




TECHNOLOGY BRIEF

IPv6 Technology Brief



Internet Scaling Problems

Executive Summary

The Roadmap to IPv6

To deal with the addressing issues of the widely used IP protocol, the IETF has proposed a new routing protocol called IPv6. Development of the IPv6 standard has been on-going for several years, with a number of large networks fully operational, and several PC applications implemented. Although the technical issues are well understood, large-scale deployment is dependent on a seamless transition from IPv4. This technology brief discusses the technology behind IPv6, its current deployment, and the path to full implementation.

CONTENTS

Internet Scaling Problems	1
Internet Protocol Addressing.....	2
IPv6 Addressing	2
Additional IPv6 Features	3
Routing Protocols	4
IPv6 Deployment	4
Transition Strategy	5
Network Designs	6
The Future of IPv6	6
3Com IPv6 Strategy	6
Glossary	7

Over the past few years, the Internet has experienced two major scaling issues as it has struggled to provide continuous and uninterrupted growth:

- The eventual exhaustion of IP version 4 (IPv4) address space
- The need to route traffic between the ever-increasing number of networks that make up the Internet

The first problem concerns the eventual depletion of the IP address space. IPv4 defines a 32-bit address which means that there are only 2^{32} (4,294,967,296) IPv4 addresses available. As the Internet continues to grow, this finite number of IP addresses will eventually be exhausted.


The address shortage problem is aggravated by the fact that portions of the IP address space have not been efficiently allocated. Also, the traditional model of classful addressing does not allow the address space to be used to its maximum potential. The Address Lifetime Expectancy (ALE) Working Group of the Internet Engineering Task Force (IETF) has expressed concerns that if the current address allocation policies are not modified, the Internet will experience a near to medium term exhaustion of its unallocated address pool. If the Internet's address supply problem is not solved, new users may be unable to connect to the global Internet.

More than half of all possible IPv4 addresses have been assigned to ISPs, corporations, and government agencies, but only an estimated 69 million addresses are actually in use.

The second problem is the result of the rapid growth in the size of the Internet routing tables. Internet backbone routers are required to maintain complete routing information for the Internet. Over recent years, routing tables have experienced exponential growth as increasing numbers of organizations connect to the Internet. In December 1990 there were 2,190 routes, in December 1995 there were more than 30,000 routes, and in December 2000 more than 100,000 routes.

Unfortunately, simply installing more router memory and increasing the size of the routing tables cannot solve the routing problem. Other factors related to the capacity problem include the growing demand for CPU horsepower to compute routing table/topology changes, the increasingly dynamic nature of WWW connections and their effect on router forwarding caches, and the sheer volume of information that needs to be managed by people and machines. If the number of entries in the global routing table is allowed to increase without bounds, core routers will be forced to drop routes and portions of the Internet will become unreachable.

Figure 1. Network Number Growth



The long-term solution to these problems is the widespread deployment of IP Next Generation (initially called IPng and now widely known as IPv6). Currently, IPv6 is implemented on a number of networks, the most advanced of which is the 6bone network. An informal collaborative project covering North America, Europe, and Japan, 6bone supports the routing of IPv6 packets, since that function has not yet been integrated into many production routers. Until

IPv6 can be deployed worldwide, IPv4 patches will need to be used and modified to continue to provide the universal connectivity users have come to expect. As discussed below, a number of solutions (such as NAT or CIDR) have delayed the depletion of IPv4 addresses. However, these are temporary solutions, and they have serious limitations, such as an inability to scale and lower overall performance.

Internet Protocol Addressing

IP uses a 32-bit address, using the dotted decimal representation 127.127.127.127. This allows up to 4 billion IP entities. Although 4 billion may seem a very large number, the classful addressing methodology is not very flexible and many addresses may remain unused. With the advent of the web, the demand for IP addresses has increased dramatically. Now, with IP phones and PDAs using

more addresses, there is a risk of addresses running out.

Several techniques have been used to improve the use of IPv4 addresses. These include Variable Length Subnet Masks (VLSM) introduced in 1987 by RFC1009, classless inter-domain routing (CIDR) introduced in 1993 by RFC 1517 to 1520, and Network Address Translation (NAT).

IPv6 Addressing

With the continued growth of the Internet and its possible extensions to more devices, such as televisions, toasters, and coffee makers, all IPv4 solutions proposed for scaling the Internet address space will only delay the inevitable. The IETF has proposed a comprehensive set of specifications to define what is commonly known as the next-generation IP protocol ("IPng" or IPv6").

IPv6 increases the address size from 32 bits to 128 bits, supporting up to 3.4×10^{38} nodes. This is enough to reach every grain of sand in our galaxy. It is represented using hexadecimal values separated by colons using the format X:X:X:X:X:X:X:, where each X refers to a four-digit hexadecimal integer (16 bits each).

One such address could be
BA98:7654:3210:FEDC:BA98:7654:3210.

