

Configuring OSPF

This chapter describes how to configure OSPF. For a complete description of the OSPF commands in this chapter, refer to the “OSPF Commands” chapter of the *Network Protocols Command Reference, Part 1*. To locate documentation of other commands that appear in this chapter, use the command reference master index or search online.

Open shortest path first (OSPF) is an 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/receiving packets.

We support RFC 1253, Open Shortest Path First (OSPF) MIB, August 1991. The OSPF MIB defines an IP routing protocol that provides management information related to OSPF and is supported by Cisco routers.

For protocol-independent features, see the chapter “Configuring IP Routing Protocol-Independent Features” in this document.

Cisco's OSPF Implementation

Cisco's implementation conforms to the OSPF Version 2 specifications detailed in the Internet RFC 1583. The list that follows outlines key features supported in Cisco's OSPF implementation:

- Stub areas—Definition of stub areas is supported.
- Route redistribution—Routes learned via any IP routing protocol can be redistributed into any other IP routing protocol. At the intradomain level, this means that OSPF can import routes learned via IGRP, RIP, and IS-IS. OSPF routes can also be exported into IGRP, RIP, and IS-IS. At the interdomain level, OSPF can import routes learned via EGP and BGP. OSPF routes can be exported into EGP and BGP.
- 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.
- NSSA areas—RFC 1587.
- OSPF over demand circuit—RFC 1793.

Note To take advantage of the OSPF stub area support, *default routing* must be used in the stub area.

OSPF Configuration Task List

OSPF typically requires coordination among many internal routers, *area border routers* (routers connected to multiple areas), and autonomous system boundary routers. 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.

To configure OSPF, complete the tasks in the following sections. Enabling OSPF is mandatory; the other tasks are optional, but might be required for your application.

- Enable OSPF
- Configure OSPF Interface Parameters
- Configure OSPF over Different Physical Networks
- Configure OSPF Area Parameters
- Configure OSPF Not So Stubby Area (NSSA)
- Configure Route Summarization between OSPF Areas
- Configure Route Summarization when Redistributing Routes into OSPF
- Create Virtual Links
- Generate a Default Route
- Configure Lookup of DNS Names
- Force the Router ID Choice with a Loopback Interface
- Control Default Metrics
- Configure OSPF on Simplex Ethernet Interfaces
- Configure Route Calculation Timers
- Configure OSPF over On Demand Circuits
- Log Neighbor Changes
- Monitor and Maintain OSPF

In addition, you can specify route redistribution; see the task “Redistribute Routing Information” in the chapter “Configuring IP Routing Protocol-Independent Features” for information on how to configure route redistribution.

Enable OSPF

As with other routing protocols, enabling OSPF requires that you create an OSPF routing process, specify the range of IP addresses to be associated with the routing process, and assign area IDs to be associated with that range of IP addresses. Perform the following tasks, starting in global configuration mode:

Task	Command
Step 1 Enable OSPF routing, which places you in router configuration mode.	router ospf <i>process-id</i>
Step 2 Define an interface on which OSPF runs and define the area ID for that interface.	network <i>address wildcard-mask area</i> <i>area-id</i>

Configure OSPF Interface Parameters

Our OSPF implementation allows you to alter certain interface-specific OSPF parameters, as needed. You are not required to alter any of these parameters, but some interface parameters must be consistent across all routers in an attached network. Those parameters are controlled by the **ip ospf hello-interval**, **ip ospf dead-interval**, and **ip ospf authentication-key** commands. Therefore, be sure that if you do configure any of these parameters, the configurations for all routers on your network have compatible values.

In interface configuration mode, specify any of the following interface parameters as needed for your network:

Task	Command
Explicitly specify the cost of sending a packet on an OSPF interface.	ip ospf cost <i>cost</i>
Specify the number of seconds between link state advertisement retransmissions for adjacencies belonging to an OSPF interface.	ip ospf retransmit-interval <i>seconds</i>
Set the estimated number of seconds it takes to transmit a link state update packet on an OSPF interface.	ip ospf transmit-delay <i>seconds</i>
Set priority to help determine the OSPF designated router for a network.	ip ospf priority <i>number</i>
Specify the length of time, in seconds, between the hello packets that the Cisco IOS software sends on an OSPF interface.	ip ospf hello-interval <i>seconds</i>
Set the number of seconds that a device's hello packets must not have been seen before its neighbors declare the OSPF router down.	ip ospf dead-interval <i>seconds</i>
Assign a specific password to be used by neighboring OSPF routers on a network segment that is using OSPF's simple password authentication.	ip ospf authentication-key <i>key</i>
Enable OSPF MD5 authentication.	ip ospf message-digest-key <i>keyid md5 key</i>

Configure OSPF over Different Physical Networks

OSPF classifies different media into the following three types of networks by default:

- Broadcast networks (Ethernet, Token Ring, FDDI)
- Nonbroadcast multiaccess networks (SMDS, Frame Relay, X.25)
- Point-to-point networks (HDLC, PPP)

You can configure your network as either a broadcast or a nonbroadcast multiaccess network.

