



Configuring OSPF

This chapter describes how to configure Open Shortest Path First (OSPF). For a complete description of the OSPF commands in this chapter, refer to the “OSPF Commands” chapter of the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols* publication. To locate documentation of other commands that appear in this chapter, use the command reference master index, or search online.

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.

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 that include OSPF, see the chapter “Configuring IP Routing Protocol-Independent Features” in this book.

To identify the hardware platform or software image information associated with a feature, use the Feature Navigator on Cisco.com to search for information about the feature or refer to the software release notes for a specific release. For more information, see the “Identifying Supported Platforms” section in the “Using Cisco IOS Software” chapter in this book.

The Cisco OSPF Implementation

The Cisco implementation conforms to the OSPF Version 2 specifications detailed in the Internet RFC 2328. The list that follows outlines key features supported in the Cisco 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, OSPF can import routes learned via Interior Gateway Routing Protocol (IGRP), Routing Information Protocol (RIP), and Intermediate System-to-Intermediate System (IS-IS). OSPF routes can also be exported into IGRP, RIP, and IS-IS. At the interdomain level, OSPF can import routes learned via Exterior Gateway Protocol (EGP) and Border Gateway Protocol (BGP). OSPF routes can be exported into BGP and EGP.
- Authentication—Plain text and Message Digest 5 (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.
- OSPF over demand circuit—RFC 1793.

OSPF Configuration Task List

OSPF typically requires coordination among many internal routers: Area Border Routers (ABRs), which are routers connected to multiple areas, and Autonomous System Boundary Routers (ASBRs). 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, perform the tasks described in the following sections. The tasks in the first section are required; the tasks in the remaining sections are optional, but might be required for your application.

- [Enabling OSPF](#) (Required)
- [Configuring OSPF Interface Parameters](#) (Optional)
- [Configuring OSPF over Different Physical Networks](#) (Optional)
- [Configuring OSPF Area Parameters](#) (Optional)
- [Configuring OSPF NSSA](#) (Optional)
- [Configuring Route Summarization Between OSPF Areas](#) (Optional)
- [Configuring Route Summarization When Redistributing Routes into OSPF](#) (Optional)
- [Creating Virtual Links](#) (Optional)
- [Generating a Default Route](#) (Optional)
- [Configuring Lookup of DNS Names](#) (Optional)
- [Forcing the Router ID Choice with a Loopback Interface](#) (Optional)
- [Controlling Default Metrics](#) (Optional)
- [Changing the OSPF Administrative Distances](#) (Optional)
- [Configuring OSPF on Simplex Ethernet Interfaces](#) (Optional)
- [Configuring Route Calculation Timers](#) (Optional)
- [Configuring OSPF over On-Demand Circuits](#) (Optional)
- [Logging Neighbors Going Up or Down](#) (Optional)
- [Changing the LSA Group Pacing](#) (Optional)
- [Blocking OSPF LSA Flooding](#) (Optional)
- [Reducing LSA Flooding](#) (Optional)
- [Ignoring MOSPF LSA Packets](#) (Optional)
- [Displaying OSPF Update Packet Pacing](#) (Optional)
- [Monitoring and Maintaining OSPF](#) (Optional)

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.

Enabling 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. To do so, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# router ospf <i>process-id</i>	Enables OSPF routing, which places you in router configuration mode.
Step 2	Router(config-router)# network <i>ip-address wildcard-mask area</i> <i>area-id</i>	Defines an interface on which OSPF runs and define the area ID for that interface.

Configuring 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** interface configuration commands. Therefore, be sure that if you do configure any of these parameters, the configurations for all routers on your network have compatible values.

To specify interface parameters for your network, use the following commands in interface configuration mode, as needed:

Command	Purpose
Router(config-if)# ip ospf cost <i>cost</i>	Explicitly specifies the cost of sending a packet on an OSPF interface.
Router(config-if)# ip ospf retransmit-interval <i>seconds</i>	Specifies the number of seconds between link-state advertisement (LSA) retransmissions for adjacencies belonging to an OSPF interface.
Router(config-if)# ip ospf transmit-delay <i>seconds</i>	Sets the estimated number of seconds required to send a link-state update packet on an OSPF interface.
Router(config-if)# ip ospf priority <i>number-value</i>	Sets priority to help determine the OSPF designated router for a network.
Router(config-if)# ip ospf hello-interval <i>seconds</i>	Specifies the length of time between the hello packets that the Cisco IOS software sends on an OSPF interface.
Router(config-if)# ip ospf dead-interval <i>seconds</i>	Sets the number of seconds that a device must wait before it declares a neighbor OSPF router down because it has not received a hello packet.
Router(config-if)# ip ospf authentication-key <i>key</i>	Assigns a password to be used by neighboring OSPF routers on a network segment that is using the OSPF simple password authentication.

Command	Purpose
Router(config-if)# ip ospf message-digest-key <i>key-id md5 key</i>	Enables OSPF MD5 authentication. The values for the <i>key-id</i> and <i>key</i> arguments must match values specified for other neighbors on a network segment.
Router(config-if)# ip ospf authentication [message-digest null]	Specifies the authentication type for an interface.

Configuring OSPF over Different Physical Networks

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

- Broadcast networks (Ethernet, Token Ring, and FDDI)
- Nonbroadcast multiaccess (NBMA) networks (Switched Multimegabit Data Service (SMDS), Frame Relay, and X.25)
- Point-to-point networks (High-Level Data Link Control [HDLC], PPP)

You can configure your network as either a broadcast or an NBMA 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. Refer to the **x25 map** and **frame-relay map** command descriptions in the *Cisco IOS Wide-Area Networking Command Reference* publication for more detail.

