



Configuring IPv6 Unicast Routing

Internet Protocol Version 6 (IPv6) is the network-layer Internet Protocol intended to replace Version 4 (IPv4) in the TCP/IP suite of protocols. This chapter describes how to configure IPv6 unicast routing on the Catalyst 3560 switch. For information about configuring IPv4 unicast routing, see [Chapter 34, “Configuring IP Unicast Routing.”](#) For information about configuring IPv6 Multicast Listener Discovery (MLD) snooping, see [Chapter 36, “Configuring IPv6 MLD Snooping.”](#) For information on configuring IPv6 access control lists (ACLs) see [Chapter 37, “Configuring IPv6 ACLs.”](#)

To use this feature, the switch must be running the advanced IP services image, which is orderable from Cisco. This image includes all IP services image (formerly known as the enhanced multilayer image [EMI]) features plus IPv6 host and unicast routing support. To enable IPv6 routing, you must also configure a switch database management (SDM) dual IPv4 and IPv6 template. See the [“SDM Templates” section on page 35-7.](#)



Note

For complete syntax and usage information for the commands used in this chapter, see the Cisco IOS documentation referenced in the procedures

This chapter consists of these sections:

- [“Understanding IPv6” section on page 35-1](#)
- [“Configuring IPv6” section on page 35-9](#)
- [“Displaying IPv6” section on page 35-21](#)

Understanding IPv6

The primary reason for using IPv6 is to increase Internet global address space to accommodate the rapidly increasing number of users and applications that require unique global IP addresses. IPv4 uses 32-bit addresses to provide approximately 4 billion available addresses. Large blocks of these addresses are allocated to government agencies and large organizations, and the number of available IP addresses is rapidly decreasing. IPv6 incorporates 128-bit source and destination addresses and can provide significantly more globally unique IP addresses than IPv4.

The architecture of IPv6 allows existing IPv4 users to transition easily to IPv6, and provides services such as end-to-end security, quality of service (QoS), and globally unique addresses. The flexibility of the IPv6 address space reduces the need for private addresses and the use of Network Address Translation (NAT) processing by border routers at the edge of networks. IPv6 provides newer unicast methods, introduces hexadecimal values into the IP address, and uses colons (:) instead of periods (.) as delimiters.

IPv6 also provides these advantages over IPv4:

- Easier address management and delegation
- Easy address autoconfiguration with *stateless autoconfiguration*, which is similar to DHCP but does not require a specified DHCP application or server
- Embedded IPsec (encrypted security)
- Routing optimized for mobile devices
- Duplicate Address Detection (DAD) feature

For information about how Cisco Systems implements IPv6, go to this URL:

<http://www.cisco.com/warp/public/732/Tech/ipv6/>

This section describes IPv6 implementation on the switch. These sections are included:

- [IPv6 Addresses, page 35-2](#)
- [Supported IPv6 Unicast Routing Features, page 35-3](#)
- [Unsupported IPv6 Unicast Routing Features, page 35-6](#)
- [Limitations, page 35-7](#)
- [SDM Templates, page 35-7](#)

IPv6 Addresses

IPv6 supports three types of addresses: unicast (one-to-one), multicast (one-to-many), and anycast (one-to-nearest). Multicast addresses replace the use of broadcast addresses. The switch supports only IPv6 unicast addresses.

The IPv6 128-bit addresses are represented as a series of eight 16-bit hexadecimal fields separated by colons in the format: x:x:x:x:x:x:x. 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, go to “Implementing Basic Connectivity for IPv6” at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ipv6_vgf.htm

In the “Information About Implementing Basic Connectivity for IPv6” section, these sections apply to the Catalyst 3560 switch:

- IPv6 Address Formats
- IPv6 Address Type: Unicast
- IPv6 Address Output Display
- Simplified IPv6 Packet Header

The switch does not support site-local unicast addresses, anycast addresses, or multicast addresses in this release.

