# IPv6 Time Is Now

*A Primer on How IPv6 Addresses Are Constructed and the Benefits of this Next-Generation Technology*

**BY RAJ GULANI**

The current Internet address protocol has a potential of about 4 billion public 32-bit addresses. Shortages, so far caused primarily by allocation, have already forced many Internet users to make do with Network Address Translation (NAT) or private addresses piggybacking on public ones. What's more, the number of 32-bit addresses, governed by IP Version 4 (IPv4), is already shortchanging some areas of the world, particularly the Far East. It will be wholly inadequate after devices needing Internet connectivity are placed in mobile phones, personal digital assistants (PDAs), vehicles, appliances, and an as-yet unimaginable list of other devices.

The solution to the shortage of addresses that has been adopted by IP standards bodies around the world is a new address format of 128 bits, part of a protocol known as Internet Protocol Version 6 (IPv6). The 128-bit length was chosen because it creates an optimum number of addresses in a header of workable size. And it will expand the number of public addresses to 340,232,366,920,938,463,374,607,431,768,211,456, or $3.4 \times 10^{38}$.

Work has been underway for a decade, and IPv6 addressing and IPv6 networks are now realities. A number of academic and government networks already carry IPv6 traffic shifting from the experimental IPv6 backbone (also known as 6Bone networks) to production dual-purpose IPv4 and IPv6 networks now under construction. Cisco has been involved in IPv6 development since its inception. In addition to being a founding member of the IPv6 Forum and co-chair of the Internet Engineering Task Force (IETF) IPv6 and Transition (NGTrans) Working Group for several years, Cisco runs a well-known 6Bone router that helps up to 70 other companies make their first IPv6 steps, and acts as an IPv6 to IPv4 relay for individual users or small companies. And its implementation of the protocol on Cisco IOS® platforms was one of the industry's first official IPv6 support in a broad product portfolio.

The same registries that now allocate IPv4 addresses—the Asia Pacific Network Information Center (APNIC) for Asia, the American Registry for Internet Numbers (ARIN) for the Americas, and the Reseaux IP Europeens-Network Coordination Center (RIPE-NCC) for Europe and the Middle East—will allocate IPv6 addresses to Internet service providers (ISPs) and National Research Networks (NRN), who will continue to assign them to end users.

This article explores the five major areas of change brought about by IPv6:

- Enhanced address space
- Header format simplification
- Enhanced routing protocol support
- Security and mandated IP Security (IPSec) support
- Enhanced support for Mobile IP

## Enhanced Address Space

The 128 bits in the IPv6 address space are broken down into eight groups of 16 bits, each of which can be represented in text as a hexadecimal value. The text address then reads x:x:x:x:x:x:x:x, where each "x" represents a given hexadecimal value. Typical addresses would be written as 2031:0000:130F:0000:0000:09C0:876A:130B.

Because of the methods used for allocating addresses, IPv6 may contain long strings of zeros. It is acceptable to substitute a pair of colons "::" anywhere in the address to represent and thus compress two or more 16-bit groups of zeros in a string. The above address would then be written: 2031:0000:130F::09C0:876A:130B. As a computer expands the "::" with 0 to get 128 bits, an address may not contain more than one set of paired colons, as it won't be possible to determine how many zeros in each segment.

RAJ GULANI

**RAJ GULANI,** CCIE®, is a manager of technology marketing at Cisco. He has presented at Networkers and various other technical forums. With a Bachelor's Degree in Engineering (Electronics) from Bombay University, India, Gulani's networking career spans more than eight years. In 1997, he started his career at Cisco in Customer Advocacy as a Technical Assistance Center (TAC) engineer focusing on WAN and broadband technologies, and moved on to lead network design engineer in the Advanced Network Services Group. He can be reached at rgulani@cisco.com.

ROB BRODMAN

An alternative form (nearly deprecated) for the address is x:x:x:x:x:x:d:d:d:d, in which the "x's" are hexadecimal values of the six higher-order 16-bit pieces and the "d's" represent the decimal values of the four lower-order 8-bit pieces. The latter portion is the standard IPv4 address, and this form would be used when configuring automatic IPv6 tunnels through an IPv4 network. This particular format has shown some operational challenges, so its use is discouraged.

An IPv6 address typically has two major 64-bit parts: the *network prefix*, which contains the registry, provider, subscriber ID, and subnet, and occupies the higher order groups of bits; and the *interface ID*, which occupies the lower ones. Bits within these groupings are allocated according to the specific requirements of various types of addresses. An IPv6 address prefix is represented as ipv6-address/prefix-length, similar to IPv4 address prefixes written in IPv4 Classless Interdomain Routing (CIDR).

