IPv6

Internet Protocol version 6 (**IPv6**) is a network layer protocol for packet-switched internetworks. It is designated as the successor of IPv4, the current version of the Internet Protocol, for general use on the Internet.

The main improvement brought by IPv6 is the increase in the number of addresses available for networked devices, allowing, for example, each mobile phone and mobile electronic device to have its own address. IPv4 supports 2^{32} (about 4.3 billion) addresses, which is inadequate for giving even one address to every living person, let alone supporting embedded and portable devices. IPv6, however, supports 2^{128} (about 340 billion billion billion billion) addresses, or approximately 5×10^{28} addresses for *each* of the roughly 6.5 billion people alive today. With such a large address space available, IPv6 nodes can have as many universally scoped addresses as they need, and network address translation is not required.

Introduction

By the early 1990s, it was clear that the change to a classless network introduced a decade earlier was not enough to prevent the IPv4 address exhaustion and that further changes to IPv4 were needed.[1] By the winter of 1992, several proposed systems were being circulated and by the fall of 1993, the IETF announced a call for white papers (RFC 1550) and the creation of the "IPng Area" of working groups.[1][2]

IPng was adopted by the Internet Engineering Task Force on July 25, 1994 with the formation of several "IP Next Generation" (IPng) working groups.[1] By 1996, a series of RFCs were released defining IPv6, starting with RFC 2460. (Incidentally, IPv5 was not a successor to IPv4, but an experimental flow-oriented streaming protocol intended to support video and audio.)

It is expected that IPv4 will be supported alongside IPv6 for the foreseeable future. However, IPv4-only clients/servers will not be able to communicate directly with IPv6 clients/servers, and will require service-specific intermediate servers or NAT-PT protocol-translation servers. Free Ipv4 adresseses will exhaust around 2010, which is within the depreciation time of equipment currently being acquired.

Features of IPv6

To a great extent, IPv6 is a conservative extension of IPv4. Most transport- and application-layer protocols need little or no change to work over IPv6; exceptions are applications protocols that embed network-layer addresses (such as FTP or NTPv3).

Applications, however, usually need small changes and a recompile in order to run over IPv6.

Larger address space

The main feature of IPv6 that is driving adoption today is the larger address space: addresses in IPv6 are 128 bits long versus 32 bits in IPv4.

The larger address space avoids the potential exhaustion of the IPv4 address space without the need for network address translation and other devices that break the end-to-end nature of Internet traffic. It also makes administration of medium and large networks simpler, by avoiding the need for complex subnetting schemes. Subnetting will, ideally, revert to its purpose of logical segmentation of an IP network for optimal routing and access.

The drawback of the large address size is that IPv6 carries some bandwidth overhead over IPv4, which may hurt regions where bandwidth is limited (header compression can sometimes be used to alleviate this problem). The address size also lacks the immediate memorability of the more familiar, shorter IPv4 address.

Stateless autoconfiguration of hosts

IPv6 hosts can be configured automatically when connected to a routed IPv6 network. When first connected to a network, a host sends a link-local multicast (broadcast) request for its configuration parameters; if configured suitably, routers respond to such a request with a *router advertisement* packet that contains network-layer configuration parameters.

If IPv6 autoconfiguration is not suitable, a host can use stateful autoconfiguration (DHCPv6) or be configured manually.

Stateless autoconfiguration is only suitable for hosts: routers must be configured manually or by other means.

Multicast

Multicast is part of the base protocol suite in IPv6. This is in opposition to IPv4, where multicast is optional.

Most environments do not currently have their network infrastructures configured to route multicast; that is — the link-scoped aspect of multicast will work but the site-scope, organization-scope and global-scope multicast will not be routed.

IPv6 does not have a link-local broadcast facility; the same effect can be achieved by multicasting to the all-hosts group (FF02::1).

The m6bone is catering for deployment of a global IPv6 Multicast network.