Supported IPv6 Unicast Routing Features

These sections describe the IPv6 protocol (RFC 2460) features supported by the switch:

- [128-Bit Wide Unicast Addresses, page 35-3](#)
- [Path MTU Discovery for IPv6 Unicast, page 35-4](#)
- [ICMPv6, page 35-4](#)
- [IPv6 Stateless Autoconfiguration and Duplicate Address Detection, page 35-5](#)
- [IPv6 Applications, page 35-5](#)
- [Dual IPv4 and IPv6 Protocol Stacks, page 35-5](#)

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 native Ethernet Inter-Switch Link (ISL) or 802.1Q trunk ports for static routes, Routing Information Protocol (RIP) for IPv6 (RFC 2080), and Open Shortest Path First (OSPF) Version 3 protocol (RFC 2740). It supports up to 16 equal-cost routes and can forward IPv4 and IPv6 frames simultaneously at line rate.

128-Bit Wide Unicast Addresses

The switch supports aggregatable global unicast addresses and link-local unicast addresses (RFC 2373). 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 universal 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 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.

See the section on IPv6 Unicast Addresses in the “Implementing Basis Connectivity for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ipv6_vgf.htm

Each IPv6 host interface can support up to three addresses in hardware (one aggregatable global unicast address, one link-local unicast address, and zero or more privacy addresses).

DNS for IPv6

IPv6 introduces new Domain Name System (DNS) record types that are supported in the DNS name-to-address and address-to-name lookup processes. The new 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 MTU to IPv6 nodes and path MTU discovery. Path MTU discovery (RFC 1981) 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 (RFC 2463) functions the same as in IPv4. ICMP 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. A value of 58 in the Next Header field of the basic IPv6 packet header identifies an IPv6 ICMP packet.

The switch also supports Neighbor Discovery Protocol (NDP) for IPv6 (RFC 2461), a protocol running on top of ICMPv6, and Static Neighbor Discovery 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), verify the reachability of the neighbor, and keep track of neighboring routers.

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. These messages are sent on the local link when a node needs to determine the link-layer address of another node on the same local link. When a destination node receives a neighbor solicitation message, it replies by sending a neighbor advertisement message, which has a value of 136 in the ICMP packet header Type field.

A value of 137 in the ICMP packet header Type field identifies an IPv6 neighbor redirect message. The switch supports ICMPv6 redirect (RFC 2463) for routes with mask lengths less than 64. ICMP redirect is not supported for host routes or for summarized routes with mask lengths greater than 64. Routers send neighbor-redirect messages to inform hosts of better first-hop nodes on the path to a destination. A router does not update its routing tables after receiving a neighbor-redirect message and hosts do not originate neighbor-redirect messages.

IPv6 Stateless Autoconfiguration and Duplicate Address Detection

IPv6 supports two types of autoconfiguration:

- Stateless autoconfiguration (RFC 2462), where a host autonomously configures its own link-local address, and booting nodes send router solicitations to request router advertisements for configuring interfaces
- Stateful autoconfiguration using Dynamic Host Configuration Protocol (DHCP) v6.

The switch supports stateless autoconfiguration to manage link, subnet, and site addressing changes, such as management of host and mobile IP addresses.

All interfaces on IPv6 nodes must have a link-local address, which is automatically configured from the identifier (router MAC address) for an interface and the link-local prefix FE80::/10. A link-local address enables a node to communicate with other nodes on the link and can be used to further configure the node. Nodes can connect to a network and automatically generate global IPv6 addresses without the need for manual configuration or the help of a server, such as a DHCP server. With IPv6, a router on the link uses router advertisement messages to advertise global prefixes and its ability to act as a default router for the link. A node on the link can automatically configure global IPv6 addresses by appending its interface identifier (64-bits) to the prefixes (64 bits) included in the router advertisement messages.

The 128-bit IPv6 addresses configured by a node are then subjected to duplicate-address detection (RFC 2462) to ensure their uniqueness on the link. If the advertised prefixes are globally unique, the IPv6 addresses configured by the node are guaranteed to be globally unique. Router solicitation messages, which have a value of 133 in the ICMP packet header Type field, are sent by hosts at system startup so that the host can be immediately autoconfigured without waiting for the next scheduled router advertisement message. IPv6 duplicate-address detection is performed on unicast addresses before they are assigned to an interface. The switch does not support automatically generated site-local IPv6 addresses.

