

Wireless LAN IPv6 Client Deployment Guide

Document ID: 113427

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Prerequisites for Wireless IPv6 Client Connectivity

- SLAAC Address Assignment
- DHCPv6 Address Assignment
- Additional Information

IPv6 Client Mobility

Support for VLAN Select (Interface Groups)

First Hop Security for IPv6 Clients

- Router Advertisement Guard
- DHCPv6 Server Guard
- IPv6 Source Guard
- IPv6 Address Accounting
- IPv6 Access Control Lists

Packet Optimization for IPv6 Clients

- Neighbor Discovery Caching
- Router Advertisement Throttling

IPv6 Guest Access

IPv6 VideoStream

IPv6 Quality of Service

IPv6 and FlexConnect

- FlexConnect Local Switching WLANs
- FlexConnect Central Switching WLANs

IPv6 Clients Visibility with NCS

- IPv6 Dashboard Items
- Monitor IPv6 Clients

Configuration for Wireless IPv6 Client Support

- Multicast Distribution Mode to APs
- Configure IPv6 Mobility
- Configure IPv6 Multicast
- Configure IPv6 RA Guard
- Configure IPv6 Access Control Lists
- Configure IPv6 Guest Access for External Web Authentication
- Configure IPv6 RA Throttling
- Configure the IPv6 Neighbor Binding Table
- Configure IPv6 VideoStream

Troubleshoot IPv6 Client Connectivity

- Certain Clients are Unable to Pass IPv6 Traffic
- Verify Successful Layer 3 Roaming for an IPv6 Client:
- Useful IPv6 CLI Commands:


Frequently Asked Questions

Related Information

Introduction

This document provides information on the theory of operation and configuration for the Cisco Unified Wireless LAN solution as it pertains to supporting IPv6 clients.

IPv6 Wireless Client Connectivity



The IPv6 feature set within the Cisco Unified Wireless Network software release v7.2 allows the wireless network to support IPv4, Dual-Stack, and IPv6-only clients on the same wireless network. The overall goal for the addition of IPv6 client support to the Cisco Unified Wireless LAN was to maintain feature parity between IPv4 and IPv6 clients including mobility, security, guest access, quality of service, and endpoint visibility.

Up to eight IPv6 client addresses can be tracked per device. This allows IPv6 clients to have a link-local, Stateless Address Auto Configuration (SLAAC) address, Dynamic Host Configuration Protocol for IPv6 (DHCPv6) address, and even addresses in alternative prefixes to be on a single interface. Work Group Bridge (WGB) clients connected to the uplink of an autonomous access point (AP) in WGB mode can also support IPv6.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on these software and hardware versions:

- Wireless LAN Controllers 2500 Series, 5500 Series, or WiSM2
- APs 1130, 1240, 1250, 1040, 1140, 1260, 3500, 3600 Series APs, and 1520 or 1550 Series Mesh APs
- IPv6-capable Router

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.


Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Prerequisites for Wireless IPv6 Client Connectivity

In order to enable wireless IPv6 client connectivity, the underlying wired network must support IPv6 routing and an address assignment mechanism such as SLAAC or DHCPv6. The wireless LAN controller must have L2 adjacency to the IPv6 router, and the VLAN needs to be tagged when the packets enter the controller. APs do not require connectivity on an IPv6 network, as all traffic is encapsulated inside the IPv4 CAPWAP tunnel between the AP and controller.

SLAAC Address Assignment




The most common method for IPv6 client address assignment is SLAAC. SLAAC provides simple plug-and-play connectivity where clients self-assign an address based on the IPv6 prefix. This process is achieved when the IPv6 router sends out periodic Router Advertisement messages which inform the client of the IPv6 prefix in use (the first 64 bits) and of the IPv6 default gateway. From that point, clients can generate the remaining 64 bits of their IPv6 address based on two algorithms: EUI-64 which is based on the MAC address of the interface, or private addresses which are randomly generated. The choice of algorithm is up to the client and is often configurable. Duplicate address detection is performed by IPv6 clients in order to ensure random addresses that are picked do not collide with other clients. The address of the router sending advertisements is used as the default gateway for the client.

These Cisco IOS® configuration commands from a Cisco-capable IPv6 router are used to enable SLAAC addressing and router advertisements:

```
ipv6 unicast-routing
interface Vlan20
description IPv6-SLAAC
ip address 192.168.20.1 255.255.255.0
ipv6 address 2001:DB8:0:20::1/64
ipv6 enable
end
```

DHCPv6 Address Assignment



The use of DHCPv6 is not required for IPv6 client connectivity if SLAAC is already deployed. There are two modes of operation for DHCPv6 called **Stateless** and **Stateful**.

The DHCPv6 **Stateless** mode is used to provide clients with additional network information not available in the router advertisement, but not an IPv6 address as this is already provided by SLAAC. This information can include the DNS domain name, DNS server(s), and other DHCP vendor-specific options. This interface configuration is for a Cisco IOS IPv6 router implementing stateless DHCPv6 with SLAAC enabled:

```
ipv6 unicast-routing
interface Vlan20
  description IPv6-DHCP-Stateless
  ip address 192.168.20.1 255.255.255.0
  ipv6 enable
  ipv6 address 2001:DB8:0:20::1/64
  ipv6 nd other-config-flag
  ipv6 dhcp relay destination 2001:DB8:0:10::100
end
```


The DHCPv6 **Stateful** option, also known as managed mode, operates similarly to DHCPv4 in that it assigns unique addresses to each client instead of the client generating the last 64 bits of the address as in SLAAC. This interface configuration is for a Cisco IOS IPv6 router implementing stateful DHCPv6 with SLAAC disabled:

```
ipv6 unicast-routing
interface Vlan20
  description IPv6-DHCP-Stateful
  ip address 192.168.20.1 255.255.255.0
  ipv6 enable
  ipv6 address 2001:DB8:0:20::1/64
  ipv6 nd prefix 2001:DB8:0:20::/64 no-advertise
  ipv6 nd managed-config-flag
  ipv6 nd other-config-flag
  ipv6 dhcp relay destination 2001:DB8:0:10::100
end
```

Additional Information

Configuring the wired network for complete IPv6 campus-wide connectivity using dual-stack or tunneling connectivity methods is out of the scope of this document. For more information refer to the Cisco validated deployment guide *Deploying IPv6 in Campus Networks*.

