Technical Guide

Internet Protocol v6 (IPv6)

FEATURE OVERVIEW AND CONFIGURATION GUIDE

Introduction

This guide describes the main features of IPv6, the switch's implementation of IPv6 and how to configure and operate IPv6 on the switch.

The following IPv6 features are discussed:

- linking together networks that run IPv6.
- allowing address autoconfiguration of hosts connected to the switch.

Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support IPv6, running version **5.4.4** or later.

However, support and implementation of IPv6 varies between products. To see whether a product supports a particular feature or command, see the following documents:

- The product's Datasheet
- The AlliedWare Plus Datasheet
- The product's Command Reference

These documents are available from the above links on our website at alliedtelesis.com.

Feature support may change in later software versions. For the latest information, see the above documents.

Content

Overview

IPv6 is the next generation of the Internet Protocol (IP). It has primarily been developed to solve the problem of the eventual exhaustion of the IPv4 address space, but also offers other enhancements. IPv6 addresses are **16** bytes long, in contrast to IPv4's 4 byte addresses. Other features of IPv6 include:

- Address structure improvements:
  - globally unique addresses with more levels of addressing hierarchy to reduce the size of routing tables
  - autoconfiguration of addresses by hosts
  - improved scalability of multicast routing by adding a "scope" field to multicast addresses
  - a new type of addressing method, the "anycast address", which sends packets to any one of a group of devices
- Removes the need for packet fragmentation en-route, by dynamic determination of the largest packet size that is supported by every link in the path. A link's MTU (Maximum Transmission Unit) must be at least 1280 bytes, compared with 576 bytes for IPv4.
- Includes a Traffic Class that allow packets to be labelled with an appropriate priority. If the network becomes congested, the lowest priority packets are dropped.
- Includes Flow labels that indicate to intermediate switches and routers that packets are part of a flow, and that a particular flow requires a particular type of service. This feature enables, for example, real-time processing of data streams. It also increases routing speed because the forwarding router or switch needs only to check the flow label, not the rest of the header. The handling indicated by the flow label can be done by the IPv6 Hop-by-Hop header, or by a separate protocol such as RSVP.

IPv6 Addresses and Prefixes

IPv6 addresses have a hexadecimal format that is made up of eight pairs of octets separated by colons. An example of a valid address is **2001:0db8:0000:0000:0260:0000:97ff:64aa**. In the interests of brevity, addresses can be abbreviated in two ways:

- Leading zeros can be omitted, so this address can be written as **2001:db8:0:0:260:0:97ff:64aa**.
- Consecutive zeros can be replaced with a double colon, so this address can be written as **2001:db8::260:0:97ff:64a**. Note that a double colon can replace any number of consecutive zeros, but an address can contain only one double colon.

Like IPv4 addresses, a proportion of the leftmost bits of the IPv6 address can be used to indicate the subnet, rather than a single node. This part of the address is called the *prefix*. Prefixes provide the equivalent functionality to a subnet mask in IPv4, allowing a subnet to be addressed, rather than a single node. If a prefix is specified, the IPv6 address is followed by a slash and the number of bits that represent the prefix. For example, **2001::/16** indicates that the first 16 bits (**2001**) of the address **2001:0:0:0:0:0:0:0** represent the prefix.

Like IPv4 addresses, IPv6 addresses are attached to interfaces. Note that IPv6 addressing is supported on PPP interfaces as well as VLAN and Eth interfaces.

Address types

IPv6 supports the following address types:

- Unicast
- Multicast
- Anycast

Unicast addresses

A unicast address is attached to a single interface and delivers packets only to that interface.

Unicast addresses can be grouped into 3 subcategories:

1. Link-local
2. Unique-local
3. Global

Link-local—these addresses start with **FE8x**: and are used in a single link or subnet. Any packets that are transmitted with a link local source/destination address are never routed out of that subnet.

Unique-local—originally called site-local, these are the equivalent of IPv4 **private** addresses (RFC1918), that are used within a local organization. Unique-local addresses cannot be routed across the global Internet IPv6 address space. L3 devices will not forward any packets with unique-local source or destination addresses outside of the private enterprise or customer site. IPv6 routing between multiple Unique-local subnets within a private enterprise is allowed.

There is a bit of history to which address ranges have become used for local addresses. Originally it was the range FEC0 : : /10 (RFC 1884). But the term 'site-local' was not well defined in the original definition of site-local addresses. The use of FEC0 : : /10 was deprecated in RFC3879. Shortly later, a new range was defined - FC00 : : /7 (RFC 4193) for Unique-local address ranges.

Global—these addresses start with either a 2xxx: or 3xxx: they are the equivalent of public IPv4 addresses. Global addresses can be routed publicly in the Internet. Any device or site that wishes to transmit packets to another site must be uniquely identified with a global address. Some global addresses are allocated to special purposes:

Reserved for documentation:

- 3FFF:FFFF::/32
- 2001:0db8:/32

Used for 6 to 4 tunneling:

- 2002::/16

Used for IPv4 mapped IPv6 addresses:

- ::ffff:0:0/96

The following special addresses have been defined:

- IPv4-compatible and IPv4-mapped addresses. IPv4-compatible addresses are used to tunnel IPv6 packets across an IPv4 network. IPv4-mapped addresses are used by an IPv6 host to communicate with an IPv4 host. The IPv6 host addresses the packet to the mapped address.
- The Loopback address, consisting of **::1**, which is the equivalent of the IPv4 loopback address and allows a host to send packets to itself.
- The Unspecified address, consisting of **::**, which is the equivalent of the IPv4 unspecified address and is used as a source address by hosts during the autoconfiguration process.

