CHAPTER 36

Configuring IP Unicast Routing

This chapter describes how to configure IP Version 4 (IPv4) unicast routing on the Catalyst 3560 switch. Basic routing functions, including static routing and the Routing Information Protocol (RIP), are available with both the IP base image (formerly known as the standard multilayer image [SMI]) and the IP services image (formerly known as the enhanced multilayer image [EMI]). To use advanced routing features and other routing protocols, you must have the IP services image installed on the switch.

Note

and configure interfaces to forward IPv6 traffic in addition to IPv4 traffic. For information about configuring IPv6 on the switch, see Chapter 37, "Configuring IPv6 Unicast Routing."

For more detailed IP unicast configuration information, see the *Cisco IOS IP Configuration Guide, Release 12.2* **Documentation > Cisco IOS Software 12.2 Mainline Configuration Guides** For complete syntax and usage information for the commands used in this chapter, see these command references from the Cisco.com page under **Documentation > Cisco IOS Software > 12.2 Mainline > Command References**:

- *Cisco IOS IP Command Reference, Volume 1 of 3: Addressing and Services, Release 12.2*

  Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols, Release 12.2

  Cisco IOS IP Command Reference, Volume 3 of 3: Multicast, Release 12.2

-
-

Note When configuring routing parameters on the switch and to allocate system resources to maximize the number of unicast routes allowed, you can use the **sdm prefer routing** to set the Switch Database Management (sdm) feature to the routing template. For more information on the SDM templates, see Chapter 8, "Configuring SDM Templates" or see the command in the command reference for this release.

Understanding IP Routing

size of the broadcast domain and keeps local traffic local. 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 36-1 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 36-1 Routing Topology Example

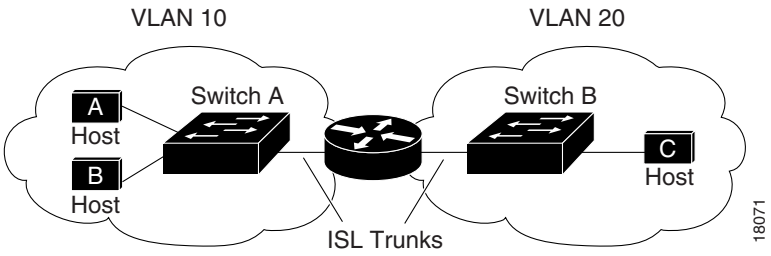

When Host A in VLAN 10 needs to communicate with Host B in VLAN 10, it sends a packet addressed to that host. Switch A forwards the packet directly to Host B, without sending it to the router.

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

-
-
-

the network, which speeds up the convergence time or time required to respond to these changes. Link-state protocols respond quickly to topology changes, but require greater bandwidth and more resources than distance-vector protocols.

Distance-vector protocols supported by the switch are Routing Information Protocol (RIP), which uses a single distance metric (cost) to determine the best path and Border Gateway Protocol (BGP), which adds a path vector mechanism. The switch also supports the Open Shortest Path First (OSPF) link-state protocol and Enhanced IGRP (EIGRP), which adds some link-state routing features to traditional Interior Gateway Routing Protocol (IGRP) to improve efficiency.

The supported protocols are determined by the software running on the switch. If the switch is running the IP base image, only default routing, static routing and RIP are supported. All other routing protocols require the IP services image.

Steps for Configuring Routing

- **no switchport**
- **interface vlan** *vlan_id*

interface port-channel *port-channel-number*

Note

-
-
-
-
-
-

Configuring IP Addressing

interface is required; the other procedures are optional.

Default Addressing Configuration

Table 36-1 Default Addressing Configuration

Feature	Default Setting

Default Addressing Configuration (continued)

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. Spanning Tree Protocol (STP): Disabled. Turbo-flood: Disabled.
IP helper address	Disabled.
IP host	Disabled.
IRDP	Disabled. Defaults when enabled: • 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

Beginning in privileged EXEC mode, follow these steps to assign an IP address and a network mask to a Layer 3 interface:

	Command	Purpose
Step 1		
Step 2		

	Command	Purpose
Step 3		
Step 4	*subnet-mask*	
Step 5		
Step 6		
Step 7		
Step 8	**copy running-config startup-config**	

Use of Subnet Zero

	Command	Purpose
Step 1	**configure terminal**	
Step 2	**ip subnet-zero**	
Step 3	**end**	
Step 4	**show running-config**	
Step 5	**copy running-config startup-config**	

