3

IPv6 Addressing

Contents

Introduction

IPv6 supports multiple addresses on an interface, and uses them in a manner comparable to subnetting an IPv4 VLAN. For example, where the switch is configured with multiple VLANs and each is connected to an IPv6 router, each VLAN will have a single link-local address and one or more global unicast addresses. This section describes IPv6 addressing and outlines the options for configuring IPv6 addressing on the switch. The configuration process includes automatically or statically creating an IPv6 address and automatically verifying the uniqueness of each.

IPv6 Address Structure and Format

Address Format

An IPv6 address is composed of 128 bits divided into eight 2-byte fields of hexadecimal values. The full format is:

XXXX : XXXX : XXXX : XXXX : XXXX : XXXX : XXXX : XXXX

where each field delimited by a colon (:) is a set of four hexadecimal digits.

For example:

2001:0db8:0000:00A9:0215:60ff:fe7a:adc0

2001:0db8:0260:0212:0000:0000:0000:01b4

The hexadecimal characters in IPv6 addresses are not case-sensitive.

Address Notation

Leading zeros in each field can be omitted as long as each field is represented by at least one value. The exception to this rule is when there is an uninterrupted series of zeros in one or more contiguous fields. In this case, the series of zeros can be replaced by "::", with the restriction that "::" can be used only once in a given address. Applying this convention to the above examples results in the following address notations:

2001:db8::a9:215:60ff:fe7a:adc0

2001:db8:260:0212::01b4

An IPv6 address includes a network prefix and an interface identifier.

Network Prefix

The network prefix (high-order bits) in an IPv6 address begins with a well-known, fixed prefix for defining the address type. Some examples of well-known, fixed prefixes are:

2000::/3global (routable) unicast address

fd08::/8 unique local unicast address

fe80::/8link-local unicast address

ff00::/8multicast address

The remainder of the network prefix depends on the prefix type, and includes information such as the subnet destination of unicast addresses or the flags and scope of multicast addresses.

In a given address, CIDR-type notation (Classless Inter-Domain Routing) is used to define the network prefix. In the following address example, the 64 bits comprising 2001:0db8:0260:0201 form the network prefix:

2001:0db8:0260:0212:0215:60ff:fe7a:adc0/64

A shorter way to show this address is to remove the leading zeros:

2001:db8:260:212:215:60ff:fe7a:adc0/64

Interface (Device) Identifier

The remaining (low-order) bits in the address comprise a unique interface identifier in an IPv6 address. In the above example, the rightmost 64 bits (215:60ff:fe7a:adc0) comprise the interface identifier. Unlike IPv4, an IPv6 identifier for a unicast or anycast address can be automatically generated from the switch's MAC address using EUI-64 (Extended Unique Identifier) format. Other methods include DHCPv6 assignments and static configuration. Interface identifiers are covered in more detail in the later sections of this chapter describing different address types.

IPv6 Addressing Options

IPv6 Address Sources

IPv6 addressing sources provide a flexible methodology for assigning addresses to VLAN interfaces on the switch. Options include:

- stateless IPv6 autoconfiguration on VLAN interfaces includes:
  - link-local unicast addresses
  - global unicast addresses
- stateful, global unicast IPv6 address configuration using DHCPv6
- static IPv6 address configuration

You can combine stateless, stateful, and static IP addressing methods on the switch as needed, according to the needs in your network. For example, if your network includes only one VLAN, you may need only stateless autoconfiguration of link-local addresses, although you could also use the static IPv6 method. (DHCPv6 does not configure link-local addresses.) Where routed traffic is used, you will also need global unicast addressing, either through stateless autoconfiguration or the other listed methods.

General IPv6 Address Types

IPv6 supports stateless and stateful address autoconfiguration, as well as static address configuration.This enables IPv6 to automatically address a device so that it can be placed in a network with or without static or DHCPv6 addressing intervention. All three of these methods can be used exclusively or in conjunction with each other, and a given IPv6 device can have multiple addresses assigned to the same interface in a manner similar to subnetting in IPv4.

Stateless Address Autoconfiguration . This method does not require the use of servers. Instead, in the default operation, the host uses its MAC address to automatically generate a link-local IPv6 address using the EUI-64 method to generate the device identifier. (Refer to “Autoconfiguring Link-Local Unicast Addresses” on page 3-13.) The scope of the link-local address enables communication with other IPv6 devices on the same VLAN. If an IPv6 router is present, an IPv6 address supporting routing is automatically generated, as well. (The switch merges a router-generated prefix received in router advertisements with the last 64 bits of thelink-local address on an interface to create the global address.) Refer to page 3-7.