Configuring Your OSPF Network Type

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

Configuring NBMA, multiaccess networks as either broadcast or nonbroadcast assumes that there are virtual circuits (VCs) 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 VCs to both routers. Note that you need 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 NBMA 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 VC failure.

To configure your OSPF network type, use the following command in interface configuration mode:

Command	Purpose
Router(config-if)# ip ospf network { broadcast non-broadcast { point-to-multipoint [non-broadcast]}}	Configures the OSPF network type for a specified interface.

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

Configuring Point-to-Multipoint, Broadcast Networks

On point-to-multipoint, broadcast networks, there is no need to specify neighbors. However, you can specify neighbors with the **neighbor** router configuration command, in which case you should specify a cost to that neighbor.

Before the **point-to-multipoint** keyword was added to the **ip ospf network** interface configuration command, some OSPF point-to-multipoint protocol traffic was treated as multicast traffic. Therefore, the **neighbor** router configuration command was not needed for point-to-multipoint interfaces because multicast took care of the traffic. Hello, update, and acknowledgment messages were sent using multicast. In particular, multicast hello messages discovered all neighbors dynamically.

On any point-to-multipoint interface (broadcast or not), the Cisco IOS software assumed that the cost to each neighbor was equal. The cost was configured with the **ip ospf cost** interface configuration command. In reality, the bandwidth to each neighbor is different, so the cost should differ. With this feature, you can configure a separate cost to each neighbor. This feature applies to point-to-multipoint interfaces only.

To treat an interface as point-to-multipoint broadcast and assign a cost to each neighbor, use the following commands beginning in interface configuration mode:

	Command	Purpose
Step 1	Router(config-if)# ip ospf network point-to-multipoint	Configures an interface as point-to-multipoint for broadcast media.
Step 2	Router(config-if)# exit	Enters global configuration mode.
Step 3	Router(config)# router ospf process-id	Configures an OSPF routing process and enters router configuration mode.
Step 4	Router(config-router)# neighbor ip-address cost number	Specifies a neighbor and assigns a cost to the neighbor.

Repeat Step 4 for each neighbor if you want to specify a cost. Otherwise, neighbors will assume the cost of the interface, based on the **ip ospf cost** interface configuration command.

Configuring OSPF for Nonbroadcast Networks

Because many routers might be attached to an OSPF network, a *designated router* is selected for the network. Special configuration parameters are needed 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 with a nonzero router priority value).

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

Command	Purpose
Router(config-router)# neighbor <i>ip-address</i> [<i>priority number</i>] [<i>poll-interval seconds</i>]	Configures a router interconnecting to nonbroadcast networks.

You can specify the following neighbor parameters, as required:

- Priority for a neighboring router
- Nonbroadcast poll interval

On point-to-multipoint, nonbroadcast networks, you now use the **neighbor** router configuration command to identify neighbors. Assigning a cost to a neighbor is optional.

Prior to Cisco IOS Release 12.0, some customers were using point-to-multipoint on nonbroadcast media (such as classic IP over ATM), so their routers could not dynamically discover their neighbors. This feature allows the **neighbor** router configuration command to be used on point-to-multipoint interfaces.

On any point-to-multipoint interface (broadcast or not), the Cisco IOS software assumed the cost to each neighbor was equal. The cost was configured with the **ip ospf cost** interface configuration command. In reality, the bandwidth to each neighbor is different, so the cost should differ. With this feature, you can configure a separate cost to each neighbor. This feature applies to point-to-multipoint interfaces only.

To treat the interface as point-to-multipoint when the media does not support broadcast, use the following commands beginning in interface configuration mode:

	Command	Purpose
Step 1	Router(config-if)# ip ospf network point-to-multipoint non-broadcast	Configures an interface as point-to-multipoint for nonbroadcast media.
Step 2	Router(config-if)# exit	Enters global configuration mode.
Step 3	Router(config)# router ospf process-id	Configures an OSPF routing process and enters router configuration mode.
Step 4	Router(config-router)# neighbor ip-address [cost number]	Specifies a neighbor and assigns a cost to the neighbor.

Repeat Step 4 for each neighbor if you want to specify a cost. Otherwise, neighbors will assume the cost of the interface, based on the **ip ospf cost** interface configuration command.

Configuring OSPF Area Parameters

Our OSPF software allows you to configure several area parameters. These area parameters, shown in the following task 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 ABR, into the stub area for destinations outside the autonomous system. To take advantage of the OSPF stub area support, *default routing* must be used in the stub area. To further reduce the number of LSAs sent into a stub area, you can configure the **no-summary** keyword of the **area stub** router configuration command on the ABR to prevent it from sending summary link advertisement (LSAs type 3) into the stub area.

To specify an area parameter for your network, use the following commands in router configuration mode as needed:

Command	Purpose
Router(config-router)# area area-id authentication	Enables authentication for an OSPF area.
Router(config-router)# area area-id authentication message-digest	Enables MD5 authentication for an OSPF area.
Router(config-router)# area area-id stub [no-summary]	Defines an area to be a stub area.
Router(config-router)# area area-id default-cost cost	Assigns a specific cost to the default summary route used for the stub area.

Configuring OSPF NSSA

The OSPF implementation of NSSA is similar to OSPF stub area. NSSA does not flood type 5 external LSAs from the core into the area, but it can import autonomous system external routes in a limited fashion within the area.

NSSA allows importing of type 7 autonomous system external routes within 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 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 could not be redistributed into 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 the remote router as an NSSA.

To specify area parameters as needed to configure OSPF NSSA, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# area area-id nssa [no-redistribution] [default-information-originate]	Defines an area to be NSSA.

