



IPv6 Neighbor Discovery

Last Updated: September 19, 2012

The IPv6 neighbor discovery process uses Internet Control Message Protocol (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 a neighbor, and track neighboring devices.

- [Finding Feature Information, page 1](#)
- [Information About IPv6 Neighbor Discovery, page 1](#)
- [How to Configure IPv6 Neighbor Discovery, page 7](#)
- [Configuration Examples for IPv6 Neighbor Discovery, page 10](#)
- [Additional References, page 12](#)
- [Feature Information for IPv6 Neighbor Discovery, page 13](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About IPv6 Neighbor Discovery

- [IPv6 Neighbor Discovery, page 1](#)

IPv6 Neighbor Discovery

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 a neighbor, and track neighboring devices.

The IPv6 static cache entry for neighbor discovery feature allows static entries to be made in the IPv6 neighbor cache. Static routing requires an administrator to manually enter IPv6 addresses, subnet masks,



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA


gateways, and corresponding Media Access Control (MAC) addresses for each interface of each device into a table. Static routing enables more control but requires more work to maintain the table. The table must be updated each time routes are added or changed.

- [IPv6 Neighbor Solicitation Message, page 2](#)
- [IPv6 Router Advertisement Message, page 3](#)
- [IPv6 Neighbor Redirect Message, page 5](#)

IPv6 Neighbor Solicitation Message

A value of 135 in the Type field of the ICMP packet header identifies a neighbor solicitation message. Neighbor solicitation messages are sent on the local link when a node wants to determine the link-layer address of another node on the same local link (see the figure below). When a node wants to determine the link-layer address of another node, the source address in a neighbor solicitation message is the IPv6 address of the node sending the neighbor solicitation message. The destination address in the neighbor solicitation message is the solicited-node multicast address that corresponds to the IPv6 address of the destination node. The neighbor solicitation message also includes the link-layer address of the source node.

Figure 1 IPv6 Neighbor Discovery: Neighbor Solicitation Message



After receiving the neighbor solicitation message, the destination node replies by sending a neighbor advertisement message, which has a value of 136 in the Type field of the ICMP packet header, on the local link. The source address in the neighbor advertisement message is the IPv6 address of the node (more specifically, the IPv6 address of the node interface) sending the neighbor advertisement message. The destination address in the neighbor advertisement message is the IPv6 address of the node that sent the neighbor solicitation message. The data portion of the neighbor advertisement message includes the link-layer address of the node sending the neighbor advertisement message.

After the source node receives the neighbor advertisement, the source node and destination node can communicate.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verify the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link. When there is such a change, the destination address for the neighbor advertisement is the all-nodes multicast address.

Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. Neighbor unreachability detection identifies the failure of a neighbor or the failure of the forward path to the neighbor, and is used for all paths between hosts and neighboring nodes (hosts or devices). Neighbor unreachability detection is performed for neighbors to which only unicast packets are being sent and is not performed for neighbors to which multicast packets are being sent.

A neighbor is considered reachable when a positive acknowledgment is returned from the neighbor (indicating that packets previously sent to the neighbor have been received and processed). A positive acknowledgment from an upper-layer protocol (such as TCP) indicates that a connection is making forward progress (reaching its destination) or the receipt of a neighbor advertisement message in response to a neighbor solicitation message. If packets are reaching the peer, they are also reaching the next-hop neighbor of the source. Therefore, forward progress is also a confirmation that the next-hop neighbor is reachable.

For destinations that are not on the local link, forward progress implies that the first-hop device is reachable. When acknowledgments from an upper-layer protocol are not available, a node probes the neighbor using unicast neighbor solicitation messages to verify that the forward path is still working.

The return of a solicited neighbor advertisement message from the neighbor is a positive acknowledgment that the forward path is still working (neighbor advertisement messages that have the solicited flag set to a value of 1 are sent only in response to a neighbor solicitation message). Unsolicited messages confirm only the one-way path from the source to the destination node; solicited neighbor advertisement messages indicate that a path is working in both directions.

**Note**

A neighbor advertisement message that has the solicited flag set to a value of 0 must not be considered as a positive acknowledgment that the forward path is still working.