There are three major types of addresses: *unicast*, *anycast*, and *multicast*. Addresses are assigned to interfaces. A unicast address governs one-to-one transmission. It identifies a single interface and is the equivalent of an IPv4 unicast address. A transmission is delivered only to that address. There are, however, several types of unicast addresses, including the global, link-local, and site-local.

The *global unicast address* is the equivalent of the IPv4 global unicast address, used on links that are aggregated upwards though an organization and eventually to the ISP. The structure of this type of address enables policies that allow aggregation of routing prefixes so as to limit the number of entries in the global routing table. The address consists of a 48-bit routing prefix managed by the provider and a 16-bit subnet ID managed by the local site.

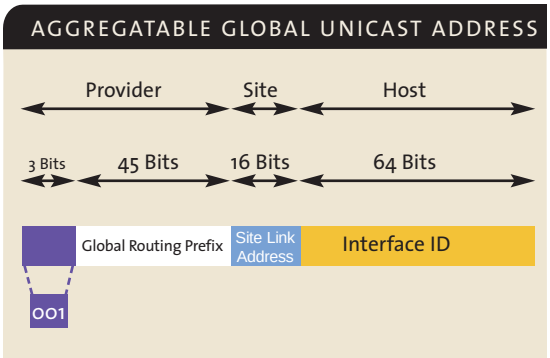


A *link-local unicast address* is used to communicate among nodes on a local link, in the neighbor discovery protocol, and in the stateless autoconfiguration process. It can be automatically configured on any interface by using the link-local prefix FE80::/10(1111 1110 10) and the interface identifier in the IEEE EUI-64 format (an EUI-64 may be derived from an EUI-48).

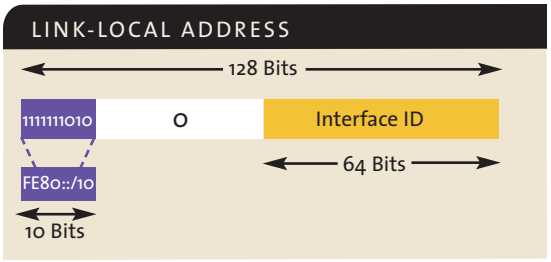


A *site-local unicast address* is similar to a private address in IPv4 format (for example, 172.16.0.0/12). Because the prefix is not propagated between routing domains, this restricts communication to a specific domain.

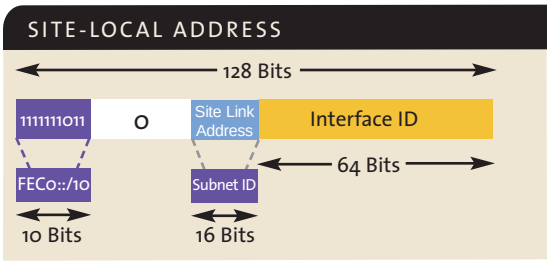


The unicast address 0:0:0:0:0:0:0:1 or ::1 is the loopback address, which should never be assigned to any interface. It performs the same function as in IPv4—identifying a transmission sent by a node back to itself.

An *anycast address* directs one-to-any transmission. Examples are the use of anycast addresses to reach a DNS server or a 6to4 Relay or Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) routers (see next page for more on 6to4 and ISATAP).

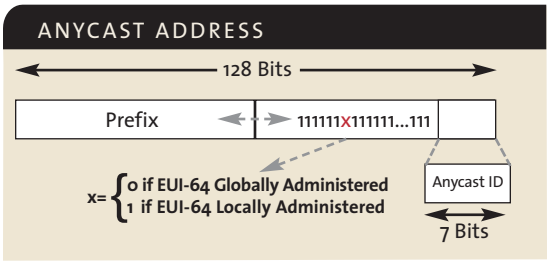


An IPv4-mapped IPv6 address is used to identify an IPv4-only node to an IPv6 node. The 16 bits of zeros before the IPv4 decimal representation in the address should be replaced with four F's. An IPv4-mapped address should never be used as an "Src" or "Dst" on the wire.