To control summarization and filtering of type 7 LSAs into type 5 LSAs, use the following command in router configuration mode on the ABR:

Command	Purpose
Router(config-router)# summary <i>address prefix mask</i> [not advertise] [tag <i>tag</i>]	Controls the summarization and filtering during the translation.

Implementation Considerations

Evaluate the following considerations before you implement 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 or 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.

Configuring 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, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# area <i>area-id</i> range <i>ip-address mask</i> [advertise not-advertise]	Specifies an address range for which a single route will be advertised.

Configuring Route Summarization When Redistributing Routes into OSPF

When routes from other protocols are redistributed into OSPF (as described in the chapter “Configuring IP Routing Protocol-Independent Features”), each route is advertised individually in an external 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, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# summary-address <i>ip-address mask prefix mask</i> [not-advertise][tag <i>tag</i>]	Specifies an address and mask that covers redistributed routes, so only one summary route is advertised. Use the optional not-advertise keyword to filter out a set of routes.

Creating 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 endpoints of a virtual link are ABRs. 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, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# area <i>area-id</i> virtual-link <i>router-id</i> [authentication [message-digest null]] [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>key-id md5 key</i>]]	Establishes a virtual link.

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.

Generating a Default Route

You can force an ASBR 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 ASBR. However, an ASBR does not, by default, generate a *default route* into the OSPF routing domain.

To force the ASBR to generate a default route, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# default-information originate [always] [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [route-map <i>map-name</i>]	Forces the autonomous system boundary router to generate a default route into the OSPF routing domain.

For a discussion of redistribution of routes, see the “Configuring IP Routing Protocol-Independent Features” chapter.

Configuring Lookup of DNS Names

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

To configure DNS name lookup, use the following command in global configuration mode:

Command	Purpose
Router(config)# ip ospf name-lookup	Configures DNS name lookup.

Forcing 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. Because 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, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# interface loopback 0	Creates a loopback interface, which places the router in interface configuration mode.
Step 2	Router(config-if)# ip address ip-address mask	Assigns an IP address to this interface.

Controlling Default Metrics

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

The OSPF metric is calculated as the *ref-bw* value divided by the *bandwidth* value, with the *ref-bw* value equal to 10⁸ by default, and the *bandwidth* value determined by the **bandwidth** interface configuration 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, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# auto-cost reference-bandwidth ref-bw	Differentiates high bandwidth links.

Changing the OSPF Administrative Distances

An administrative distance is a rating of the trustworthiness of a routing information source, such as an individual router or a group of routers. Numerically, an administrative distance is an integer from 0 to 255. In general, the higher the value, the lower the trust rating. An administrative distance of 255 means the routing information source cannot be trusted at all and should be ignored.

OSPF uses three different administrative distances: intra-area, interarea, and external. Routes within an area are intra-area; routes to another area are interarea; and routes from another routing domain learned via redistribution are external. The default distance for each type of route is 110.

To change any of the OSPF distance values, use the following command in router configuration mode:

Command	Purpose
<code>Router(config-router)# distance ospf {[intra-area dist1] [inter-area dist2] [external dist3]}</code>	Changes the OSPF distance values.

For an example of changing administrative distance, see the section “[Changing OSPF Administrative Distance Example](#)” at the end of this chapter.

Configuring 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 sending interface to be a passive interface. This configuration prevents OSPF from sending hello packets for the sending 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, use the following command in router configuration mode:

Command	Purpose
<code>Router(config-router)# passive-interface interface-type interface-number</code>	Suppresses the sending of hello packets through the specified interface.

Configuring 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 so, use the following command in router configuration mode:

Command	Purpose
<code>Router(config-router)# timers spf spf-delay spf-holdtime</code>	Configures route calculation timers.

Configuring 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 switched virtual circuits (SVCs), and dialup lines. This feature supports RFC 1793, *Extending OSPF to Support Demand Circuits*.

Prior to this feature, OSPF periodic hello and LSA 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 data link 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 send.

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 sent in order to maintain network topology integrity. Periodic refreshes that do not include changes, however, are not sent across the link.

To configure OSPF for on-demand circuits, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# router ospf <i>process-id</i>	Enables OSPF operation.
Step 2	Router(config)# interface <i>interface-type</i> <i>interface-number</i>	Enters interface configuration mode.
Step 3	Router(config-if)# ip ospf demand-circuit	Configures OSPF on an on-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.

For an example of OSPF over an on-demand circuit, see the section “[OSPF over On-Demand Routing Example](#)” at the end of this chapter.

Implementation Considerations

Evaluate the following considerations before implementing this feature:

- Because LSAs that include topology changes are flooded over an on-demand circuit, we recommend that you 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 because type 5 external LSAs are flooded throughout all areas.
- Do not implement this feature on a broadcast-based network topology because the overhead protocols (such as hello and LSA packets) cannot be successfully suppressed, which means the link will remain up.

Logging Neighbors Going Up or Down

By default, the system sends a syslog message when an OSPF neighbor goes up or down. If you turned off this feature and want to restore it, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# log-adj-changes [detail]	Sends syslog message when an OSPF neighbor goes up or down.

Configure this command if you want to know about OSPF neighbors going up or down without turning on the **debug ip ospf adjacency EXEC** command. The **log-adj-changes** router configuration command provides a higher level view of the peer relationship with less output. Configure **log-adj-changes detail** if you want to see messages for each state change.

Changing the LSA Group Pacing

The OSPF LSA group pacing feature allows the router to group OSPF LSAs and pace the refreshing, checksumming, and aging functions. The group pacing results in more efficient use of the router.

The router groups OSPF LSAs and paces the refreshing, checksumming, and aging functions so that sudden increases in CPU usage and network resources are avoided. This feature is most beneficial to large OSPF networks.