X.25 and Frame Relay provide an optional broadcast capability that can be configured in the map to allow OSPF to run as a broadcast network. See the **x25 map** and **frame-relay map** command descriptions in the *Wide-Area Networking Command Reference* for more detail.

Configure Your OSPF Network Type

You have the choice of configuring your OSPF network type as either broadcast or nonbroadcast multiaccess, regardless of the default media type. Using this feature, you can configure broadcast networks as nonbroadcast multiaccess networks when, for example, you have routers in your network that do not support multicast addressing. You also can configure nonbroadcast multiaccess networks (such as X.25, Frame Relay, and SMDS) as broadcast networks. This feature saves you from having to configure neighbors, as described in the section “Configure OSPF for Nonbroadcast Networks.”

Configuring nonbroadcast, multiaccess networks as either broadcast or nonbroadcast assumes that there are virtual circuits from every router to every router or fully meshed network. This is not true for some cases, for example, because of cost constraints, or when you have only a partially meshed network. In these cases, you can configure the OSPF network type as a point-to-multipoint network. Routing between two routers not directly connected will go through the router that has virtual circuits to both routers. Note that you must not configure neighbors when using this feature.

An OSPF point-to-multipoint interface is defined as a numbered point-to-point interface having one or more neighbors. It creates multiple host routes. An OSPF point-to-multipoint network has the following benefits compared to nonbroadcast multiaccess and point-to-point networks:

- Point-to-multipoint is easier to configure because it requires no configuration of neighbor commands, it consumes only one IP subnet, and it requires no designated router election.
- It costs less because it does not require a fully meshed topology.
- It is more reliable because it maintains connectivity in the event of virtual circuit failure.

To configure your OSPF network type, perform the following task in interface configuration mode:

Task	Command
Configure the OSPF network type for a specified interface.	ip ospf network { broadcast non-broadcast point-to-multipoint }

See the “OSPF Point-to-Multipoint Example” section at the end of this chapter for an example of an OSPF point-to-multipoint network.

Configure OSPF for Nonbroadcast Networks

Because there might be many routers attached to an OSPF network, a *designated router* is selected for the network. It is necessary to use special configuration parameters in the designated router selection if broadcast capability is not configured.

These parameters need only be configured in those devices that are themselves eligible to become the designated router or backup designated router (in other words, routers or access servers with a nonzero router priority value).

To configure routers that interconnect to nonbroadcast networks, perform the following task in router configuration mode:

Task	Command
Configure routers or access servers interconnecting to nonbroadcast networks.	neighbor <i>ip-address</i> [priority <i>number</i>] [poll-interval <i>seconds</i>]

You can specify the following neighbor parameters, as required:

- Priority for a neighboring router
- Nonbroadcast poll interval
- Interface through which the neighbor is reachable

Configure OSPF Area Parameters

Our OSPF software allows you to configure several area parameters. These area parameters, shown in the following table, include authentication, defining stub areas, and assigning specific costs to the default summary route. *Authentication* allows password-based protection against unauthorized access to an area.

Stub areas are areas into which information on external routes is not sent. Instead, there is a default external route generated by the area border router, into the stub area for destinations outside the autonomous system. To further reduce the number of link state advertisements sent into a stub area, you can configure **no-summary** on the ABR to prevent it from sending summary link advertisement (link state advertisements Type 3) into the stub area.

In router configuration mode, specify any of the following area parameters as needed for your network:

Task	Command
Enable authentication for an OSPF area.	area <i>area-id</i> authentication
Enable MD5 authentication for an OSPF area.	area <i>area-id</i> authentication message-digest
Define an area to be a stub area.	area <i>area-id</i> stub [no-summary]
Assign a specific cost to the default summary route used for the stub area.	area <i>area-id</i> default-cost <i>cost</i>

Configure OSPF Not So Stubby Area (NSSA)

NSSA area is similar to OSPF stub area. NSSA does not flood Type 5 external link state advertisements (LSAs) from the core into the area, but it has the ability of importing AS external routes in a limited fashion within the area.

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

Use NSSA to simplify administration if you are an Internet service provider (ISP), or a network administrator that must connect a central site using OSPF to a remote site that is using a different routing protocol.

Prior to NSSA, the connection between the corporate site border router and the remote router could not be run as OSPF stub area because routes for the remote site cannot be redistributed into stub area. A simple protocol like RIP is usually run and handle the redistribution. This meant maintaining two routing protocols. With NSSA, you can extend OSPF to cover the remote connection by defining the area between the corporate router and the remote router as an NSSA.

