

OSPF Configuration Guide for Cisco 8000 Series Routers, Cisco IOS XR Releases

Cisco 8000 Series Routers
Updated September 8, 2026

Topics included

1 Getting Started....4

2 OSPF fundamentals and basic configuration....6
OSPF configuration requirements for Cisco IOS XR....7
OSPF implementation features and supported versions....7
How OSPF processes work....8
OSPF Implementations....8
Enable OSPF routing....9
Verify OSPF configuration and operation....10

3 OSPF areas, neighbors, and route-control....12
OSPF areas....13
OSPF neighbors, designated routers, and default routes....13
Configure OSPF area types....14
Configure OSPF neighbors for nonbroadcast networks....15
OSPF route advertisement and timer control features....16
Configure OSPF link-state advertisement frequency....18
Redistribute routes into OSPF....19

4 OSPF security and high availability....22
OSPF authentication methods....23
Configure OSPF authentication....24
OSPF resiliency and graceful restart....26
Enable BFD strict mode for OSPF....29
Configure OSPF restart and nonstop routing features....30

5 OSPF optimization scale and fast-reroute....34
OSPF scale, optimization, and protection features....35
Configure OSPF SPF scheduling and prefix prioritization....36
Configure OSPF scale and topology controls....38
IP fast reroute loop-free alternate....40

6 OSPF advanced integrations....42
OSPF virtual links and sham links....43
OSPF integrations and management support....44
Configure LDP IGP synchronization for OSPF....45
Configure OSPF Transport and Service Integrations....45

Configure OSPF virtual and sham links.....46

1 Getting Started

Outlines the release history and update summary for the OSPF Configuration Guide. It enables users to access the latest IOS XR features and release information through a centralized, continuously updated resource.

This cumulative guide provides a single, continuously updated version that includes all the latest IOS XR features and release updates. It simplifies your experience by letting you bookmark one link and access the complete guide, instead of navigating through multiple release-specific versions.

Specific changes or updates tied to individual releases are clearly called out within the relevant sections. For a list of features introduced in a specific release, refer to the Release Notes.

The table lists the release numbers for which this document has been updated since its initial publication.

Table 1: Changes to this document

Date	Summary
September 2026	Republished for Release 26.3.1
June 2026	First published for Release 26.2.1

2 OSPF fundamentals and basic configuration

Topics:

- OSPF configuration requirements for Cisco IOS XR
- OSPF implementation features and supported versions
- How OSPF processes work
- OSPF Implementations
- Enable OSPF routing
- Verify OSPF configuration and operation

Explains OSPFv2 and OSPFv3 requirements, implementation features, and routing processes, and describes how to enable and verify OSPF routing on Cisco IOS XR routers.

OSPF configuration requirements for Cisco IOS XR

Lists the essential requirements you must meet before configuring OSPF on Cisco IOS XR routers.

Provides configuration and planning requirements for implementing OSPF on Cisco IOS XR routers.

Requirements

The following requirements apply before you configure OSPF on Cisco IOS XR routers:

- User group assignment: You must belong to a user group associated with a task group containing the required task IDs to enable OSPF commands. See the command reference for task ID details. Contact your AAA administrator if you encounter command restrictions.
- OSPFv3 and IPv6 familiarity: You should be familiar with IPv6 addressing and basic network configuration before configuring OSPFv3. For more information, see the "Implementing Network Stack IPv4 and IPv6" module.
- Interface preparation for OSPFv3: Complete your OSPF network strategy and IPv6 planning before enabling OSPFv3 on interfaces. Decide if multiple OSPF areas are required, and make sure IPv6 is enabled on each interface.
- Authentication configuration (optional): If you plan to enable authentication (IP Security), choose between plaintext or Message Digest 5 (MD5) authentication, and determine whether authentication applies to an entire area or to specific interfaces.

Use these requirements as a reference when planning, configuring, or validating OSPF on Cisco IOS XR routers.

Result

Meeting these requirements helps ensure successful preparation, configuration, and validation of OSPF on Cisco IOS XR routers.

OSPF implementation features and supported versions

Provides key facts about OSPF implementation, supported features, and protocol details for Cisco 8000 Series Routers running IOS XR.

Provides an overview of OSPF implementation, supported features, and protocol details.

The OSPF protocol is an Interior Gateway Protocol (IGP) developed by the OSPF working group of the Internet Engineering Task Force (IETF) for IP networks. OSPF supports IP subnetting, tagging of externally derived routing information, and packet authentication.

Cisco 8000 Series Routers running IOS XR support both OSPF Version 2 (OSPFv2) for IPv4 and OSPF Version 3 (OSPFv3) for IPv6. OSPFv3 adds support for IPv6 routing prefixes.

- IP subnetting and tagging of external routing information
- Packet authentication for secure communication
- Operational integration and resiliency features designed for complex network environments
- Support for both OSPFv2 (IPv4) and OSPFv3 (IPv6)

Understanding OSPF behavior, dependencies, and design considerations is recommended before configuring or verifying OSPF features.

How OSPF processes work

Describes how OSPF processes operate on Cisco 8000 Series Routers to enable link-state routing and rapid network convergence, supporting scalable and hierarchical routing on Cisco IOS XR.

Summary

OSPF (Open Shortest Path First) processes act as logical entities within Cisco 8000 Series Routers to support both OSPFv2 (for IPv4) and OSPFv3 (for IPv6). They allow routers to compute shortest paths, maintain accurate topology knowledge, and enable advanced hierarchical routing configurations across distinct domains.

The key components involved in the process are:

- OSPF process: A logical routing unit within the router that manages its own link-state database and facilitates link-state-based path calculation.
- Router ID: A unique identifier for each OSPF process, assigned manually or derived from interface addresses or checkpoint data.
- Routing domain: Represents a distinct OSPF domain; multiple OSPF processes can run on one router, each with its own database, enabling separation of routing information. Routes are shared across domains only via redistribution.

Workflow

The process involves the following stages:

1. Adjacency establishment and database synchronization: The router forms OSPF adjacencies with its neighbors and synchronizes link-state databases.
   - On broadcast and NBMA networks, a designated router or backup designated router floods Link State Advertisements (LSAs).
   - On point-to-point links, LSA flooding occurs directly between routers.
2. Link-state advertisement and route calculation: Routers exchange and update LSAs to reflect network changes. Each OSPF process uses updated link-state databases to calculate the shortest path tree for its domain.
3. Propagation of changes and convergence: When topology changes, only the updated link-state information is shared, enabling the network to quickly converge while minimizing routing overhead.

Result

OSPF processes achieve scalable, efficient routing by ensuring fast convergence and low control-plane overhead—propagating only changed routing information throughout the network.

OSPF Implementations

Describes OSPF implementations, including behavior, dependencies, and configuration context for Cisco 8000 Series Routers running IOS XR.

An OSPF implementation is a routing protocol configuration that:

- facilitates dynamic exchange of routing information between routers within a network,
- supports hierarchical network design through areas, and
- ensures rapid convergence and optimal route selection based on link-state algorithms.

Additional reference information

OSPF (Open Shortest Path First) is used to enable, configure, tune, secure, integrate, and verify routing on Cisco 8000 Series Routers running IOS XR. Before configuring or verifying OSPF, consider the protocol's requirements, operational dependencies, and design considerations such as network topology, area types, and authentication methods.

Examples

For example, configuring OSPF on Cisco 8000 Series Routers involves defining OSPF areas, setting router IDs, and applying authentication for enhanced security.

Enable OSPF routing

Configure OSPF on Cisco 8000 Series Routers following a standard, repeatable workflow.

Enable OSPF routing by completing the essential configuration steps and preparing for result verification.

Perform this task after you have identified the OSPF process, area, relevant interfaces, and your deployment’s specific routing requirements.

Before you begin

- Ensure at least one IP address is configured on the router (OSPF routing requires at least one active IP address).
- Identify the OSPF process number, area IDs, and interfaces for OSPF participation.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. Configure the minimum OSPF requirements by defining an OSPF process with a router ID, assigning a backbone or non-backbone area, and enabling OSPF on the required interfaces.

   Example

   ```
   Router# router ospf 1
   Router (config-ospf)# router-id 192.168.4.3
   Router(config-ospf)# area 0
   Router(config-ospf-ar)# interface TenGigE 0/1/0/3
   Router(config-ospf-ar-if)# log adjacency changes detail
   ```

   Area IDs can be specified in decimal (e.g., area 1000) or dotted-decimal (e.g., area 0.0.3.232). Use a consistent format; IPv4 address notation is recommended.