OSPF LSA group pacing is enabled by default. For typical customers, the default group pacing interval for refreshing, checksumming, and aging is appropriate and you need not configure this feature.

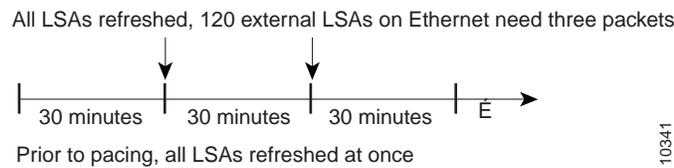
Original LSA Behavior

Each OSPF LSA has an age, which indicates whether the LSA is still valid. Once the LSA reaches the maximum age (1 hour), it is discarded. During the aging process, the originating router sends a refresh packet every 30 minutes to refresh the LSA. Refresh packets are sent to keep the LSA from expiring, whether there has been a change in the network topology or not. Checksumming is performed on all LSAs every 10 minutes. The router keeps track of LSAs it generates and LSAs it receives from other routers. The router refreshes LSAs it generated; it ages the LSAs it received from other routers.

Prior to the LSA group pacing feature, the Cisco IOS software would perform refreshing on a single timer, and checksumming and aging on another timer. In the case of refreshing, for example, the software would scan the whole database every 30 minutes, refreshing every LSA the router generated, no matter how old it was. [Figure 37](#) illustrates all the LSAs being refreshed at once. This process wasted CPU resources because only a small portion of the database needed to be refreshed. A large OSPF database (several thousand LSAs) could have thousands of LSAs with different ages. Refreshing on a single timer

resulted in the age of all LSAs becoming synchronized, which resulted in much CPU processing at once. Furthermore, a large number of LSAs could cause a sudden increase of network traffic, consuming a large amount of network resources in a short period of time.

Figure 37 OSPF LSAs on a Single Timer Without Group Pacing



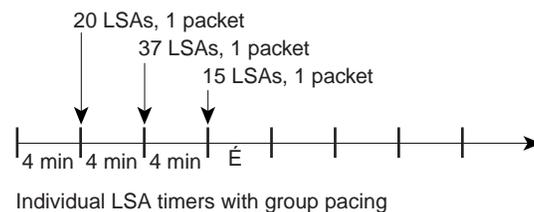
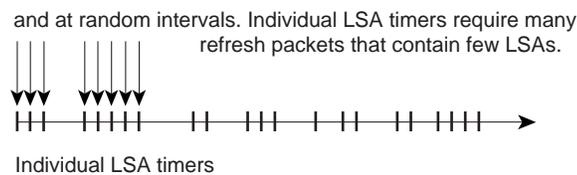
LSA Group Pacing With Multiple Timers

This problem is solved by configuring each LSA to have its own timer. To again use the example of refreshing, each LSA gets refreshed when it is 30 minutes old, independent of other LSAs. So the CPU is used only when necessary. However, LSAs being refreshed at frequent, random intervals would require many packets for the few refreshed LSAs the router must send out, which would be inefficient use of bandwidth.

Therefore, the router delays the LSA refresh function for an interval of time instead of performing it when the individual timers are reached. The accumulated LSAs constitute a group, which is then refreshed and sent out in one packet or more. Thus, the refresh packets are paced, as are the checksumming and aging. The pacing interval is configurable; it defaults to 4 minutes, which is randomized to further avoid synchronization.

Figure 38 illustrates the case of refresh packets. The first timeline illustrates individual LSA timers; the second timeline illustrates individual LSA timers with group pacing.

Figure 38 OSPF LSAs on Individual Timers with Group Pacing



The group pacing interval is inversely proportional to the number of LSAs the router is refreshing, checksumming, and aging. For example, if you have approximately 10,000 LSAs, decreasing the pacing interval would benefit you. If you have a very small database (40 to 100 LSAs), increasing the pacing interval to 10 to 20 minutes might benefit you slightly.

The default value of pacing between LSA groups is 240 seconds (4 minutes). The range is from 10 seconds to 1800 seconds (30 minutes). To change the LSA group pacing interval, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# timers lsa-group-pacing <i>seconds</i>	Changes the group pacing of LSAs.

For an example, see the section “[LSA Group Pacing Example](#)” at the end of this chapter.

Blocking OSPF LSA Flooding

By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. Some redundancy is desirable, because it ensures robust flooding. However, too much redundancy can waste bandwidth and might destabilize the network due to excessive link and CPU usage in certain topologies. An example would be a fully meshed topology.

You can block OSPF flooding of LSAs two ways, depending on the type of networks:

- On broadcast, nonbroadcast, and point-to-point networks, you can block flooding over specified OSPF interfaces.
- On point-to-multipoint networks, you can block flooding to a specified neighbor.

On broadcast, nonbroadcast, and point-to-point networks, to prevent flooding of OSPF LSAs, use the following command in interface configuration mode:

Command	Purpose
Router(config-if)# ospf database-filter all out	Blocks the flooding of OSPF LSA packets to the interface.

On point-to-multipoint networks, to prevent flooding of OSPF LSAs, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# neighbor ip-address database-filter all out	Blocks the flooding of OSPF LSA packets to the specified neighbor.

For an example of blocking LSA flooding, see the section “[Block LSA Flooding Example](#)” at the end of this chapter.

Reducing LSA Flooding

The explosive growth of the Internet has placed the focus on the scalability of IGPs such as OSPF. By design, OSPF requires LSAs to be refreshed as they expire after 3600 seconds. Some implementations have tried to improve the flooding by reducing the frequency to refresh from 30 minutes to about 50 minutes. This solution reduces the amount of refresh traffic but requires at least one refresh before the

LSA expires. The OSPF flooding reduction solution works by reducing unnecessary refreshing and flooding of already known and unchanged information. To achieve this reduction, the LSAs are now flooded with the higher bit set. The LSAs are now set as “do not age.”

