



Overview

This document describes how to configure unicast routing on the Cisco Industrial Ethernet 2000U Series Switches, hereafter referred to as IE 2000U or switch. To use unicast routing, the switch must be running the IP services image.

This chapter provides an overview of the following unicast routing features:

- [IPv4 Unicast Routing, page 1-1](#)
- [IPv6 Unicast Routing, page 1-1](#)
- [Enhanced Object Tracking, page 1-2](#)

IPv4 Unicast Routing

Routers and Layer 3 switches can route packets in the following ways:

- By using default routing—sending traffic with a destination unknown to the router to a default outlet or destination.
- By using preprogrammed static routes for the traffic

Static unicast routing forwards packets from predetermined ports through a single path into and out of a network. Static routing does not automatically respond to changes in the network and therefore, might result in unreachable destinations.

- By dynamically calculating routes by using a routing protocol

Dynamic routing protocols are used by routers to dynamically calculate the best route for forwarding traffic. Routing protocols supported by the switch are Routing Information Protocol (RIP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) protocol, Enhanced IGRP (EIGRP), System-to-Intermediate System (IS-IS), and Bidirectional Forwarding Detection (BFD).

Related Topics

[Chapter 2, “Configuring IP Unicast Routing”](#)

IPv6 Unicast Routing

IPv4 users can move to IPv6 and receive services such as end-to-end security, quality of service (QoS), and globally unique addresses. The IPv6 address space reduces the need for private addresses and Network Address Translation (NAT) processing by border routers at network edges.

IPv6 unicast routing support on the switch includes expanded address capability, header format simplification, improved support of extensions and options, and hardware parsing of the extension header. The switch supports hop-by-hop extension header packets, which are routed or bridged in software.

The switch provides IPv6 routing capability over 802.1Q trunk ports for static routes, Routing Information Protocol (RIP) for IPv6, and Open Shortest Path First (OSPF) Version 3 Protocol. It supports up to 16 equal-cost routes and can simultaneously forward IPv4 and IPv6 frames at line rate.

Related Topics

[Chapter 3, “Configuring IPv6 Unicast Routing”](#)

Enhanced Object Tracking

Enhanced object tracking on the switch provides a more complete alternative to the Hot Standby Routing Protocol (HSRP) tracking mechanism, which allows you to track the line-protocol state of an interface. If the line protocol state of an interface goes down, the HSRP priority of the interface is reduced and another HSRP device with a higher priority becomes active. The enhanced object tracking feature separates the tracking mechanism from HSRP and creates a separate, standalone tracking process that can be used by processes other than HSRP. This allows tracking other objects in addition to the interface line-protocol state.

A client process, such as HSRP or Gateway Local Balancing Protocol (GLBP), can register an interest in tracking objects and request notification when the tracked object changes state. This feature increases the availability and speed of recovery of a routing system and decreases outages and outage duration.

Related Topics

[Chapter 4, “Configuring Enhanced Object Tracking”](#)



Configuring IP Unicast Routing

This chapter describes how to configure IP Version 4 (IPv4) unicast routing on the Cisco Industrial Ethernet 2000U Series Switches, hereafter referred to as IE 2000U or switch. For information about IPv6 routing, see [Chapter 3, “Configuring IPv6 Unicast Routing.”](#)



Note

Routing is supported only on switches that are running the IP services image.

For more detailed IPv4 unicast configuration information and complete syntax and usage information for the commands used in this chapter, see documents listed in the [“Related Documents” section on page 2-135](#).

This chapter includes the following sections:

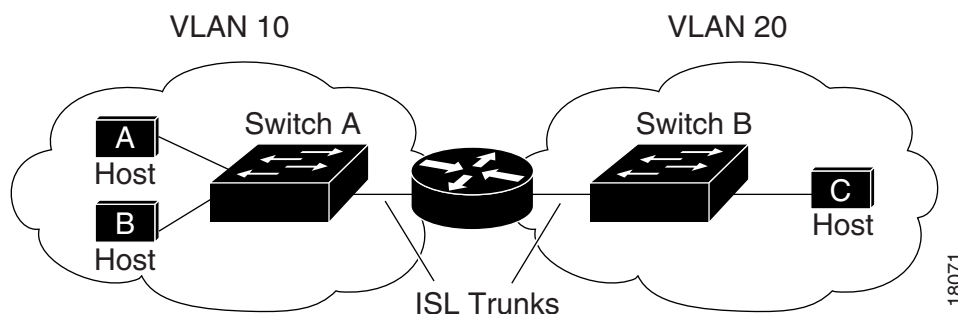
- [Information About IP Routing, page 2-2](#)
- [Prerequisites, page 2-3](#)
- [Guidelines and Limitations, page 2-3](#)
- [Configuring IP Addressing, page 2-4](#)
- [Enabling IPv4 Unicast Routing, page 2-21](#)
- [Configuring RIP, page 2-22](#)
- [Configuring OSPF, page 2-28](#)
- [Configuring EIGRP, page 2-42](#)
- [Configuring BGP, page 2-51](#)
- [Configuring ISO CLNS Routing, page 2-76](#)
- [Configuring BFD, page 2-88](#)
- [Configuring Multi-VRF CE, page 2-100](#)
- [Configuring Protocol-Independent Features, page 2-116](#)
- [Verifying Configuration, page 2-134](#)
- [Related Documents, page 2-135](#)

Information About IP Routing

In an IP network, each subnetwork is mapped to an individual VLAN. However, network devices in different VLANs cannot communicate with one another without a Layer 3 device (router) to route traffic between the VLAN, referred to as inter-VLAN routing. You configure one or more routers to route traffic to the appropriate destination VLAN.

Figure 2-1 shows a basic routing topology. Switch A is in VLAN 10, and Switch B is in VLAN 20. The router has an interface in each VLAN.

Figure 2-1 Routing Topology Example



When Host A in VLAN 10 needs to communicate with Host B in VLAN 10, it sends a packet addressed to that host. Switch A forwards the packet directly to Host B, without sending it to the router.

When Host A sends a packet to Host C in VLAN 20, Switch A forwards the packet to the router, which receives the traffic on the VLAN 10 interface. The router checks the routing table, finds the correct outgoing interface, and forwards the packet on the VLAN 20 interface to Switch B. Switch B receives the packet and forwards it to Host C.

Types of Routing

Routers and Layer 3 switches can route packets in the following ways:

- By using default routing—sending traffic with a destination unknown to the router to a default outlet or destination.
- By using preprogrammed static routes for the traffic

Static unicast routing forwards packets from predetermined ports through a single path into and out of a network. Static routing does not automatically respond to changes in the network and therefore, might result in unreachable destinations.

- By dynamically calculating routes by using a routing protocol

Dynamic routing protocols are used by routers to dynamically calculate the best route for forwarding traffic. Routing protocols supported by the switch are Routing Information Protocol (RIP), Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) protocol, Enhanced IGRP (EIGRP), System-to-Intermediate System (IS-IS), and Bidirectional Forwarding Detection (BFD).

Prerequisites

- To support VLAN interfaces, create and configure VLANs on the switch, and assign VLAN membership to Layer 2 interfaces. For more information, see the chapter “Configuring VLANs” in the *Cisco Connected Grid Switches Layer 2 Switching Software Configuration Guide*.
- By default, IPv4 routing is disabled on the switch, and you must enable it before routing can take place. See the “[Enabling IPv4 Unicast Routing](#)” section on page 2-21.
- We recommend that you configure the BFD interval parameters on an interface before configuring the routing protocol commands, especially when using EIGRP. For information about BFD, see the “[Configuring BFD](#)” section on page 2-88.

Guidelines and Limitations

- In the following procedures, the specified interface must be one of these Layer 3 interfaces:
 - A routed port: a physical port configured as a Layer 3 port by using the **no switchport** interface configuration command.
 - A switch virtual interface (SVI): a VLAN interface created by using the **interface vlan** *vlan_id* global configuration command and by default a Layer 3 interface.
 - An EtherChannel port channel in Layer 3 mode: a port-channel logical interface created by using the **interface port-channel** *port-channel-number* global configuration command and binding the Ethernet interface into the channel group. For more information, see the “Configuring Layer 3 EtherChannels” section in the “Configuring EtherChannels and Link State Tracking” chapter in the *Cisco Connected Grid Switches High Availability and Redundancy Software Configuration Guide*.
- The switch does not support tunnel interfaces for unicast routed traffic.
- All Layer 3 interfaces on which routing will occur must have IP addresses assigned to them. See the “[Assigning IP Addresses to Network Interfaces](#)” section on page 2-5.
- A Layer 3 switch can have an IP address assigned to each routed port and SVI. The number of routed ports and SVIs that you can configure is not limited by software. However, the interrelationship between this number and the number and volume of features being implemented might have an impact on CPU utilization because of hardware limitations.

To support IPv4 routing, use the **sdm prefer default** global configuration command to set the Switch Database Management (sdm) feature to balance resources. For more information on the SDM templates, see the “Configuring SDM Templates” chapter in the *Cisco Connected Grid Switches System Management Software Configuration Guide* or see the **sdm prefer** command in the command reference listed in the “[Related Documents](#)” section on page 2-135.

Steps for Configuring Routing

Configuring IPv4 routing consists of several main procedures:

- Configure Layer 3 interfaces.
- Enable IPv4 routing on the switch.
- Assign IPv4 addresses to the Layer 3 interfaces.
- Enable selected routing protocols on the switch.

- Configure routing protocol parameters (optional).

Configuring IP Addressing

IP routing requires that Layer 3 network interfaces are assigned IP addresses to enable the interfaces and to allow communication with the hosts on interfaces that use IP. These sections describe how to configure various IP addressing features. Assigning IP addresses to the interface is required; the other procedures are optional.

- [Default Addressing Configuration, page 2-4](#)
- [Assigning IP Addresses to Network Interfaces, page 2-5](#)
- [Configuring Address Resolution Methods, page 2-8](#)
- [Routing Assistance When IP Routing is Disabled, page 2-12](#)
- [Configuring Broadcast Packet Handling, page 2-14](#)
- [Monitoring and Maintaining IP Addressing, page 2-20](#)

Default Addressing Configuration

Feature	Default Setting
IP address	None defined.
ARP	No permanent entries in the Address Resolution Protocol (ARP) cache. Encapsulation: Standard Ethernet-style ARP. Timeout: 14400 seconds (4 hours).
IP broadcast address	255.255.255.255 (all ones).
IP classless routing	Enabled.
IP default gateway	Disabled.
IP directed broadcast	Disabled (all IP directed broadcasts are dropped).
IP domain	Domain list: No domain names defined. Domain lookup: Enabled. Domain name: Enabled.
IP forward-protocol	If a helper address is defined or User Datagram Protocol (UDP) flooding is configured, UDP forwarding is enabled on default ports. Any-local-broadcast: Disabled. Turbo-flood: Disabled.
IP helper address	Disabled.
IP host	Disabled.

Feature	Default Setting
IRDP	Disabled. Defaults when enabled: <ul style="list-style-type: none"> Broadcast IRDP advertisements. Maximum interval between advertisements: 600 seconds. Minimum interval between advertisements: 0.75 times max interval. Preference: 0.
IP proxy ARP	Enabled.
IP routing	Disabled.
IP subnet-zero	Disabled.

Assigning IP Addresses to Network Interfaces

An IP address identifies a location to which IP packets can be sent. An interface can have one primary IP address. A mask identifies the bits that denote the network number in an IP address. When you use the mask to subnet a network, the mask is referred to as a subnet mask.

BEFORE YOU BEGIN

To receive an assigned network number, contact your Internet service provider.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. User network interfaces (UNIs) and enhanced network interfaces (ENIs) are disabled by default; network node interfaces (NNIs) are enabled by default.
Step 4	no switchport	Remove the interface from Layer 2 configuration mode (if it is a physical interface).
Step 5	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet mask.
Step 6	end	Return to privileged EXEC mode.
Step 7	show interfaces [<i>interface-id</i>] show ip interface [<i>interface-id</i>] show running-config interface [<i>interface-id</i>]	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# interface gi0/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.1.5.1 255.255.255.0
Switch(config-if)# end
```

Enabling Subnet Zero

Enabling subnet zero provides the ability to configure and route to subnet 0 subnets.

You can use the all ones subnet (131.108.255.0) and even though it is discouraged, you can enable the use of subnet zero if you need the entire subnet space for your IP address.

BEFORE YOU BEGIN

Subnetting with a subnet address of zero is strongly discouraged because of the problems that can arise if a network and a subnet have the same addresses. For example, if network 131.108.0.0 is subnetted as 255.255.255.0, subnet zero would be written as 131.108.0.0, which is the same as the network address.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip subnet-zero	Enable the use of subnet zero for interface addresses and routing updates.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

Use the **no ip subnet-zero** global configuration command to restore the default and disable the use of subnet zero.

EXAMPLE

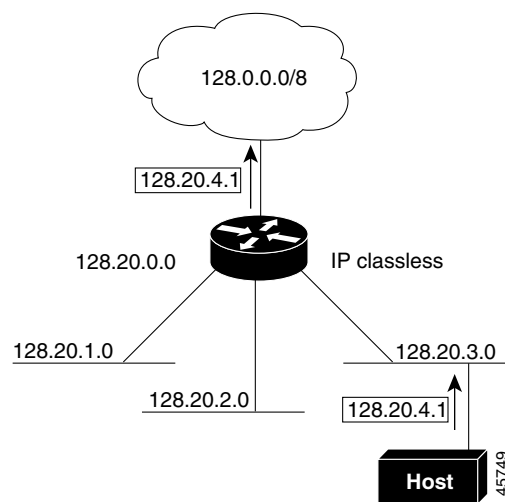
```
Switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
Switch(config)# ip subnet-zero
Switch(config)# end
```

Classless Routing

By default, classless routing behavior is enabled on the switch when it is configured to route. With classless routing, if a router receives packets for a subnet of a network with no default route, the router forwards the packet to the best supernet route. A *supernet* consists of contiguous blocks of Class C address spaces used to simulate a single, larger address space and is designed to relieve the pressure on the rapidly depleting Class B address space.

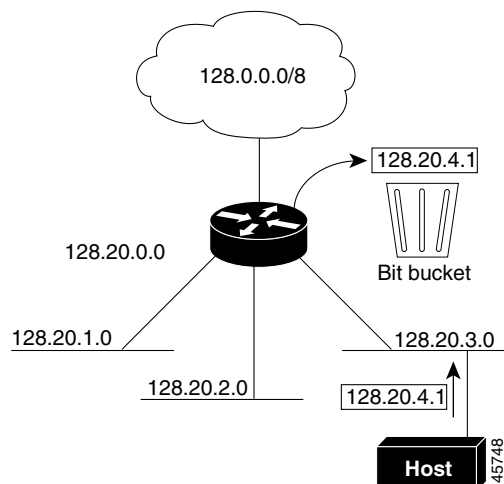
In [Figure 2-2](#), classless routing is enabled. When the host sends a packet to 120.20.4.1, instead of discarding the packet, the router forwards it to the best supernet route. If you disable classless routing and a router receives packets destined for a subnet of a network with no network default route, the router discards the packet.

Figure 2-2 IP Classless Routing



In [Figure 2-3](#), the router in network 128.20.0.0 is connected to subnets 128.20.1.0, 128.20.2.0, and 128.20.3.0. If the host sends a packet to 120.20.4.1, because there is no network default route, the router discards the packet.

Figure 2-3 No IP Classless Routing



To prevent the switch from forwarding packets destined for unrecognized subnets to the best supernet route possible, you can disable classless routing behavior.

BEFORE YOU BEGIN

Review the [“Information About IP Routing”](#) section on page 2-2.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	no ip classless	Disable classless routing behavior.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To restore the default and have the switch forward packets destined for a subnet of a network with no network default route to the best supernet route possible, use the **ip classless** global configuration command.

EXAMPLE

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# no ip classless
Switch(config)# end
```

Configuring Address Resolution Methods

You can control interface-specific handling of IP by using address resolution. A device using IP can have both a local address or MAC address, which uniquely defines the device on its local segment or LAN, and a network address, which identifies the network to which the device belongs. To communicate with a device on Ethernet, the software must learn the MAC address of the device. The process of learning the MAC address from an IP address is called *address resolution*. The process of learning the IP address from the MAC address is called *reverse address resolution*.

The switch can use these forms of address resolution:

- Address Resolution Protocol (ARP) is used to associate IP address with MAC addresses. Taking an IP address as input, ARP learns the associated MAC address and then stores the IP address/MAC address association in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of IP datagrams and ARP requests or replies on IEEE 802 networks other than Ethernet is specified by the Subnetwork Access Protocol (SNAP).
- Proxy ARP helps hosts with no routing tables learn the MAC addresses of hosts on other networks or subnets. If the switch (router) receives an ARP request for a host that is not on the same interface as the ARP request sender, and if the router has all of its routes to the host through other interfaces, it generates a proxy ARP packet giving its own local data link address. The host that sent the ARP request then sends its packets to the router, which forwards them to the intended host.

The switch also uses the Reverse Address Resolution Protocol (RARP), which functions the same as ARP does, except that the RARP packets request an IP address instead of a local MAC address. Using RARP requires a RARP server on the same network segment as the router interface. Use the **ip rarp-server address** interface configuration command to identify the server.

For more information on RARP, see the [IP Addressing: ARP Configuration Guide, Cisco IOS Release 15M&T](#).

You can perform these tasks to configure address resolution:

- [Defining a Static ARP Cache, page 2-9](#)
- [Setting ARP Encapsulation, page 2-10](#)
- [Enabling Proxy ARP, page 2-11](#)

Defining a Static ARP Cache

ARP and other address resolution protocols provide dynamic mapping between IP addresses and MAC addresses. Because most hosts support dynamic address resolution, you usually do not need to specify static ARP cache entries. If you must define a static ARP cache entry, you can do so globally, which installs a permanent entry in the ARP cache that the switch uses to translate IP addresses into MAC addresses. Optionally, you can also specify that the switch respond to ARP requests as if it were the owner of the specified IP address. If you do not want the ARP entry to be permanent, you can specify a timeout period for the ARP entry.

BEFORE YOU BEGIN

Review the [“Configuring Address Resolution Methods” section on page 2-8](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	arp ip-address hardware-address type	Globally associate an IP address with a MAC (hardware) address in the ARP cache, and specify encapsulation type as one of these: <ul style="list-style-type: none"> • arpa—ARP encapsulation for Ethernet interfaces • snap—Subnetwork Address Protocol encapsulation for Token Ring and FDDI interfaces • sap—HP’s ARP type
Step 3	arp ip-address hardware-address type [alias]	(Optional) Specify that the switch respond to ARP requests as if it were the owner of the specified IP address.
Step 4	interface interface-id	Enter interface configuration mode, and specify the interface to configure.
Step 5	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.

	Command	Purpose
Step 6	arp timeout <i>seconds</i>	(Optional) Set the length of time an ARP cache entry will stay in the cache. The default is 14400 seconds (4 hours). The range is 0 to 2147483 seconds.
Step 7	end	Return to privileged EXEC mode.
Step 8	show interfaces [<i>interface-id</i>]	Verify the type of ARP and the timeout value used on all interfaces or a specific interface.
Step 9	show arp or show ip arp	View the contents of the ARP cache.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an entry from the ARP cache, use the **no arp ip-address hardware-address type** global configuration command. To remove all nonstatic entries from the ARP cache, use the **clear arp-cache** privileged EXEC command.

EXAMPLE

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# arp 10.0.0.0 aabb.cc03.8200 arpa
Switch(config)# end
```

Setting ARP Encapsulation

By default, Ethernet ARP encapsulation (represented by the **arpa** keyword) is enabled on an IP interface. You can change the encapsulation methods to SNAP if required by your network.

BEFORE YOU BEGIN

The encapsulation type specified in this procedure should match the encapsulation type specified in the [“Defining a Static ARP Cache” procedure on page 2-9](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	arp {arpa snap}	Specify the ARP encapsulation method: <ul style="list-style-type: none"> arpa—Address Resolution Protocol snap—Subnetwork Address Protocol
Step 5	end	Return to privileged EXEC mode.

	Command	Purpose
Step 6	show interfaces [<i>interface-id</i>]	Verify ARP encapsulation configuration on all interfaces or the specified interface.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an encapsulation type, use the **no arp arpa** or **no arp snap** interface configuration command.

EXAMPLE

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gi0/2
Switch(config-if)# arp arpa
Switch(config-if)# end
```

Enabling Proxy ARP

By default, the switch uses proxy ARP to help hosts learn MAC addresses of hosts on other networks or subnets. Follow these steps to enable proxy ARP if it has been disabled.

BEFORE YOU BEGIN

Review the [“Configuring Address Resolution Methods”](#) section on page 2-8.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip proxy-arp	Enable proxy ARP on the interface.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip interface [<i>interface-id</i>]	Verify the configuration on the interface or all interfaces.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable proxy ARP on the interface, use the **no ip proxy-arp** interface configuration command.

EXAMPLE

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gi0/2
Switch(config-if)# ip proxy-arp
Switch(config-if)# end
```

Routing Assistance When IP Routing is Disabled

These mechanisms allow the switch to learn about routes to other networks when it does not have IP routing enabled:

- [Proxy ARP, page 2-12](#)
- [Default Gateway, page 2-12](#)
- [ICMP Router Discovery Protocol \(IRDP\), page 2-13](#)

Proxy ARP

Proxy ARP, the most common method for learning about other routes, enables an Ethernet host with no routing information to communicate with hosts on other networks or subnets. The host assumes that all hosts are on the same local Ethernet and that they can use ARP to learn their MAC addresses. If a switch receives an ARP request for a host that is not on the same network as the sender, the switch evaluates whether it has the best route to that host. If it does, it sends an ARP reply packet with its own Ethernet MAC address, and the host that sent the request sends the packet to the switch, which forwards it to the intended host. Proxy ARP treats all networks as if they are local and performs ARP requests for every IP address.

Proxy ARP is enabled by default. To enable it after it has been disabled, see the [“Enabling Proxy ARP” section on page 2-11](#). Proxy ARP works as long as other routers support it.

Default Gateway

Another method for locating routes is to define a default router or default gateway. All nonlocal packets are sent to this router, which either routes them appropriately or sends an IP Control Message Protocol (ICMP) redirect message back, defining which local router the host should use. The switch caches the redirect messages and forwards each packet as efficiently as possible. A limitation of this method is that there is no means of detecting when the default router has gone down or is unavailable.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip default-gateway <i>ip-address</i>	Set up a default gateway (router).
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip redirects	Display the address of the default gateway router to verify the setting.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip default-gateway** global configuration command to disable this function.

EXAMPLE

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip default-gateway 192.31.7.18
Switch(config)# end
```

ICMP Router Discovery Protocol (IRDP)

Router discovery allows the switch to dynamically learn about routes to other networks using IRDP. IRDP allows hosts to locate routers. When operating as a client, the switch generates router discovery packets. When operating as a host, the switch receives router discovery packets. The switch can also listen to Routing Information Protocol (RIP) routing updates and use this information to infer locations of routers. The switch does not actually store the routing tables sent by routing devices; it merely keeps track of which systems are sending the data. The advantage of using IRDP is that it allows each router to specify both a priority and the time after which a device is assumed to be down if no further packets are received.

Each device discovered becomes a candidate for the default router, and a new highest-priority router is selected when a higher priority router is discovered, when the current default router is declared down, or when a TCP connection is about to time out because of excessive retransmissions.

The only required task for IRDP routing on an interface is to enable IRDP processing on that interface. When enabled, the default parameters apply. You can optionally change any of these parameters.

BEFORE YOU BEGIN

- The **ip irdp multicast** command allows for compatibility with Sun Microsystems Solaris, which requires IRDP packets to be sent out as multicasts. Many implementations cannot receive these multicasts; ensure end-host ability before using this command.
- If you change the **maxadvertinterval** value, the **holdtime** and **minadvertinterval** values also change, so it is important to first change the **maxadvertinterval** value before manually changing either the **holdtime** or **minadvertinterval** values.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip irdp	Enable IRDP processing on the interface.
Step 5	ip irdp multicast	(Optional) Send IRDP advertisements to the multicast address (224.0.0.1) instead of IP broadcasts.
Step 6	ip irdp holdtime <i>seconds</i>	(Optional) Set the IRDP period for which advertisements are valid. The default is three times the maxadvertinterval value. It must be greater than maxadvertinterval and cannot be greater than 9000 seconds. If you change the maxadvertinterval value, this value also changes.
Step 7	ip irdp maxadvertinterval <i>seconds</i>	(Optional) Set the IRDP maximum interval between advertisements. The default is 600 seconds.
Step 8	ip irdp minadvertinterval <i>seconds</i>	(Optional) Set the IRDP minimum interval between advertisements. The default is 0.75 times the maxadvertinterval . If you change the maxadvertinterval , this value changes to the new default (0.75 of maxadvertinterval).

	Command	Purpose
Step 9	ip irdp preference <i>number</i>	(Optional) Set a device IRDP preference level. The allowed range is -2^{31} to 2^{31} . The default is 0. A higher value increases the router preference level.
Step 10	ip irdp address <i>address</i> [<i>number</i>]	(Optional) Specify an IRDP address and preference to proxy-advertise.
Step 11	end	Return to privileged EXEC mode.
Step 12	show ip irdp	Verify settings by displaying IRDP values.
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip irdp** interface configuration command to disable IRDP routing.

EXAMPLE

```
Switch(config)# interface ethernet 0 !Enable irdp on interface Ethernet 0.
Switch(config-if)# ip irdp
Switch(config-if)# ip irdp multicast !Send IRDP advertisements to the multicast address.
Switch(config-if)# ip irdp preference 900 !Increase router preference from 0 to 900.
Switch(config-if)# ip irdp maxadvertinterval 400 !Set maximum time between advertisements
to 400 secs.
Switch(config-if)# ip irdp minadvertinterval 100 !Set minimum time between advertisements
to 100 secs.
Switch(config-if)# ip irdp holdtime 6000 !Advertisements are good for 6000 seconds.
Switch(config-if)# ip irdp address 10.108.14.5 !Proxy-advertise 10.108.14.5 with default
router preference.
Switch(config-if)# ip irdp address 10.108.14.6 50 !Proxy-advertise 10.108.14.6 with
preference of 50.
```

Configuring Broadcast Packet Handling

After configuring an IP interface address, you can enable routing and configure one or more routing protocols, or you can configure the way the switch responds to network broadcasts. A broadcast is a data packet destined for all hosts on a physical network. The switch supports two kinds of broadcasting:

- A directed broadcast packet is sent to a specific network or series of networks. A directed broadcast address includes the network or subnet fields.
- A flooded broadcast packet is sent to every network.



Note

You can also limit broadcast, unicast, and multicast traffic on Layer 2 interfaces by using the **storm-control** interface configuration command to set traffic suppression levels. For more information, see the “Configuring Port-Based Traffic Control” chapter in the *Cisco Connected Grid Switches System Management Software Configuration Guide*.

Routers provide some protection from broadcast storms by limiting their extent to the local cable. Bridges (including intelligent bridges), because they are Layer 2 devices, forward broadcasts to all network segments, thus propagating broadcast storms. The best solution to the broadcast storm problem is to use a single broadcast address scheme on a network. In most modern IP implementations, you can set the address to be used as the broadcast address. The switch supports several addressing schemes for forwarding broadcast messages.

- [Enabling Directed Broadcast-to-Physical Broadcast Translation, page 2-15](#)
- [Forwarding UDP Broadcast Packets and Protocols, page 2-16](#)
- [Establishing an IP Broadcast Address, page 2-17](#)
- [Flooding IP Broadcasts, page 2-18](#)

Enabling Directed Broadcast-to-Physical Broadcast Translation

By default, IP-directed broadcasts are not forwarded; they are dropped to make routers less susceptible to denial-of-service attacks. You can enable forwarding of IP-directed broadcasts on an interface where the broadcast becomes a physical (MAC-layer) broadcast. Only those protocols configured by using the **ip forward-protocol** global configuration command are forwarded.

BEFORE YOU BEGIN

You can specify an access list to control which broadcasts are forwarded. Only those IP packets permitted by the access list are eligible to be translated from directed broadcasts to physical broadcasts. For more information on access lists, see the “Configuring Network Security with ACLs” chapter in the *Cisco Connected Grid Switches Security Software Configuration Guide*.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip directed-broadcast [<i>access-list-number</i>]	Enable directed broadcast-to-physical broadcast translation on the interface. You can include an access list to control which broadcasts are forwarded. When an access list is specified, only IP packets permitted by the access list are eligible to be translated.
Step 5	exit	Return to global configuration mode.
Step 6	ip forward-protocol { udp [<i>port</i>] nd sdns }	Specify which protocols and ports the router forwards when forwarding broadcast packets. <ul style="list-style-type: none"> • udp—Forward UDP datagrams. <i>port</i>: (Optional) Destination port that controls which UDP services are forwarded. • nd—Forward ND datagrams. • sdns—Forward SDNS datagrams.
Step 7	end	Return to privileged EXEC mode.

	Command	Purpose
Step 8	show ip interface <i>[interface-id]</i>	Verify the configuration on the interface or all interfaces.
	or show running-config	
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip directed-broadcast** interface configuration command to disable translation of directed broadcast to physical broadcasts. Use the **no ip forward-protocol** global configuration command to remove a protocol or port.

EXAMPLE

The following example enables forwarding of IP directed broadcasts on Ethernet interface 0. The **ip forward-protocol** command using the **udp** keyword without specifying any port numbers allows forwarding of UDP packets on the default ports.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface ethernet 0
Switch(config-if)# ip directed-broadcast
Switch(config-if)# exit
Switch(config)# ip forward-protocol udp
Switch(config)# end
```

Forwarding UDP Broadcast Packets and Protocols

User Datagram Protocol (UDP) is an IP host-to-host layer protocol that provides a low-overhead, connectionless session between two end systems and does not provide for acknowledgment of received datagrams. Network hosts occasionally use UDP broadcasts to find address, configuration, and name information. If such a host is on a network segment that does not include a server, UDP broadcasts are normally not forwarded. You can configure an interface on a router to forward certain classes of broadcasts to a helper address. You can use more than one helper address per interface.

You can specify a UDP destination port to control which UDP services are forwarded. You can specify multiple UDP protocols. You can also specify the Network Disk (ND) protocol, which is used by older diskless Sun workstations and the network security protocol SDNS.

By default, both UDP and ND forwarding are enabled if a helper address has been defined for an interface.

If you do not specify any UDP ports when you configure the forwarding of UDP broadcasts, you are configuring the router to act as a BOOTP forwarding agent. BOOTP packets carry DHCP information.

BEFORE YOU BEGIN

See the description for the **ip forward-protocol** interface configuration command in the [Cisco IOS IP Application Services Command Reference](#) for the list of ports that are forwarded by default if you do not specify any UDP ports.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip helper-address <i>address</i>	Enable forwarding and specify the destination address for forwarding UDP broadcast packets, including BOOTP.
Step 5	exit	Return to global configuration mode.
Step 6	ip forward-protocol { udp [<i>port</i>] nd sdns }	Specify which protocols the router forwards when forwarding broadcast packets.
Step 7	end	Return to privileged EXEC mode.
Step 8	show ip interface [<i>interface-id</i>] or show running-config	Verify the configuration on the interface or all interfaces.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip helper-address** interface configuration command to disable the forwarding of broadcast packets to specific addresses. Use the **no ip forward-protocol** global configuration command to remove a protocol or port.

EXAMPLE

The following example defines a helper address and uses the **ip forward-protocol** command. Using the **udp** keyword without specifying any port numbers will allow forwarding of UDP packets on the default ports.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface ethernet 0
Switch(config-if)# ip helper-address 10.24.42.2
Switch(config-if)# exit
Switch(config)# ip forward-protocol udp
Switch(config)# end
```

Establishing an IP Broadcast Address

The most popular IP broadcast address (and the default) is an address consisting of all ones (255.255.255.255). However, the switch can be configured to generate any form of IP broadcast address.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip broadcast-address <i>ip-address</i>	Enter a broadcast address different from the default, for example 128.1.255.255.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip interface [<i>interface-id</i>]	Verify the broadcast address on the interface or all interfaces.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To restore the default IP broadcast address, use the **no ip broadcast-address** interface configuration command.

EXAMPLE

The following example specifies an IP broadcast address of 0.0.0.0:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface ethernet 0
Switch(config-if)# ip broadcast-address 0.0.0.0
Switch(config-if)# end
```

Flooding IP Broadcasts

You can allow IP broadcasts to be flooded throughout your internetwork in a controlled fashion by using the database created by the bridging STP. Using this feature also prevents loops. To support this capability, bridging must be configured on each interface that is to participate in the flooding. If bridging is not configured on an interface, the interface can receive broadcasts but it never forwards the broadcasts it receives, and the router never uses that interface to send broadcasts received on a different interface.

Packets that are forwarded to a single network address using the IP helper-address mechanism can be flooded. Only one copy of the packet is sent on each network segment.

To be considered for flooding, packets must meet these criteria. (Note that these are the same conditions used to consider packet forwarding using IP helper addresses.)

- The packet must be a MAC-level broadcast.
- The packet must be an IP-level broadcast.
- The packet must be a TFTP, DNS, Time, NetBIOS, ND, or BOOTP packet, or a UDP specified by the **ip forward-protocol udp** global configuration command.
- The time-to-live (TTL) value of the packet must be at least two.

A flooded UDP datagram is given the destination address specified with the **ip broadcast-address** interface configuration command on the output interface. The destination address can be set to any address so it might change as the datagram propagates through the network. The source address is never changed. The TTL value is decremented.

When a flooded UDP datagram is sent out an interface (and the destination address possibly changed), the datagram is handed to the normal IP output routines and is, therefore, subject to access lists, if they are present on the output interface.

BEFORE YOU BEGIN

Ensure that bridging is configured on each interface that is to participate in the flooding.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip forward-protocol spanning-tree	Use the bridging spanning-tree database to flood UDP datagrams.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

Use the **no ip forward-protocol spanning-tree** global configuration command to disable the flooding of IP broadcasts.

EXAMPLE

The following example permits IP broadcasts to be flooded through the internetwork in a controlled fashion:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip forward-protocol spanning-tree
Switch(config)# end
```

Speeding up STP-Based UDP Flooding

In the switch, the majority of packets are forwarded in hardware; most packets do not go through the switch CPU. For those packets that do go to the CPU, you can speed up spanning tree-based UDP flooding by a factor of about four to five times by using turbo-flooding. This feature is supported over Ethernet interfaces configured for ARP encapsulation.

BEFORE YOU BEGIN

Enable the flooding of IP broadcasts as described in the [“Flooding IP Broadcasts” procedure on page 2-18](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode
Step 2	ip forward-protocol turbo-flood	Use the spanning-tree database to speed up flooding of UDP datagrams.
Step 3	end	Return to privileged EXEC mode.
Step 4	show running-config	Verify your entry.
Step 5	copy running-config startup-config	(Optional) Save your entry in the configuration file.

To disable this feature, use the **no ip forward-protocol turbo-flood** global configuration command.

EXAMPLE

The following example shows how to speed up the flooding of UDP packets using the spanning-tree algorithm:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip forward-protocol turbo-flood
Switch(config)# end
```

Monitoring and Maintaining IP Addressing

When the contents of a particular cache, table, or database have become or are suspected to be invalid, you can remove all its contents by using the **clear** privileged EXEC commands.

Command	Purpose
clear arp-cache	Clear the IP ARP cache and the fast-switching cache.
clear host { <i>name</i> *}	Remove one or all entries from the hostname and the address cache.
clear ip route { <i>network</i> [<i>mask</i>] *}	Remove one or more routes from the IP routing table.

You can display specific statistics, such as the contents of IP routing tables, caches, and databases; the reachability of nodes; and the routing path that packets are taking through the network.

Command	Purpose
show arp	Display the entries in the ARP table.
show hosts	Display the default domain name, style of lookup service, name server hosts, and the cached list of hostnames and addresses.
show ip aliases	Display IP addresses mapped to TCP ports (aliases).
show ip arp	Display the IP ARP cache.
show ip interface [<i>interface-id</i>]	Display the IP status of interfaces.

Command	Purpose
show ip irdp	Display IRDP values.
show ip masks <i>address</i>	Display the masks used for network addresses and the number of subnets using each mask.
show ip redirects	Display the address of a default gateway.
show ip route [<i>address</i> [<i>mask</i>]] [<i>protocol</i>]	Display the current state of the routing table.
show ip route summary	Display the current state of the routing table in summary form.

Enabling IPv4 Unicast Routing

By default, the switch is in Layer 2 switching mode and IP routing is disabled. To use the Layer 3 capabilities of the switch, you must enable IP routing.

BEFORE YOU BEGIN

Review the [“Guidelines and Limitations” section on page 2-3](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing.
Step 3	router <i>ip_routing_protocol</i>	Specify an IP routing protocol. This step might include other commands, such as specifying the networks to route with the network (RIP) router configuration command. For information on specific protocols, see sections later in this chapter.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip routing** global configuration command to disable routing.

EXAMPLE

This example shows how to enable IP routing using RIP as the routing protocol:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# router rip
Switch(config-router)# network 10.0.0.0
Switch(config-router)# end
```

Configuring RIP

The Routing Information Protocol (RIP) is an interior gateway protocol (IGP) used in small, homogeneous networks. It is a distance-vector routing protocol that uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. You can find detailed information about RIP in *IP Routing Fundamentals*, published by Cisco Press.

Using RIP, the switch sends routing information updates (advertisements) every 30 seconds. If a router does not receive an update from another router for 180 seconds or more, it marks the routes served by that router as unusable. If there is still no update after 240 seconds, the router removes all routing table entries for the non-updating router.

RIP uses hop counts to rate the value of different routes. The hop count is the number of routers that can be traversed in a route. A directly connected network has a hop count of zero; a network with a hop count of 16 is unreachable. This small range (0 to 15) makes RIP unsuitable for large networks.

If the router has a default network path, RIP advertises a route that links the router to the pseudonetwork 0.0.0.0. The 0.0.0.0 network does not exist, but is treated by RIP as a network to implement default routing. The switch advertises the default network if a default was learned by RIP or if the router has a gateway of last resort and RIP is configured with a default metric. RIP sends updates to the interfaces in specified networks. If an interface's network is not specified, it is not advertised in any RIP update.

This section includes the following topics:

- [Default RIP Configuration, page 2-22](#)
- [Configuring Basic RIP Parameters, page 2-23](#)
- [Configuring RIP Authentication, page 2-25](#)
- [Configuring Split Horizon, page 2-26](#)

Default RIP Configuration

Feature	Default Setting
Auto summary	Enabled.
Default-information originate	Disabled.
Default metric	Built-in; automatic metric translations.
IP RIP authentication key-chain	No authentication. Authentication mode: clear text.
IP RIP receive version	According to the version router configuration command.
IP RIP send version	According to the version router configuration command.
IP RIP triggered	According to the version router configuration command.
IP split horizon	Varies with media.
Neighbor	None defined.
Network	None specified.
Offset list	Disabled.
Output delay	0 milliseconds.

Feature	Default Setting
Timers basic	<ul style="list-style-type: none"> • Update: 30 seconds. • Invalid: 180 seconds. • Hold-down: 180 seconds. • Flush: 240 seconds.
Validate-update-source	Enabled.
Version	Receives RIP Version 1 and 2 packets; sends Version 1 packets.

Configuring Basic RIP Parameters

To configure RIP, you enable RIP routing for a network and optionally configure other parameters. On the switch, RIP configuration commands are ignored until you configure the network number.

BEFORE YOU BEGIN

Complete the RIP network strategy and planning for your network. For example, you must decide whether to receive and send only RIP Version 1 or RIP Version 2 packets and whether to use RIP authentication. (RIP Version 1 does not support authentication.)

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing. (Required only if IP routing is disabled.)
Step 3	router rip	Enable a RIP routing process, and enter router configuration mode.
Step 4	network <i>network number</i>	Associate a network with a RIP routing process. You can specify multiple network commands. RIP routing updates are sent and received through interfaces only on these networks. Note You must configure a network number for RIP commands to take effect.
Step 5	neighbor <i>ip-address</i>	(Optional) Define a neighboring router with which to exchange routing information. This step allows routing updates from RIP (normally a broadcast protocol) to reach nonbroadcast networks.
Step 6	offset list [<i>access-list number name</i>] { in out } <i>offset [type number]</i>	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through RIP. You can limit the offset list with an access list or an interface.