Multicast addresses

Multicast addresses start with FFxx: and they operate the same as the IPv4 multicast addresses. Interfaces can belong to one or more multicast groups and will accept a multicast packet only if they belong to the group corresponding to the packet's destination address.

There are no broadcast packets in IPv6, instead the IPv6 protocol uses IPv6 multicast packets to do the job of an IPv4 broadcast packet. Multicasting provides a much more efficient mechanism than broadcasting, which requires that every host on a link accept and process each broadcast packet.

Multicast address scopes

The scope of a multicast address is indicated by the fourth hex digit in the address, i.e. The digit 'S' in FF0S::

Table 1: Multicast address scopes

VALUE	SCOPE	MEANING
FF01::	node-local	Contained within a single device*
FF02::	link-local	Forwarded only within a subnet on an Ethernet segment
FF04::	admin-local	Forwarded within a small administratively- defined topology
FF05::	site-local	Forwarded only within a single site
FF08::	organisational-local	Forwarding can span multiple sites of a single organization
FF0E::	global	Can be sent across the Internet

* Node-local means that the scope for the address is within the node itself, e.g. a PC streams multicast data with node-local scope, thus only other applications in your PC can join/see the stream, and the stream never goes out of the PC on any of the interfaces.

Anycast addresses

An anycast address is a unicast address that is attached to more than one interface. If a packet is sent to an anycast address it is delivered to the nearest interface with that address, with the definition of "nearest" depending on the protocol used for routing. If the protocol is RIPv6, the nearest interface is the one that is the shortest number of hops away.

Anycast addresses can be assigned to routers only, and packets cannot originate from an anycast address. A router must be configured to know that it is using an anycast address because the address format cannot be distinguished from that of a unicast address.

Only one anycast address has been predefined: the subnet-router address. The subnet-router address sends messages to the nearest router on a subnet and consists of the subnet's prefix followed by zeros.

IPv6 headers

The basic unit of data sent through an Internet is called a packet in IPv6. A packet consists of a header followed by the data. The following figure shows the IPv6 packet.

Figure 1: IPv6 packet

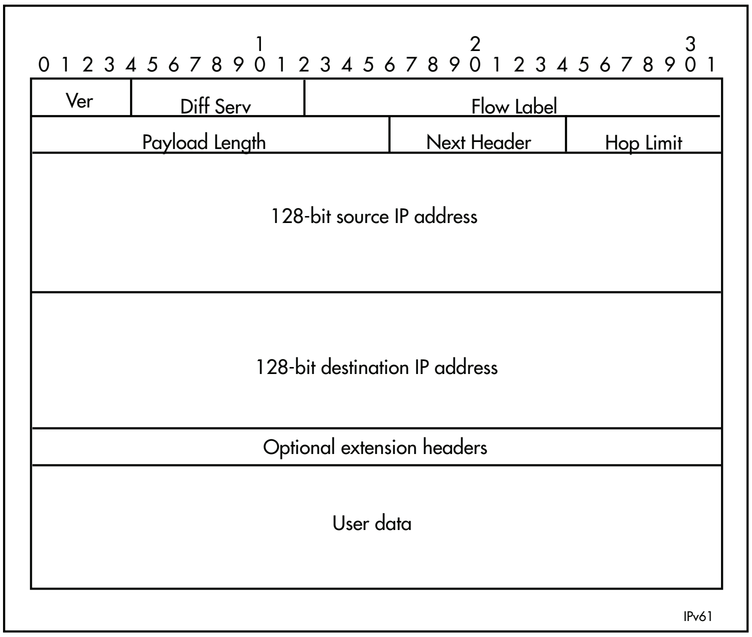

Table 2: IPv6 packet - Field Descriptions

FIELD	FUNCTION
Ver	Version of the IP protocol that created the packet. For IPv6, this field has a value of 6.
Differentiated Services	8-bit value that contains the 6-bit DSCP and is used to prioritize traffic as part of a Quality of Service system. Additional information can be found in RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers.
Flow Label	20-bit value that indicates the data flow to which this packet belongs.This flow may be handled in a particular way.
Payload Length	Length of the user data portion of the packet. If the data payload is larger than 64 kB, the length is given in the optional "Jumbo Payload" header and the Payload Length header is given a value of zero.
Next Header	Number that indicates the type of header that immediately follows the basic IP header. This header type may be an optional IPv6 extension header, a relevant IPv4 option header, or another protocol, such as TCP or ICMPv6. The IPv6 extension header values are: **0** (Hop-by-Hop Options Header) **43** (IPv6 Routing Header) **44** (IPv6 Fragment Header) **50** (Encapsulating Security Payload) **51** (IPv6 Authentication Header) **59** (No Next Header) **60** (Destination Options Header)

Table 2: IPv6 packet - Field Descriptions (Continued)

FIELD	FUNCTION
Hop Limit	Field that is the equivalent of the IPv4 Time To Live field, measured in hops.
Source IP address	128-bit IPv6 address of the sender.
Destination IP address	128-bit IPv6 address of the recipient.
Optional extension headers	Headers for less-frequently used information.
User data	Payload.

Basic IPv6 header structure

The headers contain information necessary to move the packet across the Internet. They must be able to cope with missing and duplicated packets as well as possible fragmentation (and reassembly) of the original packet.IPv6 headers are twice as long as IPv4 headers (40 bytes instead of 20 bytes) and contain four times the address space size (128 bits instead of 32 bits).