Stateful Address Autoconfiguration. This method allows use of a DHCPv6 server to automatically configure IPv6 addressing on a host in a manner similar to stateful IP addressing with a DHCPv4 server. For software release K.13.01, a DHCPv6 server can provide routable IPv6 addressing and NTP (timep) server addresses. Also, if the host acquires its IPv6 addressing through stateless or static methods, the DHCPv6 server can still be used to automatically provide other configuration information to the host. Refer to page 3-8.

Static Address Configuration. Static configuration is used instead of or in addition to stateless and stateful autoconfiguration where use of the host MAC address does not provide the desired level of address control and distribution. Refer to page 3-9.

Duplicate Address Detection (DAD). IPv6 verifies both the link-local and the global unicast address(es) on each interface for uniqueness, regardless of the method used to configure the address. If an address fails this test, it is identified as a **duplicate**, and a replacement must be configured using the static method. (To view address status, use the **show ipv6** command.) For more information on DAD, refer to "Neighbor Discovery (ND)" on page 4-17.

Developing an Addressing Plan. For small, flat networks and any environment where control of address assignments need not be restricted or tightly controlled, stateless addressing is adequate for network management and control. Where systematic and controlled addressing is needed, stateful and static addressing methods should be used. Where dual-stack operation is used in a VLAN, incorporating the local IPv4 addressing scheme into the IPv6 addresses you use can help to provide consistency and correspondence among the IPv6 and IPv4 addresses in use on the VLAN.

Related Information.

- RFC 4291: "IP Version 6 Addressing Architecture"
- RFC 2462: "IPv6 Stateless Address Autoconfiguration"
- RFC 3315: "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)"

IPv6 Address Sources

IPv6 addressing sources provide a flexible methodology for assigning addresses to VLAN interfaces on the switch. Options include:

- stateless IPv6 autoconfiguration on VLAN interfaces includes:
  - link-local unicast addresses
  - global unicast addresses
- stateful IPv6 address configuration using DHCPv6
- static IPv6 address configuration

You can combine stateless, stateful, and static IP addressing methods on the switch as needed, according to the needs in your network. For example, if your network includes only one VLAN, you may need only stateless autoconfiguration of link-local addresses, although you could also use the static IPv6 method. (DHCPv6 does not configure link-local addresses.) Where routed traffic is used, you will also need global unicast addressing, either through stateless autoconfiguration or the other listed methods.

Stateless Address Autoconfiguration (SLAAC)

On the switches covered by this guide, stateless address autoconfiguration (SLAAC) generates link-local unicast and global unicast IPv6 addresses on a VLAN interface. In all cases, the prefix is 64 bits.

Applications

Stateless autoconfiguration is suitable where a link-local or global unicast IPv6 address (if a router is present) must be unique, but the actual address used is not significant. Where a specific unicast address or a unicast address from a specific range of choices is needed on an interface, DHCPv6 or static IPv6 address configuration should be used. (Refer to pages 3-8 and 3-9.)

Preferred and Valid Lifetimes of Stateless Autoconfigured Addresses

The preferred and valid lifetimes of an autoconfigured global unicast address are set by the router advertisements (RA) used to generate the address, and are the autoconfiguration counterpart to the lease time assigned by DHCPv6

servers. These lifetimes cannot be reset using control from the switch console or SNMP methods. Refer to "Preferred and Valid Address Lifetimes" on page 3-25.

Stateful (DHCPv6) Address Configuration

Stateful addresses are defined by a system administrator or other authority, and automatically assigned to the switch and other devices through the Dynamic Host Configuration Protocol (DHCPv6). Generally, DHCPv6 should be applied when you want specific, non-default addressing to be assigned automatically. For IPv6, DHCP use is indicated for conditions such as the following:

- address conventions used in your network require defined control
- static addressing is not feasible due to the number of nodes in the network
- automatic assignment of multiple IPv6 addresses per interfaces is needed
- automatic configuration of IPv6 access to DNS, SNTP, or TimeP servers

To implement stateful address configuration:

- The DHCPv6 server must be configured and accessible to the switch, either on the same VLAN or through an IPv6 router configured with DHCP Relay to support service requests from the switch.

Note

DHCPv6 relay may not currently be available in some IPv6 routers.

- DHCPv6 addressing must be enabled per-VLAN on the switch.

Note that IPv6 router advertisements (RAs) can also include instructions to clients to use DHCPv6 resources. Refer to the documentation for your IPv6 router.