To reduce unnecessary refreshing and flooding of LSAs on your network, use the following command in interface configuration mode:

Command	Purpose
Router(config-if)# ip ospf flood-reduction	Suppresses the unnecessary flooding of LSAs in stable topologies.

Ignoring MOSPF LSA Packets

Cisco routers do not support LSA type 6 Multicast OSPF (MOSPF), and they generate syslog messages if they receive such packets. If the router is receiving many MOSPF packets, you might want to configure the router to ignore the packets and thus prevent a large number of syslog messages. To do so, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# ignore lsa mospf	Prevents the router from generating syslog messages when it receives MOSPF LSA packets.

For an example of suppressing MOSPF LSA packets, see the section “[Ignore MOSPF LSA Packets Example](#)” at the end of this chapter.

Displaying OSPF Update Packet Pacing

The former OSPF implementation for sending update packets needed to be more efficient. Some update packets were getting lost in cases where the link was slow, a neighbor could not receive the updates quickly enough, or the router was out of buffer space. For example, packets might be dropped if either of the following topologies existed:

- A fast router was connected to a slower router over a point-to-point link.
- During flooding, several neighbors sent updates to a single router at the same time.

OSPF update packets are now automatically paced so they are not sent less than 33 milliseconds apart. Pacing is also added between resends to increase efficiency and minimize lost retransmissions. Also, you can display the LSAs waiting to be sent out an interface. The benefit of the pacing is that OSPF update and retransmission packets are sent more efficiently.

There are no configuration tasks for this feature; it occurs automatically.

To observe OSPF packet pacing by displaying a list of LSAs waiting to be flooded over a specified interface, use the following command in EXEC mode:

Command	Purpose
Router# show ip ospf flood-list <i>interface-type interface-number</i>	Displays a list of LSAs waiting to be flooded over an interface.

Monitoring and Maintaining 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 that your device packets are taking through the network.

To display various routing statistics, use the following commands in EXEC mode, as needed:

Command	Purpose
Router# show ip ospf [<i>process-id</i>]	Displays general information about OSPF routing processes.
Router# show ip ospf border-routers	Displays the internal OSPF routing table entries to the ABR and ASBR.
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database	Displays lists of information related to the OSPF database.
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>database-summary</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>router</i>] [<i>self-originate</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>router</i>] [<i>adv-router</i> [<i>ip-address</i>]]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>router</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>network</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>summary</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>asbr-summary</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>external</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>nssa-external</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>opaque-link</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>opaque-area</i>] [<i>link-state-id</i>]	
Router# show ip ospf [<i>process-id</i> [<i>area-id</i>]] database [<i>opaque-as</i>] [<i>link-state-id</i>]	
Router# show ip ospf flood-list interface <i>interface-type</i>	Displays a list of LSAs waiting to be flooded over an interface (to observe OSPF packet pacing).
Router# show ip ospf interface [<i>interface-type</i>] <i>interface-number</i>	Displays OSPF-related interface information.

Command	Purpose
Router# show ip ospf neighbor [<i>interface-name</i>] [<i>neighbor-id</i>] detail	Displays OSPF neighbor information on a per-interface basis.
Router# show ip ospf request-list [<i>neighbor</i>] [<i>interface</i>] [<i>interface-neighbor</i>]	Displays a list of all LSAs requested by a router.
Router# show ip ospf retransmission-list [<i>neighbor</i>] [<i>interface</i>] [<i>interface-neighbor</i>]	Displays a list of all LSAs waiting to be resent.
Router# show ip ospf [<i>process-id</i>] summary-address	Displays a list of all summary address redistribution information configured under an OSPF process.
Router# show ip ospf virtual-links	Displays OSPF-related virtual links information.

To restart an OSPF process, use the following command in EXEC mode:

Command	Purpose
Router# clear ip ospf [<i>pid</i>] { process redistribution counters [<i>neighbor</i> [<i>neighbor-interface</i>] [<i>neighbor-id</i>]]}	Clears redistribution based on the OSPF routing process ID. If the <i>pid</i> option is not specified, all OSPF processes are cleared.

OSPF Configuration Examples

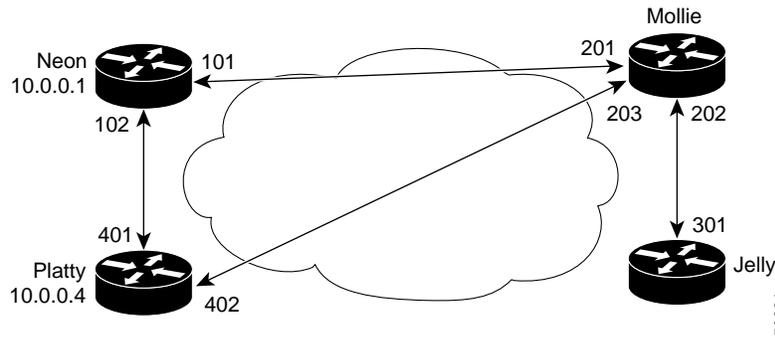
The following sections provide OSPF configuration examples:

- [OSPF Point-to-Multipoint Example](#)
- [OSPF Point-to-Multipoint, Broadcast Example](#)
- [OSPF Point-to-Multipoint, Nonbroadcast Example](#)
- [Variable-Length Subnet Masks Example](#)
- [OSPF Routing and Route Redistribution Examples](#)
- [Route Map Examples](#)
- [Changing OSPF Administrative Distance Example](#)
- [OSPF over On-Demand Routing Example](#)
- [LSA Group Pacing Example](#)
- [Block LSA Flooding Example](#)
- [Ignore MOSPF LSA Packets Example](#)

OSPF Point-to-Multipoint Example

In [Figure 39](#), the router named Mollie uses data-link connection identifier (DLCI) 201 to communicate with the router named Neon, DLCI 202 to the router named Jelly, and DLCI 203 to the router named 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). Configuration examples follow the figure.

