



OSPF-IPv4

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) developed by the OSPF working group of the Internet Engineering Task Force (IETF). Designed expressly for IP networks, OSPF supports IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets. This chapter describes the concepts and tasks you need to configure OSPF on your Cisco NCS 4000 Series Router.

- [Prerequisites for Implementing OSPF](#) , on page 1
- [Information About Implementing OSPF](#) , on page 1
- [Information About Implementing OSPF](#) , on page 15
- [How to Implement OSPF](#) , on page 24

Prerequisites for Implementing OSPF

The following are prerequisites for implementing OSPF on Cisco IOS XR software:

- You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.k
- Configuring authentication (IP Security) is an optional task. If you choose to configure authentication, you must first decide whether to configure plain text or Message Digest 5 (MD5) authentication, and whether the authentication applies to an entire area or specific interfaces.

Information About Implementing OSPF

To implement OSPF you need to understand the following concepts:

OSPF Functional Overview

OSPF is a routing protocol for IP. It is a link-state protocol, as opposed to a distance-vector protocol. A link-state protocol makes its routing decisions based on the states of the links that connect source and destination machines. The state of the link is a description of that interface and its relationship to its neighboring networking devices. The interface information includes the IP address of the interface, network mask, type of network to which it is connected, routers connected to that network, and so on. This information is propagated in various types of link-state advertisements (LSAs).

A router stores the collection of received LSA data in a link-state database. This database includes LSA data for the links of the router. The contents of the database, when subjected to the Dijkstra algorithm, extract data to create an OSPF routing table. The difference between the database and the routing table is that the database contains a complete collection of raw data; the routing table contains a list of shortest paths to known destinations through specific router interface ports.

OSPF is the IGP of choice because it scales to large networks. It uses areas to partition the network into more manageable sizes and to introduce hierarchy in the network. A router is attached to one or more areas in a network. All of the networking devices in an area maintain the same complete database information about the link states in their area only. They do not know about all link states in the network. The agreement of the database information among the routers in the area is called convergence.

At the intradomain level, OSPF can import routes learned using Intermediate System-to-Intermediate System (IS-IS). OSPF routes can also be exported into IS-IS. At the interdomain level, OSPF can import routes learned using Border Gateway Protocol (BGP). OSPF routes can be exported into BGP.



Note Following are the number of routes supported in NCS 4000 :

- 32K with NCS4K-2H10T-OP-KS line card
- 32K with NCS4K-4H-OPW-QC2 line card

Unlike Routing Information Protocol (RIP), OSPF does not provide periodic routing updates. On becoming neighbors, OSPF routers establish an adjacency by exchanging and synchronizing their databases. After that, only changed routing information is propagated. Every router in an area advertises the costs and states of its links, sending this information in an LSA. This state information is sent to all OSPF neighbors one hop away. All the OSPF neighbors, in turn, send the state information unchanged. This flooding process continues until all devices in the area have the same link-state database.

To determine the best route to a destination, the software sums all of the costs of the links in a route to a destination. After each router has received routing information from the other networking devices, it runs the shortest path first (SPF) algorithm to calculate the best path to each destination network in the database.

The networking devices running OSPF detect topological changes in the network, flood link-state updates to neighbors, and quickly converge on a new view of the topology. Each OSPF router in the network soon has the same topological view again. OSPF allows multiple equal-cost paths to the same destination. Since all link-state information is flooded and used in the SPF calculation, multiple equal cost paths can be computed and used for routing.

On broadcast and non broadcast multiaccess (NBMA) networks, the designated router (DR) or backup DR performs the LSA flooding. On point-to-point networks, flooding simply exits an interface directly to a neighbor.

OSPF runs directly on top of IP; it does not use TCP or User Datagram Protocol (UDP). OSPF performs its own error correction by means of checksums in its packet header and LSAs.

OSPF typically requires coordination among many internal routers: Area Border Routers (ABRs), which are routers attached to multiple areas, and Autonomous System Border Routers (ASBRs) that export reroutes from other sources (for example, IS-IS, BGP, or static routes) into the OSPF topology. At a minimum, OSPF-based routers or access servers can be configured with all default parameter values, no authentication, and interfaces assigned to areas. If you intend to customize your environment, you must ensure coordinated configurations of all routers.

Key Features Supported in the Cisco IOS XR Software OSPF Implementation

The Cisco IOS XR Software implementation of OSPF conforms to the OSPF Version 2 and OSPF Version 3 specifications detailed in the Internet RFC 2328 and RFC 2740, respectively.

The following key features are supported in the Cisco IOS XR Software implementation:

- Hierarchy—CLI hierarchy is supported.
- Inheritance—CLI inheritance is supported.
- Stub areas—Definition of stub areas is supported.
- NSF—Nonstop forwarding is supported.
- SPF throttling—Shortest path first throttling feature is supported.
- LSA throttling—LSA throttling feature is supported.
- Fast convergence—SPF and LSA throttle timers are set, configuring fast convergence. The OSPF LSA throttling feature provides a dynamic mechanism to slow down LSA updates in OSPF during network instability. LSA throttling also allows faster OSPF convergence by providing LSA rate limiting in milliseconds.
- Route redistribution—Routes learned using any IP routing protocol can be redistributed into any other IP routing protocol.
- Authentication—Plain text and MD5 authentication among neighboring routers within an area is supported.
- Routing interface parameters—Configurable parameters supported include interface output cost, retransmission interval, interface transmit delay, router priority, router “dead” and hello intervals, and authentication key.
- Virtual links—Virtual links are supported.
- Not-so-stubby area (NSSA)—RFC 1587 is supported.
- OSPF over demand circuit—RFC 1793 is supported.

OSPF Hierarchical CLI and CLI Inheritance

Cisco IOS XR Software introduces new OSPF configuration fundamentals consisting of hierarchical CLI and CLI inheritance.

Hierarchical CLI is the grouping of related network component information at defined hierarchical levels such as at the router, area, and interface levels. Hierarchical CLI allows for easier configuration, maintenance, and troubleshooting of OSPF configurations. When configuration commands are displayed together in their hierarchical context, visual inspections are simplified. Hierarchical CLI is intrinsic for CLI inheritance to be supported.

With CLI inheritance support, you need not explicitly configure a parameter for an area or interface. In Cisco IOS XR Software, the parameters of interfaces in the same area can be exclusively configured with a single command, or parameter values can be inherited from a higher hierarchical level—such as from the area configuration level or the router ospf configuration levels.

For example, the hello interval value for an interface is determined by this precedence “IF” statement:

If the **hello interval** command is configured at the interface configuration level, then use the interface configured value, else

If the **hello interval** command is configured at the area configuration level, then use the area configured value, else

If the **hello interval** command is configured at the router ospf configuration level, then use the router ospf configured value, else

Use the default value of the command.



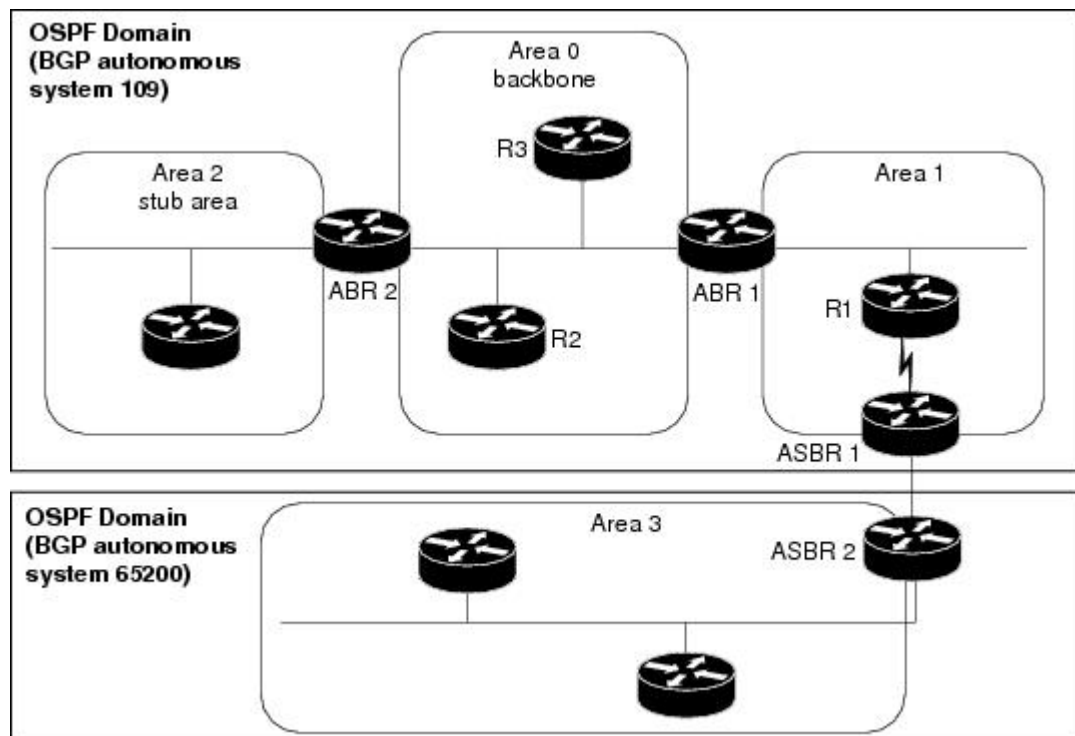
Tip Understanding hierarchical CLI and CLI inheritance saves you considerable configuration time. See [Configuring Authentication at Different Hierarchical Levels for OSPF Version 2, on page 31](#) to understand how to implement these fundamentals. In addition, Cisco IOS XR Software examples are provided in [Configuration Examples for Implementing OSPF , on page 63](#).

OSPF Routing Components

Before implementing OSPF, you must know what the routing components are and what purpose they serve. They consist of the autonomous system, area types, interior routers, ABRs, and ASBRs.

Figure 1: OSPF Routing Components

This figure illustrates the routing components in an OSPF network topology.



Autonomous Systems

The autonomous system is a collection of networks, under the same administrative control, that share routing information with each other. An autonomous system is also referred to as a routing domain. [Figure 1: OSPF Routing Components, on page 4](#) shows two autonomous systems: 109 and 65200. An autonomous system can consist of one or more OSPF areas.

Areas

Areas allow the subdivision of an autonomous system into smaller, more manageable networks or sets of adjacent networks. As shown in [Figure 1: OSPF Routing Components, on page 4](#), autonomous system 109 consists of three areas: Area 0, Area 1, and Area 2.

OSPF hides the topology of an area from the rest of the autonomous system. The network topology for an area is visible only to routers inside that area. When OSPF routing is within an area, it is called *intra-area routing*. This routing limits the amount of link-state information flood into the network, reducing routing traffic. It also reduces the size of the topology information in each router, conserving processing and memory requirements in each router.

Also, the routers within an area cannot see the detailed network topology outside the area. Because of this restricted view of topological information, you can control traffic flow between areas and reduce routing traffic when the entire autonomous system is a single routing domain.

Backbone Area

A backbone area is responsible for distributing routing information between multiple areas of an autonomous system. OSPF routing occurring outside of an area is called *interarea routing*.

The backbone itself has all properties of an area. It consists of ABRs, routers, and networks only on the backbone. As shown in [Figure 1: OSPF Routing Components, on page 4](#), Area 0 is an OSPF backbone area. Any OSPF backbone area has a reserved area ID of 0.0.0.0.

Stub Area

A stub area is an area that does not accept route advertisements or detailed network information external to the area. A stub area typically has only one router that interfaces the area to the rest of the autonomous system. The stub ABR advertises a single default route to external destinations into the stub area. Routers within a stub area use this route for destinations outside the area and the autonomous system. This relationship conserves LSA database space that would otherwise be used to store external LSAs flooded into the area. In [Figure 1: OSPF Routing Components, on page 4](#), Area 2 is a stub area that is reached only through ABR 2. Area 0 cannot be a stub area.

Not-so-Stubby Area

A Not-so-Stubby Area (NSSA) is similar to the stub area. NSSA does not flood Type 5 external LSAs from the core into the area, but can import autonomous system external routes in a limited fashion within the area.

NSSA allows importing of Type 7 autonomous system external routes within an NSSA area by redistribution. These Type 7 LSAs are translated into Type 5 LSAs by NSSA ABRs, which are flooded throughout the whole routing domain. Summarization and filtering are supported during the translation.

Use NSSA to simplify administration if you are a network administrator that must connect a central site using OSPF to a remote site that is using a different routing protocol.

Before NSSA, the connection between the corporate site border router and remote router could not be run as an OSPF stub area because routes for the remote site could not be redistributed into a stub area, and two

routing protocols needed to be maintained. A simple protocol like RIP was usually run and handled the redistribution. With NSSA, you can extend OSPF to cover the remote connection by defining the area between the corporate router and remote router as an NSSA. Area 0 cannot be an NSSA.

Routers

The OSPF network is composed of ABRs, ASBRs, and interior routers.