Neighbor solicitation messages are also used in the stateless autoconfiguration process to verify the uniqueness of unicast IPv6 addresses before the addresses are assigned to an interface. Duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed). Specifically, a node sends a neighbor solicitation message with an unspecified source address and a tentative link-local address in the body of the message. If another node is already using that address, the node returns a neighbor advertisement message that contains the tentative link-local address. If another node is simultaneously verifying the uniqueness of the same address, that node also returns a neighbor solicitation message. If no neighbor advertisement messages are received in response to the neighbor solicitation message and no neighbor solicitation messages are received from other nodes that are attempting to verify the same tentative address, the node that sent the original neighbor solicitation message considers the tentative link-local address to be unique and assigns the address to the interface.

Every IPv6 unicast address (global or link-local) must be verified for uniqueness on the link; however, until the uniqueness of the link-local address is verified, duplicate address detection is not performed on any other IPv6 addresses associated with the link-local address. The Cisco implementation of duplicate address detection in the Cisco software does not verify the uniqueness of anycast or global addresses that are generated from 64-bit interface identifiers.

IPv6 Router Advertisement Message

Router advertisement (RA) messages, which have a value of 134 in the Type field of the ICMP packet header, are periodically sent out each configured interface of an IPv6 device. For stateless autoconfiguration to work properly, the advertised prefix length in RA messages must always be 64 bits. The RA messages are sent to the all-nodes multicast address (see the figure below).

Figure 2 IPv6 Neighbor Discovery: RA Message



RA messages typically include the following information:

- One or more onlink IPv6 prefixes that nodes on the local link can use to automatically configure their IPv6 addresses
- Lifetime information for each prefix included in the advertisement
- Sets of flags that indicate the type of autoconfiguration (stateless or stateful) that can be completed
- Default device information (whether the device sending the advertisement should be used as a default device and, if so, the amount of time, in seconds, the device should be used as a default device)
- Additional information for hosts, such as the hop limit and maximum transmission unit (MTU) a host should use in packets that it originates

RAs are also sent in response to device solicitation messages. Device solicitation messages, which have a value of 133 in the Type field of the ICMP packet header, are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled RA message. Given that device solicitation messages are usually sent by hosts at system startup (the host does not have a configured unicast address), the source address in device solicitation messages is usually the unspecified IPv6 address (0:0:0:0:0:0:0:0). If the host has a configured unicast address, the unicast address of the interface sending the device solicitation message is used as the source address in the message. The destination address in device solicitation messages is the all-devices multicast address with a scope of the link. When an RA is sent in response to a device solicitation, the destination address in the RA message is the unicast address of the source of the device solicitation message.

The following RA message parameters can be configured:

- The time interval between periodic RA messages
- The “device lifetime” value, which indicates the usefulness of a device as the default device (for use by all nodes on a given link)
- The network prefixes in use on a given link
- The time interval between neighbor solicitation message retransmissions (on a given link)
- The amount of time a node considers a neighbor reachable (for use by all nodes on a given link)

The configured parameters are specific to an interface. The sending of RA messages (with default values) is automatically enabled on Ethernet and FDDI interfaces when the **ipv6 unicast-routing** command is configured. For other interface types, the sending of RA messages must be manually configured by using

the **no ipv6 nd ra suppress** command. The sending of RA messages can be disabled on individual interfaces by using the **ipv6 nd ra suppress** command.

- [Default Router Preferences for Traffic Engineering, page 5](#)

Default Router Preferences for Traffic Engineering

Hosts discover and select default devices by listening to Router Advertisements (RAs). Typical default device selection mechanisms are suboptimal in certain cases, such as when traffic engineering is needed. For example, two devices on a link may provide equivalent but not equal-cost routing, and policy may dictate that one of the devices is preferred. Some examples are as follows:


- Multiple devices that route to distinct sets of prefixes—Redirects (sent by nonoptimal devices for a destination) mean that hosts can choose any device and the system will work. However, traffic patterns may mean that choosing one of the devices would lead to considerably fewer redirects.
- Accidentally deploying a new device—Deploying a new device before it has been fully configured could lead to hosts adopting the new device as a default device and traffic disappearing. Network managers may want to indicate that some devices are more preferred than others.
- Multihomed situations—Multihomed situations may become more common, because of multiple physical links and because of the use of tunneling for IPv6 transport. Some of the devices may not provide full default routing because they route only to the 6-to-4 prefix or they route only to a corporate intranet. These situations cannot be resolved with redirects, which operate only over a single link.