2. Verify OSPF configuration.

   Example

   ```
   /* OSPF areas must be explicitly configured,
   and interfaces configured under
   the area configuration mode are explicitly bound
    to that area.
   In this example, interface 10.1.2.0/24 is bound
    to area 0 and interface 10.1.3.0/24 is bound to area 1.
   */
   Router# show running-config interface TenGigE 0/3/0/0
   interface TenGigE 0/3/0/0
   ```

```
 ip address 10.1.2.1 255.255.255.0
 negotiation auto
!
interface TenGigE 0/3/0/1
 ip address 10.1.3.1 255.255.255.0
 negotiation auto
!
router ospf 1
 router-id 10.2.3.4
area 0
 interface TenGigE 0/3/0/0
!
area 1
 interface TenGigE 0/3/0/1
!
!
```

The OSPF process is configured to redistribute the specified routes, and the operational state reflects your intended deployment.

Verify OSPF configuration and operation

Verify your router's OSPF configuration and operation to confirm accurate routing and operational readiness.

Ensure that OSPF is configured correctly and functioning as intended on your router.

Use this task to validate OSPF process status, check routing tables, LSAs, interfaces, neighbors, and clear OSPF state if needed for troubleshooting or optimal operation.

Before you begin

- Identify the OSPF or OSPFv3 process name, area, interface, and any VRF you need to verify.
- Ensure OSPF is already configured before performing these verification steps.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. (Optional) Verify the general status of OSPF routing processes.

   Example

   ```
   Router# show ospf group1
   ```

2. (Optional) Review OSPF routing table entries to ABRs and ASBRs.

   Example

   ```
   Router# show ospf group1 border-routers
   ```

3. (Optional) Examine the OSPF database and LSAs for the selected process.

Example

```
Router# show ospf group2 database
```

4. (Optional) Check for LSAs waiting to be flooded on an interface.

   Example

   ```
   Router# show ospf 100 flood-list interface TenGigE 0/3/0/0
   ```

5. (Optional) Display OSPF interface information, including state, area, cost, timers, neighbors, and authentication.

   Example

   ```
   Router# show ospf 100 interface TenGigE 0/3/0/0
   ```

6. (Optional) Display OSPF neighbor relationships.

   Example

   ```
   Router# show ospf 100 neighbor
   ```

7. If you need to reset the OSPF process for troubleshooting, reset the OSPF process.

   Example

   ```
   Router# clear ospf 100 process
   ```

8. (Optional) If necessary for troubleshooting, clear OSPF redistribution, routes, or VRF-specific state.

   Example

   ```
   Router# clear ospf 100 redistribution
   Router# clear ospf 100 routes
   Router# clear ospf 100 vrf vrf_1 process
   ```

You verify OSPF configuration and operation, and optionally clear OSPF state as needed for troubleshooting.

3 OSPF areas, neighbors, and route-control

Topics:

- OSPF areas
- OSPF neighbors, designated routers, and default routes
- Configure OSPF area types
- Configure OSPF neighbors for nonbroadcast networks
- OSPF route advertisement and timer control features
- Configure OSPF link-state advertisement frequency
- Redistribute routes into OSPF

Explains how OSPF areas, router roles, neighbors, designated routers, route advertisements, and timers support scalable routing, and describes how to configure area types, nonbroadcast neighbors, LSA frequency, default routes, and route redistribution.

OSPF areas

Explains how OSPF areas segment an autonomous system to enhance scalability, improve routing efficiency, and control traffic flow, while also describing router roles that manage internal and external connectivity.

An OSPF area is a network architecture concept that

- divides an OSPF autonomous system into multiple areas to improve scalability and reduce routing overhead,
- assigns routers specific roles—such as interior routers, area border routers (ABRs), and autonomous system boundary routers (ASBRs), and
- defines how routing information is distributed within and across areas to control traffic flow and manage external connectivity.

OSPF area types

- Backbone area (Area 0): Distributes routing information between all areas and uses the reserved area ID 0.0.0.0.
- Stub area: Does not accept external route advertisements and typically has only one router connecting it to the rest of the autonomous system.
- Not-so-Stubby Area (NSSA): Similar to a stub area, but allows limited import of external routes (Type 7 LSAs) and supports redistribution of routes from other protocols.

Router roles

- Interior router: Attaches to a single area; all interfaces reside within that area.
- Area border router (ABR): Connects multiple areas, runs separate OSPF processes for each, and summarizes area information to the backbone area.
- Autonomous system boundary router (ASBR): Connects OSPF with other autonomous systems and redistributes routes from external protocols (such as BGP) into OSPF using Type 5 and Type 7 LSAs.

Additional reference information

- Areas subdivide an OSPF autonomous system into smaller, manageable segments, which help reduce routing traffic and the size of the link-state database.
- OSPF routers within an area only see their area's topology, limiting unnecessary flooding of link-state information.
- ABRs send configuration summaries for their attached areas to the backbone, supporting efficient interarea routing.
- ASBRs connect the OSPF network to other autonomous systems, ensuring the exchange of external routing information.
- Stub ABRs advertise a single default route to stub areas, conserving database space.
- NSSA ABRs translate Type 7 LSAs into Type 5 LSAs for propagation through the routing domain.
- Autonomous systems are collections of networks under unified administrative control and may include multiple OSPF areas.

OSPF neighbors, designated routers, and default routes

Explains the core concepts of OSPF neighbors, designated routers, and default routes in OSPF-enabled networks.

OSPF neighbor

An OSPF neighbor is a router on the same OSPF-enabled network segment that:

- exchanges OSPF Hello packets with other routers,
- forms and maintains OSPF adjacencies, and
- participates in updating and synchronizing the OSPF link-state database.

Designated router

A designated router is an elected OSPF router on a broadcast or non-broadcast multi-access network that:

- coordinates OSPF link-state advertisement (LSA) exchanges,
- reduces OSPF overhead by centralizing LSA flooding, and
- maintains adjacency with all other routers on the segment (DROthers).

Default route

A default route is an entry in the OSPF routing table that:

- matches all destinations not explicitly listed by more specific routes,
- forwards unmatched packets to a specified next-hop gateway, and
- enables OSPF routers to reach external destinations with minimal route entries.

Additional reference information

- **OSPF neighbors** are established through the Hello protocol, allowing routers to verify network reachability and exchange LSAs.
- **Designated routers (DRs)** are selected based on the highest priority or, if tied, the highest Router ID, and optimize LSA flooding on multi-access networks.
- Use the **default-information originate** command to inject a default route into the OSPF domain; you can also redistribute a static default route.

Configure OSPF area types

Configure OSPF area types to optimize routing updates and control link-state advertisement (LSA) flooding.

Configure OSPF areas as stub or Not-So-Stubby Areas (NSSA) to control routing advertisements and enhance OSPF network efficiency.

Perform this task after identifying the OSPF process, area numbers, interfaces, and deployment-specific routing requirements.

Before you begin

- Ensure you have the required configuration privileges.
- Complete OSPF addressing, routing policy, and interface planning.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. Access OSPF configuration mode and set the router ID.

   Example

   ```
   Router# configure
   Router(config)# router ospf 1
   Router(config-ospf)# router-id 192.168.4.3
   ```

2. Define a stub area and reduce LSAs (if needed).

   To configure a stub area and prevent the ABR from sending summary LSAs (Type 3) into the area, use the `no-summary` keyword:

   Example

   ```
   Router(config-ospf)# area 1
   Router(config-ospf-ar)# stub no-summary
   ```

3. Configure an area as Not-So-Stubby Area (NSSA), with optional modifiers.

   You may use `no-redistribution`, `translate`, or `stub` as needed. Use the following syntax:

   Example

   ```
   Router(config-ospf-ar)# nssa no-redistribution
   Router(config-ospf-ar)# nssa translate 1 always
   Router(config-ospf-ar)# stub
   ```

   Or configure the area simply as NSSA:

   ```
   Router(config-ospf-ar)# nssa
   ```

   If you previously configured a stub or NSSA area with optional keywords, reissue the `stub` or `nssa` command without any keywords to remove options.

   Do not use the `no` form of the command to remove options such as `no-summary`, `no-redistribution`, `default-information-originate`, or `translate`.

Your OSPF area configuration is applied, and the router operates in the intended stub or NSSA mode.

Configure OSPF neighbors for nonbroadcast networks

Configure OSPF neighbors for nonbroadcast networks on Cisco 8000 Series Routers.