They no longer contain the header length, identification, flags, fragment offset, and header checksum fields. Some of these options are placed in extension headers. The Time To Live field is replaced with a hop limit, and the IPv4 Type of Service field is replaced with a Differentiated Services field.

The Differentiated Services field contains the DSCP bits, used in a Quality of Service (QoS) regime.

The following table explains IPv4 header fields that changed in IPv6:

Table 3: IPv4 header fields changed in IPv6

CHANGED FIELD	DESCRIPTION
Type of Service	The type of service that a connection should receive is indicated in IPv6 by the Flow Label field in the IPv6 header.
Fragmentation information (the Identification field, the Flags field and the Fragment Offset field)	In most cases fragmentation does not occur in IPv6. If it does, packets are fragmented at their source and not en route. Therefore, the fragmentation information is contained in an extension header to reduce the size of the basic IPv6 header.
Header Checksum	This option has not been provided in IPv6. This is because transport protocols implement checksums and because of the availability of the IPsec authentication header (AH) in IPv6.
Options	Extension headers handle all the optional values associated with IPv6 packets. The biggest advantage of this scheme is that the size of the basic IP header is a constant.

Extension headers

IPv6 implements many of the less commonly used fields in the IPv4 header (or their equivalents) as extension headers, which are placed after the basic IPv6 header. The length of each header must be a multiple of 8 bytes.

The first extension header is identified by the Next Header field in the basic IPv6 header. Any subsequent extension headers are identified by an 8-bit "Next Header" value at the beginning of the preceding extension header.

IPv6 nodes that originate packets are required to place extension headers in a specific order:

1. The basic IPv6 header. This must come immediately before the extension headers.
2. The Hop-by-Hop header. This specifies options that must be examined by every node in the routing path.
3. A Destination Options header. This is used to specify options to be processed by the first destination or final destination. The destination options header is the only extension header that may be present more than once in the IPv6 packet.
4. The Routing header. This enables a static path to be specified for the packet, if the dynamically-determined path is undesirable.
5. The Fragment header. This indicates that the source node has fragmented the packet, and contains information about the fragmentation.
6. The Authentication header (AH). This verifies the integrity of the packet and its headers.
7. The Encapsulating Security Payload (ESP) header. This encrypts a packet and verifies the integrity of its contents.
8. The Upper Layer Protocol header. This indicates which protocol a higher layer (such as the transport layer) is to process the packet with (for example, TCP).

The Internet Control Message Protocol (ICMPv6)

The Internet Control Message Protocol, ICMPv6, provides a mechanism for error reporting and route discovery and diagnostics. It also conveys information about multicast group membership, a function that is carried out by the Internet Group Management Protocol (IGMP) in IPv4, and performs address resolution, which the Address Resolution Protocol (ARP) performs in IPv4.

Significant aspects of ICMPv6 include neighbor discovery, which enables one device in a network to find out about other nearby devices; and stateless address autoconfiguration, which allows a device to dynamically determine its own IPv6 address.

ICMPv6 is also used to support the Ping v6 (Packet Internet Groper) and Trace route v6 functions that are used to verify the connections between networks and network devices. Ping is used to test the connectivity between two network devices to determine whether each network device can "see" the other device. Trace route is used to discover the route used to pass packets between two systems running the IP protocol.

Ping and Trace route operate almost identically in IPv4 and IPv6.

Neighbor discovery

Neighbor discovery is an ICMPv6 function that enables a router or a host to identify other devices on its links. This information is then used in address autoconfiguration, to redirect a node to use a more appropriate router if necessary, and to maintain reachability information with its neighbors.

The IPv6 Neighbor Discovery protocol is similar to a combination of the IPv4 protocols ARP, ICMP Router Discovery and ICMP Redirect.

The following table describes packet types involved with neighbor discovery:

Table 4: Packet types involved with neighbor discovery

PACKET TYPE	DESCRIPTION
router solicitation	Packet in which a host sends out a request for routers to generate advertisements.
router advertisement	Allows routers to advertise their presence and other network parameters. A router sends an advertisement packet in response to a solicitation packet from a host.
neighbor solicitation	Packet in which a node sends a packet to determine the link layer address of a neighbor or to verify that a neighbor is still active.
neighbor advertisement	A response to a neighbor solicitation packet. These packets are also used to notify neighbors of link layer address changes.
redirect	Informs hosts of a better first hop.

To comply with Section 6.2.1 of RFC 2461, IPv6 Neighbor Discovery, the router does not generate router advertisements by default.

The following table explains packet types and services:

Table 5: Packet types and services

PACKET TYPE	DESCRIPTION
address resolution	A method for carrying out address autoconfiguration, and is achieved using the Neighbor Solicitation Message and the Neighbor Advertisement Message.
router and prefix discovery	On connection to a link, a node needs to know the address of a router that the node can use to reach the rest of the world. The node also needs to know the prefix (or prefixes) that define the range of IP addresses on its link that it can reach without going through a router. Routers use ICMP to convey this information to hosts, by means of router advertisements. The message may have an option attached (the source link address option), which enables the receiving node to respond directly to the router, without performing a neighbor solicitation.
immediate information	The configuration of a router includes a defined frequency at which unsolicited advertisements are sent. If a node wants to obtain information about the nearest router immediately, rather than waiting for the next unsolicited advertisement, the node can send a router solicitation message. Each router that receives the solicitation message sends a router advertisement specifically to the node that sent the solicitation.
redirection	If a node is aware of more than one router that it can use to connect to wider networks, the router to which it sends packets by default does not always represent the most desirable route. ICMPv6 uses the redirect packet to communicate a more effective path to the node.
Neighbor Unreachability Detection (NUD)	A node may issue solicitation requests to determine whether a path is still viable, or may listen in on acknowledgement packets of higher layer protocols, such as TCP. If the node determines that a path is no longer viable, it attempts to establish a new link to the neighbor, or to re-establish the previous link. NUD can be used between any two devices in the network, independent of whether the devices are acting as hosts or routers.