A *multicast address* governs one-to-many transmission. The address subsumes a set of interfaces, usually

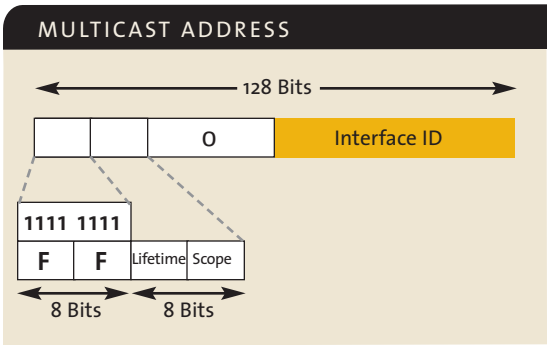


belonging to different nodes, and orders delivery of the packet to all of them. IPv6 does not include any broadcast addresses, their function being taken over by the multicast address.

### 6to4 and ISATAP Addresses

To ease the IPv6 integration on the campus and WAN, the IETF NGTrans Working Group has been working on several tunneling mechanisms that can be set when native IPv6 or dual-stack are not a solution. Automatic IPv6 over IPv4 tunnels require that an IPv4 address be extracted from the IPv6 address to establish a tunnel on request. This approach leads to a particular IPv6 unicast address called *6to4* (RFC 3056) on the WAN and *ISATAP* (draft-ietf-ngtrans-isatap) on campus.

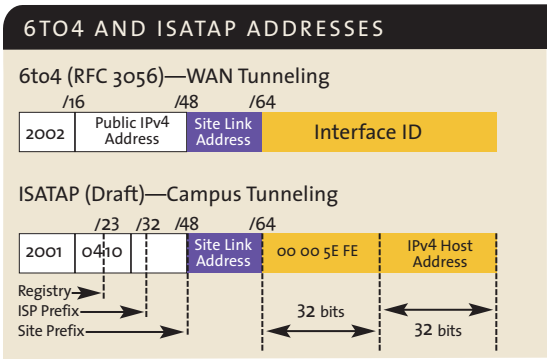


### Simplified Header Format

An IPv4 header contains at least 12 different fields: version, header length, type of service, total length, identification, flags, fragment offset, time to live, protocol, header checksum, source address, destination address, and possible options. The simpler IPv6 header has only eight: version, traffic class, flow label, payload length, next header, hop limit, source address, and destination address. They are used as follows:

- The field specifies 6 for IPv6.
- The traffic class identifies differentiated services.
- The flow label tags packets with a specific flow that differentiates packets at the network layer.
- The payload length indicates the length of the data portion of the packet.
- The "next header" field determines the type of information following this packet—perhaps a TCP or UDP packet or extension header.
- The hop limit specifies the maximum number of routers the packet can pass through before being considered invalid.
- The "source address" and "destination address" contain the appropriate IPv6 128-bit addresses.

Neighbor discovery provisions are built on top of Internet Control Message Protocol for IPv6 (ICMPv6) and are a combination of IPv4 protocols. They enable router discovery, parameters discovery, address autoconfiguration, next-hop determination, neighbor unreachability detection, and other functions, as well as determine the link-layer address of a neighbor on the same link, and identification and tracking of neighboring routers.

Stateless autoconfiguration (see Figure 1) is a key feature of IPv6.

### Enhanced Routing Protocol Support

IPv6 headers have been designed to be as compatible as possible with current standards and protocols. For example:

- They use the same "longest-prefix match" routing as IPv4 CIDR.
- The changes to existing IPv4 routing protocols for

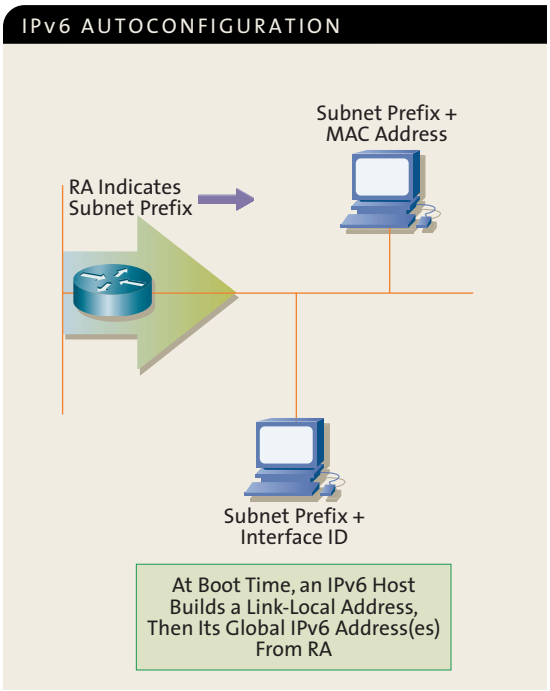


**FIGURE 1:** With IPv6 stateless autoconfiguration, the host autonomously configures its own link-local address. Booting nodes request router addresses (RAs) to help in configuring the interfaces.

handling longer addresses are done through new IETF documents.