IPv6 Applications

The switch has IPv6 support for these applications:

- Ping, traceroute, Telnet, TFTP, and FTP
- HTTP server access over IPv6 transport
- DNS resolver for AAAA over IPv4 transport

For more information about managing these applications with Cisco IOS, see the “Managing Cisco IOS Applications over IPv6” section in the *Cisco IOS IPv6 Configuration Library* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/sa_mgev6.htm


Dual IPv4 and IPv6 Protocol Stacks

One technique for transitioning to IPv6 is by using dual IPv4 and IPv6 protocol stacks. Using dual stacks enables gradual, one-by-one upgrades to applications running on nodes. Applications that are upgraded to IPv6 use the IPv6 protocol stack, and applications that are not upgraded and support only IPv4 can coexist with upgraded applications on the same node. New and upgraded applications can use both IPv4 and IPv6 protocol stacks.

The Cisco IOS software supports the dual IPv4 and IPv6 protocol stack technique. When both IPv4 and IPv6 routing are enabled and an interface is configured with both an IPv4 and IPv6 address, the interface forwards both IPv4 and IPv6 traffic.

Figure 35-1 shows a router forwarding both IPv4 and IPv6 traffic through the same interface, based on the IP packet and destination addresses.

Figure 35-1 Dual IPv4 and IPv6 Support on an Interface



The switch uses ternary content addressable memory (TCAM) to store unicast routes, MAC addresses, access control lists (ACLs), and other features, and provides the switch database management (SDM) templates to allocate memory resources depending on how the switch is used. You must use the dual IPv4 and IPv6 template templates to allocate TCAM usage to both IPv4 and IPv6 protocols. See the “[SDM Templates](#)” section on page 35-7.

Unsupported IPv6 Unicast Routing Features

The switch does not support these IPv6 features in this release:

- IPv6 policy-based routing
- IPv6 virtual private network (VPN) routing and forwarding (VRF) table support
- Web Cache Communication Protocol (WCCP) over IPv6
- Support for these IPv6 routing protocols: Multiprotocol Border Gateway Protocol (BGP), Intermediate System-to-Intermediate System (IS-IS) routing, Enhanced Interior Gateway Routing Protocol (EIGRP)
- Simple Network Management Protocol (SNMP) over IPv6 transport
- Cisco Discovery Protocol (CDP) IPv6 address family support
- IPv6 Hot Standby Router Protocol (HSRP)
- DHCPv6
- 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

Limitations

Because IPv6 is implemented in hardware in the switch, some limitations occur due to the use of IPv6 compressed addresses in the TCAM. These hardware limitations result in some loss of functionality and limits some features.

These are feature limitations.

- 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. 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.
- 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.
- The switch cannot correctly forward SNAP-encapsulated IPv6 packets. These packets are corrupted before being forwarded (bridged or routed) and reach the network as corrupted packets.



Note There is a similar limitation for IPv4 SNAP-encapsulated packets, but the packets are dropped at the switch and are not forwarded as corrupted packets.

- 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 egress RSPAN IPv6-routed packets, the source MAC address in the SPAN output packet can be corrupted.
 - When you egress RSPAN IPv6-routed packets, the destination MAC address can be corrupted. 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.

SDM Templates

Catalyst 3560 switches have one TCAM to store unicast routes, MAC addresses, ACLs, and other features. To allocate TCAM resources for different usages, the switch SDM templates prioritize system resources to optimize support for certain features. You select the template that best suits the switch environment by entering the **sdm prefer** global configuration command. For more information about SDM templates, see [Chapter 7, “Configuring SDM Templates.”](#)

If the switch is running the advanced IP services image, you can select SDM templates to support IP Version 6 (IPv6). The dual IPv4 and IPv6 templates allow the switch to be used in dual stack environments (supporting both IPv4 and IPv6).

**Note**

If you try to configure IPv6 without first selecting a dual IPv4 and IPv6 template, a warning message is generated.

- 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 and ACLs in hardware. IPv6 QoS and ACLs are not supported in this release.

**Note**

Using the dual stack templates results in less TCAM capacity allowed for each resource, so do not use if you plan to forward only IPv4 traffic.