Jumbograms

In IPv4, packets are limited to 64 KiB of payload. When used between capable communication partners and on communication links with a maximum transmission unit larger than 65,576 octets, IPv6 has optional support for packets over this limit, referred to as jumbograms which can be as large as 4 GiB. The use of jumbograms may improve performance over high-MTU networks.

Network-layer security

IPsec, the protocol for IP network-layer encryption and authentication, is an integral part of the base protocol suite in IPv6; this is unlike IPv4, where it is optional (but usually implemented). IPsec, however, is not widely deployed except for securing traffic between IPv6 BGP routers.

Mobility

Unlike mobile IPv4, Mobile IPv6 (MIPv6) avoids triangular routing and is therefore as efficient as normal IPv6. This advantage is mostly hypothetical, as neither MIP nor MIPv6 are widely deployed today.

Deployment status

As of December 2005, IPv6 accounts for a tiny percentage of the live addresses in the publicly-accessible Internet, which is still dominated by IPv4. The adoption of IPv6 has been slowed by the introduction of classless inter-domain routing (CIDR) and network address translation (NAT), each of which has partially alleviated the impact of address space exhaustion. Estimates as to when the pool of available IPv4 addresses will be exhausted vary — in 2003, Paul Wilson (director of APNIC) stated that, based on then-current rates of deployment, the available space would last until 2023,[3] while in September 2005 a report by Cisco Systems that the pool of available addresses would be exhausted in as little as 4–5 years.[4] As of November 2006, a regularly updated report projected that the IANA pool of unallocated addresses would be exhausted in May 2011, with the various Regional Internet Registries using up their allocations from IANA in August 2012.[5] This report also argues that, if assigned but unused addresses were reclaimed and used to meet continuing demand, allocation of IPv4 addresses could continue until 2024. The U.S. Government has specified that the network backbones of all federal agencies must deploy IPv6 by 2008.[6] But there are two specific challenges to this requirement. 1) There is no special federal funding available for IPv6 transitions. Thus agencies are expected to make the migration via their ongoing equipment purchases and network updates. Most agencies now have their transition plan in place, but surveys have noted that many are lagging when it comes to making that transition a reality. [7]. 2) Agency IT budgets are tight at the moment, especially since the current 2007 IT Budget has been stalled, thanks to the Continuing Resolution.

Meanwhile China is planning to get a head start implementing IPv6 with their 5 year plan for the China Next Generation Internet.

With the notable exceptions of stateless autoconfiguration, more flexible addressing and Secure Neighbor Discovery (SEND), many of the features of IPv6 have been ported to IPv4 in a more or less elegant manner. Thus IPv6 deployment is primarily driven by address space exhaustion.

Addressing

128-bit length

The primary change from IPv4 to IPv6 is the length of network addresses. IPv6 addresses are 128 bits long (as defined by RFC 4291), whereas IPv4 addresses are 32 bits; where the IPv4 address space contains roughly 4 billion addresses, IPv6 has enough room for 3.4×10^{38} unique addresses.

IPv6 addresses are typically composed of two logical parts: a 64-bit (sub-)network prefix, and a 64-bit host part, which is either automatically generated from the interface's MAC address or assigned sequentially. Because the globally unique MAC addresses offer an opportunity to track user equipment, and so users, across time and IPv6 address changes, RFC 3041 was developed to reduce the prospect of user identity being permanently tied to an IPv6 address, thus restoring some of the possibilities of anonymity existing at IPv4. RFC 3041 specifies a mechanism by which time-varying random bit strings can be used as interface circuit identifiers, replacing unchanging and traceable MAC addresses.

Notation

IPv6 addresses are normally written as eight groups of four hexadecimal digits. For example, 2001:0db8:85a3:08d3:1319:8a2e:0370:7334 is a valid IPv6 address.