no ip subnet-zero

Classless Routing

Figure 36-2 IP Classless Routing

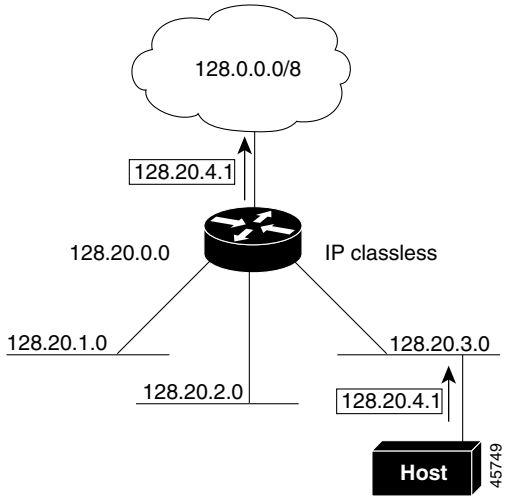

Figure 36-3 No IP Classless Routing

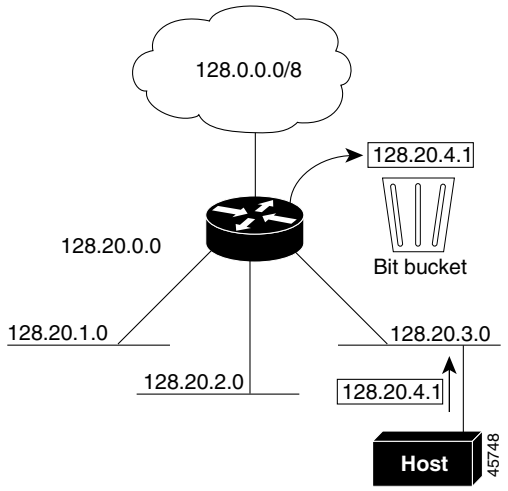

	Verify your entry.
	(Optional) Save your entry in the configuration file.

Configuring Address Resolution Methods

- 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
interface configuration command to identify the server.

For more information on RARP, see the *Cisco IOS Configuration Fundamentals Configuration Guide, Release 12.2*

, page 36-9

Set ARP Encapsulation, page 36-9

Enable Proxy ARP, page 36-10

Define a Static ARP Cache

	Command	Purpose
Step 1		
Step 2		• —ARP encapsulation for Ethernet interfaces • —Subnetwork Address Protocol encapsulation for Token Ring and FDDI interfaces • —HP's ARP type
	[ ]	(Optional) Specify that the switch respond to ARP requests as if it were the owner of the specified IP address.
		Enter interface configuration mode, and specify the interface to configure.
		(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.
		Return to privileged EXEC mode.
	[ ]	Verify the type of ARP and the timeout value used on all interfaces or a specific interface.
	or	View the contents of the ARP cache.
Step 9		

ip-address hardware-address type

Set ARP Encapsulation

	Command	Purpose
Step 1		
Step 2		
Step 3	{ | }	Specify the ARP encapsulation method: —Address Resolution Protocol —Subnetwork Address Protocol
		Return to privileged EXEC mode.
	[ ]	Verify ARP encapsulation configuration on all interfaces or the specified interface.
		(Optional) Save your entries in the configuration file.

Enable Proxy ARP

	Enter global configuration mode.
	Enter interface configuration mode, and specify the Layer 3 interface to configure.

no ip proxy-arp

Routing Assistance When IP Routing is Disabled

-
-
-

Default Gateway

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		

ICMP Router Discovery Protocol (IRDP)

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		**Note**
Step 5		
Step 6		
Step 7		
Step 8		(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.
	[ ]	(Optional) Specify an IRDP address and preference to proxy-advertise.
Step 10		Return to privileged EXEC mode.
Step 11		Verify settings by displaying IRDP values.
Step 12		(Optional) Save your entries in the configuration file.

Configuring Broadcast Packet Handling

-
-

Note

-
-
-
-

Enabling Directed Broadcast-to-Physical Broadcast Translation

	Command	Purpose
Step 1		
Step 2		

	Command	Purpose
Step 3		Note
Step 4		
Step 5	| nd | sdns	• udp • nd • sdns
Step 6	end	
Step 7	show ip interface show running-config	
Step 8	copy running-config startup-config	

Forwarding UDP Broadcast Packets and Protocols

	Command	Purpose
Step 1	configure terminal	
Step 2	interface	
Step 3	ip helper-address	
Step 4	exit	
Step 5	ip forward-protocol udp | nd | sdns	
Step 6	end	
Step 7	show ip interface show running-config	
Step 8	copy running-config startup-config	

no ip helper-address

no ip forward-protocol

Establishing an IP Broadcast Address

	Command	Purpose
Step 1	configure terminal	
Step 2	interface	
Step 3	ip broadcast-address	
Step 4	end	
Step 5	show ip interface	
Step 6	copy running-config startup-config	

no ip broadcast-address

Flooding IP Broadcasts

-
-
-
  ip forward-protocol udp
-

ip broadcast-address

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.

Beginning in privileged EXEC mode, follow these steps to use the bridging spanning-tree database to flood UDP datagrams:

	Enter global configuration mode.
	Use the bridging spanning-tree database to flood UDP datagrams.
	Return to privileged EXEC mode.
	Verify your entry.
	(Optional) Save your entry in the configuration file.

Use the global configuration command to disable the flooding of IP broadcasts.

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.

Beginning in privileged EXEC mode, follow these steps to increase spanning-tree-based flooding:

	Enter global configuration mode
	Use the spanning-tree database to speed up flooding of UDP datagrams.
	Return to privileged EXEC mode.
	Verify your entry.
	(Optional) Save your entry in the configuration file.

To disable this feature, use the global configuration command.

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 privileged EXEC commands. Table 36-2 lists the commands for clearing contents.

Commands to Clear Caches, Tables, and Databases

| *}	Remove one or all entries from the hostname and the address cache.
{ [ ] |*}	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. Table 36-3 lists the privileged EXEC commands for displaying IP statistics.