Area Border Routers

An area border routers (ABR) is a router with multiple interfaces that connect directly to networks in two or more areas. An ABR runs a separate copy of the OSPF algorithm and maintains separate routing data for each area that is attached to, including the backbone area. ABRs also send configuration summaries for their attached areas to the backbone area, which then distributes this information to other OSPF areas in the autonomous system. In [Figure 1: OSPF Routing Components, on page 4](#), there are two ABRs. ABR 1 interfaces Area 1 to the backbone area. ABR 2 interfaces the backbone Area 0 to Area 2, a stub area.

Autonomous System Boundary Routers (ASBR)

An autonomous system boundary router (ASBR) provides connectivity from one autonomous system to another system. ASBRs exchange their autonomous system routing information with boundary routers in other autonomous systems. Every router inside an autonomous system knows how to reach the boundary routers for its autonomous system.

ASBRs can import external routing information from other protocols like BGP and redistribute them as AS-external (ASE) Type 5 LSAs to the OSPF network. If the Cisco IOS XR router is an ASBR, you can configure it to advertise VIP addresses for content as autonomous system external routes. In this way, ASBRs flood information about external networks to routers within the OSPF network.

ASBR routes can be advertised as a Type 1 or Type 2 ASE. The difference between Type 1 and Type 2 is how the cost is calculated. For a Type 2 ASE, only the external cost (metric) is considered when multiple paths to the same destination are compared. For a Type 1 ASE, the combination of the external cost and cost to reach the ASBR is used. Type 2 external cost is the default and is always more costly than an OSPF route and used only if no OSPF route exists.

Interior Routers

An interior router (such as R1 in [Figure 1: OSPF Routing Components, on page 4](#)) is attached to one area (for example, all the interfaces reside in the same area).

OSPF Process and Router ID

An OSPF process is a logical routing entity running OSPF in a physical router. This logical routing entity should not be confused with the logical routing feature that allows a system administrator (known as the Cisco IOS XR Software Owner) to partition the physical box into separate routers.

A physical router can run multiple OSPF processes, although the only reason to do so would be to connect two or more OSPF domains. Each process has its own link-state database. The routes in the routing table are calculated from the link-state database. One OSPF process does not share routes with another OSPF process unless the routes are redistributed.

Each OSPF process is identified by a router ID. The router ID must be unique across the entire routing domain. OSPF obtains a router ID from the following sources, in order of decreasing preference:

- By default, when the OSPF process initializes, it checks if there is a router-id in the checkpointing database.
- The 32-bit numeric value specified by the OSPF router-id command in router configuration mode. (This value can be any 32-bit value. It is not restricted to the IPv4 addresses assigned to interfaces on this router, and need not be a routable IPv4 address.)
- The ITAL selected router-id.
- The primary IPv4 address of an interface over which this OSPF process is running. The first interface address in the OSPF interface is selected.

We recommend that the router ID be set by the **router-id** command in router configuration mode. Separate OSPF processes could share the same router ID, in which case they cannot reside in the same OSPF routing domain.

Supported OSPF Network Types

OSPF classifies different media into the following types of networks:

- NBMA networks
- Point-to-point networks (POS)
- Broadcast networks (Ten Gigabit Ethernet and Hundred Gigabit Ethernet)
- Point-to-multipoint

You can configure your Cisco IOS XR network as either a broadcast or an NBMA network.

Route Authentication Methods for OSPF

OSPF Version 2 supports two types of authentication: plain text authentication and MD5 authentication. By default, no authentication is enabled (referred to as null authentication in RFC 2178).

OSPF Version 3 supports all types of authentication except key rollover.

Plain Text Authentication

Plain text authentication (also known as Type 1 authentication) uses a password that travels on the physical medium and is easily visible to someone that does not have access permission and could use the password to infiltrate a network. Therefore, plain text authentication does not provide security. It might protect against a faulty implementation of OSPF or a misconfigured OSPF interface trying to send erroneous OSPF packets.

MD5 Authentication

MD5 authentication provides a means of security. No password travels on the physical medium. Instead, the router uses MD5 to produce a message digest of the OSPF packet plus the key, which is sent on the physical medium. Using MD5 authentication prevents a router from accepting unauthorized or deliberately malicious routing updates, which could compromise your network security by diverting your traffic.



Note MD5 authentication supports multiple keys, requiring that a key number be associated with a key.

See [OSPF Authentication Message Digest Management, on page 22](#).

Authentication Strategies

Authentication can be specified for an entire process or area, or on an interface or a virtual link. An interface or virtual link can be configured for only one type of authentication, not both. Authentication configured for an interface or virtual link overrides authentication configured for the area or process.

If you intend for all interfaces in an area to use the same type of authentication, you can configure fewer commands if you use the **authentication** command in the area configuration submode (and specify the **message-digest** keyword if you want the entire area to use MD5 authentication). This strategy requires fewer commands than specifying authentication for each interface.

Key Rollover

To support the changing of an MD5 key in an operational network without disrupting OSPF adjacencies (and hence the topology), a key rollover mechanism is supported. As a network administrator configures the new key into the multiple networking devices that communicate, some time exists when different devices are using both a new key and an old key. If an interface is configured with a new key, the software sends two copies of the same packet, each authenticated by the old key and new key. The software tracks which devices start using the new key, and the software stops sending duplicate packets after it detects that all of its neighbors are using the new key. The software then discards the old key. The network administrator must then remove the old key from each the configuration file of each router.

Neighbors and Adjacency for OSPF

Routers that share a segment (Layer 2 link between two interfaces) become neighbors on that segment. OSPF uses the hello protocol as a neighbor discovery and keep alive mechanism. The hello protocol involves receiving and periodically sending hello packets out each interface. The hello packets list all known OSPF neighbors on the interface. Routers become neighbors when they see themselves listed in the hello packet of the neighbor. After two routers are neighbors, they may proceed to exchange and synchronize their databases, which creates an adjacency. On broadcast and NBMA networks all neighboring routers have an adjacency.

Enabling strict-mode

The following procedure describes how to enable BFD strict-mode for Open Shortest Path First (OSPF) on an interface:

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0:hostname# configure	Enters global configuration mode.

	Command or Action	Purpose
Step 2	router ospf <i>process-name</i> Example: RP/0/RP0:hostname(config)# router ospf 1	Enters OSPF configuration mode, allowing you to configure the OSPF routing process. Use the show ospf command in XR EXEC mode to obtain the process-name for the current router.
Step 3	area <i>area-id</i> Example: RP/0/RP0:hostname(config-ospf) # area 0	Configures an Open Shortest Path First (OSPF) area. Replace <i>area-id</i> with the OSPF area identifier.
Step 4	interface <i>type interface-path-id</i> Example: RP/0/RP0:hostname(config-ospf-ar) # interface TenGigE 0/6/0/6.11	Enters interface configuration mode and specifies the interface name and notation <i>rack/slot/module/port</i> . The example indicates a Ten Gigabit Ethernet interface in modular services card slot 3.
Step 5	bfd fast-detect strict-mode Example: RP/0/RP0:hostname(config-ospf-ar-if) # bfd fast-detect strict-mode	Enables strict-mode to hold down neighbor session until BFD session is up.
Step 6	commit	Commits the changes to the running configuration.
Step 7	show ospf interface <i>type interface-path-id</i> Example: RP/0/RP0:hostname(config-ospf-ar-if) #show ospf interface 0/6/0/6.11	Verify that strict-mode is enabled on the appropriate interface.

BFD strict-mode: Example

The following example shows how to enable BFD strict-mode for OSPF on a Hundred Gigabit Ethernet interface and check the OSPF interface information. The value of **Mode** displays as **Strict** when BFD strict-mode is enabled. By default, the value of **Mode** displays as **Default**.

```
RP/0/RP0:hostname#configure
RP/0/RP0:hostname(config)#router ospf 0
RP/0/RP0:hostname(config-ospf)#area 0
RP/0/RP0:hostname(config-ospf-ar)#interface HundredGigE0/6/0/0.30
RP/0/RP0:hostname(config-ospf-ar-if)#bfd fast-detect strict-mode
RP/0/RP0:hostname(config-ospf-ar-if)#commit
RP/0/RP0:hostname(config-ospf-ar-if)#end
RP/0/RP0:hostname#show ospf interface HundredGigE0/6/0/0.30
```

```
HundredGigE0/6/0/0.30 is up, line protocol is up
  Internet Address 10.1.1.2/24, Area 0
  Process ID 1, Router ID 2.2.2.2, 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) 2.2.2.2, 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

```

The following example shows the output of the **show ospf neighbor** command. # indicates that the neighbor is waiting for the BFD session to come up.

```

RP/0/RP0:hostname#show ospf neighbor

Neighbors for OSPF 1

Neighbor ID      Pri   State           Dead Time   Address          Interface
1.1.1.1          0     DOWN/DROTHER    00:00:33    10.1.1.3/24     HundredGigE0/6/0/0.30#

Total neighbor count: 1

```

OSPF FIB Download Notification

OSPF FIB Download Notification feature minimizes the ingress traffic drop for a prolonged period of time after the line card reloads and this feature is enabled by default.

Open Shortest Path First (OSPF) registers with Routing Information Base (RIB) through Interface Table Attribute Library (ITAL) which keeps the interface down until all the routes are downloaded to Forwarding Information Base (FIB). OSPF gets the Interface Up notification when all the routes on the reloaded line card are downloaded through RIB/FIB.

RIB provides notification to registered clients when a:

- Node is lost.
- Node is created.
- Node's FIB upload is completed.

Designated Router (DR) for OSPF

On point-to-point and point-to-multipoint networks, the Cisco IOS XR software floods routing updates to immediate neighbors. No DR or backup DR (BDR) exists; all routing information is flooded to each router.

On broadcast or NBMA segments only, OSPF minimizes the amount of information being exchanged on a segment by choosing one router to be a DR and one router to be a BDR. Thus, the routers on the segment have a central point of contact for information exchange. Instead of each router exchanging routing updates with every other router on the segment, each router exchanges information with the DR and BDR. The DR and BDR relay the information to the other routers.

The software looks at the priority of the routers on the segment to determine which routers are the DR and BDR. The router with the highest priority is elected the DR. If there is a tie, then the router with the higher router ID takes precedence. After the DR is elected, the BDR is elected the same way. A router with a router priority set to zero is ineligible to become the DR or BDR.

Default Route for OSPF

Type 5 (ASE) LSAs are generated and flooded to all areas except stub areas. For the routers in a stub area to be able to route packets to destinations outside the stub area, a default route is injected by the ABR attached to the stub area.

The cost of the default route is 1 (default) or is determined by the value specified in the **default-cost** command.

Link-State Advertisement Types for OSPF Version 2

Each of the following LSA types has a different purpose:

- Router LSA (Type 1)—Describes the links that the router has within a single area, and the cost of each link. These LSAs are flooded within an area only. The LSA indicates if the router can compute paths based on quality of service (QoS), whether it is an ABR or ASBR, and if it is one end of a virtual link. Type 1 LSAs are also used to advertise stub networks.
- Network LSA (Type 2)—Describes the link state and cost information for all routers attached to a multiaccess network segment. This LSA lists all the routers that have interfaces attached to the network segment. It is the job of the designated router of a network segment to generate and track the contents of this LSA.
- Summary LSA for ABRs (Type 3)—Advertises internal networks to routers in other areas (interarea routes). Type 3 LSAs may represent a single network or a set of networks aggregated into one prefix. Only ABRs generate summary LSAs.
- Summary LSA for ASBRs (Type 4)—Advertises an ASBR and the cost to reach it. Routers that are trying to reach an external network use these advertisements to determine the best path to the next hop. ABRs generate Type 4 LSAs.
- Autonomous system external LSA (Type 5)—Redistributes routes from another autonomous system, usually from a different routing protocol into OSPF.
- Autonomous system external LSA (Type 7)—Provides for carrying external route information within an NSSA. Type 7 LSAs may be originated by and advertised throughout an NSSA. NSSAs do not receive or originate Type 5 LSAs. Type 7 LSAs are advertised only within a single NSSA. They are not flooded into the backbone area or into any other area by border routers.
- Intra-area-prefix LSAs (Type 9)—A router can originate multiple intra-area-prefix LSAs for every router or transit network, each with a unique link-state ID. The link-state ID for each intra-area-prefix LSA describes its association to either the router LSA or network LSA and contains prefixes for stub and transit networks.
- Area local scope (Type 10)—Opaque LSAs are not flooded past the borders of their associated area.
- Link-state (Type 11)—The LSA is flooded throughout the AS. The flooding scope of Type 11 LSAs are equivalent to the flooding scope of AS-external (Type 5) LSAs. Similar to Type 5 LSAs, the LSA is rejected if a Type 11 opaque LSA is received in a stub area from a neighboring router within the stub area. Type 11 opaque LSAs have these attributes:

- LSAs are flooded throughout all transit areas.
- LSAs are not flooded into stub areas from the backbone.
- LSAs are not originated by routers into their connected stub areas.

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.

Figure 2: Virtual Link to Area 0

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



Passive Interface

Setting an interface as passive disables the sending of routing updates for the neighbors, hence adjacencies will not be formed in OSPF. However, the particular subnet will continue to be advertised to OSPF neighbors.

Use the **passive** command in appropriate mode to suppress the sending of OSPF protocol operation on an interface.