Figure 39 OSPF Point-to-Multipoint Example



Mollie 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 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 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 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

```

OSPF Point-to-Multipoint, Broadcast Example

The following example illustrates a point-to-multipoint network with broadcast:

```

interface Serial0
 ip address 10.0.1.1 255.255.255.0
 encapsulation frame-relay
 ip ospf cost 100
 ip ospf network point-to-multipoint
 frame-relay map ip 10.0.1.3 202 broadcast
 frame-relay map ip 10.0.1.4 203 broadcast
 frame-relay map ip 10.0.1.5 204 broadcast
 frame-relay local-dlci 200
!
router ospf 1
 network 10.0.1.0 0.0.0.255 area 0
 neighbor 10.0.1.5 cost 5
 neighbor 10.0.1.4 cost 10

```

The following example shows the configuration of the neighbor at 10.0.1.3:

```

interface serial 0
 ip address 10.0.1.3 255.255.255.0
 ip ospf network point-to-multipoint
 encapsulation frame-relay
 frame-relay local-dlci 301
 frame-relay map ip 10.0.1.1 300 broadcast
 no shut
!
router ospf 1
 network 10.0.1.0 0.0.0.255 area 0

```

The output shown for neighbors in the first configuration is as follows:

```
Router# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.1.1.1	1	FULL/ -	00:01:50	10.0.1.5	Serial0
3.1.1.1	1	FULL/ -	00:01:47	10.0.1.4	Serial0
2.1.1.1	1	FULL/ -	00:01:45	10.0.1.3	Serial0

The route information in the first configuration is as follows:

```
Router# show ip route
```

```

Codes: C - connected, S - static, I - IGRP, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, E - EGP
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, * - candidate default
       U - per-user static route, o - ODR
Gateway of last resort is not set
C    1.0.0.0/8 is directly connected, Loopback0

```

```

    10.0.0.0/8 is variably subnetted, 4 subnets, 2 masks
O   10.0.1.3/32 [110/100] via 10.0.1.3, 00:39:08, Serial0
C   10.0.1.0/24 is directly connected, Serial0
O   10.0.1.5/32 [110/5] via 10.0.1.5, 00:39:08, Serial0
O   10.0.1.4/32 [110/10] via 10.0.1.4, 00:39:08, Serial0

```

OSPF Point-to-Multipoint, Nonbroadcast Example

The following example illustrates a point-to-multipoint network with nonbroadcast:

```

interface Serial0
ip address 10.0.1.1 255.255.255.0
ip ospf network point-to-multipoint non-broadcast
encapsulation frame-relay
no keepalive
frame-relay local-dlci 200
frame-relay map ip 10.0.1.3 202
frame-relay map ip 10.0.1.4 203
frame-relay map ip 10.0.1.5 204
no shut
!
router ospf 1
network 10.0.1.0 0.0.0.255 area 0
neighbor 10.0.1.3 cost 5
neighbor 10.0.1.4 cost 10
neighbor 10.0.1.5 cost 15

```

The following example is the configuration for the router on the other side:

```

interface Serial9/2
ip address 10.0.1.3 255.255.255.0
encapsulation frame-relay
ip ospf network point-to-multipoint non-broadcast
no ip mroute-cache
no keepalive
no fair-queue
frame-relay local-dlci 301
frame-relay map ip 10.0.1.1 300
no shut
!
router ospf 1
network 10.0.1.0 0.0.0.255 area 0

```

The output shown for neighbors in the first configuration is as follows:

```

Router# show ip ospf neighbor

```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.1.1.1	1	FULL/ -	00:01:52	10.0.1.5	Serial0
3.1.1.1	1	FULL/ -	00:01:52	10.0.1.4	Serial0
2.1.1.1	1	FULL/ -	00:01:52	10.0.1.3	Serial0

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 ethernet

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, ABRs, and ASBRs. At a minimum, OSPF-based routers can be configured with all default parameter values, with no authentication, and with interfaces assigned to areas.