Set up explicit OSPF neighbor relationships on nonbroadcast (NBMA) networks by using repeatable configuration steps.

Perform this task to determine your OSPF process, area, interface assignments, and routing requirements.

Before you begin

- Ensure the network is fully meshed so every router connects directly, or via virtual circuits, to every other router.

- Decide whether your NBMA network will use broadcast or nonbroadcast mode.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

Configure NBMA networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network.

Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# area 0
Router(config-ospf-ar)# network non-broadcast
Router(config-ospf-ar)# dead-interval 40
Router(config-ospf-ar)# hello-interval 10
Router(config-ospf-ar)# interface GigabitEthernet 0/2/0/0
Router(config-ospf-ar-if)# neighbor 10.20.20.1 priority 3 poll-interval 15
! or
Router(config-ospf-ar-if)# neighbor fe80::3203:a0ff:fe9d:f3fe
Router(config-ospf-ar-if)# exit
Router(config-ospf-ar)# commit
```

The router applies the OSPF configuration and establishes neighbor adjacencies on the nonbroadcast network segment.

OSPF route advertisement and timer control features

Explains OSPF route advertisement and timer control features, covering how these mechanisms manage router information, update timing, and enable accurate and efficient routing tables.

OSPF route advertisement and timer control features are routing protocol mechanisms that:

- manage the dissemination of network link-state information,
- regulate the timing of network topology updates, and
- determine how routers on Cisco 8000 Series platforms maintain accurate and efficient routing tables.

Link-State Advertisement Types for OSPF Version 2

Each Link-State Advertisement (LSA) type serves a specific purpose:

- Router LSA (Type 1): Describes the links a router has within a single area and the cost of each link. Indicates path computation capabilities, router roles (ABR or ASBR), virtual links, and advertises stub networks. Flooded within an area only.
- Network LSA (Type 2): Describes link-state and cost information for all routers attached to a multiaccess network segment. Generated by the designated router and lists all router interfaces attached.
- Summary LSA for ABRs (Type 3): Advertises internal networks to routers in other areas (interarea routes). May represent single or aggregated networks. Only ABRs generate these LSAs.
- Summary LSA for ASBRs (Type 4): Advertises an ASBR and the cost to reach it. Used to determine the best path to external networks. Generated by ABRs.

- Autonomous System External LSA (Type 5): Redistributes routes from another autonomous system, typically from a different routing protocol.
- Autonomous System External LSA (Type 7): Carries external route information within a Not-So-Stubby Area (NSSA). Originated and advertised within an NSSA only and not flooded to other areas.
- Intra-area-prefix LSA (Type 9): Originated for routers or transit networks with unique link-state IDs that describe association and contain prefixes for stub and transit networks.
- Area local scope (Type 10): Opaque LSAs not flooded past the borders of their associated area.
- Link-state (Type 11): Flooded throughout the autonomous system, similar flooding scope to Type 5 LSAs but not implemented in stub areas.

Additional reference information

The OSPF FIB Download Notification feature minimizes ingress traffic drops after line card reload and is enabled by default.

OSPF registers with the Routing Information Base (RIB) using the Interface Table Attribute Library (ITAL), which keeps interfaces down until all routes are downloaded to the Forwarding Information Base (FIB). The "Interface Up" notification occurs when routes are complete.

RIB notifications include:

- Node is lost
- Node is created
- Node's FIB upload completion

OSPFv3 Timers and SPF Throttle Values

- timers throttle lsa all: start-interval 50 ms, hold-interval 200 ms
- timers throttle spf: spf-start 50 ms, spf-hold 200 ms, spf-max-wait 5000 ms

OSPF SPF Throttling

SPF scheduling allows millisecond intervals for calculations and helps delay SPF calculations during network instability. SPF is scheduled for topology changes, and multiple events may trigger a single run.

SPF calculation intervals are chosen dynamically based on topology stability. Intervals double until reaching max-wait, then remain constant. Intervals reset based on timing parameters if events occur after max-wait begins.

Example:

```
timers spf 5 1000 90000
```

Start interval: 5 ms; initial wait: 1000 ms; maximum wait: 90,000 ms. Wait intervals double until max-wait, then reset as topology stabilizes.

Subdefinitions for OSPFv3 LSA Types

- Router LSA (Type 1): No address information; protocol independent; multiple router LSAs may need to be concatenated before SPF calculation.
- Network LSA (Type 2): No address information; protocol independent; generated by designated router.
- Interarea-prefix LSA (Type 3): Addresses are "prefix and prefix length" (default route = prefix length 0).
- Interarea-router LSA (Type 4): Advertises ASBR and cost.

- Autonomous System External LSA (Type 5): Addresses are "prefix and prefix length." Default route = prefix length 0.
- Type 7 LSA: Carries external information within NSSA.
- Link LSA (Type 8): Flooded only within their link; communicates link-local address, IPv6 prefixes, and option bits.
- Intra-area-prefix LSA (Type 9): Multiple per router/network; unique IDs describe association.

Address prefixes appear in most new LSAs as three fields: Prefix Length, Prefix Options, and Address Prefix. Inter-area and intra-area LSAs carry IPv6 prefix information that replaces the IPv4 information contained in router/network LSAs. Certain LSAs have expanded options fields (24 bits for IPv6 support). In OSPFv3, the link-state ID identifies database items, but addresses are carried in the LSA body.

Configure OSPF link-state advertisement frequency

Configure OSPF link-state advertisement frequency to optimize route convergence and control LSA flooding intervals.

Tune LSA frequency as part of OSPF deployment, especially in large or dynamic networks, to balance responsiveness with resource usage.

Perform this task to optimize OSPF or OSPFv3 route convergence by adjusting how frequently LSAs are generated and flooded, improving routing stability and efficiency.

Before you begin

- Identify the OSPF process and relevant interfaces.
- Complete OSPF addressing, routing policy, and interface planning.
- Ensure you have configuration privileges.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. Enter global configuration mode.

   Example

   ```
   Router# configure
   ```

2. Access OSPF or OSPFv3 router configuration mode.

   Example

   ```
   Router(config)# router ospf 1
   Router(config)# router ospfv3 1
   ```

3. Set the router ID (if required).

   Example

   ```
   Router(config-ospf)# router-id 192.168.4.3
   ```

4. Adjust OSPF LSA timers for desired frequency and flooding behavior.

Example

```
Router(config-ospf)# timers lsa refresh 1800
Router(config-ospf)# timers lsa min-arrival 2
Router(config-ospf)# timers lsa group-pacing 1000
```

5. Commit the configuration and exit.

Example

```
Router(config-ospf)# commit
Router(config-ospf)# end
```

The router now uses optimized LSA frequency settings for efficient, stable routing table updates.

Redistribute routes into OSPF

Configure OSPF route redistribution using a repeatable CLI workflow that follows the recommended command sequence.

Redistribute routes from an internal gateway protocol (IGP), such as BGP or another OSPF process, into your OSPF routing domain.

Perform this task after identifying the OSPF process, area numbers, interfaces, and deployment-specific routing requirements.

Before you begin

- Determine the source protocol and routing policy for redistribution.
- Review documentation on routing policy if necessary.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

Configure the OSPF process to redistribute the specified routes, and the operational state reflects your intended deployment.

Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# redistribute bgp 100
Router(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0
```

```
Router# configure
Router(config)# router ospfv3 1
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# redistribute bgp 100
Router(config-ospf)# summary-prefix 2010:11:22::/32
```

The OSPF process is configured to redistribute the specified routes, and the operational state reflects your intended deployment.

4 OSPF security and high availability

Topics:

- OSPF authentication methods
- Configure OSPF authentication
- OSPF resiliency and graceful restart
- Enable BFD strict mode for OSPF
- Configure OSPF restart and nonstop routing features

Explains how authentication, BFD strict mode, graceful restart, nonstop routing, and nonstop forwarding protect OSPF exchanges, preserve adjacencies, and minimize routing disruption during process restarts and route processor failovers.

OSPF authentication methods

Explains OSPF authentication methods, mechanisms, and strategies, focusing on protocol behavior, dependencies, and best practices for secure deployment.

A set of OSPF authentication methods is a group of routing security measures that:

- authenticate all OSPF protocol exchanges between routers,
- support multiple authentication types such as plain text and MD5 to prevent unauthorized routing updates, and
- enable configuration flexibility at the process, area, interface, or virtual link level.

Types of OSPF authentication

OSPF Version 2 supports two types of authentication: plain text and MD5. By default, authentication is disabled (null authentication per RFC 2178). OSPF Version 3 supports all authentication types except key rollover.