It is acceptable to compress the representation by using double colons to indicate groups of :0000:, and to leave out all zeroes to the left of any four-hexadecimal group. For example, FF01:0:0:0:0:0:0043 can be abbreviated to FF01::43. Finally, in mixed IPv4/IPv6 you can use a representation such as X:X:X:X:X:X:X:D.D.D.D where X represents a four-digit hexadecimal value as above and D represents the standard IPv4 decimal values. One such example would be ::FFFF:129.144.52.38.

IPv6 defines three types of addresses: Unicast, Anycast, and Multicast. Please note that there are no Broadcast addresses, thus alleviating the risk of broadcast storms. In addition, a single interface can be assigned multiple IPv6 addresses of any type.

Additional IPv6 Features

Although IPv6 appears to have functionality similar to CIDRs, there are many additional features that were not included in IPv4. These features make IPv6 much more robust and convenient. These changes include a streamlined IPv6 header, stateless configuration, built-in security, better QoS, and increased real-time performance.

Streamlined IPv6 Header

The IPv6 header has a new format that is designed to keep header overhead to a minimum. This is achieved by moving both nonessential fields and option fields to extension headers that are placed after the IPv6 header. The streamlined IPv6 header provides more efficient processing at intermediate routers.

Stateless and Stateful Configuration

IPv6 supports both stateful and stateless address configurations. IPv6 will work with or without a DHCP server. With stateless address configuration, hosts on a link automatically configure themselves with IPv6 addresses for the link (called link-local addresses) and with addresses derived from prefixes advertised by local routers. Even in the absence of a router, hosts on the same link can automatically configure themselves with link-local addresses and communicate without manual configuration.

Built-in Security

Support for IPSec is an IPv6 protocol suite requirement. This requirement provides a standards-based solution for network security and promotes interoperability between different IPv6 implementations.

Better Support for QoS

New fields in the IPv6 header define how traffic is handled and identified. Traffic identification using a flow label field in the IPv6 header allows routers to identify and provide special handling for packets belonging to a flow, which is a series of packets between a source and destination. Because the traffic is identified in the IPv6 header, support for QoS can be achieved even when the packet payload is encrypted through IPSec.

Real-Time Performance

IPv6 offers a packet prioritization feature that provides the real-time and near real-time applications an improved response time. Consequently, IPv6 will become the protocol of choice for those applications.

Improved Multicast Support

IPv6 does not support broadcasts. All functions that required broadcasts are now handled by multicasts. While IPv4 used a limited range of addresses for multicasts, IPv6 has a much broader range, and allows restriction of the scope of each multicast address. Multicast is an intrinsic part of the IPv6 protocol. It employs the Multicast Listener Discovery (MLD) protocol defined by RFC 2710. MLD uses ICMPv6 to enable IPv6 routers to discover the presence of multicast listeners. This is usable by multicast routing protocols such as PIM to ensure that multicast packets are delivered to all links where there are interested listeners.

Table 1: Comparison Between IPv4 and IPv6

Feature	IPv4	IPv6
Address Bits	32	128
Configuration	DHCP	Auto/DHCPv6
QoS	DiffServ/IntServ	DiffServ/IntServ
Security	IPSec (optional)	IPSec (mandatory)
Multicast	IGMP/PIM/MBGP	MLD/PIM/MBGP (Scope ID)

Routing Protocols

Routing Information Protocol (RIP)

A new version of RIP called RIPng was introduced by RFC 2080. It works essentially in the same way as RIPv2, but uses IPv6 addresses for RIP update messages.

Open Shortest Path First (OSPF)

Changes were made to OSPFv3 to support IPv6, as defined in RFC 2740. In addition to supporting the 128-bit IPv6 addresses, it provides per-link processing rather than the usual per-node processing, and allows for multiple instances per link.

IS-IS

IPv6 support for IS-IS is still an IETF draft described under “draft-ietf-isis-ipv6-02”. It applies to the IP or dual environment modes.

Border Gateway Protocol (BGP-4)

BGP is an Exterior Gateway Protocol (EGP), rather than an Interior Gateway Protocol (IGP) such as the three routing protocols described above. The extensions to support IPv6 are described in RFC 2545.

Internet Control Message Protocol (ICMP)

To handle the information exchanges between IPv6 routers, ICMPv6 was defined in RFC 2463

IPv6 Deployment

Initially, the first IPv6 networks were established for testing purposes only. A number of large IPv6 networks are now operational—6bone is probably the best known. 6bone has been in operation since 1996 when it began linking IPv6 laboratories in France, Denmark, and Japan. 6bone is actually layered over IPv4, using tunnels. The backbone itself is using the BGP4+ protocol. Regions that have shown an interest in IPv6 are mainly those with high population density and high technological development. They include North America, Europe, India, China, and Japan. From an application point of view, the early adopters have been carriers with large data networks.

The U.S. DoD has mandated that all its networks must be able to support IPv6. This means that all systems

acquired after January 9, 2004 must be capable of operating in an IPv6 network, while maintaining the ability to function in today’s IPv4 networks. Although specific near-term IPv6 pilots are underway, no implementation of IPv6 is permitted on operational networks at this time. The DoD expects that the transition will take place from FY05 to FY07, with the goal being to complete the transition in FY08. A few large corporations have expressed interest and are testing IPv6 in limited applications.