If a four-digit group is 0000, the zeros may be omitted and replaced with two colons(::). For example, 2001:0db8:0000:0000:0000:0000:1428:57ab can be shortened as 2001:0db8::1428:57ab. Following this rule, any number of consecutive 0000 groups may be reduced to two colons, as long as there is only one double colon used in an address. Leading zeros in a group can also be omitted. Thus, the addresses below are all valid and equivalent:

2001:0db8:0000:0000:0000:0000:1428:57ab
2001:0db8:0000:0000:0000::1428:57ab
2001:0db8:0:0:0:0:1428:57ab
2001:0db8:0:0::1428:57ab
2001:0db8::1428:57ab
2001:db8::1428:57ab

Having more than one double-colon abbreviation in an address is invalid, as it would make the notation ambiguous.

A sequence of 4 bytes at the end of an IPv6 address can also be written in decimal, using dots as separators. This notation is often used with compatibility addresses (see below). Thus, ::ffff:1.2.3.4 is the same address as ::ffff:0102:0304, and ::ffff:15.16.18.31 is the same address as ::ffff:0f10:121f.

Additional information can be found in RFC 4291 - IP Version 6 Addressing Architecture.

Literal IPv6 Addresses in URLs

In a URL the IPv6-Address is enclosed in brackets. Example:

http://[2001:0db8:85a3:08d3:1319:8a2e:0370:7344]/

This notation allows parsing a URL without confusing the IPv6 address and port number:

http://[2001:0db8:85a3:08d3:1319:8a2e:0370:7344]:443/

Additional information can be found in "RFC 2732 - Format for Literal IPv6 Addresses in URL's" and "RFC 3986 - Uniform Resource Identifier (URI): Generic Syntax"

Network notation

IPv6 networks are written using CIDR notation.

An IPv6 network (or subnet) is a contiguous group of IPv6 addresses the size of which must be a power of two; the initial bits of addresses, which are identical for all hosts in the network, are called the network's prefix.

- A network is denoted by the first address in the network and the size in bits of the prefix (in decimal), separated with a slash. For example, 2001:0db8:1234::/48 stands for the network with addresses

2001:0db8:1234:0000:0000:0000:0000:0000 through 2001:0db8:1234:FFFF:FFFF:FFFF:FFFF:FFFF

Because a single host can be seen as a network with a 128-bit prefix, you will sometimes see host addresses written followed with /128.

Kinds of IPv6 addresses

IPv6 addresses are divided into 3 categories [8] :

- Unicast Addresses
- Multicast Addresses
- Anycast Addresses

A Unicast address defines a single interface. It identifies a single network interface A packet sent to a unicast address is delivered to that specific computer.

Multicast addresses are used to define a set of interfaces that typically belong to different nodes instead of just one. When a packet is sent to a multicast address, the protocol delivers the packet to all interfaces identified by that address. Multicast addresses begin with the prefix FF00::/8, and their second octet identifies the addresses *scope*, i.e. the range over which the multicast address is propagated. Commonly used scopes include link-local (2), site-local (5) and global (E).

Anycast addresses, are also assigned to more than one interface, belonging to different nodes. However, a packet sent to an anycast address is delivered to just one of the member interfaces, typically the “nearest” according to the routing protocol’s idea of distance. Anycast addresses cannot be identified easily: they have the structure of normal unicast addresses, and differ only by being injected into the routing protocol at multiple points in the network.

Special addresses

There are a number of addresses with special meaning in IPv6:

- ::/128 — the address with all zeros is an unspecified address, and is to be used only in software.
- ::1/128 — the loopback address is a localhost address. If an application in a host sends packets to this address, the IPv6 stack will loop these packets back to the same host (corresponding to 127.0.0.1 in IPv4).
- ::/96 — the zero prefix was used for IPv4-compatible addresses; it is now obsolete.
- ::ffff:0:0/96 — this prefix is used for IPv4 mapped addresses (see *Transition mechanisms* below).
- 2001:db8::/32 — this prefix is used in documentation (RFC 3849). Anywhere where an example IPv6 address is given, addresses from this prefix should be used.
- fc00::/7 — Unique Local Addresses (ULA) are routable only within a set of cooperating sites. They were defined in RFC 4193 as a replacement for site-local addresses (see below). The addresses include a 40-bit pseudorandom number that minimizes the risk of conflicts if sites merge or packets somehow leak out. This address space is split into two parts:
    - fc00::/8 — - ULA Central, currently not used as the draft is expired.
    - fd00::/8 — - ULA, as per RFC 4193, Generator and unofficial registry.
- fe80::/64 — The link-local prefix specifies that the address only is valid in the local physical link. This is analogous to the Autoconfiguration IP address 169.254.x.x in IPv4.

- fec0::/10 — The site-local prefix specifies that the address is valid only inside the local organisation. Its use has been deprecated in September 2004 by RFC 3879 and systems must not support this special type of address.
- ff00::/8 — The multicast prefix is used for multicast addresses[9] as defined by in "IP Version 6 Addressing Architecture" (RFC 4291).

There are no address ranges reserved for broadcast in IPv6 — applications use multicast to the *all-hosts* group instead. IANA maintains the official list of the IPv6 address space. Global unicast assignments can be found at the various RIR's or at the GRH DFP pages.

IPv6 packet

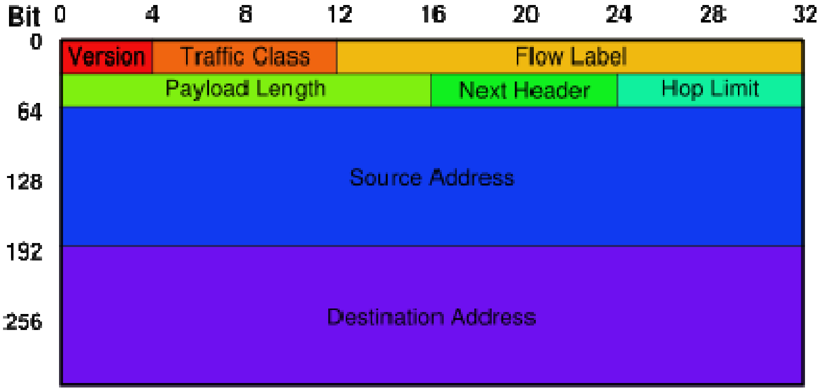

The structure of an IPv6 packet header.

The IPv6 packet is composed of two main parts: the header and the payload.

The header is in the first 40 octets/bytes of the packet and contains both source and destination addresses (128 bits each), as well as the version (4-bit IP version), traffic class (8 bits, Packet Priority), flow label (20 bits, QoS management), payload length in bytes (16 bits), next header (8 bits), and hop limit (8 bits, time to live). The payload can be up to 64KiB in size in standard mode, or larger with a "jumbo payload" option.

Fragmentation is handled only in the sending host in IPv6: routers never fragment a packet, and hosts are expected to use PMTU discovery.

The *protocol* field of IPv4 is replaced with a *Next Header* field. This field usually specifies the transport layer protocol used by a packet's payload.

In the presence of options, however, the Next Header field specifies the presence of an extra *options* header, which then follows the IPv6 header; the payload's protocol itself is specified in a field of the options header. This insertion of an extra header to carry options is analogous to the handling of AH and ESP in IPsec for both IPv4 and IPv6.

IPv6 and the Domain Name System

IPv6 addresses are represented in the Domain Name System by *AAAA records* (so-called quad-A records) for forward lookups; reverse lookups take place under ip6.arpa (previously ip6.int), where address space is delegated on nibble boundaries. This scheme, which is a straightforward adaptation of the familiar A record and *in-addr.arpa* schemes, is defined in RFC 3596.

The AAAA scheme was one of two proposals at the time the IPv6 architecture was being designed. The other proposal, designed to facilitate network renumbering, would have had *A6 records* for the forward lookup and a number of other innovations such as *bit-string labels* and *DNAME records*. It is defined in the experimental RFC 2874 and its references (with further discussion of the pros and cons of both schemes in RFC 3364).