It is recommended to use passive configuration on interfaces that are connecting LAN segments with hosts to the rest of the network, but are not meant to be transit links between routers.

OSPFv2 SPF Prefix Prioritization

The OSPFv2 SPF Prefix Prioritization feature enables an administrator to converge, in a faster mode, important prefixes during route installation.

When a large number of prefixes must be installed in the Routing Information Base (RIB) and the Forwarding Information Base (FIB), the update duration between the first and last prefix, during SPF, can be significant.

In networks where time-sensitive traffic (for example, VoIP) may transit to the same router along with other traffic flows, it is important to prioritize RIB and FIB updates during SPF for these time-sensitive prefixes.

The OSPFv2 SPF Prefix Prioritization feature provides the administrator with the ability to prioritize important prefixes to be installed, into the RIB during SPF calculations. Important prefixes converge faster among prefixes of the same route type per area. Before RIB and FIB installation, routes and prefixes are assigned to various priority batch queues in the OSPF local RIB, based on specified route policy. The RIB priority batch queues are classified as "critical," "high," "medium," and "low," in the order of decreasing priority.

When enabled, prefix alters the sequence of updating the RIB with this prefix priority:

Critical > High > Medium > Low

As soon as prefix priority is configured, /32 prefixes are no longer preferred by default; they are placed in the low-priority queue, if they are not matched with higher-priority policies. Route policies must be devised to retain /32s in the higher-priority queues (high-priority or medium-priority queues).

Priority is specified using route policy, which can be matched based on IP addresses or route tags. During SPF, a prefix is checked against the specified route policy and is assigned to the appropriate RIB batch priority queue.

These are examples of this scenario:

- If only high-priority route policy is specified, and no route policy is configured for a medium priority:
 - Permitted prefixes are assigned to a high-priority queue.
 - Unmatched prefixes, including /32s, are placed in a low-priority queue.
- If both high-priority and medium-priority route policies are specified, and no maps are specified for critical priority:
 - Permitted prefixes matching high-priority route policy are assigned to a high-priority queue.
 - Permitted prefixes matching medium-priority route policy are placed in a medium-priority queue.
 - Unmatched prefixes, including /32s, are moved to a low-priority queue.
- If both critical-priority and high-priority route policies are specified, and no maps are specified for medium priority:
 - Permitted prefixes matching critical-priority route policy are assigned to a critical-priority queue.
 - Permitted prefixes matching high-priority route policy are assigned to a high-priority queue.

- Unmatched prefixes, including /32s, are placed in a low-priority queue.
- If only medium-priority route policy is specified and no maps are specified for high priority or critical priority:
 - Permitted prefixes matching medium-priority route policy are assigned to a medium-priority queue.
 - Unmatched prefixes, including /32s, are placed in a low-priority queue.

Use the **[no] spf prefix-priority route-policy** *rpl* command to prioritize OSPFv2 prefix installation into the global RIB during SPF.

SPF prefix prioritization is disabled by default. In disabled mode, /32 prefixes are installed into the global RIB, before other prefixes. If SPF prioritization is enabled, routes are matched against the route-policy criteria and are assigned to the appropriate priority queue based on the SPF priority set. Unmatched prefixes, including /32s, are placed in the low-priority queue.

If all /32s are desired in the high-priority queue or medium-priority queue, configure this single route map:

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

Route Redistribution for OSPF

Redistribution allows different routing protocols to exchange routing information. This technique can be used to allow connectivity to span multiple routing protocols. It is important to remember that the **redistribute** command controls redistribution *into* an OSPF process and not from OSPF. See [Configuration Examples for Implementing OSPF](#), on page 63 for an example of route redistribution for OSPF.

OSPF Shortest Path First Throttling

OSPF SPF throttling makes it possible to configure SPF scheduling in millisecond intervals and to potentially delay SPF calculations during network instability. SPF is scheduled to calculate the Shortest Path Tree (SPT) when there is a change in topology. One SPF run may include multiple topology change events.

The interval at which the SPF calculations occur is chosen dynamically and based on the frequency of topology changes in the network. The chosen interval is within the boundary of the user-specified value ranges. If network topology is unstable, SPF throttling calculates SPF scheduling intervals to be longer until topology becomes stable.


SPF calculations occur at the interval set by the **timers throttle spf** command. The wait interval indicates the amount of time to wait until the next SPF calculation occurs. Each wait interval after that calculation is twice as long as the previous interval until the interval reaches the maximum wait time specified.

The SPF timing can be better explained using an example. In this example, the start interval is set at 5 milliseconds (ms), initial wait interval at 1000 ms, and maximum wait time at 90,000 ms.

```
timers spf 5 1000 90000
```

Figure 3: SPF Calculation Intervals Set by the timers spf Command

This figure shows the intervals at which the SPF calculations occur as long as at least one topology change event is received in a given wait interval.




Notice that the wait interval between SPF calculations doubles when at least one topology change event is received during the previous wait interval. After the maximum wait time is reached, the wait interval remains the same until the topology stabilizes and no event is received in that interval.

If the first topology change event is received after the current wait interval, the SPF calculation is delayed by the amount of time specified as the start interval. The subsequent wait intervals continue to follow the dynamic pattern.

If the first topology change event occurs after the maximum wait interval begins, the SPF calculation is again scheduled at the start interval and subsequent wait intervals are reset according to the parameters specified in the **timers throttle spf** command. Notice in [Figure 4: Timer Intervals Reset After Topology Change Event, on page 15](#) that a topology change event was received after the start of the maximum wait time interval and that the SPF intervals have been reset.

Figure 4: Timer Intervals Reset After Topology Change Event



Information About Implementing OSPF

To implement OSPF you need to understand the following concepts:

Warm Standby and Nonstop Routing for OSPF Version 2

OSPFv2 warm standby provides high availability across RP switchovers. With warm standby extensions, each process running on the active RP has a corresponding standby process started on the standby RP. A standby OSPF process can send and receive OSPF packets with no performance impact to the active OSPF process.

Nonstop routing (NSR) allows an RP failover, process restart, or in-service upgrade to be invisible to peer routers and ensures that there is minimal performance or processing impact. Routing protocol interactions between routers are not impacted by NSR. NSR is built on the warm standby extensions. NSR alleviates the requirement for Cisco NSF and IETF graceful restart protocol extensions.



Note It is recommended to set the hello timer interval to the default of 10 seconds. OSPF sessions may flap during switchover if hello-interval timer configured is less than default value.

Multicast-Intact Support for OSPF

The multicast-intact feature provides the ability to run multicast routing (PIM) when IGP shortcuts are configured and active on the router. Both OSPFv2 and IS-IS support the multicast-intact feature.

You can enable multicast-intact in the IGP when multicast routing protocols (PIM) are configured and IGP shortcuts are configured on the router. IGP shortcuts are MPLS tunnels that are exposed to IGP. The IGP routes IP traffic over these tunnels to destinations that are downstream from the egress router of the tunnel (from an SPF perspective). PIM cannot use IGP shortcuts for propagating PIM joins, because reverse path forwarding (RPF) cannot work across a unidirectional tunnel.

When you enable multicast-intact on an IGP, the IGP publishes a parallel or alternate set of equal-cost next hops for use by PIM. These next hops are called *mcast-intact* next hops. The mcast-intact next hops have the following attributes:

- They are guaranteed not to contain any IGP shortcuts.
- They are not used for unicast routing but are used only by PIM to look up an IPv4 next-hop to a PIM source.
- They are not published to the FIB.
- When multicast-intact is enabled on an IGP, all IPv4 destinations that were learned through link-state advertisements are published with a set equal-cost mcast-intact next hops to the RIB. This attribute applies even when the native next hops have no IGP shortcuts.

In OSPF, the max-paths (number of equal-cost next hops) limit is applied separately to the native and mcast-intact next hops. The number of equal cost mcast-intact next hops is the same as that configured for the native next hops.

Configure Prefix Suppression for OSPF

Transit-only networks that connect two routers are usually configured with routing IP addresses that are advertised in the Links State Advertisements (LSAs). However, these prefixes are not needed for data traffic. Suppressing these prefixes would reduce the number of links in LSAs, thereby improving convergence and also reducing the vulnerability of potential remote attacks.

Prefixes can be suppressed for an OSPF process, an OSPF area, or for specific interfaces of a router.

Configure Prefix Suppression for a Router Running OSPF

Use the procedure in this section to configure prefix suppression for an OSPF process on a router.

**Note**

- If you suppress prefixes for an OSPF process on a router, the suppression is valid for all interfaces and areas associated with the router.
- When prefix suppression is configured on an NSSA ASBR, all interfaces on the routers have their prefixes suppressed, and the Type 7 LSAs have a forwarding address of 0. This would stop the translation of Type 7 LSAs to Type 5 by the NSSA ABR. The workaround for this is to configure at least one loopback interface in the NSSA area, or one interface with prefix suppression disabled, so that the interface address is selected as the forwarding address for all the Type 7 LSAs.

1. Enter the global configuration mode and configure the interfaces of the router.

```
RP/0/RP0:hostname# configure
RP/0/RP0:hostname(config)# interface TenGigE0/6/0/2.10
RP/0/RP0:hostname(config-if)# ipv4 address 10.1.1.1 255.255.255.0
RP/0/RP0:hostname(config-if)# no shut
RP/0/RP0:hostname(config-if)# exit
RP/0/RP0:hostname(config)# interface Loopback 0
RP/0/RP0:hostname(config-if)# ipv4 address 10.10.10.10 255.255.255.255
RP/0/RP0:hostname(config-if)# no shut
RP/0/RP0:hostname(config-if)# exit
```

2. Configure the OSPF process with prefix suppression.

```
RP/0/RP0:hostname# configure
RP/0/RP0:hostname(config)# router ospf pfx
RP/0/RP0:hostname(config-ospf)# router-id 10.10.10.10
RP/0/RP0:hostname(config-ospf)# prefix-suppression
```

3. Add the configured interfaces to the OSPF area.

```
RP/0/RP0:hostname(config-ospf)# area 0
RP/0/RP0:hostname(config-ospf-ar)# interface Loopback 0
RP/0/RP0:hostname(config-ospf-ar-if)# exit
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
RP/0/RP0:hostname(config-ospf-ar-if)# network point-to-point
```

4. Exit the OSPF area configuration mode and commit your configuration.

```
RP/0/RP0:hostname(config-ospf-ar-if)# exit
RP/0/RP0:hostname(config-ospf-ar)# exit
RP/0/RP0:hostname(config-ospf)# exit
RP/0/RP0:hostname(config)# commit
RP/0/RP0:hostname(config)# exit
```

5. Confirm your configuration.

```
RP/0/RP0:hostname# show running-configuration
...
interface Loopback0
  ipv4 address 10.10.10.10 255.255.255.255
!
interface TenGigE0/6/0/2.10
  ipv4 address 10.1.1.1 255.255.255.0
!
router ospf pfx
router-id 10.10.10.10
prefix-suppression
  area 0
    interface TenGigE0/6/0/2.10
      network point-to-point
    !
  !
!
```

6. Verify if prefix suppression is enabled.

```
RP/0/RP0:hostname# show ospf interface
Fri Jun 17 15:13:08.470 IST

Interfaces for OSPF 1

TenGigE0/6/0/2.10 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 1, Router ID 10.10.10.10, Network Type BROADCAST, Cost: 1
```

```

Transmit Delay is 1 sec, State BDR, Priority 1, MTU 1500, MaxPktSz 1500
Designated Router (ID) 10.10.10.20, Interface address 10.1.1.2
Backup Designated router (ID) 10.10.10.30, Interface address 10.1.1.3
Primary addresses not advertised
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:06:898
Index 2/2, flood queue length 0
Next 0(0)/0(0)
Last flood scan length is 2, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
LS Ack List: current length 0, high water mark 2
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.10.10.30 (Designated Router)
Suppress hello for 0 neighbor(s)
Multi-area interface Count is 0

```

If your output verifies that primary addresses are not advertised, then you have successfully configured prefix suppression for the OSPF process on the router.

Configure Prefix Suppression for an OSPF Area

Use the procedure in this section to configure prefix suppression for an OSPF area.



Note If you suppress prefixes on an area, the suppression is valid for all interfaces associated with the area.

