redistribute {bgp id direct eigrp id isis id ospf id rip id static} route-map map-name</code>	Redistributes the selected protocol into OSPF through the configured route map. Note If you redistribute static routes, Cisco NX-OS also redistributes the default static route.
Step 4	<code>switch(config-router)# default-information originate [always] [route-map map-name]</code>	Creates a default route into this OSPF domain if the default route exists in the RIB. Use the following optional keywords: <ul style="list-style-type: none"> • always—Always generate the default route of 0.0.0. even if the route does not exist in the RIB • route-map—Generate the default route if the route map returns true. Note This command ignores match statements in the route map.
Step 5	<code>switch(config-router)# default-metric [cost]</code>	Sets the cost metric for the redistributed routes. This command does not apply to directly connected routes. Use a route map to set the default metric for directly connected routes.
Step 6	(Optional) <code>switch(config)# copy running-config startup-config</code>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

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

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# copy running-config startup-config
```

Limiting the Number of Redistributed Routes

Route redistribution can add many routes to the OSPFv2 route table. You can configure a maximum limit to the number of routes accepted from external protocols. OSPFv2 provides the following options to configure redistributed route limits:

- **Fixed limit**—Logs a message when OSPFv2 reaches the configured maximum. OSPFv2 does not accept any more redistributed routes. You can optionally configure a threshold percentage of the maximum where OSPFv2 logs a warning when that threshold is passed.
- **Warning only**—Logs a warning only when OSPFv2 reaches the maximum. OSPFv2 continues to accept redistributed routes.
- **Withdraw**—Starts the timeout period when OSPFv2 reaches the maximum. After the timeout period, OSPFv2 requests all redistributed routes if the current number of redistributed routes is less than the maximum limit. If the current number of redistributed routes is at the maximum limit, OSPFv2 withdraws all redistributed routes. You must clear this condition before OSPFv2 accepts more redistributed routes.
- You can optionally configure the timeout period.

Before you begin

Ensure that you have enabled the OSPF feature.

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

Procedure

	Command or Action	Purpose
Step 1	<code>switch# configure terminal</code>	Enters global configuration mode.
Step 2	<code>switch(config)# router ospf instance-tag</code>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	<code>switch(config-router)# redistribute {bgp id direct eigrp id isis id ospf id rip id static} route-map map-name</code>	Redistributes the selected protocol into OSPF through the configured route map.
Step 4	<code>switch(config-router)# redistribute maximum-prefix max [threshold] [warning-only withdraw [num-retries timeout]]</code>	Specifies a maximum number of prefixes that OSPFv2 distributes. The range is from 0 to 65536. Optionally specifies the following: <ul style="list-style-type: none"> • threshold—Percent of maximum prefixes that trigger a warning message. • warning-only—Logs an warning message when the maximum number of prefixes is exceeded. • withdraw—Withdraws all redistributed routes. Optionally tries to retrieve the redistributed routes. The <i>num-retries</i> range is from 1 to 12. The <i>timeout</i> is 60 to 600 seconds. The default is 300 seconds. Use

	Command or Action	Purpose
		the clear ip ospf redistribution command if all routes are withdrawn.
Step 5	(Optional) switch(config-router)# show running-config ospf	Displays the OSPFv2 configuration.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to limit the number of redistributed routes into OSPF:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# redistribute bgp route-map FilterExternalBGP
switch(config-router)# redistribute maximum-prefix 1000 75
```

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.

Before you begin

Ensure that you have enabled the OSPF feature.

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	Configure one of the following commands: <ul style="list-style-type: none"> • area area-id range ip-prefix/length [no-advertise] [cost cost] • summary-address ip-prefix/length [no-advertise tag tag] 	<p>The first command creates a summary address on an ABR for a range of addresses and optionally does not advertise this summary address in a Network Summary (type 3) LSA. The cost range is from 0 to 16777215.</p> <p>The second command 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.</p>

	Command or Action	Purpose
Step 4	(Optional) switch(config-router)# [no] discard route {internal external}	When you configure a summary address, Cisco NX-OS software automatically configures a discard route for the summary address to prevent routing black holes and route loops. You can use the no form of this command to prevent the discard routes from being created.
Step 5	(Optional) switch(config-router)# show ip ospf summary-address	Displays information about OSPF summary addresses.
Step 6	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

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

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# area 0.0.0.10 range 10.3.0.0/16
switch(config-router)# copy running-config startup-config
```

This example shows how to create summary addresses on an ASBR:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# summary-address 10.5.0.0/16
switch(config-router)# copy running-config startup-config
```

Configuring Stub Route Advertisements

Use stub route advertisements when you want to limit the OSPFv2 traffic through this router for a short time.