If you want to use DHCPv6 in a dual-stack environment, you will need both DHCPv4 and DHCPv6 server access. Also, further developments in DHCP services are likely to mean new capabilities affecting DHCPv6 deployments.

For related information, refer to:

- RFC 3315: "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)"
- RFC 3041: "Privacy Extensions for Stateless Address Autoconfiguration in IPv6"

Static Address Configuration

Generally, static address configuration should be used when you want specific, non-default addressing to be assigned to a VLAN interface. For IPv6, DHCP use is indicated for conditions such as the following:

- address conventions used in your network require defined control
- the task of static addressing is not so extensive as to be impractical due to the number of addresses and/or interfaces needing configuration

If IPv6 is not already enabled on a VLAN interface, the following is true:

- Statically configuring a link-local address on the interface also enables IPv6.
- Statically configuring a global unicast or anycast address also enables IPv6 and generates a link-local address.

Statically configured global unicast addresses can be used in addition to stateless addresses on the same interface. However, because only one link-local address is allowed on a VLAN interface (fe80::), static configuration of a link-local address automatically replaces an existing link-local address.

Note

For a statically configured global unicast address to be routable, a gateway router must be transmitting router advertisements on the VLAN that include the prefix used in the statically configured address. If the VLAN is not receiving an RA with this prefix, the address is listed as “preferred”, but is not used.

Statically configured IPv6 addresses saved to the startup-config file (by using **write memory**) remain across a reboot and are permanent, unless statically removed by **no ipv6 address < *ipv6-addr* >**.

For more information and the CLI command for static address configuration, refer to “Configuring a Static IPv6 Address on a VLAN” on page 4-11.

Address Types and Scope

Address Types

IPv6 uses these IP address types:

- **Unicast:** Identifies a specific IPv6 interface. Traffic having a unicast destination address is intended for a single interface. Like IPv4 addresses, unicast addresses can be assigned to a specific VLAN on the switch and to other IPv6 devices connected to the switch. At a minimum, a given interface must have at least a link-local address. To send or receive traffic off of a VLAN, an interface must also have one or more global unicast addresses.
- **Multicast:** Provides a single destination address for traffic intended for all members of a group, and provides a means for reducing unnecessary traffic to interfaces that do not belong to a given multicast group. Membership in a group can be determined by request or by a characteristic, such as all nodes, all routers, or all routers of a given type. Multicast traffic can be generated by a single source or multiple sources, but in either case is intended for multiple destinations.Common types of multicast traffic include streaming video and audio to multiple receivers who have joined a specific group from diverse locations.

Note

Unlike IPv4, broadcast addresses are not used in IPv6. Multicast addresses are used instead. For more on this topic, refer to “Multicast Application to IPv6 Addressing” on page 3-21.

- **Anycast:** A single address of this type can be assigned to multiple interfaces, possibly on separate devices within a defined address scope, where any of the interfaces having the anycast address can provide the desired service or response. A packet sent to a given anycast address is delivered only to the nearest interface having an instance of the address. This option is useful where multiple servers provide the same service, and it does not matter to the client which source it uses to acquire the service. Anycast usage can be of value, for example, in a network supporting multiple DNS servers. Refer to “Anycast Addresses” on page 3-20.

A given interface can have only one link-local address, but can have multiple unicast and anycast addresses.

Address Scope

The address scope determines the area (topology) in which a given IPv6 address is used. This section provides an overview of IPv6 address types. For more information, refer to the chapter titled “IPv6 Addressing”.

Link-Local Address. Limited to a given interface (VLAN). Enabling IPv6 on a given VLAN automatically generates a link-local address used for switched traffic on the VLAN.

Global Unicast Address. Applies to a unique IPv6 routable address on the internet. A unique global address has a routing prefix and a unique device identifier.When autoconfiguration is enabled on a VLAN receiving an IPv6 router advertisement (RA), the prefix specified in the RA and the device identifier specified in the link-local address are combined to create a unique, global unicast address. A global unicast address can also be statically configured to either replace or complement an automatically configured address of the same type.

Unique Local Unicast. Applies to a routable, globally unique address intended for use within an entity defined by the system adminstrator, such as a specific site or a group of related sites defined by IPv6 border routers. These addresses are intended to be routable on a local site or an organization's intranet, but are not intended to be routed on the global internet. A unique local unicast address has the same format as a global unicast address. In this guide, unless otherwise stated, information on global unicast addresses also applies to unique local unicast addresses. For more on this topic, refer to “Unique Local Unicast IPv6 Address” on page 3-19.