The default router preference (DRP) feature provides a basic preference metric (low, medium, or high) for default devices. The DRP of a default device is signaled in unused bits in RA messages. This extension is backward compatible, both for devices (setting the DRP bits) and hosts (interpreting the DRP bits). These bits are ignored by hosts that do not implement the DRP extension. Similarly, the values sent by devices that do not implement the DRP extension will be interpreted by hosts that do implement it as indicating a “medium” preference. DRPs need to be configured manually.

IPv6 Neighbor Redirect Message

A value of 137 in the type field of the ICMP packet header identifies an IPv6 neighbor redirect message. Devices send neighbor redirect messages to inform hosts of better first-hop nodes on the path to a destination (see the figure below).

Figure 3 IPv6 Neighbor Discovery: Neighbor Redirect Message



Note

A device must be able to determine the link-local address for each of its neighboring devices in order to ensure that the target address (the final destination) in a redirect message identifies the neighbor device by its link-local address. For static routing, the address of the next-hop device should be specified using the link-local address of the device; for dynamic routing, all IPv6 routing protocols must exchange the link-local addresses of neighboring devices.

After forwarding a packet, a device should send a redirect message to the source of the packet under the following circumstances:

- The destination address of the packet is not a multicast address.
- The packet was not addressed to the device.
- The packet is about to be sent out the interface on which it was received.
- The device determines that a better first-hop node for the packet resides on the same link as the source of the packet.
- The source address of the packet is a global IPv6 address of a neighbor on the same link, or a link-local address.

Use the **ipv6 icmp error-interval** command to limit the rate at which the device generates all IPv6 ICMP error messages, including neighbor redirect messages, which ultimately reduces link-layer congestion.

**Note**

A device must not update its routing tables after receiving a neighbor redirect message, and hosts must not originate neighbor redirect messages.

How to Configure IPv6 Neighbor Discovery

- [Tuning the Parameters for IPv6 Neighbor Discovery, page 7](#)
- [Customizing IPv6 ICMP Rate Limiting, page 8](#)
- [Displaying IPv6 Redirect Messages, page 9](#)

Tuning the Parameters for IPv6 Neighbor Discovery

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ipv6 nd nud retry** *base interval max-attempts*
5. **ipv6 nd cache expire** *expire-time-in-seconds* [**refresh**]
6. **ipv6 nd na glean**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface <i>type number</i> Example: Device(config)# interface GigabitEthernet 1/0/0	Specifies an interface type and number, and places the device in interface configuration mode.

Command or Action	Purpose
Step 4 <code>ipv6 nd nud retry <i>base interval max-attempts</i></code> Example: Device(config-if)# ipv6 nd nud retry 1 1000 3	Configures the number of times NUD resends neighbor solicitations.
Step 5 <code>ipv6 nd cache expire <i>expire-time-in-seconds</i> [refresh]</code> Example: Device(config-if)# ipv6 nd cache expire 7200	Configures the length of time before an IPv6 ND cache entry expires.
Step 6 <code>ipv6 nd na glean</code> Example: Device(config-if)# ipv6 nd na glean	Configures ND to glean an entry from an unsolicited NA.

Customizing IPv6 ICMP Rate Limiting

SUMMARY STEPS

1. enable
2. configure terminal
3. `ipv6 icmp error-interval milliseconds [bucketsize]`

DETAILED STEPS

Command or Action	Purpose
Step 1 <code>enable</code> Example: Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2 <code>configure terminal</code> Example: Device# configure terminal	Enters global configuration mode.

Command or Action	Purpose
Step 3 <code>ipv6 icmp error-interval</code> <i>milliseconds</i> [<i>bucketsize</i>] Example: Device(config)# <code>ipv6 icmp error-interval 50 20</code>	Customizes the interval and bucket size for IPv6 ICMP error messages.

