

IPv6 Clients

- IPv6 Client Mobility, on page 1
- Prerequisites for Configuring IPv6 Mobility, on page 1
- Restrictions on Configuring IPv6 Mobility, on page 2
- Global IPv6, on page 2
- RA Guard, on page 3
- RA Throttling, on page 4
- IPv6 Neighbor Discovery, on page 5

IPv6 Client Mobility

Internet Protocol version 6 (IPv6) is the next-generation network layer Internet protocol intended to replace version 4 (IPv4) in the TCP/IP suite of protocols. This new version increases the Internet global address space to accommodate users and applications that require unique global IP addresses. IPv6 incorporates 128-bit source and destination addresses, which provide significantly more addresses than the 32-bit IPv4 addresses.

To support IPv6 clients across controllers, ICMPv6 messages must be dealt with specially to ensure the IPv6 client remains on the same Layer 3 network. The controllers keep track of IPv6 clients by intercepting the ICMPv6 messages to provide seamless mobility and protect the network from network attacks. The ICMPv6 packets are converted from multicast to unicast and delivered individually per client. This process allows more control. Specific clients can receive specific Neighbor Discovery and Router Advertisement packets, which ensures correct IPv6 addressing and avoids unnecessary multicast traffic.

The configuration for IPv6 mobility is the same as IPv4 mobility and requires no separate software on the client side to achieve seamless roaming. The controllers must be part of the same mobility group. Both IPv4 and IPv6 client mobility are enabled by default.

Prerequisites for Configuring IPv6 Mobility

- Up to eight client addresses can be tracked per client.
- To allow stateful DHCPv6 IP addressing to operate properly, you must have a switch or router that supports the DHCP for IPv6 feature that is configured to act like a DHCPv6 server, or you need a dedicated server such as a Windows 2008 server with a built-in DHCPv6 server.

To support the seamless IPv6 Mobility, you might need to configure the following:

- Configuring RA Guard for IPv6 Clients
- Configuring RA Throttling for IPv6 Clients
- Configuring IPv6 Neighbor Discovery Caching

Restrictions on Configuring IPv6 Mobility

- The Dynamic VLAN function for IPv6 is not supported.
- Roaming of IPv6 clients that are associated with a WLAN that is mapped to an untagged interface to another WLAN that is mapped to a tagged interface is not supported.
- The controllers that have the same mobility group, same VLAN ID, and different IPv4 and IPv6 subnets, generate different IPv6 router advertisements. WLAN on these controllers is assigned to the same dynamic interface with the same VLAN ID on all the controllers. The client receives the correct IPv4 address; however, it receives a router advertisement from the different subnets that reach the other controllers. There could be an issue of no traffic from the client because the first given IPv6 address to the client does not match to the subnet for the IPv4 address. To resolve this, make sure if performing Layer 3 roams between controllers that the client is assigned to different VLANs.
- IPv6 is not supported in Flex local switching with AAA override VLAN.
- IPv6 ping from controller to a client is not supported if the client is in the management subnet.
- Controller sends all application IPv6 traffic to the gateway even if the host is in the same subnet. The gateway forwards the traffic to the host in the same subnet. If the gateway is a Cisco ASA, by default, the Cisco ASA drops traffic sent by the controller to the gateway, if traffic has to be sent to the same subnet. This is because traffic ingress and egress interface is the same. To allow Cisco ASA to forward this traffic, use the **same-security-traffic permit intra-interface** command in Cisco ASA. For more information, see https://www.cisco.com/c/en/us/td/docs/security/asa/asa92/configuration/vpn/asa-vpn-cli/vpn-params.html#56144.

Global IPv6

This section contains the following subsections:

Restrictions on Global IPv6

- IPv4 address needs to be configured on the interface prior to configuring the IPv6 address.

Configuring IPv6 Globally (GUI)

Procedure

Step 1 Choose **Controller** > **General**.

Step 2 From the Global IPv6 Config drop-down list, choose **Enabled** or **Disabled**.

Step 3 Click **Apply**.

Step 4 Click **Save Configuration**.

Configuring IPv6 Globally (CLI)

Procedure

- Enable or disable IPv6 globally by entering this command:

  config ipv6 {**enable** | **disable**}

RA Guard

IPv6 clients configure IPv6 addresses and populate their router tables based on IPv6 Router Advertisement (RA) packets. The RA Guard feature is similar to the RA guard feature of wired networks. RA Guard increases the security of the IPv6 network by dropping the unwanted or rogue RA packets that come from wireless clients. If this feature is not configured, malicious IPv6 clients could announce themselves as the router for the network, which would take higher precedence over legitimate IPv6 routers.

RA Guard occurs at the controller. You can configure the controller to drop RA messages at the access point or at the controller. By default, RA Guard is configured at the access point and also enabled in the controller. All IPv6 RA messages are dropped, which protects other wireless clients and upstream wired network from malicious IPv6 clients.

Note

- IPv6 RA guard feature works on wireless clients only. This feature does not work on wired guest access (GA).
- RA guard is also supported in FlexConnect local switching mode.

This section contains the following subsections:

Configuring RA Guard (GUI)

Procedure

Step 1 Choose **Controller > IPv6 > RA Guard** to open the IPv6 RA Guard page. By default the IPv6 RA Guard on AP is enabled.

