A Primer on IPv6

White Paper

Abstract

This paper discusses the evolution of Internet Protocol version 4 (IPv4) to Internet Protocol version 6 (IPv6). It includes an overview of the limitations of IPv4, IPv6 features, the driving forces behind the transition and key differences between the two protocols.

The Limitations of IPv4

The current version of Internet Protocol or IP (known as Version 4 or IPv4) has not been substantially changed in the past 25 years, a lifespan over which IPv4 has proven to be robust, easily implemented and interoperable, and for the most part scalable enough to accommodate the ever-expanding Internet. However, continued exponential growth of Internet-enabled devices and the evolving sensitivity for secure data transfer over the Internet are outstripping the practical capabilities of IPv4 and revealing its limitations:

- **Insufficient IP address space**
  With only 32-bit capacity, IPv4 addresses have become relatively scarce, forcing some organizations to use Network Address Translation (NAT) to map multiple private addresses to a single public IP address. While NAT promotes conservation of the public address space, it does not support standards-based network layer security or the correct mapping of all higher layer protocols and can create problems when connecting two organizations that use the same private address space. The continued expansion of Internet-connected devices and appliances continues to put greater and greater stress on the public IPv4 address space.

- **Address prefix allocation**
  Because of the way that IPv4 address prefixes have been and are currently allocated, Internet backbone routers are routinely required to maintain unreasonably large routing tables of over 85,000 specified routes. The current IPv4 Internet routing infrastructure is a combination of both flat and hierarchical routing.

- **Complexity of configuration**
  Most current IPv4 implementations must be either manually configured or use a stateful address configuration protocol such as Dynamic Host Configuration Protocol (DHCP). With more computers and devices using IP, there is a need for a simpler and more automatic configuration of addresses and other configuration settings that do not rely on the administration of a DHCP infrastructure.

- **Data security**
  Private communication over a public medium like the Internet requires encryption services that protect the data being sent from being viewed or modified in transit. Although an add-on standard now exists for providing security for IPv4 packets (known as Internet Protocol Security or IPsec), this standard is optional and proprietary alternatives are commonly used.

- **Quality of Service (QoS)**
  While standards for QoS exist for IPv4, no identification of packet flow for QoS handling by routers is present within the IPv4 header. Instead, real-time traffic support relies on the IPv4 Type of Service (ToS) field and the identification of the payload, typically using a UDP or TCP port. However, the IPv4 ToS field has limited functionality and payload identification using a TCP and UDP port is not possible when the IPv4 packet payload is encrypted.

A new suite of protocols and standards known as IP version 6 (IPv6) has been developed to address these limitations. Previously called IP-The Next Generation (IPng), IPv6 was intentionally designed to minimize impact on upper and lower layer protocols by standardizing packet header formation and making it easy to handle new data types without causing a negative impact on network performance.

IPv6 Features

The IPv6 protocol includes the following features:

- New standardized header format
- Larger address space
- Multicast and anycast
- Stateless address configuration
- Built-in security
- Better support for QoS
- Extensibility

The following sections discuss each of these new features in detail.

New Header Format
IPv6 introduces a more streamlined header format that reduces overhead processing on intermediate routers and speeds throughput. IPv4 headers and IPv6 headers are not interoperable and IPv6 is not backward compatible with IPv4. A host or router must use an implementation of both IPv4 and IPv6 (e.g., dual stack) in order to recognize and process both header formats. The new IPv6 header is only twice as large as the IPv4 header, even though IPv6 addresses are four times as large as IPv4 addresses.

Larger Address Space
IPv6 has 128-bit (16-byte) source and destination IP addresses, allowing, for example, each cell phone or mobile electronic device to be assigned a unique IP address. IPv4 supports $4.3x10^9$ (4.3 billion) addresses, which is incapable of furnishing one address to every living person. Remember, millions of people have multiple IP-enabled devices. With 128-bits, IPv6 can express over $3.4x10^{38}$ possible combinations or $5x10^{28}$ addresses for each of the roughly 6.5 billion people alive today.

Even though only a small number of the possible addresses are currently allocated for use by hosts, there are plenty of addresses available for future use. With such a large number of available addresses, address-conservation techniques, such as the deployment of NATs, are no longer necessary.

Notation
In order to represent larger addresses more compactly, IPv6 addresses are written in a hexadecimal notation system as opposed to the "dotted quad" system used in IPv4. As a result, IP addresses appear vastly different in IPv6.

Example:
IPv4 70.57.159.129
IPv6 2002.6688.9E8D.0000.0000.0000.0000.0001
Additional information on IPv6 notation can be found at http://www.ietf.org/rfc/rfc3513.txt.

Multicast and Anycast
Multimedia applications can take advantage of multicast: the transmission of a single datagram to multiple receivers. Multicast (both on the local link and across routers) is a requirement of IPv6, in contrast to IPv4, where multicast is optional and rarely deployed across routers.