1. Enter the global configuration mode and configure the interfaces of the router.

```

RP/0/RP0:hostname# configure
RP/0/RP0:hostname(config)# interface TenGigE0/6/0/2.10
RP/0/RP0:hostname(config-if)# ipv4 address 10.1.1.1 255.255.255.0
RP/0/RP0:hostname(config-if)# no shut
RP/0/RP0:hostname(config-if)# exit
RP/0/RP0:hostname(config)# interface Loopback 0
RP/0/RP0:hostname(config-if)# ipv4 address 10.10.10.10 255.255.255.255
RP/0/RP0:hostname(config-if)# no shut
RP/0/RP0:hostname(config-if)# exit

```

2. Configure the OSPF area with prefix suppression.

```

RP/0/RP0:hostname# configure
RP/0/RP0:hostname(config)# router ospf pfx
RP/0/RP0:hostname(config-ospf)# router-id 10.10.10.10
RP/0/RP0:hostname(config-ospf)# area 0
RP/0/RP0:hostname(config-ospf-ar)# prefix-suppression

```

3. Add the configured interfaces to the OSPF area.

```

RP/0/RP0:hostname(config-ospf-ar)# interface Loopback 0
RP/0/RP0:hostname(config-ospf-ar-if)# exit
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
RP/0/RP0:hostname(config-ospf-ar-if)# network point-to-point

```

4. Exit the OSPF area configuration mode and commit your configuration.

```

RP/0/RP0:hostname(config-ospf-ar-if)# exit
RP/0/RP0:hostname(config-ospf-ar)# exit
RP/0/RP0:hostname(config-ospf)# exit
RP/0/RP0:hostname(config)# commit
RP/0/RP0:hostname(config)# exit

```

5. Confirm your configuration.

```
RP/0/RP0:hostname# show running-configuration
...
interface Loopback0
  ipv4 address 10.10.10.10 255.255.255.255
!
interface TenGigE0/6/0/2.10
  ipv4 address 10.1.1.1 255.255.255.0
!
router ospf pfx
  router-id 10.10.10.10
  area 0
    prefix-suppression
    interface TenGigE0/6/0/2.10
      network point-to-point
    !
  !
!
```

6. Verify if prefix suppression is enabled.

```
RP/0/RP0:hostname# show ospf interface
Fri Jun 17 15:13:08.470 IST

Interfaces for OSPF 1

TenGigE0/6/0/2.10 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 1, Router ID 10.10.10.10, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1, MTU 1500, MaxPktSz 1500
  Designated Router (ID) 10.10.10.20, Interface address 10.1.1.2
  Backup Designated router (ID) 10.10.10.30, Interface address 10.1.1.3
  Primary addresses not advertised
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:06:898
  Index 2/2, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 2, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  LS Ack List: current length 0, high water mark 2
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.10.10.30 (Designated Router)
  Suppress hello for 0 neighbor(s)
  Multi-area interface Count is 0
```

If your output verifies that primary addresses are not advertised, then you have successfully configured prefix suppression for the OSPF area.

Configure Prefix Suppression for an OSPF Interface

Use the procedure in this section to configure prefix suppression for an OSPF interface.



Note If you suppress prefixes on an interface, suppression is valid only on that interface, and all other interfaces must be configured separately with prefix suppression.

1. Enter the global configuration mode and configure the interfaces of the router.

```
RP/0/RP0:hostname# configure
RP/0/RP0:hostname(config)# interface TenGigE0/6/0/2.10
```

```
RP/0/RP0:hostname(config-if)# ipv4 address 10.1.1.1 255.255.255.0
RP/0/RP0:hostname(config-if)# no shut
RP/0/RP0:hostname(config-if)# exit
RP/0/RP0:hostname(config)# interface Loopback 0
RP/0/RP0:hostname(config-if)# ipv4 address 10.10.10.10 255.255.255.255
RP/0/RP0:hostname(config-if)# no shut
RP/0/RP0:hostname(config-if)# exit
```

2. Configure the OSPF area.

```
RP/0/RP0:hostname# configure
RP/0/RP0:hostname(config)# router ospf pfx
RP/0/RP0:hostname(config-ospf)# router-id 10.10.10.10
RP/0/RP0:hostname(config-ospf)# area 0
```

3. Add the configured interfaces to the OSPF area, and configure prefix suppression on the required interface.

```
RP/0/RP0:hostname(config-ospf-ar)# interface Loopback 0
RP/0/RP0:hostname(config-ospf-ar-if)# exit
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
RP/0/RP0:hostname(config-ospf-ar-if)# network point-to-point
RP/0/RP0:hostname(config-ospf-ar-if)# prefix-suppression
```

4. Exit the OSPF area configuration mode and commit your configuration.

```
RP/0/RP0:hostname(config-ospf-ar-if)# exit
RP/0/RP0:hostname(config-ospf-ar)# exit
RP/0/RP0:hostname(config-ospf)# exit
RP/0/RP0:hostname(config)# commit
RP/0/RP0:hostname(config)# exit
```

5. Confirm your configuration.

```
RP/0/RP0:hostname# show running-configuration
...
interface Loopback0
  ipv4 address 10.10.10.10 255.255.255.255
!
interface TenGigE0/6/0/2.10
  ipv4 address 10.1.1.1 255.255.255.0
!
router ospf pfx
  router-id 10.10.10.10
  area 0
    interface TenGigE0/6/0/2.10
    network point-to-point
    prefix-suppression
  !
!
!
```

6. Verify if prefix suppression is enabled.

```
RP/0/RP0:hostname# show ospf interface
Fri Jun 17 15:13:08.470 IST

Interfaces for OSPF 1

TenGigE0/6/0/2.10 is up, line protocol is up
  Internet Address 10.1.1.1/24, Area 0
  Process ID 1, Router ID 10.10.10.10, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1, MTU 1500, MaxPktSz 1500
  Designated Router (ID) 10.10.10.20, Interface address 10.1.1.2
  Backup Designated router (ID) 10.10.10.30, Interface address 10.1.1.3
  Primary addresses not advertised
```

```

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:06:898
Index 2/2, flood queue length 0
Next 0(0)/0(0)
Last flood scan length is 2, maximum is 2
Last flood scan time is 0 msec, maximum is 0 msec
LS Ack List: current length 0, high water mark 2
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 10.10.10.30 (Designated Router)
Suppress hello for 0 neighbor(s)
Multi-area interface Count is 0

```

If your output verifies that primary addresses are not advertised, then you have successfully configured prefix suppression on the interface.

Multi-Area Adjacency for OSPF Version 2

The multi-area adjacency feature for OSPFv2 allows a link to be configured on the primary interface in more than one area so that the link could be considered as an intra-area link in those areas and configured as a preference over more expensive paths.

This feature establishes a point-to-point unnumbered link in an OSPF area. A point-to-point link provides a topological path for that area, and the primary adjacency uses the link to advertise the link consistent with draft-ietf-ospf-multi-area-adj-06.

The following are multi-area interface attributes and limitations:

- Exists as a logical construct over an existing primary interface for OSPF; however, the neighbor state on the primary interface is independent of the multi-area interface.
- Establishes a neighbor relationship with the corresponding multi-area interface on the neighboring router. A mixture of multi-area and primary interfaces is not supported.
- Advertises an unnumbered point-to-point link in the router link state advertisement (LSA) for the corresponding area when the neighbor state is full.
- Created as a point-to-point network type. You can configure multi-area adjacency on any interface where only two OSPF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OSPF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.
- Inherits the Bidirectional Forwarding Detection (BFD) characteristics from its primary interface. BFD is not configurable under a multi-area interface; however, it is configurable under the primary interface.

The multi-area interface inherits the interface characteristics from its primary interface, but some interface characteristics can be configured under the multi-area interface configuration mode as shown below:

```

RP/0/RP0:hostname(config-ospf-ar)# multi-area-interface TenGigE0/3/0/9.21
RP/0/RP0:hostname(config-ospf-ar-mif)# ?
 authentication          Enable authentication
 authentication-key      Authentication password (key)
 cost                    Interface cost
 cost-fallback           Cost when cumulative bandwidth goes below the threshold
 database-filter        Filter OSPF LSA during synchronization and flooding
 dead-interval          Interval after which a neighbor is declared dead
 distribute-list        Filter networks in routing updates
 hello-interval         Time between HELLO packets
 message-digest-key     Message digest authentication password (key)

```

mtu-ignore	Enable/Disable ignoring of MTU in DBD packets
packet-size	Customize size of OSPF packets upto MTU
retransmit-interval	Time between retransmitting lost link state advertisements
transmit-delay	Estimated time needed to send link-state update packet

```
RP/0/RP0:hostname(config-ospf-ar-mif)#
```

OSPF Authentication Message Digest Management

All OSPF routing protocol exchanges are authenticated and the method used can vary depending on how authentication is configured. When using cryptographic authentication, the OSPF routing protocol uses the Message Digest 5 (MD5) authentication algorithm to authenticate packets transmitted between neighbors in the network. For each OSPF protocol packet, a key is used to generate and verify a message digest that is appended to the end of the OSPF packet. The message digest is a one-way function of the OSPF protocol packet and the secret key. Each key is identified by the combination of interface used and the key identification. An interface may have multiple keys active at any time.

To manage the rollover of keys and enhance MD5 authentication for OSPF, you can configure a container of keys called a *keychain* with each key comprising the following attributes: generate/accept time, key identification, and authentication algorithm.

GTSM TTL Security Mechanism for OSPF

OSPF is a link state protocol that requires networking devices to detect topological changes in the network, flood Link State Advertisement (LSA) updates to neighbors, and quickly converge on a new view of the topology. However, during the act of receiving LSAs from neighbors, network attacks can occur, because there are no checks that unicast packets are originating from a neighbor that is one hop away or multiple hops away over virtual links.

For virtual links, OSPF packets travel multiple hops across the network; hence, the TTL value can be decremented several times. For these type of links, a minimum TTL value must be allowed and accepted for multiple-hop packets.

To filter network attacks originating from invalid sources traveling over multiple hops, the Generalized TTL Security Mechanism (GTSM), RFC 3682, is used to prevent the attacks. GTSM filters link-local addresses and allows for only one-hop neighbor adjacencies through the configuration of TTL value 255. The TTL value in the IP header is set to 255 when OSPF packets are originated, and checked on the received OSPF packets against the default GTSM TTL value 255 or the user configured GTSM TTL value, blocking unauthorized OSPF packets originated from TTL hops away.

Path Computation Element for OSPFv2

A PCE is an entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCE is accomplished when a PCE address and client is configured for MPLS-TE. PCE communicates its PCE address and capabilities to OSPF then OSPF packages this information in the PCE Discovery type-length-value (TLV) (Type 2) and reoriginates the RI LSA. OSPF also includes the Router Capabilities TLV (Type 1) in all its RI LSAs. The PCE Discovery TLV contains the PCE address sub-TLV (Type 1) and the Path Scope Sub-TLV (Type 2).

The PCE Address Sub-TLV specifies the IP address that must be used to reach the PCE. It should be a loop-back address that is always reachable, this TLV is mandatory, and must be present within the PCE Discovery TLV. The Path Scope Sub-TLV indicates the PCE path computation scopes, which refers to the PCE ability to compute or participate in the computation of intra-area, inter-area, inter-AS or inter-layer TE LSPs.

PCE extensions to OSPFv2 include support for the Router Information Link State Advertisement (RI LSA). OSPFv2 is extended to receive all area scopes (LSA Types 9, 10, and 11). However, OSPFv2 originates only area scope Type 10.

OSPF IP Fast Reroute Loop Free Alternate

The OSPF IP Fast Reroute (FRR) Loop Free Alternate (LFA) computation supports these:

- Fast rerouting capability by using IP forwarding and routing
- Handles failure in the line cards in minimum time

OSPF Over GRE Interfaces

Cisco IOS XR software provides the capability to run OSPF protocols over Generic Routing Encapsulation (GRE) tunnel interfaces.

VRF-lite Support for OSPFv2

VRF-lite capability is enabled for OSPF version 2 (OSPFv2). VRF-lite is the virtual routing and forwarding (VRF) deployment without the BGP/MPLS based backbone. In VRF-lite, individual provider edge (PE) routers are directly connected using VRF interfaces. To enable VRF-lite in OSPFv2, configure the **capability vrf-lite** command in VRF configuration mode. When VRF-lite is configured, the DN bit processing and the automatic Area Border Router (ABR) status setting are disabled.

OSPFv2 Unequal Cost Load Balancing

Unequal Cost Load Balancing feature in Cisco IOS XR OSPFv2 feature enables Unequal Cost Multipath (UCMP) calculation based on configured prefix-list and based on variance factor. UCMP path can be calculated for all prefixes or only for selected prefixes based on the configuration. Selected interfaces can be excluded to be used as a candidate for UCMP paths. The calculated UCMP paths are then installed in the routing information base (RIB) subject to the max-path limit.

The OSPFv2 interior gateway protocol is used to calculate paths to prefixes inside an autonomous system. OSPF calculates up to maximum paths (max-path) equal cost multi-paths (ECMPs) for each prefix, where max-path is either limited by the router support or is configured by the user.

UCMP Paths Calculation

In some topologies, alternate paths to prefix exist even though their metric is higher than the metric of the best path(s). These paths are called Unequal Cost Multipaths (UCMPs). These paths are guaranteed to be loop free. Users can send some portion of the traffic down these paths to better utilize the available bandwidth. However, the UCMP paths are not discovered by the traditional Dijkstra calculation. Additional computation is required to discover these paths.

Unequal Cost Multipath Load-balancing for OSPF

The unequal cost multipath (UCMP) load-balancing adds the capability with Open Shortest Path First (OSPF) to load-balance traffic proportionally across multiple paths, with different cost. Without UCMP enabled, only the best cost paths are discovered by OSPF (ECMP) and alternate higher cost paths are not computed.