Dual IPv4-and IPv6 SDM Templates

These SDM templates support IPv4 and IPv6 environments:

**Note**

This release does not support IPv6 multicast, QoS, or ACLs.

- Dual IPv4 and IPv6 default template—supports Layer 2, multicast, routing, QoS, and ACLs for IPv4, and Layer 2 and routing for IPv6 on the switch
- Dual IPv4 and IPv6 VLAN template—supports basic Layer 2, multicast, QoS, and ACLs for IPv4, and basic Layer 2 for IPv6 on the switch.

**Note**

An IPv4 route requires only one TCAM entry. Because of the hardware compression scheme used for IPv6, an IPv6 route can take more than one TCAM entry, reducing the number of entries forwarded in hardware. For example, for IPv6 directly connected IP addresses, the desktop template might allow less than two thousand entries.

[Table 35-1](#) defines the approximate feature resources allocated by each new template. Template estimations are based on a switch with eight routed interfaces and approximately one thousand VLANs.

**Note**

The current software release does not support IPv6 QoS or ACLs.

Table 35-1 Approximate Feature Resources Allowed by Dual IPv4-IPv6 Templates

Resource	IPv4-and-IPv6 Default	IPv4-and-IPv6 VLAN
Unicast MAC addresses	2 K	8 K
IPv4 IGMP groups and multicast routes	1 K	1 K
Total IPv4 unicast routes:	3 K	0
• Directly connected IPv4 hosts	2 K	0
• Indirect IPv4 routes	1 K	0
IPv6 multicast groups	1 K	1 K

Table 35-1 *Approximate Feature Resources Allowed by Dual IPv4-IPv6 Templates*

Resource	IPv4-and-IPv6 Default	IPv4-and-IPv6 VLAN
Total IPv6 unicast routes:	3 K	0
• Directly connected IPv6 addresses	2 K	0
• Indirect IPv6 unicast routes	1 K	0
IPv4 policy-based routing ACEs	0	0
IPv4 or MAC QoS ACEs (total)	512	512
IPv4 or MAC security ACEs (total)	1 K	1K
IPv6 policy-based routing ACEs	0	0
IPv6 QoS ACEs	510	510
IPv6 security ACEs	510	510

Configuring IPv6

These sections contain this IPv6 forwarding configuration information:

- [Default IPv6 Configuration, page 35-9](#)
- [Configuring IPv6 Addressing and Enabling IPv6 Routing, page 35-10](#)
- [Configuring IPv4 and IPv6 Protocol Stacks, page 35-12](#)
- [Configuring IPv6 ICMP Rate Limiting, page 35-14](#)
- [Configuring CEF for IPv6, page 35-14](#)
- [Configuring Static Routing for IPv6, page 35-15](#)
- [Configuring RIP for IPv6, page 35-17](#)
- [Configuring OSPF for IPv6, page 35-19](#)

Default IPv6 Configuration

[Table 35-2](#) shows the default IPv6 configuration.

Table 35-2 *Default IPv6 Configuration*

Feature	Default Setting
SDM template	Default
IPv6 routing	Disabled globally and on all interfaces
CEFv6 or dCEFv6	Disabled (IPv4 CEF and dCEF are enabled by default) Note When IPv6 routing is enabled, CEFv6 and dCEF6 are automatically enabled.
IPv6 addresses	None configured

Configuring IPv6 Addressing and Enabling IPv6 Routing

This section describes how to assign IPv6 addresses to individual Layer 3 interfaces and enable the forwarding of IPv6 traffic globally on the switch.



Note

In the **ipv6 address** interface configuration command, you must enter the *ipv6-address* and *ipv6-prefix* variables in the form documented in RFC 2373 where the address is 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).

For an interface to forward IPv6 traffic, you must configure an IPv6 address on the interface. Configuring a global 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 FF02:0:0:0:1:ff00::/104 for each unicast address assigned to the interface (this address is used in the neighbor discovery process.)
- all-nodes link-local multicast group FF02::1
- all-routers link-local multicast group FF02::2



Note

Before configuring IPv6 on the switch, be sure to select a dual IPv4 and IPv6 SDM template.