Unicast Address Prefixes

Traffic having a unicast destination address is intended for a single interface identified by that address. While IPv6 unicast addresses can have prefixes of varying length, a 64-bit prefix is generally adequate.

Link-Local Unicast Prefix (fe80): This well-known 64-bit fixed prefix is for a non- routable address used to identify a device on a single VLAN interface, and requires the high-order ten bits to be set to fe80 (fe80::/10). The remaining 54 bits in the prefix are set to zeros, followed by an interface ID of 64 bits.

fe80:0000:0000:0000:0215:60ff:fe7a:adc0/64

or

fe80::215:60ff:fe7a:asc0/64

In binary notation, the fixed prefix for link-local prefixes is:

1111 1110 10 = fe80/10

For more on link-local addresses, refer to “Link-Local Unicast Address” on page 3-13.

Routable Global Unicast Prefix. This well-known 3-bit fixed-prefix indicates a routable address used to identify a device on a VLAN interface that is accessible by routing from multiple networks. The complete prefix is 64 bits, followed by a 64-bit interface identifier. For example, the leading 2 in the first octet of the following address illustrates a global unicast address:

2001:db8:260:212:215:60ff:fe7a:adc0/64

In binary notation, the fixed prefix in this example appears as follows:

0010 0000 = 20/3

Unique Local Unicast Prefix (fd). This well-known fixed prefix is defined as FC00/7. However, the eighth high-order bit must also be set to 1, resulting in a fixed prefix of fd00/8. (In the future, setting the eighth high-order bit to zero may become an option.) This prefix signifies a routable address intended for use within the boundaries of a site or organization. For example, the leading fd in the first octet of this address illustrates a unique local unicast address intended to be used in a privately defined network.

fd00:00ff:0C00:000a:215:60ff:fe7a:adc0

Unique local unicast addresses are described in more detail under "Unique Local Unicast IPv6 Address" on page 3-19.

Multicast Prefix (ff). This well-known 8-bit fixed prefix signifies a permanent or temporary multicast address. The second 8 high-order bits are used for flags and scope for the multicast address. The remaining 112 bits define the multicast group identifier. For example:

ff02::1:ffc7:b5b9

For more information, refer to “Multicast Application to IPv6 Addressing” on page 3-21.

Other Prefix Types. There are other designated global unicast prefixes such as those for the following address types:

- RFC 4380: “Teredo: Tunneling IPv6 over UDP”
- RFC 3056: “Connection of IPv6 Domains via IPv4 Clouds”
- RFC 4214: “Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)”

For related information, refer also to:

- RFC 4291: "IP Version 6 Addressing Architecture

Link-Local Unicast Address

A link-local unicast address is a non-routable address for use on a single VLAN interface, and provides basic connectivity to an IPv6 network. Because the scope of a link-local address is restricted to the VLAN on which the address is used, a link-local address must be unique only for the VLAN on which it is configured. (Traffic with a link-local source or destination address cannot be routed between VLANs.)

Autoconfiguring Link-Local Unicast Addresses

Enabling IPv6 on a given VLAN automatically generates a link-local address. This address is limited in scope to that VLAN, and is usable only for switched traffic. This address has a well- known, 64-bit prefix of fe80:0000:0000:0000 (hexadecimal), or fe80::, and a 64-bit device identifier derived from the VLAN's MAC address using the Extended Unique Identifier format (EUI-64, page 3-14). For example, if the MAC address of VLAN 10 is 021560-7aadc0, the automatically generated link-local address for VLAN 10 is:

fe80:0000:0000:0000:0215:60ff:fe7a:adc0

or, in standard IPv6 notation,

fe80::215:60ff:fe7a:adc0

Note that only one link-local address is allowed on an interface. Thus, on a given interface, statically configuring a link-local address type replaces the existing link-local address.

Because all VLANs configured on the switch use the same MAC address, all automatically generated link-local addresses on the switch will have the same link-local address. However, since the scope of a link-local address includes only the VLAN on which it was generated, this should not be a problem.

For example, executing **ipv6 address dhcp full** on a VLAN for which IPv6 was not previously configured does all of the following:

- enables IPv6 on the VLAN
- causes the switch to generate a stateless link-local unicast address on the VLAN
- configures the VLAN to send DHCPv6 requests

Note

Only one link-local unicast address can exist on a VLAN interface at any time. Configuring a new address of this type on an interface on which IPv6 is already enabled replaces the previously existing link-local address with the new one.

Any link-local address must include the well-known link-local prefix fe80::/64 plus a 64-bit device identifier.