There are five IPv6 Neighbor Discovery messages that replace existing IPv4 messages:

Table 2: IPv6 replacement types

IPV6 DISCOVERY MESSAGES	ICMPV6 TYPE	REPLACE THESE IPV4 MESSAGES
Router Solicitation Router Advertisement	133 134	ICMPv4 Router Discovery
Neighbor Solicitation Neighbor Advertisement	135 136	ARP
Redirect	137	ICMPv4 Redirect

Operation of Neighbour and Router Discovery

Neighbor solicitation

IPv6's replacement for ARP is Neighbor solicitation, which uses two ICMP messages:

- Neighbor Solicitation (ICMPv6 Type 135)
- Neighbor Advertisement (ICMPv6 Type 136)

Neighbor solicitation messages perform the following functionality:

- They allow IPv6 nodes (IPv6 hosts and IPv6 routers) to resolve the Link Layer address of a neighboring node (a node on the same physical or logical link).
- When the Link Layer address of a neighboring node has changed, Neighbor discovery messages allow the other IPv6 nodes to learn that this address has changed.
- They enable IPv6 nodes to determine whether neighboring nodes are still reachable.

In the diagram below, Host A sends a multicast packet (Neighbor solicitation), and if Host C is operational it will respond to this packet with a Neighbor advertisement packet.

Figure 2: Neighbor solicitation

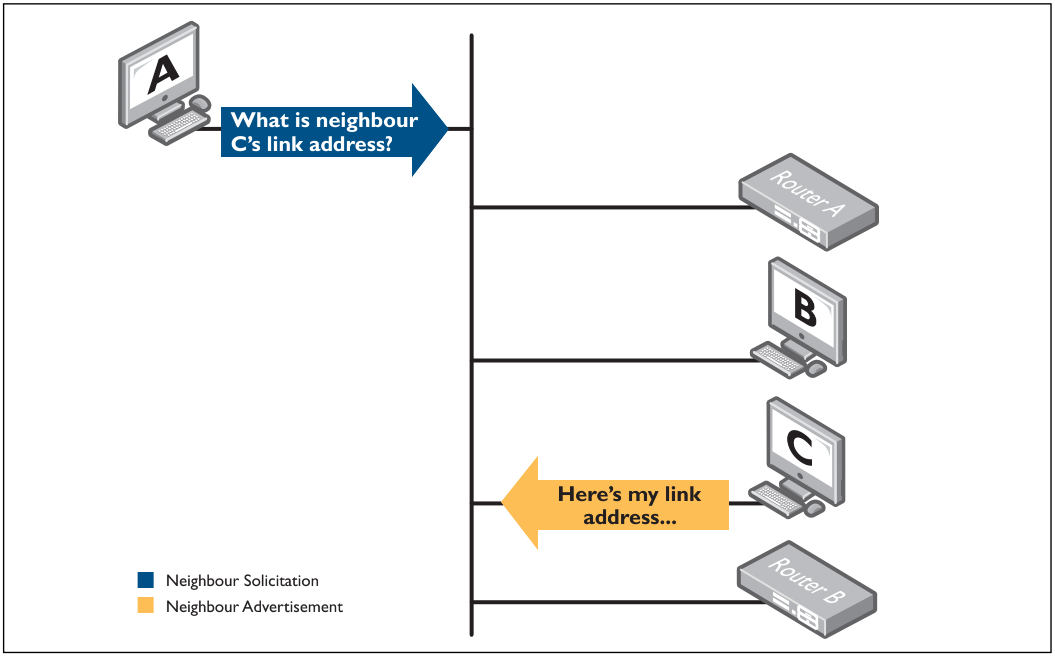

Solicited-node multicast address

When requesting the identity of the host that possesses a given IPv6 address, it is more efficient to multicast the request to potential candidates, rather than broadcast to all hosts. This means that hosts that cannot possibly possess the address do not have to process unnecessary broadcast packets. Solicited Node addresses are often flooded by switches and filtered by NIC cards drivers.

The multicast address used is called the solicited-node multicast address. It is created by attaching the last three bytes of the requested address to FF02::1:FF00:0

For example:

Address being requested: 2001: : 2AA:FF:F28:9C5A

1. Begin with FF02:0000:0000:0000:0000:0001:FF00:0000
2. Take the last 3 bytes of the requested address: 28:9C5A
3. Attached them to the address FF02::1:FF00:0

This is the solicited-node multicast address: FF02:0000:0000:0000:0000:0001:FF28:9C5A

Router discovery

IPv4 hosts need either an administrator to manually configure the default gateway or DHCP to provide this information. When IPv6 is being used, the host themselves can automatically locate routers on the LAN. The host achieves this by using two different ICMPv6 messages.