For more information about configuring IPv6 routing, see the “Implementing Basis Connectivity for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ipv6_vgf.htm



Note

Not all features discussed in this chapter are supported by the Catalyst 3560 switch. See the “[Unsupported IPv6 Unicast Routing Features](#)” section on page 35-6.

Beginning in privileged EXEC mode, follow these steps to assign an IPv6 address to a Layer 3 interface and enable IPv6 routing:

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	sdm prefer dual-ipv4-and-ipv6 { default 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. • 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 (after the switch has reloaded).

	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 eui-64</i> or ipv6 address <i>ipv6-address link-local</i> or ipv6 enable	Specify a global IPv6 address with an extended universal 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 ip v6 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.

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 fastethernet0/11
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 fastethernet0/11
FastEthernet0/11 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 IPv4 and IPv6 Protocol Stacks

When you configure an interface with both an IPv4 and IPv6 address, the interface forwards both IPv4 and IPv6 traffic and can send and receive data on both IPv4 and IPv6 networks.

Beginning in privileged EXEC mode, follow these steps to configure a Layer 3 interface to support both the IPv4 and IPv6 and enable IPv6 routing.



Note

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 | 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 for the template to take effect.

	Command	Purpose
Step 1	configure terminal	Enter global configuration mode.
Step 2	ip routing	Enable IPv4 routing on the switch.
Step 3	ipv6 unicast-routing	Enable forwarding of IPv6 data packets on the switch.
Step 4	interface interface-id	Enter interface configuration mode, and specify the Layer 3 interface to configure.
Step 5	no switchport	Remove the interface from Layer 2 configuration mode (if it is a physical interface).
Step 6	ip address ip-address mask [secondary]	Specify a primary or secondary IPv4 address for the interface.

	Command	Purpose
Step 7	ipv6 address <i>ipv6-prefix/prefix length eui-64</i>	Specify a global IPv6 address with an interface identifier 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.
	or	
	ipv6 address <i>ipv6-address link-local</i>	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.
	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.
Step 8	end	Return to privileged EXEC mode.
Step 9	show interface <i>interface-id</i> show ip interface <i>interface-id</i> show ip v6 interface <i>interface-id</i>	Verify your entries.
Step 10	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.

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 fastethernet0/11
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 IPv6 ICMP Rate Limiting

IPv6 ICMP rate limiting uses a token-bucket algorithm for limiting the rate at which IPv6 ICMP error messages are sent to the network. The interval between error messages is specified in a time interval and a bucket size. Because some applications, such as traceroute, sometimes require replies to a group of requests to be sent out in rapid succession, specifying only the interval between error messages can cause the application to fail. The token bucket allows a number of tokens, each representing the ability to send one error message, to be stored in virtual buckets. For every message to be sent, one token is removed from the bucket. If a series of error messages is generated, error messages can be sent until the bucket is empty. When the bucket is empty, IPv6 ICMP error messages are not sent until a new token is placed in the bucket. This method does not increase the average rate-limiting time interval, but it provides more flexibility than fixed-time intervals.

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.

Beginning in privileged EXEC mode, follow these steps to change the ICMP rate-limiting parameters:

	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.

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 used to optimize network performance. CEF implements an advanced IP look-up and forwarding algorithm to deliver maximum Layer 3 switching performance. It is less CPU-intensive than fast-switching route-caching, 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, you must first globally configure forwarding of IPv6 unicast packets by using the **ipv6 unicast-routing** global configuration command, and you must 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 reenable IPv6 CEF if it has been disabled, use the **ipv6 cef** global configuration command. You can verify the IPv6 state by entering the **show ipv6 cef** privileged EXEC command.

Configuring Static Routing for IPv6

Static routes are manually configured and define an explicit route between two networking devices. The benefits of static routes include security and resource efficiency. Static routes use less bandwidth than dynamic routing protocols because there is no requirement for routes to be calculated and communicated. The main disadvantage of using static routes is that static routes are not automatically updated, as with a dynamic routing protocol, and must be manually reconfigured if the network topology changes. Static routes are useful for smaller networks with only one path to an outside network or to provide security for a larger network for certain types of traffic.