	Command	Purpose
Step 7	timers basic <i>update invalid holddown flush</i>	(Optional) Adjust routing protocol timers. Valid ranges for all timers are 0 to 4294967295 seconds. <ul style="list-style-type: none"> <i>update</i>—The time between sending routing updates. The default is 30 seconds. <i>invalid</i>—The timer after which a route is declared invalid. The default is 180 seconds. <i>holddown</i>—The time before a route is removed from the routing table. The default is 180 seconds. <i>flush</i>—The amount of time for which routing updates are postponed. The default is 240 seconds.
Step 8	version {1 2}	(Optional) Configure the switch to receive and send only RIP Version 1 or RIP Version 2 packets. By default, the switch receives Version 1 and 2 but sends only Version 1. You can also use the interface commands ip rip {send receive} version 1 2 1 2 to control what versions are used for sending and receiving on interfaces.
Step 9	no auto summary	(Optional) Disable automatic summarization. By default, the switch summarizes subprefixes when crossing classful network boundaries. Disable summarization (RIP Version 2 only) to advertise subnet and host routing information to classful network boundaries.
Step 10	no validate-update-source	(Optional) Disable validation of the source IP address of incoming RIP routing updates. By default, the switch validates the source IP address of incoming RIP routing updates and discards the update if the source address is not valid. Under normal circumstances, disabling this feature is not recommended. However, if you have a router that is off-network and you want to receive its updates, you can use this command.
Step 11	output-delay <i>delay</i>	(Optional) Add interpacket delay for RIP updates sent. By default, packets in a multiple-packet RIP update have no delay added between packets. If you are sending packets to a lower-speed device, you can add an interpacket delay in the range of 8 to 50 milliseconds.
Step 12	end	Return to privileged EXEC mode.
Step 13	show ip protocols	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To turn off the RIP routing process, use the **no router rip** global configuration command.

To display the parameters and current state of the active routing protocol process, use the **show ip protocols** privileged EXEC command. Use the **show ip rip database** privileged EXEC command to display summary address entries in the RIP database.

EXAMPLE

In the following example, RIP updates are sent to all interfaces on network 10.108.0.0 except Ethernet interface 1. However, in this case, a neighbor router configuration command is included. This command permits the sending of routing updates to specific neighbors. One copy of the routing update is generated per neighbor.

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# router rip
Switch(config-router)# network 10.108.0.0
Router(config-router)# passive-interface Ethernet 1
Router(config-router)# neighbor 10.108.20.4
Router(config-router)# end
```

Configuring RIP Authentication

RIP Version 1 does not support authentication. If you are sending and receiving RIP Version 2 packets, you can enable RIP authentication on an interface. The key chain specifies the set of keys that can be used on the interface. If a key chain is not configured, no authentication is performed, not even the default. Therefore, you must also perform the tasks in the [“Managing Authentication Keys”](#) section on page 2-133.

The switch supports two modes of authentication on interfaces for which RIP authentication is enabled: plain text and MD5. The default is plain text.

BEFORE YOU BEGIN

Configure RIP as described in the [“Configuring Basic RIP Parameters”](#) procedure on page 2-23.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip rip authentication key-chain <i>name-of-chain</i>	Enable RIP authentication.
Step 5	ip rip authentication mode [text md5]	Configure the interface to use plain text authentication (the default) or MD5 digest authentication.
Step 6	end	Return to privileged EXEC mode.
Step 7	show running-config interface [<i>interface-id</i>]	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To restore clear text authentication, use the **no ip rip authentication mode** interface configuration command. To prevent authentication, use the **no ip rip authentication key-chain** interface configuration command.

EXAMPLE

The following example configures the interface to accept and send any key belonging to the key chain named trees and configures the interface to use MD5 authentication:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface ethernet 0
Switch(config-if)# ip rip authentication key-chain trees
Switch(config-if)# ip rip authentication mode md5
Switch(config-if)# end
```

Configuring Split Horizon

Routers connected to broadcast-type IP networks and using distance-vector routing protocols normally use the split-horizon mechanism to reduce the possibility of routing loops. Split horizon blocks information about routes from being advertised by a router on any interface from which that information originated. This feature can optimize communication among multiple routers when links are broken.

BEFORE YOU BEGIN

In general, Cisco does not recommend disabling split horizon unless you are certain that your application requires disabling it to properly advertise routes.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet.
Step 5	no ip split-horizon	Disable split horizon on the interface.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip interface <i>interface-id</i>	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To enable the split horizon mechanism, use the **ip split-horizon** interface configuration command.

EXAMPLE

The following simple example disables split horizon on a serial link:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface serial 0
Switch(config-if)# no ip split-horizon
Switch(config-if)# end
```

Configuring Summary Addresses

To configure an interface running RIP to advertise a summarized local IP address pool on a network access server for dial-up clients, use the **ip summary-address rip** interface configuration command.



Note

If split horizon is enabled, neither autosummary nor interface IP summary addresses are advertised.

BEFORE YOU BEGIN

If the interface is in Layer 2 mode (the default), you must enter a **no switchport** interface configuration command before entering the **ip address** interface configuration command.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip address <i>ip-address subnet-mask</i>	Configure the IP address and IP subnet.
Step 5	ip summary-address rip <i>ip address ip-network mask</i>	Configure the IP address to be summarized and the IP network mask.
Step 6	no ip split horizon	Disable split horizon on the interface.
Step 7	end	Return to privileged EXEC mode.
Step 8	show ip interface <i>interface-id</i>	Verify your entries.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IP summarization, use the **no ip summary-address rip** router configuration command.

EXAMPLE

In this example, the major net is 10.0.0.0. The summary address 10.2.0.0 overrides the autosummary address of 10.0.0.0 so that 10.2.0.0 is advertised out interface Gigabit Ethernet port 2, and 10.0.0.0 is not advertised.

```
Switch(config)# router rip
Switch(config-router)# interface gi0/2
```

```
Switch(config-if) # no switchport
Switch(config-if) # ip address 10.1.5.1 255.255.255.0
Switch(config-if) # ip summary-address rip 10.2.0.0 255.255.0.0
Switch(config-if) # no ip split-horizon
Switch(config-if) # exit
Switch(config) # router rip
Switch(config-router) # network 10.0.0.0
Switch(config-router) # neighbor 2.2.2.2 peer-group mygroup
Switch(config-router) # end
```

Configuring OSPF

Open Shortest Path First (OSPF) is an Interior Gateway Protocol (IGP) designed expressly for IP networks, supporting IP subnetting and tagging of externally derived routing information. OSPF also allows packet authentication and uses IP multicast when sending and receiving packets.

This section briefly describes how to configure OSPF. For a complete description of the OSPF commands, see the OSPF documents listed in the [“Related Documents” section on page 2-135](#).



Note

OSPF classifies different media into broadcast, nonbroadcast multiaccess (NBMA), or point-to-point networks. Broadcast and nonbroadcast networks can also be configured as point-to-multipoint networks. The switch supports all these network types.

The Cisco implementation conforms to the OSPF Version 2 specifications with these key features:

- Definition of stub areas is supported.
- Routes learned through any IP routing protocol can be redistributed into another IP routing protocol. At the intradomain level, this means that OSPF can import routes learned through EIGRP and RIP. OSPF routes can also be exported into RIP.
- Plain text and MD5 authentication among neighboring routers within an area is supported.
- Configurable routing interface parameters include interface output cost, retransmission interval, interface transmit delay, router priority, router dead and hello intervals, and authentication key.
- Virtual links are supported.
- Not-so-stubby-areas (NSSAs) per RFC 1587 are supported.

OSPF typically requires coordination among many internal routers, *area border routers* (ABRs) connected to multiple areas, and *autonomous system boundary routers* (ASBRs). The minimum configuration would use all default parameter values, no authentication, and interfaces assigned to areas. If you customize your environment, you must ensure coordinated configuration of all routers.

This section includes the following topics:

- [Default OSPF Configuration, page 2-29](#)
- [Nonstop Forwarding Awareness, page 2-30](#)
- [Configuring OSPF Interfaces, page 2-31](#)
- [Configuring OSPF Network Types, page 2-33](#)
- [Configuring OSPF Area Parameters, page 2-36](#)
- [Configuring Other OSPF Parameters, page 2-38](#)
- [Changing LSA Group Pacing, page 2-40](#)
- [Configuring a Loopback Interface, page 2-41](#)

- [Monitoring OSPF, page 2-42](#)

Default OSPF Configuration

Feature	Default Setting
Interface parameters	Cost: No default cost predefined. Retransmit interval: 5 seconds. Transmit delay: 1 second. Priority: 1. Hello interval: 10 seconds. Dead interval: 4 times the hello interval. No authentication. No password specified. MD5 authentication disabled.
Area	Authentication type: 0 (no authentication). Default cost: 1. Range: Disabled. Stub: No stub area defined. NSSA: No NSSA area defined.
Auto cost	100 Mbps.
Default-information originate	Disabled. When enabled, the default metric setting is 10, and the external route type default is Type 2.
Default metric	Built-in, automatic metric translation, as appropriate for each routing protocol.
Distance OSPF	dist1 (all routes within an area): 110. dist2 (all routes from one area to another): 110. and dist3 (routes from other routing domains): 110.
OSPF database filter	Disabled. All outgoing link-state advertisements (LSAs) are flooded to the interface.
IP OSPF name lookup	Disabled.
Log adjacency changes	Enabled.
Neighbor	None specified.
Neighbor database filter	Disabled. All outgoing LSAs are flooded to the neighbor.
Network area	Disabled.
NSF ¹ awareness	Enabled ² . Allows Layer 3 switches to continue forwarding packets from a neighboring NSF-capable router during hardware or software changes.
Router ID	No OSPF routing process defined.
Summary address	Disabled.

Feature	Default Setting
Timers LSA group pacing	240 seconds.
Timers shortest path first (spf)	spf delay: 5 seconds. spf-holdtime: 10 seconds.
Virtual link	No area ID or router ID defined. Hello interval: 10 seconds. Retransmit interval: 5 seconds. Transmit delay: 1 second. Dead interval: 40 seconds. Authentication key: no key predefined. Message-digest key (MD5): no key predefined.

1. NSF = Nonstop forwarding
2. OSPF NSF awareness is enabled for IPv4 on switches running the IP services image.

Nonstop Forwarding Awareness

The OSPF NSF Awareness feature is supported for IPv4 in the IP services image. When the neighboring router is NSF-capable, the Layer 3 switch continues to forward packets from the neighboring router during the interval between the primary Route Processor (RP) in a router crashing and the backup RP taking over, or while the primary RP is manually reloaded for a non-disruptive software upgrade.

This feature cannot be disabled. For more information about this feature, see the [“Configuring Nonstop Forwarding”](#) chapter in the *High Availability Configuration Guide, Cisco IOS Release 15S*.

Configuring Basic OSPF Parameters

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.

BEFORE YOU BEGIN

Complete the OSPF network strategy and planning for your network. For example, you must decide whether multiple areas are required.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf process-id	Enable OSPF routing, and enter router configuration mode. The process ID is an internally used identification parameter that is locally assigned and can be any positive integer. Each OSPF routing process has a unique value.

	Command	Purpose
Step 3	network <i>address wildcard-mask area area-id</i>	Define an interface on which OSPF runs and the area ID for that interface. You can use the wildcard-mask to use a single command to define one or more multiple interfaces to be associated with a specific OSPF area. The area ID can be a decimal value or an IP address.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip protocols	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To terminate an OSPF routing process, use the **no router ospf process-id** global configuration command.

EXAMPLE

This example shows how to configure an OSPF routing process and assign it a process number of 109:

```
Switch(config)# router ospf 109
Switch(config-router)# network 131.108.0.0 255.255.255.0 area 24
```

Configuring OSPF Interfaces

You can use the **ip ospf** interface configuration commands to modify interface-specific OSPF parameters. You are not required to modify any of these parameters, but some interface parameters (hello interval, dead interval, and authentication key) must be consistent across all routers in an attached network.



Note

The **ip ospf** interface configuration commands are all optional.

BEFORE YOU BEGIN

If you modify these parameters, be sure all routers in the network have compatible values.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip ospf cost	(Optional) Explicitly specify the cost of sending a packet on the interface.
Step 5	ip ospf retransmit-interval <i>seconds</i>	(Optional) Specify the number of seconds between link state advertisement transmissions. The range is 1 to 65535 seconds. The default is 5 seconds.

	Command	Purpose
Step 6	ip ospf transmit-delay <i>seconds</i>	(Optional) Set the estimated number of seconds to wait before sending a link state update packet. The range is 1 to 65535 seconds. The default is 1 second.
Step 7	ip ospf priority <i>number</i>	(Optional) Set priority to help find the OSPF designated router for a network. The range is from 0 to 255. The default is 1.
Step 8	ip ospf hello-interval <i>seconds</i>	(Optional) Set the number of seconds between hello packets sent on an OSPF interface. The value must be the same for all nodes on a network. The range is 1 to 65535 seconds. The default is 10 seconds.
Step 9	ip ospf dead-interval <i>seconds</i>	(Optional) Set the number of seconds after the last device hello packet was seen before its neighbors declare the OSPF router to be down. The value must be the same for all nodes on a network. The range is 1 to 65535 seconds. The default is 4 times the hello interval.
Step 10	ip ospf authentication-key <i>key</i>	(Optional) Assign a password to be used by neighboring OSPF routers. The password can be any string of keyboard-entered characters up to 8 bytes in length. All neighboring routers on the same network must have the same password to exchange OSPF information.
Step 11	ip ospf message-digest-key <i>keyid md5 key</i>	(Optional) Enable MDS authentication. <ul style="list-style-type: none"> <i>keyid</i>—An identifier from 1 to 255. <i>key</i>—An alphanumeric password of up to 16 bytes.
Step 12	ip ospf database-filter all out	(Optional) Block flooding of OSPF LSA packets to the interface. By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives.
Step 13	end	Return to privileged EXEC mode.
Step 14	show ip ospf interface [<i>interface-name</i>]	Display OSPF-related interface information.
Step 15	show ip ospf neighbor detail	Display NSF awareness status of neighbor switch. The output matches one of these examples: <ul style="list-style-type: none"> <i>Options is 0x52</i> <i>LLS Options is 0x1 (LR)</i> When both of these lines appear, the neighbor switch is NSF aware. <i>Options is 0x42</i>—This means the neighbor switch is not NSF aware.
Step 16	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of these commands to remove the configured parameter value or return to the default value.

EXAMPLE

The following example specifies a cost of 65 and sets the interval between link-state advertisement (LSA) retransmissions to 1 second:

```
Switch# configure terminal
Switch(config)# interface GigabitEthernet 0/0
Switch(config-if)# ip ospf cost 65
Switch(config-if)# ip ospf retransmit-interval 1
Switch(config-if)# end
```

Configuring OSPF Network Types

OSPF classifies different media into the 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 also configure network interfaces as either a broadcast or an NBMA network and as point-to-point or point-to-multipoint, regardless of the default media type.

Configuring OSPF for Nonbroadcast Networks

Because many routers might be attached to an OSPF network, a designated router is selected for the network. If broadcast capability is not configured in the network, the designated router selection requires special configuration parameters. You need to configure these parameters only for devices that are eligible to become the designated router or backup designated router (in other words, routers with a nonzero router priority value).

BEFORE YOU BEGIN

Complete the OSPF network strategy and planning for your network.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Configure an OSPF routing process and enter router configuration mode.

	Command	Purpose
Step 3	neighbor <i>ip-address</i> [priority <i>number</i>] [poll-interval <i>seconds</i>]	Specify an OSPF neighbor with neighbor parameters as required. <ul style="list-style-type: none"> <i>ip-address</i>—Enter the interface IP address of the OSPF neighbor. (Optional) priority <i>number</i>—Specify the router priority value of the nonbroadcast neighbor associated with the IP address. The range is 0 to 255; the default is 0. (Optional) poll-interval <i>seconds</i>—Specify a number that represents the poll interval time (in seconds). This value should be much larger than the hello interval. The range is 0-4294967295; the default is 120 seconds (2 minutes).
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip ospf [<i>process-id</i>]	Display OSPF-related information.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

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

EXAMPLE

The following example declares a router at address 192.168.3.4 on a nonbroadcast network, with a priority of 1 and a poll interval of 180 seconds:

```
Switch# configure terminal
Switch(config)# router ospf
Switch(config-router)# neighbor 192.168.3.4 priority 1 poll-interval 180
Switch(config-router)# end
```

Configuring Network Types for OSPF Interfaces

You can configure network interfaces as either broadcast or NBMA and as point-to-point or point-to-multipoint, regardless of the default media type.

An OSPF point-to-multipoint interface is defined as a numbered point-to-point interface with one or more neighbors. On point-to-multipoint broadcast networks, specifying neighbors is optional. When you configure an interface as point-to-multipoint when the media does not support broadcast, you should use the **neighbor** command to identify neighbors.

BEFORE YOU BEGIN

Complete the OSPF network strategy and planning for your network.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip ospf network { broadcast non-broadcast { point-to-multipoint [non-broadcast] point-to-point } }	<p>Configure the OSPF network type for the specified interface. Select one of these network types:</p> <ul style="list-style-type: none"> • broadcast—Specify an OSPF broadcast multi-access network. • non-broadcast—Specify an OSPF NBMA network. • point-to-multipoint—Specify an OSPF point-to-multipoint network. If you do not enter another keyword, the interface is point-to-multipoint for broadcast media. • point-to-multipoint non-broadcast—Specify an OSPF nonbroadcast point-to-multipoint network. • point-to-point—Specify an OSPF point-to-point network.
Step 5	exit	Return to global configuration mode.
Step 6	router ospf <i>process-id</i>	(Optional for point-to-multipoint; required for point-to-multipoint nonbroadcast) Configure an OSPF routing process and enter router configuration mode.
Step 7	neighbor <i>ip-address</i> cost <i>number</i>	<p>(Optional for point-to-multipoint; required for point-to-multipoint nonbroadcast). Specify a configured OSPF neighbor and assign a cost to the neighbor.</p> <ul style="list-style-type: none"> • <i>ip-address</i>—Enter the interface IP address of the OSPF neighbor. • cost number—Specify a cost for the neighbor as an integer from 1 to 65535. <p>Note On point-to-multipoint broadcast networks, specifying a neighbor is optional, but if you do specify a neighbor, you must specify a cost for that neighbor.</p> <p>On point-to-multipoint nonbroadcast neighbors, you must specify a neighbor, but assigning a cost to the neighbor is optional. If not specified, neighbors assume the cost of the interface, based on the ip ospf cost interface configuration command.</p>
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip ospf interface [<i>interface-id</i>]	Display OSPF-related interface information.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of the **ip ospf network** command to return to the default network type for the media.

EXAMPLE

The following example sets your OSPF network as a broadcast network:

```
interface serial 0
 ip address 192.168.77.17 255.255.255.0
 ip ospf network broadcast
 encapsulation frame-relay
```

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

```
interface serial 0
 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
```

Configuring OSPF Area Parameters

You can optionally configure several OSPF area parameters. These parameters include authentication for password-based protection against unauthorized access to an area, stub areas, and not-so-stubby-areas (NSSAs). *Stub areas* are areas into which information on external routes is not sent. Instead, the area border router (ABR) generates a default external route into the stub area for destinations outside the autonomous system (AS). An NSSA does not flood all LSAs from the core into the area, but can import AS external routes within the area by redistribution.

Route summarization is the consolidation of advertised addresses into a single summary route to be advertised by other areas. If network numbers are contiguous, you can use the **area range** router configuration command to configure the ABR to advertise a summary route that covers all networks in the range.



Note

The OSPF **area** router configuration commands are all optional.

BEFORE YOU BEGIN

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.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode.
Step 3	area <i>area-id</i> authentication	(Optional) Allow password-based protection against unauthorized access to the identified area. The identifier can be either a decimal value or an IP address.
Step 4	area <i>area-id</i> authentication message-digest	(Optional) Enable MD5 authentication on the area.
Step 5	area <i>area-id</i> stub [no-summary]	(Optional) Define an area as a stub area. The no-summary keyword prevents an ABR from sending summary link advertisements into the stub area.
Step 6	area <i>area-id</i> nssa [no-redistribution] [default-information-originate] [no-summary]	(Optional) Defines an area as a not-so-stubby-area. Every router within the same area must agree that the area is NSSA. Select one of these keywords: <ul style="list-style-type: none"> • no-redistribution—Select when the router is an NSSA ABR and you want the redistribute command to import routes into normal areas, but not into the NSSA. • default-information-originate—Select on an ABR to allow importing type 7 LSAs into the NSSA. • no-redistribution—Select to not send summary LSAs into the NSSA.
Step 7	area <i>area-id</i> range <i>address mask</i>	(Optional) Specify an address range for which a single route is advertised. Use this command only with area border routers.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip ospf [<i>process-id</i>] show ip ospf [<i>process-id</i> [<i>area-id</i>]] database	Display information about the OSPF routing process in general or for a specific process ID to verify configuration. Display lists of information related to the OSPF database for a specific router.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of these commands to remove the configured parameter value or to return to the default value.

EXAMPLE

The following example mandates authentication for areas 0 and 10.0.0.0 of OSPF routing process 201. Authentication keys are also provided.

```
interface ethernet 0
 ip address 192.168.251.201 255.255.255.0
 ip ospf authentication-key adcdefgh
!
interface ethernet 1
 ip address 10.56.0.201 255.255.0.0
 ip ospf authentication-key ijklmnop
```

```

!
router ospf 201
 network 10.0.0.0 0.255.255.255 area 10.0.0.0
 network 192.168.0.0 0.0.255.255 area 0
 area 10.0.0.0 authentication
 area 0 authentication

```

Configuring Other OSPF Parameters

You can optionally configure other OSPF parameters in router configuration mode.

- **Route summarization:** When redistributing routes from other protocols as described in the [“Using Route Maps to Redistribute Routing Information”](#) section on page 2-121, each route is advertised individually in an external LSA. To help decrease the size of the OSPF link state database, you can use the **summary-address** router configuration command to advertise a single router for all the redistributed routes included in a specified network address and mask.
- **Virtual links:** In OSPF, all areas must be connected to a backbone area. You can establish a virtual link in case of a backbone-continuity break by configuring two Area Border Routers as endpoints of a virtual link. Configuration information includes the identity of the other virtual endpoint (the other ABR) and the nonbackbone link that the two routers have in common (the transit area). Virtual links cannot be configured through a stub area.
- **Default route:** When you specifically configure redistribution of routes into an OSPF routing domain, the route automatically becomes an autonomous system boundary router (ASBR). You can force the ASBR to generate a default route into the OSPF routing domain.
- **Domain Name Server (DNS) names** for use in all OSPF **show** privileged EXEC command displays makes it easier to identify a router than displaying it by router ID or neighbor ID.
- **Default Metrics:** OSPF calculates the OSPF metric for an interface according to the bandwidth of the interface. The metric is calculated as *ref-bw* divided by bandwidth, where *ref* is 10 by default, and bandwidth (*bw*) is specified by the **bandwidth** interface configuration command. For multiple links with high bandwidth, you can specify a larger number to differentiate the cost on those links.
- **Administrative distance** is a rating of the trustworthiness of a routing information source, an integer between 0 and 255, with a higher value meaning a lower 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: routes within an area (interarea), routes to another area (interarea), and routes from another routing domain learned through redistribution (external). You can change any of the distance values.
- **Passive interfaces:** Because interfaces between two devices on an Ethernet represent only one network segment, to prevent OSPF from sending hello packets for the sending interface, you must configure the sending device to be a passive interface. Both devices can identify each other through the hello packet for the receiving interface.
- **Route calculation timers:** You can configure the delay time between when OSPF receives a topology change and when it starts the shortest path first (SPF) calculation and the hold time between two SPF calculations.
- **Log neighbor changes:** You can configure the router to send a syslog message when an OSPF neighbor state changes, providing a high-level view of changes in the router.

BEFORE YOU BEGIN

Complete the OSPF network strategy and planning for your network.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode.
Step 3	summary-address <i>address mask</i>	(Optional) Specify an address and IP subnet mask for redistributed routes so that only one summary route is advertised.
Step 4	area <i>area-id</i> virtual-link <i>router-id</i> [hello-interval <i>seconds</i>] [retransmit-interval <i>seconds</i>] [trans] [[authentication-key <i>key</i>] message-digest-key <i>keyid md5 key</i>]]	(Optional) Establish a virtual link and set its parameters. See the “ Configuring OSPF Interfaces ” section on page 2-31 for parameter definitions and the “ Default OSPF Configuration ” section on page 2-29 for virtual link defaults.
Step 5	default-information originate [always] [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [route-map <i>map-name</i>]	(Optional) Force the ASBR to generate a default route into the OSPF routing domain. Parameters are all optional.
Step 6	ip ospf name-lookup	(Optional) Configure DNS name lookup. The default is disabled.
Step 7	ip auto-cost reference-bandwidth <i>ref-bw</i>	(Optional) Specify an address range for which a single route will be advertised. Use this command only with area border routers.
Step 8	distance ospf {[inter-area <i>dist1</i>] [inter-area <i>dist2</i>] [external <i>dist3</i>]}	(Optional) Change the OSPF distance values. The default distance for each type of route is 110. The range is 1 to 255.
Step 9	passive-interface <i>type number</i>	(Optional) Suppress the sending of hello packets through the specified interface.
Step 10	timers throttle spf <i>spf-delay</i> <i>spf-holdtime</i> <i>spf-wait</i>	(Optional) Configure route calculation timers. <ul style="list-style-type: none"> <i>spf-delay</i>—Delay between receiving a change to SPF calculation. The range is from 1 to 600000 milliseconds. <i>spf-holdtime</i>—Delay between first and second SPF calculation. The range is from 1 to 600000 in milliseconds. <i>spf-wait</i>—Maximum wait time in milliseconds for SPF calculations. The range is from 1 to 600000 in milliseconds.
Step 11	ospf log-adj-changes	(Optional) Send syslog message when a neighbor state changes.
Step 12	end	Return to privileged EXEC mode.
Step 13	show ip ospf [<i>process-id</i> [<i>area-id</i>]] database	Display lists of information related to the OSPF database for a specific router. For some of the keyword options, see the “ Monitoring OSPF ” section on page 2-42.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

In the following example, the summary address 10.1.0.0 includes address 10.1.1.0, 10.1.2.0, 10.1.3.0, and so on. Only the address 10.1.0.0 is advertised in an external link-state advertisement.

```
Switch(config)# router ospf 201
Switch(config-router)# summary-address 10.1.0.0 255.255.0.0
Switch(config-router)# end
```

Changing LSA Group Pacing

The OSPF LSA group pacing feature allows the router to group OSPF LSAs and pace the refreshing, check-summing, and aging functions for more efficient router use. This feature is enabled by default with a 4-minute default pacing interval, and you will not usually need to modify this parameter. The optimum group pacing interval is inversely proportional to the number of LSAs the router is refreshing, check-summing, and aging. For example, if you have approximately 10,000 LSAs in the database, 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.

BEFORE YOU BEGIN

Do not change the packet pacing timers unless all other options to meet OSPF packet flooding requirements have been exhausted. Specifically, network operators should prefer summarization, stub area usage, queue tuning, and buffer tuning before changing the default flooding timers. Furthermore, there are no guidelines for changing timer values; each OSPF deployment is unique and should be considered on a case-by-case basis. The network operator assumes the risks associated with changing the default timer values.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Enable OSPF routing, and enter router configuration mode.
Step 3	timers pacing lsa-group <i>seconds</i>	Change the group pacing of LSAs.
Step 4	end	Return to privileged EXEC mode.
Step 5	show running-config	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default value, use the **no timers pacing lsa-group** router configuration command.

EXAMPLE

The following example configures OSPF group packet-pacing updates between LSA groups to occur in 60-second intervals for OSPF routing process 1:

```
Switch(config)# router ospf 1
Switch(config-router)# timers pacing lsa-group 60
```

Configuring a Loopback Interface

OSPF uses the highest IP address configured on the interfaces as its router ID. If this interface is down or 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, OSPF uses this IP address as its router ID, even if other interfaces have higher IP addresses. Because loopback interfaces never fail, this provides greater stability. OSPF automatically prefers a loopback interface over other interfaces, and it chooses the highest IP address among all loopback interfaces.

BEFORE YOU BEGIN

The IP address for the loopback interface must be unique and not in use by another interface.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface loopback 0	Create a loopback interface, and enter interface configuration mode.
Step 3	ip address <i>address mask</i>	Assign an IP address to this interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip interface	Verify your entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no interface loopback 0** global configuration command to disable the loopback interface.

EXAMPLE

```
Switch(config)# interface loopback 0  
Switch(config-if)# ip address 10.108.1.1 255.255.255.0
```

Monitoring OSPF

You can display specific statistics such as the contents of IP routing tables, caches, and databases.

Following are some of the privileged EXEC commands for displaying OSPF statistics. For more **show ip ospf database** privileged EXEC command options and for explanations of fields in the resulting display, see the [Cisco IOS IP Routing: OSPF Command Reference](#).

Command	Purpose
show ip ospf [<i>process-id</i>]	Display general information about OSPF routing processes.
show ip ospf [<i>process-id</i>] database [router] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [router] [self-originate] show ip ospf [<i>process-id</i>] database [router] [adv-router] [<i>ip-address</i>] show ip ospf [<i>process-id</i>] database [network] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [summary] [<i>link-state-id</i>] show ip ospf [<i>process-id</i>] database [asbr-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 lists of information related to the OSPF database.
show ip ospf border-routes	Display the internal OSPF routing ABR and ASBR table entries.
show ip ospf interface [<i>interface-name</i>]	Display OSPF-related interface information.
show ip ospf neighbor [<i>interface-name</i>] [<i>neighbor-id</i>] detail	Display OSPF interface neighbor information.
show ip ospf virtual-links	Display OSPF-related virtual links information.

Configuring EIGRP

Enhanced IGRP (EIGRP) is a Cisco proprietary enhanced version of the Interior Gateway Routing Protocol (IGRP). EIGRP uses the same distance vector algorithm and distance information as IGRP; however, the convergence properties and the operating efficiency of EIGRP are significantly improved.

The convergence technology employs an algorithm referred to as the Diffusing Update Algorithm (DUAL), which guarantees loop-free operation at every instant throughout a route computation and allows all devices involved in a topology change to synchronize at the same time. Routers that are not affected by topology changes are not involved in recomputations.

IP EIGRP provides increased network width. With RIP, the largest possible width of your network is 15 hops. Because the EIGRP metric is large enough to support thousands of hops, the only barrier to expanding the network is the transport-layer hop counter. EIGRP increments the transport control field only when an IP packet has traversed 15 routers and the next hop to the destination was learned through EIGRP.

EIGRP has these four basic components:

- *Neighbor discovery and recovery* is the process that routers use to dynamically learn of other routers on their directly attached networks. Routers must also discover when their neighbors become unreachable or inoperative. Neighbor discovery and recovery is achieved by periodically sending small hello packets. As long as hello packets are received, the neighbor is alive and functioning. When this status is determined, the neighboring routers exchange routing information.
- The *reliable transport protocol* is responsible for guaranteed, ordered delivery of EIGRP packets to all neighbors. It supports intermixed transmission of multicast and unicast packets. Some EIGRP packets must be sent reliably, and others need not be. For efficiency, reliability is provided only when necessary. For example, on a multiaccess network that has multicast capabilities, it is not necessary to send hellos reliably to all neighbors individually. Therefore, EIGRP sends a single multicast hello with an indication in the packet informing the receivers that the packet need not be acknowledged. Other types of packets (such as updates) require acknowledgment, which is shown in the packet. To ensure low convergence time, the reliable transport sends multicast packets quickly when there are unacknowledged packets pending.
- The *DUAL finite state machine* handles the decision process for all route computations. It tracks all routes advertised by all neighbors and uses the distance information (known as a metric) to select efficient, loop-free paths. DUAL selects routes to be inserted into a routing table based on feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination that is guaranteed not to be part of a routing loop.

When there are no feasible successors, but there are neighbors advertising the destination, a recomputation must occur to determine a new successor. The amount of time it takes to recompute the route affects the convergence time. When a topology change occurs, DUAL tests for feasible successors to avoid unnecessary recomputation.

- The *protocol-dependent modules* are responsible for network layer protocol-specific tasks. An example is the IP EIGRP module, which is responsible for sending and receiving EIGRP packets that are encapsulated in IP. It is also responsible for parsing EIGRP packets and informing DUAL of the new information received. Routing decisions are stored in the IP routing table. EIGRP also redistributes routes learned by other IP routing protocols.

This section includes the following topics:

- [Default EIGRP Configuration, page 2-43](#)
- [Configuring Basic EIGRP Parameters, page 2-45](#)
- [Configuring EIGRP Interfaces, page 2-46](#)
- [Configuring EIGRP Route Authentication, page 2-47](#)
- [Configuring EIGRP Stub Routing, page 2-49](#)
- [Monitoring and Maintaining EIGRP, page 2-50](#)

Default EIGRP Configuration

Feature	Default Setting
Auto summary	Enabled. Subprefixes are summarized to the classful network boundary when crossing classful network boundaries.
Default-information	Exterior routes are accepted and default information is passed between EIGRP processes when doing redistribution.

Feature	Default Setting
Default metric	Only connected routes and interface static routes can be redistributed without a default metric. The metric includes: <ul style="list-style-type: none"> Bandwidth: 0 or greater kbps. Delay (tens of microseconds): 0 or any positive number that is a multiple of 39.1 nanoseconds. Reliability: any number between 0 and 255 (255 means 100 percent reliability). Loading: effective bandwidth as a number between 0 and 255 (255 is 100 percent loading). MTU: maximum transmission unit size of the route in bytes. 0 or any positive integer.
Distance	Internal distance: 90. External distance: 170.
EIGRP log-neighbor changes	Disabled. No adjacency changes logged.
IP authentication key-chain	No authentication provided.
IP authentication mode	No authentication provided.
IP bandwidth-percent	50 percent.
IP hello interval	For low-speed nonbroadcast multiaccess (NBMA) networks: 60 seconds; all other networks: 5 seconds.
IP hold-time	For low-speed NBMA networks: 180 seconds; all other networks: 15 seconds.
IP split-horizon	Enabled.
IP summary address	No summary aggregate addresses are predefined.
Metric weights	tos: 0; k1 and k3: 1; k2, k4, and k5: 0.
Network	None specified.
NSF ¹ Awareness	Enabled ² . Allows Layer 3 switches to continue forwarding packets from a neighboring NSF-capable router during hardware or software changes.
Offset-list	Disabled.
Router EIGRP	Disabled.
Set metric	No metric set in the route map.
Traffic-share	Distributed proportionately to the ratios of the metrics.
Variance	1 (equal-cost load balancing).

1. NSF = Nonstop Forwarding

2. EIGRP NSF awareness is enabled for IPv4 on switches running the IP services image.

To create an EIGRP routing process, you must enable EIGRP and associate networks. EIGRP sends updates to the interfaces in the specified networks. If you do not specify an interface network, it is not advertised in any EIGRP update.

Nonstop Forwarding Awareness

The EIGRP NSF Awareness feature is supported for IPv4 in the IP services image. When the neighboring router is NSF-capable, the Layer 3 switch continues to forward packets from the neighboring router during the interval between the primary Route Processor (RP) in a router failing and the backup RP taking over, or while the primary RP is manually reloaded for a nondisruptive software upgrade.

This feature cannot be disabled. For more information on this feature, see the “[Configuring Nonstop Forwarding](#)” chapter in the *High Availability Configuration Guide, Cisco IOS Release 15S*.


Configuring Basic EIGRP Parameters

In this procedure, configuring the routing process is required; other steps are optional.

BEFORE YOU BEGIN

Complete the EIGRP network strategy and planning for your network.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router eigrp <i>autonomous-system</i>	Enable an EIGRP routing process, and enter router configuration mode. The AS number identifies the routes to other EIGRP routers and is used to tag routing information.
Step 3	network <i>network-number</i>	Associate networks with an EIGRP routing process. EIGRP sends updates to the interfaces in the specified networks.
Step 4	eigrp log-neighbor-changes	(Optional) Enable logging of EIGRP neighbor changes to monitor routing system stability.
Step 5	metric weights <i>tos k1 k2 k3 k4 k5</i>	(Optional) Adjust the EIGRP metric. Although the defaults have been carefully set to provide excellent operation in most networks, you can adjust them.
		 Caution Setting metrics is complex and is not recommended without guidance from an experienced network designer.
Step 6	offset list [<i>access-list number</i> <i>name</i>] { in out } <i>offset</i> [<i>type number</i>]	(Optional) Apply an offset list to routing metrics to increase incoming and outgoing metrics to routes learned through EIGRP. You can limit the offset list with an access list or an interface.
Step 7	no auto-summary	(Optional) Disable automatic summarization of subnet routes into network-level routes.
Step 8	ip summary-address eigrp <i>autonomous-system-number address mask</i>	(Optional) Configure a summary aggregate.

	Command	Purpose
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip protocols	Verify your entries. For NSF awareness, the output shows: *** IP Routing is NSF aware *** EIGRP NSF enabled
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or return the setting to the default value.

EXAMPLE

The following example configures EIGRP autonomous system 1 and establishes neighbors through networks 172.16.0.0 and 192.168.0.0:

```
Switch(config)# router eigrp 1
Switch(config-router)# network 172.16.0.0
Switch(config-router)# network 192.168.0.0
```

Configuring EIGRP Interfaces


Other optional EIGRP parameters can be configured on an interface basis.

BEFORE YOU BEGIN

Enable EIGRP as described in the [“Configuring Basic EIGRP Parameters”](#) section on page 2-45.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip bandwidth-percent eigrp <i>percent</i>	(Optional) Configure the percentage of bandwidth that can be used by EIGRP on an interface. The default is 50 percent.
Step 5	ip summary-address eigrp <i>autonomous-system-number address mask</i>	(Optional) Configure a summary aggregate address for a specified interface (not usually necessary if auto-summary is enabled).
Step 6	ip hello-interval eigrp <i>autonomous-system-number seconds</i>	(Optional) Change the hello time interval for an EIGRP routing process. The range is 1 to 65535 seconds. The default is 60 seconds for low-speed NBMA networks and 5 seconds for all other networks.

	Command	Purpose
Step 7	ip hold-time eigrp <i>autonomous-system-number seconds</i>	(Optional) Change the hold time interval for an EIGRP routing process. The range is 1 to 65535 seconds. The default is 180 seconds for low-speed NBMA networks and 15 seconds for all other networks.
		 Caution Do not adjust the hold time without consulting Cisco technical support.
Step 8	no ip split-horizon eigrp <i>autonomous-system-number</i>	(Optional) Disable split horizon to allow route information to be advertised by a router out any interface from which that information originated.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip eigrp interface	Display which interfaces EIGRP is active on and information about EIGRP relating to those interfaces.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or return the setting to the default value.

EXAMPLE

The following example allows EIGRP to use up to 75 percent (42 kbps) of a 56-kbps serial link in autonomous system 209:

```
Switch(config)# interface serial 0
Switch(config-if)# bandwidth 56
Switch(config-if)# ip bandwidth-percent eigrp 209 75
```

Configuring EIGRP Route Authentication

EIGRP route authentication provides MD5 authentication of routing updates from the EIGRP routing protocol to prevent the introduction of unauthorized or false routing messages from unapproved sources.

BEFORE YOU BEGIN

Enable EIGRP as described in the [“Configuring Basic EIGRP Parameters”](#) section on page 2-45.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.

	Command	Purpose
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	ip authentication mode eigrp <i>autonomous-system md5</i>	Enable MD5 authentication in IP EIGRP packets.
Step 5	ip authentication key-chain eigrp <i>autonomous-system key-chain</i>	Enable authentication of IP EIGRP packets.
Step 6	exit	Return to global configuration mode.
Step 7	key chain <i>name-of-chain</i>	Identify a key chain and enter key-chain configuration mode. Match the name configured in Step 4.
Step 8	key number	In key-chain configuration mode, identify the key number.
Step 9	key-string <i>text</i>	In key-chain key configuration mode, identify the key string.
Step 10	accept-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be received. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite .
Step 11	send-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be sent. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite .
Step 12	end	Return to privileged EXEC mode.
Step 13	show key chain	Display authentication key information.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** forms of these commands to disable the feature or to return the setting to the default value.

EXAMPLE

The following example configures EIGRP to apply authentication to address-family autonomous system 1 and identifies a key chain named SITE1:

```
Switch(config)# router eigrp virtual-name
Switch(config-router)# address-family ipv4 autonomous-system 1
Switch(config-router-af)# af-interface ethernet0/0
Switch(config-router-af-interface)# authentication key-chain SITE1
```