Stub route advertisements can be configured with the following optional parameters:

- **on startup**—Sends stub route advertisements for the specified announce time.
- **wait-for bgp**—Sends stub router advertisements until BGP converges.



Note You should not save the running configuration of a router when it is configured for a graceful shutdown because the router continues to advertise a maximum metric after it is reloaded.

Before you begin

Ensure that you have enabled the OSPF feature.

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# max-metric router-lsa [external-lsa [<i>max-metric-value</i>]] [include-stub] [on-startup [<i>seconds</i>]] [wait-for bgp tag] [summary-lsa [<i>max-metric-value</i>]]	Configures OSPFv2 stub route advertisements.
Step 4	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to enable the stub router advertisements on startup for the default 600 seconds:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# max-metric router-lsa on-startup
switch(config-router)# copy running-config startup-config
```

Configuring the Administrative Distance of Routes

Beginning with Cisco NX-OS Release 6.1, you can set the administrative distance of routes added by OSPFv2 into the RIB.

The administrative distance is a rating of the trustworthiness of a routing information source. A higher value indicates a lower trust rating. Typically, a route can be learned through more than one routing protocol. The administrative distance is used to discriminate between routes learned from more than one routing protocol. The route with the lowest administrative distance is installed in the IP routing table.

Before you begin

Ensure that you have enabled OSPF.

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

See the guidelines and limitations for this feature.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.

	Command or Action	Purpose
Step 3	switch(config-router)# [no] table-map <i>map-name</i> [filter]	Configures the policy for filtering or modifying OSPFv2 routes before sending them to the RIB. You can enter up to 63 alphanumeric characters for the map name. The filter keyword specifies that only routes that are permitted by the route map(<i>map-name</i>) configuration are downloaded to the routing information base (RIB).
Step 4	switch(config-router)# exit	Exits router configuration mode.
Step 5	switch(config)# route-map <i>map-name</i> [permit deny] [<i>seq</i>]	Creates a route map or enters route-map configuration mode for an existing route map. Use <i>seq</i> to order the entries in a route map. Note The permit option enables you to set the distance. If you use the deny option, the default distance is applied.
Step 6	switch(config-route-map)# match route-type <i>route-type</i>	Matches against one of the following route types: <ul style="list-style-type: none"> • external: The external route (BGP, EIGRP, and OSPF type 1 or 2) • inter-area: OSPF inter-area route • internal: The internal route (including the OSPF intra- or inter-area) • intra-area: OSPF intra-area route • nssa-external: The NSSA external route (OSPF type 1 or 2) • type-1: The OSPF external type 1 route • type-2: The OSPF external type 2 route
Step 7	switch(config-route-map)# match ip route-source prefix-list <i>name</i>	Matches the IPv4 route source address or router ID of a route to one or more IP prefix lists. Use the ip prefix-list command to create the prefix list.
Step 8	switch(config-route-map)# match ip address prefix-list <i>name</i>	Matches against one or more IPv4 prefix lists. Use the ip prefix-list command to create the prefix list.
Step 9	switch(config-route-map)# set distance <i>value</i>	Sets the administrative distance of routes for OSPFv2. The range is from 1 to 255.

	Command or Action	Purpose
Step 10	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

This example shows how to configure the OSPFv2 administrative distance for inter-area routes to 150, for external routes to 200, and for all prefixes in prefix list p1 to 190:

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# table-map foo
switch(config-router)# exit
switch(config)# route-map foo permit 10
switch(config-route-map)# match route-type inter-area
switch(config-route-map)# set distance 150
switch(config)# route-map foo permit 20
switch(config-route-map)# match route-type external
switch(config-route-map)# set distance 200
switch(config)# route-map foo permit 30
switch(config-route-map)# match ip route-source prefix-list p1
switch(config-route-map)# match ip address prefix-list p1
switch(config-route-map)# set distance 190
```

The following example shows how to configure a route map for blocking the next hops that are learned through VLAN 10:

```
switch(config)# route-map Filter-OSPF 10 deny
switch(config-route-map)# match interface VLAN 10
switch(config-route-map)# exit
switch(config)# route-map Filter-OSPF 20 permit
```

The following example shows how to configure the **table-map** command with the **filter** keyword to use a route map (Filter-OSPF) to remove the next-hop path that is learned through VLAN 10 but not the next-hop path that is learned through VLAN 20:

```
switch(config)# route ospf p1
switch(config-router)# table-map Filter-OSPF filter
```

Modifying the Default Timers

OSPFv2 includes a number of timers that control the behavior of protocol messages and shortest path first (SPF) calculations. OSPFv2 includes the following optional timer parameters:

- LSA arrival time—Sets the minimum interval allowed between LSAs that arrive from a neighbor. LSAs that arrive faster than this time are dropped.
- Pacing LSAs—Sets 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.
- Throttle LSAs—Sets the rate limits for generating LSAs. This timer controls how frequently LSAs are generated after a 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.