Three types of 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 Examples

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

```
interface ethernet 0
 ip address 10.93.1.1 255.255.255.0
 ip ospf cost 1
 !
interface ethernet 1
 ip address 10.94.1.1 255.255.255.0
 !
router ospf 9000
 network 10.93.0.0 0.0.255.255 area 0.0.0.0
 redistribute rip metric 1 subnets
 !
router rip
 network 10.94.0.0
 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, and 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 *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols* publication 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 Ethernet interface 0. Ethernet interface 0 is attached to area 10.9.50.0 only.

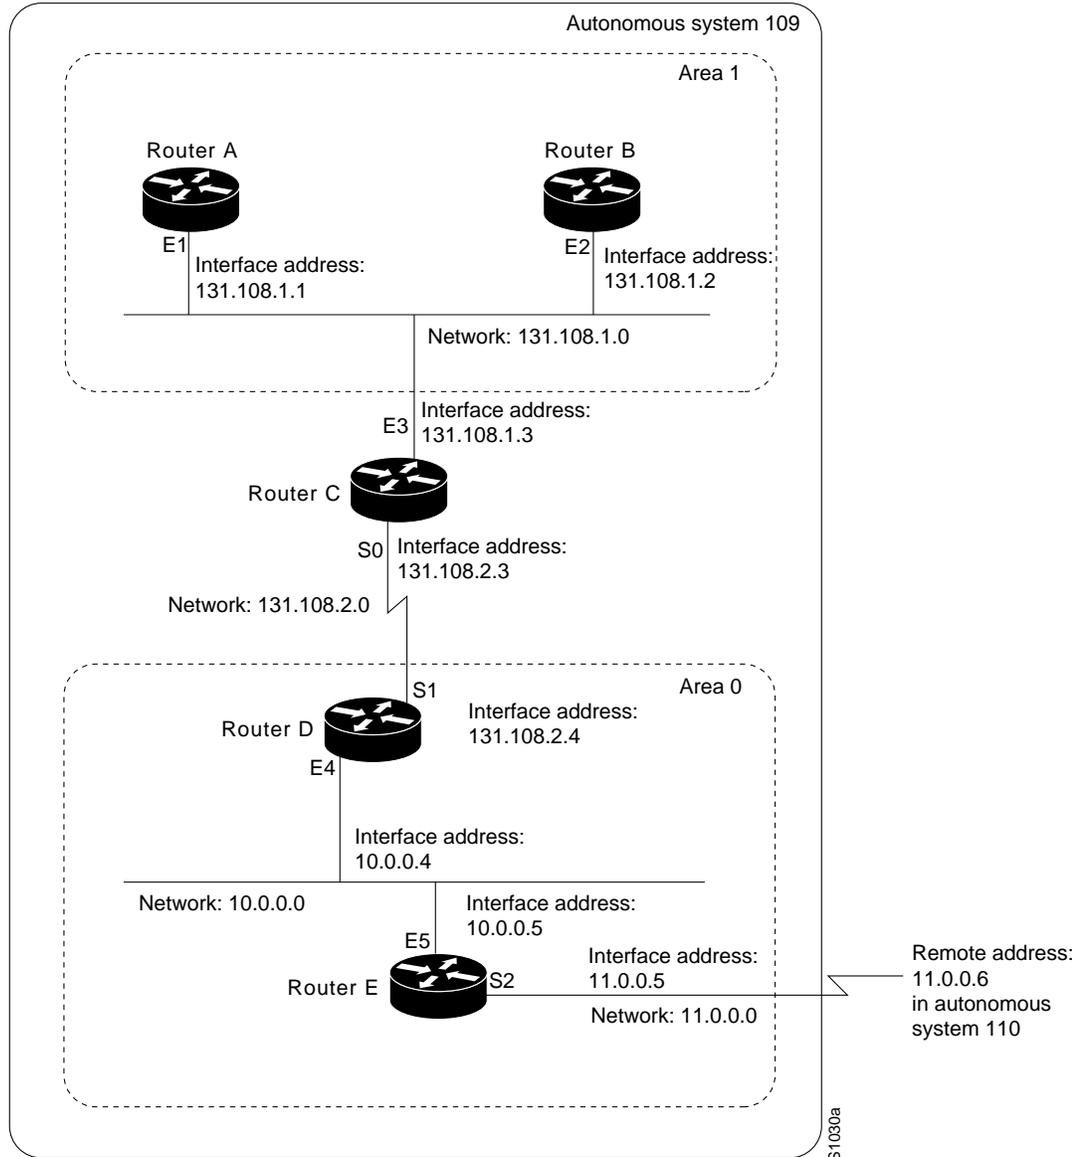
The second **network area** command is evaluated next. For area 2, the same process is then applied to all interfaces (except Ethernet interface 0). Assume that a match is determined for interface Ethernet 1. OSPF is then enabled for that interface and Ethernet interface 1 is attached to area 2.

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 40](#) provides a general network map that illustrates this example configuration.

Figure 40 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 ABR. 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 ASBR. 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 ABR (Router C) injects summary 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. Example configurations follow.

Following is the sample configuration for the general network map shown in [Figure 40](#).

Router A Configuration—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 Configuration—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 Configuration—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 Configuration—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 Configuration—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

```

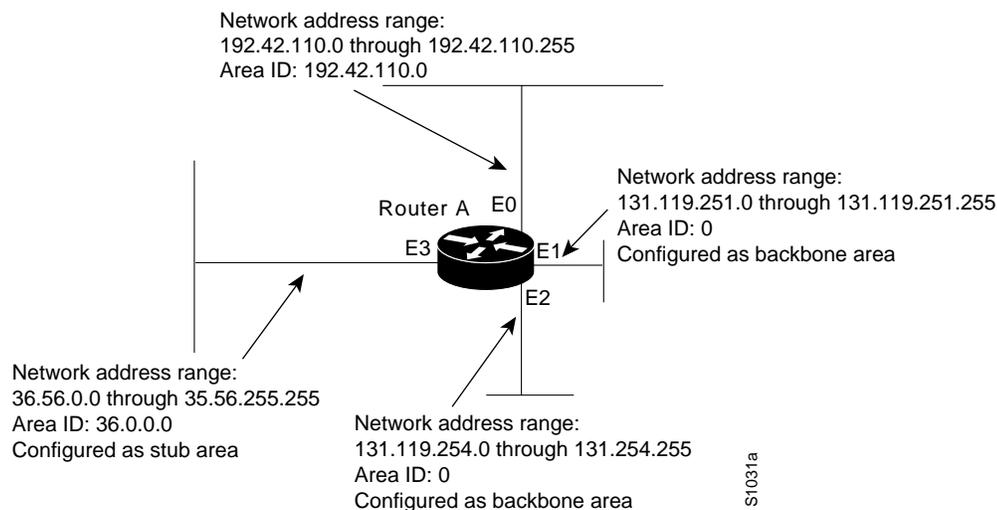
Complex OSPF Configuration for ABR Examples

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 41](#) illustrates the network address ranges and area assignments for the interfaces.