They are:

- Router Solicitation (ICMPv6 Type 133)
- Router Advertisement (ICMPv6 Type 134)

When a host is first connected to a LAN, it will send an IPv6 Router Solicitation packet to request information about routers on the network. Each router which is active on the LAN will respond to this packet by sending a Router Advertisement (RA) with its address to all nodes in the group. It informs the host what network address(es) is (are) in use on the subnet. It also informs the host if it is a default gateway.

Figure 3: Router discovery

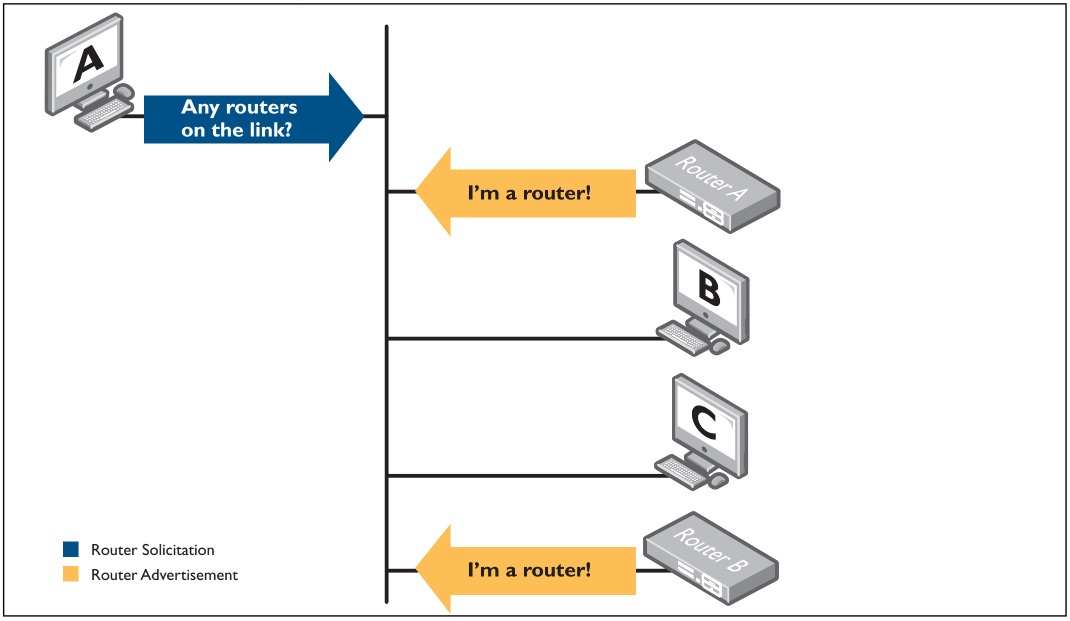

As well as responding to router solicitation events, a router will also send out router advertisements packets at regular intervals.

Configuring router advertisements on AlliedWare Plus

Router Advertisements are configured on AlliedWare Plus on a per-interface basis.

To enable RA advertisements use the command:

```
awplus(config-if)#no ipv6 ndsuppress_ra
```

The options available are:

- IPv6 nd prefix<*x:x.../N*> which sets the prefix to advertise
- IPv6 nd ra-interval <*seconds*> which sets the period of periodic advertisements
- IPv6 nd ra-lifetime <*seconds*> which sets the time for which the router will act as a default router, set this to zero to inform hosts that this is not a default router.

Redirect

Redirect uses ICMP type 137 to inform a host of a better router to use as the gateway to a given destination. If a router receives a packet and has to forward that packet to another router in the same subnet, it will also send a redirect back to the sender, telling it to send directly to the other router.

Stateless Address Autoconfiguration

Stateless address autoconfiguration allows an IPv6-aware device to be plugged into a network without manual configuration with an IP address. This plug and play functionality results in networks that are easier to set up and modify, and simplifies the process of shifting to use a new Internet Service Provider (ISP).

Stateless address autoconfiguration is achieved in a series of steps. Routers and hosts perform the first three steps, which autoconfigure a link-local address. A global address is autoconfigured in the last three steps, which only hosts perform.

On the router or host

1. During system start-up, the node begins autoconfiguration by generating a link-local address for the interface. A link-local address is formed by adding the interface ID to the link-local prefix **fe80::/10** (reference RFC 3513).

Note: Different interfaces on a device may have the same link-local address. The switch will automatically generate a link-local address for all interfaces that are using IPv6. Commands entered to configure link-local addresses that match any automatically generated link-local addresses by the switch will not be executed. Enter the show ipv6 interface command to display automatically generated link-local addresses not shown in the running-config. Automatically generated link-local addresses contain the last six hexadecimal numbers of the MAC address for a given interface.

2. The node then transmits a neighbor solicitation message to this address. If the address is already in use, the node that the address belongs to replies with a neighbor advertisement message. The autoconfiguration process stops and manual configuration of the node is then required.
3. If no neighbor advertisement is received, the node concludes that the address is available and assigns it to the chosen interface.

On the host

1. The node then sends one or more router solicitations to detect if any routers are present. Any routers present responds with a router advertisement.

   If no router advertisement is received, the node tries to use DHCP to obtain an address and other configuration information. If no DHCP server responds, the node continues using the link-level address

   If a router advertisement is received, this message informs the node how to proceed with the auto configuration process. The prefix from the router advertisement, if received, is added to the link-level address to form the global unicast IP address.
2. This address is then assigned to the network interface.

   If routers are present, the node continues to receive router advertisements. The node updates its configuration when there are changes in the router advertisements.