Any of the following commands enable IPv6 on a VLAN and automatically generate a link-local address:

- **ipv6 enable** (page 4-6)
- **ipv6 address autoconfig** (page 4-7)
- **ipv6 address dhcp full [rapid-commit]** (page 4-9)
- **ipv6 address < *network-prefix*><*device-id* >/< *prefix-length* >** (page 4-13)

Extended Unique Identifier (EUI)

When the link-local address is automatically generated, the device identifier is derived from the switch's 48- bit (hexadecimal) MAC address to create a 64-bit Extended Unique Identifier (EUI) to be appended to the fe80 link-local prefix, as follows:

- ff-fe is inserted between third and fourth bytes of MAC address
- The second low-order bit (the Universal/Local bit) in the first byte of the MAC address is complemented, which usually means the bit is originally set to 0 and is changed to 1. This indicates a globally unique IPv6 interface identifier. For example:

MAC Address	IPv6 I/F Identifier	Full Link-Local Unicast Address
00-15-60-7a-ad-c0	215:60ff:fe7a:adc0	fe80::215:60ff:fe7a:adc0/64
09-c1-8a-44-b4-9d	11c1:8aff:fe44:b49d	fe80::11c1:8aff:fe44:b49d/64
00-1a-73-5a-7e-57	21a:73ff:fe5a:7e57	fe80::21a:73ff:fe5a:7e57/64

The EUI method of generating a link-local address is automatically implemented on the switches covered by this guide when IPv6 is enabled on a VLAN interface.

If automatically generated link-local addresses are not suitable for the addressing scheme you want to use, statically assigned link-local addresses can be used instead. (Refer to “Static Address Configuration” on page 3-9.)

For related information, refer to:

- RFC 2373: “IP Version 6 Addressing Architecture”
- RFC 2464: “Transmission of IPv6 Packets Over Ethernet Networks”

Note

While only one link-local IPv6 address is allowed on an interface, multiples of other address types can exist on the same interface. Thus, an interface can have one link-local unicast address, but multiple global unicast, anycast, and unique local addresses.

Statically Configuring Link-Local Addresses

A link-local unicast address can be configured statically on a VLAN interface. If IPv6 is not already enabled on the VLAN, this action also enables IPv6 on the VLAN. Only one link-local address can exist on a VLAN at any time. If a link-local address (static or autoconfigured) already exists on the VLAN, then statically configuring a new one replaces the previously existing one. To statically configure a link-local address, refer to “Statically Configuring a Link-Local Unicast Address ” on page 4-12.

Global Unicast Address

A global unicast address is required for unicast traffic to be routed across VLANs within an organization as well as across the public internet. To support subnetting, a VLAN can be configured with multiple global unicast addresses. Any of the following methods can be used to configure this kind of address on a VLAN:

- stateless address autoconfiguration using a prefix received in an advertisement received from a router on the VLAN (page 3-7)
- stateful address configuration using DHCPv6 (page 3-8)
- static address configuration (page 3-9)

Stateless Autoconfiguration of a Global Unicast Address

If there is an IPv6-enabled router transmitting router advertisements on a VLAN interface, enabling this method generates a global, routable unicast address for the VLAN. The prefix for this address type is typically 64 bits with the three highest-order bits set to 2.

Router Advertisements. With autoconfiguration enabled, if the switch receives the same prefix from router advertisements (RAs) from multiple IPv6 routers on the same VLAN, then one global unicast address is configured with that prefix. If different prefixes are received from different routers on the same VLAN, then there will be one address configured on the VLAN for each unique prefix received. Where there are multiple routers on the VLAN, the default route for the VLAN is determined by the relative router priorities included in the RAs the VLAN receives. If the highest priority is duplicated on multiple routers, then the first RA detected on the VLAN determines the default route.

If the RA used to define the prefix for an autoconfigured address ceases to be received on the VLAN, then the address becomes deprecated. (Refer to “IPv6 Address Deprecation” on page 3-25.)

If IPv6 is not already enabled on a VLAN when you enable autoconfiguration on the VLAN, then the switch automatically generates a link-local address for the VLAN as well.

If IPv6 Is Not Already Enabled. Enabling address autoconfiguration on a VLAN when IPv6 is not already enabled on the VLAN causes the switch to:

- generate a link-local address on the VLAN as described in the preceeding section (page 3-13).
- transmit a router solicitation on the VLAN, and to listen for advertisements from any IPv6 routers on the VLAN.