Figure 41 Interface and Area Specifications for OSPF Example Configuration



The basic configuration tasks in this example are as follows:

- Configure address ranges for Ethernet interface 0 through Ethernet interface 3.
- 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

```

In the following configuration 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

```

In the following configuration 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 Connectionless Network Service (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 LSAs with a metric of 5, a 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 type is OSPF external IP routes with tag 5; these routes are inserted into Level 2 IS-IS LSPs with a metric of 5. The second type is ISO-IGRP derived CLNS prefix routes that match CLNS access list 2000; these routes 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
```

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
clns filter-set 2 permit 49.0001.0002...
```

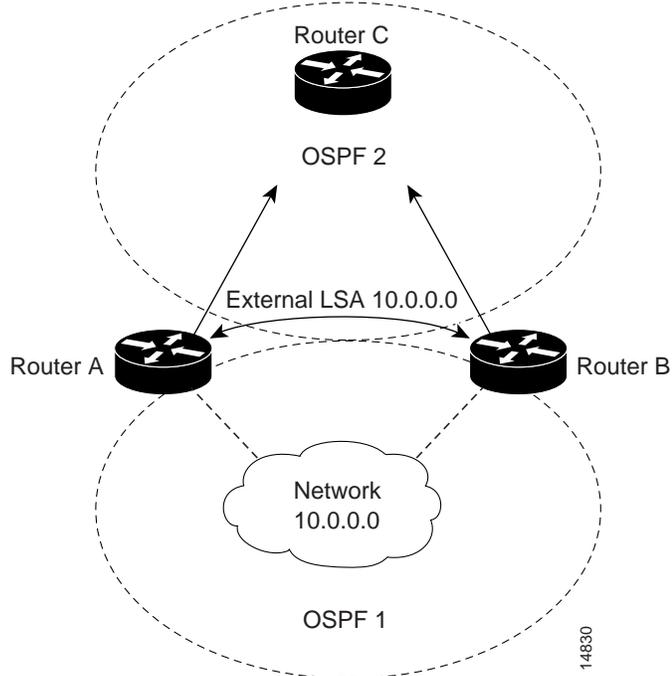
The following configuration example illustrates how a route map is referenced by the **default-information** router configuration command. This type of reference is called *conditional default origination*. OSPF will originate the default route (network 0.0.0.0) with a type 2 metric of 5 if 140.222.0.0 is in the routing table.

```
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
```

Changing OSPF Administrative Distance Example

The following configuration changes the external distance to 200, making it less trustworthy. [Figure 42](#) illustrates the example.

Figure 42 OSPF Administrative Distance

**Router A Configuration**

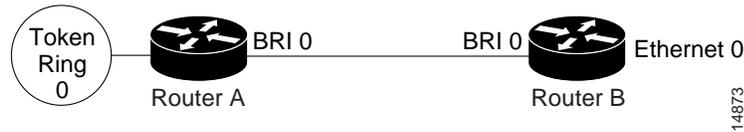
```
router ospf 1
 redistribute ospf 2 subnet
 distance ospf external 200
!
router ospf 2
 redistribute ospf 1 subnet
 distance ospf external 200
```

Router B Configuration

```
router ospf 1
 redistribute ospf 2 subnet
 distance ospf external 200
!
router ospf 2
 redistribute ospf 1 subnet
 distance ospf external 200
```

OSPF over On-Demand Routing Example

The following configuration allows OSPF over an on-demand circuit, as shown in [Figure 43](#). Note that the on-demand circuit is defined on one side only (BRI 0 on Router A). It is not required to be configured on both sides.

Figure 43 OSPF over On-Demand Circuit**Router A Configuration**

```

username RouterB password 7 060C1A2F47
isdn switch-type basic-5ess
ip routing
!
interface TokenRing0
 ip address 140.10.20.7 255.255.255.0
 no shut
!
interface BRI0
 no cdp enable
 description connected PBX 1485
 ip address 140.10.10.7 255.255.255.0
 encapsulation ppp
 ip ospf demand-circuit
 dialer map ip 140.10.10.6 name RouterB broadcast 61484
 dialer-group 1
 ppp authentication chap
 no shut
!
router ospf 100
 network 140.10.10.0 0.0.0.255 area 0
 network 140.10.20.0 0.0.0.255 area 0
!
dialer-list 1 protocol ip permit

```

Router B Configuration

```

username RouterA password 7 04511E0804
isdn switch-type basic-5ess
ip routing
!
interface Ethernet0
 ip address 140.10.60.6 255.255.255.0
 no shut
!
interface BRI0
 no cdp enable
 description connected PBX 1484
 ip address 140.10.10.6 255.255.255.0
 encapsulation ppp
 dialer map ip 140.10.10.7 name RouterA broadcast 61485
 dialer-group 1
 ppp authentication chap
 no shut
!
router ospf 100
 network 140.10.10.0 0.0.0.255 area 0
 network 140.10.60.0 0.0.0.255 area 0
!
dialer-list 1 protocol ip permit

```

LSA Group Pacing Example

The following example changes the OSPF pacing between LSA groups to 60 seconds:

```
router ospf
 timers lsa-group-pacing 60
```

Block LSA Flooding Example

The following example prevents flooding of OSPF LSAs to broadcast, nonbroadcast, or point-to-point networks reachable through Ethernet interface 0:

```
interface ethernet 0
 ospf database-filter all out
```

The following example prevents flooding of OSPF LSAs to point-to-multipoint networks to the neighbor at IP address 1.2.3.4:

```
router ospf 109
 neighbor 1.2.3.4 database-filter all out
```

Ignore MOSPF LSA Packets Example

The following example configures the router to suppress the sending of syslog messages when it receives MOSPF packets:

```
router ospf 109
 ignore lsa mospf
```