Before you begin

Ensure that you have enabled the OSPF feature.

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# timers lsa-arrival msec	Sets the LSA arrival time in milliseconds. The range is from 10 to 600000. The default is 1000 milliseconds.
Step 4	switch(config-router)# timers lsa-group-pacing seconds	Sets the interval in seconds for grouping LSAs. The range is from 1 to 1800. The default is 240 seconds.
Step 5	switch(config-router)# timers throttle lsa <i>start-time hold-interval max-time</i>	Sets the rate limit in milliseconds for generating LSAs with the following timers: <ul style="list-style-type: none"> • <i>start-time</i>—The range is from 50 to 5000 milliseconds. The default value is 50 milliseconds. • <i>hold-interval</i>—The range is from 50 to 30,000 milliseconds. The default value is 5000 milliseconds. • <i>max-time</i>—The range is from 50 to 30,000 milliseconds. The default value is 5000 milliseconds
Step 6	switch(config-router)# timers throttle spf <i>delay-time hold-time max-wait</i>	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 7	switch(config)# interface <i>type slot/port</i>	Enters interface configuration mode.
Step 8	switch(config-if)# ip ospf hello-interval <i>seconds</i>	Sets the hello interval this interface. The range is from 1 to 65535. The default is 10.

	Command or Action	Purpose
Step 9	switch(config-if)# ip ospf dead-interval <i>seconds</i>	Sets the dead interval for this interface. The range is from 1 to 65535.
Step 10	switch(config-if)# ip ospf retransmit-interval <i>seconds</i>	Sets the estimated time in seconds between LSAs transmitted from this interface. The range is from 1 to 65535. The default is 5.
Step 11	switch(config-if)# ip ospf transmit-delay <i>seconds</i>	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 12	(Optional) switch(config-if)# show ip ospf	Displays information about OSPF.
Step 13	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

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

```
switch# configure terminal
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 OSPFv2 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 OSPFv2 instance. OSPFv2 does not participate in the graceful restart of a neighbor.
- **Planned graceful restart only**—Configures OSPFv2 to support graceful restart only in the event of a planned restart.

Before you begin

Ensure that you have enabled OSPF.

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).

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# router ospf instance-tag	Creates a new OSPFv2 instance with the configured instance tag.
Step 3	switch(config-router)# graceful-restart	Enables a graceful restart. A graceful restart is enabled by default.
Step 4	(Optional) switch(config-router)# graceful-restart grace-period seconds	Sets the grace period, in seconds. The range is from 5 to 1800. The default is 60 seconds.
Step 5	(Optional) switch(config-router)# graceful-restart helper-disable	Disables helper mode. This feature is enabled by default.
Step 6	(Optional) switch(config-router)# graceful-restart planned-only	Configures a graceful restart for planned restarts only.
Step 7	(Optional) switch(config-if)# show ip ospf instance-tag	Displays OSPF information.
Step 8	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

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

```
switch# configure terminal
switch(config)# router ospf 201
switch(config-router)# graceful-restart
switch(config-router)# graceful-restart grace-period 120
switch(config-router)# copy running-config startup-config
```

Restarting an OSPFv2 Instance

You can restart an OSPFv2 instance. This action clears all neighbors for the instance.

To restart an OSPFv2 instance and remove all associated neighbors, use the following command:

Command	Purpose
restart ospf instance-tag	Restarts the OSPFv2 instance and removes all neighbors.

Configuring OSPFv2 with Virtualization

You can configure multiple OSPFv2 instances in each VDC. You can also create multiple VRFs within each VDC and use the same or multiple OSPFv2 instances in each VRF. You assign an OSPFv2 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.

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

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# vrf context <i>vrf-name</i>	Creates a new VRF and enters VRF configuration mode.
Step 3	switch(config)# router ospf <i>instance-tag</i>	Creates a new OSPFv2 instance with the configured instance tag.
Step 4	switch(config-router)# vrf <i>vrf-name</i>	Enters VRF configuration mode.
Step 5	(Optional) switch(config-router-vrf)# maximum-paths <i>path</i>	Configures the maximum number of equal OSPFv2 paths to a destination in the route table for this VRF. Used for load balancing.
Step 6	switch(config-router-vrf)# interface <i>interface-type slot/port</i>	Enters interface configuration mode.
Step 7	switch(config-if)# vrf member <i>vrf-name</i>	Adds this interface to a VRF.
Step 8	switch(config-if)# ip address <i>ip-prefix/length</i>	Configures an IP address for this interface. You must do this step after you assign this interface to a VRF.
Step 9	switch(config-if)# ip router ospf <i>instance-tag area area-id</i>	Assigns this interface to the OSPFv2 instance and area configured.
Step 10	(Optional) switch(config)# copy running-config startup-config	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Example

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

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config)# router ospf 201
switch(config)# interface ethernet 1/2

switch(config-if)# vrf member NewVRF
switch(config-if)# ip address 192.0.2.1/16
switch(config-if)# ip router ospf 201 area 0
switch(config)# copy running-config startup-config
```