There are types of static routes:

- Directly attached static routes—Only the output interface is specified because the destination is assumed to be directly attached to this interface. The packet destination is used as the next hop address. A directly attached static route is valid only when the specified interface is IPv6-enabled and is up.
- Recursive static routes—Only the next hop is specified, and the output interface is derived from the next hop. A recursive static route is valid only when the specified next hop results in a valid IPv6 output interface, the route does not self-recur, and the recursion depth does not exceed the maximum IPv6 forwarding recursion depth.
- Fully specified static routes—Both the output interface and the next hop are specified. The next hop is assumed to be directly attached to the specified output interface. A fully specified route is valid when the specified IPv6 interface is IPv6-enabled and up.
- Floating static routes—Any of the three types of static routes can be floating static routes, used to back up dynamic routes learned through configured routing protocols. A floating static route is configured with a less efficient administrative distance than the routing protocol it is backing up. Therefore, the dynamic route is always used for routing traffic in preference to the floating static route. If the dynamic route is lost, the floating static route is used in its place.



Note

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, and enable IPv6 on at least one Layer 3 interface by configuring an IPv6 address on the interface.

Beginning in privileged EXEC mode, follow these steps to configure an IPv6 static route:

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>]	Configure a static IPv6 route. <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 of the address 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 IPv6 address of the next hop need not be directly connected; recursion is done to find the IPv6 address of the directly connected next hop. The address must be in the form documented in RFC 2373, specified in hexadecimal using 16-bit values between colons. • <i>interface-id</i>—Specify direct static routes from point-to-point and broadcast interfaces. With point-to-point interfaces, there is no need to specify the IPv6 address of the next hop. With 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 also 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 any other type of route except 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 or	show ipv6 static [<i>ipv6-address / ipv6-prefix/prefix length</i>] [interface interface-id] [recursive] [detail]	Verify your entries by displaying the contents of the IPv6 routing table. <ul style="list-style-type: none"> • interface interface-id—(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 included 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 [<i>ipv6-address / ipv6-prefix/prefix length</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.

This example shows how to configure a floating static route to an interface with an administrative distance of 130:

```
Switch(config)# ipv6 route 2001:0DB8::/32 gigabitethernet0/1 130
```

For more information about configuring static IPv6 routing, see the “Implementing Static Routes for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ipv6_vgf.htm

Configuring RIP for IPv6

Routing Information Protocol (RIP) for IPv6 is a distance-vector protocol that uses hop count as a routing metric. IPv6 RIP functions the same and offers the same benefits as RIP in IPv4. IPv6 RIP enhancements include support for IPv6 addresses and prefixes and the use of the all-RIP-routers multicast group address FF02::9 as the destination address for RIP update messages.

Each IPv6 RIP process maintains a local routing table, referred to as a Routing Information Database (RIB), that contains a set of best-cost IPv6 routes learned from all its neighboring networking devices. If IPv6 RIP learns the same route from two different neighbors, but with different costs, it stores only the lowest-cost route in the local RIB. The RIB also stores any expired routes that the RIP process is advertising to its neighbors that are running RIP. If the same route is learned from a different routing protocol with a better administrative distance than IPv6 RIP, the RIP route is not added to the IPv6 RIB, but the route still exists in the IPv6 RIP RIB.



Note

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, and enable IPv6 on any Layer 3 interfaces on which IPv6 RIP is to be enabled.

Beginning in privileged EXEC mode, follow these required and optional steps to configure IPv6 RIP:

	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.
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 four paths.
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 [<i>ipv6-address</i> / <i>ipv6-prefix/prefix</i> <i>length</i> / <i>protocol</i>]	<p>Display information about current IPv6 RIP processes.</p> <p>Display the current contents of the IPv6 routing table.</p>
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.

