



IPv6 Neighbor Discovery

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.

- [Information About IPv6 Neighbor Discovery, on page 1](#)
- [How to Configure IPv6 Neighbor Discovery, on page 5](#)
- [Configuration Examples for IPv6 Neighbor Discovery, on page 6](#)
- [Additional References, on page 7](#)

Information About IPv6 Neighbor Discovery

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, 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

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 router. 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 router information (whether the router sending the advertisement should be used as a default router and, if so, the amount of time (in seconds) the router should be used as a default router)
- Additional information for hosts, such as the hop limit and MTU a host should use in packets that it originates

RAs are also sent in response to router solicitation messages.

The following RA message parameters can be configured:

- The time interval between periodic RA messages
- The "router lifetime" value, which indicates the usefulness of a router as the default router (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 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 rasuppress** command.

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 Advertisement Message

The IPv6 neighbor advertisement message is a response to the IPv6 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.

IPv6 Router Solicitation Message

When a host does not have a configured unicast address, for example at system startup, it sends a router solicitation message. A router solicitation enables the host to autoconfigure itself quickly without having to wait for the next scheduled router advertisement message.

Given that router solicitation messages are usually sent by hosts at system startup (the host does not have a configured unicast address), the source address in router 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 router solicitation message is used as the source address in the message. The destination address in router solicitation messages is the all-routers multicast address (FF02::2) with the link-local scope. When an RA is sent in response to a router solicitation, the destination address in the RA message is the unicast address of the source of the router solicitation message.

How to Configure IPv6 Neighbor Discovery

Tuning the Parameters for IPv6 Neighbor Discovery

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.

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

Configuration Examples for IPv6 Neighbor Discovery

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 nud retry 1 1000 3 1000
ipv6 nd na glean
ipv6 nd cache expire 7200
no ipv6 redirects
```

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>

Related Topic	Document Title
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