Displaying IPv6 Redirect Messages

SUMMARY STEPS

1. `enable`
2. `show ipv6 interface` [*brief*] [*type number*] [*prefix*]
3. `show ipv6 neighbors` [*interface-type interface-number* | *ipv6-address* | *ipv6-hostname*] *statistics*
4. `show ipv6 route` [*ipv6-address* | *ipv6-prefix/prefix-length* | *protocol* | *interface-type interface-number*]
5. `show ipv6 traffic`
6. `show hosts` [*vrf vrf-name* | *all* | *hostname* | *summary*]
7. `enable`
8. `show running-config`

DETAILED STEPS

Command or Action	Purpose
Step 1 <code>enable</code> Example: Device# <code>enable</code>	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2 <code>show ipv6 interface</code> [<i>brief</i>] [<i>type number</i>] [<i>prefix</i>] Example: Device# <code>show ipv6 interface gigabitethernet 0/0/0</code>	Displays the usability status of interfaces configured for IPv6.
Step 3 <code>show ipv6 neighbors</code> [<i>interface-type interface-number</i> <i>ipv6-address</i> <i>ipv6-hostname</i>] <i>statistics</i> Example: Device# <code>show ipv6 neighbors gigabitethernet 2/0/0</code>	Displays IPv6 Neighbor Discovery cache information.

Command or Action	Purpose
Step 4 <code>show ipv6 route</code> [<i>ipv6-address</i> <i>ipv6-prefix/prefix-length</i> <i>protocol</i> <i>interface-type interface-number</i>] Example: Device# <code>show ipv6 route</code>	(Optional) Displays the current contents of the IPv6 routing table.
Step 5 <code>show ipv6 traffic</code> Example: Device# <code>show ipv6 traffic</code>	(Optional) Displays statistics about IPv6 traffic.
Step 6 <code>show hosts</code> [<i>vrf vrf-name</i> all <i>hostname</i> summary] Example: Device# <code>show hosts</code>	Displays the default domain name, the style of name lookup service, a list of name server hosts, and the cached list of hostnames and addresses.
Step 7 <code>enable</code> Example: Device> <code>enable</code>	Enables privileged EXEC mode. Enter your password if prompted.
Step 8 <code>show running-config</code> Example: Device# <code>show running-config</code>	Displays the current configuration running on the device.

Configuration Examples for IPv6 Neighbor Discovery

- [Example: Customizing the Parameters for IPv6 Neighbor Discovery, page 10](#)
- [Example: IPv6 ICMP Rate Limiting Configuration, page 11](#)
- [Example: Displaying Information About ICMP Rate-Limited Counters, page 11](#)
- [Example: Displaying IPv6 Interface Statistics, page 11](#)

Example: Customizing the Parameters for IPv6 Neighbor Discovery

In the following example, IPv6 ND NA gleaning is enabled and the IPv6 ND cache expiry is set to 7200 seconds (2 hours):

```
interface Port-channel189
```

```

no ip address
ipv6 address FC07::789:1:0:0:3/64
ipv6 nd reachable-time 2700000
ipv6 nd na glean
ipv6 nd cache expire 7200
no ipv6 redirects
standby version 2
standby 2 ipv6 FC07::789:1:0:0:1/64
standby 2 priority 150
standby 2 preempt

```

Example: IPv6 ICMP Rate Limiting Configuration

The following example shows an interval of 50 milliseconds and a bucket size of 20 tokens being configured for IPv6 ICMP error messages:

```

ipv6 icmp error-interval 50 20

```

Example: Displaying Information About ICMP Rate-Limited Counters

In the following example, information about ICMP rate-limited counters is displayed:

```

Device# show ipv6 traffic

```

```

ICMP statistics:
  Rcvd: 188 input, 0 checksum errors, 0 too short
        0 unknown info type, 0 unknown error type
        unreachable: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
        parameter: 0 error, 0 header, 0 option
        0 hopcount expired, 0 reassembly timeout, 0 too big
        0 echo request, 0 echo reply
        0 group query, 0 group report, 0 group reduce
        1 router solicit, 175 router advert, 0 redirects
        0 neighbor solicit, 12 neighbor advert
  Sent: 7376 output, 56 rate-limited
        unreachable: 0 routing, 15 admin, 0 neighbor, 0 address, 0 port
        parameter: 0 error, 0 header, 0 option
        0 hopcount expired, 0 reassembly timeout, 0 too big
        15 echo request, 0 echo reply
        0 group query, 0 group report, 0 group reduce
        0 router solicit, 7326 router advert, 0 redirects
        2 neighbor solicit, 22 neighbor advert