RFC 3484 specifies how applications should select an IPv6 or IPv4 address for use, including addresses retrieved from DNS.

AAAA record fields

NAME	Domain name
TYPE	AAAA (28)
CLASS	Internet (1)
TTL	Time to live in seconds
RDLENGTH	Length of RDATA field
RDATA	String form of the IPV6 address as described in RFC 3513

IPv6 and DNS RFCs

- DNS Extensions to support IP version 6 - RFC 1886
- DNS Extensions to Support IPv6 Address Aggregation and Renumbering - RFC 2874
- Tradeoffs in Domain Name System (DNS) Support for Internet Protocol version 6 (IPv6) - RFC 3364
- Default Address Selection for Internet Protocol version 6 (IPv6) - RFC 3484
- Internet Protocol Version 6 (IPv6) Addressing Architecture - RFC 3513
- DNS Extensions to Support IP Version 6 (Obsoletes 1886 and 3152) - RFC 3596

IPv6 scope

IPv6 defines 3 unicast address scopes: global, site, and link.. Site-local addresses are non-link-local addresses that are valid within the scope of an administratively-defined site and cannot be exported beyond it.

Site-local addresses are deprecated by RFC 3879. Note that this does not deprecate other site-scoped address types (e.g. site-scoped multicast).

Companion IPv6 specifications further define that only link-local addresses can be used when generating ICMP Redirect Messages [ND] and as next-hop addresses in most routing protocols.

These restrictions do imply that an IPv6 router must have a link-local next-hop address for all directly connected routes (routes for which the given router and the next-hop router share a common subnet prefix).

IPv6 deployment

In February 1999, The IPv6 Forum was founded by the IETF Deployment WG to drive deployment worldwide creating by now over 30 IPv6 Country Fora and IPv6 Task Forces [10]. On 20 July 2004 ICANN announced[11] that the root DNS servers for the Internet had been modified to support both IPv6 and IPv4.

A global view into the IPv6 routing tables, which displays also which ISPs are already deploying IPv6, can be found by looking at the SixXS Ghost Route Hunter pages: these pages display a list of all allocated IPv6 prefixes and give colors to the ones that are actually being announced in BGP. When a prefix is announced, that means that the ISP at least can receive IPv6 packets for their prefix. They might then actually also offer IPv6 services, maybe even to end users/sites directly.

ISPs that provide IPv6 connectivity to their customers can be found in the Where can I get native IPv6 FAQ.

The mandate by the United States Government to move to an IPv6 platform for all civilian and defense vendors by summer 2008 will greatly boost deployment. The awarding of over $150 billion in contracts in spring of 2007 by the General Services Administration will in itself come close to the total amount spent on the Y2K upgrade of the previous decade, and total cost will swell far beyond that, to as much as $500 billion.[12]

Transition mechanisms

Until IPv6 completely supplants IPv4, which is not likely to happen in the foreseeable future, a number of so-called *transition mechanisms* are needed to enable IPv6-only hosts to reach IPv4 services and to allow isolated IPv6 hosts and networks to reach the IPv6 Internet over the IPv4 infrastructure. [13] contains an overview of the below mentioned transition mechanisms.

Dual stack

Since IPv6 is a conservative extension of IPv4, it is relatively easy to write a network stack that supports both IPv4 and IPv6 while sharing most of the code. Such an implementation is called a *dual stack*, and a host implementing a dual stack is called a *dual-stack host*. This approach is described in RFC 4213.

Most current implementations of IPv6 use a dual-stack. Some early experimental implementations used independent IPv4 and IPv6 stacks. There are no known implementations that implement IPv6 only.

Tunneling

In order to reach the IPv6 Internet, an isolated host or network must be able to use the existing IPv4 infrastructure to carry IPv6 packets. This is done using a technique somewhat misleadingly known as *tunnelling* which consists in encapsulating IPv6 packets within IPv4, in effect using IPv4 as a link layer for IPv6.