	Display the entries in the ARP table.
	Display the default domain name, style of lookup service, name server hosts, and the cached list of hostnames and addresses.
	Display IP addresses mapped to TCP ports (aliases).
	Display the IP ARP cache.
[ ]	Display the IP status of interfaces.
	Display IRDP values.
show ip masks	
show ip redirects	
show ip route	
show ip route summary	

Enabling IP Unicast Routing


```
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
                       end
```

Table 36-4 Default RIP Configuration

Feature	Default Setting
	• • • •

Configuring Basic RIP Parameters

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		 Note
Step 5		
Step 6		
Step 7		 • • • •

{ }	(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 { | } | | } to control what versions are used for sending and receiving on interfaces.
	(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.
	(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.
	(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.
	Return to privileged EXEC mode.
	Verify your entries.
	(Optional) Save your entries in the configuration file.

Configuring RIP Authentication

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 36-99.

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

Beginning in privileged EXEC mode, follow these steps to configure RIP authentication on an interface:

	Enter global configuration mode.
	Enter interface configuration mode, and specify the interface to configure.

Step 3		
Step 4	**md5**	
	end	
	show running-config interface	
	copy running-config startup-config	

To restore clear text authentication, use the command. To prevent authentication, use the configuration command.

interface configuration
interface

Configuring Summary Addresses and Split Horizon

Note

Note

configure terminal	
interface	
ip address	
ip summary-address rip	
no ip split horizon	
end	
show ip interface	
copy running-config startup-config	

```
router rip
       interface gi0/2
   ip address 10.1.5.1 255.255.255.0
   ip summary-address rip 10.2.0.0 255.255.0.0
   no ip split-horizon
   exit
router rip
       network 10.0.0.0
       neighbor 2.2.2.2 peer-group mygroup
       end
```


Configuring OSPF

Note

-
-

-
-
-
-

-
-
-
-
-
-
-
-

Note

Default OSPF Configuration

Table 36-5 Default OSPF Configuration

	[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.

1. NSF = Nonstop forwarding
2. OSPF NSF awareness is enabled for IPv4 on Catalyst 3550, 3560 and 3750 switches running the IP services image.

OSPF NSF Awareness

OSPF Nonstop Forwarding (NSF) Awareness Feature Guide

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_white_paper09186a0080153edd.shtml

Configuring Basic OSPF Parameters

	Enter global configuration mode.
	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.
	Define an interface on which OSPF runs and the area ID for that interface. You can use the wildcard mask as a single command to define one or more interfaces to be associated with a specific OSPF area. The area ID can be a decimal value or an IP address.
	Return to privileged EXEC mode.

To end an OSPF routing process, use the global configuration command.

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

```
router ospf 109
       network 131.108.0.0 255.255.255.0 area 24
```

	Options is 0x52 *LLS Options is 0x1 (LR)* *Options is 0x42*

process-id	
area-id	
area-id	
area-id	
area-id	
area-id *address mask*	
process-id *process-id* *area-id*	

ref-bw *ref*
bw

process-id	
address mask	
area-id *seconds* *seconds* *key* *keyid* *key*	
metric-value *type-value* *map-name*	

Step 7		
Step 8		
Step 9		
Step 10		• • •
Step 11	**ospf log-adj-changes**	
	end	
	show ip ospf database	
	copy running-config startup-config	

Changing LSA Group Pacing

configure terminal	
router ospf	
timers lsa-group-pacing	
end	
show running-config	
copy running-config startup-config	

configure terminal	
interface loopback 0	
ip address	
end	
show ip interface	
copy running-config startup-config	

no interface loopback 0

show ip ospf

database

Show IP OSPF Statistics Commands

show ip ospf	
show ip ospf database router **show ip ospf database router self-originate** **show ip ospf database router adv-router** **show ip ospf database network** **show ip ospf database summary** **show ip ospf database asbr-summary** **show ip ospf database external** **show ip ospf database database-summary**	
show ip ospf border-routes	
show ip ospf interface	

interface-name neighbor-id	

Configuring EIGRP

-
-
-
-
-
-
-
-

- *The reliable transport protocol*

The DUAL finite state machine

protocol-dependent modules

Default EIGRP Configuration

Table 36-7 Default EIGRP Configuration

EIGRP NSF awareness is enabled for IPv4 on Catalyst 3550, 3560 and 3750 switches running the IP services image.

If you have routers on your network that are configured for IGRP, and you want to change to EIGRP, you must designate transition routers that have both IGRP and EIGRP configured. In these cases, perform Steps 1 through 3 in the next section and also see the "Configuring Split Horizon" section on page 36-23. You must use the same AS number for routes to be automatically redistributed.

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 *EIGRP Nonstop Forwarding (NSF) Awareness Feature Guide* at this URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a0080160010.html

Beginning in privileged EXEC mode, follow these steps to configure EIGRP. Configuring the routing process is required; other steps are optional:

		Enter global configuration mode.
	autonomous-system number	Enable an EIGRP routing process, and enter router configuration mode. The AS number identifies the routes to other EIGRP routers and tags routing information.
	network-number	Associate networks with an EIGRP routing process. EIGRP sends updates to the interfaces in the specified networks.
		(Optional) Enable logging of EIGRP neighbor changes to monitor routing system stability.
	tos k1 k2 k3 k4 k5	(Optional) Adjust the EIGRP metric. Although the defaults have been carefully set to provide excellent operation in most networks, you can adjust them. **Caution**
Step 6	**offset list [in out**	
Step 7	**no auto-summary**	
Step 8	**ip summary-address eigrp**	

Step 9	**end**	
Step 10	**show ip protocols**	
Step 11	**show ip protocols**	*** IP Routing is NSF aware *** EIGRP NSF enabled

no

Configuring EIGRP Interfaces

	and the earliest acceptable date as January 1, 1993. The default and is .
{ }	(Optional) Specify the time period during which the key can be sent. The and syntax can be either : : or : : . The default is forever with the default and the earliest acceptable date as January 1, 1993. The default and is .
	Return to privileged EXEC mode.