```
Switch(config-router-af-interface)# authentication mode md5
```

Configuring EIGRP Stub Routing

The EIGRP stub routing feature reduces resource utilization by moving routed traffic closer to the end user. In a network using EIGRP stub routing, the only allowable route for IP traffic to the user is through a switch that is configured with EIGRP stub routing. The switch sends the routed traffic to interfaces that are configured as user interfaces or are connected to other devices.

When using EIGRP stub routing, you need to configure the distribution and remote routers to use EIGRP and to configure only the switch as a stub. Only specified routes are propagated from the switch. The switch responds to all queries for summaries, connected routes, and routing updates.



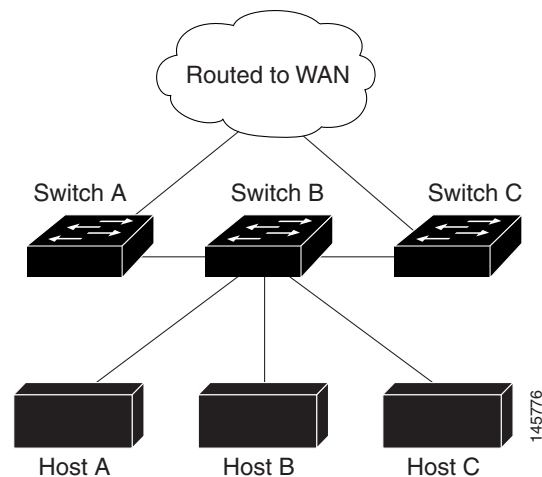
Note

EIGRP stub routing only advertises connected or summary routes from the routing tables to other switches in the network. The switch uses EIGRP stub routing at the access layer to eliminate the need for other types of routing advertisements. If you try to configure multi-VRF-CE and EIGRP stub routing at the same time, the configuration is not allowed.

Any neighbor that receives a packet informing it of the stub status does not query the stub router for any routes, and a router that has a stub peer does not query that peer. The stub router depends on the distribution router to send the proper updates to all peers.

In [Figure 2-4](#), switch B is configured as an EIGRP stub router. Switches A and C are connected to the rest of the WAN. Switch B advertises connected, static, redistribution, and summary routes to switch A and C. Switch B does not advertise any routes learned from switch A (and the reverse).

Figure 2-4 EIGRP Stub Router Configuration



For more information about EIGRP stub routing, see the [IP Routing: EIGRP Configuration Guide, Cisco IOS Release 15M&T](#).

BEFORE YOU BEGIN

Complete the EIGRP network strategy and planning for your network.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router eigrp 1	Configure a remote or distribution router to run an EIGRP process and enter router configuration mode.
Step 3	network <i>network-number</i>	Associate networks with an EIGRP routing process.
Step 4	eigrp stub [receive-only connected static summary]	Configure a remote router as an EIGRP stub router. The keywords have these meanings: <ul style="list-style-type: none"> Enter receive-only to set the router as a receive-only neighbor. Enter connected to advertise connected routes. Enter static to advertise static routes. Enter summary to advertise summary routes.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip eigrp neighbor detail	Verify that a remote router has been configured as a stub router with EIGRP. The last line of the output shows the stub status of the remote or spoke router.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Enter the **show ip eigrp neighbor detail** privileged EXEC command from the distribution router to verify the configuration.

EXAMPLE

In the following example, the **eigrp stub** command is used to configure the router as a stub that advertises connected and summary routes:

```
Switch(config)# router eigrp 1
Switch(config-router)# network 10.0.0.0
Switch(config-router)# eigrp stub
```

Monitoring and Maintaining EIGRP

You can delete neighbors from the neighbor table. You can also display various EIGRP routing statistics.

Command	Purpose
clear ip eigrp neighbors [<i>if-address</i> <i>interface</i>]	Delete neighbors from the neighbor table.
show ip eigrp interface [<i>interface</i>] [<i>as number</i>]	Display information about interfaces configured for EIGRP.
show ip eigrp neighbors [<i>type-number</i>]	Display EIGRP discovered neighbors.

Command	Purpose
show ip eigrp topology [<i>autonomous-system-number</i>] [[<i>ip-address</i>] <i>mask</i>]]	Display the EIGRP topology table for a given process.
show ip eigrp traffic [<i>autonomous-system-number</i>]	Display the number of packets sent and received for all or a specified EIGRP process.

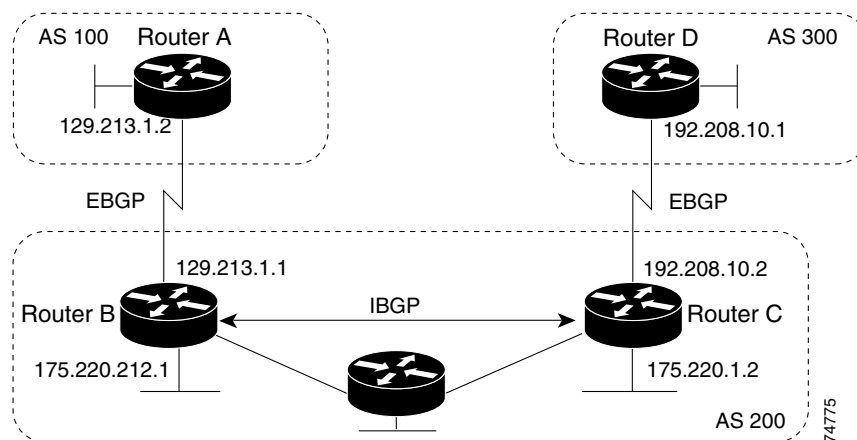
Configuring BGP

The Border Gateway Protocol (BGP) is an exterior gateway protocol used to set up an interdomain routing system for loop-free exchanges of routing information between autonomous systems. Autonomous systems are made up of routers that operate under the same administration and that run Interior Gateway Protocols (IGPs), such as RIP or OSPF, within their boundaries and that interconnect by using an Exterior Gateway Protocol (EGP). BGP Version 4 is the standard EGP for interdomain routing in the Internet.

For details about BGP configuration and commands, see the BGP documents listed in the [“Related Documents”](#) section on page 2-135.

Routers that belong to the same autonomous system (AS) and that exchange BGP updates run *internal BGP* (IBGP), and routers that belong to different autonomous systems and that exchange BGP updates run *external BGP* (EBGP). Most configuration commands are the same for configuring EBGP and IBGP. The difference is that the routing updates are exchanged either between autonomous systems (EBGP) or within an AS (IBGP). Figure 2-5 shows a network that is running both EBGP and IBGP.

Figure 2-5 EBGP, IBGP, and Multiple Autonomous Systems



Before exchanging information with an external AS, BGP ensures that networks within the AS can be reached by defining internal BGP peering among routers within the AS and by redistributing BGP routing information to IGPs that run within the AS, such as IGRP and OSPF.

Routers that run a BGP routing process are often referred to as BGP *speakers*. BGP uses the Transmission Control Protocol (TCP) as its transport protocol (specifically port 179). Two BGP speakers that have a TCP connection to each other for exchanging routing information are known as *peers* or *neighbors*. In Figure 2-5, Routers A and B are BGP peers, as are Routers B and C and Routers C and D. The routing information is a series of AS numbers that describe the full path to the destination network. BGP uses this information to construct a loop-free map of autonomous systems.

The network has these characteristics:

- Routers A and B are running EBGp, and Routers B and C are running IBGP. Note that the EBGp peers are directly connected and that the IBGP peers are not. As long as there is an IGP running that allows the two neighbors to reach one another, IBGP peers do not have to be directly connected.
- All BGP speakers within an AS must establish a peer relationship with each other. That is, the BGP speakers within an AS must be fully meshed logically. BGP4 provides two techniques that reduce the requirement for a logical full mesh: *confederations* and *route reflectors*.
- AS 200 is a *transit AS* for AS 100 and AS 300—that is, AS 200 is used to transfer packets between AS 100 and AS 300.

BGP peers initially exchange their full BGP routing tables and then send only incremental updates. BGP peers also exchange keepalive messages (to ensure that the connection is up) and notification messages (in response to errors or special conditions).

In BGP, each route consists of a network number, a list of autonomous systems that information has passed through (the *autonomous system path*), and a list of other *path attributes*. The primary function of a BGP system is to exchange network reachability information, including information about the list of AS paths, with other BGP systems. This information can be used to determine AS connectivity, to prune routing loops, and to enforce AS-level policy decisions.

A router or switch running Cisco IOS does not select or use an IBGP route unless it has a route available to the next-hop router and it has received synchronization from an IGP (unless IGP synchronization is disabled). When multiple routes are available, BGP bases its path selection on *attribute* values. See the “[Configuring BGP Decision Attributes](#)” section on page 2-59 for information about BGP attributes.

BGP Version 4 supports classless interdomain routing (CIDR) so you can reduce the size of your routing tables by creating aggregate routes, resulting in *supernets*. CIDR eliminates the concept of network classes within BGP and supports the advertising of IP prefixes.

This section includes the following topics:

- [Default BGP Configuration, page 2-53](#)
- [Enabling BGP Routing, page 2-55](#)
- [Managing Routing Policy Changes, page 2-58](#)
- [Configuring BGP Decision Attributes, page 2-59](#)
- [Configuring BGP Filtering with Route Maps, page 2-62](#)
- [Configuring BGP Filtering by Neighbor, page 2-63](#)
- [Configuring Prefix Lists for BGP Filtering, page 2-64](#)
- [Configuring BGP Community Filtering, page 2-66](#)
- [Configuring BGP Neighbors and Peer Groups, page 2-68](#)
- [Configuring Aggregate Addresses, page 2-70](#)
- [Configuring Routing Domain Confederations, page 2-72](#)
- [Configuring BGP Route Reflectors, page 2-73](#)
- [Configuring Route Dampening, page 2-74](#)
- [Monitoring and Maintaining BGP, page 2-75](#)

Default BGP Configuration

Feature	Default Setting
Aggregate address	Disabled: None defined.
AS path access list	None defined.
Auto summary	Enabled.
Best path	<ul style="list-style-type: none"> The router considers <i>as-path</i> in choosing a route and does not compare similar routes from external BGP peers. Compare router ID: Disabled.
BGP community list	<ul style="list-style-type: none"> Number: None defined. When you permit a value for the community number, the list defaults to an implicit deny for everything else that has not been permitted. Format: Cisco default format (32-bit number).
BGP confederation identifier/peers	<ul style="list-style-type: none"> Identifier: None configured. Peers: None identified.
BGP Fast external fallover	Enabled.
BGP local preference	100. The range is 0 to 4294967295 with the higher value preferred.
BGP network	None specified; no backdoor route advertised.
BGP route dampening	Disabled by default. When enabled: <ul style="list-style-type: none"> Half-life is 15 minutes. Re-use is 750 (10-second increments). Suppress is 2000 (10-second increments). Max-suppress-time is 4 times half-life; 60 minutes.
BGP router ID	The IP address of a loopback interface if one is configured or the highest IP address configured for a physical interface on the router.
Default information originate (protocol or network redistribution)	Disabled.
Default metric	Built-in, automatic metric translations.
Distance	<ul style="list-style-type: none"> External route administrative distance: 20 (acceptable values are from 1 to 255). Internal route administrative distance: 200 (acceptable values are from 1 to 255). Local route administrative distance: 200 (acceptable values are from 1 to 255).
Distribute list	<ul style="list-style-type: none"> In (filter networks received in updates): Disabled. Out (suppress networks from being advertised in updates): Disabled.
Internal route redistribution	Disabled.
IP prefix list	None defined.

Feature	Default Setting
Multi exit discriminator (MED)	<ul style="list-style-type: none"> Always compare: Disabled. Does not compare MEDs for paths from neighbors in different autonomous systems. Best path compare: Disabled. MED missing as worst path: Disabled. Deterministic MED comparison is disabled.
Neighbor	<ul style="list-style-type: none"> Advertisement interval: 30 seconds for external peers; 5 seconds for internal peers. Change logging: Enabled. Conditional advertisement: Disabled. Default originate: No default route is sent to the neighbor. Description: None. Distribute list: None defined. External BGP multihop: Only directly connected neighbors are allowed. Filter list: None used. Maximum number of prefixes received: No limit.
Neighbor	<ul style="list-style-type: none"> Next hop (router as next hop for BGP neighbor): Disabled. Password: Disabled. Peer group: None defined; no members assigned. Prefix list: None specified. Remote AS (add entry to neighbor BGP table): No peers defined. Private AS number removal: Disabled. Route maps: None applied to a peer. Send community attributes: None sent to neighbors. Shutdown or soft reconfiguration: Not enabled. Timers: keepalive: 60 seconds; holdtime: 180 seconds. Update source: Best local address. Version: BGP Version 4. Weight: Routes learned through BGP peer: 0; routes sourced by the local router: 32768.
NSF ¹ Awareness	Disabled ² . Allows Layer 3 switches to continue forwarding packets from a neighboring NSF-capable router during hardware or software changes.
Route reflector	None configured.
Synchronization (BGP and IGP)	Enabled.
Table map update	Disabled.
Timers	Keepalive: 60 seconds; holdtime: 180 seconds.

1. NSF = Nonstop Forwarding

2. BGP NSF Awareness can be enabled for IPv4 on switches with the IP services image by enabling Graceful Restart.

Nonstop Forwarding Awareness

The BGP NSF Awareness feature is supported for IPv4 in the IP services image. To enable this feature with BGP routing, you need to enable Graceful Restart. When the neighboring router is NSF-capable, and this feature is enabled, the Layer 3 switch continues to forward packets from the neighboring router during the interval between the primary Route Processor (RP) in a router failing and the backup RP taking over, or while the primary RP is manually reloaded for a nondisruptive software upgrade. For more information, see the *IP Routing: BGP Configuration Guide, Cisco IOS Release 15M&T*.

Enabling BGP Routing

To enable BGP routing, you establish a BGP routing process and define the local network. Because BGP must completely recognize the relationships with its neighbors, you must also specify a BGP neighbor.

BGP supports two kinds of neighbors: internal and external. *Internal neighbors* are in the same AS; *external neighbors* are in different autonomous systems. External neighbors are usually adjacent to each other and share a subnet, but internal neighbors can be anywhere in the same AS.

The switch supports the use of private AS numbers, usually assigned by service providers and given to systems whose routes are not advertised to external neighbors. The private AS numbers are from 64512 to 65535. You can configure external neighbors to remove private AS numbers from the AS path by using the **neighbor remove-private-as** router configuration command. Then when an update is passed to an external neighbor, if the AS path includes private AS numbers, these numbers are dropped.

If your AS must pass traffic through it from another AS to a third AS, it is important to be consistent about the routes it advertises. If BGP advertises a route before all routers in the network learn about the route through the IGP, the AS might receive traffic that some routers can not yet route. To prevent this from happening, BGP must wait until the IGP has propagated information across the AS so that BGP is *synchronized* with the IGP. Synchronization is enabled by default. If your AS does not pass traffic from one AS to another AS, or if all routers in your autonomous systems are running BGP, you can disable synchronization, which allows your network to carry fewer routes in the IGP and allows BGP to converge more quickly.

BEFORE YOU BEGIN

You should know your network design and how you want traffic to flow through it before configuring BGP. Gather the network requirements you need, which should include the following:

- Whether you need to run IBGP for internal connectivity
- External connectivity to the service provider network
- Configuration parameters such as neighbor IP addresses and their AS number, and which networks you will advertise through BGP

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing (required only if IP routing is disabled).
Step 3	router bgp <i>autonomous-system</i>	Enable a BGP routing process, assign it an AS number, and enter router configuration mode. The AS number can be from 1 to 65535, with 64512 to 65535 designated as private autonomous numbers.
Step 4	network <i>network-number</i> [mask <i>network-mask</i>] [route-map <i>route-map-name</i>]	Configure a network as local to this AS, and enter it in the BGP table.
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i>	Add an entry to the BGP neighbor table specifying that the neighbor identified by the IP address belongs to the specified AS. For EBGP, neighbors are usually directly connected, and the IP address is the address of the interface at the other end of the connection. For IBGP, the IP address can be the address of any of the router interfaces.
Step 6	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remove-private-as	(Optional) Remove private AS numbers from the AS-path in outbound routing updates.
Step 7	no synchronization	(Optional) Disable synchronization between BGP and an IGP.
Step 8	no auto-summary	(Optional) Disable automatic network summarization. By default, when a subnet is redistributed from an IGP into BGP, only the network route is inserted into the BGP table.
Step 9	bgp fast-external-fallover	(Optional) Automatically reset a BGP session when a link between external neighbors goes down. By default, the session is not immediately reset.
Step 10	bgp graceful-restart	(Optional) Enable NSF awareness on switch. By default, NSF awareness is disabled.
Step 11	end	Return to privileged EXEC mode.

	Command	Purpose
Step 12	show ip bgp network <i>network-number</i> or show ip bgp neighbor	Verify the configuration. Verify that NSF awareness (Graceful Restart) is enabled on the neighbor. If NSF awareness is enabled on the switch and the neighbor, this message appears: <i>Graceful Restart Capability: advertised and received</i> If NSF awareness is enabled on the switch, but not on the neighbor, this message appears: <i>Graceful Restart Capability: advertised</i>
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp** *autonomous-system* global configuration command to remove a BGP AS. Use the **no network** *network-number* router configuration command to remove the network from the BGP table. Use the **no neighbor** {*ip-address* | *peer-group-name*} **remote-as** *number* router configuration command to remove a neighbor. Use the **no neighbor** {*ip-address* | *peer-group-name*} **remove-private-as** router configuration command to include private AS numbers in updates to a neighbor. Use the **synchronization** router configuration command to re-enable synchronization.

EXAMPLE

These examples show how to configure BGP on the routers in [Figure 2-5](#).

Router A:

```
Switch(config)# router bgp 100
Switch(config-router)# neighbor 129.213.1.1 remote-as 200
```

Router B:

```
Switch(config)# router bgp 200
Switch(config-router)# neighbor 129.213.1.2 remote-as 100
Switch(config-router)# neighbor 175.220.1.2 remote-as 200
```

Router C:

```
Switch(config)# router bgp 200
Switch(config-router)# neighbor 175.220.212.1 remote-as 200
Switch(config-router)# neighbor 192.208.10.1 remote-as 300
```

Router D:

```
Switch(config)# router bgp 300
Switch(config-router)# neighbor 192.208.10.2 remote-as 200
```

To verify that BGP peers are running, use the **show ip bgp neighbors** privileged EXEC command. This is the output of this command on Router A:

```
Switch# show ip bgp neighbors

BGP neighbor is 129.213.1.1, remote AS 200, external link
BGP version 4, remote router ID 175.220.212.1
BGP state = established, table version = 3, up for 0:10:59
Last read 0:00:29, hold time is 180, keepalive interval is 60 seconds
Minimum time between advertisement runs is 30 seconds
Received 2828 messages, 0 notifications, 0 in queue
```

```
Sent 2826 messages, 0 notifications, 0 in queue
Connections established 11; dropped 10
```

Anything other than *state = established* means that the peers are not running. The remote router ID is the highest IP address on that router (or the highest loopback interface). Each time the table is updated with new information, the table version number increments. A table version number that continually increments means that a route is flapping, causing continual routing updates.

For exterior protocols, a reference to an IP network from the **network** router configuration command controls only which networks are advertised. This is in contrast to Interior Gateway Protocols (IGPs), such as EIGRP, which also use the **network** command to specify where to send updates.

Managing Routing Policy Changes

Routing policies for a peer include all the configurations that might affect inbound or outbound routing table updates. When you have defined two routers as BGP neighbors, they form a BGP connection and exchange routing information. If you later change a BGP filter, weight, distance, version, or timer, or make a similar configuration change, you must reset the BGP sessions so that the configuration changes take effect.

There are two types of reset: hard reset and soft reset. The switch supports a soft reset without any prior configuration when both BGP peers support the soft route refresh capability, which is advertised in the OPEN message sent when the peers establish a TCP session. A soft reset allows the dynamic exchange of route refresh requests and routing information between BGP routers and the subsequent re-advertisement of the respective outbound routing table.

- When soft reset generates inbound updates from a neighbor, it is called *dynamic inbound soft reset*.
- When soft reset sends a set of updates to a neighbor, it is called *outbound soft reset*.

A soft inbound reset causes the new inbound policy to take effect. A soft outbound reset causes the new local outbound policy to take effect without resetting the BGP session. As a new set of updates is sent during outbound policy reset, a new inbound policy can also take effect.

Table 2-1 Advantages and Disadvantages of Hard and Soft Resets

Type of Reset	Advantages	Disadvantages
Hard reset	No memory overhead.	The prefixes in the BGP, IP, and FIB tables provided by the neighbor are lost. Not recommended.
Outbound soft reset	No configuration; no storing of routing table updates.	Does not reset inbound routing table updates.
Dynamic inbound soft reset	Does not clear the BGP session and cache. Does not require storing of routing table updates and has no memory overhead.	Both BGP routers must support the route refresh capability.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	show ip bgp neighbors	Display whether a neighbor supports the route refresh capability. When supported, this message appears for the router: <i>Received route refresh capability from peer.</i>
Step 2	clear ip bgp { * <i>address</i> <i>peer-group-name</i> }	Reset the routing table on the specified connection. <ul style="list-style-type: none"> • Enter an asterisk (*) to specify that all connections be reset. • Enter an IP <i>address</i> to specify the connection to be reset. • Enter a peer group name to reset the peer group.
Step 3	clear ip bgp { * <i>address</i> <i>peer-group-name</i> } soft out	(Optional) Perform an outbound soft reset to reset the inbound routing table on the specified connection. Use this command if route refresh is supported. <ul style="list-style-type: none"> • Enter an asterisk (*) to specify that all connections be reset. • Enter an IP <i>address</i> to specify the connection to be reset. • Enter a peer group name to reset the peer group.
Step 4	show ip bgp show ip bgp neighbors	Verify the reset by checking information about the routing table and about BGP neighbors.

EXAMPLE

In the following example, an outbound soft reset is initiated for sessions with all routers in the autonomous system numbered 35700:

```
Switch# clear ip bgp 35700 soft out
```

Configuring BGP Decision Attributes

When a BGP speaker receives updates from multiple autonomous systems that describe different paths to the same destination, it must choose the single best path for reaching that destination. The decision is based on the value of attributes that the update contains and other BGP-configurable factors. The selected path is entered into the BGP routing table and propagated to its neighbors.

When a BGP peer learns two EBGP paths for a prefix from a neighboring AS, it chooses the best path and inserts that path in the IP routing table. If BGP multipath support is enabled and the EBGP paths are learned from the same neighboring autonomous systems, multiple paths are installed in the IP routing table. Then, during packet switching, per-packet or per-destination load balancing is performed among the multiple paths. The **maximum-paths** router configuration command controls the number of paths allowed.

These factors summarize the order in which BGP evaluates the attributes for choosing the best path:

1. If the path specifies a next hop that is inaccessible, drop the update. The BGP next-hop attribute, automatically determined by the software, is the IP address of the next hop that is going to be used to reach a destination. For EBGP, this is usually the IP address of the neighbor specified by the **neighbor remote-as** router configuration command. You can disable next-hop processing by using route maps or the **neighbor next-hop-self** router configuration command.

2. Prefer the path with the largest weight (a Cisco proprietary parameter). The weight attribute is local to the router and not propagated in routing updates. By default, the weight attribute is 32768 for paths that the router originates and zero for other paths. You can use access lists, route maps, or the **neighbor weight** router configuration command to set weights.
3. Prefer the route with the highest local preference. Local preference is part of the routing update and exchanged among routers in the same AS. The default value of the local preference attribute is 100. You can set local preference by using the **bgp default local-preference** router configuration command or by using a route map.
4. Prefer the route that was originated by BGP running on the local router.
5. Prefer the route with the shortest AS path.
6. Prefer the route with the lowest origin type. An interior route or IGP is lower than a route learned by EGP, and an EGP-learned route is lower than one of unknown origin or learned in another way.
7. Prefer the route with the lowest multi-exit discriminator (MED) metric attribute if the neighboring AS is the same for all routes considered. You can configure the MED by using route maps or by using the **default-metric** router configuration command. When an update is sent to an IBGP peer, the MED is included.
8. Prefer the external (EBGP) path over the internal (IBGP) path.
9. Prefer the route that can be reached through the closest IGP neighbor (the lowest IGP metric). This means that the router will prefer the shortest internal path within the AS to reach the destination (the shortest path to the BGP next-hop).
10. If these conditions are all true, insert the route for this path into the IP routing table:
 - Both the best route and this route are external.
 - Both the best route and this route are from the same neighboring autonomous system.
 - Maximum-paths is enabled.
11. If multipath is not enabled, prefer the route with the lowest IP address value for the BGP router ID. The router ID is usually the highest IP address on the router or the loopback (virtual) address, but might be implementation-specific.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enable a BGP routing process, assign it an AS number, and enter router configuration mode.
Step 3	bgp best-path as-path ignore	(Optional) Configure the router to ignore AS path length in selecting a route.
Step 4	neighbor { <i>ip-address</i> <i>peer-group-name</i> } next-hop-self	(Optional) Disable next-hop processing on BGP updates to a neighbor by entering a specific IP address to be used instead of the next-hop address.

	Command	Purpose
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } weight <i>weight</i>	(Optional) Assign a weight to a neighbor connection. Acceptable values are from 0 to 65535; the largest weight is the preferred route. Routes learned through another BGP peer have a default weight of 0; routes sourced by the local router have a default weight of 32768.
Step 6	default-metric <i>number</i>	(Optional) Set a MED metric to set preferred paths to external neighbors. All routes without a MED will also be set to this value. The range is 1 to 4294967295. The lowest value is the most desirable.
Step 7	bgp bestpath med missing-as-worst	(Optional) Configure the switch to consider a missing MED as having a value of infinity, making the path without a MED value the least desirable path.
Step 8	bgp always-compare med	(Optional) Configure the switch to compare MEDs for paths from neighbors in different autonomous systems. By default, MED comparison is only done among paths in the same AS.
Step 9	bgp bestpath med confed	(Optional) Configure the switch to consider the MED in choosing a path from among those advertised by different subautonomous systems within a confederation.
Step 10	bgp deterministic med	(Optional) Configure the switch to consider the MED variable when choosing among routes advertised by different peers in the same AS.
Step 11	bgp default local-preference <i>value</i>	(Optional) Change the default local preference value. The range is 0 to 4294967295; the default value is 100. The highest local preference value is preferred.
Step 12	maximum-paths <i>number</i>	(Optional) Configure the number of paths to be added to the IP routing table. The default is to only enter the best path in the routing table. The range is from 1 to 8. Having multiple paths allows load balancing among the paths.
Step 13	end	Return to privileged EXEC mode.
Step 14	show ip bgp show ip bgp neighbors	Verify the reset by checking information about the routing table and about BGP neighbors.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no** form of each command to return to the default state.

EXAMPLE

The following example forces all updates destined for 10.108.1.1 to advertise this router as the next hop:

```
Switch(config)# router bgp 109
```

```
Switch(config-router)# neighbor 10.108.1.1 next-hop-self
```

In the following example, the local BGP routing process is configured to compare the MED from alternative paths, regardless of the autonomous system from which the paths are received:

```
Switch(config)# router bgp 500000
Switch(config-router)# bgp always-compare-med
```

Configuring BGP Filtering with Route Maps

Within BGP, you can use route maps to control and to modify routing information and to define the conditions by which routes are redistributed between routing domains. See the [“Using Route Maps to Redistribute Routing Information” section on page 2-121](#) for more information about route maps. Each route map has a name that identifies the route map (*map tag*) and an optional sequence number.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>map-tag</i> [[permit deny] sequence-number]	Create a route map, and enter route-map configuration mode.
Step 3	set ip next-hop <i>ip-address</i> [... <i>ip-address</i>] [peer-address]	(Optional) Set a route map to disable next-hop processing. <ul style="list-style-type: none"> In an inbound route map, set the next hop of matching routes to be the neighbor peering address, overriding third-party next hops. In an outbound route map of a BGP peer, set the next hop to the peering address of the local router, disabling the next-hop calculation.
Step 4	end	Return to privileged EXEC mode.
Step 5	show route-map [<i>map-name</i>]	Display all route maps configured or only the one specified to verify configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no route-map** *map-tag* command to delete the route map. Use the **no set ip next-hop** *ip-address* command to re-enable next-hop processing.

EXAMPLE

In the following example, the inbound route map named rmap sets the next hop:

```
Switch(config)# route-map rmap permit 10
Switch(config-route-map)# set ip next-hop 10.2.0.1
```

Configuring BGP Filtering by Neighbor

You can filter BGP advertisements by using AS-path filters, such as the **as-path access-list** global configuration command and the **neighbor filter-list** router configuration command. You can also use access lists with the **neighbor distribute-list** router configuration command. Distribute-list filters are applied to network numbers. See the [“Controlling Advertising and Processing in Routing Updates” section on page 2-130](#) for information about the **distribute-list** command.

You can use route maps on a per-neighbor basis to filter updates and to modify various attributes. A route map can be applied to either inbound or outbound updates. Only the routes that pass the route map are sent or accepted in updates. On both inbound and outbound updates, matching is supported based on AS path, community, and network numbers. Autonomous-system path matching requires the **match as-path access-list** route-map command, community-based matching requires the **match community-list** route-map command, and network-based matching requires the **ip access-list** global configuration command.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enable a BGP routing process, assign it an AS number, and enter router configuration mode.
Step 3	neighbor { <i>ip-address</i> <i>peer-group name</i> } distribute-list { <i>access-list-number</i> <i>name</i> } { in out }	(Optional) Filter BGP routing updates to or from neighbors as specified in an access list. Note You can also use the neighbor prefix-list router configuration command to filter updates, but you cannot use both commands to configure the same BGP peer.
Step 4	neighbor { <i>ip-address</i> <i>peer-group name</i> } route-map <i>map-tag</i> { in out }	(Optional) Apply a route map to filter an incoming or outgoing route.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp neighbors	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no neighbor distribute-list** command to remove the access list from the neighbor. Use the **no neighbor route-map** *map-tag* router configuration command to remove the route map from the neighbor.

EXAMPLE

The following router configuration mode example applies list 39 to incoming advertisements from neighbor 172.16.4.1. List 39 permits the advertisement of network 10.109.0.0.

```
Switch(config)# router bgp 109
Switch(config-router)# network 10.108.0.0
Switch(config-router)# neighbor 172.16.4.1 distribute-list 39 in
```

Configuring BGP Filtering By Access Lists

Another method of filtering is to specify an access list filter on both incoming and outbound updates, based on the BGP autonomous system paths. Each filter is an access list based on regular expressions. (See [Using Regular Expressions in BGP](#) for more information on forming regular expressions.) To use this method, define an autonomous system path access list, and apply it to updates to and from particular neighbors.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip as-path access-list <i>access-list-number</i> { permit deny } <i>as-regular-expressions</i>	Define a BGP-related access list.
Step 3	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 4	neighbor { <i>ip-address</i> <i>peer-group name</i> } filter-list { <i>access-list-number</i> <i>name</i> } { in out } [weight weight]	Establish a BGP filter based on an access list.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp neighbors [<i>paths</i> <i>regular-expression</i>]	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

In the following example, an autonomous system path access list (number 500) is defined to configure the router to not advertise any path through or from autonomous system 65535 to the 10.20.2.2 neighbor:

```
Switch(config)# ip as-path access-list 500 deny _65535_
Switch(config)# ip as-path access-list 500 deny ^65535$
Switch(config)# router bgp 50000
Switch(config-router)# neighbor 192.168.1.1 remote-as 65535
Switch(config-router)# neighbor 10.20.2.2 remote-as 40000
Switch(config-router)# neighbor 10.20.2.2 filter-list 500 out
Switch(config-router)# end
```

Configuring Prefix Lists for BGP Filtering

You can use prefix lists as an alternative to access lists in many BGP route filtering commands, including the **neighbor distribute-list** router configuration command. Filtering by a prefix list involves matching the prefixes of routes with those listed in the prefix list, as when matching access lists. When there is a match, the route is used. Whether a prefix is permitted or denied is based upon these rules:

- An empty prefix list permits all prefixes.

- An implicit deny is assumed if a given prefix does not match any entries in a prefix list.
- When multiple entries of a prefix list match a given prefix, the sequence number of a prefix list entry identifies the entry with the lowest sequence number.

By default, sequence numbers are generated automatically and incremented in units of five. If you disable the automatic generation of sequence numbers, you must specify the sequence number for each entry. You can specify sequence values in any increment. If you specify increments of one, you cannot insert additional entries into the list; if you choose very large increments, you might run out of values.

You do not need to specify a sequence number when removing a configuration entry. **Show** commands include the sequence numbers in their output.

Before using a prefix list in a command, you must set up the prefix list.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip prefix-list <i>list-name</i> [seq <i>seq-value</i>] deny permit <i>network/len</i> [ge <i>ge-value</i>] [le <i>le-value</i>]	Create a prefix list with an optional sequence number to deny or permit access for matching conditions. You must enter at least one permit or deny clause. <ul style="list-style-type: none"> • <i>network/len</i> is the network number and length (in bits) of the network mask. • (Optional) ge and le values specify the range of the prefix length to be matched. The specified <i>ge-value</i> and <i>le-value</i> must satisfy this condition: $len < ge-value < le-value < 32$
Step 3	ip prefix-list <i>list-name</i> seq <i>seq-value</i> deny permit <i>network/len</i> [ge <i>ge-value</i>] [le <i>le-value</i>]	(Optional) Add an entry to a prefix list, and assign a sequence number to the entry.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip prefix list [detail summary] <i>name</i> [<i>network/len</i>] [seq <i>seq-num</i>] [longer] [first-match]	Verify the configuration by displaying information about a prefix list or prefix list entries.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete a prefix list and all of its entries, use the **no ip prefix-list** *list-name* global configuration command. To delete an entry from a prefix list, use the **no ip prefix-list seq** *seq-value* global configuration command. To disable automatic generation of sequence numbers, use the **no ip prefix-list sequence number** command; to reenable automatic generation, use the **ip prefix-list sequence number** command. To clear the hit-count table of prefix list entries, use the **clear ip prefix-list** privileged EXEC command.

EXAMPLE

In the following example, a prefix list is configured to deny the default route 0.0.0.0/0:

```
Switch(config)# ip prefix-list RED deny 0.0.0.0/0
```

In the following example, a prefix list is configured to permit traffic from the 172.16.1.0/24 subnet:

```
Switch(config)# ip prefix-list BLUE permit 172.16.1.0/24
```

In the following example, a prefix list is configured to permit routes from the 10.0.0.0/8 network that have a mask length that is less than or equal to 24 bits:

```
Switch(config)# ip prefix-list YELLOW permit 10.0.0.0/8 le 24
```

In the following example, a prefix list is configured to deny routes from the 10.0.0.0/8 network that have a mask length that is greater than or equal to 25 bits:

```
Switch(config)# ip prefix-list PINK deny 10.0.0.0/8 ge 25
```

In the following example, a prefix list is configured to permit routes from any network that have a mask length from 8 to 24 bits:

```
Switch(config)# ip prefix-list GREEN permit 0.0.0.0/0 ge 8 le 24
```

In the following example, a prefix list is configured to deny any route with any mask length from the 10.0.0.0/8 network:

```
Switch(config)# ip prefix-list ORANGE deny 10.0.0.0/8 le 32
```

Configuring BGP Community Filtering

One way that BGP controls the distribution of routing information based on the value of the COMMUNITIES attribute. A *community* is a group of destinations that share some common attribute. Each destination can belong to multiple communities. AS administrators can define to which communities a destination belongs. By default, all destinations belong to the general Internet community. The community is identified by the COMMUNITIES attribute, an optional, transitive, global attribute in the numerical range from 1 to 4294967200. These are some predefined, well-known communities:

- **internet**—Advertise this route to the Internet community. All routers belong to it.
- **no-export**—Do not advertise this route to EBGp peers.
- **no-advertise**—Do not advertise this route to any peer (internal or external).
- **local-as**—Do not advertise this route to peers outside the local autonomous system.

Based on the community, you can control which routing information to accept, prefer, or distribute to other neighbors. A BGP speaker can set, append, or modify the community of a route when learning, advertising, or redistributing routes. When routes are aggregated, the resulting aggregate has a COMMUNITIES attribute that contains all communities from all the initial routes.

You can use community lists to create groups of communities to use in a match clause of a route map. As with an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded.

To set the COMMUNITIES attribute and match clauses based on communities, see the **match community-list** and **set community** route-map configuration commands in the [“Using Route Maps to Redistribute Routing Information”](#) section on page 2-121.

By default, no COMMUNITIES attribute is sent to a neighbor. You can specify that the COMMUNITIES attribute be sent to the neighbor at an IP address by using the **neighbor send-community** router configuration command.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip community-list <i>community-list-number</i> { permit deny } <i>community-number</i>	Create a community list, and assign it a number. <ul style="list-style-type: none"> The <i>community-list-number</i> is an integer from 1 to 99 that identifies one or more permit or deny groups of communities. The <i>community-number</i> is the number configured by a set community route-map configuration command.
Step 3	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 4	neighbor { <i>ip-address</i> <i>peer-group name</i> } send-community	Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
Step 5	set comm-list <i>list-num</i> delete	(Optional) Remove communities from the community attribute of an inbound or outbound update that match a standard or extended community list specified by a route map.
Step 6	exit	Return to global configuration mode.
Step 7	ip bgp-community new-format	(Optional) Display and parse BGP communities in the format AA:NN. A BGP community is displayed in a two-part format 2 bytes long. The Cisco default community format is in the format NNAA. In the most recent RFC for BGP, a community takes the form AA:NN, where the first part is the AS number and the second part is a 2-byte number.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ip bgp community	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

In the following example, a standard community list is configured that permits routes from network 10 in autonomous system 50000:

```
Router(config)# ip community-list 1 permit 50000:10
```

In the following router configuration mode example, the router belongs to autonomous system 109 and is configured to send the communities attribute to its neighbor at IP address 172.16.70.23:

```
Switch(config)# router bgp 109
Switch(config-router)# neighbor 172.16.70.23 send-community
```

In the following example, a router that uses the 32-bit number community format is upgraded to use the AA:NN format:

```
Switch(config)# ip bgp-community new-format
```

The following sample output shows how BGP community numbers are displayed when the **ip bgp-community new-format** command is enabled:

```
Switch# show ip bgp 10.0.0.0
BGP routing table entry for 10.0.0.0/8, version 4
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
    10.0.33.35
    35
    10.0.33.35 from 10.0.33.35 (192.168.3.3)
      Origin incomplete, metric 10, localpref 100, valid, external
      Community: 1:1
  Local
    0.0.0.0 from 0.0.0.0 (10.0.33.34)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
```

Configuring BGP Neighbors and Peer Groups

Often many BGP neighbors are configured with the same update policies (that is, the same outbound route maps, distribute lists, filter lists, update source, and so on). Neighbors with the same update policies can be grouped into peer groups to simplify configuration and to make updating more efficient. When you have configured many peers, we recommend this approach.

To configure a BGP peer group, you create the peer group, assign options to the peer group, and add neighbors as peer group members. You configure the peer group by using the **neighbor** router configuration commands. By default, peer group members inherit all the configuration options of the peer group, including the remote-as (if configured), version, update-source, out-route-map, out-filter-list, out-dist-list, minimum-advertisement-interval, and next-hop-self. All peer group members also inherit changes made to the peer group. Members can also be configured to override the options that do not affect outbound updates.

To assign configuration options to an individual neighbor, specify any of these router configuration commands by using the neighbor IP address. To assign the options to a peer group, specify any of the commands by using the peer group name. You can disable a BGP peer or peer group without removing all the configuration information by using the **neighbor shutdown** router configuration command.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	neighbor <i>peer-group-name</i> peer-group	Create a BGP peer group.
Step 4	neighbor <i>ip-address</i> peer-group <i>peer-group-name</i>	Make a BGP neighbor a member of the peer group.
Step 5	neighbor { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i>	Specify a BGP neighbor. If a peer group is not configured with a remote-as <i>number</i> , use this command to create peer groups containing EBGP neighbors. The range is 1 to 65535.

	Command	Purpose
Step 6	neighbor { <i>ip-address</i> <i>peer-group-name</i> } description <i>text</i>	(Optional) Associate a description with a neighbor.
Step 7	neighbor { <i>ip-address</i> <i>peer-group-name</i> } default-originate [route-map <i>map-name</i>]	(Optional) Allow a BGP speaker (the local router) to send the default route 0.0.0.0 to a neighbor for use as a default route.
Step 8	neighbor { <i>ip-address</i> <i>peer-group-name</i> } send-community	(Optional) Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
Step 9	neighbor { <i>ip-address</i> <i>peer-group-name</i> } update-source <i>interface</i>	(Optional) Allow internal BGP sessions to use any operational interface for TCP connections.
Step 10	neighbor { <i>ip-address</i> <i>peer-group-name</i> } ebgp-multihop	(Optional) Allow BGP sessions, even when the neighbor is not on a directly connected segment. The multihop session is not established if the only route to the multihop peer's address is the default route (0.0.0.0).
Step 11	neighbor { <i>ip-address</i> <i>peer-group-name</i> } local-as <i>number</i>	(Optional) Specify an AS number to use as the local AS. The range is 1 to 65535.
Step 12	neighbor { <i>ip-address</i> <i>peer-group-name</i> } advertisement-interval <i>seconds</i>	(Optional) Set the minimum interval between sending BGP routing updates.
Step 13	neighbor { <i>ip-address</i> <i>peer-group-name</i> } maximum-prefix <i>maximum</i> [<i>threshold</i>]	(Optional) Control how many prefixes can be received from a neighbor. The range is 1 to 4294967295. The <i>threshold</i> (optional) is the percentage of maximum at which a warning message is generated. The default is 75 percent.
Step 14	neighbor { <i>ip-address</i> <i>peer-group-name</i> } next-hop-self	(Optional) Disable next-hop processing on the BGP updates to a neighbor.
Step 15	neighbor { <i>ip-address</i> <i>peer-group-name</i> } password <i>string</i>	(Optional) Set MD5 authentication on a TCP connection to a BGP peer. The same password must be configured on both BGP peers, or the connection between them is not made.
Step 16	neighbor { <i>ip-address</i> <i>peer-group-name</i> } route-map <i>map-name</i> { in out }	(Optional) Apply a route map to incoming or outgoing routes.
Step 17	neighbor { <i>ip-address</i> <i>peer-group-name</i> } send-community	(Optional) Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
Step 18	neighbor { <i>ip-address</i> <i>peer-group-name</i> } timers <i>keepalive holdtime</i>	(Optional) Set timers for the neighbor or peer group. <ul style="list-style-type: none"> The <i>keepalive</i> interval is the time within which keepalive messages are sent to peers. The range is 1 to 4294967295 seconds; the default is 60. The <i>holdtime</i> is the interval after which a peer is declared inactive after not receiving a keepalive message from it. The range is 1 to 4294967295 seconds; the default is 180.
Step 19	neighbor { <i>ip-address</i> <i>peer-group-name</i> } weight <i>weight</i>	(Optional) Specify a weight for all routes from a neighbor.
Step 20	neighbor { <i>ip-address</i> <i>peer-group-name</i> } distribute-list { <i>access-list-number</i> <i>name</i> } { in out }	(Optional) Filter BGP routing updates to or from neighbors, as specified in an access list.

	Command	Purpose
Step 21	neighbor { <i>ip-address</i> <i>peer-group-name</i> } filter-list <i>access-list-number</i> { in out weight <i>weight</i> }	(Optional) Establish a BGP filter.
Step 22	neighbor { <i>ip-address</i> <i>peer-group-name</i> } version <i>value</i>	(Optional) Specify the BGP version to use when communicating with a neighbor.
Step 23	neighbor { <i>ip-address</i> <i>peer-group-name</i> } soft-reconfiguration inbound	(Optional) Configure the software to start storing received updates.
Step 24	end	Return to privileged EXEC mode.
Step 25	show ip bgp neighbors	Verify the configuration.
Step 26	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an existing BGP neighbor or neighbor peer group, use the **neighbor shutdown** router configuration command. To enable a previously existing neighbor or neighbor peer group that had been disabled, use the **no neighbor shutdown** router configuration command.

EXAMPLE

The following example configures a peer group and sets the minimum time between sending BGP routing updates to 10 seconds for the peer group:

```
Switch(config)# router bgp 45000
Switch(config-router)# neighbor mygroup peer-group
Switch(config-router)# neighbor 192.168.1.2 remote-as 40000
Switch(config-router)# neighbor 192.168.3.2 remote-as 50000
Switch(config-router)# neighbor 192.168.1.2 peer-group mygroup
Switch(config-router)# neighbor 192.168.3.2 peer-group mygroup
Switch(config-router)# neighbor mygroup advertisement-interval 10
```

Configuring Aggregate Addresses

Classless interdomain routing (CIDR) enables you to create aggregate routes (or *supernets*) to minimize the size of routing tables. You can configure aggregate routes in BGP either by redistributing an aggregate route into BGP or by creating an aggregate entry in the BGP routing table. An aggregate address is added to the BGP table when there is at least one more specific entry in the BGP table.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.

	Command	Purpose
Step 3	aggregate-address <i>address mask</i>	Create an aggregate entry in the BGP routing table. The aggregate route is advertised as coming from the AS, and the atomic aggregate attribute is set to indicate that information might be missing.
Step 4	aggregate-address <i>address mask as-set</i>	(Optional) Generate AS set path information. This command creates an aggregate entry following the same rules as the previous command, but the advertised path will be an AS_SET consisting of all elements contained in all paths. Do not use this keyword when aggregating many paths because this route must be continually withdrawn and updated.
Step 5	aggregate-address <i>address-mask summary-only</i>	(Optional) Advertise summary addresses only.
Step 6	aggregate-address <i>address mask suppress-map map-name</i>	(Optional) Suppress selected, more specific routes.
Step 7	aggregate-address <i>address mask advertise-map map-name</i>	(Optional) Generate an aggregate based on conditions specified by the route map.
Step 8	aggregate-address <i>address mask attribute-map map-name</i>	(Optional) Generate an aggregate with attributes specified in the route map.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip bgp neighbors [advertised-routes]	Verify the configuration.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an aggregate entry, use the **no aggregate-address** *address mask* router configuration command. To return options to the default values, use the command with keywords.

EXAMPLE

In the following example, an aggregate BGP address is created in router configuration mode. The path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized.