IPv6 Client Mobility




In order to deal with roaming IPv6 clients across controllers, the ICMPv6 messages such as Neighbor Solicitation (NS), Neighbor Advertisement (NA), Router Advertisement (RA), and Router Solicitation (RS) must be dealt with specially in order to ensure a client remains on the same Layer 3 network. The configuration for IPv6 mobility is the same as for IPv4 mobility and requires no separate software on the client side to achieve seamless roaming. The only required configuration is that the controllers must be part of the same mobility group/domain.

Here is the process for IPv6 client mobility across controllers:

1. If both controllers have access to the same VLAN the client was originally on, the roam is simply a Layer 2 roaming event where the client record is copied to the new controller and no traffic is tunneled back to the anchor controller.
2. If the second controller does not have access to the original VLAN the client was on, a Layer 3 roaming event will occur, meaning all traffic from the client must be tunneled via the mobility tunnel (Ethernet over IP) to the anchor controller.
 - a. In order to ensure the client retains its original IPv6 address, the RAs from the original VLAN are sent by the anchor controller to the foreign controller where they are delivered to the client using L2 unicast from the AP.
 - b. When the roamed client goes to renew its address via DHCPv6 or generate a new address via SLAAC, the RS, NA, and NS packets continue to be tunneled to the original VLAN so the client will receive an IPv6 address that is applicable to that VLAN.

Note: Mobility for IPv6-only clients is based on VLAN information. This means that IPv6-only client mobility is not supported on untagged VLANs.


Support for VLAN Select (Interface Groups)



The interface groups feature allows an organization to have a single WLAN with multiple VLANs configured on the controller in order to permit load balancing of wireless clients across these VLANs. This feature is commonly used to keep IPv4 subnet sizes small while enabling a WLAN to scale to thousands of users across multiple VLANs in the group. In order to support IPv6 clients with interface groups, no additional configuration is required as the system automatically sends the correct RA to the correct clients via L2 wireless unicast. By unicasting the RA, clients on the same WLAN, but a different VLAN, do not receive the incorrect RA.

First Hop Security for IPv6 Clients

Router Advertisement Guard



The RA Guard feature increases the security of the IPv6 network by dropping RAs coming from wireless clients. Without this feature, misconfigured or malicious IPv6 clients could announce themselves as a router for the network, often with a high priority which could take precedence over legitimate IPv6 routers.

By default, RA Guard is enabled at the AP (but can be disabled at the AP) and is always enabled on the controller. Dropping RAs at the AP is preferred as it is a more scalable solution and provides enhanced per-client RA drop counters. In all cases, the IPv6 RA will be dropped at some point, protecting other wireless clients and upstream wired network from malicious or misconfigured IPv6 clients.

DHCPv6 Server Guard

The DHCPv6 Server Guard feature prevents wireless clients from handing out IPv6 addresses to other wireless clients or wired clients upstream. In order to prevent DHCPv6 addresses from being handed out, any DHCPv6 advertise packets from wireless clients are dropped. This feature operates on the controller, requires no configuration and is enabled automatically.

IPv6 Source Guard

The IPv6 Source Guard feature prevents a wireless client spoofing an IPv6 address of another client. This feature is analogous to IPv4 Source Guard. IPv6 Source Guard is enabled by default but can be disabled via the CLI.

IPv6 Address Accounting

For RADIUS authentication and accounting, the controller sends back one IP address using the Framed-IP-address attribute. The IPv4 address is used in this case.

The Calling-Station-ID attribute uses this algorithm in order to send back an IP address when the Call Station ID Type on the controller is configured to IP Address :

1. IPv4 address
2. Global Unicast IPv6 Address
3. Link Local IPv6 Address

Since client IPv6 addresses can change often (temporary or private addresses), it is important to track them over time. Cisco NCS records all IPv6 addresses in use by each client and historically logs them each time the client roams or establishes a new session. These records can be configured at NCS to be held for up to a year.

Note: The default value for the Call Station ID Type on the controller has been changed to System MAC Address in version 7.2. When upgrading, this should be changed to allow unique tracking of clients by MAC address as IPv6 addresses may change mid-session and cause issues in accounting if the Calling-Station-ID

is set to IP address.

IPv6 Access Control Lists


In order to restrict access to certain upstream wired resources or block certain applications, IPv6 Access Control Lists (ACLs) can be used to identify traffic and permit or deny it. IPv6 ACLs support the same options as IPv4 ACLs including source, destination, source port, and destination port (port ranges are also supported). Pre-authentication ACLs are also supported to support IPv6 guest authentication using an external web server. The wireless controller supports up to 64 unique IPv6 ACLs with 64 unique rules in each. The wireless controller continues to support an additional 64 unique IPv4 ACLs with 64 unique rules in each for a total of 128 ACLs for a dual-stack client.

AAA Override for IPv6 ACLs

In order to support centralized access control through a centralized AAA server such as the Cisco Identity Services Engine (ISE) or ACS, the IPv6 ACL can be provisioned on a per-client basis using AAA Override attributes. In order to use this feature, the IPv6 ACL must be configured on the controller and the WLAN must be configured with the AAA Override feature enabled. The actual named AAA attribute for an IPv6 ACL is *Airespace-IPv6-ACL-Name* similar to the *Airespace-ACL-Name* attribute used for provisioning an IPv4-based ACL. The AAA attribute returned contents should be a string equal to the name of the IPv6 ACL as configured on the controller.

Packet Optimization for IPv6 Clients

Neighbor Discovery Caching



The IPv6 neighbor discovery protocol (NDP) utilizes NA and NS packets in place of Address Resolution Protocol (ARP) in order to allow IPv6 clients to resolve the MAC address of other clients on the network. The NDP process can be very chatty as it initially uses multicast addresses to perform address resolution; this can consume valuable wireless airtime as the multicast packets are sent to all clients on the network segment.