Step 2 From the drop-down list, choose **Disable** to disable RA Guard. The controller also displays the clients that have been identified as sending RA packets.

Step 3 Click **Apply** to commit your changes.

Step 4 Click **Save Configuration** to save your changes.

Configuring RA Guard (CLI)

Procedure

- Configure RA Guard by entering this command:

  config ipv6 ra-guard ap {**enable** | **disable**}

RA Throttling

RA throttling allows the controller to enforce limits to RA packets headed toward the wireless network. By enabling RA throttling, routers that send many RA packets can be trimmed to a minimum frequency that will still maintain an IPv6 client connectivity. If a client sends an RS packet, then an RA is sent back to the client. This is allowed through the controller and unicasted to the client. This process ensures that the new clients or roaming clients are not affected by the RA throttling.

This section contains the following subsections:

Configuring RA Throttling (GUI)

Procedure

Step 1 Choose **Controller > IPv6 > RA Throttle Policy** page. By default the IPv6 RA Throttle Policy is disabled. Unselect the check box to disable RA throttle policy.

Step 2 Configure the following parameters:

- **Throttle period**—The period of time for throttling. RA throttling takes place only after the Max Through limit is reached for the VLAN or the Allow At-Most value is reached for a particular router. The range is from 10 seconds to 86400 seconds. The default is 600 seconds.
- **Max Through**—The maximum number of RA packets on a VLAN that can be sent before throttling takes place. The No Limit option allows an unlimited number of RA packets through with no throttling. The range is from 0 to 256 RA packets. The default is 10 RA packets.
- **Interval Option**—This option allows the controller to act differently based on the RFC 3775 value set in IPv6 RA packets.
  - **Passthrough**— Allows any RA messages with the RFC 3775 interval option to go through without throttling.
  - **Ignore**—Causes the RA throttle to treat packets with the interval option as a regular RA and subject to throttling if in effect.
  - **Throttle**—Causes the RA packets with the interval option to always be subject to rate limiting.

- **Allow At-least**—The minimum number of RA packets per router that can be sent as multicast before throttling takes place. The range is from 0 to 32 RA packets.
- **Allow At-most**—The maximum number of RA packets per router that can be sent as multicast before throttling takes place. The No Limit option allows an unlimited number of RA packets through the router. The range is from 0 to 256 RA packets.

Note
When RA throttling occurs, only the first IPv6 capable router is allowed through. For networks that have multiple IPv6 prefixes being served by different routers, you should disable RA throttling.

Step 3 Save the configuration.

Configuring the RA Throttle Policy (CLI)

Procedure

Configure the RA throttle policy by entering this command:

config ipv6 neigbhor-binding ra-throttle {**allow at-least** *at-least-value* | **enable** | **disable** | **interval-option** { **ignore** | **passthrough** | **throttle**} | **max-through** {**max-through-value** | **no-limit**}}

IPv6 Neighbor Discovery

IPv6 Neighbor Discovery is a set of messages and processes that determine relationships between neighboring nodes. Neighbor Discovery replaces ARP, ICMP Router Discovery, and ICMP Redirect used in IPv4.

At any given time, only eight IPv6 addresses are supported per client. When the ninth IPv6 address is encountered, the controller removes the oldest stale entry and accommodates the latest one.

IPv6 Neighbor Discovery inspection analyzes neighbor discovery messages in order to build a trusted binding table database, and IPv6 neighbor discovery packets that do not comply are dropped. The neighbor binding table in the controller track each IPv6 address and its associated MAC address. Clients are expired from the table according to Neighbor Binding timers.

This section contains the following subsections:

Configuring Neighbor Binding (GUI)

Procedure

Step 1 Choose **Controller > IPv6 > Neighbor Binding** page.

Step 2 Configure the following:

- Down–Lifetime—Specifies how long IPv6 cache entries are kept if the interface goes down. The range is from 0 to 86400 seconds.
- Reachable–Lifetime—Specifies how long IPv6 addresses are active. The range is from 0 to 86400 seconds.
- Stale–Lifetime—Specifies how long to keep IPv6 addresses in the cache. The range is from 0 to 86400 seconds.

Step 3 Enable or disable the Unknown Address Multicast NS Forwarding.

Step 4 Enable or disable NA Multicast Forwarding.

If you enable NA Multicast Forwarding, all unsolicited multicast NA from Wired/Wireless is not forwarded to Wireless.

Step 5 Click **Apply**.

Step 6 Click **Save Configuration**.

Configuring Neighbor Binding (CLI)

Procedure

- Configure the neighbor binding parameters by entering this command:

  config ipv6 neighbor-binding timers {**down-lifetime** | **reachable-lifetime** | **stale-lifetime**} {**enable** | **disable**}
- Configure the Unknown Address Multicast NS Forwarding by entering this command:

  config ipv6 ns-mcast-fwd {**enable** | **disable**}
- Configure NA Multicast Forwarding by entering this command:

  config ipv6 na-mcast-fwd {**enable** | **disable**}

  If you enable NA Multicast Forwarding, all unsolicited multicast NA from Wired/Wireless is not forwarded to Wireless.
- See the status of neighbor binding data that are configured on the controller by entering this command:

  show ipv6 neighbor-binding summary