Generally, higher bandwidth links have lower IGP metrics configured, so that they form the shortest IGP paths. With the UCMP load-balancing enabled, IGP can use even lower bandwidth links or higher cost links for traffic, and can install these paths to the forwarding information base (FIB). OSPF installs multiple paths to the same destination in FIB, but each path will have a 'load metric/weight' associated with it. FIB uses this load metric/weight to decide the amount of traffic that needs to be sent on a higher bandwidth path and the amount of traffic that needs to be sent on a lower bandwidth path.

The UCMP computation is provided under OSPF VRF context, enabling UCMP computation for a particular VRF. For default VRF the configuration is done under the OSPF global mode. The UCMP configuration is also provided with a prefix-list option, which would limit the UCMP computation only for the prefixes present in the prefix-list. If prefix-list option is not provided, UCMP computation is done for the reachable prefixes in OSPF. The number of UCMP paths to be considered and installed is controlled using the **variance** configuration. Variance value identifies the range for the UCMP path metric to be considered for installation into routing information base (RIB/FIB) and is defined in terms of a percentage of the primary path metric. Total number of paths, including ECMP and UCMP paths together is limited by the max-path configuration or by the max-path capability of the platform.

There is an option to exclude an interface from being used for UCMP computation. If it is desired that a particular interface should not be considered as a UCMP nexthop, for any prefix, then use the UCMP **exclude interface** command to configure the interface to be excluded from UCMP computation.

Enabling the UCMP configuration indicates that OSPF should perform UCMP computation for the all the reachable OSPF prefixes or all the prefixes permitted by the prefix-list, if the prefix-list option is used. The UCMP computation happens only after the primary SPF and route calculation is completed. There would be a configurable delay (default delay is 100 ms) from the time primary route calculation is completed and UCMP computation is started. Use the UCMP **delay-interval** command to configure the delay between primary SPF completion and start of UCMP computation. UCMP computation will be done during the fast re-route computation (IPFRR does not need to be enabled for UCMP computation to be performed). If IPFRR is enabled, the fast re-route backup paths will be calculated for both the primary equal cost multipath (ECMP) paths and the UCMP paths.

To manually adjust UCMP ratio, use any command that changes the metric of the link.

- By using the bandwidth command in interface configuration mode
- By adjusting the OSPF interface cost on the link

How to Implement OSPF

This section contains the following procedures:

Enabling OSPF

This task explains how to perform the minimum OSPF configuration on your router that is to enable an OSPF process with a router ID, configure a backbone or nonbackbone area, and then assign one or more interfaces on which OSPF runs.

Before you begin

Although you can configure OSPF before you configure an IP address, no OSPF routing occurs until at least one IP address is configured.

Procedure

Step 1 **configure**

Step 2 **router ospf**

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id { router-id }**

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IP address as the router ID.

Step 4 **area area-id**

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures an area for the OSPF process.

- Backbone areas have an area ID of 0.
- Nonbackbone areas have a nonzero area ID.
- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 5 **interface type interface-path-id**

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.

Step 6 Repeat Step 5 for each interface that uses OSPF.

—

Step 7 **log adjacency changes [detail] [enable | disable]**

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# log adjacency changes detail
```

(Optional) Requests notification of neighbor changes.

- By default, this feature is enabled.
- The messages generated by neighbor changes are considered notifications, which are categorized as severity Level 5 in the **logging console** command. The **logging console** command controls which severity level of messages are sent to the console. By default, all severity level messages are sent.

Step 8 **commit**

Configuring Stub and Not-So-Stubby Area Types

This task explains how to configure the stub area and the NSSA for OSPF.

Procedure**Step 1** **configure****Step 2** **router ospf** *process-name***Example:**

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }**Example:**

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IP address as the router ID.

Step 4 **area** *area-id***Example:**

```
RP/0/RP0:hostname(config-ospf)# area 1
```

Enters area configuration mode and configures a nonbackbone area for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 5 Do one of the following:

- **stub** [**no-summary**]
- **nssa** [**no-redistribution**] [**default-information-originate**] [**no-summary**]

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# stub no summary
```

or

```
RP/0/RP0:hostname(config-ospf-ar)# nssa no-redistribution
```

Defines the nonbackbone area as a stub area.

- Specify the **no-summary** keyword to further reduce the number of LSAs sent into a stub area. This keyword prevents the ABR from sending summary link-state advertisements (Type 3) in the stub area.

or

Defines an area as an NSSA.

Step 6 Do one of the following:

- **stub**
- **nssa**

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# stub
```

or

```
RP/0/RP0:hostname(config-ospf-ar)# nssa
```

(Optional) Turns off the options configured for stub and NSSA areas.

- If you configured the stub and NSSA areas using the optional keywords (**no-summary** , **no-redistribution** , **default-information-originate** , and **no-summary**) in Step 5, you must now reissue the **stub** and **nssa** commands without the keywords—rather than using the **no** form of the command.
- For example, the **no nssa default-information-originate** form of the command changes the NSSA area into a normal area that inadvertently brings down the existing adjacencies in that area.

Step 7 **default-cost** *cost*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)#default-cost 15
```

(Optional) Specifies a cost for the default summary route sent into a stub area or an NSSA.

- Use this command only on ABRs attached to the NSSA. Do not use it on any other routers in the area.
- The default cost is 1.

Step 8 **commit**

Step 9 Repeat this task on all other routers in the stub area or NSSA.

Configuring Neighbors for Nonbroadcast Networks

This task explains how to configure neighbors for a nonbroadcast network. This task is optional.

Before you begin

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

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IP address as the router ID.

Step 4 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures an area for the OSPF process.

- The example configures a backbone area.
- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 5 **network** { **broadcast** | **non-broadcast** | { **point-to-multipoint** [**non-broadcast**] | **point-to-point** }

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# network non-broadcast
```

Configures the OSPF network type to a type other than the default for a given medium.

- The example sets the network type to NBMA.

Step 6 **dead-interval** *seconds*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# dead-interval 40
```

(Optional) Sets the time to wait for a hello packet from a neighbor before declaring the neighbor down.

Step 7 **hello-interval** *seconds*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# hello-interval 10
```

(Optional) Specifies the interval between hello packets that OSPF sends on the interface.

Note It is recommended to set the hello timer interval to the default of 10 seconds. OSPF sessions may flap during switchover if hello-interval timer configured is less than default value.

Step 8 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.

- In this example, the interface inherits the nonbroadcast network type and the hello and dead intervals from the areas because the values are not set at the interface level.

Step 9 **neighbor** *ip-address* [**priority** *number*] [**poll-interval** *seconds*] [**cost** *number*]

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# neighbor 10.20.20.1 priority 3 poll-interval 15
```

Configures the IPv4 address of OSPF neighbors interconnecting to nonbroadcast networks.

- The *ipv6-link-local-address* argument must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.
- The **priority** keyword notifies the router that this neighbor is eligible to become a DR or BDR. The priority value should match the actual priority setting on the neighbor router. The neighbor priority default value is zero. This keyword does not apply to point-to-multipoint interfaces.
- The **poll-interval** keyword does not apply to point-to-multipoint interfaces. RFC 1247 recommends that this value be much larger than the hello interval. The default is 120 seconds (2 minutes).
- Neighbors with no specific cost configured assumes the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost number** is the only keyword and argument combination that works. The **cost** keyword does not apply to NBMA networks.

- The **database-filter** keyword filters outgoing LSAs to an OSPF neighbor. If you specify the **all** keyword, incoming and outgoing LSAs are filtered. Use with extreme caution since filtering may cause the routing topology to be seen as entirely different between two neighbors, resulting in an unwanted traffic drop or routing loops.

Step 10 Repeat Step 9 for all neighbors on the interface.

—

Step 11 exit

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# exit
```

Enters area configuration mode.

Step 12 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/3/0/5.20
```

Enters interface configuration mode and associates one or more interfaces for the area configured in Step 4.

- In this example, the interface inherits the nonbroadcast network type and the hello and dead intervals from the areas because the values are not set at the interface level.

Step 13 **neighbor** *ip-address* [**priority** *number*] [**poll-interval** *seconds*] [**cost** *number*] [**database-filter** [**all**]]

Example:

```
RP/0/  
/CPU0:router(config-ospf-ar)# neighbor 10.34.16.6
```

Configures the IPv4 address of OSPF neighbors interconnecting to nonbroadcast networks.

- The *ipv6-link-local-address* argument must be in the form documented in RFC 2373 in which the address is specified in hexadecimal using 16-bit values between colons.
- The **priority** keyword notifies the router that this neighbor is eligible to become a DR or BDR. The priority value should match the actual priority setting on the neighbor router. The neighbor priority default value is zero. This keyword does not apply to point-to-multipoint interfaces.
- The **poll-interval** keyword does not apply to point-to-multipoint interfaces. RFC 1247 recommends that this value be much larger than the hello interval. The default is 120 seconds (2 minutes).
- Neighbors with no specific cost configured assumes the cost of the interface, based on the **cost** command. On point-to-multipoint interfaces, **cost number** is the only keyword and argument combination that works. The **cost** keyword does not apply to NBMA networks.
- The **database-filter** keyword filters outgoing LSAs to an OSPF neighbor. If you specify the **all** keyword, incoming and outgoing LSAs are filtered. Use with extreme caution since filtering may cause the routing topology to be seen as entirely different between two neighbors, resulting in an unwanted traffic drop or routing loops.

Step 14 Repeat Step 13 for all neighbors on the interface.

Step 15 `commit`

Configuring Authentication at Different Hierarchical Levels for OSPF Version 2

This task explains how to configure MD5 (secure) authentication on the OSPF router process, configure one area with plain text authentication, and then apply one interface with clear text (null) authentication.



Note Authentication configured at the interface level overrides authentication configured at the area level and the router process level. If an interface does not have authentication specifically configured, the interface inherits the authentication parameter value from a higher hierarchical level. See [OSPF Hierarchical CLI and CLI Inheritance, on page 3](#) for more information about hierarchy and inheritance.

Before you begin

If you choose to configure authentication, you must first decide whether to configure plain text or MD5 authentication, and whether the authentication applies to all interfaces in a process, an entire area, or specific interfaces. See [Route Authentication Methods for OSPF, on page 7](#) for information about each type of authentication and when you should use a specific method for your network.

Procedure

Step 1 `router ospf process-name`

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 2 `router-id { router-id }`

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Step 3 `authentication [message-digest | null]`

Example:

```
RP/0/RP0:hostname(config-ospf)#authentication message-digest
```

Enables MD5 authentication for the OSPF process.

- This authentication type applies to the entire router process unless overridden by a lower hierarchical level such as the area or interface.

Step 4 **message-digest-key** *key-id* **md5** { *key* | **clear** *key* | **encrypted** *key* | **LINE**}

Example:

```
RP/0/RP0:hostname(config-ospf)#message-digest-key 4 md5 yourkey
```

Specifies the MD5 authentication key for the OSPF process.

- The neighbor routers must have the same key identifier.

Step 5 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures a backbone area for the OSPF process.

Step 6 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/3/0/5.20
```

Enters interface configuration mode and associates one or more interfaces to the backbone area.

- All interfaces inherit the authentication parameter values specified for the OSPF process (Step 4, Step 5, and Step 6).

Step 7 Repeat Step 7 for each interface that must communicate, using the same authentication.

—

Step 8 **exit**

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# exit
```

Enters area OSPF configuration mode.

Step 9 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 1
```

Enters area configuration mode and configures a nonbackbone area 1 for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 10 **authentication** [**message-digest** | **null**]

Example:


```
RP/0/RP0:hostname(config-ospf-ar)# authentication
```

Enables Type 1 (plain text) authentication that provides no security.

- The example specifies plain text authentication (by not specifying a keyword). Use the **authentication-key** command in interface configuration mode to specify the plain text password.

Step 11 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/3/0/9.21
```

Enters interface configuration mode and associates one or more interfaces to the nonbackbone area 1 specified in Step 7.

- All interfaces configured inherit the authentication parameter values configured for area 1.

Step 12 Repeat Step 12 for each interface that must communicate, using the same authentication.

—

Step 13 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to a different authentication type.

Step 14 **authentication** [**message-digest** | **null**]

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# authentication null
```

Specifies no authentication on Ten Gigabit Ethernet interface 0/6/0/2.10, overriding the plain text authentication specified for area 1.

- By default, all of the interfaces configured in the same area inherit the same authentication parameter values of the area.

Step 15 **commit**

Controlling the Frequency That the Same LSA Is Originated or Accepted for OSPF

This task explains how to tune the convergence time of OSPF routes in the routing table when many LSAs need to be flooded in a very short time interval.

Procedure

Step 1 **configure**

Step 2 `router ospf process-name`**Example:**

```
RP/0/RP0:hostname:router(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The `process-name` argument is any alphanumeric string no longer than 40 characters.

Step 3 `router-id {router-id}`**Example:**

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IP address as the router ID.

Step 4 Perform Step 5 or Step 6 or both to control the frequency that the same LSA is originated or accepted.

—

Step 5 `timers lsa refresh seconds`**Example:**

```
RP/0/RP0:hostname(config-ospf)# timers lsa refresh 1800
```

Sets how often self-originated LSAs should be refreshed, in seconds.

- The default is 1800 seconds for both OSPF.