In order to increase the efficiency of the NDP process, neighbor discovery caching allows the controller to act as a proxy and respond back to NS queries that it can resolve. Neighbor discovery caching is made possible by the underlying neighbor binding table present in the controller. The neighbor binding table keeps track of each IPv6 address and its associated MAC address. When an IPv6 client attempts to resolve another client's link layer address, the NS packet is intercepted by the controller which responds back with a NA packet.

Router Advertisement Throttling

Router Advertisement Throttling allows the controller to enforce rate limiting of RAs headed towards the wireless network. By enabling RA throttling, routers which are configured to send RAs very often (for example, every three seconds) can be trimmed back to a minimum frequency that will still maintain IPv6 client connectivity. This allows airtime to be optimized by reducing the number of multicast packets that must be sent. In all cases, if a client sends an RS, then an RA will be allowed through the controller and unicast to the requesting client. This is to ensure that new clients or roaming clients are not negatively impacted by RA throttling.

IPv6 Guest Access


The wireless and wired guest features present for IPv4 clients work in the same manner for dual-stack and IPv6-only clients. Once the guest user associates, they are placed in a WEB_AUTH_REQ run state until the client is authenticated via the IPv4 or IPv6 captive portal. The controller will intercept both IPv4 and IPv6 HTTP/HTTPS traffic in this state and redirect it to the virtual IP address of the controller. Once the user is authenticated via the captive portal, their MAC address is moved to the run state and both IPv4 and IPv6 traffic is allowed to pass. For external web authentication, pre-authentication ACL allows an external web server to be used.

In order to support the redirection of IPv6-only clients, the controller automatically creates an IPv6 virtual address based off of the IPv4 virtual address configured on the controller. The virtual IPv6 address follows the convention of [::ffff:<virtual IPv4 address>]. For example, a virtual IP address of 1.1.1.1 would translate to [::ffff:1.1.1.1].

When using a trusted SSL certificate for guest access authentication, make sure that both the IPv4 and IPv6 virtual address of the controller is defined in DNS to match the SSL certificates hostname. This ensures that clients do not receive a security warning stating that the certificate does not match the hostname of the device.


Note: The controller's auto-generated SSL certificate does not contain the IPv6 virtual address. This may cause some web browsers to present a security warning. Using a trusted SSL certificate for guest access is recommended.

IPv6 VideoStream




VideoStream enables reliable and scalable wireless multicast video delivery, sending each client the stream in a unicast format. The actual multicast to unicast conversion (of L2) occurs at the AP providing a scalable solution. The controller sends the IPv6 video traffic inside an IPv4 CAPWAP multicast tunnel which allows efficient network distribution to the AP.

IPv6 Quality of Service



IPv6 packets use a similar marking to IPv4's use of DSCP values supporting up to 64 different traffic classes (0–63). For downstream packets from the wired network, the IPv6 Traffic Class value is copied to the header of the CAPWAP tunnel in order to ensure that QoS is preserved end-to-end. In the upstream direction, the same occurs as client traffic marked at Layer 3 with IPv6 traffic class will be honored by marking the CAPWAP packets destined for the controller.

IPv6 and FlexConnect



FlexConnect Local Switching WLANs

FlexConnect in local switching mode supports IPv6 clients by bridging the traffic to the local VLAN, similar to IPv4 operation. Client mobility is supported for Layer 2 roaming across the FlexConnect group.

These IPv6-specific features are supported in FlexConnect local switching mode:

- IPv6 RA Guard
- IPv6 Bridging
- IPv6 Guest Authentication (controller-hosted)

These IPv6-specific features are not supported in FlexConnect local switching mode:

- Layer 3 Mobility
- IPv6 VideoStream
- IPv6 Access Control Lists
- IPv6 Source Guard
- DHCPv6 Server Guard

- Neighbor Discovery Caching
- Router Advertisement Throttling

FlexConnect Central Switching WLANs

For APs in FlexConnect mode using central switching (tunneling traffic back to the controller), the controller must be set to **Multicast – Unicast Mode** for the **AP Multicast Mode**. Since FlexConnect APs do not join the CAPWAP multicast group of the controller, multicast packets must be replicated at the controller and unicast to each AP individually. This method is less efficient than **Multicast – Multicast Mode** and places additional load on the controller.

This IPv6-specific feature is not supported in FlexConnect central switching mode:

- IPv6 VideoStream

Note: Centrally switched WLANs running IPv6 are not supported on the Flex 7500 Series Controller.


IPv6 Clients Visibility with NCS

With the release of NCS v1.1, many additional IPv6 specific capabilities are added to monitor and manage a network of IPv6 clients on both wired and wireless networks.


IPv6 Dashboard Items

In order to view what types of clients are present on the network, a **Dashlet** in NCS is available in order to provide insight into IPv6 specific statistics and offer the capability to drill down into IPv6 clients.


IP Address Type Dashlet – Displays the types of IP clients on the network:



Client Count by IP Address Type – Displays the IP client type over time:




Client Traffic by IP Address Type – Displays the traffic from each type of client. Clients in the dual-stack category include both IPv4 and IPv6 traffic:



IPv6 Address Assignment – Displays the method of address assignment for each client as one of these four categories:

- DHCPv6 – For clients with addresses assigned by a central server. The client may also have a SLAAC address as well.
- SLAAC or Static – For clients using stateless address auto assignment or using statically configured addresses.
- Unknown – In some cases, the IPv6 address assignment cannot be discovered.
 - ◆ This condition only occurs on wired clients in NCS as some switches do not snoop IPv6 address assignment information.
- Self-Assigned – For clients with only a Link-local address which is entirely self-assigned.
 - ◆ Clients in this category can have IPv6 connectivity issues since they lack a Global Unique or Local Unique address.

Each of the sections of the pie chart is clickable, which allows the administrator to drill down to a list of clients.



