Internet Protocol version 6

What is IPv6?

IPv6 is short for “Internet Protocol version 6.” IPv6 is designed to replace the longtime standard network layer protocol, Internet Protocol version 4 (IPv4). With an increasing number of devices becoming network capable (mobile phones, laptops, PDAs, TVs, etc.), the number of available IPv4 addresses will soon run out. Acknowledging the addressing limitations of IPv4, the Internet Engineering Task Force (IETF) began forming the “IP Next Generation” working group in 1994, and produced what is now titled IPv6.

Enhancements over IPv4

To alleviate the limited address space of IPv4, the IETF working group designed IPv6 addresses to be 128 bits long. IPv4 addresses are only 32 bits long. This increase in address space yields about $34x10^{37}$ IPv6 addresses, while IPv4 only has about 4.3 billion addresses. Having a large quantity of IPv6 addresses allows every current and future device to have its own address. The increased address space offers other benefits as well.

For instance, Network Address Translation (NAT) is heavily used in IPv4 networks today to mask multiple internal IP addresses behind a single, external IP address. NAT was created to combat the shortage of IP addresses in IPv4, but with the number of addresses available with IPv6, the use of NAT will no longer be required for IP address conservation. This also restores direct end-to-end network connectivity that NAT disrupts.

IPv6 provides a superior multicasting ability to that of IPv4. Multicasting allows one device to communicate to multiple devices at once without resending the same data to each. It is often used to stream real-time services such as music performances or live video conferencing. IPv6, unlike IPv4, has multicasting built in, which allows for simpler administration and routing.

Another protocol that is optional for IPv4, but is built into IPv6 is IP Security (IPSec). IPv4 does not have any built-in security features like those that IPSec provides for IPv6. IPSec provides security to IPv4 and IPv6 via authenticated and encrypted end-to-end network traffic. However, IPSec relies on the availability of end-to-end connectivity in order to provide this security. With medium and large networks utilizing NAT, the end-to-end connectivity is broken, thus preventing use of the security IPSec provides. By using IPv6 and eliminating NAT, both end-to-end connectivity and IPSec security can be restored.

Building an IPv6 address

A typical IPv6 address consists of three parts: the "global routing prefix," the "subnet identifier," and the "interface identifier." The global routing prefix is used to identify a special address, such as multicast, or an address range assigned to a site. A subnet identifier, also known as “subnet prefix” or “subnet,” is used to identify a network within a site. The global routing prefix combined with the subnet identifier is collectively referred to as the 'network prefix.'

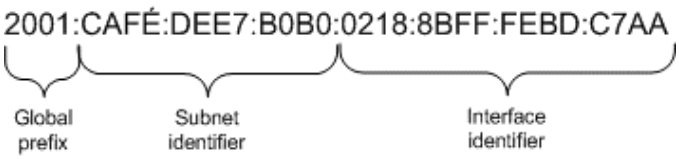

Figure 1: IPv6 address components

The interface identifier is used to identify a device interface on a network link and needs to be unique to that link. Unless manually configured, the interface identifier is automatically derived from the

48 bit MAC address of the device's network adapter. To derive a 64 bit interface identifier from the 48 bit MAC address the hexadecimal digits FFFE are inserted between the third and fourth bytes of the MAC address. A "Universal/Local" bit (the second bit of the first byte of the resulting interface identifier) is also complemented.

For example, if the starting MAC address is 00:18:8B:BD:C7:AA then:

- Hexadecimal digits FFFE are inserted between 8B and BD creating the 64 bit address 0018:8BFF:FEBD:C7AA
- The Universal/Local bit, the second bit of 00 is complimented, making it 02.

The resulting example IPv6 interface identifier is then 0218:8BFF:FEBD:C7AA.

A host can generate its own *link-local* address by pre-pending its interface identifier with a network prefix of FE80 (so for the example above the local address would be FE80::0218:8BFF:FEBD:C7AA). This *link-local* address is sufficient for communication between hosts on the same network link, but it can't be routed. *Link-local* addresses are used in Stateless Autoconfiguration and for creating local networks without the presence of a router.

New features of IPv6

IPv6 offers a variety of new features over IPv4, but two important ones stand out: Stateless Address Autoconfiguration and Router Discovery.

Stateless Address Autoconfiguration can greatly reduce the time required for network configuration. To accomplish autoconfiguration, a node (a device like a host or router) creates a tentative *link-local* address. The node then sends a Neighbor Solicitation message with its tentative *link-local* address as the target address to the *solicited-node* multicast group. This process checks whether the tentative *link-local* address is unique on the link. If autoconfiguration fails then the *link-local* address is not unique and manual intervention is required. However, if it succeeds, the *link-local* address is unique and the node assigns the *link-local* address to its interface and IP connectivity over the local link is possible.

After establishing a *link-local* address, a host does not have network connectivity beyond its local link. To gain this connectivity, the host must locate a default router and obtain a *global address* (an address that is globally unique to the entire network). This is normally done through Router Discovery.

As part of Router Discovery, a host sends a Router Solicitation message to the *all-router* multicast group. All routers on the host's network link will reply with a Router Advertisement message containing important details about the router and link. One or more routers are typically configured to advertise the appropriate network prefix. The host combines this network prefix with its interface identifier, creating a tentative *global address*. The host then sends a Neighbor Solicitation message with its tentative *global address* as the target address to the *solicited-node* multicast group. This process checks whether the *global address* is unique in the network. If autoconfiguration fails then the *global address* is not unique and manual intervention is required. However, if it succeeds, the *global address* is unique and the host assigns the *global address* to its interface and routed IP connectivity over the entire network is possible.

Coming to a network near you

In short, IPv6 is not a "stop gap" protocol whose only function is to fix the IPv4 address space problem. IPv6 has been designed to fix both the scaling and security problems of IPv4. With features such as stateless address autoconfiguration and IPSec built into the protocol, IPv6 deals with several of the problems currently encountered by network administrators and security professionals.