	Display authentication key information.
	(Optional) Save your entries in the configuration file.

Use the forms of these commands to disable the feature or to return the setting to the default value.

The EIGRP stub routing feature, available in all images, reduces resource utilization by moving routed traffic closer to the end user.

The IP base image contains only EIGRP stub routing capability, which 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. For enhanced capability and complete EIGRP routing, the switch must be running the IP services image.
On a switch running the IP base image, if you try to configure multi-VRF-CE and EIGRP stub routing at the same time, the configuration is not allowed.

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.

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 36-4, 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).

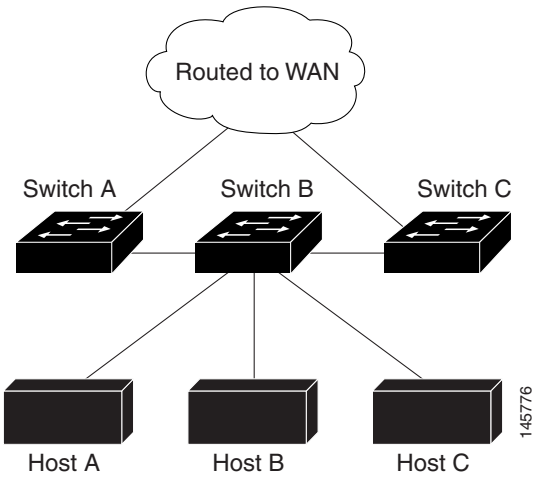

Monitoring and Maintaining EIGRP

Table 36-8 IP EIGRP Clear and Show Commands

Command	Purpose

Configuring BGP

BGP
external BGP

EBGP, IBGP, and Multiple Autonomous Systems

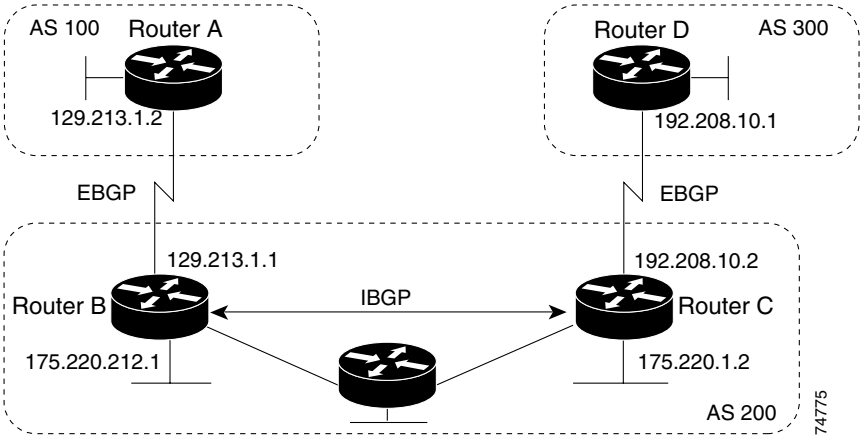

Table 36-9 Default BGP Configuration

Table 36-9 Default BGP Configuration (continued)

NSF Awareness can be enabled for IPv4 on Catalyst 3550, 3560, and 3750 switches with the IP services image by enabling Graceful Restart.

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 e at this URL:

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a008015fede.html

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. are in the same AS; 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 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 will be passing traffic through it from another AS to a third AS, it is important to be consistent about the routes it advertises. If BGP advertised a route before all routers in the network had learned about the route through the IGP, the AS might receive traffic that some routers could 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*

autonomous-system	
network-number *network-mask* *route-map-name*	

ip-address peer-group-name number	
ip-address peer-group-name	

network-number	`Graceful Restart Capability: advertised and received`
	
	`Graceful Restart Capability: advertised`

autonomous-system
network-number
ip-address peer-group-name number
ip-address peer-group-name

```
        neighbor 175.220.1.2 remote-as 200


router bgp 200
        neighbor 175.220.212.1 remote-as 200
        neighbor 192.208.10.1 remote-as 300


router bgp 300
        neighbor 192.208.10.2 remote-as 200
```

```
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
```

state = established

Table 36-10 Advantages and Disadvantages of Hard and Soft Resets

Type of Reset	Advantages	Disadvantages

	Command	Purpose
Step 1		
Step 2	*	• • •
Step 3	**clear ip bgp** * **soft out**	• • •
Step 4	**show ip bgp** **show ip bgp neighbors**	

Configuring BGP Decision Attributes

maximum-paths

1.
2.
3.
4.
5.
6.
7.
8.
9.
10.
    -
    -
    -
11.