Monitor IPv6 Clients

Cisco Prime Network Control System

Virtual Domain: ROOT-DOMAIN root Log Out

Home Monitor Configure Services Reports Administration

Clients and Users

Troubleshoot Test Disable Remove More Track Clients Identify Unknown Users

MAC Address	Vendor	IP Address	IP Type	Link Local	Router Advertisements Dropped
00:21:6a:a7:4f:ee	Intel	2001:db8:0:20:3057:534d:587d:73ae	IPv6	fe80::3057-534d:587d:73ae	0
00:21:6a:a7:54:88	Intel	192.168.20.21	Dual-Stack	fe80::5dda:a8e0:a969:fde6	0
00:24:d7:99:97:08	Intel	192.168.20.23	Dual-Stack	fe80::224:d7ff:fe99:9708	70
00:21:6a:5a:86:70	Intel	192.168.20.30	Dual-Stack	fe80::221:6aff:fe5a:8670	0
00:21:6a:67:31:48	Intel	192.168.20.25	Dual-Stack	fe80::acec:d514:2a14:ca7d	0
00:21:6a:a7:54:4e	Intel	192.168.20.22	Dual-Stack	fe80::1981:6f73:e618:32bd	0
fb:1e:df:e5:5b:03	Apple	192.168.20.29	Dual-Stack	fe80::fa1e:dfff:fee5:5b03	0
fb:1e:df:e3:0a:76	Apple	192.168.20.28	Dual-Stack	fe80::fa1e:dfff:fee3:a76	0
00:21:6a:a7:78:64	Intel	192.168.20.27	Dual-Stack	fe80::b5ba:eb3d:848d:ab6a	0

In order to monitor and manage IPv6 client information, these columns were added to the Clients and Users page:

- **IP Type** – The type of client based on what IP addresses have been seen from the client. The possible options are IPv4, IPv6, or Dual-Stack which signifies a client with both IPv4 and IPv6 addresses.
- **IPv6 Assignment Type** – The method of address assignment is detected by NCS as either SLAAC or Static, DHCPv6, Self-Assigned, or Unknown.
- **Global Unique** – The most recent IPv6 global address used by the client. A mouse-over on column contents reveals any additional IPv6 global unique addresses used by the client.
- **Local Unique** – The most recent IPv6 local unique address used by the client. A mouse over on column contents reveals any additional IPv6 global unique addresses used by the client.
- **Link Local** – The IPv6 address of the client which is self-assigned and used for communication before any other IPv6 address is assigned.
- **Router Advertisements Dropped** – The number of router advertisements sent by the client and dropped at the AP. This column can be used to track down clients that may be misconfigured or maliciously configured to act like an IPv6 router. This column is sortable, which allows offending clients to be identified easily.

Client IPv6 Addresses for: 00:21:6a:a7:54:4e


IP Address	Scope	Assignment	Discovery Time
2001:db8:0:25:1981:6f73:e618:32bd	Global Unique	NDP	2011-Oct-07, 18:47:58 UTC
2001:db8:0:25:4df2:542d:76b3:d9a6	Global Unique	NDP	2011-Oct-07, 18:47:58 UTC
2001:db8:0:25:6edc:f72b:39c:cd39	Global Unique	DHCP	2011-Oct-07, 18:47:58 UTC
2001:db8:0:25:9120:37c4:d14e:4cb6	Global Unique	NDP	2011-Oct-07, 18:47:58 UTC
fe80::1981:6f73:e618:32bd	Link Local	NDP	2011-Oct-07, 18:47:58 UTC

In addition to displaying IPv6 specific columns, the IP Address column will show the current IP address of the client with a priority to display the IPv4 address first (in the case of a Dual-Stack client) or the IPv6 Global Unique address in the case of an IPv6-only client.

Configuration for Wireless IPv6 Client Support

Multicast Distribution Mode to APs

The Cisco Unified Wireless Network supports two methods of multicast distribution to APs associated to the controller. In both modes, the original multicast packet from the wired network is encapsulated inside a Layer 3 CAPWAP packet sent via either CAPWAP Unicast or Multicast to the AP. Since the traffic is CAPWAP encapsulated, APs do not have to be on the same VLAN as the client traffic. The two methods of Multicast distribution are compared here:




	Multicast–Unicast Mode	Multicast–Multicast Mode
Delivery Mechanism	The controller replicates the multicast packet and sends it to each AP in a Unicast CAPWAP Tunnel	The controller sends one copy of the multicast packet
Supported AP Modes	FlexConnect and Local	Local Mode Only
Requires L3 Multicast Routing on Wired Network	No	Yes
Controller Loading	High	Low
Wired Network Loading	High	Low

Configure Multicast–Multicast Distribution Mode

Multicast–multicast mode is the recommended option for scalability and wired bandwidth efficiency reasons.

Note: This step is only absolutely required for the 2500 Series Wireless Controller, but it enables more efficient multicast transmission and is recommended for all controller platforms.

Go to the **Controller** tab under the **General** page and make sure the AP Multicast Mode is configured to use **Multicast** mode and that a valid group address is configured. The group address is an IPv4 multicast group and is recommended to be in the 239.X.X.X–239.255.255.255 range, which is scoped for private multicast applications.




Note: Do not use the 224.X.X.X, 239.0.0.X, or the 239.128.0.X address ranges for the multicast group address. Addresses in these ranges overlap with the link local MAC addresses and flood all switch ports, even with IGMP snooping enabled.

Configure Multicast–Unicast Distribution Mode

If the wired network is not properly configured to deliver the CAPWAP multicast between the controller and AP or FlexConnect mode, and APs will be used for centrally switched WLANs supporting IPv6, then unicast mode is required.

1. Go to the **Controller** tab under the General page, and make sure the AP Multicast Mode is configured to use **Unicast** mode.




2. Connect an IPv6 capable client to the wireless LAN. Validate that the client receives an IPv6 address by navigating to the **Monitor** tab and then the **Clients** menu.



Configure IPv6 Mobility

There is no specific configuration for IPv6 mobility except to place controllers in the same mobility group or within the same mobility domain. This allows up to 72 total controllers to participate in a mobility domain providing seamless mobility for even the largest of campuses.

Go to the **Controller** tab > **Mobility Groups**, and add each controller by MAC Address and IP address into the group. This must be done on all controllers in the mobility group.



Configure IPv6 Multicast

The controller supports MLDv1 snooping for IPv6 multicast which allows it to intelligently keep track of and deliver multicast flows to clients that request them.

Note: Unlike previous versions of releases, IPv6 unicast traffic support does not mandate that Global Multicast Mode be enabled on the controller. IPv6 unicast traffic support is enabled automatically.