```
Switch(config)# router bgp 50000
Switch(config-router)# aggregate-address 10.0.0.0 255.0.0.0 as-set
```

In the following example, a route map called MAP-ONE is created to match on an AS-path access list. The path advertised for this route will be an AS_SET consisting of elements contained in paths that are matched in the route map.

```
Switch(config)# ip as-path access-list 1 deny ^1234_
Switch(config)# ip as-path access-list 1 permit .*
Switch(config)# !
Switch(config)# route-map MAP-ONE
Switch(config-route-map)# match ip as-path 1
Switch(config-route-map)# exit
Switch(config)# router bgp 50000
Switch(config-router)# address-family ipv4
Switch(config-router-af)# aggregate-address 10.0.0.0 255.0.0.0 as-set advertise-map
MAP-ONE
Switch(config-router-af)# end
```

Configuring Routing Domain Confederations

One way to reduce the IBGP mesh is to divide an autonomous system into multiple subautonomous systems and to group them into a single confederation that appears as a single autonomous system. Each autonomous system is fully meshed within itself and has a few connections to other autonomous systems in the same confederation. Even though the peers in different autonomous systems have EBGP sessions, they exchange routing information as if they were IBGP peers. Specifically, the next hop, MED, and local preference information is preserved. You can then use a single IGP for all of the autonomous systems.

To configure a BGP confederation, you must specify a confederation identifier that acts as the autonomous system number for the group of autonomous systems.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	bgp confederation identifier <i>autonomous-system</i>	Configure a BGP confederation identifier.
Step 4	bgp confederation peers <i>autonomous-system</i> [<i>autonomous-system</i> ...]	Specify the autonomous systems that belong to the confederation and that will be treated as special EBGP peers.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp neighbor show ip bgp network	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

In the following example, the routing domain is divided into autonomous systems 50001, 50002, 50003, 50004, 50005, and 50006 and is identified by the confederation identifier 50007. Neighbor 10.2.3.4 is a peer inside of the routing domain confederation. Neighbor 10.4.5.6 is a peer outside of the routing domain confederation. To external peers and routing domains, the confederation appears as a single autonomous system with the number 50007.

```
router bgp 50000
  bgp confederation identifier 50007
  bgp confederation peers 50001 50002 50003 50004 50005 50006
  neighbor 10.2.3.4 remote-as 50001
  neighbor 10.4.5.6 remote-as 40000
end
```

Configuring BGP Route Reflectors

BGP requires that all of the IBGP speakers be fully meshed. When a router receives a route from an external neighbor, it must advertise it to all internal neighbors. To prevent a routing information loop, all IBGP speakers must be connected. The internal neighbors do not send routes learned from internal neighbors to other internal neighbors.

With route reflectors, all IBGP speakers need not be fully meshed because another method is used to pass learned routes to neighbors. When you configure an internal BGP peer to be a *route reflector*, it is responsible for passing IBGP learned routes to a set of IBGP neighbors. The internal peers of the route reflector are divided into two groups: *client peers* and *nonclient peers* (all the other routers in the autonomous system). A route reflector reflects routes between these two groups. The route reflector and its client peers form a *cluster*. The nonclient peers must be fully meshed with each other, but the client peers need not be fully meshed. The clients in the cluster do not communicate with IBGP speakers outside their cluster.

When the route reflector receives an advertised route, it takes one of these actions, depending on the neighbor:

- A route from an external BGP speaker is advertised to all clients and nonclient peers.
- A route from a nonclient peer is advertised to all clients.
- A route from a client is advertised to all clients and nonclient peers. Hence, the clients need not be fully meshed.

Usually a cluster of clients have a single route reflector, and the cluster is identified by the route reflector router ID. To increase redundancy and to avoid a single point of failure, a cluster might have more than one route reflector. In this case, all route reflectors in the cluster must be configured with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster. All the route reflectors serving a cluster should be fully meshed and should have identical sets of client and nonclient peers.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	neighbor <i>ip-address</i> <i>peer-group-name</i> route-reflector-client	Configure the local router as a BGP route reflector and the specified neighbor as a client.
Step 4	bgp cluster-id <i>cluster-id</i>	(Optional) Configure the cluster ID if the cluster has more than one route reflector.
Step 5	no bgp client-to-client reflection	(Optional) Disable client-to-client route reflection. By default, the routes from a route reflector client are reflected to other clients. However, if the clients are fully meshed, the route reflector does not need to reflect routes to clients.
Step 6	end	Return to privileged EXEC mode.

	Command	Purpose
Step 7	show ip bgp	Verify the configuration. Display the originator ID and the cluster-list attributes.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

In the following router configuration mode example, the local router is a route reflector. It passes learned IBGP routes to the neighbor at 172.16.70.24.

```
router bgp 5
 neighbor 172.16.70.24 route-reflector-client
```

Configuring Route Dampening

Route flap dampening minimizes the propagation of flapping routes across an internetwork. A route is considered to be flapping when it is repeatedly available, then unavailable, then available, then unavailable, and so on. When route dampening is enabled, a numeric *penalty* value is assigned to a route when it flaps. When a route's accumulated penalties reach a configurable limit, BGP suppresses advertisements of the route, even if the route is running. The *reuse limit* is a configurable value that is compared with the penalty. If the penalty is less than the reuse limit, a suppressed route that is up is advertised again.

Dampening is not applied to routes that are learned by IBGP. This policy prevents the IBGP peers from having a higher penalty for routes external to the AS.

BEFORE YOU BEGIN

Enable BGP routing as described in the [“Enabling BGP Routing” procedure on page 2-55](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system</i>	Enter BGP router configuration mode.
Step 3	bgp dampening	Enable BGP route dampening.
Step 4	bgp dampening <i>half-life reuse suppress max-suppress</i> [route-map <i>map</i>]	(Optional) Change the default values of route dampening factors.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ip bgp flap-statistics [{ regex <i>regex</i> } { filter-list <i>list</i> } { address mask [longer-prefix]}]	(Optional) Monitor the flaps of all paths that are flapping. The statistics are deleted when the route is not suppressed and is stable.
Step 7	show ip bgp dampened-paths	(Optional) Display the dampened routes, including the time remaining before they are suppressed.
Step 8	clear ip bgp flap-statistics [{ regex <i>regex</i> } { filter-list <i>list</i> } { address mask [longer-prefix]}]	(Optional) Clear BGP flap statistics to make it less likely that a route will be dampened.

	Command	Purpose
Step 9	clear ip bgp dampening	(Optional) Clear route dampening information, and unsuppress the suppressed routes.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable flap dampening, use the **no bgp dampening** router configuration command without keywords. To set dampening factors back to the default values, use the **no bgp dampening** router configuration command with values.

EXAMPLE

In the following example, BGP dampening is applied to prefixes filtered through the route-map named BLUE:

```
Switch(config)# ip prefix-list RED permit 10.0.0.0/8
Switch(config)# !
Switch(config)# route-map BLUE

Switch(config-route-map)# match ip address ip prefix-list RED
Switch(config-route-map)# exit
Switch(config)# router bgp 50000

Switch(config-router)# address-family ipv4
Switch(config-router-af)# bgp dampening route-map BLUE
Switch(config-router-af)# end
```

Monitoring and Maintaining BGP

You can remove all contents of a particular cache, table, or database. This might be necessary when the contents of the particular structure have become or are suspected to be invalid.

You can display specific statistics, such as the contents of BGP routing tables, caches, and databases. You can use the information to get 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.

Command	Purpose
clear ip bgp address	Reset a particular BGP connection.
clear ip bgp *	Reset all BGP connections.
clear ip bgp peer-group tag	Remove all members of a BGP peer group.
show ip bgp prefix	Display peer groups and peers not in peer groups to which the prefix has been advertised. Also display prefix attributes such as the next hop and the local prefix.
show ip bgp cidr-only	Display all BGP routes that contain subnet and supernet network masks.
show ip bgp community [community-number] [exact]	Display routes that belong to the specified communities.
show ip bgp community-list community-list-number [exact-match]	Display routes that are permitted by the community list.

Command	Purpose
show ip bgp filter-list <i>access-list-number</i>	Display routes that are matched by the specified AS path access list.
show ip bgp inconsistent-as	Display the routes with inconsistent originating autonomous systems.
show ip bgp regexp <i>regular-expression</i>	Display the routes that have an AS path that matches the specified regular expression entered on the command line.
show ip bgp	Display the contents of the BGP routing table.
show ip bgp neighbors [<i>address</i>]	Display detailed information on the BGP and TCP connections to individual neighbors.
show ip bgp neighbors [<i>address</i>] [advertised-routes dampened-routes flap-statistics paths <i>regular-expression</i> received-routes routes]	Display routes learned from a particular BGP neighbor.
show ip bgp paths	Display all BGP paths in the database.
show ip bgp peer-group [<i>tag</i>] [summary]	Display information about BGP peer groups.
show ip bgp summary	Display the status of all BGP connections.

You can also enable the logging of messages generated when a BGP neighbor resets, comes up, or goes down by using the **bgp log-neighbor changes** router configuration command.

Configuring ISO CLNS Routing

The International Organization for Standardization (ISO) Connectionless Network Service (CLNS) protocol is a standard for the network layer of the Open System Interconnection (OSI) model. Addresses in the ISO network architecture are referred to as network service access point (NSAP) addresses and network entity titles (NETs). Each node in an OSI network has one or more NETs. In addition, each node has many NSAP addresses.

When you enable connectionless routing on the switch by using the **clns routing** global configuration command, the switch makes only forwarding decisions, with no routing-related functionality. For dynamic routing, you must also enable a routing protocol. The switch supports the Intermediate System-to-Intermediate System (IS-IS) dynamic routing protocols for ISO CLNS networks. This routing protocol supports the concept of *areas*. Within an area, all routers know how to reach all the system IDs. Between areas, routers know how to reach the proper area. IS-IS supports two levels of routing: *station routing* (within an area) and *area routing* (between areas).

The key difference between the ISO IGRP and IS-IS NSAP addressing schemes is in the definition of area addresses. Both use the system ID for Level 1 routing (routing within an area). However, they differ in the way addresses are specified for area routing. An ISO IGRP NSAP address includes three separate fields for routing: the *domain*, *area*, and *system ID*. An IS-IS address includes two fields: a single continuous *area* field (comprising the domain and area fields) and the *system ID*.

For more detailed information about ISO CLNS, see the ISO CLNS documents listed in the “[Related Documents](#)” section on page 2-135.

Configuring IS-IS Dynamic Routing

IS-IS is an ISO dynamic routing protocol. Enabling IS-IS requires that you create an IS-IS routing process and assign it to a specific interface, rather than to a network. You can specify more than one IS-IS routing process per Layer 3 switch or router by using the multiarea IS-IS configuration syntax. You then configure the parameters for each instance of the IS-IS routing process.

Small IS-IS networks are built as a single area that includes all the routers in the network. As the network grows larger, it is usually reorganized into a backbone area made up of the connected set of all Level 2 routers from all areas, which is in turn connected to local areas. Within a local area, routers know how to reach all system IDs. Between areas, routers know how to reach the backbone, and the backbone routers know how to reach other areas.

Routers establish Level 1 adjacencies to perform routing within a local area (station routing). Routers establish Level 2 adjacencies to perform routing between Level 1 areas (area routing).

A single Cisco router can participate in routing in up to 29 areas and can perform Level 2 routing in the backbone. In general, each routing process corresponds to an area. By default, the first instance of the routing process configured performs both Level 1 and Level 2 routing. You can configure additional router instances, which are automatically treated as Level 1 areas. You must configure the parameters for each instance of the IS-IS routing process individually.

For IS-IS multiarea routing, you can configure only one process to perform Level 2 routing, although you can define up to 29 Level 1 areas for each Cisco unit. If Level 2 routing is configured on any process, all additional processes are automatically configured as Level 1. You can configure this process to perform Level 1 routing at the same time. If Level 2 routing is not desired for a router instance, remove the Level 2 capability using the **is-type** global configuration command. Use the **is-type** command also to configure a different router instance as a Level 2 router.

This section briefly describes how to configure IS-IS routing. For more detailed information about IS-IS, see the IS-IS documents listed in the [“Related Documents” section on page 2-135](#).

This section includes the following topics:

- [Default IS-IS Configuration, page 2-77](#)
- [Nonstop Forwarding Awareness, page 2-78](#)
- [Configuring IS-IS Global Parameters, page 2-81](#)
- [Configuring IS-IS Interface Parameters, page 2-84](#)

Default IS-IS Configuration

Feature	Default Setting
Ignore link-state PDU (LSP) errors	Enabled.
IS-IS type	Conventional IS-IS: the router acts as both a Level 1 (station) and a Level 2 (area) router. Multiarea IS-IS: the first instance of the IS-IS routing process is a Level 1-2 router. Remaining instances are Level 1 routers.
Default-information originate	Disabled.
Log IS-IS adjacency state changes.	Disabled.

Feature	Default Setting
LSP generation throttling timers	Maximum interval between two consecutive occurrences: 5 seconds. Initial LSP generation delay: 50 ms. Hold time between the first and second LSP generation: 5000 ms.
LSP maximum lifetime (without a refresh)	1200 seconds (20 minutes) before the LSP packet is deleted.
LSP refresh interval	Send LSP refreshes every 900 seconds (15 minutes).
Maximum LSP packet size	1497 bytes.
NSF ¹ Awareness	Enabled ² . Allows Layer 3 switches to continue forwarding packets from a neighboring NSF-capable router during hardware or software changes.
Partial route computation (PRC) throttling timers	Maximum PRC wait interval: 5 seconds. Initial PRC calculation delay after a topology change: 2000 ms. Hold time between the first and second PRC calculation: 5000 ms.
Partition avoidance	Disabled.
Password	No area or domain password is defined, and authentication is disabled.
Set-overload-bit	Disabled. When enabled, if no arguments are entered, the overload bit is set immediately and remains set until you enter the no set-overload-bit command.
Shortest path first (SPF) throttling timers	Maximum interval between consecutive SFPS: 10 seconds. Initial SFP calculation after a topology change: 5500 ms. Holdtime between the first and second SFP calculation: 5500 ms.
Summary-address	Disabled.

1. NSF = Nonstop Forwarding

2. IS-IS NSF awareness is enabled for IPv4 on switches running the IP services image.

Nonstop Forwarding Awareness

The integrated IS-IS NSF Awareness feature is supported for IPv4 in the IP services image. The feature allows customer premises equipment (CPE) routers that are NSF-aware to help NSF-capable routers perform nonstop forwarding of packets. The local router is not necessarily performing NSF, but its awareness of NSF allows the integrity and accuracy of the routing database and link-state database on the neighboring NSF-capable router to be maintained during the switchover process.

This feature is automatically enabled and requires no configuration. For more information on this feature, see the “[Configuring Nonstop Forwarding](#)” chapter in the *High Availability Configuration Guide, Cisco IOS Release 15S*.

Enabling IS-IS Routing

To enable IS-IS, you specify a name and NET for each routing process. You then enable IS-IS routing on the interface and specify the area for each instance of the routing process.

BEFORE YOU BEGIN

You should know your network design and how you want traffic to flow through it before configuring IS-IS. Define areas, prepare an addressing plan for the devices (including defining the NETs), and determine the interfaces that will run integrated IS-IS. To facilitate verification, a matrix of adjacencies should be prepared before you configure your devices, showing what neighbors should be expected in the adjacencies table.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	clns routing	Enable ISO connectionless routing on the switch.
Step 3	router isis [<i>area tag</i>]	Enable the IS-IS routing for the specified routing process and enter IS-IS routing configuration mode. (Optional) Use the <i>area tag</i> argument to identify the area to which the IS-IS router is assigned. You must enter a value if you are configuring multiple IS-IS areas. The first IS-IS instance configured is Level 1-2 by default. Later instances are automatically Level 1. You can change the level of routing by using the is-type global configuration command.
Step 4	net <i>network-entity-title</i>	Configure the NETs for the routing process. If you are configuring multiarea IS-IS, specify a NET for each routing process. You can specify a name for a NET and for an address.
Step 5	is-type { level-1 level-1-2 level-2-only }	(Optional) You can configure the router to act as a Level 1 (station) router, a Level 2 (area) router for multi-area routing, or both (the default): <ul style="list-style-type: none">• level-1—act as a station router only• level-1-2—act as both a station router and an area router• level 2—act as an area router only
Step 6	exit	Return to global configuration mode.
Step 7	interface <i>interface-id</i>	Specify an interface to route IS-IS, and enter interface configuration mode. If the interface is not already configured as a Layer 3 interface, enter the no switchport command to put it into Layer 3 mode.
Step 8	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 9	ip router isis [<i>area tag</i>]	Configure an IS-IS routing process for ISO CLNS on the interface and attach an area designator to the routing process.
Step 10	clns router isis [<i>area tag</i>]	Enable ISO CLNS on the interface.

	Command	Purpose
Step 11	ip address <i>ip-address-mask</i>	Define the IP address for the interface. An IP address is required on all interfaces in an area enabled for IS-IS if any one interface is configured for IS-IS routing.
Step 12	end	Return to privileged EXEC mode.
Step 13	show isis [<i>area tag</i>] database detail	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IS-IS routing, use the **no router isis** *area-tag* router configuration command.

EXAMPLE

This example shows how to configure three routers to run conventional IS-IS as an IP routing protocol. In conventional IS-IS, all routers act as Level 1 and Level 2 routers (by default).

Router A:

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# net 49.0001.0000.0000.000a.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config-router)# exit
```

Router B:

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# net 49.0001.0000.0000.000b.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config-router)# exit
```

Router C:

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# net 49.0001.0000.0000.000c.00
Switch(config-router)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# ip router isis
Switch(config-if)# clns router isis
Switch(config-router)# exit
```

Configuring IS-IS Global Parameters

These are some optional IS-IS global parameters that you can configure:

- You can force a default route into an IS-IS routing domain by configuring a default route controlled by a route map. You can also specify other filtering options configurable under a route map.
- You can configure the router to ignore IS-IS LSPs that are received with internal checksum errors or to purge corrupted LSPs, which causes the initiator of the LSP to regenerate it.
- You can assign passwords to areas and domains.
- You can create aggregate addresses that are represented in the routing table by a summary address (route-summarization). Routes learned from other routing protocols can also be summarized. The metric used to advertise the summary is the smallest metric of all the specific routes.
- You can set an overload bit.
- You can configure the LSP refresh interval and the maximum time that an LSP can remain in the router database without a refresh
- You can set the throttling timers for LSP generation, shortest path first computation, and partial route computation.
- You can configure the switch to generate a log message when an IS-IS adjacency changes state (up or down).
- If a link in the network has a maximum transmission unit (MTU) size of less than 1500 bytes, you can lower the LSP MTU so that routing will still occur.
- The partition avoidance router configuration command prevents an area from becoming partitioned when full connectivity is lost among a Level1-2 border router, adjacent Level 1 routers, and end hosts.

BEFORE YOU BEGIN

Enable IS-IS routing as described in the [“Enabling IS-IS Routing” procedure on page 2-79](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	cls routing	Enable ISO connectionless routing on the switch.
Step 3	router isis	Specify the IS-IS routing protocol and enter router configuration mode.
Step 4	default-information originate [route-map <i>map-name</i>]	(Optional) Force a default route into the IS-IS routing domain. If you enter route-map <i>map-name</i> , the routing process generates the default route if the route map is satisfied.
Step 5	ignore-lsp-errors	(Optional) Configure the router to ignore LSPs with internal checksum errors, instead of purging the LSPs. This command is enabled by default (corrupted LSPs are dropped). To purge the corrupted LSPs, enter the no ignore-lsp-errors router configuration command.
Step 6	area-password <i>password</i>	(Optional) Configure the area authentication password, which is inserted in Level 1 (station router level) LSPs.

	Command	Purpose
Step 7	domain-password <i>password</i>	(Optional) Configure the routing domain authentication password, which is inserted in Level 2 (area router level) LSPs.
Step 8	summary-address <i>address mask</i> [level-1 level-1-2 level-2]	(Optional) Create a summary of addresses for a given level.
Step 9	set-overload-bit [on-startup { <i>seconds</i> wait-for-bgp }]	<p>(Optional) Set an overload bit (a hippity bit) to allow other routers to ignore the router in their shortest path first (SPF) calculations if the router is having problems.</p> <ul style="list-style-type: none"> • (Optional) on-startup—sets the overload bit only on startup. If on-startup is not specified, the overload bit is set immediately and remains set until you enter the no set-overload-bit command. If on-startup is specified, you must enter a number of seconds or wait-for-bgp. • <i>seconds</i>—When the on-startup keyword is configured, causes the overload bit to be set upon system startup and remain set for this number of seconds. The range is from 5 to 86400 seconds. • wait-for-bgp—When the on-startup keyword is configured, causes the overload bit to be set upon system startup and remain set until BGP has converged. If BGP does not signal IS-IS that it is converged, IS-IS will turn off the overload bit after 10 minutes.
Step 10	lsp-refresh-interval <i>seconds</i>	(Optional) Set an LSP refresh interval in seconds. The range is from 1 to 65535 seconds. The default is to send LSP refreshes every 900 seconds (15 minutes).
Step 11	max-lsp-lifetime <i>seconds</i>	(Optional) Set the maximum time that LSP packets remain in the router database without being refreshed. The range is from 1 to 65535 seconds. The default is 1200 seconds (20 minutes). After the specified time interval, the LSP packet is deleted.
Step 12	lsp-gen-interval [level-1 level-2] <i>lsp-max-wait</i> [<i>lsp-initial-wait</i> <i>lsp-second-wait</i>]	<p>(Optional) Set the IS-IS LSP generation throttling timers:</p> <ul style="list-style-type: none"> • <i>lsp-max-wait</i>—the maximum interval (in seconds) between two consecutive occurrences of an LSP being generated. The range is 1 to 120, the default is 5. • <i>lsp-initial-wait</i>—the initial LSP generation delay (in milliseconds). The range is 1 to 10000; the default is 50. • <i>lsp-second-wait</i>—the hold time between the first and second LSP generation (in milliseconds). The range is 1 to 10000; the default is 5000.

	Command	Purpose
Step 13	spf-interval [level-1 level-2] <i>spf-max-wait</i> [<i>spf-initial-wait</i> <i>spf-second-wait</i>]	(Optional) Sets IS-IS shortest path first (SPF) throttling timers. <ul style="list-style-type: none"> • <i>spf-max-wait</i>—the maximum interval between consecutive SFPs (in seconds). The range is 1 to 120, the default is 10. • <i>spf-initial-wait</i>—the initial SFP calculation after a topology change (in milliseconds). The range is 1 to 10000; the default is 5500. • <i>spf-second-wait</i>—the holdtime between the first and second SFP calculation (in milliseconds). The range is 1 to 10000; the default is 5500.
Step 14	prc-interval <i>prc-max-wait</i> [<i>prc-initial-wait</i> <i>prc-second-wait</i>]	(Optional) Sets IS-IS partial route computation (PRC) throttling timers. <ul style="list-style-type: none"> • <i>prc-max-wait</i>—the maximum interval (in seconds) between two consecutive PRC calculations. The range is 1 to 120; the default is 5. • <i>prc-initial-wait</i>—the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 10,000; the default is 2000. • <i>prc-second-wait</i>—the hold time between the first and second PRC calculation (in milliseconds). The range is 1 to 10,000; the default is 5000.
Step 15	log-adjacency-changes [detail]	(Optional) Set the router to log IS-IS adjacency state changes. Enter detail to include all changes generated by events that are not related to the Intermediate System-to-Intermediate System Hellos, including End System-to-Intermediate System PDUs and link state packets (LSPs).
Step 16	lsp-mtu <i>size</i>	(Optional) Specify the maximum LSP packet size in bytes. The range is 128 to 4352; the default is 1497 bytes. Note If any link in the network has a reduced MTU size, you must change the LSP MTU size on all routers in the network.
Step 17	partition avoidance	(Optional) Causes an IS-IS Level 1-2 border router to stop advertising the Level 1 area prefix into the Level 2 backbone when full connectivity is lost among the border router, all adjacent level 1 routers, and end hosts.
Step 18	end	Return to privileged EXEC mode.
Step 19	show clns	Verify your entries.
Step 20	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable default route generation, use the **no default-information originate** router configuration command. Use the **no area-password** or **no domain-password** router configuration command to disable passwords. To disable LSP MTU settings, use the **no lsp mtu** router configuration command. To return to the default conditions for summary addressing, LSP refresh interval, LSP lifetime, LSP timers, SFP timers, and PRC timers, use the **no** form of the commands. Use the **no partition avoidance** router configuration command to disable the output format.

EXAMPLE

```
Switch(config)# clns routing
Switch(config)# router isis
Switch(config-router)# set-overloadbit on-startup 360
Switch(config-router)# log-adjacency-changes
Switch(config-router)# ignore-lsp-errors
Switch(config-router)# max-lsp-lifetime 65535
Switch(config-router)# lsp-refresh-interval 65000
Switch(config-router)# spf-interval 5 1 50
Switch(config-router)# prc-interval 5 1 50
Switch(config-router)# lsp-gen-interval 5 1 50
Switch(config-router)# end
```

Configuring IS-IS Interface Parameters

You can optionally configure certain interface-specific IS-IS parameters, independently from other attached routers. However, if you change some values from the defaults, such as multipliers and time intervals, it makes sense to also change them on multiple routers and interfaces. Most of the interface parameters can be configured for level 1, level 2, or both.

These are some interface level parameters you can configure:

- The default metric on the interface, which is used as a value for the IS-IS metric and assigned when there is no quality of service (QoS) routing performed.
- The hello interval (length of time between hello packets sent on the interface) or the default hello packet multiplier used on the interface to determine the hold time sent in IS-IS hello packets. The hold time determines how long a neighbor waits for another hello packet before declaring the neighbor down. This determines how quickly a failed link or neighbor is detected so that routes can be recalculated. Change the hello multiplier in circumstances where hello packets are lost frequently and IS-IS adjacencies are failing unnecessarily. You can raise the hello multiplier and lower the hello interval correspondingly to make the hello protocol more reliable without increasing the time required to detect a link failure.
- Other time intervals:
 - Complete sequence number PDU (CSNP) interval. CSNPs are sent by the designated router to maintain database synchronization.
 - Retransmission interval. This is the time between retransmission of IS-IS LSPs for point-to-point links.
 - IS-IS LSP retransmission throttle interval. This is the maximum rate (number of milliseconds between packets) at which IS-IS LSPs are re-sent on point-to-point links. This interval is different from the retransmission interval, which is the time between successive retransmissions of the *same* LSP.
- Designated router election priority, which allows you to reduce the number of adjacencies required on a multiaccess network, which in turn reduces the amount of routing protocol traffic and the size of the topology database.
- The interface circuit type, which is the type of adjacency desired for neighbors on the specified interface.
- Password authentication for the interface.

BEFORE YOU BEGIN

Enable IS-IS routing as described in the [“Enabling IS-IS Routing” procedure on page 2-79](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify the interface to be configured and enter interface configuration mode. If the interface is not already configured as a Layer 3 interface, enter the no switchport command to put it into Layer 3 mode.
Step 3	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 4	isis metric <i>default-metric</i> [level-1 level-2]	(Optional) Configure the metric (or cost) for the specified interface. The range is from 0 to 63. The default is 10. If no level is entered, the default is to apply to both Level 1 and Level 2 routers.
Step 5	isis hello-interval { <i>seconds</i> minimal } [level-1 level-2]	(Optional) Specify the length of time between hello packets sent by the switch. By default, a value three times the hello interval <i>seconds</i> is advertised as the <i>holdtime</i> in the hello packets sent. With smaller hello intervals, topological changes are detected faster, but there is more routing traffic. <ul style="list-style-type: none"> • minimal—causes the system to compute the hello interval based on the hello multiplier so that the resulting hold time is 1 second. • <i>seconds</i>—the range is from 1 to 65535. The default is 10 seconds.
Step 6	isis hello-multiplier <i>multiplier</i> [level-1 level-2]	(Optional) Specify the number of IS-IS hello packets a neighbor must miss before the router should declare the adjacency as down. The range is from 3 to 1000. The default is 3. Using a smaller hello multiplier causes fast convergence, but can result in more routing instability.
Step 7	isis csnp-interval <i>seconds</i> [level-1 level-2]	(Optional) Configure the IS-IS complete sequence number PDU (CSNP) interval for the interface. The range is from 0 to 65535. The default is 10 seconds.
Step 8	isis retransmit-interval <i>seconds</i>	(Optional) Configure the number of seconds between retransmission of IS-IS LSPs for point-to-point links. The value you specify should be an integer greater than the expected round-trip delay between any two routers on the network. The range is from 0 to 65535. The default is 5 seconds.
Step 9	isis retransmit-throttle-interval <i>milliseconds</i>	(Optional) Configure the IS-IS LSP retransmission throttle interval, which is the maximum rate (number of milliseconds between packets) at which IS-IS LSPs will be re-sent on point-to-point links. The range is from 0 to 65535. The default is determined by the isis lsp-interval command.
Step 10	isis priority <i>value</i> [level-1 level-2]	(Optional) Configure the priority to use for designated router election. The range is from 0 to 127. The default is 64.

	Command	Purpose
Step 11	isis circuit-type { level-1 level-1-2 level-2-only }	(Optional) Configure the type of adjacency desired for neighbors on the specified interface (specify the interface circuit type). <ul style="list-style-type: none"> • level-1—a Level 1 adjacency is established if there is at least one area address common to both this node and its neighbors. • level-1-2—a Level 1 and 2 adjacency is established if the neighbor is also configured as both Level 1 and Level 2 and there is at least one area in common. If there is no area in common, a Level 2 adjacency is established. This is the default. • level 2—a Level 2 adjacency is established. If the neighbor router is a Level 1 router, no adjacency is established.
Step 12	isis password <i>password</i> [level-1 level-2]	(Optional) Configure the authentication password for an interface. By default, authentication is disabled. Specifying Level 1 or Level 2 enables the password only for Level 1 or Level 2 routing, respectively. If you do not specify a level, the default is Level 1 and Level 2.
Step 13	end	Return to privileged EXEC mode.
Step 14	show clns interface <i>interface-id</i>	Verify your entries.
Step 15	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default settings, use the **no** forms of the commands.

EXAMPLE

The following configuration example for an IS-IS routing process called *area1* sets a global default metric of 111 for the IS-IS interfaces:

```
interface Ethernet3/1
ip address 172.16.10.2 255.255.0.0
ip router isis area1
no ip route-cache
duplex half
!
interface Ethernet3/2
ip address 192.168.242.2 255.255.255.0
ip router isis area1
no ip route-cache
duplex half
router isis area1
net 01.0000.0309.1234.00
metric-style wide
metric 111
```

Monitoring and Maintaining IS-IS

You can remove all contents of a CLNS cache or remove information for a particular neighbor or route. You can display specific CLNS or IS-IS statistics, such as the contents of routing tables, caches, and databases. You can also display information about specific interfaces, filters, or neighbors.

Command	Purpose
clear clns cache	Clear and reinitialize the CLNS routing cache.
clear clns es-neighbors	Remove end system (ES) neighbor information from the adjacency database.
clear clns is-neighbors	Remove intermediate system (IS) neighbor information from the adjacency database.
clear clns neighbors	Remove CLNS neighbor information from the adjacency database.
clear clns route	Remove dynamically derived CLNS routing information.
show clns	Display information about the CLNS network.
show clns cache	Display the entries in the CLNS routing cache.
show clns es-neighbors	Display ES neighbor entries, including the associated areas.
show clns filter-expr	Display filter expressions.
show clns filter-set	Display filter sets.
show clns interface <i>[interface-id]</i>	Display the CLNS-specific or ES-IS information about each interface.
show clns neighbor	Display information about IS-IS neighbors.
show clns protocol	List the protocol-specific information for each IS-IS or ISO IGRP routing process in this router.
show clns route	Display all the destinations to which this router knows how to route CLNS packets.
show clns traffic	Display information about the CLNS packets this router has seen.
show ip route isis	Display the current state of the IS-IS IP routing table.
show isis database	Display the IS-IS link-state database.
show isis routes	Display the IS-IS Level 1 routing table.
show isis spf-log	Display a history of the shortest path first (SPF) calculations for IS-IS.
show isis topology	Display a list of all connected routers in all areas.
show route-map	Display all route maps configured or only the one specified.
trace clns <i>destination</i>	Discover the paths taken to a specified destination by packets in the network.
which-route <i>{nsap-address clns-name}</i>	Display the routing table in which the specified CLNS destination is found.

Configuring BFD

The Bidirectional Forwarding Detection (BFD) Protocol quickly detects forwarding-path failures for a variety of media types, encapsulations, topologies, and routing protocols. It operates in a unicast, point-to-point mode on top of any data protocol being forwarded between two systems to track IPv4 connectivity between directly connected neighbors. BFD packets are encapsulated in UDP packets with a destination port number of 3784 or 3785.

In EIGRP, IS-IS, and OSPF deployments, the closest alternative to BFD is the use of modified failure-detection mechanisms. Although reducing the EIGRP, IS-IS, and OSPF timers can result in a failure-detection rate of 1 to 2 seconds, BFD can provide failure detection in less than 1 second. BFD can be less CPU-intensive than the reduced timers and, because it is not tied to any particular routing protocol, it can be used as a generic and consistent failure detection mechanism for multiple routing protocols.

To create a BFD session, you must configure BFD on both systems (BFD peers). Enabling BFD at the interface and routing protocol level on BFD peers creates a BFD session. BFD timers are negotiated and the BFD peers send control packets to each other at the negotiated intervals. If the neighbor is not directly connected, BFD neighbor registration is rejected.

Figure 2-6 shows a simple network with two routers running OSPF and BFD. When OSPF discovers a neighbor (1), it sends a request to the BFD process to initiate a BFD neighbor session with the neighbor OSPF router (2), establishing the BFD neighbor session (3).

Figure 2-6 Establishing a BFD Session

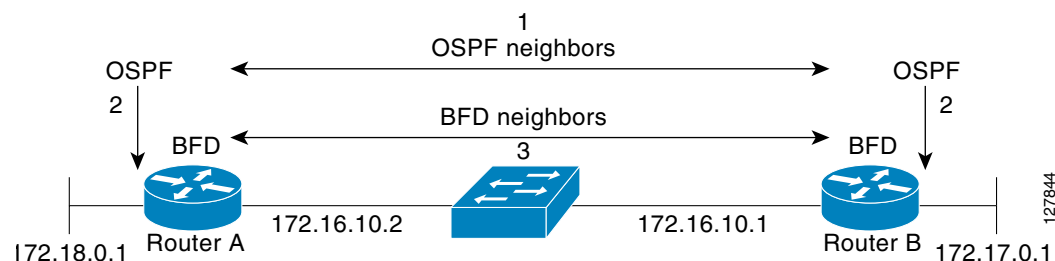
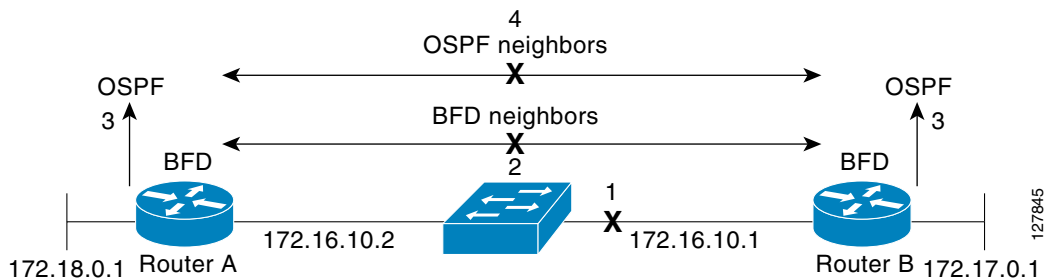


Figure 2-7 shows what happens when a failure occurs in the network (1). The BFD neighbor session with the OSPF neighbor closes (2). BFD notifies the OSPF process that the BFD neighbor is no longer reachable, and the OSPF process breaks the OSPF neighbor relationship (4). If an alternative path is available, the routers start converging on it.

Figure 2-7 Breaking an OSPF Neighbor Relationship



BFD clients are routing protocols that register neighbors with BFD. The switch supports IS-IS, OSPF v1 and v2, BGP, EIGRP, and HSRP clients. You can use one BFD session for multiple client protocols. For example, if a network is running OSPF and EIGRP across the same link to the same peer, you need to create only one BFD session, and information is shared with both routing protocols.

The switch supports BFD version 0 and version 1. BFD neighbors automatically negotiate the version and the protocol always runs at the higher version. The default version is version 1.

By default, BFD neighbors exchange both control packets and echo packets for detecting forwarding failures. The switch sends echo packets at the configured BFD interval rate (from 50 to 999 ms), and control packets at the BFD slow-timer rate (from 1000 to 3000 ms).

Failure-rate detection can be faster in BFD echo mode, which is enabled by default when you configure BFD session. In this mode, the switch sends echo packets from the BFD software layer, and the BFD neighbor responds to the echo packets through its fast-switching layer. The echo packets do not reach the BFD neighbor software layer, but are reflected back over the forwarding path for failure detection. You configure the rate at which each BFD interface sends BFD echo packets by entering the **bfd interval** interface configuration command.