In router configuration mode, specify the following area parameters as needed to configure OSPF NSSA:

Task	Command
Define an area to be NSSA.	area <i>area-id</i> nssa [no-redistribution] [default-information-originate]

In router configuration mode on the ABR, specify the following command to control summarization and filtering of Type 7 LSA into Type 5 LSA:

Task	Command
(Optional) Control the summarization and filtering during the translation.	summary <i>address prefix mask</i> [not advertise] [tag <i>tag</i>]

Implementation Considerations

Evaluate the following considerations before implementing this feature:

- You can set a Type 7 default route that can be used to reach external destinations. When configured, the router generates a Type 7 default into the NSSA by the NSSA ABR.
- Every router within the same area must agree that the area is NSSA; otherwise, the routers will not be able to communicate with each other.

If possible, avoid using explicit redistribution on NSSA ABR because confusion may result over which packets are being translated by which router.

Configure Route Summarization between OSPF Areas

Route summarization is the consolidation of advertised addresses. This feature causes a single summary route to be advertised to other areas by an ABR. In OSPF, an ABR will advertise networks in one area into another area. If the network numbers in an area are assigned in a way such that they are contiguous, you can configure the ABR to advertise a summary route that covers all the individual networks within the area that fall into the specified range.

To specify an address range, perform the following task in router configuration mode:

Task	Command
Specify an address range for which a single route will be advertised.	area <i>area-id</i> range <i>address mask</i> [advertise not-advertise]

Configure Route Summarization when Redistributing Routes into OSPF

When redistributing routes from other protocols into OSPF (as described in the chapter “Configuring IP Routing Protocol-Independent Features”), each route is advertised individually in an external link state advertisement (LSA). However, you can configure the Cisco IOS software to advertise a single route for all the redistributed routes that are covered by a specified network address and mask. Doing so helps decrease the size of the OSPF link state database.

To have the software advertise one summary route for all redistributed routes covered by a network address and mask, perform the following task in router configuration mode:

Task	Command
Specify an address and mask that covers redistributed routes, so only one summary route is advertised.	summary-address <i>address mask</i>

Create Virtual Links

In OSPF, all areas must be connected to a backbone area. If there is a break in backbone continuity, or the backbone is purposefully partitioned, you can establish a *virtual link*. The two end points of a virtual link are Area Border Routers. The virtual link must be configured in both routers. The configuration information in each router consists of the other virtual endpoint (the other ABR), and the nonbackbone area that the two routers have in common (called the *transit area*). Note that virtual links cannot be configured through stub areas.

To establish a virtual link, perform the following task in router configuration mode:

Task	Command
Establish a virtual link.	area <i>area-id</i> virtual-link <i>router-id</i> [hello-interval <i>seconds</i>] [retransmit-interval <i>seconds</i>] [transmit-delay <i>seconds</i>] [dead-interval <i>seconds</i>] [[authentication-key <i>key</i>] [message-digest-key <i>keyid md5 key</i>]]

To display information about virtual links, use the **show ip ospf virtual-links** EXEC command. To display the router ID of an OSPF router, use the **show ip ospf** EXEC command.

Generate a Default Route

You can force an autonomous system boundary router to generate a default route into an OSPF routing domain. Whenever you specifically configure redistribution of routes into an OSPF routing domain, the router automatically becomes an autonomous system boundary router. However, an autonomous system boundary router does not, by default, generate a *default route* into the OSPF routing domain.

To force the autonomous system boundary router to generate a default route, perform the following task in router configuration mode:

Task	Command
Force the autonomous system boundary router to generate a default route into the OSPF routing domain.	default-information originate [always] [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [route-map <i>map-name</i>]

See the discussion of redistribution of routes in the “Configuring IP Routing Protocol-Independent Features” chapter.

Configure Lookup of DNS Names

You can configure OSPF to look up Domain Naming System (DNS) names for use in all OSPF **show** command displays. This feature makes it easier to identify a router, because it is displayed by name rather than by its router ID or neighbor ID.

To configure DNS name lookup, perform the following task in global configuration mode:

Task	Command
Configure DNS name lookup.	ip ospf name-lookup

Force the Router ID Choice with a Loopback Interface

OSPF uses the largest IP address configured on the interfaces as its router ID. If the interface associated with this IP address is ever brought down, or if the address is removed, the OSPF process must recalculate a new router ID and resend all its routing information out its interfaces.

If a loopback interface is configured with an IP address, the Cisco IOS software will use this IP address as its router ID, even if other interfaces have larger IP addresses. Since loopback interfaces never go down, greater stability in the routing table is achieved.

OSPF automatically prefers a loopback interface over any other kind, and it chooses the highest IP address among all loopback interfaces. If no loopback interfaces are present, the highest IP address in the router is chosen. You cannot tell OSPF to use any particular interface.