1. Go to the **Controller** tab > **Multicast** page and **Enable MLD Snooping** in order to support multicast IPv6 traffic. In order for IPv6 Multicast to be enabled, the **Global Multicast Mode** of the controller must be enabled as well.

Controller

- General
- Inventory
- Interfaces
- Interface Groups
- Multicast**
- Network Routes
- Internal DHCP Server
- Mobility Management
- Ports

Multicast

- Enable Global Multicast Mode
- Enable IGMP Snooping
- IGMP Timeout (seconds) 60
- IGMP Query Interval (seconds) 20
- Enable MLD Snooping
- MLD Timeout (seconds) 60
- MLD Query Interval (seconds) 20

Note: Global Multicast Mode, IGMP, and MLD snooping should be enabled if peer-to-peer discovery applications such as Apple's Bonjour are required.

2. In order to verify that IPv6 multicast traffic is being snooped, go to the **Monitor** tab and the **Multicast** page. Notice that both IPv4 (IGMP) and IPv6 (MLD) multicast groups are listed. Click the MGID in order to view the wireless clients joined to that group address.

Monitor

- Summary
- Access Points
- Cisco CleanAir
- Statistics
- CDP
- Rogues
- Clients
- Multicast**

Multicast Groups

Layer3 MGID(Multicast Group ID) Mapping

Group address	Vlan	MGID	IGMP/MLD
224.0.0.251	20	1106	IGMP
224.0.0.252	20	1101	IGMP
239.255.255.250	20	1103	IGMP
ff02::c	20	1102	MLD
ff02::fb	20	1105	MLD
ff02::1:3	20	1100	MLD
ff02::2:fb5:a199	20	1110	MLD


Configure IPv6 RA Guard

Navigate to the **Controller** tab and then **IPv6 > RA Guard** on the left-hand menu. **Enable** IPv6 RA Guard on AP. RA Guard on the controller cannot be disabled. In addition to RA Guard configuration, this page also showcases any clients that have been identified as sending RAs.




Configure IPv6 Access Control Lists


1. Go to the **Security** tab, open **Access Control Lists**, and click **New**.




2. Enter a unique name for the ACL, change the ACL Type to **IPv6**, and click **Apply**.




3. Click on the new ACL that was created in the above steps.



- Click **Add New Rule**, enter the desired parameters for the rule, and click **Apply**. Leave the sequence number blank in order to place the rule at the end of the list. The **Direction** option of **Inbound** is used for traffic coming from the wireless network, and **Outbound** for traffic destined for wireless clients. Remember, the last rule in an ACL is an implicit deny—all. Use a prefix length of 64 in order to match an entire IPv6 subnet, and a prefix length of 128 for uniquely restricting access to an individual address.




- IPv6 ACLs are applied on a per-WLAN/SSID basis and can be used on multiple WLANs concurrently. Navigate to the **WLANs** tab and click the **WLAN ID** of the SSID in question in order to apply the IPv6 ACL. Click the **Advanced** tab and change the **Override Interface ACL** for IPv6 to the ACL name.




Configure IPv6 Guest Access for External Web Authentication

1. Configure the IPv4 and IPv6 pre-authentication ACL for the webserver. This allows traffic to and from the external server before the client is fully authenticated.




For more information on the operation of external web access, refer to External Web Authentication with Wireless LAN Controllers Configuration Example.

2. Configure the Guest WLAN by browsing to the WLANs tab at the top. Create the Guest SSID and use a Layer 3 web policy. The pre-authentication ACLs defined in Step 1 are selected for IPv4 and IPv6. Check the Over-ride Global Config section and select **External** from the Web Auth type drop-down box. Enter the URL of the web server. The host name of the external server should be resolvable in IPv4 and IPv6 DNS.



Configure IPv6 RA Throttling

1. Navigate to the **Controller** top-level menu and click the **IPv6 > RA Throttle Policy** option on the left-hand side. Enable RA Throttling by clicking the checkbox.



Note: When RA Throttling occurs, only the first IPv6 capable router is allowed through. For networks with multiple IPv6 prefixes being served by different routers, RA throttling should be disabled.

2. Adjust the throttle period and other options only under advisement from TAC. However, the default is recommended for most deployments. The various configuration options of the RA Throttling policy should be adjusted with this in mind:

- ◆ The numerical values of Allow At-least should be less than Allow At-most which should be less than Max Through .
- ◆ The RA throttle policy should not use a throttle period that is more than 1800 seconds as this is the default lifetime of most RAs.

Each RA Throttling option is described below:

- Throttle Period – The period of time that throttling takes place. RA throttling takes effect only after the Max Through limit is reached for the VLAN.
- Max Through – This is the maximum number of RAs per VLAN before throttling kicks in. The No Limit option allows an unlimited amount of RAs through with no throttling.
- Interval Option – The interval option allows the controller to act differently based on the RFC 3775 value set in the IPv6 RA.
 - ◆ Passthrough – This value allows any RAs with an RFC3775 interval option to go through without throttling.
 - ◆ Ignore – This value will cause the RA throttler to treat packets with the interval option as a regular RA and subject to throttling if in effect.
 - ◆ Throttle – This value will cause the RAs with the interval option to always be subject to rate limiting.
- Allow At-least – The minimum number of RAs per router that will be sent as multicast.
- Allow At-most – The maximum number of RAs per router that will be sent as multicast before throttling takes effect. The No Limit option will allow an unlimited number of RAs through for that router.

Configure the IPv6 Neighbor Binding Table

1. Go to the Controller top-level menu and click **IPv6 > Neighbor Binding Timers** on the left-hand menu.

The screenshot shows the Cisco Controller configuration interface. The top navigation bar includes 'MONITOR', 'WLANS', 'CONTROLLER', and 'WIRELESS'. The left-hand menu is expanded to 'IPv6', with 'Neighbor Binding Timers' selected. The main content area, titled 'Neighbor Binding Timers', contains three input fields: 'Down Lifetime (0-86400)' with a value of 30, 'Reachable Lifetime (0-86400)' with a value of 300, and 'Stale Lifetime (0-86400)' with a value of 86400. A red box highlights these three fields.