Verifying the OSPFv2 Configuration

To display the OSPFv2 configuration, perform one of the following tasks:

Command	Purpose
show ip ospf [<i>instance-tag</i>] [vrf <i>vrf-name</i>]	Displays the information about one or more OSPFv2 routing instances. The output includes the following area-level counts: <ul style="list-style-type: none"> • Interfaces in this area—A count of all interfaces added to this area (configured interfaces). • Active interfaces—A count of all interfaces considered to be in router link states and SPF (UP interfaces). • Passive interfaces—A count of all interfaces considered to be OSPF passive (no adjacencies will be formed). • Loopback interfaces—A count of all local loopback interfaces.
show ip ospf border-routers [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 link-state database summary.
show ip ospf interface <i>number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 interface configuration.
show ip ospf lsa-content-changed-list <i>neighbor-id interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 LSAs that have changed.
show ip ospf neighbors [<i>neighbor-id</i>] [detail] [<i>interface-type number</i>] [vrf { <i>vrf-name</i> all default management }] [summary]	Displays the list of OSPFv2 neighbors.
show ip ospf request-list <i>neighbor-id interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the list of OSPFv2 link-state requests.

Command	Purpose
show ip ospf retransmission-list <i>neighbor-id interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the list of OSPFv2 link-state retransmissions.
show ip ospf route [<i>ospf-route</i>] [summary] [vrf { <i>vrf-name</i> all default management }]	Displays the internal OSPFv2 routes.
show ip ospf summary-address [vrf { <i>vrf-name</i> all default management }]	Displays information about the OSPFv2 summary addresses.
show ip ospf virtual-links [brief] [vrf { <i>vrf-name</i> all default management }]	Displays information about OSPFv2 virtual links.
show ip ospf vrf { <i>vrf-name</i> all default management }	Displays information about VRF-based OSPFv2 configuration.
show running-configuration ospf	Displays the current running OSPFv2 configuration.

Monitoring OSPFv2

To display OSPFv2 statistics, use the following commands:

Command	Purpose
show ip ospf policy statistics area <i>area-id filter-list</i> { in out } [vrf <i>vrf-name</i> all default management]	Displays the OSPFv2 route policy statistics for an area.
show ip ospf policy statistics redistribute { bgp <i>id</i> direct eigrp <i>id</i> isis <i>id</i> ospf <i>id</i> rip <i>id</i> static } [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 route policy statistics.
show ip ospf statistics [vrfnumber [vrf { <i>vrf-name</i> all default management }]]	Displays the OSPFv2 event counters.
show ip ospf traffic <i>interface-type number</i> [vrf { <i>vrf-name</i> all default management }]	Displays the OSPFv2 packet counters.

Configuration Examples for OSPFv2

```
feature ospf
router ospf 201
  router-id 290.0.2.1
interface ethernet 1/2

ip router ospf 201 area 0.0.0.10
ip ospf authentication
ip ospf authentication-key 0 mypass
```

Related Documents for OSPFv2

For more information related to OSPFv2 CLI commands, see the *Cisco Nexus 5000 Series Command Reference*

Feature History for OSPFv2

This table includes only the updates for those releases that have resulted in additions or changes to the feature.

Table 2: Feature History for OSPFv2

Feature Name	Release	Feature Information
OSPF Packet-size	8.3(2)	Added support for configuring OSPF packet-size.
OSPF—Distribute List to Filter Paths	6.2(6a)	Added support for filtering next-hop paths for an OSPF route to prevent the path from being added to the RIB.
Administrative distance of routes	6.2(2)	Added the filter keyword to the table-map command to specify that only routes permitted by the route map are downloaded to the RIB.
Route summarization	6.2(2)	Added the ability to prevent discard routes from being created
OSPFv2	6.2(2)	Added support for the optional name lookup parameter for OSPFv2 instances.
OSPFv2	6.1(1)	Added support for more than four process instances for OSPFv2 per VDC.
OSPFv2	6.1(1)	Added support for configuring the administrative distance of routes for OSPFv2.
Passive interface	5.2(1)	Added support for setting the passive interface mode on all interfaces in the router or VRF.
OSPFv2	5.1(2)	Added options for the max-metric router-lsa command.
BFD	5.0(2)	Added support for BFD. See the <i>Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide</i> for more information.
OSPFv2	4.0(1)	This feature was introduced.