- Use of Routing Information Protocol (RIP) is achieved with RIPng similar to RIPv2 on IPv4.
- A new version of Open Shortest Path First (OSPF)—OSPFv3—is defined for IPv6.
- IPv6 address family in now in Intermediate System to Intermediate System (IS-IS) and Multiprotocol Border Gateway Protocol v4 (MP-BGP4).
- Headers with anycast destination addresses can be used to route packets through particular regions for provider selection and performance, compliance with policies, among other reasons.

### Mandated IPSec, Enhanced Support for Mobile IP

Compliance with the IPSec standard is mandatory in IPv6, not optional as in IPv4, and is part of the IPv6 protocol suite. Network administrators could enable IPSec in every node, if desired, potentially making networks more secure, but IPv6 does not specify how the key can be distributed. IPv6 provides security header extensions that make it easy to implement encryption, authentication, and virtual private networks (VPNs). The security extensions help to prevent hacking and ensure data integrity. Nevertheless, when a centralized security mechanism (for example, firewalls) is preferred, there is no definition—yet—on how to make both IPSec and firewalls coexist on the network.

IPv6 also offers other features that facilitate mobility. Mobility is essentially built in, and any node can support it as needed, in contrast to IPv4, in which mobility must be specifically provided for.

IPv6 packets addressed to the home address of a mobile node are transparently and automatically routed to its care-of address because both addresses are bound and cached together. On Mobile IPv4, a correspondent node can't talk directly with a mobile node; it has to go through the home agent. But a mobile node can talk directly with a correspondent node. This is called triangular routing. On Mobile IPv6 (see Figure 2), a correspondent node can talk directly with a mobile node after a binding update process.

◆ ◆ ◆

Growth in mobile and fixed applications that require Internet-connected devices is just one factor putting a strain on the current IPv4 32-bit address pool. The time for IPv6 has come, and IT professionals need to begin familiarizing themselves now with IPv6 addresses and their benefits. ▲▲

ENHANCED SUPPORT FOR MOBILITY WITH IPv6

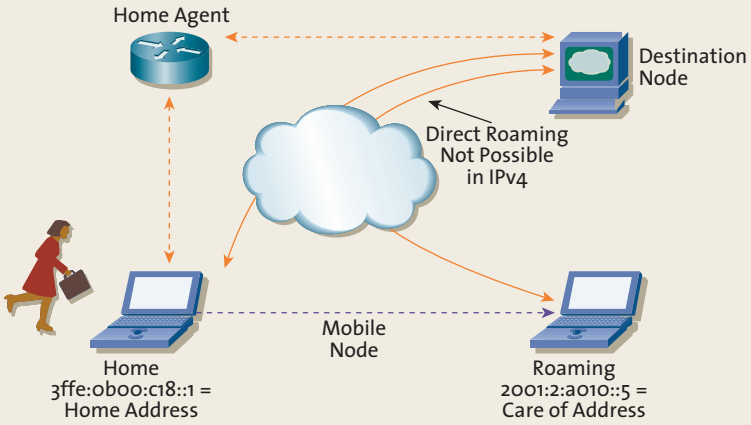


**FIGURE 2:** With IPv6, use of the routing header for Mobile IP enables the service to avoid "triangle routing," making the service much more efficient than in IPv4.

## The IPv4 to IPv6 Transition

A number of techniques have been identified and developed to make the transition from IPv4 to IPv6. They fall into three categories:

- *Dual-stack techniques,* which allow both protocols to coexist in the same devices and networks
- *Tunneling techniques,* to avoid order dependencies when upgrading hosts, routers, or regions
- *Translation techniques,* to allow IPv6-only devices to communicate with IPv4-only devices

These three techniques, which probably will be used in various combinations, also lay the foundation for IPv6 deployment scenarios. To view a dual-stack IPv4/IPv6 campus and ISP deployment scenario, go to cisco.com/go/packet/dual-stack at *Packet Online*.

### FURTHER READING

- **Cisco IPv6:** cisco.com/ipv6
- **IPv6 Forum:** ipv6forum.org
- **IPv6 Information Page:** ipv6.org
- **IETF IPv6 Working Group:** ietf.org/html.charters/ipv6-charter.html