2. Adjust Down Lifetime, Reachable Lifetime, and Stale Lifetime as needed. For deployments with clients that are highly mobile, the timers for a stale address timer should be tweaked. Recommended values are:

- ◆ Down Lifetime – 30 seconds
- ◆ Reachable Lifetime – 300 seconds
- ◆ State Lifetime – 86400 seconds

Each lifetime timer refers to the state that an IPv6 address can be in:


- ◆ **Down Lifetime** – The down timer specifies how long IPv6 cache entries should be kept if the controller's uplink interface goes down.
- ◆ **Reachable Lifetime** – This timer specifies how long an IPv6 address will be marked active which means traffic has been received from this address recently. Once this timer expires, the address is moved to the Stale status.
- ◆ **Stale Lifetime** – This timer specifies how long to keep IPv6 addresses in the cache which have not been seen within the Reachable Lifetime. After this lifetime, the address is removed from the binding table.

Configure IPv6 VideoStream

1. Ensure Global VideoStream features are enabled on the Controller. Refer to Cisco Unified Wireless Network Solution: VideoStream Deployment Guide for information on enabling VideoStream on the 802.11a/g/n network as well as the WLAN SSID.
2. Go to the **Wireless** tab on the controller, and on the left-hand menu, choose **Media Stream > Streams**. Click **Add New** in order to create a new stream.



3. Name the stream and enter the start and end IPv6 addresses. When using only a single stream, the start and end addresses are equal. After adding the addresses, click **Apply** in order to create the stream.



Troubleshoot IPv6 Client Connectivity

Certain Clients are Unable to Pass IPv6 Traffic

Some client IPv6 networking stack implementations do not properly announce themselves when coming onto the network and therefore their address is not snooped appropriately by the controller for placement in the neighbor binding table. Any addresses not present in the neighbor binding table are blocked according to the IPv6 source guard feature. In order to permit these clients to pass traffic, these options need to be configured:

1. Disable the IPv6 Source Guard Feature through the CLI:


```
config network ip-mac-binding disable
```
2. Enable Multicast Neighbor Solicitation Forwarding through the CLI:

```
config ipv6 ns-mcast-fwd enable
```

Verify Successful Layer 3 Roaming for an IPv6 Client:

Issue these **debug** commands on both the anchor and foreign controller:

```
debug client <mac address>
```

debug mobility handoff enable

debug mobility packet enable

Debug Results on Anchor Controller:

```
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) State Update from Mobility-Complete to
  Mobility-Incomplete
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Setting handles to 0x00000000
00:21:6a:a7:4f:ee pemApfDeleteMobileStation2: APF_MS_PEM_WAIT_L2_AUTH_COMPLETE =
  0.
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Deleted mobile LWAPP rule on AP
  [04:fe:7f:49:03:30]
00:21:6a:a7:4f:ee Updated location for station old AP 04:fe:7f:49:03:30-1, new
  AP 00:00:00:00:00:00-0
00:21:6a:a7:4f:ee Stopping deletion of Mobile Station: (callerId: 42)
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) State Update from Mobility-Incomplete to
  Mobility-Complete, mobility role=Anchor, client state=APF_MS_STATE_ASSOCIATED
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Change state to RUN (20) last state RUN (20)
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Reached PLUMBFASPATH: from line 4968
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Adding Fast Path rule type = Airespace AP
  Client on AP 00:00:00:00:00:00, slot 0, interface = 13, QOS = 0
  IPv4 ACL ID = 255, IPv6 ACL ID = 255,
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Fast Path rule (contd...) 802.1P = 0, DSCP =
  0, TokenID = 7006 Local Bridging Vlan = 20, Local Bridging intf id = 13
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Successfully plumbed mobile rule (IPv4 ACL ID
  255, IPv6 ACL ID 255)
00:21:6a:a7:4f:ee 0.0.0.0 Removed NPU entry.
00:21:6a:a7:4f:ee Set symmetric mobility tunnel for 00:21:6a:a7:4f:ee as in
  Anchor role
00:21:6a:a7:4f:ee 0.0.0.0 Added NPU entry of type 1, dtlFlags 0x1
00:21:6a:a7:4f:ee Pushing IPv6: fe80:0000:0000:0000: 3057:534d:587d:73ae , and
  MAC: 00:21:6A:A7:4F:EE , Binding to Data Plane. SUCCESS !!
00:21:6a:a7:4f:ee Pushing IPv6: 2001:0db8:0000:0020: 3057:534d:587d:73ae , and
  MAC: 00:21:6A:A7:4F:EE , Binding to Data Plane. SUCCESS !!
00:21:6a:a7:4f:ee 0.0.0.0, VLAN Id 20 Not sending gratuitous ARP
00:21:6a:a7:4f:ee Copy AP LOCP - mode:0 slotId:0, apMac 0x0:0:0:0:0:0
00:21:6a:a7:4f:ee Copy WLAN LOCP EssIndex:3 aid:0 ssid: Roam
00:21:6a:a7:4f:ee Copy Security LOCP ecypher:0x0 ptype:0x2, p:0x0, eaptype:0x6
  w:0x1 aalg:0x0, PMState: RUN
00:21:6a:a7:4f:ee Copy 802.11 LOCP a:0x0 b:0x0 c:0x0 d:0x0 e:0x0 protocol2:0x5
  statusCode 0, reasoncode 99, status 3
00:21:6a:a7:4f:ee Copy CCX LOCP 4
00:21:6a:a7:4f:ee Copy e2e LOCP 0x1
00:21:6a:a7:4f:ee Copy MobilityData LOCP status:2, anchorip:0xac14e2c6
00:21:6a:a7:4f:ee Copy IPv6 LOCP: fe80::3057:534d:587d:73ae
```