This example shows how to enable the RIP routing process *cisco*, with a maximum of eight equal-cost routes and enable it on an interface:

```
Switch(config)# ipv6 router rip cisco
Switch(config-router)# maximum-paths 8
Switch(config)# exit
Switch(config)# interface fastethernet0/11
Switch(config-if)# ipv6 rip cisco enable
```

For more information about configuring RIP routing for IPv6, see the “Implementing RIP for IPv6” chapter in the *Cisco IOS IPv6 Configuration Library* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ipv6_vgf.htm

Configuring OSPF for IPv6

Open Shortest Path First (OSPF) is a link-state protocol for IP, which means that routing decisions are based on the states of the links that connect the source and destination devices. The state of a link is a description of the interface and its relationship to its neighboring networking devices. Interface information, which is propagated in link-state advertisements (LSAs), includes the IPv6 prefix of the interface, the network mask, the type of network it is connected to, the routers connected to that network, and so on. LSA data is stored in a link-state database, which is used to create the OSPF routing table. While the database contains a collection of raw data, the routing table contains a list of shortest paths to known destinations that use specific Layer 3 ports. OSPF Version 2 (RFC 2740) supports IPv6.

OSPF for IPv6 is largely the same as OSPF Version 2 (for IPv4), expanded to provide support for IPv6 routing prefixes and the larger size of IPv6 addresses. However, there are these differences:

- Enabling OSPF for IPv6 on an interface automatically creates a routing process and its associated configuration; you do not need to explicitly create a routing process as in IPv4.
- In OSPF for IPv6, you must enable OSPF on each interface by using commands in interface configuration mode. In OSPF Version 2, interfaces are indirectly enabled by using router configuration mode.
- In IPv6, you can configure many address prefixes on an interface. All address prefixes configured on an interface are included by default; you cannot select a subset of address prefixes to import.
- Unlike OSPF Version 2, multiple instances of IPv6 can run on a link.
- OSPF Version 2 uses the 32-bit IPv4 address configured on the interface to choose an IPv4 address to use as the router ID. When you enable OSPF for IPv6 on an interface, if an IPv4 address is configured on the interface that IP address is used for the IPv6 router ID. If no IPv4 address is configured on the interface, you must use the **router-id** router configuration command to configure a router ID before the OSPF process is started.

OSPF automatically chooses a loopback interface over other interfaces and chooses the highest IP address among all loopback interfaces. If no loopback interfaces are present, OSPF selects the highest IP address in the router. You cannot configure OSPF to use any particular interface.

You can customize OSPF for IPv6 for your network, but you will most likely not need to. The defaults for OSPF in IPv6 are set to meet the requirements of most customers and features.

**Note**

Be careful when changing the defaults for IPv6 commands. Changing the defaults might adversely affect OSPF for the IPv6 network.

**Note**

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, and enable IPv6 on Layer 3 interfaces on which you are enabling IPv6 OSPF.

Beginning in privileged EXEC mode, follow these required and optional steps to configure IPv6 OSPF:

	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 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.
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 *Cisco IOS IPv6 Configuration Library* at this URL:

http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t13/ipv6_vgf.htm

Displaying IPv6

Table 35-3 shows the privileged EXEC commands for monitoring IPv6 on the switch.

Table 35-3 Commands for Monitoring IPv6

Command	Purpose
show ipv6 access-list	Display a summary of access lists.
show ipv6 cef	Display Cisco Express Forwarding for IPv6.
show ipv6 interface interface-id	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 a list of 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 the IPv6 route table entries.
show ipv6 routers	Display the local IPv6 routers.
show ipv6 static	Display IPv6 static routes.
show ipv6 traffic	Display IPv6 traffic statistics.

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 FastEthernet1/0/11
3FFE:C000:111:1:20B:46FF:FE2F:D945/128
  receive
3FFE:C000:168:1::/64
  attached to FastEthernet2/0/43
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
    FastEthernet0/4
    FastEthernet0/11
    FastEthernet0/12
  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
    FastEthernet0/4
    FastEthernet0/11
    FastEthernet0/12
  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 Fa1/0/13
```

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 ::, FastEthernet0/11
L    3FFE:C000:111:1:20B:46FF:FE2F:D945/128 [0/0]
C    3FFE:C000:168:1::/64 [0/0]
     via ::, FastEthernet0/4
L    3FFE:C000:168:1:20B:46FF:FE2F:D94B/128 [0/0]
     via ::, FastEthernet0/4
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
      unreach: 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
      unreach: 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
```