For each unique router advertisement (RA) the switch receives from any router(s), the switch configures a unique, global unicast address. This address type is composed of a 64-bit network prefix specified by the router advertisement, plus a device identifier generated in the same way as described in the preceeding section for link-local addresses (using the EUI algorithm). For example, suppose the following is true:

- IPv6 is not enabled on VLAN 1.
- The MAC address for VLAN 1 is 00-15-60-7a-ad-c0.
- A router on the same VLAN transmits router advertisements that assign the prefix 2001:0:260:212/64, plus a 64-bit interface identifier generated using the EUI format.

In this case, enabling IPv6 address autoconfiguration on VLAN 1 generates the following address assignments on VLAN 1:

- link-local unicast: fe80::215:60ff:fe7a:adc0/64
- global unicast:2001:0:260:212:215:60ff:fe7a:adc0/64

IPv6 Already Enabled. Enabling address autoconfiguration on a VLAN when IPv6 is already enabled on the VLAN creates a global unicast address in the same way as described above, except that the device identifier applied to the new global address is a duplicate of the 64-bit identifier in the current link-local address.

Note

After a global unicast address has been configured, its device identifier will not be changed by any later changes to the link-local address.

Static Configuration of a Global Unicast Address

A global unicast address can be configured statically on a VLAN interface. If IPv6 is not already enabled on a VLAN, then statically configuring a global unicast address automatically generates a link-local unicast address on the VLAN, as described in the pdreceeding section. To statically configure a global unicast address, refer to “Statically Configuring A Global Unicast Address” on page 4-13.

Prefixes in Routable IPv6 Addresses

In routable IPv6 addresses, the prefix uniquely identifies an entity and a unicast subnet within that entity, and is defined by a length value specifying the number of leftmost contiguous (high-order) bits comprising the prefix. For an automatically generated global unicast address, the default prefix length is 64 bits. (Pratically speaking, the entire prefix in a /64 address defines the subnet.) Prefixes configured through stateful or static methods can be any length compatible with the local network application.

In the following example, the leftmost 64 bits of the address comprise the prefix:

2001:0db8:0000:0212:0215:60ff:fe7a:adc0/64

or

2001:db8::212:215:60ff:fe7a:adc0/64

In this case, the prefix is read as:

2001:0db8:0000:0212::

or

2001:db8::212::

All bits to the right of 0212 comprise the device identifier in the unicast address.

For related information, refer to:

- RFC 3177: “IAB/IESG Recommendations on IPv6 Address Allocations to Sites”
- RFC 4291: “IP Version 6 Addressing Architecture”

Unique Local Unicast IPv6 Address

A unique local unicast address is an address that falls within a specific range, but is used only as a global unicast address within an organization. Traffic having a source address within the defined range should not be allowed beyond the borders of the intended domain or onto the public internet.

The current prefix for specifically identifying unique local unicast addresses is fd00/8. The leftmost 64 bits of a unique local unicast address include:

- the well-known prefix “fd”
- a 40-bit global identifier
- a 16-bit subnet identifier

For example:

fd73:110:255:23:215:60ff:fe7a:adc0/64

In the above case, the following values are used with the well-known prefix and L-bit setting:

- global identifier: 0073:110:255
- subnet identifier: 23
- interface identifier: 215:60ff:fe7a:adc0

Unique local unicast addresses can be assigned by router advertisements, DHCPv6 servers, or static configuration. The boundaries for unique local unicast address are set by border routers. Unique local unicast addresses can be assigned in DNS servers supporting an internal network, but should not be included in global DNS assignments.

For related information, refer to:

- RFC 4193: “Unique Local IPv6 Unicast Addresses”

Anycast Addresses

Network size, traffic loads and the potential for network changes make it desirable to build in redundancy for some network services to provide increased service reliability. Anycast addressing provides this capability for applications where it does not matter which source is actually used to provide a service that is offered on multiple sources. Some applications that can benefit from anycast addressing include:

- DNS (UDP)
- time servers
- multicast rendezvous
- syslog devices
- gateways to a common network area.

Similarly, it is also useful in some cases to economically provide redundant paths to a given entity, such as a specific service provider. With IPv6 this can be done efficiently using the anycast address capability to assign the same address to multiple devices providing access to the desired services. An added benefit of utilizing anycast addresses is to reduce the need to configure clients with the addresses of multiple devices offering the same service.

An anycast address is an identifier for a set of interfaces typically belonging to different nodes. Packets sent to an anycast address are delivered to one of the interfaces identified as the “nearest” address, according to the routing protocol's measure of distance.

Note Equal-Cost paths between a host and multiple instances of the same anycast address can result in different packets in the same communication session to be sent to different destinations, and should be avoided.