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		
Step 9		
Step 10		
Step 11		
Step 12		
Step 13		
Step 14		
Step 15		

Configuring BGP Filtering with Route Maps

	Command	Purpose
Step 1		
Step 2	*sequence-number*	
	ip-address ...ip-address	
	map-name	

Configuring BGP Filtering by Neighbor

autonomous-system	
ip-address peer-group name *access-list-number name*	
ip-address peer-group name *map-tag*	

map-tag

Cisco IOS Dial Technologies Command Reference, Release 12.2

access-list-number *as-regular-expressions*	
autonomous-system	
ip-address peer-group name *access-list-number name* *weight*	
regular-expression	

Configuring Prefix Lists for BGP Filtering

-
-
-

configure terminal	
ip prefix-list seq deny permit ge le	**deny permit permit deny** **ge le** < < < 32
| / [ ] [ ]	(Optional) Add an entry to a prefix list, and assign a sequence number to the entry.
	Return to privileged EXEC mode.
[|] [/] [ ] [ ] [ ]	Verify the configuration by displaying information about a prefix list or prefix list entries.
	(Optional) Save your entries in the configuration file.

ip prefix-list seq

no ip prefix-list

command; to reenable automatic generation, use the command. To clear the hit-count table of prefix list entries, use the privileged EXEC command.

One way that BGP controls the distribution of routing information based on the value of the COMMUNITIES attribute. The attribute is a way to groups destinations into communities and to apply routing decisions based on the communities. This method simplifies configuration of a BGP speaker to control distribution of routing information.

A 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:

—Advertise this route to the Internet community. All routers belong to it.

—Do not advertise this route to EBGP peers.

—Do not advertise this route to any peer (internal or external).

—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 and route-map configuration commands in the "Using Route Maps to Redistribute Routing Information" section on page 36-90.

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 router configuration command.

Beginning in privileged EXEC mode, follow these steps to create and to apply a community list:

	Enter global configuration mode.
{ | }	Create a community list, and assign it a number. The is an integer from 1 to 99 that identifies one or more permit or deny groups of communities. The is the number configured by a route-map configuration command.
	Enter BGP router configuration mode.
{ | }	Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.

	(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.
	Return to global configuration mode.
	(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.
	Return to privileged EXEC mode.
	Verify the configuration.
	(Optional) Save your entries in the configuration file.

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 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 router configuration command.

Beginning in privileged EXEC mode, use these commands to configure BGP peers:

	Enter global configuration mode.
	Enter BGP router configuration mode.
	Create a BGP peer group.
	Make a BGP neighbor a member of the peer group.
{ | }	Specify a BGP neighbor. If a peer group is not configured with a , use this command to create peer groups containing EBGP neighbors. The range is 1 to 65535.
{ | }	(Optional) Associate a description with a neighbor.

{ | } [ ]	(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.
{ | }	(Optional) Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
{ | }	(Optional) Allow internal BGP sessions to use any operational interface for TCP connections.
{ | }	(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).
{ | }	(Optional) Specify an AS number to use as the local AS. The range is 1 to 65535.
{ | }	(Optional) Set the minimum interval between sending BGP routing updates.
{ | } [ ]	(Optional) Control how many prefixes can be received from a neighbor. The range is 1 to 4294967295. The (optional) is the percentage of maximum at which a warning message is generated. The default is 75 percent.
{ | }	(Optional) Disable next-hop processing on the BGP updates to a neighbor.
{ | }	(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.
{ | } { | }	(Optional) Apply a route map to incoming or outgoing routes.
{ | }	(Optional) Specify that the COMMUNITIES attribute be sent to the neighbor at this IP address.
{ | }	(Optional) Set timers for the neighbor or peer group. The interval is the time within which keepalive messages are sent to peers. The range is 1 to 4294967295 seconds; the default is 60. The 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.
{ | }	(Optional) Specify a weight for all routes from a neighbor.
{ | } { { | } { | }	(Optional) Filter BGP routing updates to or from neighbors, as specified in an access list.
{ | } { | | }	(Optional) Establish a BGP filter.
{ | }	(Optional) Specify the BGP version to use when communicating with a neighbor.

{ | }	(Optional) Configure the software to start storing received updates.
	Return to privileged EXEC mode.
	Verify the configuration.
	(Optional) Save your entries in the configuration file.

To disable an existing BGP neighbor or neighbor peer group, use the router configuration command. To enable a previously existing neighbor or neighbor peer group that had been disabled, use the router configuration command.

Classless interdomain routing (CIDR) enables you to create aggregate routes (or) 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.

Beginning in privileged EXEC mode, use these commands to create an aggregate address in the routing table:

	Enter global configuration mode.
	Enter BGP router configuration mode.
	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.
	(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.
	(Optional) Advertise summary addresses only.
	(Optional) Suppress selected, more specific routes.
	(Optional) Generate an aggregate based on conditions specified by the route map.
	(Optional) Generate an aggregate with attributes specified in the route map.
	Return to privileged EXEC mode.
[ ]	Verify the configuration.
	(Optional) Save your entries in the configuration file.

To delete an aggregate entry, use the router configuration command. To return options to the default values, use the command with keywords.

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.

Beginning in privileged EXEC mode, use these commands to configure a BGP confederation:

	Enter global configuration mode.
	Enter BGP router configuration mode.
	Configure a BGP confederation identifier.
[...]	Specify the autonomous systems that belong to the confederation and that will be treated as special EBGP peers.
	Return to privileged EXEC mode.
	Verify the configuration.
	(Optional) Save your entries in the configuration file.

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 IBPG 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 , 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: and 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 . 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.