To configure an IP address on a loopback interface, perform the following tasks, starting in global configuration mode:

Task	Command
Step 1 Create a loopback interface, which places you in interface configuration mode.	interface loopback 0
Step 2 Assign an IP address to this interface.	ip address <i>address mask</i>

Control Default Metrics

In Cisco IOS Release 10.3 and later, by default, OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface. For example, a 64K link gets a metric of 1562, while a T1 link gets a metric of 64.

The OSPF metric is calculated as *ref-bw* divided by *bandwidth*, with *ref-bw* equal to 10^8 by default, and *bandwidth* determined by the **bandwidth** command. The calculation gives FDDI a metric of 1. If you have multiple links with high bandwidth, you might want to specify a larger number to differentiate the cost on those links. To do so, perform the following task in router configuration mode:

Task	Command
Differentiate high bandwidth links.	ospf auto-cost reference-bandwidth <i>ref-bw</i>

Configure OSPF on Simplex Ethernet Interfaces

Because simplex interfaces between two devices on an Ethernet represent only one network segment, for OSPF you must configure the transmitting interface to be a passive interface. This prevents OSPF from sending hello packets for the transmitting interface. Both devices are able to see each other via the hello packet generated for the receiving interface.

To configure OSPF on simplex Ethernet interfaces, perform the following task in router configuration mode:

Task	Command
Suppress the sending of hello packets through the specified interface.	passive-interface <i>type number</i>

Configure Route Calculation Timers

You can configure the delay time between when OSPF receives a topology change and when it starts a shortest path first (SPF) calculation. You can also configure the hold time between two consecutive SPF calculations. To do this, perform the following task in router configuration mode:

Task	Command
Configure route calculation timers.	timers spf <i>spf-delay spf-holdtime</i>

Configure OSPF over On Demand Circuits

The OSPF on demand circuit is an enhancement to the OSPF protocol that allows efficient operation over on demand circuits like ISDN, X.25 SVCs and dial-up lines. This feature supports RFC 1793, *Extending OSPF to Support Demand Circuits*.

Prior to this feature, OSPF periodic hello and link state advertisements (LSAs) updates would be exchanged between routers that connected the on demand link, even when no changes occurred in the hello or LSA information.

With this feature, periodic hellos are suppressed and the periodic refreshes of LSAs are not flooded over the demand circuit. These packets bring up the link only when they are exchanged for the first time, or when a change occurs in the information they contain. This operation allows the underlying datalink layer to be closed when the network topology is stable.

This feature is useful when you want to connect telecommuters or branch offices to an OSPF backbone at a central site. In this case, OSPF for on demand circuits allows the benefits of OSPF over the entire domain, without excess connection costs. Periodic refreshes of hello updates, LSA updates, and other protocol overhead are prevented from enabling the on demand circuit when there is no “real” data to transmit.

Overhead protocols such as hellos and LSAs are transferred over the on demand circuit only upon initial setup and when they reflect a change in the topology. This means that critical changes to the topology that require new SPF calculations are transmitted in order to maintain network topology integrity. Periodic refreshes that do not include changes, however, are not transmitted across the link.

To configure OSPF for on demand circuits, perform the following tasks, beginning in global configuration mode:

Task	Command
Step 1 Enable OSPF operation.	router ospf <i>process-id</i>
Step 2 Configure OSPF on an on demand circuit.	ip ospf demand-circuit

If the router is part of a point-to-point topology, then only one end of the demand circuit must be configured with this command. However, all routers must have this feature loaded.

If the router is part of a point-to-multipoint topology, only the multipoint end must be configured with this command.

Implementation Considerations

Evaluate the following considerations before implementing this feature:

- Because LSAs that include topology changes are flooded over an on demand circuit, it is advised to put demand circuits within OSPF stub areas, or within NSSAs to isolate the demand circuits from as many topology changes as possible.
- To take advantage of the on demand circuit functionality within a stub area or NSSA, every router in the area must have this feature loaded. If this feature is deployed within a regular area, all other regular areas must also support this feature before the demand circuit functionality can take effect. This is because type 5 external LSAs are flooded throughout all areas.
- You do not want to do on a broadcast-based network topology because the overhead protocols (such as hellos and LSAs) cannot be successfully suppressed, which means the link will remain up.

Log Neighbor Changes

To configure the router to send a syslog message when an OSPF neighbor state changes, perform the following task in router configuration mode:

Task	Command
Send syslog message when a neighbor state changes.	ospf log-adj-changes

Configure this command if you want to know about OSPF neighbor changes without turning on the debugging command **debug ip ospf adjacency**. The **ospf log-adj-changes** command provides a higher level view of changes to the state of the peer relationship with less output.

Monitor and Maintain OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases. Information provided can be used to determine resource utilization and solve network problems. You can also display information about node reachability and discover the routing path your device's packets are taking through the network.