In addition, IPv6 defines a new broadcasting method termed "anycast." Like multicast, anycast has groups of nodes that send and receive data packets; however, when a packet is sent to an anycast group, it is only delivered to one of the group members, thereby limiting the data flooding that characterizes IPv4 networks. IPv6 eliminates Broadcast packets – allowing greater use of switches instead of routers, flattening networks and improving performance at the physical level.

Stateless Address Configuration
To simplify host configuration, IPv6 supports both stateful address configuration (with DHCP) and stateless address configuration (auto-configuration without DHCP). With stateless address configuration, IPv6 hosts can configure themselves automatically. In this scenario, when first connected to a routed IPv6 network, a host sends a link-local multicast request for its configuration parameters. An IPv6 router on the network will hear this request and respond appropriately with an advertisement packet containing the address. If stateless configuration is not suitable, a host can still use stateful configuration or be configured manually, just as with IPv4 networks.

Built-in Security
Support for IPsec is an IPv6 requirement. This requirement provides a standards-based solution for network security needs 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, 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.

Extensibility
IPv6 can easily be extended for new features by adding extension headers after the IPv6 header. Unlike options in the IPv4 header, which can only support 40 bytes of options, the size of IPv6 extension headers is only constrained by the size of the IPv6 packet.

Key Differences between IPv4 and IPv6

IPv4	IPv6
Source and destination addresses are 32 bits (4 bytes) in length.	Source and destination addresses are 128 bits (16 bytes) in length.
IPsec support is optional.	IPsec support is required.
No identification of packet flow for QoS handling by routers is present within the IPv4 header.	Packet flow identification for QoS handling by routers is included in the IPv6 header using the Flow Label field.
Header includes options.	All optional data is moved to IPv6 extension headers; header length is standardized, and header overhead reduced, allowing for significantly more efficient packet handling.
Broadcast addresses are used to send traffic to all nodes on a subnet.	There are no IPv6 broadcast addresses. Instead, a link-local scope all-nodes multicast address is used, eliminating broadcast floods and allowing flatter network design.
Must be configured either manually or through DHCP.	Does not require manual configuration or DHCP. Supports stateless configuration.

Transition from IPv4 to IPv6

Keeping in mind the existence of a large installed base of devices that are not IPv6 compatible and are unlikely to be replaced, as well as the fact that IPv4 has already been deployed successfully in most existing applications, two key drivers are likely to influence a transition to global IPv6 deployment.

U.S. Government
The U.S. Government has specified that all federal agencies must support IPv6 on their networks by 2008. Already, IPv6 support is a requirement in many Department of Defense (DoD) hardware bidding specifications, but only a handful of projects will deploy it anytime soon. Projects involving mobile IP or remote data acquisition are likely to specify IPv6, as this will enable the DoD to realize its goal of flattening network architecture and simplifying infrastructure and maintenance.

Growth in Asia
The United States was the beneficiary of most of the allotted IPv4 addresses, leaving a relatively small number of available addresses for an area of the world that has the fastest growing population and several of the fastest growing economies. Asia, and in particular China, are proponents of IPv6 and the additional address capacity it delivers. Their desire for this additional address capacity is driven both by the innovation that it can accommodate (e.g., more IP-enabled devices), as well as the additional capacity to monitor information flows over the Internet.

Transition Mechanisms

Unless IPv6 completely supplants IPv4, which is not likely to happen in the foreseeable future, a number of transition mechanisms will be employed to enable IPv4-IPv6 interoperability.

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. Most current implementations of IPv6 provide a dual stack. Some early experimental implementations used independent IPv4 and IPv6 stacks.

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

Automatic Tunneling
Automatic tunneling refers to a technique where the tunnel endpoints are automatically determined by the routing infrastructure. 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.

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. Configured tunneling is usually more deterministic and easier to debug than automatic tunneling, and is therefore recommended for large, complex networks.

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 most widely supported form of translation 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.

Summary

While IPv4 has proven to have tremendous durability in an increasingly networked world, it exhibits some basic limitations that make the features of IPv6 ever more relevant. The most noteworthy of those features is the increased IP address space made possible in the IPv6 addressing scheme. As the installed base of Internet-enabled devices continues to expand worldwide and the high-growth, population-dense regions where IPv4 addresses are in short supply continue to expand, the need for the flexibility offered by IPv6 will become even more important.

91001386
B1/1115

DIGI INTERNATIONAL WORLDWIDE HQ
877-912-3444 / 952-912-3444 / www.digi.com

DIGI INTERNATIONAL FRANCE
+33-1-55-61-98-98 / www.digi.fr

DIGI INTERNATIONAL JAPAN
+81-3-5428-0261 / www.digi-intl.co.jp

DIGI INTERNATIONAL SINGAPORE
+65-6213-5380

DIGI INTERNATIONAL CHINA
+86-21-50492199 / www.digi.com.cn