Configuring Stateless Address Autoconfiguration

Stateless Address Autoconfiguration (SLAAC)

- Allows an IPv6 aware device to be plugged into a network without manual configuration of an IP address.
- Has plug and play functionality which makes networks much easier to set up.
- Simplifies the process of moving to a new Internet Service Provider (ISP).

This process is described below in the section: "Setting up an IPv6 interface using the EUI-64 algorithm" on page 17.

There are two halves to the SLAAC process—the client side and the router side.

- **Client side of stateless address autoconfiguration**—when an AlliedWare Plus switch is performing as the client, a VLAN interface hasn't been configured with an IPv6 address but instead learns an IPv6 address by SLAAC. To configure this mode on an interface, use the command:

  ```
  awplus(config-if)#ipv6 address autoconfig
  ```

- **Router side of stateless address autoconfiguration**—when a client is attached to the router and hasn't been configured with an IPv6 address, the router can be configured to send out the network information in an RA (Router Advertisement) so the client is able to get an address and communicate on the LAN.

  For example, on a router's VLAN interface which has the client attached, the following configuration could be used to send the Prefix:

  ```
  awplus(config)#int vlan10
  awplus(config-if)#ipv6 address 2001:1db9:1:2::/64 eui64
  awplus(config-if)#ipv6 nd ra-interval 10
  awplus(config-if)#ipv6 nd prefix 2001:1db9:1:2::/64
  awplus(config-if)#no ipv6 ndsuppress_ra
  ```

Note: AlliedWare Plus supports both the client and the router side of stateless address autoconfiguration.

Setting up an IPv6 interface using the EUI-64 algorithm

Here is an overview of the steps that occur when a switch performs SLAAC using the EUI-64 algorithm:

1. Generate a 64 bit interface identifier using the EUI-64 algorithm.

The host has to create its own host portion of its IPv6 address. It can create a unique address from its MAC address by using the EUI-64 algorithm, here is how it works:

Table 6: Generate a 64 bit interface identifier

STEP	ADDRESS
1. Start with the MAC address	0012.7FEB.6B40
2. Split the MAC address in half	0012:7F EB:6B40
3. Insert FF:EE into the MAC address	0012:7FFF:FEEB:6B40
4. Change the 7th bit to '1'	0212:7FFF:FEEB:6B40

2. Create a Link-local node address.

When IPv6 has been configured on an interface, the switch will automatically assign a link-local address to that interface. Link-local addresses are used as the source address for packets that stay within the subnet, for example:

- automatic address configuration
- neighbor discovery
- OSPF exchanges etc.

Any packets that are transmitted with a link-local source/destination address are never routed out of that subnet and are assigned the fe80::/10 prefix, equivalent to the IPv4 address block 169.x.x.x.

The link-local address for an interface is created by combining the EUI-64 host address to the network address FF80::64

FF80:0000:0000:0000: + 0212: 7FFF:FEEB:6B40= FF80 : : 0212: 7FFF:FEEB:6B40

3. Send router solicitation messages to all routers on the local link multicast address. If there is no response, SLAAC ends with only a link-local address generated.

Note: If a periodic RA containing the appropriate information is received at any time, SLAAC will use it immediately. Also, if an RA is received that reduces the lifetime of a prefix to zero, SLAAC will immediately deprecate the address (the system will then cease using it for new connections, but existing ones will continue).

4. **Once a prefix is learnt by RA, prepend the prefix to the EUI-64 interface ID, to create the full IPv6 address.**

2001:639A:1234:5678:: + 0212: 7FFF:FEEB:6B40

=

2001:639A:1234:5678:0212:7FFF:FEEB:6B40

5. **Find default gateway (default routers). On receipt of a valid Router Advertisement, a host extracts the source address of the packet and does the following:**

- If the address is not already present in the host's Default Router List, and the advertisement's Router Lifetime is non-zero, it creates a new entry in the list and initializes its invalidation timer value from the advertisement's Router Lifetime field.
- If the address is already present in the host's Default Router List as a result of a previously received advertisement, it resets its invalidation timer to the Router Lifetime value in the newly received advertisement.
- If the address is already present in the host's Default Router List and the received Router Lifetime value is zero, it immediately times-out the entry as specified.

To limit the storage needed for the Default Router List, a host may choose not to store all of the router addresses discovered via advertisements. However, a host must retain at least two router addresses and should retain more. Default router selections are made whenever communication to a destination appears to be failing. Thus, the more routers on the list, the more likely an alternative working router can be found quickly (without having to wait for the next advertisement to arrive).

Encryption and authentication in IPv6

IPv4 protocols such as OSPFv2, have authentication incorporated into their own protocol header.

In IPv6, authentication and encryption are performed by separate IP headers, completely independent to the enclosed protocol.

AH – Authentication Header – commonly MD5 or SHA

The authentication information for the Authentication Header is calculated using all the fields of the datagram that do not change in transit.

This header can be used as part of IPSec to authenticate end point to end point packets. This can be used to protect protocols like OSPFv3, IPv6, BGP, RADIUS, TACACS+, and RIPng.

ESP – Encapsulated Security Payload – commonly 3DE S or AES

ESP is used to convey the encrypted data of the IP datagram. The encrypted data is obtained by applying a specified encryption transform to the data and requires the use of a key in order to return to plain text.