Step 6 `timers lsa min-arrival seconds`**Example:**

```
RP/0/RP0:hostname(config-ospf)# timers lsa min-arrival 2
```

Limits the frequency that new processes of any particular OSPF Version 2 LSA can be accepted during flooding.

- The default is 1 second.

Step 7 `timers lsa group-pacing seconds`**Example:**

```
RP/0/  
/CPU0:router(config-ospf)# timers lsa group-pacing 1000
```

Changes the interval at which OSPF link-state LSAs are collected into a group for flooding.

- The default is 240 seconds.

Step 8 `commit`

Creating a Virtual Link with MD5 Authentication to Area 0 for OSPF

This task explains how to create a virtual link to your backbone (area 0) and apply MD5 authentication. You must perform the steps described on both ABRs, one at each end of the virtual link. To understand virtual links, see [Virtual Link and Transit Area for OSPF, on page 12](#).



Note After you explicitly configure area parameter values, they are inherited by all interfaces bound to that area—unless you override the values and configure them explicitly for the interface. An example is provided in [Virtual Link Configured with MD5 Authentication for OSPF Version 2: Example, on page 65](#).

Before you begin

The following prerequisites must be met before creating a virtual link with MD5 authentication to area 0:

- You must have the router ID of the neighbor router at the opposite end of the link to configure the local router. You can execute the **show ospf** command on the remote router to get its router ID.
- For a virtual link to be successful, you need a stable router ID at each end of the virtual link. You do not want them to be subject to change, which could happen if they are assigned by default. (See [OSPF Process and Router ID, on page 6](#) for an explanation of how the router ID is determined.) Therefore, we recommend that you perform one of the following tasks before configuring a virtual link:
 - Use the **router-id** command to set the router ID. This strategy is preferable.
 - Configure a loopback interface so that the router has a stable router ID.
- Before configuring your virtual link for OSPF Version 2, you must decide whether to configure plain text authentication, MD5 authentication, or no authentication (which is the default). Your decision determines whether you need to perform additional tasks related to authentication.



Note If you decide to configure plain text authentication or no authentication, see the **authentication** command provided in chapter *OSPF - IPv4 Commands on OTN and WDM Command Reference for Cisco NCS 4000 Series*

Procedure

Step 1 `show ospf [process-name]`

Example:

```
RP/0//CPU0:router# show ospf
```

(Optional) Displays general information about OSPF routing processes.

- The output displays the router ID of the local router. You need this router ID to configure the other end of the link.

Step 2 `configure`

Step 3 **router ospf** *process-name***Example:**

```
RP/0//CPU0:router(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 4 **router-id** { *router-id* }**Example:**

```
RP/0//CPU0:router(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 5 **area** *area-id***Example:**

```
RP/0//CPU0:router(config-ospf)# area 1
```

Enters area configuration mode and configures a nonbackbone area for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 6 **virtual-link** *router-id***Example:**

```
RRP/0//CPU0:router(config-ospf-ar)# virtual-link 10.3.4.5
```

Defines an OSPF virtual link.

- See .

Step 7 **authentication message-digest****Example:**

```
RP/0//CPU0:router(config-ospf-ar-vl)#authentication message-digest
```

Selects MD5 authentication for this virtual link.

Step 8 **message-digest-key** *key-id* **md5** { *key* | **clear** *key* | **encrypted** *key* }**Example:**

```
RP/0//CPU0:router(config-ospf-ar-vl)#message-digest-key 4 md5 yourkey
```

Defines an OSPF virtual link.

- See to understand a virtual link.

- The *key-id* argument is a number in the range from 1 to 255. The *key* argument is an alphanumeric string of up to 16 characters. The routers at both ends of the virtual link must have the same key identifier and key to be able to route OSPF traffic.
- Once the key is encrypted it must remain encrypted.

Step 9 Repeat all of the steps in this task on the ABR that is at the other end of the virtual link. Specify the same key ID and key that you specified for the virtual link on this router.

Step 10 **commit**

Step 11 **show ospf** [*process-name*] [*area-id*] **virtual-links**

Example:

```
RP/0//CPU0:router# show ospf 1 2 virtual-links
```

(Optional) Displays the parameters and the current state of OSPF virtual links.

Summarizing Subnetwork LSAs on an OSPF ABR

If you configured two or more subnetworks when you assigned your IP addresses to your interfaces, you might want the software to summarize (aggregate) into a single LSA all of the subnetworks that the local area advertises to another area. Such summarization would reduce the number of LSAs and thereby conserve network resources. This summarization is known as interarea route summarization. It applies to routes from within the autonomous system. It does not apply to external routes injected into OSPF by way of redistribution.

This task configures OSPF to summarize subnetworks into one LSA, by specifying that all subnetworks that fall into a range are advertised together. This task is performed on an ABR only.

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4 **area** *area-id***Example:**

```
RP/0/RP0:hostname(config-ospf)# area
```

Enters area configuration mode and configures a nonbackbone area for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 5 Do one of the following:

- **range** *ip-address mask* [**advertise** | **not-advertise**]
- **range** *ipv6-prefix / prefix-length* [**advertise** | **not-advertise**]

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# range 192.168.0.0 255.255.0.0 advertise
```

or

```
RP/0/RP0:hostname(config-ospf-ar)# range 4004:f000::/32 advertise
```

Consolidates and summarizes OSPF routes at an area boundary.

- The **advertise** keyword causes the software to advertise the address range of subnetworks in a Type 3 summary LSA.
- The **not-advertise** keyword causes the software to suppress the Type 3 summary LSA, and the subnetworks in the range remain hidden from other areas.
- In the first example, all subnetworks for network 192.168.0.0 are summarized and advertised by the ABR into areas outside the backbone.
- In the second example, two or more IPv4 interfaces are covered by a 192.x.x network.

Step 6 **interface** *type interface-path-id***Example:**

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 7 **commit**

Redistribute Routes into OSPF

This task redistributes routes from an IGP (could be a different OSPF process) into OSPF.

Before you begin

For information about configuring routing policy, see *Implementing Routing Policy on*

Procedure**Step 1**

configure

Step 2

router ospf *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3

router-id { *router-id* }

Example:

```
RRP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4

redistribute *protocol* [*process-id*] { **level-1** | **level-1-2** | **level-2** } [**metric** *metric-value*] [**metric-type** *type-value*] [**match** { **external** [**1** | **2**] } [**tag** *tag-value*] [**route-policy** *policy-name*]

Example:

```
RP/0/RP0:hostname(config-ospf)# redistribute bgp 100
```

or

```
RP/0/RP0:hostname(config-router)# redistribute bgp 110
```

Redistributes OSPF routes from one routing domain to another routing domain.

- This command causes the router to become an ASBR by definition.
- OSPF tags all routes learned through redistribution as external.
- The protocol and its process ID, if it has one, indicate the protocol being redistributed into OSPF.
- The metric is the cost you assign to the external route. The default is 20 for all protocols except BGP, whose default metric is 1.
- The OSPF example redistributes BGP autonomous system 1, Level 1 routes into OSPF as Type 2 external routes.

Step 5

Do one of the following:

- **summary-prefix** *address mask* [**not-advertise**] [**tag** *tag*]
- **summary-prefix** *ipv6-prefix / prefix-length* [**not-advertise**] [**tag** *tag*]

Example:

```
RP/0/RP0:hostname(config-ospf)# summary-prefix 10.1.0.0 255.255.0.0
```

or

```
RP/0/RP0:hostname(config-router)# summary-prefix 2010:11:22::/32
```

(Optional) Creates aggregate addresses for OSPF.

- This command provides external route summarization of the non-OSPF routes.
- External ranges that are being summarized should be contiguous. Summarization of overlapping ranges from two different routers could cause packets to be sent to the wrong destination.
- This command is optional. If you do not specify it, each route is included in the link-state database and advertised in LSAs.
- In the OSPFv2 example, the summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external LSA.

Step 6 **commit**

Configuring OSPF Shortest Path First Throttling

This task explains how to configure SPF scheduling in millisecond intervals and potentially delay SPF calculations during times of network instability. This task is optional.

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4 **timers throttle spf** *spf-start spf-hold spf-max-wait*

Example:


```
RP/0/RP0:hostname(config-ospf)# timers throttle spf 10 4800 90000
```

Sets SPF throttling timers.

Step 5 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures a backbone area.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 6 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 7 **commit**

Step 8 **show ospf** [*process-name*]

Example:

```
RP/0/RP0:hostname# show ospf 1
```

(Optional) Displays SPF throttling timers.

Examples

In the following example, the **show ospf XR EXEC** command is used to verify that the initial SPF schedule delay time, minimum hold time, and maximum wait time are configured correctly. Additional details are displayed about the OSPF process, such as the router type and redistribution of routes.

```
show ospf 1
```

```
Routing Process "ospf 1" with ID 192.168.4.3
  Supports only single TOS(TOS0) routes
  Supports opaque LSA
  It is an autonomous system boundary router
  Redistributing External Routes from,
    ospf 2
  Initial SPF schedule delay 5 msec
  Minimum hold time between two consecutive SPF's 100 msec
  Maximum wait time between two consecutive SPF's 1000 msec
  Minimum LSA interval 5 secs. Minimum LSA arrival 1 sec
  Number of external LSA 0. Checksum Sum 00000000
  Number of opaque AS LSA 0. Checksum Sum 00000000
  Number of DCbitless external and opaque AS LSA 0
  Number of DoNotAge external and opaque AS LSA 0
  Number of areas in this router is 1. 1 normal 0 stub 0 nssa
```

```
External flood list length 0
Non-Stop Forwarding enabled
```



Note For a description of each output display field, see the **show ospf** command in the *OSPF-IPv4 Commands on OTN and WDM Command Reference for Cisco NCS 4000 Series*

Configuring Nonstop Forwarding Specific to Cisco for OSPF Version 2

This task explains how to configure OSPF NSF specific to Cisco on your NSF-capable router. This task is optional.

Before you begin

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.



Note The following are restrictions when configuring nonstop forwarding:

- OSPF Cisco NSF for virtual links is not supported.
 - Neighbors must be NSF aware.
-

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4 Do one of the following:

- **nsf cisco**
- **nsf cisco enforce global**

Example:

```
RP/0/RP0:hostname(config-ospf)# nsf cisco enforce global
```

Enables Cisco NSF operations for the OSPF process.

- Use the **nsf cisco** command without the optional **enforce** and **global** keywords to terminate the NSF restart mechanism on the interfaces of detected non-NSF neighbors and allow NSF neighbors to function properly.
- Use the **nsf cisco** command with the optional **enforce** and **global** keywords if the router is expected to perform NSF during restart. However, if non-NSF neighbors are detected, NSF restart is canceled for the entire OSPF process.

Step 5 **nsf interval** *seconds*

Example:

```
RP/0/RP0:hostname(config-ospf)# nsf interval 120
```

Sets the minimum time between NSF restart attempts.

Note When you use this command, the OSPF process must be up for at least 90 seconds before OSPF attempts to perform an NSF restart.

Step 6 **nsfflush-delay-time***seconds*

Example:

```
RP/0/RP0:hostname(config-ospf)#nsf flush-delay-time 1000
```

Sets the maximum time allowed for external route learning in seconds.

Step 7 **nsflifetime***seconds*

Example:

```
RP/0/RP0:hostname(config-ospf)#nsf lifetime 90
```

Sets the maximum route lifetime of NSF following a restart in seconds.

Step 8 **nsfietf**

Example:

```
RP/0/RP0:hostname(config-ospf)#nsf ietf
```

Enables ietf graceful restart.

Step 9 **commit**

Configuring OSPF Version 2 for MPLS Traffic Engineering

This task explains how to configure OSPF for MPLS TE. This task is optional.

Before you begin

Your network must support the following features before you enable MPLS TE for OSPF on your router:

- MPLS
- IP Cisco Express Forwarding (CEF)



Note You must enter the commands in the following task on every OSPF router in the traffic-engineered portion of your network.

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4 **mpls traffic-eng router-id** *interface-type interface-instance*

Example:

```
RP/0/RP0:hostname(config-ospf)# mpls traffic-eng router-id loopback 0
```

(Optional) Specifies that the traffic engineering router identifier for the node is the IP address associated with a given interface.

- This IP address is flooded to all nodes in TE LSAs.
- For all traffic engineering tunnels originating at other nodes and ending at this node, you must set the tunnel destination to the traffic engineering router identifier of the destination node because that is the address that the traffic engineering topology database at the tunnel head uses for its path calculation.

- We recommend that loopback interfaces be used for MPLS TE router ID because they are more stable than physical interfaces.

Step 5 `area area-id`

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures an area for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.

Step 6 `mpls traffic-eng`

Example:

```
RP/0/RP0:hostname(config-ospf)# mpls traffic-eng
```

Configures the MPLS TE under the OSPF area.

Step 7 `interface type interface-path-id`

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface interface loopback0
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 8 `commit`

Step 9 `show ospf [process-name] [area-id] mpls traffic-eng { link | fragment }`

Example:

```
RP/0/RP0:hostname# show ospf 1 0 mpls traffic-eng link
```

(Optional) Displays information about the links and fragments available on the local router for MPLS TE.