Operating Systems support for IPv6 is now available from Microsoft, Sun, Hewlett-Packard/Compaq, IBM, Linux, Apple, and others. Microsoft has released an IPv6-enabled file- and print-sharing application in Microsoft Windows 2003 Server.

Transition Strategy

The main obstacle to the wide deployment of IPv6 is the problem of moving from a very large installed base of IPv4 networks and applications.

Dual Stack

One way to make this transition is the dual stack approach, which can be used in both the network nodes (workstations or servers) and the routers. In order to work effectively, the dual stack must be implemented in all the routers in a network. This solution works by using two addressing schemes in parallel. There is no communication between the IPv4 world and the IPv6 world—applications must be able to support both modes. The dual stack approach is used frequently today, but requires that all network resources have enough processing power and memory to support two different IP stacks. And dual management is also required.

Tunneling


The tunneling solution encapsulates one protocol type in another protocol. This requires at least dual stacks at each end of the tunnel. Tunneling can be used to pass IPv6 over an existing IPv4 network, or later to pass IPv4 traffic over IPv6 networks. Please note that tunneling can also be used to pass other protocols (IPX or AppleTalk) over IPv6. The routers involved in the tunneling must be able to map the end addresses to each other. Because of the complexity, single tunnels are acceptable, but large-scale deployment is usually too difficult. Manually configured IPv6 tunnels are described by RFC 2893. It is also possible to use IPv6 over IPv4 GRE tunnels. Other mechanisms include 6to4 tunnels, described by RFC 3056, and ISATAP, described by draft-ietf-ngtrans-isatap-04.

As an alternative, a number of service providers or large enterprises choose to tunnel IPv6 over MPLS. Although the provider edge routers need to be able to map the IPv6 addresses to generate the appropriate labels, this method has the advantage of providing a fast and QoS-enabled transmission path.

Translation

According to the IETF, translation is the tool of last resort. These schemes are inherently complex. They need to be used any time dual stack is not an option, such as in a case where there are simple workstations that cannot be upgraded, or nodes that are accessing the network through a NAT port. RFC 2765, Stateless IP/ICMP Translation Algorithm (SIIT), describes such a process, which can be used in separate applications such as NAT-PT (Network Address Translation-Protocol Translation) or BIS (Bump-in-the-Stack).

Figure 2: Dual Stack Mechanism



Network Designs

Organizations that are considering the impact of IPv6 on their networks will likely fall into one of the following categories:

Corporate Network Running IPv4 Using an IPv6 Service Provider

This is probably the simplest case of all. Typically, the organization would deploy WAN routers using dual stack at the WAN interface.

Corporate Network Running Both IPv4 and IPv6 Applications

This is the situation where dual stack for all routers and applications will make the most sense. For those nodes that cannot support dual stack, some translation method may be required. The core network will probably use IPv6, or possibly MPLS.

Corporate Network Running IPv6 That Must Communicate with IPv4 Networks and Applications

This would be critical at the start of the deployment phase, when there are still IPv4-only applications, which have not yet been upgraded. In this situation, either dual-stack clients are recommended, so that the client can use IPv4 for the IPv4-only applications, or a translation method may be needed.

The Future of IPv6

IPv6 has been in the making for almost eight years now. Support is readily available from both applications and network providers. Still, the sheer number of applications currently running over IPv4 is such that the general expectation is that the full conversion may take another decade. Just consider all the web URLs, older PCs, and so forth that will need

updating. Large corporations, service providers, and governmental agencies will be the first to complete the transition. Smaller organizations will lag by many years, unless a “killer application” requiring IPv6 emerges that is appealing to the wider public. Such an application might be as simple as the wide distribution of music to networked MP3 players.

3Com IPv6 Strategy

3Com has been an early supporter of the IPv6 initiative and is a founding member of the IPv6 forum. We fully intend to provide standard-compliant, simple, and cost-effective implementations of IPv6. Following the path of current IPv6 deployment, we will initially introduce IPv6 dual stacks in our WAN routers. The next step will

include our Core Modular product line for large Enterprise. At that time, we will offer tunneling capabilities in addition to the dual stack. Over time all 3Com products will converge to IPv6. Because we understand the needs of our users, we will provide effective tools to help the transition.

Glossary

BGP	Border Gateway Protocol
CIDR	Classless Inter Domain Routing
DHCP	Dynamic Host Configuration Protocol
DoD	Department of Defense
EGP	Exterior Gateway Protocol
GRE	Generic Routing Encapsulation
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IP	Internet Protocol
IS-IS	Intermediate System to Intermediate System
MPLS	Multi Protocol Label Switching
MLD	Multicast Listener Discovery
MTU	Maximum Transmission Unit
NAT	Network Address Translation
OSI	Open Systems Interconnect
OSPF	Open Shortest Path First
PDA	Personal Data Assistant
RFC	Request For Comment
RIP	Routing Information Protocol
VLSM	Variable Length Subnet Mask
WAN	Wide Area Network