To display various routing statistics, perform the following tasks in EXEC mode:

Task	Command
Display general information about OSPF routing processes.	show ip ospf [<i>process-id</i>]

Task	Command
Display lists of information related to the OSPF database.	show ip ospf [<i>process-id area-id</i>] database show ip ospf [<i>process-id area-id</i>] database [router] [<i>link-state-id</i>] show ip ospf [<i>process-id area-id</i>] database [network] [<i>link-state-id</i>] show ip ospf [<i>process-id area-id</i>] database [summary] [<i>link-state-id</i>] show ip ospf [<i>process-id area-id</i>] database [asb-summary] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [external] [<i>link-state-id</i>] show ip ospf [<i>process-id area-id</i>] database [database-summary]
Display the internal OSPF routing table entries to Area Border Router (ABR) and Autonomous System Boundary Router (ASBR).	show ip ospf border-routers
Display OSPF-related interface information.	show ip ospf interface [<i>interface-name</i>]
Display OSPF-neighbor information on a per-interface basis.	show ip ospf neighbor [<i>interface-name</i>] [<i>neighbor-id</i>] detail
Display a list of all LSAs requested by a router.	show ip ospf request-list [<i>nbr</i>] [<i>intf</i>] [<i>intf-nbr</i>]
Display a list of all LSAs waiting to be retransmitted.	show ip ospf retransmission-list [<i>nbr</i>] [<i>intf</i>] [<i>intf-nbr</i>]
Display OSPF-related virtual links information.	show ip ospf virtual-links

OSPF Configuration Examples


The following sections provide OSPF configuration examples:

- OSPF Point-to-Multipoint Example
- Variable-Length Subnet Masks Example
- OSPF Routing and Route Redistribution Examples
- Route Map Examples

OSPF Point-to-Multipoint Example

In Figure 20, Mollie uses DLCI 201 to communicate with Neon, DLCI 202 to Jelly, and DLCI 203 to Platty. Neon uses DLCI 101 to communicate with Mollie and DLCI 102 to communicate with Platty. Platty communicates with Neon (DLCI 401) and Mollie (DLCI 402). Jelly communicates with Mollie (DLCI 301).

Figure 20 OSPF Point-to-Multipoint Example



Mollie's Configuration

```
hostname mollie
!
interface serial 1
 ip address 10.0.0.2 255.0.0.0
 ip ospf network point-to-multipoint
 encapsulation frame-relay
 frame-relay map ip 10.0.0.1 201 broadcast
 frame-relay map ip 10.0.0.3 202 broadcast
 frame-relay map ip 10.0.0.4 203 broadcast
!
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
```

Neon's Configuration

```
hostname neon
!
interface serial 0
 ip address 10.0.0.1 255.0.0.0
 ip ospf network point-to-multipoint
 encapsulation frame-relay
 frame-relay map ip 10.0.0.2 101 broadcast
 frame-relay map ip 10.0.0.4 102 broadcast
!
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
```

Platty's Configuration

```
hostname platty
!
interface serial 3
 ip address 10.0.0.4 255.0.0.0
 ip ospf network point-to-multipoint
 encapsulation frame-relay
 clock rate 1000000
 frame-relay map ip 10.0.0.1 401 broadcast
 frame-relay map ip 10.0.0.2 402 broadcast
!
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
```

Jelly's Configuration

```
hostname jelly
!
interface serial 2
 ip address 10.0.0.3 255.0.0.0
 ip ospf network point-to-multipoint
 encapsulation frame-relay
 clock rate 2000000
 frame-relay map ip 10.0.0.2 301 broadcast
!
router ospf 1
 network 10.0.0.0 0.0.0.255 area 0
```

Variable-Length Subnet Masks Example

OSPF, static routes, and IS-IS support variable-length subnet masks (VLSMs). With VLSMs, you can use different masks for the same network number on different interfaces, which allows you to conserve IP addresses and more efficiently use available address space.

In the following example, a 30-bit subnet mask is used, leaving two bits of address space reserved for serial line host addresses. There is sufficient host address space for two host endpoints on a point-to-point serial link.

```
interface ethernet 0
 ip address 131.107.1.1 255.255.255.0
! 8 bits of host address space reserved for ethernetets

interface serial 0
 ip address 131.107.254.1 255.255.255.252
! 2 bits of address space reserved for serial lines

! Router is configured for OSPF and assigned AS 107
router ospf 107
! Specifies network directly connected to the router
 network 131.107.0.0 0.0.255.255 area 0.0.0.0
```

OSPF Routing and Route Redistribution Examples

OSPF typically requires coordination among many internal routers, area border routers, and autonomous system boundary routers. At a minimum, OSPF-based routers can be configured with all default parameter values, with no authentication, and with interfaces assigned to areas.

Three examples follow:

- The first is a simple configuration illustrating basic OSPF commands.
- The second example illustrates a configuration for an internal router, ABR, and ASBRs within a single, arbitrarily assigned, OSPF autonomous system.
- The third example illustrates a more complex configuration and the application of various tools available for controlling OSPF-based routing environments.