Examples

This section provides the following output examples:

Sample Output for the show ospf Command Before Configuring MPLS TE

In the following example, the `show route ospf XR EXEC` command verifies that Ten Gigabit Ethernet interface 0/6/0/2.10 exists and MPLS TE is not configured:

```
show route ospf 1
```

```

O   11.0.0.0/24 [110/15] via 0.0.0.0, 3d19h, tunnel-te1
O   192.168.0.12/32 [110/11] via 11.1.0.2, 3d19h, TenGigE0/6/0/2.10
O   192.168.0.13/32 [110/6] via 0.0.0.0, 3d19h, tunnel-te1

```

Sample Output for the show ospf mpls traffic-eng Command

In the following example, the **show ospf mpls traffic-eng XR EXEC** command verifies that the MPLS TE fragments are configured correctly:

```

show ospf 1 mpls traffic-eng fragment

OSPF Router with ID (192.168.4.3) (Process ID 1)

Area 0 has 1 MPLS TE fragment. Area instance is 3.
MPLS router address is 192.168.4.2
Next fragment ID is 1

Fragment 0 has 1 link. Fragment instance is 3.
Fragment has 0 link the same as last update.
Fragment advertise MPLS router address
  Link is associated with fragment 0. Link instance is 3
    Link connected to Point-to-Point network
      Link ID :55.55.55.55
      Interface Address :192.168.50.21
      Neighbor Address :192.168.4.1
      Admin Metric :0
      Maximum bandwidth :19440000
      Maximum global pool reservable bandwidth :25000000
      Maximum sub pool reservable bandwidth :3125000
      Number of Priority :8
      Global pool unreserved BW
      Priority 0 : 25000000 Priority 1 : 25000000
      Priority 2 : 25000000 Priority 3 : 25000000
      Priority 4 : 25000000 Priority 5 : 25000000
      Priority 6 : 25000000 Priority 7 : 25000000
      Sub pool unreserved BW
      Priority 0 : 3125000 Priority 1 : 3125000
      Priority 2 : 3125000 Priority 3 : 3125000
      Priority 4 : 3125000 Priority 5 : 3125000
      Priority 6 : 3125000 Priority 7 : 3125000
      Affinity Bit :0

```

In the following example, the **show ospf mpls traffic-eng XR EXEC** command verifies that the MPLS TE links on area instance 3 are configured correctly:

```

show ospf mpls traffic-eng link

OSPF Router with ID (192.168.4.1) (Process ID 1)

Area 0 has 1 MPLS TE links. Area instance is 3.

Links in hash bucket 53.
Link is associated with fragment 0. Link instance is 3
  Link connected to Point-to-Point network
    Link ID :192.168.50.20
    Interface Address :192.168.20.50
    Neighbor Address :192.168.4.1
    Admin Metric :0
    Maximum bandwidth :19440000
    Maximum global pool reservable bandwidth :25000000
    Maximum sub pool reservable bandwidth :3125000
    Number of Priority :8
    Global pool unreserved BW
    Priority 0 : 25000000 Priority 1 : 25000000
    Priority 2 : 25000000 Priority 3 : 25000000

```

```

Priority 4 : 25000000 Priority 5 : 25000000
Priority 6 : 25000000 Priority 7 : 25000000
Sub pool unreserved BW
Priority 0 : 3125000 Priority 1 : 3125000
Priority 2 : 3125000 Priority 3 : 3125000
Priority 4 : 3125000 Priority 5 : 3125000
Priority 6 : 3125000 Priority 7 : 3125000
Affinity Bit :0

```

Sample Output for the show ospf Command After Configuring MPLS TE

In the following example, the **show route ospf** XR EXEC command verifies that the MPLS TE tunnels replaced Ten Gigabit Ethernet interface 0/6/0/2.10 and that configuration was performed correctly:

```

show route ospf 1

O E2 192.168.10.0/24 [110/20] via 0.0.0.0, 00:00:15, tunnel2
O E2 192.168.11.0/24 [110/20] via 0.0.0.0, 00:00:15, tunnel2
O E2 192.168.1244.0/24 [110/20] via 0.0.0.0, 00:00:15, tunnel2
O 192.168.12.0/24 [110/2] via 0.0.0.0, 00:00:15, tunnel2

```

Enabling Nonstop Routing for OSPFv2

This optional task describes how to enable nonstop routing (NSR) for OSPFv2 process. NSR is disabled by default. When NSR is enabled, OSPF process on the active RP synchronizes all necessary data and states with the OSPF process on the standby RP. When the switchover happens, OSPF process on the newly active RP has all the necessary data and states to continue running and does not require any help from its neighbors.

Procedure

Step 1 configure

Enter the global configuration mode.

Step 2 router ospf *instance-id*

Example:

```
RP/0/RP0:hostname(config)# router ospf isp
```

Enable OSPF routing for the specified routing process. In this example, the OSPF instance is called isp.

Step 3 nsr

Example:

```
RP/0/RP0:hostname(config-ospf)# nsr
```

Enable NSR for the OSPFv2 process.

Step 4 commit

Commit your configuration.

Configuring OSPFv2 OSPF SPF Prefix Prioritization

Perform this task to configure OSPFv2 OSPF SPF (shortest path first) prefix prioritization.

Procedure

Step 1 **configure**

Step 2 **prefix-set** *prefix-set name*

Example:

```
RP/0/RP0:hostname(config)#prefix-set ospf-critical-prefixes
RP/0/RP0:hostname(config-pfx)#66.0.0.0/16
RP/0/RP0:hostname(config-pfx)#end-set
```

Configures the prefix set.

Step 3 **route-policy** *route-policy name* **if destination in** *prefix-set name* **then set spf-priority** {critical | high | medium} **endif**

Example:

```
RP/0/RP0:hostname#route-policy ospf-spf-priority
RP/0/RP0:hostname(config-rpl)#if destination in ospf-critical-prefixes then
  set spf-priority critical
endif
RP/0/RP0:hostname(config-rpl)#end-policy
```

Configures route policy and sets OSPF SPF priority.

Step 4 **router ospf** *ospf-name*

Example:

```
RP/0/RP0:hostname# router ospf 1
```

Enters Router OSPF configuration mode.

Step 5 **router ospf** *ospf name*

Example:

```
RP/0/RP0:hostname# router ospf 1
```

Enters Router OSPF configuration mode.

Step 6 **spf prefix-priority route-policy** *route-policy name*

Example:

```
RP/0/RP0:hostname(config-ospf)# spf prefix-priority route-policy ospf-spf-priority
```

Configures SPF prefix-priority for the defined route policy.

Note Configure the **spf prefix-priority** command under router OSPF.

Step 7 **commit**

Step 8 `show rpl route-policy route-policy name detail`**Example:**

```
RP/0/RP0:hostname#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
!
```

Displays the set SPF prefix priority.

Enabling Multicast-intact for OSPFv2

This optional task describes how to enable multicast-intact for OSPFv2 routes that use IPv4 addresses.

Procedure

Step 1 `configure`**Step 2** `router ospf instance-id`**Example:**

```
RP/0/RP0:hostname(config)# router ospf isp
```

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.

Step 3 `mpls traffic-eng multicast-intact`**Example:**

```
RP/0/RP0:hostname(config-ospf)# mpls traffic-eng multicast-intact
```

Enables multicast-intact.

Step 4 `commit`

Associating Interfaces to a VRF

This task explains how to associate an interface with a VPN Routing and Forwarding (VRF) instance.

Procedure

- Step 1** **configure**
Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

- Step 3** **vrf** *vrf-name*

Example:

```
RP/0/RP0:hostname(config-ospf)# vrf vrf1
```

Creates a VRF instance and enters VRF configuration mode.

- Step 4** **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# area 0
```

Enters area configuration mode and configures an area for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.

- Step 5** **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the VRF.

- Step 6** **commit**
-

Configuring OSPF as a Provider Edge to Customer Edge (PE-CE) Protocol

Procedure

- Step 1** **configure**
Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **vrf** *vrf-name*

Example:

```
RP/0/RP0:hostname(config-ospf)# vrf vrf1
```

Creates a VRF instance and enters VRF configuration mode.

Step 4 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# router-id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 5 **redistribute** *protocol* [*process-id*] { **level-1** | **level-1-2** | **level-2** } [**metric** *metric-value*] [**metric-type** *type-value*] [**match** { **external** [**1** | **2**] }] [**tag** *tag-value*] **route-policy** *policy-name*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# redistribute bgp 1 level-1
```

Redistributes OSPF routes from one routing domain to another routing domain.

- This command causes the router to become an ASBR by definition.
- OSPF tags all routes learned through redistribution as external.
- The protocol and its process ID, if it has one, indicate the protocol being redistributed into OSPF.
- The metric is the cost you assign to the external route. The default is 20 for all protocols except BGP, whose default metric is 1.
- The example shows the redistribution of BGP autonomous system 1, Level 1 routes into OSPF as Type 2 external routes.

Step 6 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# area 0
```

Enters area configuration mode and configures an area for the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.

Step 7 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the VRF.

Step 8 `exit`

Example:

```
RP/0/RP0:hostname(config-if)# exit
```

Exits interface configuration mode.

Step 9 `domain-id [secondary] type { 0005 / 0105 / 0205 / 8005 } value value`

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# domain-id type 0105 value 1AF234
```

Specifies the OSPF VRF domain ID.

- The *value* argument is a six-octet hex number.

Step 10 `domain-tag tag`

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# domain-tag 234
```

Specifies the OSPF VRF domain tag.

- The valid range for *tag* is 0 to 4294967295.

Step 11 `disable-dn-bit-check`

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# disable-dn-bit-check
```

Specifies that down bits should be ignored.

Step 12 `commit`

Creating Multiple OSPF Instances (OSPF Process and a VRF)

This task explains how to create multiple OSPF instances. In this case, the instances are a normal OSPF instance and a VRF instance.

Procedure

Step 1 `configure`

Step 2 `router ospf process-name`

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures a backbone area.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 4 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 5 **exit**

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# exit
```

Enters OSPF configuration mode.

Step 6 **vrf** *vrf-name*

Example:

```
RP/0/RP0:hostname(config-ospf)# vrf vrf1
```

Creates a VRF instance and enters VRF configuration mode.

Step 7 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# area 0
```

Enters area configuration mode and configures an area for a VRF instance under the OSPF process.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area.

Step 8 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-vrf)# interface TenGigE0/3/0/5.20
```

Enters interface configuration mode and associates one or more interfaces to the VRF.

Step 9 **commit**

Configuring Multi-area Adjacency

This task explains how to create multiple areas on an OSPF primary interface.

Before you begin



Note You can configure multi-area adjacency on any interface where only two OSF speakers are attached. In the case of native broadcast networks, the interface must be configured as an OPSF point-to-point type using the **network point-to-point** command to enable the interface for a multi-area adjacency.

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 0
```

Enters area configuration mode and configures a backbone area.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 4 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface Serial 0/1/0/3
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 5 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 1
```

Enters area configuration mode and configures an area used for multiple area adjacency.

- The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 6 **multi-area-interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# multi-area-interface Serial 0/1/0/3
```

Enables multiple adjacencies for different OSPF areas and enters multi-area interface configuration mode

Step 7 **commit**

Configuring Authentication Message Digest Management for OSPF

This task explains how to manage authentication of a keychain on the OSPF interface.

Before you begin

A valid keychain must be configured before this task can be attempted.

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router id 192.168.4.3
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 1
```

Enters area configuration mode.

The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 5 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 6 **authentication message-digest keychain** *keychain*

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# authentication message-digest keychain ospf_int1
```

Configures an MD5 keychain.

Note In the example, the *ospf_int1* keychain must be configured before you attempt this step.

Step 7 **commit**

Examples

The following example shows how to configure the keychain *ospf_intf_1* that contains five key IDs. Each key ID is configured with different **send-lifetime** values; however, all key IDs specify the same text string for the key.

```
key chain ospf_intf_1
key 1
send-lifetime 11:30:30 May 1 2007 duration 600
cryptographic-algorithm MD5T
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
```

The following example shows that keychain authentication is enabled on the TenGigE0/6/0/2.10 interface:

```
show ospf 1 interface TenGigE0/3/0/5.20
```



```
TenGigE0/3/0/5.20 is up, line protocol is up
Internet Address 100.10.10.2/24, Area 0
Process ID 1, Router ID 2.2.2.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 2.2.2.1, Interface address 100.10.10.2
Backup Designated router (ID) 1.1.1.1, Interface address 100.10.10.1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:02
Index 3/3, flood queue length 0
Next 0(0)/0(0)
Last flood scan length is 2, maximum is 16
Last flood scan time is 0 msec, maximum is 0 msec
Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 1.1.1.1 (Backup Designated Router)
Suppress hello for 0 neighbor(s)
Keychain-based authentication enabled
  Key id used is 3
Multi-area interface Count is 0
```

The following example shows output for configured keys that are active:

```
show key chain ospf_intf_1

Key-chain: ospf_intf_1/ -

Key 1 -- text "0700325C4836100B0314345D"
  cryptographic-algorithm -- MD5
  Send lifetime: 11:30:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
Key 2 -- text "10411A0903281B051802157A"
  cryptographic-algorithm -- MD5
  Send lifetime: 11:40:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
Key 3 -- text "06091C314A71001711112D5A"
  cryptographic-algorithm -- MD5
  Send lifetime: 11:50:30, 01 May 2007 - (Duration) 600 [Valid now]
  Accept lifetime: Not configured
Key 4 -- text "151D181C0215222A3C350A73"
  cryptographic-algorithm -- MD5
  Send lifetime: 12:00:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
Key 5 -- text "151D181C0215222A3C350A73"
  cryptographic-algorithm -- MD5
  Send lifetime: 12:10:30, 01 May 2007 - (Duration) 600
  Accept lifetime: Not configured
```

Configuring Generalized TTL Security Mechanism (GTSM) for OSPF

This task explains how to set the security time-to-live mechanism on an interface for GTSM.