Debug Results on Foreign Controller:

```
00:21:6a:a7:4f:ee Adding mobile on LWAPP AP f0:25:72:3c:0f:20(1)
00:21:6a:a7:4f:ee Reassociation received from mobile on AP f0:25:72:3c:0f:20
00:21:6a:a7:4f:ee 0.0.0.0 START (0) Changing IPv4 ACL 'none' (ACL ID 255) ==>
  'none' (ACL ID 255) --- (caller apf_policy.c:1697)
00:21:6a:a7:4f:ee 0.0.0.0 START (0) Changing IPv6 ACL 'none' (ACL ID 255) ==>
  'none' (ACL ID 255) --- (caller apf_policy.c:1864)
00:21:6a:a7:4f:ee Applying site-specific Local Bridging override for station
  00:21:6a:a7:4f:ee - vapId 3, site 'default-group', interface 'client-b1'
00:21:6a:a7:4f:ee Applying Local Bridging Interface Policy for station
  00:21:6a:a7:4f:ee - vlan 25, interface id 12, interface 'client-b1'
00:21:6a:a7:4f:ee processSsidIE statusCode is 0 and status is 0
00:21:6a:a7:4f:ee processSsidIE ssid_done_flag is 0 finish_flag is 0
00:21:6a:a7:4f:ee STA - rates (8): 140 18 152 36 176 72 96 108 0 0 0 0 0 0
*apfMsConnTask_4: Jan 22 20:37:45.370: 00:21:6a:a7:4f:ee suppRates statusCode
```

is 0 and gotSuppRatesElement is 1
00:21:6a:a7:4f:ee Processing RSN IE type 48, length 22 for mobile
00:21:6a:a7:4f:ee
00:21:6a:a7:4f:ee 0.0.0.0 START (0) Initializing policy
00:21:6a:a7:4f:ee 0.0.0.0 START (0) Change state to AUTHCHECK (2) last state
AUTHCHECK (2)
00:21:6a:a7:4f:ee 0.0.0.0 AUTHCHECK (2) Change state to 8021X_REQD (3) last
state 8021X_REQD (3)
00:21:6a:a7:4f:ee 0.0.0.0 8021X_REQD (3) DHCP Not required on AP
f0:25:72:3c:0f:20 vapId 3 apVapId 3for this client
00:21:6a:a7:4f:ee Not Using WMM Compliance code qosCap 00
00:21:6a:a7:4f:ee 0.0.0.0 8021X_REQD (3) Plumbed mobile LWAPP rule on AP
f0:25:72:3c:0f:20 vapId 3 apVapId 3
00:21:6a:a7:4f:ee apfMsAssoStateInc
00:21:6a:a7:4f:ee apfPemAddUser2 (apf_policy.c:268) Changing state for mobile
00:21:6a:a7:4f:ee on AP f0:25:72:3c:0f:20 from Idle to Associated
00:21:6a:a7:4f:ee Scheduling deletion of Mobile Station: (callerId: 49) in 1800
seconds
00:21:6a:a7:4f:ee Sending Assoc Response to station on BSSID f0:25:72:3c:0f:20
(status 0) ApVapId 3 Slot 1
00:21:6a:a7:4f:ee apfProcessAssocReq (apf_80211.c:6290) Changing state for
mobile 00:21:6a:a7:4f:ee on AP f0:25:72:3c:0f:20 from Associated to Associated
<&SNIP&>
00:21:6a:a7:4f:ee 0.0.0.0 8021X_REQD (3) Change state to L2AUTHCOMPLETE (4) last
state L2AUTHCOMPLETE (4)
00:21:6a:a7:4f:ee 0.0.0.0 L2AUTHCOMPLETE (4) DHCP Not required on AP
f0:25:72:3c:0f:20 vapId 3 apVapId 3for this client
00:21:6a:a7:4f:ee Not Using WMM Compliance code qosCap 00
00:21:6a:a7:4f:ee 0.0.0.0 L2AUTHCOMPLETE (4) Plumbed mobile LWAPP rule on AP
f0:25:72:3c:0f:20 vapId 3 apVapId 3
00:21:6a:a7:4f:ee 0.0.0.0 L2AUTHCOMPLETE (4) Change state to DHCP_REQD (7) last
state DHCP_REQD (7)
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) pemAdvanceState2 5253, Adding TMP rule
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Adding Fast Path rule
type = Airespace AP - Learn IP address
on AP f0:25:72:3c:0f:20, slot 1, interface = 13, QOS = 0
IPv4 ACL ID = 255, IP
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Fast Path rule (contd...) 802.1P = 0,
DSCP = 0, TokenID = 7006 Local Bridging Vlan = 25, Local Bridging intf id =
12
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Successfully plumbed mobile rule (IPv4
ACL ID 255, IPv6 ACL ID 255)
00:21:6a:a7:4f:ee Stopping retransmission timer for mobile 00:21:6a:a7:4f:ee
00:21:6a:a7:4f:ee 0.0.0.0 Added NPU entry of type 9, dtlFlags 0x0
00:21:6a:a7:4f:ee Sent an XID frame
00:21:6a:a7:4f:ee Username entry () already exists in name table, length = 253
00:21:6a:a7:4f:ee Username entry () created in mscb for mobile, length = 253
00:21:6a:a7:4f:ee Applying post-handoff policy for station 00:21:6a:a7:4f:ee -
valid mask 0x1000
00:21:6a:a7:4f:ee QOS Level: -1, DSCP: -1, dot1p: -1, Data Avg: -1, realtime
Avg: -1, Data Burst -1, Realtime Burst -1
00:21:6a:a7:4f:ee Session: -1, User session: -1, User elapsed -1 Interface:
N/A, IPv4 ACL: N/A, IPv6 ACL:
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Change state to DHCP_REQD (7) last state
DHCP_REQD (7)
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) pemCreateMobilityState 6370, Adding TMP
rule
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Replacing Fast Path rule type =
Airespace AP - Learn IP address on AP f0:25:72:3c:0f:20, slot 1, interface =
13, QOS = 0 IPv4 ACL ID = 255,
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Fast Path rule (contd...) 802.1P = 0,
DSCP = 0, TokenID = 7006 Local Bridging Vlan = 25, Local Bridging intf id =
12
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Successfully plumbed mobile rule (IPv4
ACL ID 255, IPv6 ACL ID 255)
00:21:6a:a7:4f:ee Scheduling deletion of Mobile Station: (callerId: 55) in 1800

```

seconds
00:21:6a:a7:4f:ee Pushing IPv6: fe80:0000:0000:0000: 3057:534d:587d:73ae , and
MAC: 00:21:6A:A7:4F:EE , Binding to Data Plane. SUCCESS !!
00:21:6a:a7:4f:ee apfMsRunStateInc
00:21:6a:a7:4f:ee 0.0.0.0 DHCP_REQD (7) Change state to RUN (20) last state RUN
(20)
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Reached PLUMBFASTPATH: from line 5776
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Change state to RUN (20) last state RUN (20)
00:21:6a:a7:4f:ee Pushing IPv6: 2001:0db8:0000:0020: 3057:534d:587d:73ae , and
MAC: 00:21:6A:A7:4F:EE , Binding to Data Plane. SUCCESS !!
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) State Update from Mobility-Incomplete to
Mobility-Complete, mobility role=Foreign, client state=APF_MS_STATE_ASSOCIATED
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Change state to RUN (20) last state RUN (20)
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Reached PLUMBFASTPATH: from line 4968
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Replacing Fast Path rule
type = Airespace AP Client
on AP f0:25:72:3c:0f:20, slot 1, interface = 13, QOS = 0
IPv4 ACL ID = 255, IPv6 ACL ID = 25
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Fast Path rule (contd...) 802.1P = 0, DSCP =
0, TokenID = 7006 Local Bridging Vlan = 25, Local Bridging intf id = 12
00:21:6a:a7:4f:ee 0.0.0.0 RUN (20) Successfully plumbed mobile rule (IPv4 ACL ID
255, IPv6 ACL ID 255)
00:21:6a:a7:4f:ee 0.0.0.0 Added NPU entry of type 9, dtlFlags 0x0
00:21:6a:a7:4f:ee Set symmetric mobility tunnel for 00:21:6a:a7:4f:ee as in
Foreign role
00:21:6a:a7:4f:ee 0.0.0.0 Added NPU entry of type 1, dtlFlags 0x1
00:21:6a:a7:4f:ee Pushing IPv6: fe80:0000:0000:0000: 3057:534d:587d:73ae , and
MAC: 00:21:6A:A7:4F:EE , Binding to Data Plane. SUCCESS !!
00:21:6a:a7:4f:ee Pushing IPv6: 2001:0db8:0000:0020: 3057:534d:587d:73ae , and
MAC: 00:21:6A:A7:4F:EE , Binding to Data Plane. SUCCESS !!
00:21:6a:a7:4f:ee Copy AP LOCP - mode:0 slotId:1, apMac 0xf0:25:72:3c:f:20
00:21:6a:a7:4f:ee Copy WLAN LOCP EssIndex:3 aid:1 ssid: Roam
00:21:6a:a7:4f:ee Copy Security LOCP ecypher:0x0 ptype:0x2, p:0x0, eaptype:0x6
w:0x1 aalg:0x0, PMState: RUN
00:21:6a:a7:4f:ee Copy 802.11 LOCP a:0x0 b:0x0 c:0x0 d:0x0 e:0x0 protocol2:0x7
statuscode 0, reasoncode 99, status 3
00:21:6a:a7:4f:ee Copy CCX LOCP 4
00:21:6a:a7:4f:ee Copy e2e LOCP 0x1
00:21:6a:a7:4f:ee Copy MobilityData LOCP status:3, anchorip:0xac14e2c5
00:21:6a:a7:4f:ee Copy IPv6 LOCP: fe80::3057:534d:587d:73ae
00:21:6a:a7:4f:ee Copy IPv6 LOCP: 2001:db8:0:20:3057:534d:587d:73ae