Configuring authentication strategies

Authentication can be specified for an OSPF process or area, or more granularly on an interface or virtual link. Each interface or virtual link uses only one authentication type at a time; interface-level settings override area or process configuration. Configuring authentication at the area level (using a command such as `message-digest` for MD5 or HMAC SHA-256) enables all interfaces in an area to use the same method, simplifying the configuration.

Key rollover

Key rollover enables you to change MD5 keys in an active network without disrupting OSPF adjacencies. During the transition, routers send duplicate packets authenticated with both old and new keys. After all routers use the new key, remove the old key from each device's configuration.

Neighbors and adjacency in OSPF

Routers on the same Layer 2 segment become neighbors. OSPF uses the hello protocol to discover and maintain neighbor relationships by periodically sending hello packets. When a router sees itself in a neighbor's hello packet, they form a relationship and synchronize their databases (adjacency). On broadcast and non-broadcast multi-access (NBMA) networks, all neighbors become adjacent.

OSPF strict-mode support for BFD dampening

Strict mode is a BFD operation mode in OSPF where a neighbor remains down until the BFD session is up. Strict and non-strict modes are incompatible; both routers must use strict mode to form an adjacency. If other BFD clients have already established a session, OSPF may form a neighbor relationship by design. Strict-mode can delay neighbor establishment due to BFD dependency.

Feature Name	Release Information	Feature Description
MD5 Authentication	Release 25.1.1	Introduced for Fixed Systems (8010 [ASIC: A100]) (select variants only). Supported on Cisco 8011-4G24Y4H-I routers.

Feature Name	Release Information	Feature Description
MD5 Authentication	Release 24.1.1	Introduced for Fixed Systems (8200 [ASIC: P100], 8700 [ASIC: P100, K100]) and Modular Systems (8800 [LC ASIC: P100]) (select variants only). MD5 generates a message digest for each OSPF packet, preventing unauthorized updates. Supported on 8212-48FH-M, 8711-32FH-M, 8712-MOD-M, 88-LC1-36EH, 88-LC1-12TH24FH-E, 88-LC1-52Y8H-EM.

OSPF routing exchanges are authenticated according to the configured method. Cryptographic authentication uses MD5 to generate a message digest with a key, which prevents unauthorized updates. Keys are associated with interfaces and can be managed using keychains with attributes such as activation times, IDs, and algorithms.

The key rollover feature allows MD5 key updates without disrupting adjacencies. Routers send both old and new key-authenticated packets during migration. After all neighbors accept the new key, only the new key is used and the old key should be removed from devices.

Configure OSPF authentication

Configure OSPF authentication using MD5 or plain text methods.

OSPF authentication helps prevent unauthorized routers from injecting incorrect routing information. Interface-level configuration takes precedence over area or process-level settings. If authentication is not set on the interface, it inherits the configuration from the area or process.

Enforce secure OSPF routing protocol operations by implementing authentication at the process, area, and interface levels to prevent unauthorized routing updates.

Before you begin

- Determine whether to apply authentication globally, per OSPF area, or on specific interfaces.
- Decide on the OSPF authentication method: MD5 (preferred) or plain text.
- Ensure a valid OSPF keychain is configured if using MD5 authentication.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. Configure MD5 message digest authentication for the OSPF process and apply it to the required area or interface.

   Example

```
Router(config)# router ospf 1
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# area 1
Router(config-ospf-ar)# interface GigabitEthernet0/4/0/1
Router(config-ospf-ar-if)# authentication message-digest keychain ospf_intf_1
```

2. Define the OSPF keychain with multiple keys and send-lifetime values.

   Example

```
key chain ospf_intf_1
key 1
send-lifetime 11:30:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 2
send-lifetime 11:40:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 3
send-lifetime 11:50:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 4
send-lifetime 12:00:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
key 5
send-lifetime 12:10:30 May 1 2007 duration 600
cryptographic-algorithm MD5
key-string clear ospf_intf_1
```

3. Verify that keychain authentication is enabled on the interface.

   Example

```
Router# show ospf 1 interface GigabitEthernet0/4/0/1

Keychain-based authentication enabled
Key id used is 3
```

4. Verify that the configured OSPF keys are valid and active.

   Example

```
Router# show key chain ospf_intf_1
Key-chain: ospf_intf_1/ -
Key 1 -- ...
Key 2 -- ...
Key 3 -- ... [Valid now]
...
```

5. (Optional) Configure additional OSPF security features as needed, such as router ID, adjacency logging, NSF, throttle timers, or TTL hops.

   Example

```
Router(config)# router ospf 1
```

```
Router(config-ospf)# router-id 10.10.10.100
Router(config-ospf-ar-if)# log adjacency changes detail
Router(config-ospf)# nsf ietf
Router(config-ospf)# timers throttle spf 500 500 10000
Router(config-ospf-ar)# interface GigabitEthernet0/5/0/0
Router(config-ospf-ar-if)# security ttl hops 2
Router# show ospf 1 interface GigabitEthernet0/5/0/0
```

6. (Optional) Configure plain text or null authentication at the area or interface level, if needed.

   Example

```
Router(config)# router ospf 1
Router(config-ospf)# authentication message-digest
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# message-digest-key 4 md5 yourkey
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface GigabitEthernet0/1/0/3
Router(config-ospf-ar)# exit
Router(config-ospf-ar)# interface GigabitEthernet0/3/0/0
Router(config-ospf-ar-if)# authentication null
```

7. Verify that the GTSM security TTL value is configured on an OSPF interface.

   Example

```
show ospf 1 interface GigabitEthernet0/5/0/0

TTL security enabled, hop count 2
```

OSPF authentication is enabled and verified. The router securely exchanges OSPF routing information, ensuring only devices with the correct authentication settings can participate.

OSPF resiliency and graceful restart

Explains OSPF resiliency and graceful restart mechanisms for Cisco 8000 Series Routers running IOS XR, covering behaviors, dependencies, and design considerations.

OSPF resiliency and graceful restart is a routing protocol enhancement that

- maintains neighbor relationships during failures or process restarts,
- minimizes downtime and routing disruption through mechanisms such as graceful restart and nonstop routing, and
- provides operational modes and configuration parameters to support high availability on Cisco 8000 Series Routers.

OSPF strict-mode support for BFD dampening:

Strict-mode is an OSPF BFD operation mode that keeps the neighbor in a down state until the BFD session is up. The neighbor status displays as "awaiting BFD session up" in the `show ospf neighbor` command. This ensures that client protocols do not operate independently of the declared BFD state.

Additional reference information

Multicast-intact feature

Attributes of multicast-intact next hops

OSPFv3 warm standby and nonstop routing

Protocol restoration and shutdown procedures

Protocol shutdown mode

Helper and restart modes

OSPFv3 Graceful Shutdown feature

Nonstop Forwarding (NSF) for OSPFv2

Prerequisites and requirements for OSPFv3 graceful restart

- Routers' neighbors must cooperate as helpers during a graceful restart.
- All neighbors must support graceful restart.
- Graceful restart does not occur at initial router startup.
- Simultaneous graceful restart sessions on multiple routers in a single segment are not supported.

Behavioral notes and operational details

- OSPFv3 neighbor and database information are not checkpointed.
- OSPFv3 reestablishes adjacencies after a restart and requires consistent pre- and post-restart configurations.
- IPv6 FIB tables remain initially unchanged, but routes are eventually marked stale for convergence.
- OSPFv3 must send hellos within the dead interval after restart (default is 40 seconds), or adjacencies will drop.
- The feature manages OSPFv3 route purge timers and associated grace link-scope LSAs.
- Adjacency creation times are stored for post-restart uptime reporting.

Restrictions

- Strict-mode and non-strict-mode are incompatible; mixing them prevents OSPF neighbor relationships from forming.
- Both BFD neighbors must run IOS XR images supporting strict-mode. Mixing strict and non-strict is not supported.
- If BFD sessions are established by other clients, a neighbor relationship may still form.
- The multicast-intact feature enables multicast routing (PIM) when IGP shortcuts are configured and active.
- Both OSPFv2 and IS-IS support multicast-intact.
- When enabled, IGP provides alternate multicast-intact next hops for use by PIM.
- These next hops do not include IGP shortcuts and are exclusively for PIM.
- They are not published to the FIB or used for unicast routing.
- All link-state learned IPv4 destinations are published with multicast-intact next hops to the RIB.