An anycast address is formatted the same as a unicast address. For this reason, configuring an anycast address on the switch includes using an **anycast** keyword as part of the command. The prefix for an anycast address should include all areas of the network in which the address is used. For information on configuring an anycast address on the switches covered by this guide, refer to “Statically Configuring An Anycast Address” on page 4-14.

Note Duplicate Address Detection (DAD) does not apply to anycast addresses.

For related information, refer to:

- RFC 4291: "IP Version 6 Addressing Archetecture"
- RFC 2526: "Reserved IPv6 Subnet Anycast Addresses"

Multicast Application to IPv6 Addressing

Multicast is used to reduce traffic for applications that have more than one recipient for the same data. IPv6 also uses multicast for purposes such as providing a more defined control of administrative traffic on a VLAN interface than can be achieved with the broadcast method used by IPv4. This approach improves traffic control for such purposes as neighbor and router solicitations, router advertisements, and responses to DAD messages. It also avoids the bandwidth consumption used for broadcasts by narrowing the scope of possibly interested destinations for various types of messages.

Overview of the Multicast Operation in IPv6

When IPv6 is enabled on a VLAN interface on the switch, the interface automatically joins the *All-Nodes* and *Solicited-Node* multicast address groups for each of its configured unicast and anycast addresses. The interface also attempts to learn of other devices by sending solicitations to additional, well-known multicast groups, such as the following:

- all routers
- all MLDv2-capable routers, if multicast listener discovery (MLD) is enabled on the interface
- all DHCP agents (if DHCP is enabled on the interface)

There is a separate, *solicited node multicast group* for each IPv6 unicast and anycast address configured on a given interface. These automatically generated groups are limited in scope to the VLANs on which the node resides. Where multiple IPv6 unicast or anycast addresses on the same node differ only in their prefixes, they join the same solicited-node multicast group. Solicited-Node multicast groups are used, for example, in autoconfiguration. In this case, a node attempting to autoconfigure a link-local address computes the solicited-node multicast address for the proposed link-local address, then sends a Neighbor solicitation to this solicited-node multicast address. If there is no response from another node, the proposed address is available for use.

For more on Neighbor Discovery, refer to "Neighbor Discovery (ND)" on page 4-17.

For information on Multicast Listener Discovery (MLD) refer to the chapter titled “Multicast Listener Discovery (MLD) Snooping”.

When MLD is enabled on an interface, you can use **show ipv6 mld [vlan < *vid* >]** to list the active multicast group activity the switch has detected per interface from other devices.

IPv6 Multicast Address Format

The multicast address format has three principal sections in the leading 16 bits:

- identifier: ff (bits 1-8)
- flags: 0xxx (bits 9-12)
- scope: 0001 - 1110 (bits 13-16)

For related information, refer to RFC 4291.

Multicast Group Identification

Multicast ID, Flags and Scope (16 bits)	Group Identifier (112 bits)
1111 1111 0*xxx xxxx*:	*x...x* : *x...x* : *x...x* : *x...x* : *x...x* : *x...x* : *x...x* : *x...x*

- **multicast identifier:** The first eight high-order bits, set to ff, identify the address as multicast.
- **multicast flags:** Bits 9-12 are multicast flags that provide additional information about the multicast address, as follows:

Bit ID	Options	Use
9	0	reserved
10 (R)	0	multicast address without PIM-SM rendezvous point
	1	multicast address with PIM-SM rendezvous point
11 (P)	0	multicast address without prefix information from the originating network
	1	multicast address with prefix information from the originating network
12 (T)	0	multicast address is permanent (well-known, and not restricted by scope value)
	1	multicast address is temporary (and used only within an identified scope)

- **multicast scope:** Bits 13-16 set boundaries on multicast traffic distribution, such as the interface defined by the link-local unicast address of an area, or the network boundaries of an organization. Because IPv6 uses multicast technology in place of the broadcast technology used in IPv4, the multicast scope field also controls the boundaries for broadcast-type traffic sent in multicast packets.

Bit	Use
0	reserved
1	interface-local (loopback)
2	link-local (same topology as the corresponding link-local unicast scope)
3	reserved
4	admin-local (smallest administratively configured scope)
5	site-local (single site)
6	*unassigned*
7	*unassigned*
8	organization-local (multiple sites within the same organization)
9	*unassigned*
A	*unassigned*
B	*unassigned*
C	*unassigned*
D	*unassigned*
E	global
F	reserved

  For example, the following prefix indicates multicast traffic with a temporary multicast address and a link-local scope:

  ff12 or (binary) 1111 1111 0001 0010

- **group identifier:** This field includes the last 112 bits of the multicast address and contains the actual multicast group identity. (Refer to RFCs 3306, 4291, and 2375.)