```

Useful IPv6 CLI Commands:

```
Show ipv6 neighbor-binding summary
```

```
Debug ipv6 neighbor-binding filter client <Client MAC> enable
```

```
Debug ipv6 neighbor-binding filter errors enable
```

Frequently Asked Questions

Q: What is the optimum IPv6 prefix size in order to limit the broadcast domain?

A: Although an IPv6 subnet can be sub-divided below a /64, this configuration will break SLAAC and cause issues with client connectivity. If segmentation is needed in order to reduce the number of hosts, the Interface Groups feature can be used to load balance clients among different back-end VLANs, each using a different IPv6 prefix.

Q: Are there any scalability limitations when it comes to supporting IPv6 clients?

A: The major scalability limitation for IPv6 client support is the neighbor binding table which keeps track of all wireless client IPv6 addresses. This table is scaled per controller platform in order to support the maximum number of clients multiplied by eight (the maximum number of addresses per client). The addition of the IPv6 binding table can raise the memory usage of the controller up roughly 10–15% under full load, depending on the platform.

Wireless Controller	Maximum Number of Clients	IPv6 Neighbor Binding Table Size
2500	500	4,000
5500	7,000	56,000
WiSM2	15,000	120,000

Q: What is the impact of IPv6 features on the controller's CPU and memory?

A: The impact is minimal as the CPU has multiple cores for processing the control plane. When tested with maximum supported clients, each with 8 IPv6 addresses, the CPU usage was below 30%, and memory usage was below 75%.

Q: Can IPv6 client support be disabled?

A: For customers who want to enable only IPv4 in their network and block IPv6, an IPv6 ACL of deny–all traffic can be used and applied on a per–WLAN basis.

Q: Is it possible to have one WLAN for IPv4 and another for IPv6?

A: It is not possible to have the same SSID name and security type for two different WLANs operating on the same AP. For segmentation of IPv4 clients from IPv6 clients, two WLANs must be created. Each WLAN must be configured with an ACL that blocks all IPv4 or IPv6 traffic respectively.

Q: Why is it important to support multiple IPv6 addresses per client?

A: Clients can have multiple IPv6 addresses per interface which can be static, SLAAC or DHCPv6 assigned in addition to always having a self–assigned Link–Local address. Clients can also have additional addresses using different IPv6 prefixes.

Q: What are IPv6 private addresses and why are they important to track?

A: Private (also known as temporary) addresses are randomly generated by the client when SLAAC address assignment is in use. These addresses are often rotated at a frequency of a day or so, as to prevent host traceability that would come from using the same host postfix (last 64 bits) at all times. It is important to track these private addresses for auditing purposes such as tracing copyright infringement. Cisco NCS records all IPv6 addresses in use by each client and historically logs them each time the client roams or establishes a new session. These records can be configured at NCS to be held for up to a year.

Related Information

- [Technical Support & Documentation – Cisco Systems](#)