Beginning in privileged EXEC mode, use these commands to configure a route reflector and clients:

	Enter global configuration mode.
	Enter BGP router configuration mode.
	Configure the local router as a BGP route reflector and the specified neighbor as a client.
	(Optional) Configure the cluster ID if the cluster has more than one route reflector.
	(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.
	Return to privileged EXEC mode.
	Verify the configuration. Display the originator ID and the cluster-list attributes.
	(Optional) Save your entries in the configuration file.

Route flap dampening is a BGP feature designed to minimize 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 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 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.

Beginning in privileged EXEC mode, use these commands to configure BGP route dampening:

	Enter global configuration mode.
	Enter BGP router configuration mode.
	Enable BGP route dampening.
[ ]	(Optional) Change the default values of route dampening factors.
	Return to privileged EXEC mode.
[{ } | { } | { [ ]}]	(Optional) Monitor the flaps of all paths that are flapping. The statistics are deleted when the route is not suppressed and is stable.
	(Optional) Display the dampened routes, including the time remaining before they are suppressed.
[{ } | { } | { [ ]}	(Optional) Clear BGP flap statistics to make it less likely that a route will be dampened.
	(Optional) Clear route dampening information, and unsuppress the suppressed routes.
	(Optional) Save your entries in the configuration file.

To disable flap dampening, use the router configuration command without keywords. To set dampening factors back to the default values, use the router configuration command with values.

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.

Table 36-8 lists the privileged EXEC commands for clearing and displaying BGP. For explanations of the display fields, see the from the Cisco.com page under > > > .

	Reset a particular BGP connection.
	Reset all BGP connections.
	Remove all members of a BGP peer group.
	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.

Command	Purpose

Configuring ISO CLNS Routing

Note

DECnet, ISO CLNS and XNS Configuration Guide, Release 12.2
Cisco IOS Apollo Domain, Banyan VINES,
DECnet, ISO CLNS and XNS Command Reference, Release 12.2,

Configuring IS-IS Dynamic Routing

Note

-
-

-
-

Default IS-IS Configuration

Feature	Default Setting
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 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.

IS-IS NSF awareness is enabled for IPv4 on switches running Cisco IOS Release 12.2(25)SEG or later.

The integrated IS-IS NSF Awareness feature is supported for IPv4, beginning with Cisco IOS Release 12.2(25)SEG. 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 at this URL: http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_white_paper09186a00801541c7.shtml

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.

Beginning in privileged EXEC mode, follow these steps to enable IS-IS and specify the area for each instance of the IS-IS routing process:

	Enter global configuration mode.
	Enable ISO connectionless routing on the switch.
[ ]	Enable the IS-IS routing for the specified routing process and enter IS-IS routing configuration mode. (Optional) Use the 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 global configuration command.
	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.
{ | | }	(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): —act as a station router only —act as both a station router and an area router —act as an area router only
	Return to global configuration mode.
	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 command to put it into Layer 3 mode.
[ ]	Configure an IS-IS routing process for ISO CLNS on the interface and attach an area designator to the routing process.

clns router isis	
ip address	
end	
show isis database detail	
copy running-config startup-config	

To disable IS-IS routing, use the router configuration command.

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

Router B

Router C

Configuring IS-IS Global Parameters

-
-
-
-
-
-
-
-
-
-

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		

Step 9		• • •
Step 10		
Step 11		
Step 12		• • • *lsp-second-wait*—
	spf-max-wait spf-initial-wait spf-second-wait	*spf-max-wait* *spf-initial-wait* *spf-second-wait*

[ ]	(Optional) Sets IS-IS partial route computation (PRC) throttling timers. the maximum interval (in seconds) between two consecutive PRC calculations. The range is 1 to 120; the default is 5. the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 10,000; the default is 2000. the hold time between the first and second PRC calculation (in milliseconds). The range is 1 to 10,000; the default is 5000.
[ ]	(Optional) Set the router to log IS-IS adjacency state changes. Enter 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).
	(Optional) Specify the maximum LSP packet size in bytes. The range is 128 to 4352; the default is 1497 bytes. If any link in the network has a reduced MTU size, you must change the LSP MTU size on all routers in the network.
	(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.
	Return to privileged EXEC mode.
	Verify your entries.
	(Optional) Save your entries in the configuration file.

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:

–

–

–

	minimal— **—**
isis hello-multiplier **level-1 level-2**	
isis csnp-interval level-1 **level-2**	

isis retransmit-interval	
isis retransmit-throttle-interval	**isis lsp-interval**
isis priority level-1 level-2	
isis circuit-type level-1 level-1-2 level-2-only	**level-1** **level-1-2** **level 2**
isis password level-1 level-2	
end	
show clns interface	
copy running-config startup-config	

