

# Chapter 1

## IPv6 Overview

Internet Protocol version 6 (IPv6) is the new version of the Internet Protocol (IP). The Internet Protocol allows numerous nodes on different networks to interoperate seamlessly. Internet Protocol version 4 (IPv4) is currently used in intranets and private networks, as well as the Internet. IPv6 is the successor to IPv4, and is based for the most part on IPv4.

This chapter discusses the following topics:

IPv6 Standards on page 3

IPv6 Introduction on page 3

IPv6 Packet Headers on page 4

IPv6 Addressing on page 5

### IPv6 Standards

IPv6 is defined in the following document:

RFC 2460, *Internet Protocol, Version 6 (IPv6)*

RFC 2373, *IP Version 6 Addressing Architecture*

To access Internet Requests for Comments (RFCs) and drafts, see <http://www.ietf.org>.

### IPv6 Introduction

IPv4 has been widely deployed and used to network the Internet today. With the rapid growth of the Internet, enhancements to IPv4 are needed to support the influx of new subscribers, Internet-enabled devices, and applications. IPv6 is designed to enable the global expansion of the Internet.

IPv6 builds upon the functionality of IPv4, providing improvements to addressing, configuration and maintenance, and security.

IPv6 offers the following benefits:

Expanded addressing capabilities—IPv6 provides a larger address space. IPv6 addresses consist of 128 bits, while IPv4 addresses consist of 32 bits. 128-bit addressing increases the address space by approximately  $10^{29}$  unique addresses, enough to last for the foreseeable future.

Header format simplification—IPv6 packet header format is designed to be efficient. IPv6 standardizes the size of the packet header to 40 bytes, divided into 8 fields.

Improved support for extensions and options—Extension headers carry Internet-layer information and have a standard size and structure.

Flow labeling capability—Flow labels provide consistent handling of packets belonging to the same flow.

Improved privacy and security—IPv6 supports extensions for authentication and data integrity, which enhances privacy and security.

## IPv6 Packet Headers

IPv6 introduces a new packet header structure.

This section discusses the following topics that provide background information about IPv6 headers:

Header Structure on page 4

Extension Headers on page 5

### **Header Structure**

IPv6 packet headers contain many of the fields found in IPv4 packet headers; some of these fields have been modified from IPv4. The 40-byte IPv6 header consists of the following 8 fields:

Traffic Class—Class-of-service (CoS) priority of the packet. Previously the type-of-service (ToS) field in IPv4. However, the semantics of this field (for example, diffserv code points) are identical to IPv4.

Destination address—Final destination node address for the packet.

Flow label—Packet flows requiring a specific CoS. The flow label identifies all packets belonging to a specific flow, and routers can identify these packets and handle them in a similar fashion.

Hop limit—Maximum number of hops allowed. Previously the time-to-live (TTL) field in IPv4.

Next header—Next extension header to examine. Previously the protocol field in IPv4.

Payload length—Length of the IPv6 payload. Previously the total length field in IPv4.

Source address—Address of the source node sending the packet.

Version—Version of the Internet Protocol.

## Extension Headers

In IPv6, *extension headers* are used to encode optional Internet-layer information.

Extension headers are placed between the IPv6 header and the upper layer header in a packet.

Extension headers are chained together using the next header field in the IPv6 header. The next header field indicates to the router which extension header to expect next. If there are no more extension headers, the next header field indicates the upper layer header (TCP header, User Datagram Protocol [UDP] header, ICMPv6 header, an encapsulated IP packet, or other items).

## IPv6 Addressing

IPv6 introduces a new 128-bit addressing model. This creates a much larger address space than IPv4 addresses, which are made up of 32 bits. IPv6 addresses also contain a scope field that categorizes what types of applications are suitable for the address. IPv6 does not support broadcast addresses, but instead uses multicast addresses to serve this role. In addition, IPv6 also defines a new type of address called *anycast*.

This section discusses the following topics that provide background information about IPv6 addressing:

Address Representation on page 5

Address Types on page 6

Address Scope on page 6

Address Structure on page 6

## Address Representation

IPv6 addresses consist of 8 groups of 16-bit hexadecimal values separated by colons (:). The IPv6 address format is as follows:

*aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa*

*aaaa* is a 16-bit hexadecimal value, and *a* is a 4-bit hexadecimal value. Following is an example of an actual IPv6 address:

3FFE:0000:0000:0001:0200:F8FF:FE75:50DF

You can omit the leading zeros, as shown:

```
3FFE:0:0:1:200:F8FF:FE75:50DF
```

You can compress 16-bit groups of zeros to “::”, as shown here, but only once per address:

```
3FFE::1:200:F8FF:FE75:50DF
```

## Address Types

There are three types of IPv6 addresses:

Unicast—For a single interface.

Multicast—For a set of interfaces on the same physical medium. A packet is sent to all of the interfaces associated with the address.

Anycast—For a set of interfaces on different physical mediums. A packet is sent to only one of the interfaces associated with this address, not to all the interfaces.

## Address Scope

IPv6 addresses have *scope*, which identifies the application suitable for the address. Unicast and multicast addresses support scoping.

Unicast addresses support two types of scope: *global scope* and *local scope*. There are two types of local scope: *link-local* addresses and *site-local* addresses. Link-local unicast addresses are used within a single network link. The first 10 bits of the prefix identify the address as a link-local address. Link-local addresses cannot be used outside a network link. Site-local unicast addresses are used within a site or intranet. A site consists of multiple network links, and site-local addresses identify nodes inside the intranet. Site-local addresses cannot be used outside the site.

Multicast addresses support 16 different types of scope, including node, link, site, organization, and global scope. A four-bit field in the prefix identifies the scope.

## Address Structure

Unicast addresses identify a single interface. The address consists of  $n$  bits for the prefix, and  $128-n$  bits for the interface ID.

Multicast addresses identify a set of interfaces. The address is made up of the first 8 bits of all ones, a 4-bit flags field, a 4-bit scope field, and a 112-bit group ID:

```
11111111 | flags | scop | group ID
```

The first octet of ones identifies the address as a multicast address. The flags field identifies whether the multicast address is a well-known address or a transient multicast address. The scope field identifies the scope of the multicast address. The 112-bit group ID identifies the multicast group.

Similar to multicast addresses, anycast addresses identify a set of interfaces. However, packets are sent to only one of the interfaces, not to all interfaces. Anycast addresses are allocated from the normal unicast address space and cannot be distinguished from a unicast address in format. Therefore, each member of an anycast group must be configured to recognize certain addresses as anycast addresses.