- OSPF applies the max-paths limit separately to native and multicast-intact next hops.
- Due to BFD dependency, OSPF in strict-mode may have delayed neighbor establishment and full adjacency.
- Warm standby enables a standby OSPF process on the backup Route Processor (RP) to initialize before failure, reducing switchover downtime.
- Warm standby is always enabled for OSPFv3, supporting high availability.
- Nonstop routing (NSR) makes process restarts or upgrades invisible to peer routers, eliminating the need for traditional graceful restart protocol extensions.
- To restore OSPFv3 after a shutdown triggered by the `protocol shutdown` command, use `no protocol shutdown`. If shutdown was due to critical memory, wait for memory recovery and manually restart OSPFv3.
- Upon restoration, OSPFv3 brings up all interfaces with hello packets, rebuilds and advertises local LSAs, responds to neighbor messages, and installs learned routes in the RIB.

Events on OSPFv3 restoration:

1. All OSPFv3 interfaces are activated and exchange database information.
2. Local router and link LSAs are rebuilt and advertised.
3. The router replies to incoming OSPFv3 control messages.
4. Learned routes are installed in the RIB.

In shutdown mode, OSPFv3 disables itself by flushing self-originated LSAs, taking down interfaces, and clearing the LSDB. Shutdown may be manual (`protocol shutdown`) or due to memory exhaustion.

- All local LSAs are flushed.
- Incomplete adjacencies are dropped.
- After a delay, empty hello packets are sent, ceasing after the dead interval.
- The LSDB and OSPFv3 RIB routes are purged after the dead interval.
- The router ignores OSPF packets while shut down.

Helper Mode: Enabled by default. A router enters helper mode when it receives a grace LSA Type 11 from a neighbor in graceful restart.

- If helper mode is disabled with `graceful-restart helper disable`, LSA packets are dropped.
- The helper function operates for a set duration based on the grace LSA, stopping upon successful restart or expiration.
- The dead timer is ignored in helper mode.
- The helper role ends upon reestablishing full adjacency or timer expiry.

Restart Mode: OSPFv3 checks if a restart is needed on process start, based on prior enablement (not on first boot). When enabled, it adjusts the RIB purge timer.

- During restart, routes are not populated in the RIB; full adjacencies from before are restored.
- Repeated restarts are prevented by setting a minimum interval with `graceful-restart interval`.
- Restart duration can be set with `graceful-restart lifetime`; a grace LSA Type 11 is sent for neighbor notification and helper mode entry.
- Preserves data plane capability during:

  - RP failure and switchover
  - Planned or unplanned process restarts (such as upgrades or process crashes)
  - Critical memory events triggering autonomous shutdown
- Enables nonstop forwarding on established routes during OSPFv3 restart, improving IPv6 forwarding high availability.
- NSF enables OSPF routers to continue forwarding packets during process restarts or RP failover.
- Peer devices do not experience routing flaps.
- RPs synchronize FIBs and maintain neighbor and link-state information through restart events.
- After restart, neighbor relationships and OSPF databases are rebuilt, stale routes are cleared, and the RIB/FIB are updated once full synchronization completes.

Enable BFD strict mode for OSPF

Enable BFD strict-mode for an OSPF interface so that the OSPF neighbor remains down until the BFD session is up.

Strict-mode is a BFD operation mode in OSPF where a neighbor remains down until the BFD session is up. When strict-mode is active, the neighbor state can display as `awaiting BFD session up` in the `show ospf neighbor` command output.

Strict-mode and non-strict-mode are incompatible. Both BFD neighbors must run IOS XR images that support strict-mode and must use strict-mode to form an adjacency.

Use this task to configure BFD strict-mode for OSPF on an interface and verify that the OSPF interface reports strict-mode operation.

Before you begin

- Ensure at least one IP address is configured on the router (OSPF routing requires at least one active IP address).
- Identify the OSPF process number, area IDs, and interfaces for OSPF participation.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. Enter global configuration mode.

   Example

```
Router# configure
Router(config)# router ospf 0
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface GigabitEthernet0/3/0/1
Router(config-ospf-ar-if)# bfd fast-detect strict-mode
Router(config-ospf-ar-if)# commit
Router(config-ospf-ar-if)# end
Router# show ospf interface GigabitEthernet0/3/0/1
```

2. Verify the OSPF interface information.

Example

```
Router# show ospf interface gigabitEthernet 0/3/0/1

GigabitEthernet0/3/0/1 is up, line protocol is up
  Internet Address 10.1.1.2/24, Area 0
  Process ID 1, Router ID 10.2.2.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1, MTU 1500, MaxPktSz 1500
  BFD enabled, BFD interval 150 msec, BFD multiplier 3, Mode: Strict
  Designated Router (ID) 10.2.2.2, Interface address 10.1.1.2
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:07:358
  Index 1/1, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 1, maximum is 1
  Last flood scan time is 0 msec, maximum is 0 msec
  LS Ack List: current length 0, high water mark 1
  Neighbor Count is 1, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
  Multi-area interface Count is 0
```

3. Verify the output of the `show ospf neighbor` command. **#** in the output indicates that the neighbor is waiting for the BFD session to come up.

Example

```
Router# show ospf neighbor

Neighbors for OSPF 1

Neighbor ID     Pri   State           Dead Time   Address         Interface
10.1.1.1         0     DOWN/DROTHER    00:00:33    10.1.1.3/24
GigabitEthernet0/3/0/1#
Total neighbor count: 1
```

When BFD strict-mode is enabled, the value of `Mode` displays as `Strict`. By default, the value of `Mode` displays as `Default`.

BFD strict-mode is enabled for the OSPF interface.

Configure OSPF restart and nonstop routing features

Configure OSPF restart and nonstop routing features through a focused workflow that preserves the original OSPF command sequence and result checks. Use this task when you need repeatable configuration steps.

Provide a structured and repeatable workflow for configuring OSPF restart and nonstop routing features, ensuring that the original OSPF command sequence is preserved and results can be verified reliably. Use this task when consistent and accurate configuration steps are required.

Procedure

---

1. Configure nonstop routing for OSPFv2.

Example

```
Router# configure
Router(config)# router ospf isp
Router(config-ospf)# nsr
Router(config-ospf)# commit
```

Enters mode.

2. Configure nonstop routing for OSPFv3

Example

```
Router# configure
Router(config)# router ospfv3 isp
Router(config-ospf)# nsr
Router(config-ospf)# commit
```

Enables OSPF routing for the specified routing process, and places the router in router configuration mode. In this example, the OSPF instance is called isp.

3. Configure a graceful restart for an OSPFv3 process.

Example

```
Router# configure
Router(config)# router ospfv3 test
Router(config-ospfv3)# graceful-restart
Router(config-ospfv3)# graceful-restart lifetime 120
Router(config-ospfv3)# graceful-restart interval 120
Router(config-ospfv3)# graceful-restart helper disable
Router# show ospfv3 1 database grace
```

4. Verify the state of the graceful restart capability on the local router:

Example

```
Router# show ospfv3 1

Routing Process “ospfv3 1” with ID 198.51.100.1
Initial SPF schedule delay 5000 msecs
Minimum hold time between two consecutive SPFs 10000 msecs
Maximum wait time between two consecutive SPFs 10000 msecs
Initial LSA throttle delay 0 msecs
Minimum hold time for LSA throttle 5000 msecs
Maximum wait time for LSA throttle 5000 msecs
Minimum LSA arrival 1000 msecs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Maximum number of configured interfaces 255
Number of external LSA 0. Checksum Sum 00000000
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Graceful Restart enabled, last GR 11:12:26 ago (took 6 secs)
Area BACKBONE(0)
Number of interfaces in this area is 1
SPF algorithm executed 1 times
Number of LSA 6. Checksum Sum 0x0268a7
Number of DCbitless LSA 0
```

```
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
```

5. Verify the link state for an OSPFv3 instance.

   Example

```
Router# show ospfv3 1 database grace

OSPFv3 Router with ID (192.0.2.1) (Process ID 1)
Grace (Type-11) Link States (Area 0)

  LS age: 2
  LS Type: Grace Links
  Link State ID: 34
  Advertising Router: 192.0.2.1
  LS Seq Number: 80000001
  Checksum: 0x7a4a
  Length: 36
      Grace Period : 90
      Graceful Restart Reason : Software reload/upgrade
```

6. Configure OSPF NSF specific to Cisco on your NSF-capable router.

   Note

   OSPF Cisco NSF for virtual links is not supported.