no

ISO CLNS and IS-IS Clear and Show Commands

Command	Purpose

Configuring Multi-VRF CE

Note

Understanding Multi-VRF CE

Note

-

Switches Acting as Multiple Virtual CEs

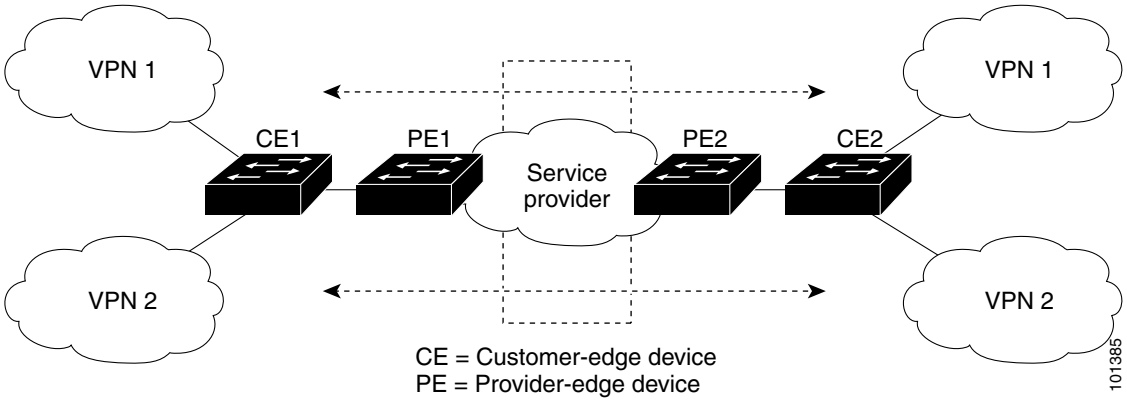

-
-

-
-
-
-

-
-
-

Default Multi-VRF CE Configuration

Table 36-14 Default VRF Configuration

Feature	Default Setting

Multi-VRF CE Configuration Guidelines

Note

-
-
-
-
-
-

-
-
-
-
-
-
-
-
-
-
-
-
-

Configuring VRFs

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		

	Command	Purpose
Step 5		
Step 6		
Step 7		
Step 8		
Step 9		
Step 10		
Step 11		

Configuring Multicast VRFs

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		

	Command	Purpose
Step 9		
Step 10		
Step 11		
Step 12		
Step 13		
Step 14		

Configuring VRF-Aware Services

-
-

-
-
-
-
-
-
-
-

User Interface for ARP

Command	Purpose

User Interface for PING

Command	Purpose

User Interface for SNMP

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5	snmp-server host < vrf informs	
Step 6	snmp-server user remote vrf	
Step 7	end	

User Interface for HSRP

	Command	Purpose
Step 1	configure terminal	
Step 2	interface	
Step 3	no switchport	
Step 4	ip vrf forwarding	
Step 5	ip address	

	Command	Purpose
Step 6	standby 1 ip	
Step 7	end	

User Interface for uRPF

	Command	Purpose
Step 1	configure terminal	
Step 2	interface	
Step 3	no switchport	
Step 4	ip vrf forwarding	
Step 5	ip address	
Step 6	ip verify unicast reverse-path	
Step 7	end	

User Interface for Syslog

	Command	Purpose
Step 1	configure terminal	
Step 2	logging on	
Step 3	logging host *ip address* *vrf name*	

traceroute vrf	

ip [t]ftp

source-interface E1/0

ip ftp source-interface
no

configure terminal	
ip ftp source-interface	
end	

ip tftp

source-interface **no**

configure terminal	
ip tftp source-interface	
end	

autonomous-system

configure terminal	
router ospf vrf	
log-adjacency-changes	
redistribute bgp subnets	
network area	
end	
show ip ospf	
copy running-config startup-config	

no router ospf vrf

configure terminal	
router bgp	
network mask	
redistribute ospf match internal	
network area	
address-family ipv4 vrf	
neighbor remote-as	
neighbor activate	
end	