Basic OSPF Configuration Example

The following example illustrates a simple OSPF configuration that enables OSPF routing process 9000, attaches Ethernet 0 to area 0.0.0.0, and redistributes RIP into OSPF, and OSPF into RIP:

```
interface ethernet 0
 ip address 130.93.1.1 255.255.255.0
 ip ospf cost 1
```

```
!
interface ethernet 1
 ip address 130.94.1.1 255.255.255.0
!
router ospf 9000
 network 130.93.0.0 0.0.255.255 area 0.0.0.0
 redistribute rip metric 1 subnets
!
router rip
 network 130.94.0.0
 redistribute ospf 9000
 default-metric 1
```

Basic OSPF Configuration Example for Internal Router, ABR, and ASBRs

The following example illustrates the assignment of four area IDs to four IP address ranges. In the example, OSPF routing process 109 is initialized, and four OSPF areas are defined: 10.9.50.0, 2, 3, and 0. Areas 10.9.50.0, 2, and 3 mask specific address ranges, while Area 0 enables OSPF for *all other* networks.

```
router ospf 109
 network 131.108.20.0 0.0.0.255 area 10.9.50.0
 network 131.108.0.0 0.0.255.255 area 2
 network 131.109.10.0 0.0.0.255 area 3
 network 0.0.0.0 255.255.255.255 area 0
!
! Interface Ethernet0 is in area 10.9.50.0:
interface ethernet 0
 ip address 131.108.20.5 255.255.255.0
!
! Interface Ethernet1 is in area 2:
interface ethernet 1
 ip address 131.108.1.5 255.255.255.0
!
! Interface Ethernet2 is in area 2:
interface ethernet 2
 ip address 131.108.2.5 255.255.255.0
!
! Interface Ethernet3 is in area 3:
interface ethernet 3
 ip address 131.109.10.5 255.255.255.0
!
! Interface Ethernet4 is in area 0:
interface ethernet 4
 ip address 131.109.1.1 255.255.255.0
!
! Interface Ethernet5 is in area 0:
interface ethernet 5
 ip address 10.1.0.1 255.255.0.0
```

Each **network area** router configuration command is evaluated sequentially, so the order of these commands in the configuration is important. The Cisco IOS software sequentially evaluates the *address/wildcard-mask* pair for each interface. See the “OSPF Commands” chapter of the *Network Protocols Command Reference, Part 1* for more information.

Consider the first **network area** command. Area ID 10.9.50.0 is configured for the interface on which subnet 131.108.20.0 is located. Assume that a match is determined for interface Ethernet 0. Interface Ethernet 0 is attached to Area 10.9.50.0 only.

The second **network area** command is evaluated next. For Area 2, the same process is then applied to all interfaces (except interface Ethernet 0). Assume that a match is determined for interface Ethernet 1. OSPF is then enabled for that interface and Ethernet 1 is attached to Area 2.

This process of attaching interfaces to OSPF areas continues for all **network area** commands. Note that the last **network area** command in this example is a special case. With this command, all available interfaces (not explicitly attached to another area) are attached to Area 0.

Complex Internal Router, ABR, and ASBRs Example

The following example outlines a configuration for several routers within a single OSPF autonomous system. Figure 21 provides a general network map that illustrates this example configuration.

Figure 21 Sample OSPF Autonomous System Network Map

In this configuration, five routers are configured in OSPF autonomous system 109:

- Router A and Router B are both internal routers within Area 1.
- Router C is an OSPF area border router. Note that for Router C, Area 1 is assigned to E3 and Area 0 is assigned to S0.
- Router D is an internal router in Area 0 (backbone area). In this case, both **network** router configuration commands specify the same area (Area 0, or the backbone area).
- Router E is an OSPF autonomous system boundary router. Note that BGP routes are redistributed into OSPF and that these routes are advertised by OSPF.

Note It is not necessary to include definitions of all areas in an OSPF autonomous system in the configuration of all routers in the autonomous system. You must only define the *directly* connected areas. In the example that follows, routes in Area 0 are learned by the routers in Area 1 (Router A and Router B) when the area border router (Router C) injects summary link state advertisements (LSAs) into Area 1.

Autonomous system 109 is connected to the outside world via the BGP link to the external peer at IP address 11.0.0.6.