Procedure

-
- Step 1** **configure**
 - Step 2** **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf 1
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Note The *process-name* argument is any alphanumeric string no longer than 40 characters.

Step 3 **router-id** { *router-id* }

Example:

```
RP/0/RP0:hostname(config-ospf)# router id 10.10.10.100
```

Configures a router ID for the OSPF process.

Note We recommend using a stable IPv4 address as the router ID.

Step 4 **log adjacency changes** [**detail** | **disable**]

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# log adjacency changes detail
```

(Optional) Requests notification of neighbor changes.

- By default, this feature is enabled.
- The messages generated by neighbor changes are considered notifications, which are categorized as severity Level 5 in the **logging console** command. The **logging console** command controls which severity level of messages are sent to the console. By default, all severity level messages are sent.

Step 5 **nsf** { **cisco** [**enforce global**] | **ietf** [**helper disable**] }

Example:

```
RP/0/RP0:hostname(config-ospf)# nsf ietf
```

(Optional) Configures NSF OSPF protocol.

The example enables graceful restart.

Step 6 **timers throttle spf** *spf-start spf-hold spf-max-wait*

Example:

```
RP/0/RP0:hostname(config-ospf)# timers throttle spf 500 500 10000
```

(Optional) Sets SPF throttling timers.

Step 7 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)# area 1
```

Enters area configuration mode.

The *area-id* argument can be entered in dotted-decimal or IPv4 address notation, such as area 1000 or area 0.0.3.232. However, you must choose one form or the other for an area. We recommend using the IPv4 address notation.

Step 8 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 9 **security ttl** [**disable** | **hops** *hop-count*]

Example:

```
RP/0/RP0:hostname(config-ospf-ar-if)# security ttl hopes 2
```

Sets the security TTL value in the IP header for OSPF packets.

Step 10 **commit**

Step 11 **show ospf** [*process-name*] [*area-id*] **interface** [*type interface-path-id*]

Example:

```
RP/0/RP0:hostname# show ospf 1 interface TenGigE0/6/0/2.10
```

Displays OSPF interface information.

Examples

The following is sample output that displays the GTSM security TTL value configured on an OSPF interface:

```
show ospf 1 interface TenGigE0/6/0/2.10

TenGigE0/6/0/2.10 is up, line protocol is up
  Internet Address 120.10.10.1/24, Area 0
  Process ID 1, Router ID 100.100.100.100, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  TTL security enabled, hop count 2
  Designated Router (ID) 102.102.102.102, Interface address 120.10.10.3
  Backup Designated router (ID) 100.100.100.100, Interface address 120.10.10.1
  Flush timer for old DR LSA due in 00:02:36
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:05
  Index 1/1, flood queue length 0
  Next 0(0)/0(0)
  Last flood scan length is 1, maximum is 4
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 102.102.102.102 (Designated Router)
  Suppress hello for 0 neighbor(s)
  Multi-area interface Count is 0
```

Verifying OSPF Configuration and Operation

This task explains how to verify the configuration and operation of OSPF.

Procedure

Step 1 `show { ospf } [process-name]`

Example:

```
RP/0/RP0:hostname# show ospf group1
```

(Optional) Displays general information about OSPF routing processes.

Step 2 `show { ospf } [process-name] border-routers [router-id]`

Example:

```
RP/0/RP0:hostname# show ospf group1 border-routers
```

(Optional) Displays the internal OSPF routing table entries to an ABR and ASBR.

Step 3 `show { ospf } [process-name] database`

Example:

```
RP/0/RP0:hostname# show ospf group2 database
```

(Optional) Displays the lists of information related to the OSPF database for a specific router.

- The various forms of this command deliver information about different OSPF LSAs.

Step 4 `show { ospf } [process-name] [area-id] flood-list interface type interface-path-id`

Example:

```
RP/0/RP0:hostname# show ospf 100 flood-list interface TenGigE0/6/0/2.10
```

(Optional) Displays a list of OSPF LSAs waiting to be flooded over an interface.

Step 5 `show { ospf } [process-name] [area-id] interface [type interface-path-id]`

Example:

```
RP/0/RP0:hostname# show ospf 100 interface TenGigE0/6/0/2.10
```

(Optional) Displays OSPF interface information.

Step 6 `show { ospf } [process-name] [area-id] neighbor [type interface-path-id] [neighbor-id] [detail]`

Example:

```
RP/0/RP0:hostname# show ospf 100 neighbor
```

(Optional) Displays OSPF neighbor information on an individual interface basis.

Step 7 `clear { ospf } [process-name] process`

Example:

```
RP/0/  
/CPU0:router# clear ospf 100 process
```

(Optional) Resets an OSPF router process without stopping and restarting it.

Step 8 **clear {ospf[*process-name*] redistribution**

Example:

```
RP/0/RP0:hostname#clear ospf 100 redistribution
```

Clears OSPF route redistribution.

Step 9 **clear {ospf[*process-name*] routes**

Example:

```
RP/0/RP0:hostname#clear ospf 100 routes
```

Clears OSPF route table.

Step 10 **clear {ospf[*process-name*] vrf [*vrf-name* | all] {process | redistribution | routes | statistics [interface *type* | interface-path-id | message-queue | neighbor] }**

Example:

```
RP/0/RP0:hostname#clear ospf 100 vrf vrf_1 process
```

Clears OSPF route table.

Step 11 **clear { ospf } [*process-name*] statistics [neighbor [*type* | interface-path-id] [ip-address]]**

Example:

```
RP/0/RP0:hostname# clear ospf 100 statistics
```

(Optional) Clears the OSPF statistics of neighbor state transitions.

Configuring IP Fast Reroute Loop-free Alternate

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.

To enable protection on broadcast links, IPFRR and bidirectional forwarding detection (BFD) must be enabled on the interface under OSPF.

Enabling IPFRR LFA

Procedure

Step 1 **configure**

Step 2 **router ospf *process-name***

Example:

```
RP/0/RP0:hostname(config)# router ospf
```

Enables OSPF routing for the specified routing process and places the router in router configuration mode.

Step 3 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config-ospf)#area 1
```

Enters area configuration mode.

Step 4 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# interface TenGigE0/6/0/2.10
```

Enters interface configuration mode and associates one or more interfaces to the area. .

Step 5 **fast-reroute per-link** { **enable** | **disable** }

Example:

```
RP/0/RP0:hostname(config-ospf-ar)#fast-reroute per-link enable
```

Enables or disables per-link LFA computation for the interface.

Step 6 **commit**

Excluding an Interface From IP Fast Reroute Per-link Computation

Procedure

Step 1 **configure**

Step 2 **router ospf** *process-name*

Example:

```
RP/0/RP0:hostname(config)# router ospf
```

Enables the OSPF routing for the specified routing process and places the router in router configuration mode.

Step 3 **area** *area-id*

Example:

```
RP/0/RP0:hostname(config)#area area-id
```

Enters area configuration mode.

Step 4 **interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf)#interface type interface-path-id
```

Enters interface configuration mode and associates one or more interfaces to the area.

Step 5 **fast-reroute per-link exclude interface** *type interface-path-id*

Example:

```
RP/0/RP0:hostname(config-ospf-ar)# fast-reroute per-link exclude interface TenGigE0/6/0/2.10
```

Excludes an interface from IP fast reroute per-link computation.

Step 6 **commit**

Enabling OSPF Interaction with SRMS Server

To enable OSPF interaction with SRMS server:

Procedure

Step 1 **configure**

Step 2 **router ospf** *instance-id*

Example:

```
RP/0/RP0:hostname(config)# router ospf isp
```

Enables OSPF routing for the specified routing instance, and places the router in router configuration mode.

Step 3 **segment-routing mpls**

Example:

```
RP/0/RP0:hostname(config-ospf)# segment-routing mpls
```

Step 4 **segment-routing forwarding mpls**

Example:

```
RP/0/RP0:hostname(config-ospf)# segment-routing forwarding mpls
```

Enables SR forwarding on all interfaces where this instance OSPF is enabled.

Step 5 **segment-routing prefix-sid-map advertise local**

Example:

```
RP/0/RP0:hostname(config-ospf)# segment-routing
prefix-sid-map advertise local
```

Enables server functionality and allows OSPF to advertise the local mapping entries using area-scope flooding. The flooding is limited to areas where segment-routing is enabled. Disabled by default.

Step 6 **segment-routing sr-prefer prefix-list** [*acl-name*]

Example:

```
RP/0/RP0:hostname(config-ospf)# segment-routing
sr-prefer prefix-list foo
```

Configuration Examples for Implementing OSPF

This section provides the following configuration examples:

Cisco IOS XR Software for OSPF Version 2 Configuration: Example

The following example shows how an OSPF interface is configured for an area in Cisco IOS XR Software.

area 0 must be explicitly configured with the **area** command and all interfaces that are in the range from 10.1.2.0 to 10.1.2.255 are bound to area 0. Interfaces are configured with the **interface** command (while the router is in area configuration mode) and the **area** keyword is not included in the interface statement.

Cisco IOS XR Software Configuration

```
interface TenGigE0/3/0/2.10
 ip address 10.1.2.1 255.255.255.255
 negotiation auto
!
router ospf 1
router-id 10.2.3.4
 area 0
  interface TenGigE0/3/0/2.10
!
!
```

The following example shows how OSPF interface parameters are configured for an area in Cisco IOS XR software.

In Cisco IOS XR software, OSPF interface-specific parameters are configured in interface configuration mode and explicitly defined for area 0. In addition, the **ip ospf** keywords are no longer required.

Cisco IOS XR Software Configuration

```
interface TenGigE0/3/0/2.10
 ip address 10.1.2.1 255.255.255.0
 negotiation auto
!
router ospf 1
router-id 10.2.3.4
area 0
 interface TenGigE0/3/0/2.10
  cost 77
  mtu-ignore
  authentication message-digest
  message-digest-key 1 md5 0 test
!
!
```

The following example shows the hierarchical CLI structure of Cisco IOS XR software:

In Cisco IOS XR software, 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.

Cisco IOS XR Software Configuration

```
interface TenGigE0/3/0/2.10
 ip address 10.1.2.1 255.255.255.0
 negotiation auto
!
interface TenGigE0/3/0/5.20
 ip address 10.1.3.1 255.255.255.0
 negotiation auto
```



```

!
router ospf 1
  router-id 10.2.3.4
  area 0
    interface TenGigE0/3/0/2.10
  !
  area 1
    interface TenGigE0/3/0/5.20
  !
!

```

MPLS TE for OSPF Version 2: Example

The following example shows how to configure the OSPF portion of MPLS TE. However, you still need to build an MPLS TE topology and create an MPLS TE tunnel.

In this example, loopback interface 0 is associated with area 0 and MPLS TE is configured within area 0.

```

interface Loopback 0
  address 10.10.10.10 255.255.255.0
!
interface TenGigE0/3/0/2.10
  address 10.1.2.2 255.255.255.0
!
router ospf 1
  router-id 10.10.10.10
  nsf
  auto-cost reference-bandwidth 10000
  mpls traffic-eng router-id Loopback 0
  area 0
    mpls traffic-eng
    interface TenGigE0/3/0/2.10
    interface Loopback 0

```

Virtual Link Configured with MD5 Authentication for OSPF Version 2: Example

The following examples show how to configure a virtual link to your backbone and apply MD5 authentication. You must perform the steps described on both ABRs at each end of the virtual link.

After you explicitly configure the ABRs, the configuration is inherited by all interfaces bound to that area—unless you override the values and configure them explicitly for the interface.

To understand virtual links, see [Virtual Link and Transit Area for OSPF, on page 12](#).

In this example, all interfaces on router ABR1 use MD5 authentication:

```

router ospf ABR1
  router-id 10.10.10.10
  authentication message-digest
  message-digest-key 100 md5 0 cisco
  area 0
    interface TenGigE0/3/0/2.10
    interface TenGigE0/6/0/5.20
  area 1
    interface TenGigE0/14/0/4.40
    virtual-link 10.10.5.5
!
!

```

In this example, only area 1 interfaces on router ABR3 use MD5 authentication:

```
router ospf ABR2
router-id 10.10.5.5
area 0
area 1
  authentication message-digest
  message-digest-key 100 md5 0 cisco
interface TenGigE0/3/0/5.20
virtual-link 10.10.10.10
area 3
interface Loopback 0
interface TenGigE0/3/0/2.10
!
```