To reduce bandwidth consumption, you can disable the sending of echo packets by entering the **no bfd echo** interface configuration command. When echo mode is disabled, control packets are used to detect forwarding failures. Control packets are exchanged at the configured slow-timer rate, which could result in longer failure-detection time. You configure this rate by entering the **bfd slow-timer** global configuration command. The range is from 1000 to 3000 ms; the default rate is every 1000 ms.

You can enable or disable echo processing at a switch interface independent of the BFD neighbor configuration. Disabling echo mode only disables the sending of echo packets by the interface. The fast-switching layer that receives an echo packet always reflects it back to the sender.

To run BFD on a switch, you need to configure basic BFD interval parameters on BFD interfaces, enable routing on the switch, and enable one or more one routing protocol clients for BFD. You also need to confirm that Cisco Express Forwarding (CEF) is enabled (the default) on participating switches.

For more information on the configuration and commands, see the BFD documents listed in the [“Related Documents” section on page 2-135](#).

This section includes the following topics:

- [Default BFD Configuration, page 2-89](#)
- [Default BFD Configuration Guidelines, page 2-90](#)
- [Configuring BFD Session Parameters on an Interface, page 2-90](#)
- [Enabling BFD Routing Protocol Clients, page 2-91](#)

Default BFD Configuration

- No BFD sessions are configured. BFD is disabled on all interfaces.
- When configured, BFD version 1 is the default, but switches negotiate for version. Version 0 is also supported.
- Standby BFD (for HSRP) is enabled by default.
- Asynchronous BFD echo mode is enabled when a BFD session is configured.

Default BFD Configuration Guidelines

The switch supports a maximum of 28 BFD sessions at one time.

To run BFD on a switch:

- Configure basic BFD interval parameters on each interface over which you want to run BFD sessions.
- Enable routing on the switch. You can configure BFD without enabling routing, but BFD sessions do not become active unless routing is enabled on the switch and on the BFD interfaces.
- Enable one or more one routing protocol clients for BFD. You should implement fast convergence for the routing protocol that you are using.

**Note**

We recommend that you configure the BFD interval parameters on an interface before configuring the routing protocol commands, especially when using EIGRP.

Confirm that CEF is enabled on participating switches (the default) as well as IP routing.

BFD is supported on physical interfaces that are configured as routing interfaces. It is not supported on Layer 2 interfaces, pseudowires, static routes, SVI interfaces, or port channels.

Although you can configure BFD interface commands on a Layer 2 port, BFD sessions do not operate on the interface unless it is configured as a Layer 3 interface (no switchport) and assigned an IP address.

In HSRP BFD, standby BFD is enabled globally by default and on all interfaces. If you disable it on an interface, you then must disable and reenabling it globally for BFD sessions to be active.

When using BFD echo mode (the default), you should disable sending of ICMP redirect messages by entering the **no ip redirects** interface configuration command on the BFD interface.

Configuring BFD Session Parameters on an Interface

Before you can start a BFD session on an interface, you must put the interface into Layer 3 mode and set the baseline BFD parameters on it.

**Note**

Although you can configure BFD on Layer 2 interfaces, a BFD session cannot start until both interfaces are in Layer 3 mode and routing is enabled on the switch.

BEFORE YOU BEGIN

See the [“Default BFD Configuration Guidelines”](#) section on page 2-90.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface for a BFD session, and enter interface configuration mode. Only physical interfaces support BFD.

	Command	Purpose
Step 3	no shutdown	Enable the interface if necessary. User network interfaces (UNIs) and enhanced network interfaces (ENIs) are disabled by default; network node interfaces (NNIs) are enabled by default.
Step 4	no switchport	Remove the interface from Layer 2 configuration mode.
Step 5	ip address <i>ip-address</i> <i>subnet-mask</i>	Configure the IP address and IP subnet mask.
Step 6	bfd interval <i>milliseconds</i> min_rx <i>milliseconds</i> multiplier <i>value</i>	<p>Set BFD parameters for echo packets on the interface.</p> <ul style="list-style-type: none"> interval—Specify the rate at which BFD echo packets are sent to BFD peers. The range is from 50 to 999 milliseconds (ms). min_rx—Specify the rate at which BFD echo packets are expected to be received from BFD peers. The range is from 50 to 999 ms. multiplier—Specify the number of consecutive BFD echo packets that must be missed from a BFD peer before BFD declares that it is unavailable and informs the other BFD peer of the failure. The range is from 3 to 50. <p>Note There are no baseline BFD parameter defaults.</p>
Step 7	end	Return to privileged EXEC mode.
Step 8	show running-config	Verify your entries.
Step 9	show bfd neighbor detail	(Optional) Display the final configured or negotiated values when the session is created with a neighbor.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the BFD parameter configuration, enter the **no bfd interval** interface configuration command.

EXAMPLE

```
Switch(config)# interface FastEthernet 6/0
Switch(config-if)# no switchport
Switch(config-if)# ip address 10.201.201.1 255.255.255.0
Switch(config-if)# bfd interval 50 min_rx 50 multiplier 5
Switch(config-if)# end
```

Enabling BFD Routing Protocol Clients

After you configure BFD parameters on an interface, you can start a BFD session for one or more routing protocols. You must first enable routing by entering the **ip routing** global configuration command on the switch. Note that there can be more than one way to start a BFD session on an interface, depending on the routing protocol.

- [Configuring BFD for OSPF, page 2-92](#)
- [Configuring BFD for IS-IS, page 2-93](#)
- [Configuring BFD for BGP, page 2-95](#)

- [Configuring BFD for EIGRP, page 2-97](#)
- [Configuring BFD for HSRP, page 2-98](#)

Configuring BFD for OSPF

When you start BFD sessions for OSPF, OSPF must be running on all participating devices. You can enable BFD support for OSPF by enabling it globally on all OSPF interfaces or by enabling it on one or more interfaces.

Configuring BFD for OSPF Globally

BEFORE YOU BEGIN

- Configure BFD parameters as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Configure OSPF as described in the [“Configuring OSPF” section on page 2-28](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Specify an OSPF process, and enter router configuration mode.
Step 3	bfd all-interfaces	Enable BFD globally on all interfaces associated with the OSPF routing process.
Step 4	exit	(Optional) Return to global configuration mode if you want to disable BFD on one or more OSPF interfaces.
Step 5	interface <i>interface-id</i>	(Optional) Specify an interface, and enter interface configuration mode.
Step 6	ip ospf bfd disable	(Optional) Disable BFD on the specified OSPF interface. Repeat Steps 5 and 6 for all OSPF interfaces on which you do not want to run BFD sessions.
Step 7	end	Return to privileged EXEC mode.
Step 8	show bfd neighbors [detail]	Verify the configuration.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable OSPF BFD on all interfaces, enter the **no bfd all-interfaces** router configuration command. To disable it on an interface, enter the **no ip ospf bfd** or the **ip ospf bfd disable** interface configuration command on the interface.

If you want to run OSPF BFD on only one or a few interfaces, you can enter the **ip ospf bfd** interface configuration command on those interfaces instead of enabling it globally. See the next procedure.



Note

If you try to configure OSPF BFD on a Layer 2 interface, the configuration is not recognized.

EXAMPLE

This is an example of enabling BFD for OSPF on all OSPF interfaces:

```
Switch(config)# router ospf 109
Switch(config-router)# bfd all-interfaces
Switch(config-router)# exit
```

Configuring BFD for OSPF on an Interface**BEFORE YOU BEGIN**

- Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Configure OSPF as described in the [“Configuring OSPF” section on page 2-28](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i>	Specify an OSPF process, and enter router configuration mode.
Step 3	exit	Return to global configuration mode.
Step 4	interface <i>interface-id</i>	Specify an interface, and enter interface configuration mode.
Step 5	ip ospf bfd	Enable BFD on the specified OSPF interface. Repeat Steps 3 and 4 for all OSPF interfaces on which you want to run BFD sessions.
Step 6	end	Return to privileged EXEC mode.
Step 7	show bfd neighbors [detail]	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable OSPF BFD on an interface, enter the **no ip ospf bfd** or the **ip ospf bfd disable** interface configuration command on the interface.

EXAMPLE

This is an example of enabling BFD for OSPF on a single interface:

```
Switch(config)# router ospf 109
Switch(config-router)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ip ospf bfd
```

Configuring BFD for IS-IS

When you start BFD sessions for IS-IS, IS-IS must be running on all devices participating in BFD. You can enable BFD support for IS-IS by enabling it globally on all IS-IS interfaces or by enabling it on one or more interfaces.

Configuring BFD for IS-IS Globally

BEFORE YOU BEGIN

- Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Configure IS-IS as described in the [“Configuring IS-IS Dynamic Routing” section on page 2-77](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router is-is <i>area-tag</i>	Specify an IS-IS process and enter router configuration mode.
Step 3	bfd all-interfaces	Enable BFD globally on all interfaces associated with the IS-IS routing process.
Step 4	exit	(Optional) Return to global configuration mode if you want to disable BFD on one or more IS-IS interfaces.
Step 5	interface <i>interface-id</i>	(Optional) Specify an interface and enter interface configuration mode.
Step 6	ip router isis	(Optional) Enable IPv4 IS-IS routing on the interface.
Step 7	isis bfd disable	(Optional) Disable BFD on the IS-IS interface. Repeat Steps 5 through 7 for all IS-IS interfaces on which you do not want to run BFD sessions.
Step 8	end	Return to privileged EXEC mode.
Step 9	show bfd neighbors [detail]	Verify the configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IS-IS BFD on all interfaces, enter the **no bfd all-interfaces** router configuration command. To disable it on the specified interface, enter the **no isis bfd** or the **isis bfd disable** interface configuration command on the interface.

If you only want to run IS-IS BFD on a few interfaces, instead of enabling it globally, you can enter the **isis bfd** interface configuration command on those interfaces. See the next procedure.



Note

Although IS-IS BFD operates only on Layer 3 interfaces, you can configure it on interfaces in Layer 2 or Layer 3 mode. When you enable it, you see this message:

```
%ISIS BFD is reverting to router mode configuration, and remains disabled.
```

EXAMPLE

This is an example of setting fast convergence and enabling BFD for IS-IS on all IS-IS interfaces:

```
Switch(config)# router is-is tag1
Switch(config-router)# bfd all-interfaces
Switch(config-router)# exit
```

Configuring BFD for IS-IS on an Interface

BEFORE YOU BEGIN

- Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Configure IS-IS as described in the [“Configuring IS-IS Dynamic Routing” section on page 2-77](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router is-is <i>area-tag</i>	Specify an IS-IS process and enter router configuration mode.
Step 3	exit	Return to global configuration mode.
Step 4	interface <i>interface-id</i>	Specify an interface, and enter interface configuration mode.
Step 5	isis bfd	Enable BFD on the specified IS-IS interface. Repeat Steps 3 and 4 for all IS-IS interfaces on which you want to run BFD sessions.
Step 6	end	Return to privileged EXEC mode.
Step 7	show bfd neighbors [detail]	Verify the configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IS-IS BFD on an interface, enter the **no isis bfd** or the **isis bfd disable** interface configuration command on the interface.

EXAMPLE

This is an example of enabling BFD for IS-IS on a single interface:

```
Switch(config)# router is-is tag1
Switch(config-router)# exit
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# isis bfd
```

Configuring BFD for BGP

When you start BFD sessions for BGP, BGP must be running on all participating devices. You enter the IP address of the BFD neighbor to enable BFD for BGP.

BEFORE YOU BEGIN

- Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Configure BGP as described in the [“Configuring BGP” section on page 2-51](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>as-tag</i>	Specify a BGP autonomous system, and enter router configuration mode.
Step 3	neighbor <i>ip-address</i> fall-over bfd	Enable BFD support for fallover on the BFD neighbor.
Step 4	end	Return to privileged EXEC mode.
Step 5	show bfd neighbors [detail] > show ip bgp neighbor	Verify the configuration. Display information about BGP connections to neighbors.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable BGP BFD, enter the **no neighbor** *ip-address* **fall-over bfd** router configuration command.

EXAMPLE

The following example shows how to configure BFD in a BGP network. In the following example, the simple BGP network consists of Router A and Router B. Fast Ethernet interface 0/1 on Router A is connected to the same network as Fast Ethernet interface 6/0 in Router B. The example, starting in global configuration mode, shows the configuration of BFD.

Router A:

```
!
interface Fast Ethernet 0/1
ip address 172.16.10.1 255.255.255.0
bfd interval 50 min_rx 50 multiplier 3
!
interface Fast Ethernet 3/0.1
ip address 172.17.0.1 255.255.255.0
!!
router bgp 40000
bgp log-neighbor-changes
neighbor 172.16.10.2 remote-as 45000
neighbor 172.16.10.2 fall-over bfd
!
address-family ipv4
neighbor 172.16.10.2 activate
no auto-summary
no synchronization
network 172.18.0.0 mask 255.255.255.0
exit-address-family
!
```

Router B:

```
!
interface Fast Ethernet 6/0
ip address 172.16.10.2 255.255.255.0
bfd interval 50 min_rx 50 multiplier 3
!
interface Fast Ethernet 6/1
ip address 172.18.0.1 255.255.255.0
```

```

!
router bgp 45000
  bgp log-neighbor-changes
  neighbor 172.16.10.1 remote-as 40000
  neighbor 172.16.10.1 fall-over bfd
!
address-family ipv4
  neighbor 172.16.10.1 activate
  no auto-summary
  no synchronization
  network 172.17.0.0 mask 255.255.255.0
exit-address-family
!

```

Configuring BFD for EIGRP

When you start BFD sessions for EIGRP, EIGRP must be running on all participating devices. You can enable BFD support for EIGRP by globally enabling it on all EIGRP interfaces or by enabling it on one or more interfaces.

BEFORE YOU BEGIN

- Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Configure EIGRP as described in the [“Configuring EIGRP” section on page 2-42](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router eigrp <i>as-number</i>	Specify an EIGRP autonomous system number, and enter router configuration mode.
Step 3	log-adjacency changes [detail]	Configure the switch to send a system logging message when an EIGRP neighbor goes up or down.
Step 4	bfd { all-interfaces interface <i>interface-id</i> }	Enable BFD for EIGRP. <ul style="list-style-type: none"> • Enter all-interfaces to globally enable BFD on all interfaces associated with the EIGRP routing process. • Enter interface <i>interface-id</i> to enable BFD on a per-interface basis for one or more interfaces associated with the EIGRP routing process.
Step 5	end	Return to privileged EXEC mode.
Step 6	show bfd neighbors [detail]	Verify the configuration.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable EIGRP BFD on all interfaces, enter the **no bfd all-interfaces** router configuration command. To disable it on an interface, enter the **no bfd interface** *interface-id* router configuration command.

EXAMPLE

The following example shows how to enable BFD for all EIGRP neighbors, using the `bfd` command in address family interface configuration mode:

```
Switch# configure terminal
Switch(config)# router eigrp my_eigrp
Switch(config-router)# address family ipv4 autonomous-system 100
Switch(config-router-af)# af-interface FastEthernet 0/0
Switch(config-router-af)# bfd
```

Configuring BFD for HSRP

HSRP supports BFD by default; it is globally enabled on all interfaces. If HSRP support has been manually disabled, you can reenable it in interface or global configuration mode.

BEFORE YOU BEGIN

- Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).
- Ensure that all participating devices have HSRP enabled and CEF enabled (the default).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Specify an interface for a BFD session, and enter interface configuration mode. Only physical interfaces support BFD.
Step 3	ip address <i>ip-address</i> <i>subnet-mask</i>	Configure the IP address and IP subnet mask for the interface.
Step 4	standby [<i>group-number</i>] ip [<i>ip-address</i>] [secondary]	Activate HSRP.
Step 5	standby bfd	(Optional) Enable HSRP support for BFD on the interface.
Step 6	exit	Return to global configuration mode.
Step 7	standby bfd all-interfaces	(Optional) Enable HSRP support for BFD on all interfaces.
Step 8	end	Return to privileged EXEC mode.
Step 9	show standby neighbors	Verify your entries.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable HSRP support for BFD on all interfaces, enter the **no standby bfd all-interfaces** global configuration command. To disable it on an interface, enter the **no standby bfd** interface configuration command.

**Note**

If you disable standby BFD on an interface by entering the **no standby bfd** interface configuration command, to activate BFD sessions on other interfaces, you must disable and reenabling it globally by entering the **no standby bfd all-interfaces** global configuration command followed by the **standby bfd all-interfaces** global configuration command.

EXAMPLE

The following example shows how to reenabling HSRP BFD peering if it has been disabled on a switch:

```
Switch(config)# standby bfd all-interfaces
```

Disabling BFD Echo Mode

When you configure a BFD session, BFD echo mode is enabled by default on BFD interfaces. You can disable echo mode on an interface so it sends no echo packets and but only sends back echo packets received from a neighbor. When echo mode is disabled, control packets are used to detect forwarding failures. You can configure slow timers to reduce the frequency of BFD control packets.

BEFORE YOU BEGIN

Configure BFD parameters on the interface as described in the [“Configuring BFD Session Parameters on an Interface” procedure on page 2-90](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter a BFD interface and enter interface configuration mode.
Step 3	no bfd echo	Disable BFD echo mode on the interface. It is enabled by default, but can be disabled independently on BFD neighbors.
Step 4	exit	Return to global configuration mode.
Step 5	bfd slow-timer [<i>milliseconds</i>]	(Optional) Configure a BFD slow-timer value. The range is from 1000 to 30000 milliseconds. The default is 1000 milliseconds.
Step 6	end	Return to privileged EXEC mode.
Step 7	show bfd neighbors detail	Verify your entries.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To reenabling echo mode on the switch, enter the **bfd echo** global configuration command.

EXAMPLE

The following example disables echo mode between BFD neighbors:

```
Switch# configure terminal
Switch(config)# interface Ethernet 0/1
Switch(config-if)# no bfd echo
```

Configuring Multi-VRF CE

Virtual Private Networks (VPNs) provide a secure way for customers to share bandwidth over an ISP backbone network. A VPN is a collection of sites sharing a common routing table. A customer site is connected to the service-provider network by one or more interfaces, and the service provider associates each interface with a VPN routing table, called a VPN routing/forwarding (VRF) table.

The switch supports multiple VPN routing/forwarding (multi-VRF) instances in customer edge (CE) devices (multi-VRF CE). With multi-VRF CE, a service provider can support two or more VPNs with overlapping IP addresses.

**Note**

The switch does not use Multiprotocol Label Switching (MPLS) to support VPNs. For information about MPLS VRF, refer to the [MPLS: Layer 3 VPNs Configuration Guide, Cisco IOS Release 15M&T](#).

- [Information About Multi-VRF CE, page 2-100](#)
- [Default Multi-VRF CE Configuration, page 2-102](#)
- [Multi-VRF CE Configuration Guidelines, page 2-102](#)
- [Configuring VRFs, page 2-103](#)
- [Configuring VRF-Aware Services, page 2-104](#)
- [Configuring a VPN Routing Session, page 2-110](#)
- [Configuring BGP PE to CE Routing Sessions, page 2-111](#)
- [Multi-VRF CE Configuration Example, page 2-112](#)
- [Displaying Multi-VRF CE Status, page 2-116](#)

Information About Multi-VRF CE

Multi-VRF CE allows a service provider to support two or more VPNs, where IP addresses can be overlapped among the VPNs. Multi-VRF CE uses input interfaces to distinguish routes for different VPNs and forms virtual packet-forwarding tables by associating one or more Layer 3 interfaces with each VRF. Interfaces in a VRF can be either physical, such as Ethernet ports, or logical, such as VLAN SVIs, but an interface cannot belong to more than one VRF at any time.

**Note**

Multi-VRF CE interfaces must be Layer 3 interfaces.

Multi-VRF CE includes these devices:

- Customer edge (CE) devices provide customers access to the service-provider network over a data link to one or more provider edge routers. The CE device advertises the site local routes to the router and learns the remote VPN routes from it. The Cisco Connected Grid switch can be a CE.
- Provider edge (PE) routers exchange routing information with CE devices by using static routing or a routing protocol such as BGP, RIPv2, OSPF, or EIGRP. The PE is only required to maintain VPN routes for those VPNs to which it is directly attached, eliminating the need for the PE to maintain all of the service-provider VPN routes. Each PE router maintains a VRF for each of its directly connected sites. Multiple interfaces on a PE router can be associated with a single VRF if all of these

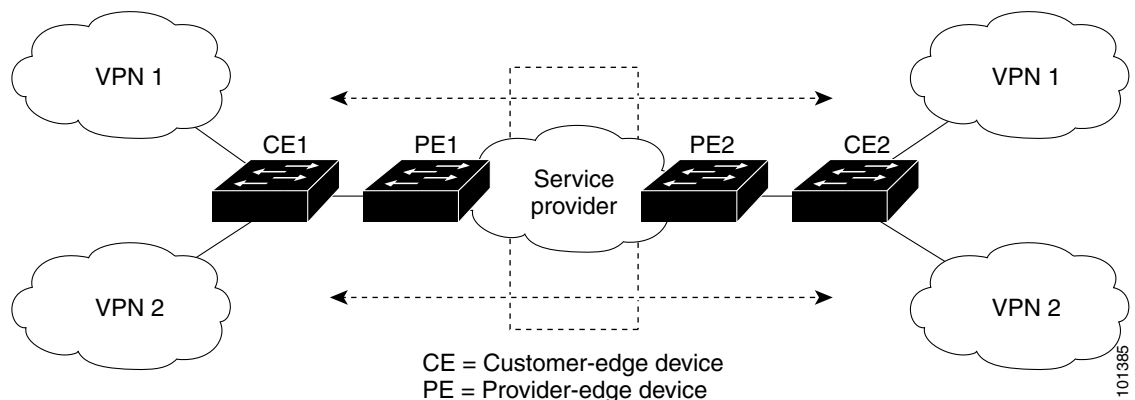
sites participate in the same VPN. Each VPN is mapped to a specified VRF. After learning local VPN routes from CEs, a PE router exchanges VPN routing information with other PE routers by using internal BGP (IBPG).

- Provider routers or core routers are any routers in the service provider network that do not attach to CE devices.

With multi-VRF CE, multiple customers can share one CE, and only one physical link is used between the CE and the PE. The shared CE maintains separate VRF tables for each customer and switches or routes packets for each customer based on its own routing table. Multi-VRF CE extends limited PE functionality to a CE device, giving it the ability to maintain separate VRF tables to extend the privacy and security of a VPN to the branch office.

Figure 2-8 shows a configuration using Cisco Connected Grid switches as multiple virtual CEs. This scenario is suited for customers who have low bandwidth requirements for their VPN service, for example, small companies. In this case, multi-VRF CE support is required in the Cisco Connected Grid switches. Because multi-VRF CE is a Layer 3 feature, each interface in a VRF must be a Layer 3 interface.

Figure 2-8 Switches Acting as Multiple Virtual CEs



When the CE switch receives a command to add a Layer 3 interface to a VRF, it sets up the appropriate mapping between the VLAN ID and the policy label (PL) in multi-VRF-CE-related data structures and adds the VLAN ID and PL to the VLAN database.

When multi-VRF CE is configured, the Layer 3 forwarding table is conceptually partitioned into two sections:

- The multi-VRF CE routing section contains the routes from different VPNs.
- The global routing section contains routes to non-VPN networks, such as the Internet.

VLAN IDs from different VRFs are mapped into different policy labels, which are used to distinguish the VRFs during processing. If no route is found in the multi-VRF CE section of the Layer 3 forwarding table, the global routing section is used to determine the forwarding path. For each new VPN route learned, the Layer 3 setup function retrieves the policy label by using the VLAN ID of the ingress port and inserts the policy label and new route to the multi-VRF CE routing section. If the packet is received from a routed port, the port internal VLAN ID number is used; if the packet is received from an SVI, the VLAN number is used.

This is the packet-forwarding process in a multi-VRF-CE-enabled network:

- When the switch receives a packet from a VPN, the switch looks up the routing table based on the input policy label number. When a route is found, the switch forwards the packet to the PE.
- When the ingress PE receives a packet from the CE, it performs a VRF lookup. When a route is found, the router adds a corresponding MPLS label to the packet and sends it to the MPLS network.
- When an egress PE receives a packet from the network, it strips the label and uses the label to identify the correct VPN routing table. Then it performs the normal route lookup. When a route is found, it forwards the packet to the correct adjacency.
- When a CE receives a packet from an egress PE, it uses the input policy label to look up the correct VPN routing table. If a route is found, it forwards the packet within the VPN.

To configure VRF, you create a VRF table and specify the Layer 3 interface associated with the VRF. Then configure the routing protocols in the VPN and between the CE and the PE. BGP is the preferred routing protocol used to distribute VPN routing information across the provider's backbone. The multi-VRF CE network has three major components:

- VPN route target communities—lists of all other members of a VPN community. You need to configure VPN route targets for each VPN community member.
- Multiprotocol BGP peering of VPN community PE routers—propagates VRF reachability information to all members of a VPN community. You need to configure BGP peering in all PE routers within a VPN community.
- VPN forwarding—transports all traffic between all VPN community members across a VPN service-provider network.

Default Multi-VRF CE Configuration

Feature	Default Setting
VRF	Disabled. No VRFs are defined.
Maps	No import maps, export maps, or route maps are defined.
VRF maximum routes	5000
Forwarding table	The default for an interface is the global routing table.

Multi-VRF CE Configuration Guidelines

These are considerations when configuring VRF in your network:

- A switch with multi-VRF CE is shared by multiple customers, and each customer has its own routing table.
- Because customers use different VRF tables, the same IP addresses can be reused. Overlapped IP addresses are allowed in different VPNs.
- Multi-VRF CE lets multiple customers share the same physical link between the PE and the CE. Trunk ports with multiple VLANs separate packets among customers. Each customer has its own VLAN.
- Multi-VRF CE does not support all MPLS-VRF functionality. It does not support label exchange, LDP adjacency, or labeled packets.

- For the PE router, there is no difference between using multi-VRF CE or using multiple CEs. In [Figure 2-8](#), multiple virtual Layer 3 interfaces are connected to the multi-VRF CE device.
- The switch supports configuring VRF by using physical ports, VLAN SVIs, or a combination of both. The SVIs can be connected through an access port or a trunk port.
- A customer can use multiple VLANs as long as they do not overlap with those of other customers. A customer's VLANs are mapped to a specific routing table ID that is used to identify the appropriate routing tables stored on the switch.
- The switch supports one global network and up to 26 VRFs.
- Most routing protocols (BGP, OSPF, RIP, EIGRP, and static routing) can be used between the CE and the PE. However, we recommend using external BGP (EBGP) for these reasons:
 - BGP does not require multiple algorithms to communicate with multiple CEs.
 - BGP is designed for passing routing information between systems run by different administrations.
 - BGP makes it easy to pass attributes of the routes to the CE.
- Multi-VRF CE does not affect the packet switching rate.
- If no VRFs are configured, up to 105 policies can be configured.
- If even one VRF is configured then 41 policies can be configured.
- If more than 41 policies are configured then VRF cannot be configured.
- VRF and private VLANs are mutually exclusive. You cannot enable VRF on a private VLAN. Similarly, you cannot enable private VLAN on a VLAN with VRF configured on the VLAN interface.
- VRF and policy-based routing (PBR) are mutually exclusive on a switch interface. You cannot enable VRF when PBR is enabled on an interface. In contrast, you cannot enable PBR when VRF is enabled on an interface.

Configuring VRFs

Follow the steps in this procedure to configure one or more VRFs.

BEFORE YOU BEGIN

See the [“Multi-VRF CE Configuration Guidelines”](#) section on page 2-102.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IP routing.
Step 3	ip vrf <i>vrf-name</i>	Name the VRF, and enter VRF configuration mode.
Step 4	rd <i>route-distinguisher</i>	Create a VRF table by specifying a route distinguisher. Enter either an AS number and an arbitrary number (xxx:y) or an IP address and arbitrary number (A.B.C.D:y).

	Command	Purpose
Step 5	route-target { export import both } <i>route-target-ext-community</i>	Create a list of import, export, or import and export route target communities for the specified VRF. Enter either an AS system number and an arbitrary number (xxx:y) or an IP address and an arbitrary number (A.B.C.D:y). The <i>route-target-ext-community</i> should be the same as the <i>route-distinguisher</i> entered in Step 4.
Step 6	import map <i>route-map</i>	(Optional) Associate a route map with the VRF.
Step 7	interface <i>interface-id</i>	Specify the Layer 3 interface to be associated with the VRF, and enter interface configuration mode. The interface can be a routed port or SVI.
Step 8	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 9	ip vrf forwarding <i>vrf-name</i>	Associate the VRF with the Layer 3 interface.
Step 10	end	Return to privileged EXEC mode.
Step 11	show ip vrf [brief detail interfaces] [<i>vrf-name</i>]	Verify the configuration. Display information about the configured VRFs.
Step 12	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip vrf vrf-name** global configuration command to delete a VRF and to remove all interfaces from it. Use the **no ip vrf forwarding** interface configuration command to remove an interface from the VRF.

EXAMPLE

The following example shows how to import a route map to a VRF instance named VPN1:

```
Switch(config)# ip vrf vpn1
Switch(config-vrf)# rd 100:2
Switch(config-vrf)# route-target both 100:2
Switch(config-vrf)# route-target import 100:1
```

Configuring VRF-Aware Services

IP services can be configured on global interfaces, and these services run within the global routing instance. IP services are enhanced to run on multiple routing instances; they are VRF-aware. Any configured VRF in the system can be specified for a VRF-aware service.

VRF-aware services are implemented in platform-independent modules. VRF means multiple routing instances in Cisco IOS. Each platform has its own limit on the number of VRFs it supports.

VRF-aware services have the following characteristics:

- The user can ping a host in a user-specified VRF.
- ARP entries are learned in separate VRFs. The user can display Address Resolution Protocol (ARP) entries for specific VRFs.

These services are VRF-aware:

- ARP
- Ping

- Simple Network Management Protocol (SNMP)
- Hot Standby Router Protocol (HSRP)
- Syslog
- Traceroute
- FTP and TFTP

**Note**

VRF-aware services are not supported for Unicast Reverse Path Forwarding (uRPF).

User Interface for ARP

Use the **arp** command in global configuration mode to add a VRF to the ARP cache.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	arp vrf <i>vrf-name</i> <i>hardware-address</i> <i>encap-type</i> [<i>interface-type</i>] [<i>alias</i>]	Add a VRF instance. The <i>vrf-name</i> argument is the name of the VRF table.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

```
switch(config)# arp vrf vpn1 0800.0900.1834
```

User Interface for PING

To check if a configured VRF is working, you can use the **ping vrf** command.

When attempting to ping from a provider edge (PE) router to a customer edge (CE) router, or from a PE router to PE router, the standard ping command will not usually work. The **ping vrf** command allows you to ping the IP addresses of LAN interfaces on CE routers.

If you are on a PE router, be sure to indicate the specific VRF (VPN) name, as shown in the “Examples” section.

If all required information is not provided at the command line, the system will enter the interactive dialog (extended mode) for ping.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

	Command	Purpose
Step 1	ping vrf <i>vrf-name</i> ip-host	Tests a connection in the context of a specific VPN connection.

EXAMPLE

In the following example, the target host in the domain 209.165.201.1 is pinged (using IP/ICMP) in the context of the “CustomerA” VPN connection:

```
Switch# ping vrf CustomerA 209.165.201.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 209.165.201.1, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 176/264/576 ms
```

User Interface for SNMP

Follow the steps in this procedure to configure VRF-aware services for SNMP.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	snmp-server trap authentication vrf	Enable VRF instance context authentication notifications.
Step 3	snmp-server engineID remote <host> vrf <vpn instance> <engine-id string>	Configure a name for the remote SNMP engine on a switch.
Step 4	snmp-server host <host> vrf <vpn instance> traps <community>	Specify the recipient of an SNMP trap operation and specify the VRF table to be used for sending SNMP traps.
Step 5	snmp-server host <host> vrf <vpn instance> informs <community>	Specify the recipient of an SNMP inform operation and specify the VRF table to be used for sending SNMP informs.
Step 6	snmp-server user <user> <group> remote <host> vrf <vpn instance> <security model>	Add a user to an SNMP group for a remote host on a VRF for SNMP access.
Step 7	end	Return to privileged EXEC mode.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

The following example specifies the SNMP engine ID and configures the VRF name traps-vrf for SNMP communications with the remote device at 172.16.20.3:

```
Switch(config)# snmp-server engineID remote 172.16.20.3 vrf trap-vrf
80000009030000B064EFE100
```

The following example shows how to send all SNMP notifications to example.com over the VRF named trap-vrf using the community string public:

```
Switch(config)# snmp-server host example.com vrf trap-vrf public
```

User Interface for HSRP

Hot Standby Router Protocol (HSRP) support for VRFs ensures that HSRP virtual IP addresses are added to the correct IP routing table.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 3	no switchport	Remove the interface from Layer 2 configuration mode if it is a physical interface.
Step 4	ip vrf forwarding < <i>vrf-name</i> >	Configure VRF on the interface. Executing this command on an interface removes the IP address.
Step 5	ip address <i>ip address</i>	Enter the IP address for the interface.
Step 6	standby 1 ip <i>ip address</i>	Enable HSRP and configure the virtual IP address.
Step 7	end	Return to privileged EXEC mode.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

```
Switch(config)# interface ethernet 0
Switch(config-if)# no switchport
Switch(config-if)# ip vrf forwarding vpn1
Switch(config-if)# ip address 172.16.1.3
Switch(config-if)# standby 1 ip
```

User Interface for Syslog

Follow the steps in this procedure to configure VRF-aware services for Syslog.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	logging on	Enable or temporarily disable logging of storage router event message.
Step 3	logging host <i>ip address</i> vrf <i>vrf name</i>	Specify the host address of the syslog server where logging messages are to be sent.
Step 4	logging buffered <i>logging buffered size</i> debugging	Log messages to an internal buffer.
Step 5	logging trap debugging	Limit the logging messages sent to the syslog server.
Step 6	logging facility <i>facility</i>	Send system logging messages to a logging facility.
Step 7	end	Return to privileged EXEC mode.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

The following example specifies a VRF that connects to the syslog server host:

```
Switch(config)# logging host 192.168.200.225 vrf vpn1
```

User Interface for Traceroute

Follow the steps in this procedure to find the destination address in a VRF.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

Command	Purpose
traceroute vrf <i>vrf-name</i> <i>ipaddress</i>	Specify the name of a VPN VRF in which to find the destination address.

EXAMPLE

The following example displays output of the traceroute command with the vrf keyword. Output includes the incoming VRF name/tag and the outgoing VRF name/tag.

```
Switch# traceroute vrf red 10.0.10.12
Type escape sequence to abort.
Tracing the route to 10.0.10.12
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.13.15 (red/13,red/13) 0 msec
   10.1.16.16 (red/13,red/13) 0 msec
   10.1.13.15 (red/13,red/13) 1 msec
 2 10.1.8.13 (red/13,red/13) 0 msec
   10.1.7.13 (red/13,red/13) 0 msec
```

```

10.1.8.13 (red/13,red/13) 0 msec
3 10.1.2.11 (red/13,blue/10) 1 msec 0 msec 0 msec
4 * * *

```

User Interface for FTP and TFTP

FTP and TFTP are VRF-aware, which means that file transfer is supported across an interface within a VRF instance. To specify a VRF as a source for FTP or TFTP connections, the VRF must be associated with the same interface that you configure with the **ip ftp source-interface** command. In this configuration, FTP looks for the destination IP address for file transfer in the specified VRF table. If the specified source interface is not up, Cisco IOS software selects the address of the interface closest to the destination as the source address.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip ftp source-interface <i>interface-type interface-number</i>	Specify the source IP address for FTP connections.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To specify the IP address of an interface as the source address for TFTP connections, use the **ip tftp source-interface** show mode command. To return to the default, use the **no** form of this command.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip tftp source-interface <i>interface-type interface-number</i>	Specify the source IP address for TFTP connections.
Step 3	end	Return to privileged EXEC mode.
Step 4	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

The following example shows how to configure the switch to use the VRF table named vpn1 to look for the destination IP address for the transfer of FTP packets:

```

Switch# configure terminal
Switch(config)# ip ftp source-interface ethernet 0
Switch(config)# ip vrf vpn1
Switch(config-vrf)# rd 200:1
Switch(config-vrf)# route-target both 200:1
Switch(config-vrf)# interface ethernet 0
Switch(config-if)# ip vrf forwarding vpn1
Switch(config-if)# end

```

User Interface for VRF-Aware RADIUS

To configure VRF-aware RADIUS, you must first enable AAA on a RADIUS server. The switch supports the **ip vrf forwarding vrf-name** server-group configuration and the **ip radius source-interface** global configuration commands.

Configuring a VPN Routing Session

Routing within the VPN can be configured with any supported routing protocol (RIP, OSPF, EIGRP, or BGP) or with static routing. The configuration shown here is for OSPF, but the process is the same for other protocols.



Note

To configure an EIGRP routing process to run within a VRF instance, you must configure an autonomous-system number by entering the **autonomous-system** *autonomous-system-number* address-family configuration mode command.

BEFORE YOU BEGIN

Configure a VRF as described in the [“Configuring VRFs” procedure on page 2-103](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router ospf <i>process-id</i> vrf <i>vrf-name</i>	Enable OSPF routing, specify a VPN forwarding table, and enter router configuration mode.
Step 3	log-adjacency-changes	(Optional) Log changes in the adjacency state. This is the default state.
Step 4	redistribute bgp <i>autonomous-system-number</i> subnets	Set the switch to redistribute information from the BGP network to the OSPF network.
Step 5	network <i>network-number</i> area <i>area-id</i>	Define a network address and mask on which OSPF runs and the area ID for that network address.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip ospf <i>process-id</i>	Verify the configuration of the OSPF network.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router ospf** *process-id* **vrf** *vrf-name* global configuration command to disassociate the VPN forwarding table from the OSPF routing process.

EXAMPLE

This example shows a basic OSPF configuration using the **router ospf** command to configure OSPF VRF processes for the VRFs first, second, and third:

```
Switch# configure terminal
Switch(config)# router ospf 12 vrf first
Switch(config)# router ospf 13 vrf second
Switch(config)# router ospf 14 vrf third
Switch(config)# exit
```

Configuring BGP PE to CE Routing Sessions

BEFORE YOU BEGIN

- Complete the BGP network strategy and planning for your network.
- Configure OSPF as described in the [“Configuring OSPF”](#) section on page 2-28.
- Configure a VRF as described in the [“Configuring VRFs”](#) procedure on page 2-103.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>autonomous-system-number</i>	Configure the BGP routing process with the AS number passed to other BGP routers, and enter router configuration mode.
Step 3	network network-number mask <i>network-mask</i>	Specify a network and mask to announce using BGP.
Step 4	redistribute ospf process-id match internal	Set the switch to redistribute OSPF internal routes.
Step 5	network network-number area <i>area-id</i>	Define a network address and mask on which OSPF runs and the area ID for that network address.
Step 6	address-family ipv4 vrf vrf-name	Define BGP parameters for PE to CE routing sessions, and enter VRF address-family mode.
Step 7	neighbor address remote-as <i>as-number</i>	Define a BGP session between PE and CE routers.
Step 8	neighbor address activate	Activate the advertisement of the IPv4 address family.
Step 9	end	Return to privileged EXEC mode.
Step 10	show ip bgp [ipv4] [neighbors]	Verify BGP configuration.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no router bgp autonomous-system-number** global configuration command to delete the BGP routing process. Use the command with keywords to delete routing characteristics.