   OSPF NSF requires that all neighbor networking devices be NSF aware, which happens automatically after you install the Cisco IOS XR software image on the router. If an NSF-capable router discovers that it has non-NSF-aware neighbors on a particular network segment, it disables NSF capabilities for that segment. Other network segments composed entirely of NSF-capable or NSF-aware routers continue to provide NSF capabilities.

   Neighbors must be NSF aware.

   Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# nsf cisco enforce global
Router(config-ospf)# nsf interval 120
Router(config-ospf)# nsf flush-delay-time 1000
Router(config-ospf)# nsf lifetime 90
Router(config-ospf)# nsf ietf
Router# commit
```

The OSPF configuration is committed or verified, and the router reflects the intended operational state.

What to do next

Continue with related OSPF verification or tuning tasks that match your deployment workflow.

5 OSPF optimization scale and fast-reroute

Topics:

- OSPF scale, optimization, and protection features
- Configure OSPF SPF scheduling and prefix prioritization
- Configure OSPF scale and topology controls
- IP fast reroute loop-free alternate

Describes how SPF scheduling, prefix prioritization, topology controls, route summarization, resource-protection mechanisms, and loop-free alternates improve OSPF scalability, convergence, and traffic protection during topology changes and link failures.

OSPF scale, optimization, and protection features

Describes OSPF scaling, optimization, and protection features for Cisco 8000 Series Routers running IOS XR, including OSPF behavior, dependencies, and key design considerations.

OSPF scale, optimization, and protection features are advanced routing mechanisms that

- improve the scalability and convergence of OSPF on Cisco 8000 Series Routers,
- optimize route installation and prefix handling for efficient operations, and
- protect the network from resource exhaustion caused by excessive routing or link-state advertisements.

OSPF SPF Prefix Prioritization enables administrators to prioritize the installation of important prefixes, ensuring faster convergence for critical traffic (such as VoIP), especially in networks with a large number of routes. Prefixes are assigned to priority queues ("critical," "high," "medium," "low") based on route policy. This mechanism helps time-sensitive routes converge quickly.

Usage guidelines and limitations

- The features support only IPv4 addresses.
- To use multi-area interfaces, associate the loopback interface with an area of an ABR and reuse it as a multi-area interface elsewhere.
- These features are limited to a single OSPF instance and a single VRF. Multiple instances or VRFs are not supported.

Feature introduction and release summary
Table 2: Reference Table

Feature Name	Release	Description
Limiting the Maximum Redistributed Type-3 LSA Prefixes in OSPF	Release 25.1.1	Introduced on select fixed and modular systems; supported on 8712-MOD-M, 8011-4G24Y4H-I.
Limiting the Maximum Redistributed Type-3 LSA Prefixes in OSPF	Release 24.4.1	Available on fixed (8200/8700) and modular (8800) systems; supported across additional models such as 8212-48FH-M, 8711-32FH-M, etc.
Limiting the Maximum Redistributed Type-3 LSA Prefixes in OSPF	Release 7.9.1	Default limit set to 100,000 per process. Halts further redistribution and generates logs when exceeded.
Limiting LSA numbers in a OSPF Link-State Database	Release 25.1.1	Enabled by default from 7.9.1. Limits nonself-generated LSAs to 500,000 per process; supported on key platforms.
Autonomous System Boundary Router Isolation and Adjacency Control	Release 25.1.1	Provides protection by isolating ASBRs in LSA overload scenarios.
Autonomous System Boundary Router Isolation and Adjacency Control	Release 7.10.1	Prevents traffic disruption during LSA overload by isolating ASBRs and controlling adjacency recovery.

If only a high-priority route policy is configured, permitted prefixes are assigned to the high queue; unmatched prefixes (including /32s) are placed in the low-priority queue.

Example for configuring all /32s as medium priority:

- **prefix-set ospf-medium-prefixes**
- **0.0.0.0/0 ge 32**
- **end-set**

Usage notes

- Always configure route policies thoughtfully to align network priorities with traffic needs.
- Peer routers must support OSPF protection features to prevent LSA leaks.
- Use prefix suppression to reduce unnecessary route advertisements and bolster convergence and security.

Additional key behaviors

- Overload protection: OSPF can disable adjacencies and log events if LSA or prefix limits are reached, preventing overall router instability.
- ASBR isolation: When LSA limits are exceeded, neighboring routers can automatically isolate a misbehaving ASBR, keeping the rest of the OSPF domain operational.
- Timers and thresholds: You can configure thresholds, ignore times, and recovery times for granular control over protection mechanisms.

Configure OSPF SPF scheduling and prefix prioritization

Configure OSPF SPF scheduling and prefix prioritization to optimize route convergence.

Use this task to adjust OSPFv2 and OSPFv3 SPF calculation intervals and prioritize specific prefixes when calculating routes, reducing convergence times during network changes.

Optimize OSPF convergence and ensure critical prefixes receive higher priority during SPF calculations.

Before you begin

- Ensure you have configuration privileges.
- Verify that OSPF addressing, routing policy, and interface planning are complete.

Follow these steps to configure and verify OSPF areas and neighbours:

Procedure

1. Configure OSPF SPF scheduling intervals.
   - Enter global configuration mode and OSPF process configuration.
   - Set the OSPF router ID.
   - Configure SPF throttle timers to specify calculation delays.

   Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# router-id 192.168.4.3
```

```
Router(config-ospf)# timers throttle spf 10 4800 90000
Router(config-ospf)# area 0
Router(config-ospf-ar)# interface GigabitEthernet 0/1/0/3
```

2. Configure prefix prioritization for critical OSPF prefixes.
   - Define a prefix-set for critical prefixes.
   - Create a route-policy to assign high SPF priority to these prefixes.
   - Apply the route-policy to OSPF SPF prefix prioritization.

   Example

```
Router# configure
Router(config)# prefix-set ospf-critical-prefixes
Router(config-pfx)# 66.0.0.0/16
Router(config-pfx)# end-set
Router# route-policy ospf-spf-priority
Router(config-rpl)# if destination in ospf-critical-prefixes then
   set spf-priority critical
endif
Router(config-rpl)# end-policy
Router# router ospf 1
Router(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
Router(config-ospf)# commit
```

3. Verify SPF prefix prioritization configuration.

   Check the prefix-set and route-policy configuration reflects the intended critical prefixes.

   Example

```
Router# show rpl route-policy ospf-spf-priority detail
prefix-set ospf-critical-prefixes
    66.0.0.0/16
end-set
!
route-policy ospf-spf-priority
    if destination in ospf-critical-prefixes then
    set spf-priority critical
    endif
end-policy
!
```

The OSPF configuration is committed and the router operational state reflects optimized SPF scheduling and prioritized critical prefixes.

What to do next

Continue with related OSPF verification or perform further tuning tasks as needed for your deployment workflow.

Configure OSPF scale and topology controls

Set up multiple OSPF areas on a primary interface and summarize routes to reduce LSA count and optimize network resource usage.

When subnetworks are configured on different interfaces, OSPF can summarize these into a single LSA, reducing LSA numbers and conserving network resources. Interarea route summarization applies only to internal routes within your autonomous system and must be performed on the ABR. Summarization does not affect external routes injected via redistribution. Multi-area adjacency is supported where only two OSPF speakers are attached; for broadcast networks, configure the interface as OSPF point-to-point.

Enforce secure OSPF routing protocol operations by implementing authentication at the process, area, and interface levels to prevent unauthorized routing updates.

Before you begin

Ensure the interface used for multi-area adjacency has only two OSPF speakers attached.

For native broadcast networks, configure the interface as point-to-point using the **network point-to-point** command.

Procedure

1. Configure the primary loopback interface IPv4 address.

   Example

   ```
   Router(config)# interface Loopback0
   Router(config-if)# ipv4 address 192.0.2.2/32
   Router(config-if)# commit
   ```

2. Associate the loopback interface with the first OSPF area.

   Example

   ```
   Router(config)# router ospf 100
   Router(config-ospf)# router-id 10.10.10.1
   Router(config-ospf)# area 1
   Router(config-ospf-ar)# interface Loopback0
   Router(config-ospf-ar-if)# commit
   ```

3. Reuse the same loopback interface in another OSPF area using the **multi-area-interface** command.

   Example

   ```
   Router(config)# router ospf 100
   Router(config-ospf)# router-id 10.10.10.1
   Router(config-ospf)# area 0
   Router(config-ospf-ar)# multi-area-interface Loopback0
   Router(config-ospf-ar-mif)# commit
   ```

4. Verify the configuration using **show running-config**.

Example

```
Router# show running-config
!
interface Loopback0
ipv4 address 192.0.2.2/32
!
router ospf 100
router-id 10.10.10.1
area 0
multi-area-interface Loopback0
area 1
interface Loopback0
!
```

5. Check how the ABR identifies Loopback0 in area 1 using **show ospf routes**.

Example

```
Router# show ospf routes 192.0.2.2/32 backup-path detail
...
Route type: Intra-area
10.10.10.1, directly connected, via Loopback0
LSA: 1/10.10.10.1/10.10.10.1, Area: 1
```

You can verify configuration changes at the RI LSA level.