IPv6 packets can be directly encapsulated within IPv4 packets using protocol number 41. They can also be encapsulated within UDP packets e.g. in order to cross a router or NAT device that blocks protocol 41 traffic. They can of course also use generic encapsulation schemes, such as AYIYA or GRE.

Automatic tunneling

Automatic tunneling refers to a technique where the tunnel endpoints are automatically determined by the routing infrastructure. The recommended technique for automatic tunneling is 6to4[14] tunneling, which uses protocol 41 encapsulation. Tunnel endpoints are determined by using a well-known IPv4 anycast address on the remote side, and embedding IPv4 address information within IPv6 addresses on the local side. 6to4 is widely deployed today.

Teredo [15] is an automatic tunneling technique that uses UDP encapsulation and is claimed to be able to cross multiple NAT boxes. Teredo is not widely deployed today, but an experimental version of Teredo is installed with the Windows XP SP2 IPv6 stack. IPv6, 6to4 and Teredo are enabled by default in Windows Vista [16].

Configured tunneling

Configured tunneling is a technique where the tunnel endpoints are configured explicitly, either by a human operator or by an automatic service known as a Tunnel Broker[17]. Configured tunneling is usually more deterministic and easier to debug than automatic tunneling, and is therefore recommended for large, well-administered networks.

Configured tunneling typically uses either protocol 41 (recommended) or raw UDP encapsulation.

Proxying and translation

When an IPv6-only host needs to access an IPv4-only service (for example a web server), some form of translation is necessary. The one form of translation that actually works is the use of a dual-stack application-layer proxy, for example a web proxy.

Techniques for application-agnostic translation at the lower layers have also been proposed, but they have been found to be too unreliable in practice due to the wide range of functionality required by common application-layer protocols, and are commonly considered to be obsolete. See for example SIIT[18], NAT-PT[19], TCP-UDP Relay[20], Socks-based Gateway[21], Bump-in-the-Stack or Bump-in-the-API[22].

Major IPv6 announcements and availability

- ICANN announced on 20 July 2004 that the IPv6 AAAA records for the Japan (.jp) and Korea (.kr) country code Top Level Domain (ccTLD) nameservers became visible in the DNS root server zone files with serial number 2004072000. The IPv6 records for France (.fr) were added a little later. This made IPv6 operational in a public fashion.
- Apple Mac OS X v10.3 "Panther" (2003) has IPv6 supported and enabled by default.[23]
- Microsoft Research[24] first released an experimental IPv6 stack in 1998. This support is not intended for use in a production environment.
- Microsoft Windows NT 4.0 and Windows 2000 SP1 had limited IPv6 support for research and testing since at least 2002.
- Microsoft Windows XP (2001) had IPv6 support for developmental purposes. In Windows XP SP1 (2002) and Windows Server 2003, IPv6 is included as a core networking technology, suitable for commercial deployment.[25]
- Microsoft Windows Vista (2007) has IPv6 supported and enabled by default.[25]
- Production-quality BSD support for IPv6 has been generally available since early to mid-2000 in FreeBSD, OpenBSD, and NetBSD via the KAME project[26].

- Linux support has been available since version 2.1.8, released in 1996. As of kernel 2.6.10, the Linux IPv6 stack was approved by the IPv6 Forum in the IPv6 Ready Logo Phase-1 Program. Development still continues on improving the stack.[27]
- In the end of 1997 IBM's AIX 4.3 was the first commercial platform that supported IPv6 [28][29]
- Apple's AirPort Extreme 802.11n base station is an IPv6 gateway in its default configuration. It uses 6to4 tunneling and can optionally route through a manually configured IPv4 tunnel.[30]
- Sun Solaris has IPv6 support since version 8 [31]

http://en.wikipedia.org/wiki/IPv6