EXAMPLE

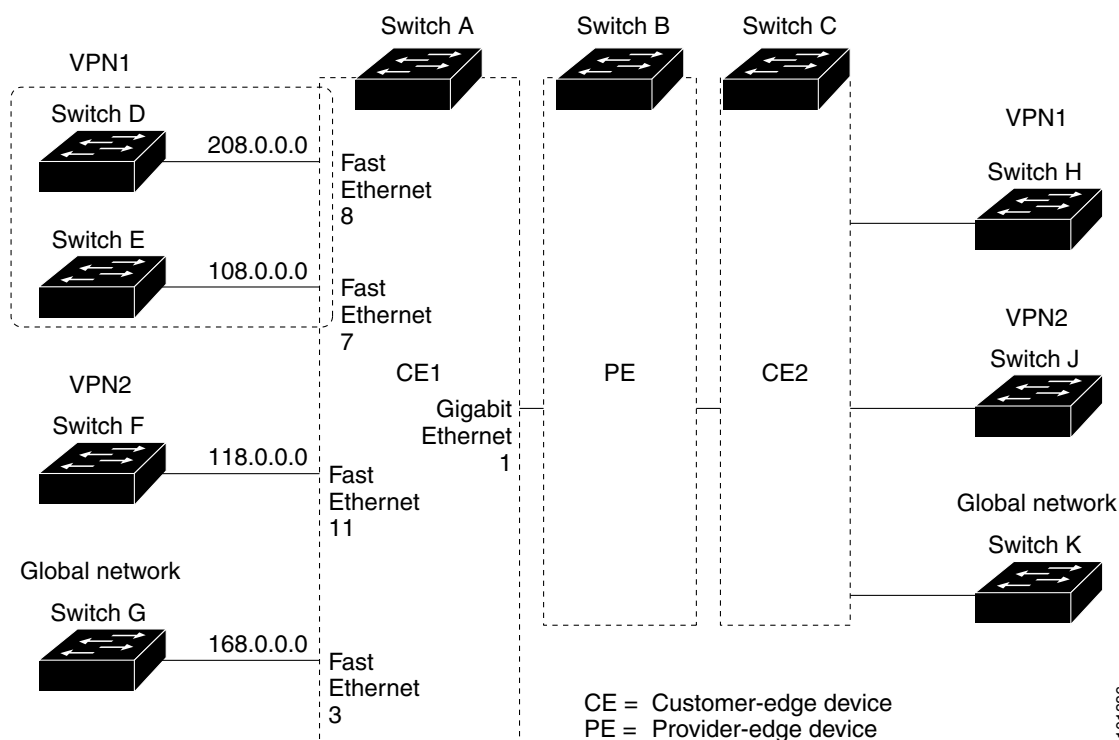
The following example configures BGP for CE to PE routing:

```
Switch(config)# router bgp 800
Switch(config-router)# address-family ipv4 vrf v12
Switch(config-router-af)# redistribute ospf 2 match internal
Switch(config-router-af)# neighbor 83.0.0.3 remote-as 100
Switch(config-router-af)# neighbor 83.0.0.3 activate
Switch(config-router-af)# network 8.8.2.0 mask 255.255.255.0
Switch(config-router-af)# exit
Switch(config-router)# address-family ipv4 vrf v11
Switch(config-router-af)# redistribute ospf 1 match internal
Switch(config-router-af)# neighbor 38.0.0.3 remote-as 100
Switch(config-router-af)# neighbor 38.0.0.3 activate
Switch(config-router-af)# network 8.8.1.0 mask 255.255.255.0
Switch(config-router-af)# end
```

Multi-VRF CE Configuration Example

Figure 2-9 is a simplified example of the physical connections in a network similar to that in Figure 2-8. OSPF is the protocol used in VPN1, VPN2, and the global network. BGP is used in the CE to PE connections. The examples following the illustration show how to configure a Cisco Connected Grid switch as CE Switch A, and the VRF configuration for customer switches D and F. Commands for configuring CE Switch C and the other customer switches are not included but would be similar. The example also includes commands for configuring traffic to Switch A for a Catalyst 6000 or Catalyst 6500 switch acting as a PE router.

Figure 2-9 Multi-VRF CE Configuration Example



101386

Configuring Switch A

On Switch A, enable routing and configure VRF:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# ip vrf v11
Switch(config-vrf)# rd 800:1
Switch(config-vrf)# route-target export 800:1
Switch(config-vrf)# route-target import 800:1
Switch(config-vrf)# exit
Switch(config)# ip vrf v12
Switch(config-vrf)# rd 800:2
Switch(config-vrf)# route-target export 800:2
Switch(config-vrf)# route-target import 800:2
Switch(config-vrf)# exit
```

Configure the loopback and physical interfaces on Switch A. Gigabit Ethernet port 1 is a trunk connection to the PE. Fast Ethernet ports 8 and 11 connect to VPNs:

```
Switch(config)# interface loopback1
Switch(config-if)# ip vrf forwarding v11
Switch(config-if)# ip address 8.8.1.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface loopback2
Switch(config-if)# ip vrf forwarding v12
Switch(config-if)# ip address 8.8.2.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface gigabitethernet0/5
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# no ip address
Switch(config-if)# exit

Switch(config)# interface fastethernet0/8
Switch(config-if)# no shutdown
Switch(config-if)# switchport access vlan 208
Switch(config-if)# no ip address
Switch(config-if)# exit

Switch(config)# interface fastethernet0/11
Switch(config-if)# no shutdown
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# no ip address
Switch(config-if)# exit
```

Configure the VLANs used on Switch A. VLAN 10 is used by VRF 11 between the CE and the PE. VLAN 20 is used by VRF 12 between the CE and the PE. VLANs 118 and 208 are used for the VPNs that include Switch F and Switch D, respectively:

```
Switch(config)# interface vlan10
Switch(config-if)# ip vrf forwarding v11
Switch(config-if)# ip address 38.0.0.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface vlan20
Switch(config-if)# ip vrf forwarding v12
Switch(config-if)# ip address 83.0.0.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface vlan118
Switch(config-if)# ip vrf forwarding v12
Switch(config-if)# ip address 118.0.0.8 255.255.255.0
Switch(config-if)# exit

Switch(config)# interface vlan208
Switch(config-if)# ip vrf forwarding v11
Switch(config-if)# ip address 208.0.0.8 255.255.255.0
Switch(config-if)# exit
```

Configure OSPF routing in VPN1 and VPN2:

```
Switch(config)# router ospf 1 vrf v11
Switch(config-router)# redistribute bgp 800 subnets
Switch(config-router)# network 208.0.0.0 0.0.0.255 area 0
Switch(config-router)# exit
Switch(config)# router ospf 2 vrf v12
Switch(config-router)# redistribute bgp 800 subnets
```

```
Switch(config-router)# network 118.0.0.0 0.0.0.255 area 0
Switch(config-router)# exit
```

Configure BGP for CE to PE routing:

```
Switch(config)# router bgp 800
Switch(config-router)# address-family ipv4 vrf v12
Switch(config-router-af)# redistribute ospf 2 match internal
Switch(config-router-af)# neighbor 83.0.0.3 remote-as 100
Switch(config-router-af)# neighbor 83.0.0.3 activate
Switch(config-router-af)# network 8.8.2.0 mask 255.255.255.0
Switch(config-router-af)# exit
```

```
Switch(config-router)# address-family ipv4 vrf v11
Switch(config-router-af)# redistribute ospf 1 match internal
Switch(config-router-af)# neighbor 38.0.0.3 remote-as 100
Switch(config-router-af)# neighbor 38.0.0.3 activate
Switch(config-router-af)# network 8.8.1.0 mask 255.255.255.0
Switch(config-router-af)# end
```

Configuring Switch D

Switch D belongs to VPN 1. Configure the connection to Switch A by using these commands:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# interface fastethernet0/2
Switch(config-if)# no shutdown
Switch(config-if)# no switchport
Switch(config-if)# ip address 208.0.0.20 255.255.255.0
Switch(config-if)# exit

Switch(config)# router ospf 101
Switch(config-router)# network 208.0.0.0 0.0.0.255 area 0
Switch(config-router)# end
```

Configuring Switch F

Switch F belongs to VPN 2. Configure the connection to Switch A by using these commands:

```
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# ip routing
Switch(config)# interface fastethernet0/1
Switch(config-if)# no shutdown
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# no ip address
Switch(config-if)# exit

Switch(config)# interface vlan118
Switch(config-if)# ip address 118.0.0.11 255.255.255.0
Switch(config-if)# exit

Switch(config)# router ospf 101
Switch(config-router)# network 118.0.0.0 0.0.0.255 area 0
Switch(config-router)# end
```

Configuring the PE Switch B

On Switch B (the PE router), these commands configure only the connections to the CE device, Switch A:

```
Router# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Router(config)# ip vrf v1
Router(config-vrf)# rd 100:1
Router(config-vrf)# route-target export 100:1
Router(config-vrf)# route-target import 100:1
Router(config-vrf)# exit

Router(config)# ip vrf v2
Router(config-vrf)# rd 100:2
Router(config-vrf)# route-target export 100:2
Router(config-vrf)# route-target import 100:2
Router(config-vrf)# exit

Router(config)# ip cef
Router(config)# interface Loopback1
Router(config-if)# ip vrf forwarding v1
Router(config-if)# ip address 3.3.1.3 255.255.255.0
Router(config-if)# exit

Router(config)# interface Loopback2
Router(config-if)# ip vrf forwarding v2
Router(config-if)# ip address 3.3.2.3 255.255.255.0
Router(config-if)# exit

Router(config)# interface gigabitethernet0.10
Router(config-if)# encapsulation dot1q 10
Router(config-if)# ip vrf forwarding v1
Router(config-if)# ip address 38.0.0.3 255.255.255.0
Router(config-if)# exit

Router(config)# interface gigabitethernet0.20
Router(config-if)# encapsulation dot1q 20
Router(config-if)# ip vrf forwarding v2
Router(config-if)# ip address 83.0.0.3 255.255.255.0
Router(config-if)# exit

Router(config)# router bgp 100
Router(config-router)# address-family ipv4 vrf v2
Router(config-router-af)# neighbor 83.0.0.8 remote-as 800
Router(config-router-af)# neighbor 83.0.0.8 activate
Router(config-router-af)# network 3.3.2.0 mask 255.255.255.0
Router(config-router-af)# exit
Router(config-router)# address-family ipv4 vrf v1
Router(config-router-af)# neighbor 38.0.0.8 remote-as 800
Router(config-router-af)# neighbor 38.0.0.8 activate
Router(config-router-af)# network 3.3.1.0 mask 255.255.255.0
Router(config-router-af)# end
```

Displaying Multi-VRF CE Status

You can use the following privileged EXEC commands to display information about multi-VRF CE configuration and status.

Command	Purpose
show ip protocols vrf <i>vrf-name</i>	Display routing protocol information associated with a VRF.
show ip route vrf <i>vrf-name</i> [connected] [<i>protocol</i> [<i>as-number</i>]] [list] [mobile] [odr] [profile] [static] [summary] [supernets-only]	Display IP routing table information associated with a VRF.
show ip vrf [brief detail interfaces] [<i>vrf-name</i>]	Display information about the defined VRF instances.

Configuring Protocol-Independent Features

This section describes how to configure IP routing protocol-independent features. For a complete description of the IP routing protocol-independent commands in this chapter, see the [Cisco IOS IP Routing: Protocol-Independent Command Reference](#).

This section includes the following topics:

- [Configuring Cisco Express Forwarding, page 2-116](#)
- [Configuring the Number of Equal-Cost Routing Paths, page 2-118](#)
- [Configuring Static Unicast Routes, page 2-119](#)
- [Specifying Default Routes and Networks, page 2-120](#)
- [Using Route Maps to Redistribute Routing Information, page 2-121](#)
- [Configuring Policy-Based Routing, page 2-126](#)
- [Filtering Routing Information, page 2-129](#)
- [Managing Authentication Keys, page 2-133](#)

Configuring Cisco Express Forwarding

Cisco Express Forwarding (CEF) is a Layer 3 IP switching technology used to optimize network performance. CEF implements an advanced IP look-up and forwarding algorithm to deliver maximum Layer 3 switching performance. CEF is less CPU-intensive than fast switching route caching, allowing more CPU processing power to be dedicated to packet forwarding. In dynamic networks, fast switching cache entries are frequently invalidated because of routing changes, which can cause traffic to be process switched using the routing table, instead of fast switched using the route cache. CEF uses the Forwarding Information Base (FIB) lookup table to perform destination-based switching of IP packets.

The two main components in CEF are the distributed FIB and the distributed adjacency tables.

- The FIB is similar to a routing table or information base and maintains a mirror image of the forwarding information in the IP routing table. When routing or topology changes occur in the network, the IP routing table is updated, and those changes are reflected in the FIB. The FIB

maintains next-hop address information based on the information in the IP routing table. Because the FIB contains all known routes that exist in the routing table, CEF eliminates route cache maintenance, is more efficient for switching traffic, and is not affected by traffic patterns.

- Nodes in the network are said to be adjacent if they can reach each other with a single hop across a link layer. CEF uses adjacency tables to prepend Layer 2 addressing information. The adjacency table maintains Layer 2 next-hop addresses for all FIB entries.

Because the switch uses Application Specific Integrated Circuits (ASICs) to achieve Gigabit-speed line rate IP traffic, CEF forwarding applies only to the software-forwarding path, that is, traffic that is forwarded by the CPU.

CEF is enabled globally by default. If for some reason it is disabled, you can re-enable it by using the **ip cef** global configuration command.

The default configuration is CEF enabled on all Layer 3 interfaces. Entering the **no ip route-cache cef** interface configuration command disables CEF for traffic that is being forwarded by software. This command does not affect the hardware forwarding path. Disabling CEF and using the **debug ip packet detail** privileged EXEC command can be useful to debug software-forwarded traffic. To enable CEF on an interface for the software-forwarding path, use the **ip route-cache cef** interface configuration command.



Caution

Although the **no ip route-cache cef** interface configuration command to disable CEF on an interface is visible in the CLI, we strongly recommend that you do not disable CEF on interfaces except for debugging purposes.

BEFORE YOU BEGIN

- Cisco Express Forwarding requires a software image that includes Cisco Express Forwarding and IP routing enabled on the switch.
- If you enable Cisco Express Forwarding and then create an access list that uses the log keyword, the packets that match the access list are not Cisco Express Forwarding switched. They are process switched. Logging disables Cisco Express Forwarding.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip cef	Enable CEF operation.
Step 3	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 4	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 5	ip route-cache cef	Enable CEF on the interface for software-forwarded traffic.
Step 6	end	Return to privileged EXEC mode.
Step 7	show ip cef	Display the CEF status on all interfaces.
Step 8	show cef linecard [detail]	Display CEF-related interface information.

	Command	Purpose
Step 9	show cef interface <i>[interface-id]</i>	Display detailed CEF information for all interfaces or the specified interface.
Step 10	show adjacency	Display CEF adjacency table information.
Step 11	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

```
Switch(config)# ip cef
Switch(config)# interface ethernet 0
Switch(config-if)# ip route-cache cef
Switch(config-if)# end
```

Configuring the Number of Equal-Cost Routing Paths

When a router has two or more routes to the same network with the same metrics, these routes can be thought of as having an equal cost. The term *parallel path* is another way to see occurrences of equal-cost routes in a routing table. If a router has two or more equal-cost paths to a network, it can use them concurrently. Parallel paths provide redundancy in case of a circuit failure and also enable a router to load balance packets over the available paths for more efficient use of available bandwidth.

Although the router automatically learns about and configures equal-cost routes, you can control the maximum number of parallel paths supported by an IP routing protocol in its routing table.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp rip ospf eigrp}	Enter router configuration mode.
Step 3	maximum-paths <i>maximum</i>	Set the maximum number of parallel paths for the protocol routing table. The range is from 1 to 8; the default is 4 for most IP routing protocols, but only 1 for BGP.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip protocols	Verify the setting in the <i>Maximum path</i> field.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no maximum-paths** router configuration command to restore the default value.

EXAMPLE

The following example shows how to allow a maximum of 16 paths to a destination in an OSPF routing process:

```
Switch(config)# router ospf 3
Switch(config-router)# maximum-paths 16
```

Configuring Static Unicast Routes

Static unicast routes are user-defined routes that cause packets moving between a source and a destination to take a specified path. Static routes can be important if the router cannot build a route to a particular destination and are useful for specifying a gateway of last resort to which all unroutable packets are sent.

The switch retains static routes until you remove them. However, you can override static routes with dynamic routing information by assigning administrative distance values. Each dynamic routing protocol has a default administrative distance, as listed in [Table 2-2](#). If you want a static route to be overridden by information from a dynamic routing protocol, set the administrative distance of the static route higher than that of the dynamic protocol.

Table 2-2 *Dynamic Routing Protocol Default Administrative Distances*

Route Source	Default Distance
Connected interface	0
Static route	1
Enhanced IGRP summary route	5
External BGP	20
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
Internal BGP	200
Unknown	225

Static routes that point to an interface are advertised through RIP, IGRP, and other dynamic routing protocols, whether or not static **redistribute** router configuration commands were specified for those routing protocols. These static routes are advertised because static routes that point to an interface are considered in the routing table to be connected and hence lose their static nature. However, if you define a static route to an interface that is not one of the networks defined in a network command, no dynamic routing protocols advertise the route unless a **redistribute** static command is specified for these protocols.

When an interface goes down, all static routes through that interface are removed from the IP routing table. When the software can no longer find a valid next hop for the address specified as the forwarding router's address in a static route, the static route is also removed from the IP routing table.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip route <i>prefix mask</i> { <i>address</i> <i>interface</i> } [<i>distance</i>]	Establish a static route.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show ip route	Display the current state of the routing table to verify the configuration.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip route** *prefix mask {address | interface}* global configuration command to remove a static route.

EXAMPLE

The following example shows how to choose an administrative distance of 110. In this case, packets for network 10.0.0.0 will be routed to a router at 172.31.3.4 if dynamic information with an administrative distance less than 110 is not available.

```
ip route 10.0.0.0 255.0.0.0 172.31.3.4 110
```

Specifying Default Routes and Networks

A router might not be able to learn the routes to all other networks. To provide complete routing capability, you can use some routers as smart routers and give the remaining routers default routes to the smart router. (Smart routers have routing table information for the entire internetwork.) These default routes can be dynamically learned or can be configured in the individual routers. Most dynamic interior routing protocols include a mechanism for causing a smart router to generate dynamic default information that is then forwarded to other routers.

If a router has a directly connected interface to the specified default network, the dynamic routing protocols running on that device generate a default route. In RIP, it advertises the pseudonetwork 0.0.0.0.s

A router that is generating the default for a network also might need a default of its own. One way a router can generate its own default is to specify a static route to the network 0.0.0.0 through the appropriate device.

When default information is passed through a dynamic routing protocol, no further configuration is required. The system periodically scans its routing table to choose the optimal default network as its default route. In IGRP networks, there might be several candidate networks for the system default. Cisco routers use administrative distance and metric information to set the default route or the gateway of last resort.

If dynamic default information is not being passed to the system, candidates for the default route are specified with the **ip default-network** global configuration command. If this network appears in the routing table from any source, it is flagged as a possible choice for the default route. If the router has no interface on the default network, but does have a path to it, the network is considered as a possible candidate, and the gateway to the best default path becomes the gateway of last resort.

BEFORE YOU BEGIN

The **ip default-network** command is a classful command. It is effective only if the network mask of the network that you wish to configure as a candidate route for computing the gateway of last resort matches the network mask in the Routing Information Base (RIB).

For example, if you configure **ip default-network 10.0.0.0**, then the mask considered by the routing protocol is 10.0.0.0/8, as it is a Class A network. The gateway of last resort is set only if the RIB contains a 10.0.0.0/8 route.

If you need to use the **ip default-network** command, ensure that the RIB contains a network route that matches the major mask of the network class.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip default-network <i>network number</i>	Specify a default network.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ip route	Display the selected default route in the gateway of last resort display.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ip default-network** *network number* global configuration command to remove the route.

EXAMPLE

The following example defines a static route to network 10.0.0.0 as the static default route:

```
ip route 10.0.0.0 255.0.0.0 10.108.3.4
ip default-network 10.0.0.0
```

Using Route Maps to Redistribute Routing Information

The switch can run multiple routing protocols simultaneously, and it can redistribute information from one routing protocol to another. Redistributing information from one routing protocol to another applies to all supported IP-based routing protocols.

You can also conditionally control the redistribution of routes between routing domains by defining enhanced packet filters or route maps between the two domains. The **match** and **set** route-map configuration commands define the condition portion of a route map. The **match** command specifies that a criterion must be matched. The **set** command specifies an action to be taken if the routing update meets the conditions defined by the match command. Although redistribution is a protocol-independent feature, some of the **match** and **set** route-map configuration commands are specific to a particular protocol.

One or more **match** commands and one or more **set** commands follow a **route-map** command. If there are no **match** commands, everything matches. If there are no **set** commands, nothing is done, other than the match. Therefore, you need at least one **match** or **set** command.



Note

A route map with no **set** route-map configuration commands is sent to the CPU, which causes high CPU utilization.

You can also identify route-map statements as **permit** or **deny**. If the statement is marked as a deny, the packets meeting the match criteria are sent back through the normal forwarding channels (destination-based routing). If the statement is marked as permit, set clauses are applied to packets meeting the match criteria. Packets that do not meet the match criteria are forwarded through the normal routing channel.

You can use the BGP route map **continue** clause to execute additional entries in a route map after an entry is executed with successful match and set clauses. You can use the **continue** clause to configure and organize more modular policy definitions so that specific policy configurations need not be repeated within the same route map. The switch supports the **continue** clause for outbound policies. For more information about using the route map **continue** clause, see the “BGP Route-Map Continue” section in the *IP Routing: BGP Configuration Guide, Cisco IOS Release 15M&T*.

**Note**

Although each of Steps 3 through 14 in the following section is optional, you must enter at least one **match** route-map configuration command and one **set** route-map configuration command.

BEFORE YOU BEGIN

You should know your network design and how you want traffic to flow through it before configuring route redistribution or policy-based routing.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>map-tag</i> [permit deny] [<i>sequence number</i>]	Define any route maps used to control redistribution and enter route-map configuration mode. <ul style="list-style-type: none"> <i>map-tag</i>—A meaningful name for the route map. The redistribute router configuration command uses this name to reference this route map. Multiple route maps might share the same map tag name. (Optional) If permit is specified and the match criteria are met for this route map, the route is redistributed as controlled by the set actions. If deny is specified, the route is not redistributed. <i>sequence number</i> (Optional)— Number that indicates the position a new route map is to have in the list of route maps already configured with the same name.
Step 3	match as-path <i>path-list-number</i>	Match a BGP AS path access list.
Step 4	match community-list <i>community-list-number</i> [exact]	Match a BGP community list.
Step 5	match ip address { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match a standard access list by specifying the name or number. It can be an integer from 1 to 199.

	Command	Purpose
Step 6	match metric <i>metric-value</i>	Match the specified route metric. The <i>metric-value</i> can be an EIGRP metric with a specified value from 0 to 4294967295.
Step 7	match ip next-hop { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match a next-hop router address passed by one of the access lists specified (numbered from 1 to 199).
Step 8	match tag <i>tag value</i> [... <i>tag-value</i>]	Match the specified tag value in a list of one or more route tag values. Each can be an integer from 0 to 4294967295.
Step 9	match interface <i>type number</i> [... <i>type number</i>]	Match the specified next hop route out one of the specified interfaces.
Step 10	match ip route-source { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	Match the address specified by the specified advertised access lists.
Step 11	match route-type { local internal external [type-1 type-2]}	Match the specified route-type : <ul style="list-style-type: none"> • local—Locally generated BGP routes. • internal—OSPF intra-area and interarea routes or EIGRP internal routes. • external—OSPF external routes (Type 1 or Type 2) or EIGRP external routes.
Step 12	set dampening <i>halflife reuse suppress max-suppress-time</i>	Set BGP route dampening factors.
Step 13	set local-preference <i>value</i>	Assign a value to a local BGP path.
Step 14	set origin { igp egp <i>as</i> incomplete }	Set the BGP origin code.
Step 15	set as-path { tag prepend <i>as-path-string</i> }	Modify the BGP autonomous system path.
Step 16	set level { level-1 level-2 level-1-2 stub-area backbone }	Set the level for routes that are advertised into the specified area of the routing domain. The stub-area and backbone are OSPF NSSA and backbone areas.
Step 17	set metric <i>metric value</i>	Set the metric value to give the redistributed routes (for EIGRP only). The <i>metric value</i> is an integer from -294967295 to 294967295.

	Command	Purpose
Step 18	set metric <i>bandwidth delay reliability loading mtu</i>	Set the metric value to give the redistributed routes (for EIGRP only): <ul style="list-style-type: none"> <i>bandwidth</i>—Metric value or IGRP bandwidth of the route in kilobits per second in the range 0 to 4294967295 <i>delay</i>—Route delay in tens of microseconds in the range 0 to 4294967295. <i>reliability</i>—Likelihood of successful packet transmission expressed as a number between 0 and 255, where 255 means 100 percent reliability and 0 means no reliability. <i>loading</i>—Effective bandwidth of the route expressed as a number from 0 to 255 (255 is 100 percent loading). <i>mtu</i>—Minimum maximum transmission unit (MTU) size of the route in bytes in the range 0 to 4294967295.
Step 19	set metric-type { type-1 type-2 }	Set the OSPF external metric type for redistributed routes.
Step 20	set metric-type internal	Set the multi-exit discriminator (MED) value on prefixes advertised to external BGP neighbor to match the IGP metric of the next hop.
Step 21	set weight	Set the BGP weight for the routing table. The value can be from 1 to 65535.
Step 22	end	Return to privileged EXEC mode.
Step 23	show route-map	Display all route maps configured or only the one specified to verify configuration.
Step 24	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete an entry, use the **no route-map** *map tag* global configuration command or the **no match** or **no set** route-map configuration commands.

EXAMPLE

The following example shows how to redistribute Routing Information Protocol (RIP) routes with a hop count equal to 1 to Open Shortest Path First (OSPF). These routes will be redistributed to OSPF as external link-state advertisements (LSAs) with a metric of 5, metric type of Type 1, and a tag equal to 1.

```
Switch(config)# router ospf 109
Switch(config-router)# redistribute rip route-map rip-to-ospf
Switch(config-router)# exit
Switch(config)# route-map rip-to-ospf permit
Switch(config-route-map)# match metric 1
Switch(config-route-map)# set metric 5
Switch(config-route-map)# set metric-type type1
Switch(config-route-map)# set tag 1
```

Controlling Route Redistribution

You can distribute routes from one routing domain into another and control route distribution. Note that the keywords in this procedure are the same as defined in the previous procedure.

The metrics of one routing protocol do not necessarily translate into the metrics of another. In these situations, an artificial metric is assigned to the redistributed route. Uncontrolled exchanging of routing information between different routing protocols can create routing loops and seriously degrade network operation.

If you have not defined a default redistribution metric that replaces metric conversion, some automatic metric translations occur between routing protocols:

- RIP can automatically redistribute static routes. It assigns static routes a metric of 1 (directly connected).
- Any protocol can redistribute other routing protocols if a default mode is in effect.

BEFORE YOU BEGIN

Review the usage guidelines and additional examples for the **redistribute** command in the [Cisco IOS IP Routing: Protocol-Independent Command Reference](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp rip ospf eigrp}	Enter router configuration mode.
Step 3	redistribute <i>protocol</i> [<i>process-id</i>] { level-1 level-1-2 level-2 } [metric <i>metric-value</i>] [metric-type <i>type-value</i>] [match internal external <i>type-value</i>] [tag <i>tag-value</i>] [route-map <i>map-tag</i>] [weight <i>weight</i>] [subnets]	Redistribute routes from one routing protocol to another routing protocol. If no route-maps are specified, all routes are redistributed. If the keyword route-map is specified with no <i>map-tag</i> , no routes are distributed.
Step 4	default-metric <i>number</i>	Cause the current routing protocol to use the same metric value for all redistributed routes (BGP, RIP, and OSPF).
Step 5	default-metric <i>bandwidth delay reliability loading mtu</i>	Cause the EIGRP routing protocol to use the same metric value for all non-EIGRP redistributed routes.
Step 6	end	Return to privileged EXEC mode.
Step 7	show route-map	Display all route maps configured or only the one specified to verify configuration.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable redistribution, use the **no** form of the commands.

EXAMPLE

Given 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 link-state PDU with metric 5:

```
router isis
 redistribute rip route-map ourmap
 redistribute iso-igrp remote route-map ourmap
 route-map ourmap permit
 match ip address 1
 match clns address ourprefix
 set metric 5
 set level level-2
 access-list 1 permit 160.89.0.0 0.0.255.255
 clns filter-set ourprefix permit 49.0001.0002...
```

Configuring Policy-Based Routing

You can use policy-based routing (PBR) to configure a defined policy for traffic flows. By using PBR, you can have more control over routing by reducing the reliance on routes derived from routing protocols. PBR can specify and implement routing policies that allow or deny paths based on:

- Identity of a particular end system
- Application
- Protocol

You can use PBR to provide equal-access and source-sensitive routing, routing based on interactive versus batch traffic, or routing based on dedicated links. For example, you could transfer stock records to a corporate office on a high-bandwidth, high-cost link for a short time while transmitting routine application data such as e-mail over a low-bandwidth, low-cost link.

With PBR, you classify traffic using access control lists (ACLs) and then make traffic go through a different path. PBR is applied to incoming packets. All packets received on an interface with PBR enabled are passed through route maps. Based on the criteria defined in the route maps, packets are forwarded (routed) to the appropriate next hop.

- If packets do not match any route map statements, all set clauses are applied.
- If a statement is marked as permit and the packets do not match any route-map statements, the packets are sent through the normal forwarding channels, and destination-based routing is performed.
- For PBR, route-map statements marked as deny are not supported.

For more information about configuring route maps, see the [“Using Route Maps to Redistribute Routing Information” section on page 2-121](#).

You can use standard IP ACLs to specify match criteria for a source address or extended IP ACLs to specify match criteria based on an application, a protocol type, or an end station. The process proceeds through the route map until a match is found. If no match is found, normal destination-based routing occurs. There is an implicit deny at the end of the list of match statements.

If match clauses are satisfied, you can use a set clause to specify the IP addresses identifying the next hop router in the path.

For details about PBR commands and keywords, see the [IP Routing: Protocol-Independent Configuration Guide, Cisco IOS Release 15M&T](#).

PBR Configuration Guidelines

Before configuring PBR, you should be aware of this information:

- Multicast traffic is not policy-routed. PBR applies only to unicast traffic.
- You can enable PBR on a routed port or an SVI.
- The switch does not support **route-map deny** statements for PBR.
- You can apply a policy route map to an EtherChannel port channel in Layer 3 mode, but you cannot apply a policy route map to a physical interface that is a member of the EtherChannel. If you try to do so, the command is rejected. When a policy route map is applied to a physical interface, that interface cannot become a member of an EtherChannel.
- You can define a maximum of 246 IP policy route maps on the switch.
- You can define a maximum of 512 access control entries (ACEs) for PBR on the switch.
- When configuring match criteria in a route map, follow these guidelines:
 - Do not match ACLs that permit packets destined for a local address. PBR would forward these packets, which could cause ping or Telnet failure or route protocol flapping.
 - Do not match ACLs with deny ACEs. Packets that match a deny ACE are sent to the CPU, which could cause high CPU utilization.
- To use PBR, you must first enable the default template by using the **sdm prefer default** global configuration command. PBR is not supported with the Layer 2 template. For more information on the SDM templates, see the chapter “Configuring SDM Templates” in the *Cisco Connected Grid Switches System Management Software Configuration Guide*.
- VRF and PBR are mutually-exclusive on a switch interface. You cannot enable VRF when PBR is enabled on an interface. In contrast, you cannot enable PBR when VRF is enabled on an interface.
- The number of TCAM entries used by PBR depends on the route map itself, the ACLs used, and the order of the ACLs and route-map entries.
- Policy-based routing based on packet length, IP precedence and TOS, set interface, set default next hop, or set default interface are not supported. Policy maps with no valid set actions or with set action set to *Don't Fragment* are not supported.

Enabling PBR

By default, PBR is disabled on the switch. To enable PBR, you must create a route map that specifies the match criteria and the resulting action if all of the match clauses are met. Then, you must enable PBR for that route map on an interface. All packets arriving on the specified interface matching the match clauses are subject to PBR.

PBR can be fast-switched or implemented at speeds that do not slow down the switch. Fast-switched PBR supports most match and set commands. PBR must be enabled before you enable fast-switched PBR. Fast-switched PBR is disabled by default.

Packets that are generated by the switch, or local packets, are not normally policy-routed. When you globally enable local PBR on the switch, all packets that originate on the switch are subject to local PBR. Local PBR is disabled by default.

BEFORE YOU BEGIN

See the “[PBR Configuration Guidelines](#)” section on page 2-127.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	route-map <i>map-tag</i> [permit] [<i>sequence number</i>]	<p>Define any route maps used to control where packets are output, and enter route-map configuration mode.</p> <ul style="list-style-type: none"> <i>map-tag</i>—A meaningful name for the route map. The ip policy route-map interface configuration command uses this name to reference the route map. Multiple route maps might share the same map tag name. (Optional) If permit is specified and the match criteria are met for this route map, the route is policy-routed as controlled by the set actions. <p>Note The route-map deny statement is not supported in PBR route maps to be applied to an interface.</p> <ul style="list-style-type: none"> <i>sequence number</i> (Optional)— Number that shows the position of a new route map in the list of route maps already configured with the same name.
Step 3	match ip address { <i>access-list-number</i> <i>access-list-name</i> } [... <i>access-list-number</i> ... <i>access-list-name</i>]	<p>Match the source and destination IP address that is permitted by one or more standard or extended access lists.</p> <p>Note Do not enter an ACL with a deny ACE or an ACL that permits a packet destined for a local address.</p> <p>If you do not specify a match command, the route map applies to all packets.</p>
Step 4	set ip next-hop <i>ip-address</i> [... <i>ip-address</i>]	Specify the action to take on the packets that match the criteria. Set next hop to which to route the packet (the next hop must be adjacent).
Step 5	exit	Return to global configuration mode.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 7	no shutdown	Enable the interface if necessary. By default, UNIs and ENIs are disabled and NNIs are enabled.
Step 8	ip policy route-map <i>map-tag</i>	<p>Enable PBR on a Layer 3 interface, and identify the route map to use. You can configure only one route map on an interface. However, you can have multiple route map entries with different sequence numbers. These entries are evaluated in sequence number order until the first match. If there is no match, packets are routed as usual.</p> <p>Note If the IP policy route map contains a deny statement, the configuration fails.</p>
Step 9	ip route-cache policy	(Optional) Enable fast-switching PBR. You must first enable PBR before enabling fast-switching PBR.
Step 10	exit	Return to global configuration mode.

	Command	Purpose
Step 11	ip local policy route-map <i>map-tag</i>	(Optional) Enable local PBR to perform policy-based routing on packets originating at the switch. This applies to packets generated by the switch and not to incoming packets.
Step 12	end	Return to privileged EXEC mode.
Step 13	show route-map [<i>map-name</i>]	(Optional) Display all route maps configured or only the one specified to verify configuration.
Step 14	show ip policy	(Optional) Display policy route maps attached to interfaces.
Step 15	show ip local policy	(Optional) Display whether or not local policy routing is enabled and, if so, the route map being used.
Step 16	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no route-map** *map-tag* global configuration command or the **no match** or **no set** route-map configuration commands to delete an entry. Use the **no ip policy route-map** *map-tag* interface configuration command to disable PBR on an interface. Use the **no ip route-cache policy** interface configuration command to disable fast-switching PBR. Use the **no ip local policy route-map** *map-tag* global configuration command to disable policy-based routing on packets originating on the switch.

EXAMPLE

The following example sends packets with the destination IP address of 172.21.16.18 to a router at IP address 172.30.3.20:

```
interface serial 0
 ip policy route-map wethersfield
!
route-map wethersfield
 match ip address 172.21.16.18
 set ip next-hop 172.30.3.20
```

Filtering Routing Information

You can filter routing protocol information by performing the tasks described in this section.



Note

When routes are redistributed between OSPF processes, no OSPF metrics are preserved.

Setting Passive Interfaces

To prevent other routers on a local network from dynamically learning about routes, you can use the **passive-interface** router configuration command to keep routing update messages from being sent through a router interface. When you use this command in the OSPF protocol, the interface address you specify as passive appears as a stub network in the OSPF domain. OSPF routing information is neither sent nor received through the specified router interface.

In networks with many interfaces, to avoid having to manually set them as passive, you can set all interfaces to be passive by default by using the **passive-interface default** router configuration command and manually setting interfaces where adjacencies are desired.

BEFORE YOU BEGIN

You should know your network design and how you want traffic to flow through it before filtering routing information.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp rip ospf eigrp}	Enter router configuration mode.
Step 3	passive-interface <i>interface-id</i>	Suppress sending routing updates through the specified Layer 3 interface.
Step 4	passive-interface default	(Optional) Set all interfaces as passive by default.
Step 5	no passive-interface <i>interface type</i>	(Optional) Activate only those interfaces that need to have adjacencies sent.
Step 6	network <i>network-address</i>	(Optional) Specify the list of networks for the routing process. The <i>network-address</i> is an IP address.
Step 7	end	Return to privileged EXEC mode.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use a network monitoring privileged EXEC command such as **show ip ospf interface** to verify the interfaces that you enabled as passive, or use the **show ip interface** privileged EXEC command to verify the interfaces that you enabled as active.

To re-enable the sending of routing updates, use the **no passive-interface *interface-id*** router configuration command.

EXAMPLE

The following example sends EIGRP updates to all interfaces on network 10.108.0.0 except Ethernet interface 1:

```
router eigrp 109
 network 10.108.0.0
 passive-interface ethernet 1
```

The following example sets all interfaces as passive and then activates Ethernet interface 0:

```
router ospf 100
 passive-interface default
 no passive-interface ethernet0
 network 10.108.0.1 0.0.0.255 area 0
```

Controlling Advertising and Processing in Routing Updates

You can use the **distribute-list** router configuration command with access control lists to suppress routes from being advertised in routing updates and to prevent other routers from learning one or more routes. When used in OSPF, this feature applies to only external routes, and you cannot specify an interface name.

You can also use a **distribute-list** router configuration command to avoid processing certain routes listed in incoming updates. (This feature does not apply to OSPF.)

BEFORE YOU BEGIN

Configure an access list defining which networks are to be sent or received and which are to be suppressed in routing updates.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp rip eigrp}	Enter router configuration mode.
Step 3	distribute-list {access-list-number access-list-name} out [interface-name routing process autonomous-system-number]	Permit or deny routes from being advertised in routing updates, depending upon the action listed in the access list.
Step 4	distribute-list {access-list-number access-list-name} in [type-number]	Suppress processing in routes listed in updates.
Step 5	end	Return to privileged EXEC mode.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no distribute-list in** router configuration command to change or cancel a filter. To cancel suppression of network advertisements in updates, use the **no distribute-list out** router configuration command.

EXAMPLE

In the following example, a prefix list and distribute list are defined to configure the BGP routing process to accept traffic from only network 10.1.1.0/24, network 192.168.1.0, and network 10.108.0.0. An inbound route refresh is initiated to activate the distribute-list.

```
Switch(config)# ip prefix-list RED permit 10.1.1.0/24
Switch(config)# ip prefix-list RED permit 10.108.0.0/16
Switch(config)# ip prefix-list RED permit 192.168.1.0/24
Switch(config)# router bgp 50000
Switch(config-router)# network 10.108.0.0
Switch(config-router)# distribute-list prefix RED in
Switch(config-router)# end
Switch# clear ip bgp in
```

Filtering Sources of Routing Information

Because some routing information might be more accurate than others, you can use filtering to prioritize information coming from different sources. An *administrative distance* is a rating of the trustworthiness of a routing information source, such as a router or group of routers. In a large network, some routing protocols can be more reliable than others. By specifying administrative distance values, you enable the router to intelligently discriminate between sources of routing information. The router always picks the route whose routing protocol has the lowest administrative distance. [Table 2-2 on page 2-119](#) shows the default administrative distances for various routing information sources.

Because each network has its own requirements, there are no general guidelines for assigning administrative distances.

BEFORE YOU BEGIN

- Always set the administrative distance from the least to the most specific network.
- Review the usage guidelines and additional examples for the **distance** command in the [Cisco IOS IP Routing: Protocol-Independent Command Reference](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router {bgp rip ospf eigrp}	Enter router configuration mode.
Step 3	distance weight {ip-address {ip-address mask}} [ip access list]	Define an administrative distance. <ul style="list-style-type: none"> • <i>weight</i>—The administrative distance as an integer from 10 to 255. Used alone, <i>weight</i> specifies a default administrative distance that is used when no other specification exists for a routing information source. Routes with a distance of 255 are not installed in the routing table. • (Optional) <i>ip access list</i>—An IP standard or extended access list to be applied to incoming routing updates.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ip protocols	Display the default administrative distance for a specified routing process.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a distance definition, use the **no distance** router configuration command.

EXAMPLE

In the following example, the **route eigrp** global configuration command sets up EIGRP routing in autonomous system number 109. The network router configuration commands specify EIGRP routing on networks 192.168.7.0 and 172.16.0.0. The first distance command sets the administrative distance to 90 for all routers on the Class C network 192.168.7.0. The second distance command sets the administrative distance to 120 for the router with the address 172.16.1.3.

```
Switch# configure terminal
Switch(config)# router eigrp 109
Switch(config-router)# network 192.168.7.0
Switch(config-router)# network 172.16.0.0
Switch(config-router)# distance 90 192.168.7.0 0.0.0.255
Switch(config-router)# distance 120 172.16.1.3 0.0.0.255
Switch(config-router)# end
```

In the following example, the set distance is from the least to the most specific network:

```
Switch# configure terminal
Switch(config)# router eigrp 109
Switch(config-router)# distance 22 10.0.0.0 0.0.0.255
Switch(config-router)# distance 33 10.11.0.0 0.0.0.255
Switch(config-router)# distance 44 10.11.12.0 0.0.0.255
Switch(config-router)# end
```

Managing Authentication Keys

Key management is a method of controlling authentication keys used by routing protocols. Not all protocols can use key management. Authentication keys are available for EIGRP and RIP Version 2.

To manage authentication keys, define a key chain, identify the keys that belong to the key chain, and specify how long each key is valid. Each key has its own key identifier (specified with the **key number** key chain configuration command), which is stored locally. The combination of the key identifier and the interface associated with the message uniquely identifies the authentication algorithm and Message Digest 5 (MD5) authentication key in use.