6. After enabling Loopback0 for area 1, verify that the summary LSA is not advertised by the ABR in area 0 with **show ospf database summary**.

Example

```
Router# show ospf database summary 10.10.10.1
...
Summary Net Link States (Area 0)
Link State ID: 10.10.10.1
Advertising Router: 10.0.0.1
```

7. Verify the RI LSA for area 0 with **show ospf database opaque-area**. Confirm the multi-area loopback interface is enabled.

Example

```
Router# show ospf database opaque-area ext-prefix 192.0.2.2/32
...
Type-10 Opaque Link Area Link States (Area 0)
Route-type: 1
Prefix: 10.10.10.1/32
SID sub-TLV: SID Index: 4
SID sub-TLV: SID Index: 2004
```

The OSPF configuration is complete and verified. The router reflects the intended operational state with summarized LSAs and correct area interface assignments.

IP fast reroute loop-free alternate

Describes IP fast reroute loop-free alternate for Cisco 8000 Series Routers running IOS XR. Use this topic to understand the OSPF behavior, dependencies, and design considerations before you configure or verify the feature.

IP fast reroute loop-free alternate is a concept that explains OSPF behavior, dependencies, and operational context for Cisco 8000 Series Routers.

This task describes how to enable the IP fast reroute (IPFRR) per-link loop-free alternate (LFA) computation to converge traffic flows around link failures.

This task describes how to enable the IP fast reroute per-link loop-free alternate computation so that traffic can converge around link failures.

To enable protection on broadcast links, IP fast reroute and Bidirectional Forwarding Detection (BFD) must be enabled on the interface under OSPF.

6 OSPF advanced integrations

Topics:

- OSPF virtual links and sham links
- OSPF integrations and management support
- Configure LDP IGP synchronization for OSPF
- Configure OSPF Transport and Service Integrations
- Configure OSPF virtual and sham links

Explains how virtual and sham links maintain OSPF connectivity and optimize MPLS VPN routing, and describes integrations with LDP, MPLS Traffic Engineering, VRFs, route redistribution, management interfaces, and security mechanisms.

OSPF virtual links and sham links

Describes OSPF virtual links and sham links for Cisco 8000 Series Routers running IOS XR. Use this topic to understand the OSPF behavior, dependencies, and design considerations before you configure or verify the feature.

OSPF virtual links and sham links is a concept that explains OSPF behavior, dependencies, and operational context for Cisco 8000 Series Routers.

Virtual Link and Transit Area for OSPF

In OSPF, routing information from all areas is first summarized to the backbone area by ABRs. The same ABRs, in turn, propagate such received information to their attached areas. Such hierarchical distribution of routing information requires that all areas be connected to the backbone area (Area 0). Occasions might exist for which an area must be defined, but it cannot be physically connected to Area 0. Examples of such an occasion might be if your company makes a new acquisition that includes an OSPF area, or if Area 0 itself is partitioned.

In the case in which an area cannot be connected to Area 0, you must configure a virtual link between that area and Area 0. The two endpoints of a virtual link are ABRs, and the virtual link must be configured in both routers. The common nonbackbone area to which the two routers belong is called a transit area. A virtual link specifies the transit area and the router ID of the other virtual endpoint (the other ABR).

A virtual link cannot be configured through a stub area or NSSA.

This figure illustrates a virtual link from Area 3 to Area 0.

Equal Cost Multipath (ECMP) mechanism is used to load-balance traffic on the Sham-link if there are multiple iBGP path for a prefix. If the sham link path and the backdoor path have the same cost, ECMP between the sham link path and backdoor path is not supported.

OSPFv3 treats the sham link as any other interface during the switch-over or process restart. OSPFv3 assumes that all the configured sham links are UP and tries to form an adjacency over them.

If the sham link is down prior to the switch-over, OSPFv3 sends the Hello packets to the remote endpoint. Once the final convergence signal is received from the RIB, OSPFv3 keeps the sham link either up or down based on the BGP route for each configured sham link in the RIB.

OSPFv3 installs the high AD routes over the sham link only after the BGP convergence is complete.

OSPFv3 sham link represents the VPN backbone as a single point-to-point connection between the two PEs. OSPFv3 treats the sham link as a point-to-point unnumbered interface, similar to virtual-link. When OSPFv3 sham link is configured, ensure that the route to the remote endpoint of the sham-link exists in the VRF RIB.

If the route to the remote endpoint exists, sham link interface is brought up. If the route to the remote endpoint of the sham-link is removed from the VRF RIB, OSPFv3 receives redistribution callback and brings the sham link down.

In an MPLS VPN environment, several VPN client sites can be connected in the same OSPF area. If these sites are connected over a backdoor link (intra-area link) and connected over the VPN backbone, all traffic passes over the backdoor link instead of over the VPN backbone, because provider edge routers advertise OSPF routes learned over the VPN backbone as inter-area or external routes that are less preferred than intra-area routes advertised over backdoor links.

To correct this default OSPF behavior in an MPLS VPN, configure a sham link between two provider edge (PE) routers to connect the sites through the MPLS VPN backbone. A sham link represents an intra-area (unnumbered point-to-point) connection between PE routers. All other routers in the area see the sham link and use it to calculate intra-area shortest path first (SPF) routes to the remote site. A cost must be configured with each sham link to determine whether traffic is sent over the backdoor link or sham link.

Configured source and destination addresses serve as the endpoints of the sham link. The source and destination IP addresses must belong to the VRF and must be advertised by Border Gateway Protocol (BGP) as host routes to remote PE routers. The sham-link endpoint addresses should not be advertised by OSPF.

For example, the figure above shows three client sites, each with backdoor links. Because each site runs OSPF within Area 1 configuration, all routing between the sites follows the intra-area path across the backdoor links instead of over the MPLS VPN backbone.

If the backdoor links between the sites are used only for backup purposes, default route selection over the backbone link is not acceptable as it creates undesirable traffic flow. To establish the desired path selection over the MPLS backbone, an additional OSPF intra-area (sham link) link between the ingress and egress PErouters must be created.

A sham link is required between any two VPN sites that belong to the same OSPF area and share an OSPF backdoor link. If no backdoor link exists between sites, no sham link is required.

The figure above shows an MPLS VPN topology where a sham link configuration is necessary. A VPN client has three sites, each with a backdoor link. Two sham links are configured, one between PE-1 and PE-2 and another between PE-2 and PE-3. A sham link is not required between PE-1 and PE-3, because there is no backdoor link between these sites.

When a sham link is configured between the PE routers, the PE routers can populate the virtual routing and forwarding (VRF) table with the OSPF routes learned over the sham link. These OSPF routes have a larger administrative distance than BGP routes. If BGP routes are available, they are preferred over these OSPF routes with the high administrative distance.

OSPF integrations and management support

Explains OSPF integrations and management support options for Cisco 8000 Series Routers running IOS XR.

An OSPF integration and management support feature set is a routing solution that:

- enables interoperability with multiple routing protocols through route redistribution,
- supports robust monitoring and management via Management Information Base (MIB) standards, and
- provides additional capabilities such as VRF-lite, Path Computation Element (PCE), security mechanisms, and automated LDP configuration.

Route Redistribution for OSPF: Redistribution allows different routing protocols to exchange routing information, enabling connectivity to span multiple routing protocols. The redistribute command controls redistribution into an OSPF process but not from OSPF.

VRF-lite capability in OSPFv2: OSPFv2 supports VRF-lite, which enables virtual routing and forwarding without using a BGP/MPLS backbone. In VRF-lite deployments, provider edge (PE) routers use VRF interfaces for direct connectivity. Configure VRF-lite capability for OSPFv2 using the `capability vrf-lite` command in VRF configuration mode. When VRF-lite is configured, DN bit processing and automatic Area Border Router (ABR) status setting are disabled.