There are two modes used for ESP:

- Tunnel Mode
- Transport mode

Tunnel Mode, where the **entire IP packet** is encrypted and/or authenticated. It is then encapsulated into a new IP packet with a new IP header. ESP Tunnel mode encrypts the whole IP datagram:

Figure 4: Tunnel mode ESP

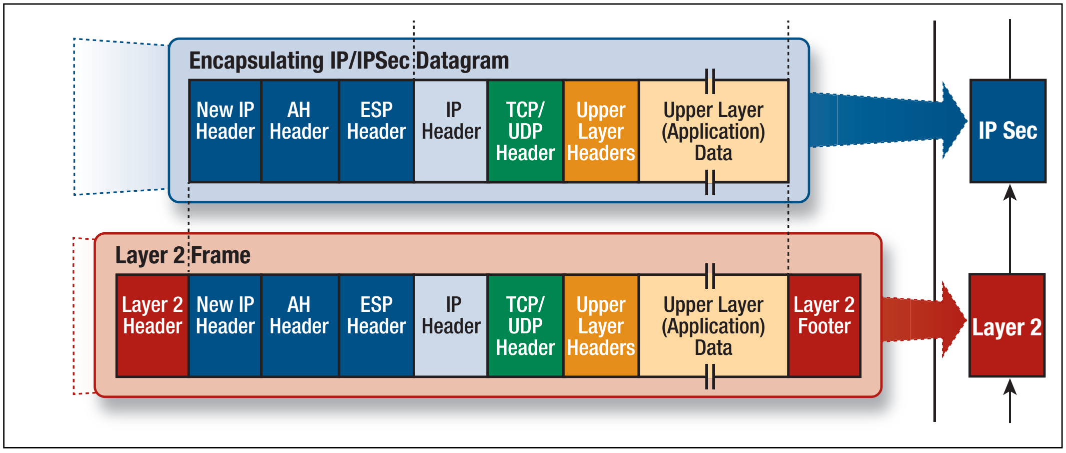

In ESP Tunnel mode, the Authentication Header appears as an extension header of the new IP datagram that encapsulates the original one being tunneled.

Transport Mode, where **only the payload** of the IP packet is encrypted and/or authenticated – not the original IPv6 Header. ESP Transport mode encrypts only the payload (Transport Layer message of the IPv6 datagram):

Figure 5: Transport mode ESP

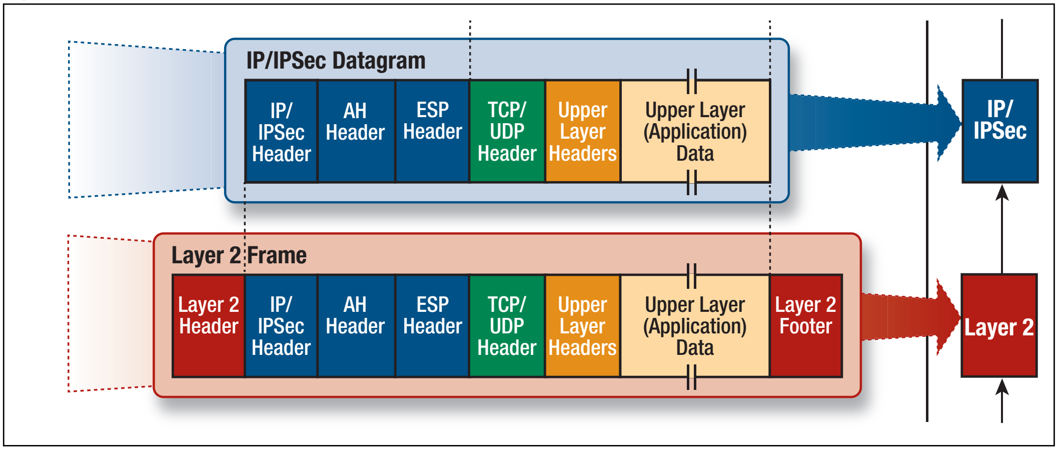

In ESP transport mode, the Authentication Header is placed into the main IP Header before any Destination Options header and before an ESP header.

The extension headers used to secure the IP communication between two hosts, Authentication and Encapsulating Security Payload Headers, are ignored by the intermediary network devices while forwarding traffic. These Extension Headers are relevant only to the source and destination of the IP packet.

All information following the ESP Header is encrypted and not available for inspection by an intermediary device.

The QoS Flow Label

The QoS Flow Label is a 20 bit field in the IPv6 packet header which provides an efficient way for packet marking, flow identification, and flow state lookup.

This field can be used by a source to label a set of packets belonging to the same flow. The switch must process the packets in the same flow in the same manner. When a flow-label aware router receives the first packet of a new flow, it sets up a new flow entry using the information carried by the IPv6 header, Routing header, and Hop-by-Hop extension headers, and stores the result.

It then uses the flow entry to route all other packets belonging to the same flow – which will have the same source address and the same Flow Label.

IPv6 routing

Routing in IPv6 is almost identical to IPv4 routing under CIDR, except that the addresses are **128-bit** IPv6 addresses instead of 32-bit IPv4 addresses.

Routing Information Protocol (RIPv6)

RIP is a simple distance vector protocol that defines networks based on how many hops they are from the router. When a network is more than 15 hops away (one hop is one link), it is not included in the routing table.