	Command	Purpose
Step 10		
Step 11		

no router bgp

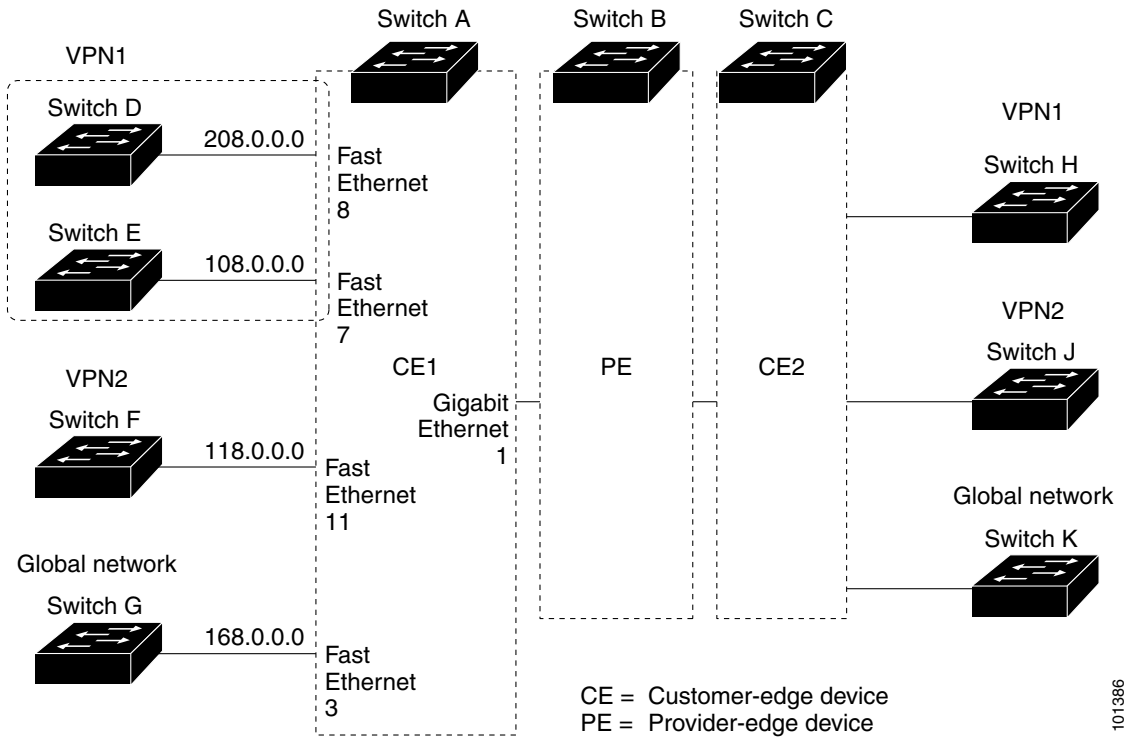

Configuring Switch A

```
ip vrf v11
    rd 800:1
    route-target export 800:1
    route-target import 800:1
    exit
ip vrf v12
    rd 800:2
    route-target export 800:2
    route-target import 800:2
    exit



interface loopback1
   ip vrf forwarding v11
   ip address 8.8.1.8 255.255.255.0
   exit

interface loopback2
   ip vrf forwarding v12
   ip address 8.8.2.8 255.255.255.0
   exit

interface gigabitethernet0/5
   switchport trunk encapsulation dot1q
   switchport mode trunk
   no ip address
   exit

interface fastethernet0/8
   switchport access vlan 208
   no ip address
   exit

interface fastethernet0/11
   switchport trunk encapsulation dot1q
   switchport mode trunk
   no ip address
   exit




interface vlan10
   ip vrf forwarding v11
   ip address 38.0.0.8 255.255.255.0
   exit

interface vlan20
   ip vrf forwarding v12
   ip address 83.0.0.8 255.255.255.0
   exit

interface vlan118
   ip vrf forwarding v12
   ip address 118.0.0.8 255.255.255.0
   exit

interface vlan208
   ip vrf forwarding v11
   ip address 208.0.0.8 255.255.255.0
```

```
   exit


router ospf 1 vrf vl1
       redistribute bgp 800 subnets
       network 208.0.0.0 0.0.0.255 area 0
       exit
router ospf 2 vrf vl2
       redistribute bgp 800 subnets
       network 118.0.0.0 0.0.0.255 area 0
       exit


router bgp 800
       address-family ipv4 vrf vl2
          redistribute ospf 2 match internal
          neighbor 83.0.0.3 remote-as 100
          neighbor 83.0.0.3 activate
          network 8.8.2.0 mask 255.255.255.0
          exit

       address-family ipv4 vrf vl1
          redistribute ospf 1 match internal
          neighbor 38.0.0.3 remote-as 100
          neighbor 38.0.0.3 activate
          network 8.8.1.0 mask 255.255.255.0
          end
```

Configuring Switch D

```
configure terminal

       ip routing
       interface fastethernet0/2
          no switchport
          ip address 208.0.0.20 255.255.255.0
          exit

       router ospf 101
              network 208.0.0.0 0.0.0.255 area 0
              end
```

Configuring Switch F

Configuring the PE Switch B

```
interface Loopback1
   ip vrf forwarding v1
   ip address 3.3.1.3 255.255.255.0
   exit

interface Loopback2
   ip vrf forwarding v2
   ip address 3.3.2.3 255.255.255.0
   exit

interface gigabitthernet1/0.10
   encapsulation dot1q 10
   ip vrf forwarding v1
   ip address 38.0.0.3 255.255.255.0
   exit

interface gigabitethernet1/0.20
   encapsulation dot1q 20
   ip vrf forwarding v2
   ip address 83.0.0.3 255.255.255.0
   exit

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

Displaying Multi-VRF CE Status

Command	Purpose

Configuring Protocol-Independent Features

-
-
-
-
-
-
-
-

Configuring Cisco Express Forwarding

-
-

Caution

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		
Step 9		
Step 10		

Configuring the Number of Equal-Cost Routing Paths

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		

Configuring Static Unicast Routes

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		

Route Source	Default Distance

router's address in a static route, the static route is also removed from the IP routing table.

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.

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		

Using Route Maps to Redistribute Routing Information

Note

Note

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		

	Command	Purpose
Step 9		
Step 10		
Step 11		• • •
Step 12		
Step 13		
Step 14		
Step 15		
Step 16		
Step 17		
Step 18		• • • • •
Step 19		
Step 20		

	Command	Purpose
Step 21		
Step 22		
Step 23		
Step 24		

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		

-
-

Configuring Policy-Based Routing

-
-
-

-
-

-

Note

PBR Configuration Guidelines

-
-
-
-

-
-
-
-
-
-
-
- *Don't Fragment*

map-tag *sequence number*	*map-tag* *sequence number*
access-list-number *access-list-name* *..access-list-number* *...access-list-name*	
ip-address *..ip-address*	
interface-id	
map-tag	

map-tag	
map-name	

map-tag

map-tag

map-tag

interface-id	
interface type	
network-address	*network-address*

interface-id

access-list-number *access-list-name* *interface-name* *routing process* *autonomous-system-number*	
access-list-number *access-list-name* *type-number*	

administrative distance

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		

Managing Authentication Keys

	Command	Purpose
Step 1		
Step 2		
Step 3		
Step 4		
Step 5		
Step 6		
Step 7		
Step 8		
Step 9		

Monitoring and Maintaining the IP Network

Command	Purpose

Command	Purpose