Additional reference information

- OSPF MIB support: Cisco IOS XR supports full MIBs and traps for OSPFv2 (RFC 4750) and OSPFv3 (RFC 5643). OSPFv2 and OSPFv3 MIB implementations follow relevant IETF drafts and RFCs. For more on MIB details, use the MIB Locator and review the Cisco ASR 9000 Series Router MIB Specification Guide.
- Multiple OSPFv3 instances: SNMPv3 supports "contexts" to create MIB views for multiple OSPFv3 instances within the same system.
- Path Computation Element (PCE) integration: A Path Computation Element (PCE) is an entity (component, application, or network node) that calculates network paths based on a network graph and constraints. PCE integration in OSPF is accomplished by configuring a PCE address and client for MPLS-TE. OSPF uses the PCE Discovery TLV in the Router Information LSA (RI LSA) to communicate PCE capabilities. Relevant IETF drafts include "draft-ietf-ospf-cap-09" and "draft-ietf-pce-disco-proto-ospf-00". The PCE Address Sub-TLV must specify a reachable (loopback) IP address. The Path Scope Sub-TLV indicates the scope of path computation.
- OSPF security with GTSM: OSPF, as a link state protocol, requires fast convergence and the flooding of LSA updates. The Generalized TTL Security Mechanism (GTSM, RFC 3682) protects OSPF adjacencies by requiring that packets sent over point-to-point links use a TTL value of 255. This approach filters out OSPF packets that originate from more than one hop away, preventing certain types of network attacks.

- LDP IGP auto-configuration: LDP IGP auto-configuration streamlines the process to enable LDP on multiple interfaces used by an IGP instance such as OSPF. The feature supports the IPv4 unicast address family for the default VRF. You can disable auto-configuration for specific interfaces using the `igp auto-config disable` command.

Configure LDP IGP synchronization for OSPF

Configure Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) synchronization for OSPF to prevent traffic loss and improve network stability during network changes.

Configure Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) synchronization for OSPF to prevent traffic loss and improve network stability during network changes.

By default, there is no synchronization between LDP and IGPs. Without synchronization, OSPF may forward traffic before LDP label exchange is complete, which can cause blackholing. LDP IGP synchronization ensures OSPF waits for label distribution, resulting in a more stable network.

Before you begin

- Log in to the router CLI with appropriate privileges.
- Verify OSPF is already configured on the device.

Follow these steps to configure and verify LDP IGP synchronization for OSPF:

Procedure

Configure LDP IGP Synchronization under OSPF:

Example

```
Router# configure terminal
Router(config)# router ospf 100
Router(config-ospf)# mpls ldp sync
Router(config-ospf)# area 1 mpls ldp sync
Router(config-ospf)# area 1 interface GigabitEthernet0/4/0/2 mpls ldp sync
```

LDP IGP synchronization is now active for the selected OSPF scope(s). You can verify configuration with MPLS LDP IGP synchronization show commands.

Configure OSPF Transport and Service Integrations

Provides configure procdures for enabling OSPF Version 2 on Cisco 8000 Series Routers, integrating advanced features such as MPLS Traffic Engineering, LDP auto-configuration, and VRF-based OSPF instances using a repeatable workflow. This process enables scalable, service-ready routing in complex networks.

Set up OSPF Version 2 with advanced integrations for MPLS Traffic Engineering, LDP auto-configuration, and VRF-based OSPF instances on Cisco 8000 Series Routers.

Use this task when deploying or expanding Cisco 8000 Series Routers to support enhanced OSPF capabilities. These integrations ensure optimal scalability and readiness for managed services and MPLS environments.

Before you begin

- Ensure that MPLS and IP Cisco Express Forwarding (CEF) are enabled on your router.

- Verify that your network supports MPLS Traffic Engineering for OSPF operations.

Follow these steps to configure and verify OSPF transport and service integrations:

Procedure

1. Configure OSPF Version 2 for MPLS Traffic Engineering.

   Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# router-id 192.168.4.3
Router(config-ospf)# mpls traffic-eng router-id loopback 0
Router(config-ospf)# area 0
Router(config-ospf)# mpls traffic-eng
Router(config-ospf-ar)# interface loopback0
```

2. Enable Label Distribution Protocol (LDP) IGP auto-configuration for OSPF.

   Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# mpls ldp auto-config
```

3. Associate an interface with a VPN Routing and Forwarding (VRF) instance in OSPF.

   Example

```
Router# configure
Router(config)# router ospf 1
Router(config-ospf)# vrf vrf1
Router(config-ospf-vrf)# area 0
Router(config-ospf-vrf-ar)# interface GigabitEthernet0/0/0/0
```

Your Cisco 8000 Series Router is now running OSPF Version 2, integrated with MPLS Traffic Engineering, LDP auto-configuration, and VRF-aware OSPF. The router is ready for scalable, service-rich routing as per your network requirements.

What to do next

- Verify OSPF and MPLS configurations using `show ospf`, `show mpls ldp neighbors`, and confirm integration.
- Adjust OSPF areas, VRFs, or interfaces as required for your deployment.

Configure OSPF virtual and sham links

Provides configuration procedures for enabling OSPF virtual and sham links to maintain backbone connectivity and optimize routing on Cisco 8000 Series Routers.

This task provides the procedures required to maintain OSPF network integrity and routing efficiency through the implementation of virtual links and sham links.

- Virtual Links: Establishes a logical connection between an Area Border Router (ABR) and the backbone area (Area 0) in scenarios where a direct physical connection is not feasible. This ensures continuous reachability and maintains the OSPF area hierarchy.
- Sham Links: Configures an intra-area link between Provider Edge (PE) routers within an MPLS VPN environment. This ensures that OSPF treats the path across the MPLS backbone as an intra-area route, thereby preventing suboptimal routing and maintaining consistent OSPF adjacency for enterprise and service provider networks.

Perform this task on the CSC-PE to configure advanced L3VPN label allocation and BGP settings required for Carrier Supporting Carrier.

Before you begin

- Configure the sham link under OSPF area submode as needed.
- To create a virtual link with MD5 authentication to area 0:
  - Obtain the neighbor router's router ID by running `show ospf` or `show ospfv3` on the remote router.
  - Set a stable router ID at each end. Avoid default-assigned router IDs to prevent link disruption.
  - Assign router IDs manually using the `router-id` command.
  - Optionally, configure a loopback interface for persistent router IDs.

Follow these steps to configure and verify Carrier Supporting Carrier for L3VPN on CSC-PE:

Procedure

1. Configure a virtual link with MD5 authentication to area 0.

   Example

   ```
   Router# configure
   Router(config)# router ospf 1
   Router(config)# router ospfv3 1
   Router(config-ospf)# router-id [local-router-id]
   Router(config-ospf)# area 1
   Router(config-ospf-ar)# virtual-link [neighbor-router-id]
   Router(config-ospf-ar-vl)# authentication message-digest
   Router(config-ospf-ar-vl)# message-digest-key [key-id] md5 [yourkey]
   Router# show ospf 1 2 virtual-links
   Router# show ospfv3 1 virtual-links
   Router# show ospf
   Router# show ospfv3
   ```

   Enter global configuration mode and configure the OSPF or OSPFv3 instance. Manually assign a stable router ID. Specify the area (other than 0) and configure the virtual link using the neighbor router's router ID. Enable MD5 authentication and set the message-digest key. Optionally, run `show ospf` or `show ospfv3` commands to verify OSPF processes and router ID.

2. Configure an OSPFv2 sham link.

Example

```
Router# configure
Router(config)# interface loopback [number]
Router(config-if)# vrf [vrf-name]
Router(config-if)# ipv4 address [address] [mask]
Router(config-if)# exit
Router(config)# router ospf [process-name]
Router(config-ospf)# vrf [vrf-name]
Router(config-ospf-vrf)# router-id [router-id]
Router(config-ospf-vrf)# redistribute bgp [asn]
Router(config-ospf-vrf)# area 0
Router(config-ospf-vrf-ar)# sham-link [source-addr] [dest-addr]
Router(config-ospf-vrf-ar-sl)# cost [value]
Router(config-ospf-vrf-ar-sl)# commit
```

Configure a loopback interface within the appropriate VRF and assign an IP address. Enter OSPF router configuration in the desired VRF and assign a stable router ID. Redistribute BGP as required. Enter area 0, then configure the sham link between the source and destination addresses. Set the cost value for the sham link and commit your changes. Ensure the source and destination addresses belong to the VRF and are advertised by BGP as host routes to remote PE routers. Explicitly set the cost to override defaults.

Your OSPF configuration is committed and verified. The router uses the intended OSPF topology with both virtual and sham links.