RIPv6, also referred to as RIPng (for "next generation") is similar to RIPv2. Extensions to RIPv2 to support IPv6 are:

- the address field of a routing entry is expanded to 128 bits to allow IPv6 prefixes
- the 32-bit RIPv2 subnet mask field is replaced by an 8-bit prefix length field
- authentication is removed in RIPv6
- the size of a routing packet is no longer arbitrarily limited
- RIPv6 specifies the next hop instead of simply allowing the recipient of the update to set the next hop to the sender of the update.

In RIPv6, each router uses a routing table to keep track of every destination that is reachable throughout the system. Each entry in the routing table contains:

- the IPv6 prefix of the destination
- a metric, which represents the total cost of getting a packet from the router to that destination
- the IPv6 address of the next router along the path to the destination
- a flag to indicate that information about the route has changed recently
- various timers associated with the route.

Integration of IPv4 and IPv6

IPv6 has been designed in such a way that a smooth transition from IPv4 is possible. The most effective way to ensure this is to use a *dual IP stack*. A node configured as a dual stack system has both a 128-bit IPv6 address and a 32-bit IPv4 address, and so can communicate with nodes running IPv4 and those running IPv6.

IPv6 on your Switch

This section describes the switch's support for IPv6, and how to configure IPv6 on the switch.

Enabling IPv6

IPv6 Layer 3 forwarding is disabled by default. To enable IPv6 forwarding, use the **ipv6 forwarding** command.

To display information about IPv6 settings, use the **show ipv6 interface brief** command.

Because AlliedWare Plus implements IPv6 as a dual stack, implementing IPv6 does not affect IPv4 functionality.

IPv6 Stateless Address Autoconfiguration (SLAAC)

The AlliedWare Plus implementation of IPv6 supports SLAAC on an interface. To enable IPv6 SLAAC on an interface, use the **ipv6 address autoconfig** command. SLAAC automatically applies the MAC address of the interface to an IPv6 address for the interface specified.

The **ipv6 address autoconfig** command enables automatic configuration of IPv6 addresses on an interface using stateless autoconfiguration, and enables IPv6 processing on an interface.

IPv6 EUI-64 addressing

The AlliedWare Plus implementation of IPv6 supports EUI-64 addressing. EUI-64 applies an IPv6 address that is based on the MAC address of the interface. The EUI-64 identifiers from the MAC address are used as the least significant 64 bits of a unicast address.

To enable IPv6 EUI-64, use the **ipv6 address** command, and specify the optional **eui64** parameter for an interface.

When configuring SLAAC you must ensure that you set the prefix length to 64 bits on the switch that is advertising the RAs used for address configuration via SLAAC.

Prefix information received in an RA (Router Advertisement) will not be applied to form an IPv6 address via SLAAC unless the prefix length is 64. Since the EUI is 64 bits long, the IPv6 prefix of the advertising device must also be 64 bits. This prefix length setting and behavior is in accordance with RFC 4864, section 5.5.3.

IPv6 link-local addresses

The AlliedWare Plus implementation of IPv6 supports IPv6 link-local addresses without global addresses for communications within the local subnetwork. Switches do not forward packets to link-local addresses. To enable IPv6 link-local addresses, use the **ipv6 enable command**. This command automatically configures an IPv6 link-local address on the interface and enables IPv6 processing on the interface.

Note that link-local addresses are retained in the system until they are negated by using the **no** variant of the command that established them. Also note that the link-local address is retained if the global address is removed using a command that was not used to establish the link-local address. For example, if a link local address is established with the **ipv6 enable** command then it will not be removed using a **no ipv6 address** command.

RA Guard

Router Advertisements (RA) and Router Redirects are key to the Network Discovery Protocol (NDP) used to manage IPv6 networks. RA messages advertise a router's presence and specify network parameters that are used by hosts as part of address auto-configuration and next hop routers for particular destinations.

Subverting this process can severely disrupt the operation of an IPv6 network. RA Guard is a feature that protects the RA process from being subverted.

RA Guard:

- is positioned in between routers and hosts, and acts as an authorisation proxy.
- drops bad RAs before they reach hosts.
- operates on all AlliedWare Plus Layer 3 switches, including stacked environments.

Rogue RAs

A rogue RA is an RA that contains invalid information that could cause unwanted changes in the network configuration. These could be generated unintentionally through misconfiguration or maliciously by someone wanting to disrupt or gain access to the network.

A switch can be configured to be selective about the RA and redirect packets it will accept. Ports are configured to trust or not trust the RA and redirect packets they receive.

RA Guard on AlliedWare Plus switches

Ports can be configured to be RA untrusted ports, i.e. RA Guard is applied to ports on a per-interface basis and can be enabled on the following:

- Standalone ports.
- Individual ports in a dynamic (LACP) aggregator, but is not supported on the dynamic aggregator itself.
- A static aggregator, but is not supported on individual ports in a static aggregator.

RA Guard is enabled on an interface as follows:

```
awplus#conf t
awplus(config)#int port1.0.2
awplus(config-if)#ipv6 nd raguard
```

Note: This feature is disabled by default.

RA Guard classifiers

The actual security enforcement of RA Guard is handled through hardware classifiers, which are dynamically added when a port is marked as trusted or untrusted.

RA Guard blocks RAs and router redirects on untrusted ports with filters for ICMPv6 type 134 and 137.

Enabling IPv6 RA Guard

IPv6 RA Guard is disabled by default. To enable IPv6 RA Guard on a port to block RAs from an untrusted host, use the **ipv6 nd raguard** command. Disable IPv6 RA Guard to allow RAs on a port using the **no ipv6 nd raguard** command.

Allied Telesis™ the **solution :** the **network**

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com