Router A—Internal Router

```
interface ethernet 1
 ip address 131.108.1.1 255.255.255.0

router ospf 109
 network 131.108.0.0 0.0.255.255 area 1
```

Router B—Internal Router

```
interface ethernet 2
 ip address 131.108.1.2 255.255.255.0

router ospf 109
 network 131.108.0.0 0.0.255.255 area 1
```

Router C—ABR

```
interface ethernet 3
 ip address 131.108.1.3 255.255.255.0

interface serial 0
 ip address 131.108.2.3 255.255.255.0

router ospf 109
 network 131.108.1.0 0.0.0.255 area 1
 network 131.108.2.0 0.0.0.255 area 0
```

Router D—Internal Router

```
interface ethernet 4
 ip address 10.0.0.4 255.0.0.0

interface serial 1
 ip address 131.108.2.4 255.255.255.0
```



```
router ospf 109
 network 131.108.2.0 0.0.0.255 area 0
 network 10.0.0.0 0.255.255.255 area 0
```

Router E—ASBR

```
interface ethernet 5
 ip address 10.0.0.5 255.0.0.0

interface serial 2
 ip address 11.0.0.5 255.0.0.0

router ospf 109
 network 10.0.0.0 0.255.255.255 area 0
 redistribute bgp 109 metric 1 metric-type 1

router bgp 109
 network 131.108.0.0
 network 10.0.0.0
 neighbor 11.0.0.6 remote-as 110
```

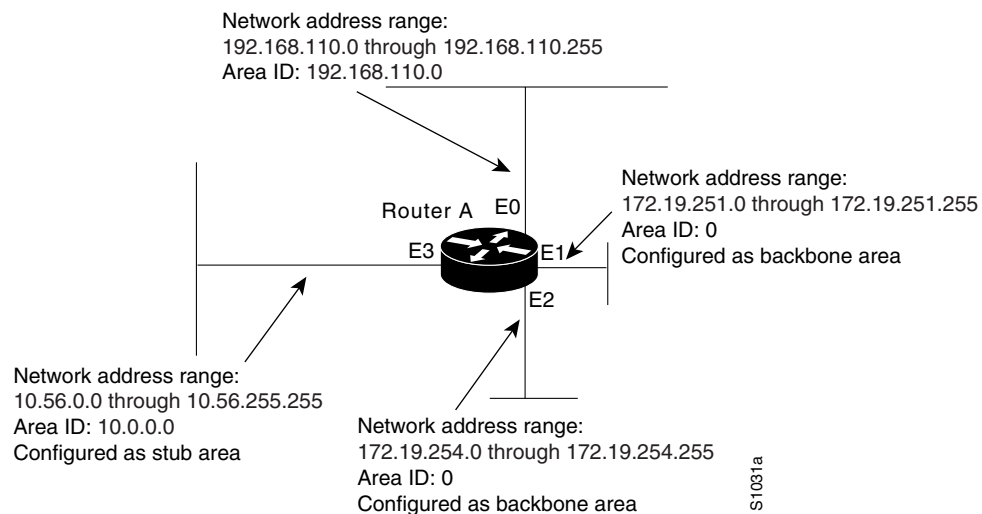
Complex OSPF Configuration for ABR Examples

The following example configuration accomplishes several tasks in setting up an ABR. These tasks can be split into two general categories:

- Basic OSPF configuration
- Route redistribution

The specific tasks outlined in this configuration are detailed briefly in the following descriptions. Figure 22 illustrates the network address ranges and area assignments for the interfaces.

Figure 22 Interface and Area Specifications for OSPF Example Configuration



The basic configuration tasks in this example are as follows:

- Configure address ranges for Ethernet 0 through Ethernet 3 interfaces.
- Enable OSPF on each interface.
- Set up an OSPF authentication password for each area and network.
- Assign link state metrics and other OSPF interface configuration options.
- Create a stub area with area id 36.0.0.0. (Note that the **authentication** and **stub** options of the **area** router configuration command are specified with separate **area** command entries, but can be merged into a single **area** command.)
- Specify the backbone area (Area 0).

Configuration tasks associated with redistribution are as follows:

- Redistribute IGRP and RIP into OSPF with various options set (including **metric-type**, **metric**, **tag**, and **subnet**).
- Redistribute IGRP and OSPF into RIP.

The following is an example OSPF configuration:

```
interface ethernet 0
 ip address 192.42.110.201 255.255.255.0
 ip ospf authentication-key abcdefgh
 ip ospf cost 10
!
interface ethernet 1
 ip address 131.119.251.201 255.255.255.0
 ip ospf authentication-key ijklmnop
 ip ospf cost 20
 ip ospf retransmit-interval 10
 ip ospf transmit-delay 2
 ip ospf priority 4
!
interface ethernet 2
 ip address 131.119.254.201 255.255.255.0
 ip ospf authentication-key abcdefgh
 ip ospf cost 10
!
interface ethernet 3
 ip address 36.56.0.201 255.255.0.0
 ip ospf authentication-key ijklmnop
 ip ospf cost 20
 ip ospf dead-interval 80
```