You can configure multiple keys with life times. Only one authentication packet is sent, regardless of how many valid keys exist. The software examines the key numbers in order from lowest to highest, and uses the first valid key it encounters. The lifetimes allow for overlap during key changes. Note that the router must know these lifetimes.

BEFORE YOU BEGIN

Before you manage authentication keys, you must enable authentication. See the appropriate protocol section to see how to enable authentication for that protocol.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	key chain <i>name-of-chain</i>	Identify a key chain, and enter key chain configuration mode.
Step 3	key <i>number</i>	Identify the key number. The range is 0 to 2147483647.
Step 4	key-string <i>text</i>	Identify the key string. The string can contain from 1 to 80 uppercase and lowercase alphanumeric characters, but the first character cannot be a number.
Step 5	accept-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be received. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite .

	Command	Purpose
Step 6	send-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }	(Optional) Specify the time period during which the key can be sent. The <i>start-time</i> and <i>end-time</i> syntax can be either <i>hh:mm:ss Month date year</i> or <i>hh:mm:ss date Month year</i> . The default is forever with the default <i>start-time</i> and the earliest acceptable date as January 1, 1993. The default <i>end-time</i> and duration is infinite.
Step 7	end	Return to privileged EXEC mode.
Step 8	show key chain	Display authentication key information.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove the key chain, use the **no key chain** *name-of-chain* global configuration command.

EXAMPLE

The following example configures a key chain named chain1. The key named key1 will be accepted from 1:30 p.m. to 3:30 p.m. and be sent from 2:00 p.m. to 3:00 p.m. The key named key2 will be accepted from 2:30 p.m. to 4:30 p.m. and be sent from 3:00 p.m. to 4:00 p.m. The overlap allows for migration of keys or a discrepancy in the set time of the router. There is a 30-minute leeway on each side to handle time differences.

```
Router(config)# interface ethernet 0
Router(config-if)# ip rip authentication key-chain chain1
Router(config-if)# ip rip authentication mode md5
!
Router(config)# router rip
Router(config-router)# network 172.19.0.0
Router(config-router)# version 2
!
Router(config)# key chain chain1
Router(config-keychain)# key 1
Router(config-keychain-key)# key-string key1
Router(config-keychain-key)# accept-lifetime 13:30:00 Jan 25 1996 duration 7200
Router(config-keychain-key)# send-lifetime 14:00:00 Jan 25 1996 duration 3600
Router(config-keychain-key)# exit
Router(config-keychain)# key 2
Router(config-keychain-key)# key-string key2
Router(config-keychain-key)# accept-lifetime 14:30:00 Jan 25 1996 duration 7200
Router(config-keychain-key)# send-lifetime 15:00:00 Jan 25 1996 duration 3600
```

Verifying Configuration

You can remove all contents of a particular cache, table, or database. You can also display specific statistics.

Command	Purpose
clear ip route { <i>network</i> [<i>mask</i> *]}	Clear one or more routes from the IP routing table.
show ip protocols	Display the parameters and state of the active routing protocol process.
show ip route [<i>address</i> [<i>mask</i>] [longer-prefixes]] [<i>protocol</i> [<i>process-id</i>]]	Display the current state of the routing table.
show ip route summary	Display the current state of the routing table in summary form.
show ip route supernets-only	Display supernets.
show ip cache	Display the routing table used to switch IP traffic.
show route-map [<i>map-name</i>]	Display all route maps configured or only the one specified.

Related Documents

- [Cisco IOS Master Command List, All Releases](#)
- [IP Addressing: ARP Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco Connected Grid Switches Layer 2 Switching Software Configuration Guide](#)
- [Cisco Connected Grid Switches High Availability and Redundancy Software Configuration Guide](#)
- [Cisco Connected Grid Switches System Management Software Configuration Guide](#)
- [Cisco Connected Grid Switches Security Software Configuration Guide](#)
- [Cisco IOS IP Routing: RIP Command Reference](#)
- [IP Routing: RIP Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco IOS IP Routing: OSPF Command Reference](#)
- [IP Routing: OSPF Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco IOS IP Routing: EIGRP Command Reference](#)
- [IP Routing: EIGRP Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco IOS IP Routing: BGP Command Reference](#)
- [IP Routing: BGP Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco IOS ISO CLNS Command Reference](#)
- [ISO CLNS Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco IOS IP Routing: ISIS Command Reference](#)
- [IP Routing: ISIS Configuration Guide, Cisco IOS Release 15M&T](#)
- [IP Routing: BFD Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco IOS IP Routing: Protocol-Independent Command Reference](#)
- [IP Routing: Protocol-Independent Configuration Guide, Cisco IOS Release 15M&T](#)
- *Internet Routing Architectures*, published by Cisco Press



Configuring IPv6 Unicast Routing

This chapter describes how to configure IPv6 unicast routing on the Cisco Industrial Ethernet 2000U Series Switches, hereafter referred to as IE 2000U or switch.

For information about configuring IPv4 unicast routing, see [Chapter 2, “Configuring IP Unicast Routing.”](#) For information on configuring IPv6 access control lists (ACLs) see the “Configuring IPv6 ACLs” chapter in the *Cisco Connected Grid Switches Security Software Configuration Guide*.

To use this feature, the switch must be running the IP services image. To enable IPv6 routing, you must configure the switch to use a dual IPv4 and IPv6 switch database management (SDM) template. See the “Dual IPv4 and IPv6 Protocol Stacks” section on [page 3-5](#).



Note

For complete syntax and usage information for the commands used in this chapter, see the Cisco IOS documentation listed in the “[Related Documents](#)” section on [page 3-29](#).

- [Information About IPv6, page 3-1](#)
- [Prerequisites, page 3-7](#)
- [Guidelines and Limitations, page 3-7](#)
- [Default Settings, page 3-8](#)
- [Configuring IPv6, page 3-8](#)
- [Verifying Configuration, page 3-25](#)
- [Configuration Example, page 3-26](#)
- [Related Documents, page 3-29](#)

Information About IPv6

IPv4 users can move to IPv6 and receive services such as end-to-end security, quality of service (QoS), and globally unique addresses. The IPv6 address space reduces the need for private addresses and Network Address Translation (NAT) processing by border routers at network edges.

This section describes IPv6 implementation on the switch and includes the following topics:

- [IPv6 Addresses, page 3-2](#)
- [Supported IPv6 Unicast Routing Features, page 3-2](#)
- [Unsupported IPv6 Unicast Routing Features, page 3-7](#)

IPv6 Addresses

The switch supports only IPv6 unicast addresses. It does not support site-local unicast addresses, anycast addresses, or multicast addresses.

The IPv6 128-bit addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons in the format: n:n:n:n:n:n:n:n. This is an example of an IPv6 address:

```
2031:0000:130F:0000:0000:09C0:080F:130B
```

For easier implementation, leading zeros in each field are optional. This is the same address without leading zeros:

```
2031:0:130F:0:0:9C0:80F:130B
```

You can also use two colons (::) to represent successive hexadecimal fields of zeros, but you can use this short version only once in each address:

```
2031:0:130F::09C0:080F:130B
```

For more information about IPv6 address formats, address types, and the IPv6 packet header, see the [IPv6 Addressing and Basic Connectivity Configuration Guide, Cisco IOS Release 15M&T](#) in the [IPv6 Configuration Library, Cisco IOS Release 15M&T](#).

In the “Information About Implementing Basic Connectivity for IPv6” chapter, these sections apply to the switch:

- IPv6 Address Formats
- IPv6 Address Type: Unicast
- IPv6 Address Output Display
- Simplified IPv6 Packet Header

Supported IPv6 Unicast Routing Features

Support on the switch includes expanded address capability, header format simplification, improved support of extensions and options, and hardware parsing of the extension header. The switch supports hop-by-hop extension header packets, which are routed or bridged in software.

The switch provides IPv6 routing capability over 802.1Q trunk ports for static routes, Routing Information Protocol (RIP) for IPv6, and Open Shortest Path First (OSPF) Version 3 Protocol. It supports up to 16 equal-cost routes and can simultaneously forward IPv4 and IPv6 frames at line rate.

**Note**

For more information about the IPv6 unicast routing features described in this section, see the [IPv6 Configuration Library, Cisco IOS Release 15M&T](#) and [IPv6 Implementation Guide, Cisco IOS Release 15.2M&T](#).

- [128-Bit Unicast Addresses, page 3-3](#)
- [DNS for IPv6, page 3-3](#)
- [Path MTU Discovery for IPv6 Unicast, page 3-3](#)
- [ICMPv6, page 3-4](#)
- [Neighbor Discovery, page 3-4](#)
- [Default Router Preference, page 3-4](#)

- [IPv6 Stateless Autoconfiguration and Duplicate Address Detection, page 3-4](#)
- [IPv6 Applications, page 3-4](#)
- [Dual IPv4 and IPv6 Protocol Stacks, page 3-5](#)
- [DHCP for IPv6 Address Assignment, page 3-5](#)
- [Static Routes for IPv6, page 3-5](#)
- [RIP for IPv6, page 3-6](#)
- [OSPF for IPv6, page 3-6](#)
- [EIGRP IPv6, page 3-6](#)
- [Multiprotocol BGP for IPv6, page 3-6](#)
- [SNMP and Syslog Over IPv6, page 3-6](#)
- [HTTP\(S\) Over IPv6, page 3-7](#)

128-Bit Unicast Addresses

The switch supports aggregatable global unicast addresses and link-local unicast addresses. It does not support site-local unicast addresses.

- Aggregatable global unicast addresses are IPv6 addresses from the aggregatable global unicast prefix. The address structure enables strict aggregation of routing prefixes and limits the number of routing table entries in the global routing table. These addresses are used on links that are aggregated through organizations and eventually to the Internet service provider.

These addresses are defined by a global routing prefix, a subnet ID, and an interface ID. Current global unicast address allocation uses the range of addresses that start with binary value 001 (2000::/3). Addresses with a prefix of 2000::/3(001) through E000::/3(111) must have 64-bit interface identifiers in the extended unique identifier (EUI)-64 format.

- Link local unicast addresses can be automatically configured on any interface by using the link-local prefix FE80::/10(1111 1110 10) and the interface identifier in the modified EUI format. Link-local addresses are used in the neighbor discovery protocol (NDP) and the stateless autoconfiguration process. Nodes on a local link use link-local addresses and do not require globally unique addresses to communicate. IPv6 routers do not forward packets with link-local source or destination addresses to other links.

DNS for IPv6

IPv6 supports Domain Name System (DNS) record types in the DNS name-to-address and address-to-name lookup processes. The DNS AAAA resource record types support IPv6 addresses and are equivalent to an A address record in IPv4. The switch supports DNS resolution for IPv4 and IPv6.

Path MTU Discovery for IPv6 Unicast

The switch supports advertising the system maximum transmission unit (MTU) to IPv6 nodes and path MTU discovery. Path MTU discovery allows a host to dynamically discover and adjust to differences in the MTU size of every link along a given data path. In IPv6, if a link along the path is not large enough to accommodate the packet size, the source of the packet handles the fragmentation. The switch does not support path MTU discovery for multicast packets.

ICMPv6

The Internet Control Message Protocol (ICMP) in IPv6 generates error messages, such as ICMP destination unreachable messages, to report errors during processing and other diagnostic functions. In IPv6, ICMP packets are also used in the neighbor discovery protocol and path MTU discovery.

Neighbor Discovery

The switch supports NDP for IPv6, a protocol running on top of ICMPv6, and static neighbor entries for IPv6 stations that do not support NDP. The IPv6 neighbor discovery process uses ICMP messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), to verify the reachability of the neighbor, and to keep track of neighboring routers.

The switch supports ICMPv6 redirect for routes with mask lengths less than 64 bits. ICMP redirect is not supported for host routes or for summarized routes with mask lengths greater than 64 bits.

Neighbor discovery throttling ensures that the switch CPU is not unnecessarily burdened while it is in the process of obtaining the next hop forwarding information to route an IPv6 packet. The switch drops any additional IPv6 packets whose next hop is the same neighbor that the switch is actively trying to resolve. This drop avoids further load on the CPU.

Default Router Preference

The switch supports IPv6 default router preference (DRP), an extension in router advertisement messages. DRP improves the ability of a host to select an appropriate router, especially when the host is multihomed and the routers are on different links. The switch does not support the Route Information Option in RFC 4191.

An IPv6 host maintains a default router list from which it selects a router for traffic to offlink destinations. The selected router for a destination is then cached in the destination cache. NDP for IPv6 specifies that routers that are reachable or probably reachable are preferred over routers whose reachability is unknown or suspect. For reachable or probably reachable routers, NDP can either select the same router every time or cycle through the router list. By using DRP, you can configure an IPv6 host to prefer one router over another, provided both are reachable or probably reachable.

IPv6 Stateless Autoconfiguration and Duplicate Address Detection

The switch uses stateless autoconfiguration to manage link, subnet, and site addressing changes, such as management of host and mobile IP addresses. A host autonomously configures its own link-local address, and booting nodes send router solicitations to request router advertisements for configuring interfaces.

IPv6 Applications

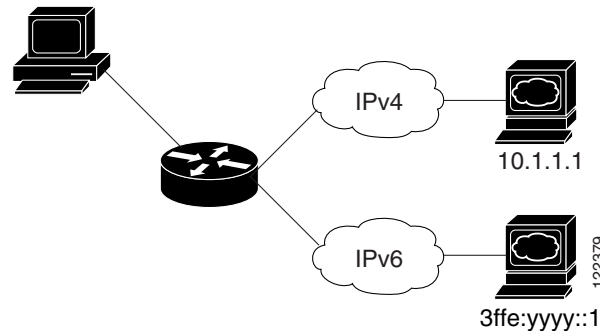
- Ping, traceroute, Telnet, TFTP, and FTP
- Secure Shell (SSH) over an IPv6 transport
- HTTP server access over IPv6 transport
- DNS resolver for AAAA over IPv4 transport
- Cisco Discovery Protocol (CDP) support for IPv6 addresses

Dual IPv4 and IPv6 Protocol Stacks

You must use the dual IPv4 and IPv6 template to allocate hardware memory usage to both IPv4 and IPv6 protocols.

Figure 3-1 shows a router forwarding both IPv4 and IPv6 traffic through the same interface, based on the IP packet and destination addresses.

Figure 3-1 Dual IPv4 and IPv6 Support on an Interface



Use the dual IPv4 and IPv6 switch database management (SDM) template to enable IPv6 routing dual stack environments (supporting both IPv4 and IPv6). For more information about the dual IPv4 and IPv6 SDM template, see the “Configuring SMD Templates” chapter in the *Cisco Connected Grid Switches System Management Software Configuration Guide*.

- If you try to configure IPv6 without first selecting a dual IPv4 and IPv6 template, a warning message appears.
- In IPv4-only environments, the switch routes IPv4 packets and applies IPv4 QoS and ACLs in hardware. IPv6 packets are not supported.
- In dual IPv4 and IPv6 environments, the switch routes both IPv4 and IPv6 packets and applies IPv4 QoS in hardware.
- IPv6 QoS is not supported.
- If you do not plan to use IPv6, do not use the dual stack template because it results in less hardware memory availability for each resource.

DHCP for IPv6 Address Assignment

DHCPv6 enables DHCP servers to pass configuration parameters, such as IPv6 network addresses, to IPv6 clients. The address assignment feature manages nonduplicate address assignment in the correct prefix based on the network where the host is connected. Assigned addresses can be from one or multiple prefix pools. Additional options, such as default domain and DNS name-server address, can be passed back to the client. Address pools can be assigned for use on a specific interface, on multiple interfaces, or the server can automatically find the appropriate pool.

Static Routes for IPv6

Static routes are manually configured and define an explicit route between two networking devices. Static routes are useful for smaller networks with only one path to an outside network or to provide security for certain types of traffic in a larger network.

RIP for IPv6

Routing Information Protocol (RIP) for IPv6 is a distance-vector protocol that uses hop count as a routing metric. It includes support for IPv6 addresses and prefixes and the all-RIP-routers multicast group address FF02::9 as the destination address for RIP update messages.

OSPF for IPv6

The switch supports Open Shortest Path First (OSPF) for IPv6, a link-state protocol for IP.

EIGRP IPv6

The switch supports Enhanced Interior Gateway Routing Protocol (EIGRP) for IPv6. It is configured on the interfaces on which it runs and does not require a global IPv6 address.

Before running, an instance of EIGRP IPv6 requires an implicit or explicit router ID. An implicit router ID is derived from a local IPv4 address, so any IPv4 node always has an available router ID. However, EIGRP IPv6 might be running in a network with only IPv6 nodes and therefore might not have an available IPv4 router ID.

Multiprotocol BGP for IPv6

Multiprotocol Border Gateway Protocol (BGP) is the supported exterior gateway protocol for IPv6. Multiprotocol BGP extensions for IPv6 support the same features and functionality as IPv4 BGP. IPv6 enhancements to multiprotocol BGP include support for IPv6 address family and network layer reachability information (NLRI) and next-hop (the next router in the path to the destination) attributes that use IPv6 addresses.

The switch does not support multicast BGP or non-stop forwarding (NSF) for IPv6 or for BGP IPv6.

SNMP and Syslog Over IPv6

To support both IPv4 and IPv6, IPv6 network management requires both IPv6 and IPv4 transports. Syslog over IPv6 supports address data types for these transports.

SNMP and syslog over IPv6 provide these features:

- Support for both IPv4 and IPv6
- IPv6 transport for SNMP and to modify the SNMP agent to support traps for an IPv6 host
- SNMP- and syslog-related MIBs to support IPv6 addressing
- Configuration of IPv6 hosts as trap receivers

For support over IPv6, SNMP modifies the existing IP transport mapping to simultaneously support IPv4 and IPv6. These SNMP actions support IPv6 transport management:

- Opens User Datagram Protocol (UDP) SNMP socket with default settings
- Provides a new transport mechanism called *SR_IPV6_TRANSPORT*
- Sends SNMP notifications over IPv6 transport
- Supports SNMP-named access lists for IPv6 transport
- Supports SNMP proxy forwarding using IPv6 transport
- Verifies SNMP Manager feature works with IPv6 transport

HTTP(S) Over IPv6

The HTTP client sends requests to both IPv4 and IPv6 HTTP servers, which respond to requests from both IPv4 and IPv6 HTTP clients. URLs with literal IPv6 addresses must be specified in hexadecimal using 16-bit values between colons.

The accept socket call chooses an IPv4 or IPv6 address family. The accept socket is either an IPv4 or IPv6 socket. The listening socket waits for both IPv4 and IPv6 signals that indicate a connection. The IPv6 listening socket is bound to an IPv6 wildcard address.

The underlying TCP/IP stack supports a dual-stack environment. HTTP relies on the TCP/IP stack and the sockets for processing network-layer interactions.

Basic network connectivity (**ping**) must exist between the client and the server hosts before HTTP connections can be made.

Unsupported IPv6 Unicast Routing Features

- IPv6 policy-based routing
- IPv6 virtual private network (VPN) routing and forwarding (VRF) table support
- Support for Intermediate System-to-Intermediate System (IS-IS) routing
- IPv6 packets destined to site-local addresses
- Tunneling protocols, such as IPv4-to-IPv6 or IPv6-to-IPv4
- The switch as a tunnel endpoint supporting IPv4-to-IPv6 or IPv6-to-IPv4 tunneling protocols
- IPv6 unicast reverse-path forwarding
- IPv6 general prefixes
- HSRP for IPv6

Prerequisites

Select a dual IPv4 and IPv6 template as described in the [“Dual IPv4 and IPv6 Protocol Stacks” section on page 3-5](#).

Guidelines and Limitations

Because IPv6 is implemented in switch hardware, some limitations occur due to the IPv6 compressed addresses in the hardware memory. This results in some loss of functionality and some feature limitations.

- When using user-network interface (UNI) or enhanced network interface (ENI) ports for any IPv6-related features, you must first globally enable IP routing and IPv6 routing on the switch by entering the **ip routing** and **ipv6 unicast-routing** global configuration commands even if you are not using IPv6 routing.
- ICMPv6 redirect functionality is not supported for IPv6 host routes (routes used to reach a specific host) or for IPv6 routes with masks greater than 64 bits. The switch cannot redirect hosts to a better first-hop router for a specific destination that is reachable through a host route or through a route with masks greater than 64 bits.

- Load balancing using equal cost and unequal cost routes is not supported for IPv6 host routes or for IPv6 routes with a mask greater than 64 bits.
- The switch cannot forward SNAP-encapsulated IPv6 packets.



Note There is a similar limitation for IPv4 SNAP-encapsulated packets, but the packets are dropped at the switch.

- The switch routes IPv6-to-IPv4 and IPv4-to-IPv6 packets in hardware, but the switch cannot be an IPv6-to-IPv4 or IPv4-to-IPv6 tunnel endpoint.
- Bridged IPv6 packets with hop-by-hop extension headers are forwarded in software. In IPv4, these packets are routed in software but bridged in hardware.
- In addition to the normal SPAN and RSPAN limitations defined in the software configuration guide, these limitations are specific to IPv6 packets:
 - When you send RSPAN IPv6-routed packets, the source MAC address in the SPAN output packet might be incorrect.
 - When you send RSPAN IPv6-routed packets, the destination MAC address might be incorrect. Normal traffic is not affected.
- The switch cannot apply QoS classification or policy-based routing on source-routed IPv6 packets in hardware.
- The switch cannot generate ICMPv6 *Packet Too Big* messages for multicast packets.

Default Settings

Feature	Default Setting
SDM template	Default.
IPv6 routing	Disabled globally and on all interfaces.
CEFv6	Disabled (IPv4 CEF is enabled by default). Note When IPv6 routing is enabled, CEFv6 is automatically enabled.
IPv6 addresses	None configured.

Configuring IPv6

- [Configuring IPv6 Addressing and Enabling IPv6 Routing, page 3-9](#)
- [Configuring Default Router Preference, page 3-11](#)
- [Configuring IPv4 and IPv6 Protocol Stacks, page 3-12](#)
- [Configuring DHCP for IPv6 Address Assignment, page 3-13](#)
- [Configuring IPv6 ICMP Rate Limiting, page 3-17](#)
- [Configuring CEF for IPv6, page 3-17](#)
- [Configuring Static Routing for IPv6, page 3-18](#)

- [Configuring RIP for IPv6, page 3-20](#)
- [Configuring OSPF for IPv6, page 3-22](#)
- [Configuring EIGRP for IPv6, page 3-23](#)
- [Configuring BGP for IPv6, page 3-24](#)

Configuring IPv6 Addressing and Enabling IPv6 Routing

To forward IPv6 traffic on an interface, you must configure a global IPv6 address on that interface. Configuring an IPv6 address on an interface automatically configures a link-local address and activates IPv6 for the interface. The configured interface automatically joins these required multicast groups for that link:

- solicited-node multicast group FF02:0:0:0:0:1:ff00::/104 for each unicast address assigned to the interface (the address for the neighbor discovery process)
- all-nodes link-local multicast group FF02::1
- all-routers link-local multicast group FF02::2

For more information about configuring IPv6 routing, see the “Implementing Addressing and Basic Connectivity for IPv6” chapter in the *IPv6 Implementation Guide, Cisco IOS Release 15.2M&T*.

BEFORE YOU BEGIN

- Be sure to select a dual IPv4 and IPv6 SDM template.
- Not all features discussed in this chapter are supported by the switch. See the “[Unsupported IPv6 Unicast Routing Features](#)” section on page 3-7.
- In the **ipv6 address** interface configuration command, you must enter the *ipv6-address* and *ipv6-prefix* variables with the address specified in hexadecimal using 16-bit values between colons. The *prefix-length* variable (preceded by a slash [/]) is a decimal value that shows how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	sdm prefer dual-ipv4-and-ipv6 { default routing vlan }	Select an SDM template that supports IPv4 and IPv6. <ul style="list-style-type: none"> • default—Set the switch to the default template to balance system resources. • routing—Set the switch to the routing template to support IPv4 and IPv6 routing, including IPv4 policy-based routing. • vlan—Maximize VLAN configuration on the switch with no routing supported in hardware.
Step 3	end	Return to privileged EXEC mode.
Step 4	reload	Reload the operating system.
Step 5	configure terminal	Enter global configuration mode.

	Command	Purpose
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure. The interface can be a physical interface, a switch virtual interface (SVI), or a Layer 3 EtherChannel.
Step 7	no switchport	Remove the interface from Layer 2 configuration mode (if it is a physical interface).
Step 8	ipv6 address <i>ipv6-prefix/prefix length</i> eui-64 or ipv6 address <i>ipv6-address</i> link-local or ipv6 enable	Specify a global IPv6 address with an extended unique identifier (EUI) in the low-order 64 bits of the IPv6 address. Specify only the network prefix; the last 64 bits are automatically computed from the switch MAC address. This enables IPv6 processing on the interface. Specify a link-local address on the interface to be used instead of the link-local address that is automatically configured when IPv6 is enabled on the interface. This command enables IPv6 processing on the interface. Automatically configure an IPv6 link-local address on the interface, and enable the interface for IPv6 processing. The link-local address can only be used to communicate with nodes on the same link.
Step 9	exit	Return to global configuration mode.
Step 10	ip routing	Enable IP routing on the switch.
Step 11	ipv6 unicast-routing	Enable forwarding of IPv6 unicast data packets.
Step 12	end	Return to privileged EXEC mode.
Step 13	show ipv6 interface <i>interface-id</i>	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove an IPv6 address from an interface, use the **no ipv6 address** *ipv6-prefix/prefix length* **eui-64** or **no ipv6 address** *ipv6-address* **link-local** interface configuration command. To remove all manually configured IPv6 addresses from an interface, use the **no ipv6 address** interface configuration command without arguments. To disable IPv6 processing on an interface that has not been explicitly configured with an IPv6 address, use the **no ipv6 enable** interface configuration command. To globally disable IPv6 routing, use the **no ipv6 unicast-routing** global configuration command.

EXAMPLE

This example shows how to enable IPv6 with both a link-local address and a global address based on the IPv6 prefix 2001:0DB8:c18:1::/64. The EUI-64 interface ID is used in the low-order 64 bits of both addresses. Output from the **show ipv6 interface** EXEC command is included to show how the interface ID (20B:46FF:FE2F:D940) is appended to the link-local prefix FE80::/64 of the interface.

```
Switch(config)# sdm prefer dual-ipv4-and-ipv6 default
Switch(config)# ipv6 unicast-routing
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# no switchport
Switch(config-if)# ipv6 address 2001:0DB8:c18:1::/64 eui 64
Switch(config-if)# end
Switch# show ipv6 interface gigabitethernet0/11
GigabitEthernet0/2 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::20B:46FF:FE2F:D940
```

```

Global unicast address(es):
2001:0DB8:c18:1:20B:46FF:FE2F:D940, subnet is 2001:0DB8:c18:1::/64 [EUI]
Joined group address(es):
  FF02::1
  FF02::2
  FF02::1:FF2F:D940
MTU is 1500 bytes
ICMP error messages limited to one every 100 milliseconds
ICMP redirects are enabled
ND DAD is enabled, number of DAD attempts: 1
ND reachable time is 30000 milliseconds
ND advertised reachable time is 0 milliseconds
ND advertised retransmit interval is 0 milliseconds
ND router advertisements are sent every 200 seconds
ND router advertisements live for 1800 seconds
Hosts use stateless autoconfig for addresses.

```

Configuring Default Router Preference

Router advertisement messages are sent with the default router preference (DRP) configured by the **ipv6 nd router-preference** interface configuration command. If no DRP is configured, router advertisements are sent with a medium preference.

A DRP is useful when two routers on a link might provide equivalent, but not equal-cost routing, and policy might dictate that hosts should prefer one of the routers.

BEFORE YOU BEGIN

Complete the [“Configuring IPv6 Addressing and Enabling IPv6 Routing” procedure on page 3-9](#).

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and enter the Layer 3 interface on which you want to specify the DRP.
Step 3	ipv6 nd router-preference { high medium low }	Specify a DRP for the router on the switch interface.
Step 4	end	Return to privileged EXEC mode.
Step 5	show ipv6 interface	Verify the configuration.
Step 6	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no ipv6 nd router-preference** interface configuration command to disable an IPv6 DRP.

EXAMPLE

This example shows how to configure a DRP of *high* for the router on an interface:

```

Switch# configure terminal
Switch(config)# interface gigabitethernet0/1
Switch(config-if)# ipv6 nd router-preference high
Switch(config-if)# end

```

Configuring IPv4 and IPv6 Protocol Stacks

Follow this procedure to configure a Layer 3 interface to support both IPv4 and IPv6 and to enable IPv6 routing.

BEFORE YOU BEGIN

Before configuring IPv6 routing, you must select an SDM template that supports IPv4 and IPv6. If not already configured, use the **sdm prefer dual-ipv4-and-ipv6 {default | routing | vlan} global** configuration command to configure a template that supports IPv6. When you select a new template, you must reload the switch by using the **reload** privileged EXEC command so that the template takes effect.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	sdm prefer dual-ipv4-and-ipv6 {default routing vlan}	Select an SDM template that supports IPv4 and IPv6. <ul style="list-style-type: none"> default—Set the switch to the default template to balance system resources. routing—Set the switch to the routing template to support IPv4 and IPv6 routing, including IPv4 policy-based routing. vlan—Maximize VLAN configuration on the switch with no routing supported in hardware.
Step 3	end	Return to privileged EXEC mode.
Step 4	reload	Reload the operating system.
Step 5	configure terminal	Enter global configuration mode.
Step 6	ip routing	Enable IPv4 routing on the switch.
Step 7	ipv6 unicast-routing	Enable forwarding of IPv6 data packets on the switch.
Step 8	interface interface-id	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 9	no switchport	Remove the interface from Layer 2 configuration mode (if it is a physical interface).
Step 10	ip address ip-address mask [secondary]	Specify a primary or secondary IPv4 address for the interface.
Step 11	ipv6 address ipv6-prefix/prefix length eui-64	Specify a global IPv6 address. Specify only the network prefix; the last 64 bits are automatically computed from the switch MAC address.
	or	
	ipv6 address ipv6-address link-local	Specify a link-local address on the interface to be used instead of the automatically configured link-local address when IPv6 is enabled on the interface.
	or	
	ipv6 enable	Automatically configure an IPv6 link-local address on the interface, and enable the interface for IPv6 processing. The link-local address can only be used to communicate with nodes on the same link.

	Command	Purpose
Step 12	end	Return to privileged EXEC mode.
Step 13	show interface <i>interface-id</i> show ip interface <i>interface-id</i> show ipv6 interface <i>interface-id</i>	Verify your entries.
Step 14	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable IPv4 routing, use the **no ip routing** global configuration command. To disable IPv6 routing, use the **no ipv6 unicast-routing** global configuration command. To remove an IPv4 address from an interface, use the **no ip address** *ip-address mask* interface configuration command. To remove an IPv6 address from an interface, use the **no ipv6 address** *ipv6-prefix/prefix length eui-64* or **no ipv6 address** *ipv6-address link-local* interface configuration command. To remove all manually configured IPv6 addresses from an interface, use the **no ipv6 address** interface configuration command without arguments. To disable IPv6 processing on an interface that has not been explicitly configured with an IPv6 address, use the **no ipv6 enable** interface configuration command.

EXAMPLE

This example shows how to enable IPv4 and IPv6 routing on an interface:

```
Switch(config)# sdm prefer dual-ipv4-and-ipv6 default
Switch(config)# ip routing
Switch(config)# ipv6 unicast-routing
Switch(config)# interface gigabitethernet0/2
Switch(config-if)# no switchport
Switch(config-if)# ip address 192.168.99.1 244.244.244.0
Switch(config-if)# ipv6 address 2001:0DB8:c18:1::/64 eui 64
Switch(config-if)# end
```

Configuring DHCP for IPv6 Address Assignment

This document describes only the DHCPv6 address assignment. For more information about configuring the DHCPv6 client, server, or relay agent functions, see the “Implementing DHCP for IPv6” chapter in the *IPv6 Implementation Guide, Cisco IOS Release 15.2M&T*.

- [Default DHCPv6 Address Assignment Configuration, page 3-13](#)
- [DHCPv6 Address Assignment Configuration Guidelines, page 3-13](#)
- [Enabling the DHCPv6 Server Function, page 3-14](#)
- [Enabling the DHCPv6 Client Function, page 3-16](#)

Default DHCPv6 Address Assignment Configuration

By default, no Dynamic Host Configuration Protocol for IPv6 (DHCPv6) features are configured on the switch.

DHCPv6 Address Assignment Configuration Guidelines

When configuring a DHCPv6 address assignment, consider these guidelines:

- In the procedures, the specified interface must be one of these Layer 3 interfaces:

- DHCPv6 IPv6 routing must be enabled on a Layer 3 interface.
- SVI: a VLAN interface created by using the **interface vlan** *vlan_id* command.
- EtherChannel port channel in Layer 3 mode: a port-channel logical interface created by using the **interface port-channel port-channel-number** command.
- Before configuring DHCPv6, you must select a Switch Database Management (SDM) template that supports IPv4 and IPv6.
- The switch can act as a DHCPv6 client, server, or relay agent. The DHCPv6 client, server, and relay function are mutually exclusive on an interface.

Enabling the DHCPv6 Server Function

BEFORE YOU BEGIN

See the “[DHCPv6 Address Assignment Configuration Guidelines](#)” section on page 3-13.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 dhcp pool <i>poolname</i>	Enter DHCP pool configuration mode, and define the name for the IPv6 DHCP pool. The pool name can be a symbolic string (such as Engineering) or an integer (such as 0).
Step 3	address prefix <i>IPv6-prefix</i> lifetime { <i>t1 t1</i> infinite }	(Optional) Specify an address prefix for address assignment. This address must be in hexadecimal, using 16-bit values between colons. <ul style="list-style-type: none"> • lifetime <i>t1 t1</i>—Specify a time interval (in seconds) that an IPv6 address prefix remains in the valid state. The range is 5 to 4294967295 seconds. Specify infinite for no time interval.
Step 4	link-address <i>IPv6-prefix</i>	(Optional) Specify a link-address IPv6 prefix. When an address on the incoming interface or a link-address in the packet matches the specified IPv6 prefix, the server uses the configuration information pool. This address must be in hexadecimal, using 16-bit values between colons.
Step 5	vendor-specific <i>vendor-id</i>	(Optional) Enter vendor-specific configuration mode, and enter a vendor-specific identification number. This number is the vendor IANA Private Enterprise Number. The range is 1 to 4294967295.

	Command	Purpose
Step 6	suboption <i>number</i> { address <i>IPv6-address</i> ascii <i>ASCII-string</i> hex <i>hex-string</i> }	(Optional) Enter a vendor-specific suboption number. The range is 1 to 65535. Enter an IPv6 address, ASCII text, or a hex string as defined by the suboption parameters.
Step 7	exit	Return to DHCP pool configuration mode.
Step 8	exit	Return to global configuration mode.
Step 9	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 10	ipv6 dhcp server [<i>poolname</i> automatic] [rapid-commit] [preference <i>value</i>] [allow-hint]	Enable the DHCPv6 server function on an interface. <ul style="list-style-type: none"> <i>poolname</i>—(Optional) User-defined name for the IPv6 DHCP pool. The pool name can be a symbolic string (such as Engineering) or an integer (such as 0). automatic—(Optional) Enables the system to automatically determine which pool to use when allocating addresses for a client. rapid-commit—(Optional) Allow two-message exchange method. preference value—(Optional) The preference value carried in the preference option in the advertise message sent by the server. The range is from 0 to 255. The preference value default is 0. allow-hint—(Optional) Specifies whether the server should consider client suggestions in the SOLICIT message. By default, the server ignores client hints.
Step 11	end	Return to privileged EXEC mode.
Step 12	show ipv6 dhcp pool or show ipv6 dhcp interface	Verify DHCPv6 pool configuration. Verify that the DHCPv6 server function is enabled on an interface.
Step 13	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To delete a DHCPv6 pool, use the **no ipv6 dhcp pool** *poolname* global configuration command. Use the **no** form of the DHCP pool configuration mode commands to change the DHCPv6 pool characteristics. To disable the DHCPv6 server function on an interface, use the **no ipv6 dhcp server** interface configuration command.

EXAMPLE

This example shows how to configure a pool called *engineering* with an IPv6 address prefix:

```
Switch# configure terminal
Switch(config)# ipv6 dhcp pool engineering
Switch(config-dhcpv6)#address prefix 2001:1000::0/64
Switch(config-dhcpv6)# end
```

This example shows how to configure a pool called *testgroup* with three link-addresses and an IPv6 address prefix:

```
Switch# configure terminal
Switch(config)# ipv6 dhcp pool testgroup
Switch(config-dhcpv6)# link-address 2001:1001::0/64
Switch(config-dhcpv6)# link-address 2001:1002::0/64
Switch(config-dhcpv6)# link-address 2001:2000::0/48
Switch(config-dhcpv6)# address prefix 2001:1003::0/64
Switch(config-dhcpv6)# end
```

This example shows how to configure a pool called *350* with vendor-specific options:

```
Switch# configure terminal
Switch(config)# ipv6 dhcp pool 350
Switch(config-dhcpv6)# address prefix 2001:1005::0/48
Switch(config-dhcpv6)# vendor-specific 9
Switch(config-dhcpv6-vs)# suboption 1 address 1000:235D::1
Switch(config-dhcpv6-vs)# suboption 2 ascii "IP-Phone"
Switch(config-dhcpv6-vs)# end
```

Enabling the DHCPv6 Client Function

BEFORE YOU BEGIN

See the [“DHCPv6 Address Assignment Configuration Guidelines”](#) section on page 3-13.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	interface <i>interface-id</i>	Enter interface configuration mode, and specify the interface to configure.
Step 3	ipv6 address dhcp [rapid-commit]	Enable the interface to acquire an IPv6 address from the DHCPv6 server. rapid-commit —(Optional) Allow two-message exchange method for address assignment.
Step 4	ipv6 dhcp client request [vendor-specific]	(Optional) Enable the interface to request the vendor-specific option.
Step 5	end	Return to privileged EXEC mode.
Step 6	show ipv6 dhcp interface	Verify that the DHCPv6 client is enabled on an interface.

To disable the DHCPv6 client function, use the **no ipv6 address dhcp** interface configuration command. To remove the DHCPv6 client request, use the **no ipv6 address dhcp client request** interface configuration command.

EXAMPLE

This example shows how to acquire an IPv6 address and to enable the rapid-commit option:

```
Switch(config)# interface gigabitethernet0/1
```

```
Switch(config-if)# ipv6 address dhcp rapid-commit
```

Configuring IPv6 ICMP Rate Limiting

ICMP rate limiting is enabled by default with a default interval between error messages of 100 milliseconds and a bucket size (maximum number of tokens to be stored in a bucket) of 10.

BEFORE YOU BEGIN

Complete the “Configuring IPv6 Addressing and Enabling IPv6 Routing” procedure on page 3-9.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 icmp error-interval <i>interval</i> [<i>bucketsize</i>]	Configure the interval and bucket size for IPv6 ICMP error messages: <ul style="list-style-type: none"> <i>interval</i>—The interval (in milliseconds) between tokens being added to the bucket. The range is from 0 to 2147483647 milliseconds. <i>bucketsize</i>—(Optional) The maximum number of tokens stored in the bucket. The range is from 1 to 200.
Step 3	end	Return to privileged EXEC mode.
Step 4	show ipv6 interface [<i>interface-id</i>]	Verify your entries.
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To return to the default configuration, use the **no ipv6 icmp error-interval** global configuration command.

EXAMPLE

This example shows how to configure an IPv6 ICMP error message interval of 50 milliseconds and a bucket size of 20 tokens:

```
Switch(config)#ipv6 icmp error-interval 50 20
```

Configuring CEF for IPv6

Cisco Express Forwarding (CEF) is a Layer 3 IP switching technology, allowing more CPU processing power to be dedicated to packet forwarding. IPv4 CEF is enabled by default. IPv6 CEF is disabled by default, but automatically enabled when you configure IPv6 routing.

To route IPv6 unicast packets, first globally configure forwarding of IPv6 unicast packets by using the **ipv6 unicast-routing** global configuration command. You must also configure an IPv6 address and IPv6 processing on an interface by using the **ipv6 address** interface configuration command.

To disable IPv6 CEF, use the **no ipv6 cef** global configuration command. To reenabling IPv6 CEF, use the **ipv6 cef** global configuration command. You can verify the IPv6 state by entering the **show ipv6 cef** privileged EXEC command.

For more information about configuring CEF, see the “Implementing IPv6 Addressing and Basic Connectivity” chapter in the *IPv6 Implementation Guide, Cisco IOS Release 15.2M&T*.