```

Example: Displaying IPv6 Interface Statistics

In the following example, the **show ipv6 interface** command is used to verify that IPv6 addresses are configured correctly for FastEthernet interface 1/0. Information may also be displayed about the status of IPv6 neighbor redirect messages, IPv6 neighbor discovery messages, stateless autoconfiguration, and MTU size.

```

Device# show ipv6 interface fastethernet 1/0

```

```

Ethernet0 is up, line protocol is up
IPv6 is stalled, link-local address is FE80::1
Global unicast address(es):
  2001:DB8:2000::1, subnet is 2001:DB8:2000::/64
  2001:DB8:3000::1, subnet is 2001:DB8:3000::/64
Joined group address(es):
  FF02::1
  FF02::2
  FF02::1:FF00:1
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.

```

Additional References

Related Documents

Related Topic	Document Title
IPv6 addressing and connectivity	<i>IPv6 Configuration Guide</i>
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
IPv6 commands	<i>Cisco IOS IPv6 Command Reference</i>
Cisco IOS IPv6 features	Cisco IOS IPv6 Feature Mapping

Standards and RFCs

Standard/RFC	Title
RFCs for IPv6	<i>IPv6 RFCs</i>

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for IPv6 Neighbor Discovery

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1 Feature Information for IPv6 Neighbor Discovery

Feature Name	Releases	Feature Information
IPv6 Neighbor Discovery	12.0(22)S	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 a neighbor, and track neighboring devices.
	12.2(2)T	
	12.2(14)S	
	12.2(17a)SX1	
	12.2(25)SG	
	12.2(28)SB	
	12.2(33)SRA	The following commands were introduced or modified: ipv6 nd cache expire , ipv6 nd na glean , ipv6 nd nud retry .
	Cisco IOS XE Release 2.1 3.2.0SG	
IPv6: ICMPv6	12.0(22)S	ICMP in IPv6 functions the same as ICMP in IPv4. ICMP generates error messages, such as ICMP destination unreachable messages, and informational messages, such as ICMP echo request and reply messages.
	12.2(2)T	
	12.2(14)S	
	12.2(17a)SX1	
	12.2(25)SG	
	12.2(28)SB	
	12.2(33)SRA	No commands were introduced or modified.
	Cisco IOS XE Release 2.1	
IPv6: ICMPv6 Redirect	12.0(22)S	The IPv6 Redirect Messages feature enables a device to send ICMP IPv6 neighbor redirect messages to inform hosts of better first-hop nodes on the path to a destination.
	12.2(4)T	
	12.2(14)S	
	12.2(17a)SX1	
	12.2(25)SG	
	12.2(28)SB	
	12.2(33)SRA	The following commands were introduced or modified: show ipv6 interface , show ipv6 neighbors , show ipv6 route , show ipv6 traffic .
	Cisco IOS XE Release 2.1	

Feature Name	Releases	Feature Information
IPv6 ICMP Rate Limiting	12.2(8)T Cisco IOS XE Release 2.1	The IPv6 ICMP Rate Limiting feature implements a token bucket algorithm for limiting the rate at which IPv6 ICMP error messages are sent out on the network. The following command was introduced or modified: ipv6 icmp error-interval .
IPv6: Neighbor Discovery Duplicate Address Detection	12.0(22)S 12.2(4)T 12.2(17a)SX1 12.2(14)S 12.2(25)SG 12.2(28)SB 12.2(33)SRA Cisco IOS XE Release 2.1	IPv6 neighbor discovery duplicate address detection is performed first on a new, link-local IPv6 address before the address is assigned to an interface (the new address remains in a tentative state while duplicate address detection is performed).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2012 Cisco Systems, Inc. All rights reserved.