OSPF is on network 131.119.0.0:

```
router ospf 201
 network 36.0.0.0 0.255.255.255 area 36.0.0.0
 network 192.42.110.0 0.0.0.255 area 192.42.110.0
 network 131.119.0.0 0.0.255.255 area 0
 area 0 authentication
 area 36.0.0.0 stub
 area 36.0.0.0 authentication
 area 36.0.0.0 default-cost 20
 area 192.42.110.0 authentication
 area 36.0.0.0 range 36.0.0.0 255.0.0.0
 area 192.42.110.0 range 192.42.110.0 255.255.255.0
 area 0 range 131.119.251.0 255.255.255.0
 area 0 range 131.119.254.0 255.255.255.0
```

```

redistribute igrp 200 metric-type 2 metric 1 tag 200 subnets
redistribute rip metric-type 2 metric 1 tag 200

```

IGRP autonomous system 200 is on 131.119.0.0:

```

router igrp 200
 network 131.119.0.0
!
! RIP for 192.42.110
!
router rip
 network 192.42.110.0
 redistribute igrp 200 metric 1
 redistribute ospf 201 metric 1

```

Route Map Examples

The examples in this section illustrate the use of redistribution, with and without route maps. Examples from both the IP and CLNS routing protocols are given.

The following example redistributes all OSPF routes into IGRP:

```

router igrp 109
 redistribute ospf 110

```

The following example redistributes RIP routes with a hop count equal to 1 into OSPF. These routes will be redistributed into OSPF as external link state advertisements with a metric of 5, metric type of Type 1, and a tag equal to 1.

```

router ospf 109
 redistribute rip route-map rip-to-ospf
!
route-map rip-to-ospf permit
 match metric 1
 set metric 5
 set metric-type type1
 set tag 1

```

The following example redistributes OSPF learned routes with tag 7 as a RIP metric of 15:

```

router rip
 redistribute ospf 109 route-map 5
!
route-map 5 permit
 match tag 7
 set metric 15

```

The following example redistributes OSPF intra-area and interarea routes with next-hop routers on serial interface 0 into BGP with an INTER_AS metric of 5:

```

router bgp 109
 redistribute ospf 109 route-map 10
!
route-map 10 permit
 match route-type internal
 match interface serial 0
 set metric 5

```

The following example redistributes two types of routes into the integrated IS-IS routing table (supporting both IP and CLNS). The first are OSPF external IP routes with tag 5; these are inserted into Level 2 IS-IS LSPs with a metric of 5. The second are ISO-IGRP derived CLNS prefix routes that match CLNS access list 2000. These will be redistributed into IS-IS as Level 2 LSPs with a metric of 30.

```
router isis
 redistribute ospf 109 route-map 2
 redistribute iso-igrp nsfnet route-map 3
!
route-map 2 permit
 match route-type external
 match tag 5
 set metric 5
 set level level-2
!
route-map 3 permit
 match address 2000
 set metric 30
```

With the following configuration, OSPF external routes with tags 1, 2, 3, and 5 are redistributed into RIP with metrics of 1, 1, 5, and 5, respectively. The OSPF routes with a tag of 4 are not redistributed.

```
router rip
 redistribute ospf 109 route-map 1
!
route-map 1 permit
 match tag 1 2
 set metric 1
!
route-map 1 permit
 match tag 3
 set metric 5
!
route-map 1 deny
 match tag 4
!
route map 1 permit
 match tag 5
 set metric 5
```

The following configuration sets the condition that if there is an OSPF route to network 140.222.0.0, generate the default network 0.0.0.0 into RIP with a metric of 1:

```
router rip
 redistribute ospf 109 route-map default
!
route-map default permit
 match ip address 1
 set metric 1
!
access-list 1 permit 140.222.0.0 0.0.255.255
access-list 2 permit 0.0.0.0 0.0.0.0
```

In the following configuration, a RIP learned route for network 160.89.0.0 and an ISO-IGRP learned route with prefix 49.0001.0002 will be redistributed into an IS-IS Level 2 LSP with a metric of 5:

```
router isis
 redistribute rip route-map 1
 redistribute iso-igrp remote route-map 1
!
route-map 1 permit
 match ip address 1
 match clns address 2
```

```
set metric 5
set level level-2
!
access-list 1 permit 160.89.0.0 0.0.255.255
clsns filter-set 2 permit 49.0001.0002...
```

The following configuration example illustrates how a route map is referenced by the **default-information** router configuration command. This is called *conditional default origination*. OSPF will originate the default route (network 0.0.0.0) with a Type 2 metric of 5 if 140.222.0.0, with network 0.0.0.0 in the routing table. Extended access-lists cannot be used in a route map for *conditional default origination*.

Note Only routes external to the OSPF process can be used for tracking, such as non-OSPF routes or OSPF routes from a separate OSPF process.

```
route-map ospf-default permit
match ip address 1
set metric 5
set metric-type type-2
!
access-list 1 140.222.0.0 0.0.255.255
!
router ospf 109
default-information originate route-map ospf-default
```