Configuring Static Routing for IPv6

BEFORE YOU BEGIN

Before configuring a static IPv6 route, you must:

- Enable routing by using the **ip routing** global configuration command.
- Enable the forwarding of IPv6 packets by using the **ipv6 unicast-routing** global configuration command.
- Enable IPv6 on at least one Layer 3 interface by configuring an IPv6 address on the interface.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 route <i>ipv6-prefix/prefix length</i> { <i>ipv6-address</i> <i>interface-id</i> [<i>ipv6-address</i>] } [<i>administrative distance</i>]	<p>Configure a static IPv6 route.</p> <ul style="list-style-type: none"> • <i>ipv6-prefix</i>—The IPv6 network that is the destination of the static route. It can also be a hostname when static host routes are configured. • <i>/prefix length</i>—The length of the IPv6 prefix. A decimal value that shows how many of the high-order contiguous bits comprise the prefix (the network portion of the address). A slash mark must precede the decimal value. • <i>ipv6-address</i>—The IPv6 address of the next hop that can be used to reach the specified network. The next hop does not need to be directly connected; recursion finds the IPv6 address of the directly connected next hop. The address must be specified in hexadecimal using 16-bit values between colons. • <i>interface-id</i>—Specify direct static routes from point-to-point and broadcast interfaces. On point-to-point interfaces, you do not need to specify the IPv6 address of the next hop. On broadcast interfaces, you should always specify the IPv6 address of the next hop, or ensure that the specified prefix is assigned to the link, specifying a link-local address as the next hop. You can optionally specify the IPv6 address of the next hop to which packets are sent. <p>Note You must specify an <i>interface-id</i> when using a link-local address as the next hop. The link-local next hop must be an adjacent router.</p> <ul style="list-style-type: none"> • <i>administrative distance</i>—(Optional) An administrative distance. The range is 1 to 254; the default value is 1, which gives static routes precedence over all but connected routes. To configure a floating static route, use an administrative distance greater than that of the dynamic routing protocol.
Step 3	end	Return to privileged EXEC mode.

	Command	Purpose
Step 4	show ipv6 static [<i>ipv6-address</i> <i>ipv6-prefix/prefix length</i>] [interface <i>interface-id</i>] [recursive] [detail]	Verify your entries by displaying the IPv6 routing table. <ul style="list-style-type: none"> • interface <i>interface-id</i>—(Optional) Display only those static routes with the specified interface as an egress interface. • recursive—(Optional) Display only recursive static routes. The recursive keyword is mutually exclusive with the interface keyword, but it can be used with or without the IPv6 prefix in the command syntax. • detail—(Optional) Display this additional information: <ul style="list-style-type: none"> – For valid recursive routes, the output path set, and maximum resolution depth. – For invalid routes, the reason why the route is not valid.
	or show ipv6 route static [<i>updated</i>]	
Step 5	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To remove a configured static route, use the **no ipv6 route** *ipv6-prefix/prefix length* {*ipv6-address* | *interface-id* [*ipv6-address*]} [*administrative distance*] global configuration command.

For more information about configuring static IPv6 routing, see the “Implementing Static Routes for IPv6” chapter in the *IPv6 Implementation Guide, Cisco IOS Release 15.2M&T*.

EXAMPLE

This example shows how to configure a floating static route to an interface. The route has an administrative distance of 130:

```
Switch(config)# ipv6 route 2001:0DB8::/32 gigabitethernet0/1 130
```

Configuring RIP for IPv6

BEFORE YOU BEGIN

Before configuring the switch to run IPv6 RIP, you must:

- Enable routing by using the **ip routing** global configuration command.
- Enable the forwarding of IPv6 packets by using the **ipv6 unicast-routing** global configuration command.
- Enable IPv6 on any Layer 3 interfaces on which IPv6 RIP is to be enabled.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 router rip <i>name</i>	Configure an IPv6 RIP routing process, and enter router configuration mode for the process.

	Command	Purpose
Step 3	maximum-paths <i>number-paths</i>	(Optional) Define the maximum number of equal-cost routes that IPv6 RIP can support. The range is from 1 to 64, and the default is 4 routes.
Step 4	exit	Return to global configuration mode.
Step 5	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 6	ipv6 rip <i>name</i> enable	Enable the specified IPv6 RIP routing process on the interface.
Step 7	ipv6 rip <i>name</i> default-information { only originate }	<p>(Optional) Originate the IPv6 default route (::/0) into the RIP routing process updates sent from the specified interface.</p> <p>Note To avoid routing loops after the IPv6 default route (::/0) is originated from any interface, the routing process ignores all default routes received on any interface.</p> <ul style="list-style-type: none"> • only—Select to originate the default route, but suppress all other routes in the updates sent on this interface. • originate—Select to originate the default route in addition to all other routes in the updates sent on this interface.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ipv6 rip [<i>name</i>] [interface <i>interface-id</i>] [database] [next-hops] or show ipv6 route rip [<i>updated</i>]	Display information about current IPv6 RIP processes. Display the current contents of the IPv6 routing table.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable a RIP routing process, use the **no ipv6 router rip** *name* global configuration command. To disable the RIP routing process for an interface, use the **no ipv6 rip** *name* interface configuration command.

For more information about configuring RIP routing for IPv6, see the “Implementing RIP for IPv6” chapter in the *IPv6 Implementation Guide, Cisco IOS Release 15.2M&T*.

EXAMPLE

This example shows how to enable the RIP routing process *cisco* with a maximum of eight equal-cost routes and to enable it on an interface:

```
Switch(config)# ipv6 router rip cisco
Switch(config-router)# maximum-paths 8
Switch(config)# exit
Switch(config)# interface gigabitethernet0/3
Switch(config-if)# ipv6 rip cisco enable
```

Configuring OSPF for IPv6

You can customize OSPF for IPv6 for your network. However, the defaults are set to meet the requirements of most customers and features.

Be careful when changing the defaults for IPv6 commands. Doing so might adversely affect OSPF for the IPv6 network.

BEFORE YOU BEGIN

Before you enable IPv6 OSPF on an interface, you must:

- Enable routing by using the **ip routing** global configuration command.
- Enable the forwarding of IPv6 packets by using the **ipv6 unicast-routing** global configuration command.
- Enable IPv6 on Layer 3 interfaces on which you are enabling IPv6 OSPF.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ipv6 router ospf <i>process-id</i>	Enable OSPF router configuration mode for the process. The process ID is the number assigned administratively when enabling the OSPF for IPv6 routing process. It is locally assigned and can be a positive integer from 1 to 65535.
Step 3	area <i>area-id</i> range {<i>ipv6-prefix/prefix length</i>} [advertise not-advertise] [cost <i>cost</i>]	(Optional) Consolidate and summarize routes at an area boundary. <ul style="list-style-type: none"> • <i>area-id</i>—Identifier of the area about which routes are to be summarized. It can be specified as either a decimal value or as an IPv6 prefix. • <i>ipv6-prefix/prefix length</i>—The destination IPv6 network and a decimal value that shows how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark (/) must precede the decimal value. • advertise—(Optional) Set the address range status to advertise and to generate a Type 3 summary link-state advertisement (LSA). • not-advertise—(Optional) Set the address range status to DoNotAdvertise. The Type 3 summary LSA is suppressed, and component networks remain hidden from other networks. • cost <i>cost</i>—(Optional) Metric or cost for this summary route, which is used during OSPF SPF calculation to determine the shortest paths to the destination. The value can be 0 to 16777215.

	Command	Purpose
Step 4	maximum paths <i>number-paths</i>	(Optional) Define the maximum number of equal-cost routes to the same destination that IPv6 OSPF should enter in the routing table. The range is from 1 to 64, and the default is 16 paths.
Step 5	exit	Return to global configuration mode.
Step 6	interface <i>interface-id</i>	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 7	ipv6 ospf <i>process-id</i> area <i>area-id</i> [instance <i>instance-id</i>]	Enable OSPF for IPv6 on the interface. <ul style="list-style-type: none"> instance <i>instance-id</i>—(Optional) Instance identifier.
Step 8	end	Return to privileged EXEC mode.
Step 9	show ipv6 ospf [<i>process-id</i>] [<i>area-id</i>] interface [<i>interface-id</i>] or show ipv6 ospf [<i>process-id</i>] [<i>area-id</i>]	Display information about OSPF interfaces. Display general information about OSPF routing processes.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

To disable an OSPF routing process, use the **no ipv6 router ospf** *process-id* global configuration command. To disable the OSPF routing process for an interface, use the **no ipv6 ospf** *process-id* **area** *area-id* interface configuration command.

For more information about configuring OSPF routing for IPv6, see the “Implementing OSPF for IPv6” chapter in the [IPv6 Implementation Guide, Cisco IOS Release 15.2M&T](#).

Configuring EIGRP for IPv6

By default, EIGRP for IPv6 is disabled. You can configure EIGRP for IPv6 on an interface. After configuring the router and the interface for EIGRP, enter the **no shutdown** privileged EXEC command to start EIGRP.



Note

If EIGRP for IPv6 is not in shutdown mode, EIGRP might start running before you enter the EIRGP router-mode commands to configure the router and the interface.

To set an explicit router ID, use the **show ipv6 eigrp** command to see the configured router IDs, and then use the **router-id** command.

As with EIGRP IPv4, you can use EIGRPv6 to specify your EIGRP IPv4 interfaces and to select a subset of those as passive interfaces. Use the **passive-interface default** command to make all interfaces passive, and then use the **no passive-interface** command on selected interfaces to make them active. EIGRP IPv6 does not need to be configured on a passive interface.

For more configuration procedures, see the “Implementing EIGRP for IPv6” chapter in the [IPv6 Implementation Guide, Cisco IOS Release 15.2M&T](#).

Configuring BGP for IPv6

When configuring multiprotocol BGP extensions for IPv6, you must create the BGP routing process, configure peering relationships, and customize BGP for your particular network. Note that BGP functions the same in IPv6 as in IPv4.

BEFORE YOU BEGIN

Before configuring the router to run BGP for IPv6, you must use the **ipv6 unicast-routing** command to globally enable IPv6 routing.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	router bgp <i>as-number</i>	Configure a BGP routing process, and enter BGP router configuration mode for the autonomous system number.
Step 3	no bgp default ipv4-unicast	Disable the IPv4 unicast address family for the BGP routing process specified in the previous step. Routing information for the IPv4 unicast address family is advertised by default for each BGP routing session unless you enter this command before configuring the neighbor remote-as command.
Step 4	bgp router-id <i>ip-address</i>	(Optional) Configure a fixed 32-bit router ID as the identifier of the local router running BGP. By default, the router ID is the IPv4 address of a router loopback interface. On a router enabled only for IPv6 (no IPv4 address), you must manually configure the BGP router ID. Note Configuring a router ID by using this command resets all active BGP peering sessions.
Step 5	neighbor { <i>ip-address</i> <i>ipv6-address</i> [%] <i>interface-type interface-number</i> <i>peer-group-name</i> } remote-as <i>as-number</i>	Add the IPv6 address of the neighbor in the specified autonomous system to the IPv6 multiprotocol BGP neighbor table of the local router. Note The <i>ipv6-address</i> must be in hexadecimal, using 16-bit values between colons.
Step 6	address-family ipv6	Specify the IPv6 address family and enter address family configuration mode
Step 7	neighbor { <i>ip-address</i> <i>peer-group-name</i> <i>ipv6-address</i> } activate	Enable the neighbor to exchange prefixes for the IPv6 address family with the local router.
Step 8	end	Return to privileged EXEC mode.
Step 9	show bgp ipv6	Display information about IPv6 BGP configuration.
Step 10	copy running-config startup-config	(Optional) Save your entries in the configuration file.

For more configuration procedures, see the “Implementing Multiprotocol BGP for IPv6” chapter in the *IPv6 Implementation Guide, Cisco IOS Release 15.2M&T*.

The switch does not support multicast IPv6 BGP, nonstop forwarding (NSF) for IPv6 BGP, 6PE multipath (EoMPLS), or IPv6 VRF.

EXAMPLE

```
router bgp 1
  bgp router-id 1.1.1.1
  no bgp default ipv4-unicast
!--- Without configuring '"no bgp default ipv4-unicast"' only IPv4 will be
!--- advertised
  bgp log-neighbor-changes
  neighbor 2010:AB8:0:2:C601:10FF:FE58:0 remote-as 2
  !
  address-family ipv6
    neighbor 2010:AB8:0:2:C601:10FF:FE58:0 activate
    network 2010:AB8:2::/48
    network 2010:AB8:3::/48
  exit-address-family
!
```

Verifying Configuration

Command	Purpose
show bgp ipv6	Display BGP IPv6 configuration and routing tables.
show ipv6 access-list	Display IPv6 access lists.
show ipv6 cef	Display Cisco Express Forwarding for IPv6.
show ipv6 interface <i>interface-id</i>	Display IPv6 interface status and configuration.
show ipv6 mtu	Display IPv6 MTU per destination cache.
show ipv6 neighbors	Display IPv6 neighbor cache entries.
show ipv6 ospf	Display IPv6 OSPF information.
show ipv6 prefix-list	Display IPv6 prefix lists.
show ipv6 protocols	Display IPv6 routing protocols on the switch.
show ipv6 rip	Display IPv6 RIP routing protocol status.
show ipv6 route	Display IPv6 route table entries.
show ipv6 routers	Display local IPv6 routers.
show ipv6 static	Display IPv6 static routes.
show ipv6 traffic	Display IPv6 traffic statistics.

Command	Purpose
show ipv6 eigrp <i>[as-number] interface</i>	Display information about interfaces configured for EIGRP IPv6.
show ipv6 eigrp <i>[as-number] neighbor</i>	Display the neighbors discovered by EIGRP IPv6.

Command	Purpose
show ipv6 eigrp [<i>as-number</i>] <i>traffic</i>	Display the number of EIGRP IPv6 packets sent and received.
show ipv6 eigrp topology [<i>as-number</i> <i>ipv6-address</i>] [active all-links detail-links pending summary zero-successors]	Display EIGRP entries in the IPv6 topology table.

Command	Purpose
show ip http server history	Display the previous 20 connections to the HTTP server, including the IP address accessed and the time when the connection was closed.
show ip http server connection	Display the current connections to the HTTP server, including the local and remote IP addresses being accessed.
show ip http client connection	Display the configuration values for HTTP client connections to HTTP servers.
show ip http client history	Display a list of the last 20 requests made by the HTTP client to the server.

Configuration Example

This is an example of the output from the **show ipv6 interface** privileged EXEC command:

```
Switch# show ipv6 interface
Vlan1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::20B:46FF:FE2F:D940
  Global unicast address(es):
    3FFE:C000:0:1:20B:46FF:FE2F:D940, subnet is 3FFE:C000:0:1::/64 [EUI]
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF2F:D940
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
<output truncated>
```

This is an example of the output from the **show ipv6 cef** privileged EXEC command:

```
Switch# show ipv6 cef
::/0
  nexthop 3FFE:C000:0:7::777 Vlan7
3FFE:C000:0:1::/64
  attached to Vlan1
3FFE:C000:0:1:20B:46FF:FE2F:D940/128
  receive
3FFE:C000:0:7::/64
  attached to Vlan7
```

```

3FFE:C000:0:7::777/128
    attached to Vlan7
3FFE:C000:0:7:20B:46FF:FE2F:D97F/128
    receive
3FFE:C000:111:1::/64
    attached to GigabitEthernet0/11
3FFE:C000:111:1:20B:46FF:FE2F:D945/128
    receive
3FFE:C000:168:1::/64
    attached to GigabitEthernet0/43
3FFE:C000:168:1:20B:46FF:FE2F:D94B/128
    receive
3FFE:C000:16A:1::/64
    attached to Loopback10
3FFE:C000:16A:1:20B:46FF:FE2F:D900/128
    receive

<output truncated>

```

This is an example of the output from the **show ipv6 protocols** privileged EXEC command:

```

Switch# show ipv6 protocols
IPv6 Routing Protocol is "connected"
IPv6 Routing Protocol is "static"
IPv6 Routing Protocol is "rip fer"
  Interfaces:
    Vlan6
  GigabitEthernet0/4
  GigabitEthernet0/11
  GigabitEthernet0/12
  Redistribution:
    None

```

This is an example of the output from the **show ipv6 rip** privileged EXEC command:

```

Switch# show ipv6 rip
RIP process "fer", port 521, multicast-group FF02::9, pid 190
  Administrative distance is 120. Maximum paths is 16
  Updates every 30 seconds, expire after 180
  Holddown lasts 0 seconds, garbage collect after 120
  Split horizon is on; poison reverse is off
  Default routes are not generated
  Periodic updates 9040, trigger updates 60
  Interfaces:
    Vlan6
  GigabitEthernet0/4
  GigabitEthernet0/11
  GigabitEthernet0/12
  Redistribution:
    None

```

This is an example of the output from the **show ipv6 neighbor** privileged EXEC command:

```

Switch# show ipv6 neighbors
IPv6 Address                               Age Link-layer Addr State Interface
3FFE:C000:0:7::777                         - 0007.0007.0007 REACH V17
3FFE:C101:113:1::33                       - 0000.0000.0033 REACH Gi0/13

```

This is an example of the output from the **show ipv6 static** privileged EXEC command:

```

Switch# show ipv6 static
IPv6 Static routes
Code: * - installed in RIB
* ::/0 via nexthop 3FFE:C000:0:7::777, distance 1

```

This is an example of the output from the **show ipv6 route** privileged EXEC command:

```
Switch# show ipv6 route
IPv6 Routing Table - 21 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S   ::/0 [1/0]
    via 3FFE:C000:0:7::777
C   3FFE:C000:0:1::/64 [0/0]
    via ::, Vlan1
L   3FFE:C000:0:1:20B:46FF:FE2F:D940/128 [0/0]
    via ::, Vlan1
C   3FFE:C000:0:7::/64 [0/0]
    via ::, Vlan7
L   3FFE:C000:0:7:20B:46FF:FE2F:D97F/128 [0/0]
    via ::, Vlan7
C   3FFE:C000:111:1::/64 [0/0]
    via ::, GigabitEthernet0/11
L   3FFE:C000:111:1:20B:46FF:FE2F:D945/128 [0/0]
C   3FFE:C000:168:1::/64 [0/0]
    via ::, GigabitEthernet0/4
L   3FFE:C000:168:1:20B:46FF:FE2F:D94B/128 [0/0]
    via ::, GigabitEthernet0/4
C   3FFE:C000:16A:1::/64 [0/0]
    via ::, Loopback10
L   3FFE:C000:16A:1:20B:46FF:FE2F:D900/128 [0/0]
    via ::, Loopback10

<output truncated>
```

This is an example of the output from the **show ipv6 traffic** privileged EXEC command.

```
Switch# show ipv6 traffic
IPv6 statistics:
  Rcvd: 1 total, 1 local destination
        0 source-routed, 0 truncated
        0 format errors, 0 hop count exceeded
        0 bad header, 0 unknown option, 0 bad source
        0 unknown protocol, 0 not a router
        0 fragments, 0 total reassembled
        0 reassembly timeouts, 0 reassembly failures
  Sent: 36861 generated, 0 forwarded
        0 fragmented into 0 fragments, 0 failed
        0 encapsulation failed, 0 no route, 0 too big
        0 RPF drops, 0 RPF suppressed drops
  Mcast: 1 received, 36861 sent

ICMP statistics:
  Rcvd: 1 input, 0 checksum errors, 0 too short
        0 unknown info type, 0 unknown error type
        unreachable: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
        parameter: 0 error, 0 header, 0 option
        0 hopcount expired, 0 reassembly timeout, 0 too big
        0 echo request, 0 echo reply
        0 group query, 0 group report, 0 group reduce
        1 router solicit, 0 router advert, 0 redirects
        0 neighbor solicit, 0 neighbor advert
  Sent: 10112 output, 0 rate-limited
        unreachable: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
        parameter: 0 error, 0 header, 0 option
        0 hopcount expired, 0 reassembly timeout, 0 too big
```

```
0 echo request, 0 echo reply
0 group query, 0 group report, 0 group reduce
0 router solicit, 9944 router advert, 0 redirects
84 neighbor solicit, 84 neighbor advert
```

UDP statistics:

```
Rcvd: 0 input, 0 checksum errors, 0 length errors
      0 no port, 0 dropped
Sent: 26749 output
```

TCP statistics:

```
Rcvd: 0 input, 0 checksum errors
Sent: 0 output, 0 retransmitted
```

Related Documents

For information about how Cisco Systems implements IPv6:

- http://www.cisco.com/en/US/products/ps6553/products_ios_technology_home.html

For information about IPv6 and other features in this chapter:

- [IPv6 Configuration Library, Cisco IOS Release 15M&T](#)
- [IPv6 Implementation Guide, Cisco IOS Release 15.2M&T](#)
- [Cisco Connected Grid Switches Security Software Configuration Guide](#)
- [Cisco Connected Grid Switches System Management Software Configuration Guide](#)



Configuring Enhanced Object Tracking

This chapter describes how to configure enhanced object tracking on the Cisco Industrial Ethernet 2000U Series Switches, hereafter referred to as IE 2000U or switch. This feature provides a more complete alternative to the Hot Standby Routing Protocol (HSRP) tracking mechanism, which allows you to track the line-protocol state of an interface. If the line protocol state of an interface goes down, the HSRP priority of the interface is reduced and another HSRP device with a higher priority becomes active. The enhanced object tracking feature separates the tracking mechanism from HSRP and creates a separate, standalone tracking process that can be used by processes other than HSRP. This allows tracking other objects in addition to the interface line-protocol state.

A client process, such as HSRP or Gateway Local Balancing Protocol (GLBP), can register an interest in tracking objects and request notification when the tracked object changes state. This feature increases the availability and speed of recovery of a routing system and decreases outages and outage duration.

For more information about enhanced object tracking and the commands used to configure it, see the [“Related Documents” section on page 4-13](#).

The chapter includes these sections:

- [Information About Enhanced Object Tracking, page 4-1](#)
- [Prerequisites, page 4-2](#)
- [Guidelines and Limitations, page 4-2](#)
- [Default Settings, page 4-2](#)
- [Configuring Enhanced Object Tracking Features, page 4-2](#)
- [Verifying Configuration, page 4-10](#)
- [Configuration Example, page 4-11](#)
- [Related Documents, page 4-13](#)

Information About Enhanced Object Tracking

Each tracked object has a unique number that is specified in the tracking command-line interface (CLI). Client processes use this number to track a specific object. The tracking process periodically polls the tracked object for value changes and sends any changes (as up or down values) to interested client processes, either immediately or after a specified delay. Several clients can track the same object, and can take different actions when the object changes state.

You can also track a combination of objects in a list by using either a weight threshold or a percentage threshold to measure the state of the list. You can combine objects using Boolean logic. A tracked list with a Boolean “AND” function requires that each object in the list be in an up state for the tracked object to be up. A tracked list with a Boolean “OR” function needs only one object in the list to be in the up state for the tracked object to be up.

Prerequisites

Your IP network is operational and you can access the destination device.

Guidelines and Limitations

Although up to 500 tracked objects can be configured, each tracked object uses CPU resources. The amount of available CPU resources on a switch is dependent upon variables such as traffic load and how other protocols are configured and run. The ability to use 500 tracked objects is dependent upon the available CPU. Testing should be conducted on site to ensure that the service works under the specific site traffic conditions.

Default Settings

No type of object tracking is configured.

Configuring Enhanced Object Tracking Features

- [Tracking Interface Line-Protocol or IP Routing State, page 4-2](#)
- [Configuring a Tracked List, page 4-3](#)
- [Configuring HSRP Object Tracking, page 4-7](#)
- [Configuring Other Tracking Characteristics, page 4-10](#)

Tracking Interface Line-Protocol or IP Routing State

You can track either the interface line protocol state or the interface IP routing state.

When you track the IP routing state, these three conditions are required for the object to be up:

- IP routing must be enabled and active on the interface.
- The interface line-protocol state must be up.
- The interface IP address must be known.

If all three of these conditions are not met, the IP routing state is down.

BEFORE YOU BEGIN

An object must exist before it can be added to a tracked list.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	track <i>object-number</i> interface <i>interface-id</i> line-protocol	(Optional) Create a tracking list to track the line-protocol state of an interface and enter tracking configuration mode. <ul style="list-style-type: none"> The <i>object-number</i> identifies the tracked object. The range is from 1 to 500. The interface <i>interface-id</i> is the interface being tracked.
Step 3	delay {up <i>seconds</i> [down <i>seconds</i>] [up <i>seconds</i>] down <i>seconds</i>}	(Optional) Specify a period of time in seconds to delay communicating state changes of a tracked object. The range is from 1 to 180 seconds.
Step 4	exit	Return to global configuration mode.
Step 5	track <i>object-number</i> interface <i>interface-id</i> ip routing	(Optional) Create a tracking list to track the IP routing state of an interface, and enter tracking configuration mode. IP-route tracking tracks an IP route in the routing table and the ability of an interface to route IP packets. <ul style="list-style-type: none"> The <i>object-number</i> identifies the tracked object. The range is from 1 to 500. The interface <i>interface-id</i> is the interface being tracked.
Step 6	delay {up <i>seconds</i> [down <i>seconds</i>] [up <i>seconds</i>] down <i>seconds</i>}	(Optional) Specify a period of time in seconds to delay communicating state changes of a tracked object. The range is from 1 to 180 seconds.
Step 7	end	Return to privileged EXEC mode.
Step 8	show track <i>object-number</i>	Verify that the specified objects are being tracked.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

This example configures the tracking of an interface line-protocol state and verifies the configuration:

```
Switch(config)# track 33 interface gigabitethernet0/1 line-protocol
Switch(config-track)# end
Switch# show track 33
Track 33
  Interface GigabitEthernet0/1 line-protocol
  Line protocol is Down (hw down)
    1 change, last change 00:18:28
```

Configuring a Tracked List

You can configure a tracked list of objects with a Boolean expression, a weight threshold, or a percentage threshold. A tracked list contains one or more objects. An object must exist before it can be added to the tracked list.

- You configure a Boolean expression to specify calculation by using either “AND” or “OR” operators.

- When you measure the tracked list state by a weight threshold, you assign a weight number to each object in the tracked list. The state of the tracked list is determined by whether or not the threshold was met. The state of each object is determined by comparing the total weight of all objects against a threshold weight for each object.
- When you measure the tracked list by a percentage threshold, you assign a percentage threshold to all objects in the tracked list. The state of each object is determined by comparing the assigned percentages of each object to the list.

Configuring a Tracked List with a Boolean Expression

Configuring a tracked list with a Boolean expression enables calculation by using either “AND” or “OR” operators. For example, when tracking two interfaces using the “AND” operator, *up* means that both interfaces are up, and *down* means that either interface is down.

BEFORE YOU BEGIN

An object must exist before it can be added to a tracked list.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	track <i>track-number</i> list boolean { and or }	Configure a tracked list object, and enter tracking configuration mode. The <i>track-number</i> can be from 1 to 500. <ul style="list-style-type: none"> • boolean—Specify the state of the tracked list based on a Boolean calculation. • and—Specify that the list is up if all objects are up or down if one or more objects are down. • or—Specify that the list is up if one object is up or down if all objects are down.
Step 3	object <i>object-number</i> [not]	Specify the object to be tracked. The range is from 1 to 500. The keyword not negates the state of the object, which means that when the object is up, the tracked list detects the object as down.
Step 4	delay { up <i>seconds</i> [down <i>seconds</i>] [up <i>seconds</i>] down <i>seconds</i> }	(Optional) Specify a period of time in seconds to delay communicating state changes of a tracked object. The range is from 1 to 180 seconds.
Step 5	end	Return to privileged EXEC mode.
Step 6	show track <i>object-number</i>	Verify that the specified objects are being tracked.
Step 7	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no track** *track-number* global configuration command to delete the tracked list.

EXAMPLE

This example configures track list 4 with a Boolean AND expression that contains two objects with one object state negated. If the list is up, the list detects that object 2 is down:

```
Switch(config)# track 4 list boolean and
Switch(config-track)# object 1
Switch(config-track)# object 2 not
Switch(config-track)# exit
```

Configuring a Tracked List with a Weight Threshold

To track by weight threshold, configure a tracked list of objects, specify that weight is used as the threshold, and configure a weight for each of its objects. The state of each object is determined by comparing the total weight of all objects that are up against a threshold weight for each object.

You cannot use the Boolean “NOT” operator in a weight threshold list.

BEFORE YOU BEGIN

An object must exist before it can be added to a tracked list.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	track <i>track-number</i> list threshold weight	Configure a tracked list object and enter tracking configuration mode. The <i>track-number</i> can be from 1 to 500. <ul style="list-style-type: none"> threshold—Specify the state of the tracked list based on a threshold. weight—Specify that the threshold is based on weight.
Step 3	object <i>object-number</i> [weight <i>weight-number</i>]	Specify the object to be tracked. The range is from 1 to 500. <ul style="list-style-type: none"> weight <i>weight-number</i>—(Optional) Specify a threshold weight for the object. The range is from 1 to 255.
Step 4	threshold weight { up <i>number</i> [down <i>number</i>]}	Specify the threshold weight. <ul style="list-style-type: none"> up <i>number</i>—The valid range is from 1 to 255. down <i>number</i>—(Optional) The range depends on the number selected for the up <i>number</i>. If you configure the up <i>number</i> as 25, the range shown for the down number is 0 to 24.
Step 5	delay { up <i>seconds</i> [down <i>seconds</i>] [up <i>seconds</i>] down <i>seconds</i> }	(Optional) Specify a period of time in seconds to delay communicating state changes of a tracked object. The range is from 1 to 180 seconds.
Step 6	end	Return to privileged EXEC mode.
Step 7	show track <i>object-number</i>	Verify that the specified objects are being tracked.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no track** *track-number* global configuration command to delete the tracked list.

EXAMPLE

This example configures track list 4 to track by weight threshold. If object 1 and object 2 are down, then track list 4 is up because object 3 satisfies the up threshold value of 30. But if object 3 is down, both objects 1 and 2 must be up in order to satisfy the threshold weight.

```
Switch(config)# track 4 list threshold weight
Switch(config-track)# object 1 weight 15
Switch(config-track)# object 2 weight 20
Switch(config-track)# object 3 weight 30
Switch(config-track)# threshold weight up 30 down 10
Switch(config-track)# exit
```

This configuration can be useful if object 1 and object 2 represent two small bandwidth connections and object 3 represents one large bandwidth connection. The configured **down 10** value means that once the tracked object is up, it will not go down until the threshold value is equal to or lower than 10, which in this example means that all connections are down.

Configuring a Tracked List with a Percentage Threshold

To track by percentage threshold, configure a tracked list of objects, specify that a percentage will be used as the threshold, and specify a percentage for all objects in the list. The state of the list is determined by comparing the assigned percentage of each object to the list.

You cannot use the Boolean “NOT” operator in a percentage threshold list.

BEFORE YOU BEGIN

An object must exist before it can be added to a tracked list.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	track <i>track-number</i> list threshold percentage	Configure a tracked list object and enter tracking configuration mode. The <i>track-number</i> can be from 1 to 500. <ul style="list-style-type: none"> threshold—Specify the state of the tracked list based on a threshold. percentage—Specify that the threshold is based on percentage.
Step 3	object <i>object-number</i>	Specify the object to be tracked. The range is from 1 to 500. Note An object must exist before you can add it to a tracked list.
Step 4	threshold percentage { up <i>number</i> [down <i>number</i>]} <i>number</i>	Specify the threshold percentage. <ul style="list-style-type: none"> up <i>number</i>—The valid range is from 1 to 100. down <i>number</i>—(Optional) The range depends on the number selected for the up <i>number</i>. If you configure the up <i>number</i> as 25, the range shown for the down number is 0 to 24.
Step 5	delay { up <i>seconds</i> [down <i>seconds</i>] [up <i>seconds</i>] down <i>seconds</i> }	(Optional) Specify a period of time in seconds to delay communicating state changes of a tracked object. The range is from 1 to 180 seconds.

	Command	Purpose
Step 6	end	Return to privileged EXEC mode.
Step 7	show track <i>object-number</i>	Verify that the specified objects are being tracked.
Step 8	copy running-config startup-config	(Optional) Save your entries in the configuration file.

Use the **no track** *track-number* global configuration command to delete the tracked list.

EXAMPLE

This example configures tracked list 4 with three objects and specified percentages to measure the state of the list:

```
Switch(config)# track 4 list threshold percentage
Switch(config-track)# object 1
Switch(config-track)# object 2
Switch(config-track)# object 3
Switch(config-track)# threshold percentage up 51 down 10
Switch(config-track)# exit
```

Configuring HSRP Object Tracking

Follow this procedure to configure the Hot Standby Router Protocol (HSRP) to track an object and change the Hot Standby priority on the basis of the state of the object.

BEFORE YOU BEGIN

An object must exist before it can be added to a tracked list.

DETAILED STEPS

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	track <i>object-number</i> { interface <i>interface-id</i> { line-protocol ip routing } ip route <i>ip-address/prefix-length</i> { metric threshold reachability } list { boolean { and or } } { threshold { weight percentage } } }	<p>(Optional) Create a tracking list to track the configured state and enter tracking configuration mode.</p> <ul style="list-style-type: none"> • The <i>object-number</i> range is from 1 to 500. • Enter interface <i>interface-id</i> to select an interface to track. • Enter line-protocol to track the interface line protocol state or enter ip routing to track the interface IP routing state. • Enter ip route <i>ip-address/prefix-length</i> to track the state of an IP route. • Enter metric threshold to track the threshold metric or enter reachability to track if the route is reachable. <p>The default up threshold is 254 and the default down threshold is 255.</p> <ul style="list-style-type: none"> • Enter list to track objects grouped in a list. Configure the list as described on the previous pages. <ul style="list-style-type: none"> – For boolean, see the “Configuring a Tracked List with a Boolean Expression” section on page 4-4 – For threshold weight, see the “Configuring a Tracked List with a Weight Threshold” section on page 4-5 – For threshold percentage, see the “Configuring a Tracked List with a Percentage Threshold” section on page 4-6 <p>Note Repeat this step for each interface to be tracked.</p>
Step 3	exit	Return to global configuration mode.
Step 4	interface <i>interface-id</i>	Enter interface configuration mode.
Step 5	standby [<i>group-number</i>] ip [<i>ip-address</i> [secondary]]	<p>Create (or enable) the HSRP group by using its number and virtual IP address.</p> <ul style="list-style-type: none"> • (Optional) <i>group-number</i>—Enter a group number on the interface for which HSRP is being enabled. The range is 0 to 255; the default is 0. If there is only one HSRP group, you do not need to enter a group number. • (Optional on all but one interface) <i>ip-address</i>—Specify the virtual IP address of the hot standby router interface. You must enter the virtual IP address for at least one of the interfaces; it can be learned on the other interfaces. • (Optional) secondary—Specify that the IP address is a secondary hot standby router interface. If this keyword is omitted, the configured address is the primary IP address.

	Command	Purpose
Step 6	standby [<i>group-number</i>] track <i>object-number</i> [decrement [<i>priority-decrement</i>]]	Configure HSRP to track an object and change the hot standby priority based on the state of the object. <ul style="list-style-type: none"> (Optional) <i>group-number</i>—Enter the group number to which the tracking applies. <i>object-number</i>—Enter a number representing the object to be tracked. The range is from 1 to 500; the default is 1. (Optional) decrement <i>priority-decrement</i>—Specify the amount by which the hot standby priority for the router is decremented (or incremented) when the tracked object goes down (or comes back up). The range is from 1 to 255; the default is 10.
Step 7	end	Return to privileged EXEC mode.
Step 8	show standby	Verify the standby router IP address and tracking states.
Step 9	copy running-config startup-config	(Optional) Save your entries in the configuration file.

EXAMPLE

In the following example, the tracking process is configured to track the line-protocol state of GigabitEthernet interface 1/0/0. HSRP on GigabitEthernet interface 0/0/0 then registers with the tracking process to be informed of any changes to the line-protocol state of GigabitEthernet interface 1/0/0. If the line protocol on GigabitEthernet interface 1/0/0 goes down, the priority of the HSRP group is reduced by 10.

Router A:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 line-protocol
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.21 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 110
Device(config-if)# standby 1 track 100 decrement 10
```

Router B:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 line-protocol
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.22 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 105
Device(config-if)# standby 1 track 100 decrement 10
```

In the following example, the tracking process is configured to track the IP-routing capability of GigabitEthernet interface 1/0/0. HSRP on GigabitEthernet interface 0/0/0 then registers with the tracking process to be informed of any changes to the IP-routing state of GigabitEthernet interface 1/0/0. If the IP-routing state on GigabitEthernet interface 1/0/0 goes down, the priority of the HSRP group is reduced by 10.

If both serial interfaces are operational, Router A will be the HSRP active router because it has the higher priority. However, if IP on GigabitEthernet interface 1/0/0 in Router A fails, the HSRP group priority will be reduced and Router B will take over as the active router, thus maintaining a default virtual gateway service to hosts on the 10.1.0.0 subnet.

Router A:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 ip routing
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.21 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 110
Device(config-if)# standby 1 track 100 decrement 10
```

Router B:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 ip routing
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.22 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 105
Device(config-if)# standby 1 track 100 decrement 10
```

Configuring Other Tracking Characteristics

You can also use the enhanced object tracking for tracking other characteristics.

- You can track the reachability of an IP route by using the **track ip route reachability** global configuration command.
- You can use the **track ip route metric threshold** global configuration command to determine if a route is above or below threshold.
- You can use the **track resolution** global configuration command to change the metric resolution default values for routing protocols.
- You can use the **track timer** tracking configuration command to configure the tracking process to periodically poll tracked objects.

Verifying Configuration

Command	Purpose
show ip route track table	Display information about the IP route track table.
show track [<i>object-number</i>]	Display information about the all tracking lists or the specified list.
show track brief	Display a single line of tracking information output.

Command	Purpose
show track interface [brief]	Display information about tracked interface objects.
show track ip [object-number] [brief] route	Display information about tracked IP-route objects.
show track resolution	Display the resolution of tracked parameters.
show track timers	Display tracked polling interval timers.

Configuration Example

This example configures the tracking of an interface line-protocol state and verifies the configuration:

```
Switch(config)# track 33 interface gigabitethernet0/1 line-protocol
Switch(config-track)# end
Switch# show track 33
Track 33
  Interface GigabitEthernet0/1 line-protocol
  Line protocol is Down (hw down)
    1 change, last change 00:18:28
```

This example configures track list 4 with a Boolean AND expression that contains two objects with one object state negated. If the list is up, the list detects that object 2 is down:

```
Switch(config)# track 4 list boolean and
Switch(config-track)# object 1
Switch(config-track)# object 2 not
Switch(config-track)# exit
```

This example configures track list 4 to track by weight threshold. If object 1 and object 2 are down, then track list 4 is up because object 3 satisfies the up threshold value of 30. But if object 3 is down, both objects 1 and 2 must be up in order to satisfy the threshold weight.

```
Switch(config)# track 4 list threshold weight
Switch(config-track)# object 1 weight 15
Switch(config-track)# object 2 weight 20
Switch(config-track)# object 3 weight 30
Switch(config-track)# threshold weight up 30 down 10
Switch(config-track)# exit
```

This configuration can be useful if object 1 and object 2 represent two small bandwidth connections and object 3 represents one large bandwidth connection. The configured **down 10** value means that once the tracked object is up, it will not go down until the threshold value is equal to or lower than 10, which in this example means that all connections are down.

This example configures tracked list 4 with three objects and specified percentages to measure the state of the list:

```
Switch(config)# track 4 list threshold percentage
Switch(config-track)# object 1
Switch(config-track)# object 2
Switch(config-track)# object 3
Switch(config-track)# threshold percentage up 51 down 10
Switch(config-track)# exit
```

In the following example, the tracking process is configured to track the line-protocol state of GigabitEthernet interface 1/0/0. HSRP on GigabitEthernet interface 0/0/0 then registers with the tracking process to be informed of any changes to the line-protocol state of GigabitEthernet interface 1/0/0. If the line protocol on GigabitEthernet interface 1/0/0 goes down, the priority of the HSRP group is reduced by 10.

Router A:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 line-protocol
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.21 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 110
Device(config-if)# standby 1 track 100 decrement 10
```

Router B:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 line-protocol
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.22 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 105
Device(config-if)# standby 1 track 100 decrement 10
```

In the following example, the tracking process is configured to track the IP-routing capability of GigabitEthernet interface 1/0/0. HSRP on GigabitEthernet interface 0/0/0 then registers with the tracking process to be informed of any changes to the IP-routing state of GigabitEthernet interface 1/0/0. If the IP-routing state on GigabitEthernet interface 1/0/0 goes down, the priority of the HSRP group is reduced by 10.

If both serial interfaces are operational, Router A will be the HSRP active router because it has the higher priority. However, if IP on GigabitEthernet interface 1/0/0 in Router A fails, the HSRP group priority will be reduced and Router B will take over as the active router, thus maintaining a default virtual gateway service to hosts on the 10.1.0.0 subnet.

Router A:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 ip routing
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.21 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 110
Device(config-if)# standby 1 track 100 decrement 10
```

Router B:

```
Device(config)# track 100 interface GigabitEthernet1/0/0 ip routing
!
Device(config)# interface GigabitEthernet0/0/0
Device(config-if)# ip address 10.1.0.22 255.255.0.0
Device(config-if)# standby 1 preempt
Device(config-if)# standby 1 ip 10.1.0.1
Device(config-if)# standby 1 priority 105
Device(config-if)# standby 1 track 100 decrement 10
```

Related Documents

- [IP Application Services Configuration Guide, Cisco IOS Release 15M&T](#)
- [Cisco Connected Grid Switches System Management Software Configuration Guide](#)