Solicited-Node Multicast Address Format

The solicited-node multicast address the switch generates for a configured unicast or anycast address is composed of a unique, 104-bit multicast prefix (ff02:0:0:0:0:1:ff) and the last 24 bits of the subject address. For example, if a VLAN interface is configured with a link-local address of

fe90::215:60ff:fe7a:adc0

then the corresponding solicited-node multicast address is

ff02:0:0:0:0:1:ff7a:adc0

For related information, refer to:

- RFC 2375: IPv6 Multicast Address Assignments
- RFC 3306: Unicast-Prefix-based IPv6 Multicast Addresses
- RFC 3956: Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address
- RFC 3177: IAB/IESG Recommendations on IPv6 Address Allocations to Sites
- RFC 4007: IPv6 Scoped Address Architecture
- RFC 4291: IP Version 6 Addressing Architecture
- "Internet Protocol Version 6 Multicast Addresses" (at **www.iana.org**)
- RFC 2710: Multicast Listener Discovery (MLD) for IPv6
- RFC 3810: Multicast Listener Discovery Version 2 (MLDv2) for IPv6 (Updates RFC 2710.)

Loopback Address

The IPv6 loopback address is a link-local unicast address that enables a device to send traffic to itself for self-testing purposes. The loopback address does not have a physical interface assignment. If an IPv6 packet destined for the loopback address is received on a switch interface, it must be dropped. The IPv6 loopback address is never used as the source IPv6 address for any packet that is sent out of a device, and the switch drops any traffic it receives with a loopback address destination. An example use case is:

```
ProCurve# ping6 ::1

0000:0000:0000:0000:0000:0000:0000:0001 is alive, time = 1 ms
```

The Unspecified Address

The “unspecified” address is defined as 0.0.0.0.0.0.0.0 (::/128, or just ::). It can be used, for example, as a temporary source address in multicast traffic sent by an interface that has not yet acquired its own address. The unspecified address cannot be statically configured on the switch, or used as a destination address.

IPv6 Address Deprecation

Preferred and Valid Address Lifetimes

Autoconfigured IPv6 global unicast addresses acquire their valid and preferred lifetime assignments from router advertisements. A *valid* lifetime is the time period during which an address is allowed to remain available and usable on an interface. A *preferred* lifetime is the length of time an address is intended for full use on an interface, and must be less than or equal to the address's valid lifetime.

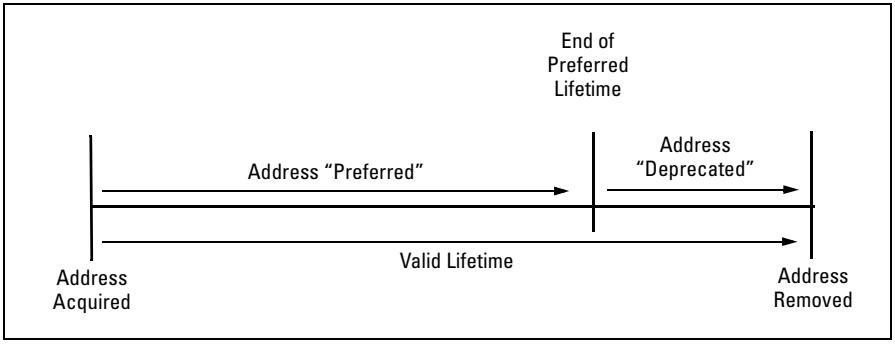

Figure 3-1. Valid and Preferred Lifetimes

When the preferred lifetime expires, the address becomes *deprecated*, meaning that the address should no longer be used as a source address (except for existing exchanges that began before the timeout occurred), but can still be used as a destination. When the timeout arrives for the valid lifetime, the address becomes unusable.

Notes

Preferred and valid lifetimes on a VLAN interface are determined by the router advertisements received on the interface. These values are not affected by the lease time assigned to an address by a DHCPv6 server. That is, lease expiration on a DHCPv6-assigned address terminates use of the address, regardless of the status of the RA-assigned lifetime, and router-assigned lifetime expiration of a leased address terminates the switch's use of the address. (The router-assigned lifetime can be extended by receipt of a new router advertisement.)

Statically configured IPv6 addresses are regarded as permanent addresses, and do not expire.

Related Information

- RFC 2462: "IPv6 Stateless Address Autoconfiguration"
- RFC 4291: "IP Version 6 Addressing Architecture"