Migrating to an IPv6 Internet while preserving IPv4 addresses

Technology White Paper

Serge-Paul Carrasco

Abstract

The Internet is running out of addresses! Depending on how long the present pool of IPv4 addresses can be preserved, the exhaustion could begin as soon as Spring 2011.

IPv6 is the "only" solution to the exhaustion of Internet addresses. Unfortunately, there is little worldwide network deployment of IPv6. And, without any migration from IPv4 to IPv6, the Internet will run out of addresses.

IPv6 introduces a number of features and benefits to the Internet Protocol for mobility, routing, QoS and security. In the meantime, a lot of mechanisms must be implemented to preserve the present pool of IPv4 addresses.

The Internet Running Out of Addresses

The Internet is running out of addresses! Each computing device connected to the Internet must have a unique worldwide Internet address. Soon, the 4.3 billion addresses that are possible with the 32-bit packet header of the Internet Protocol version 4 (IPv4) will be exhausted! Presently, only 16.4% of those 4.3 billion addresses are available. Depending on how long the present pool of IPv4 addresses can be preserved, the exhaustion could begin as soon as Spring 2011.

The Urgent Need for More IP Addresses

A larger address space for the Internet is required first by a growing Internet population of users from Asia and South America. China itself would need half of the present 705 million available IPv4 address just for its student population! Second, the growth of new computing devices in particular mobile Internet devices for 3G networks such as smart phones (iPhone, BlackBerry...) and mobile Internet devices in cars and planes. And third, the growth of new kinds of computing appliances for home networks and for industry specific applications.
Internet addresses are managed by the Internet Assigned Numbers Authority (IANA) and allocated by the Regional Internet Registries (RIRs).

IPv6 The "Only" Solution to the Exhaustion of Internet Addresses

The first standard of IP Next Generation (IPng) later called IP version 6 (IPv6) to solve the limitation of the Internet addressing was proposed in 1995 (IETF RFC 1883). IPv6 128-bit packet header will enable 340 undecillion addresses (or 1,030 IP addresses per person on Earth)! Besides addressing, IPv6 introduced a number of new features for mobility, Internet routing, security and QoS. But a lot of those features have been retrofitted into IPv4 so the real value of IPv6 lies currently only in its address space.

Unfortunately, there is little worldwide network deployment of IPv6. And, without any migration from IPv4 to IPv6, the Internet will run out of addresses.

Government institutions in Japan, US and China have mandated the transition to IPv6. Since August 2005, the US Office of Management and Budget (OMB) government has required that US Federal Agencies must use IPv6 by June 2008. Those requirements have pushed software and equipment vendors (Microsoft, Cisco, IBM...) to provide IPv6 solutions for a long time. Even Mac (Mac OS X) and PC (Windows XP and Vista) computers can be automatically configured for IPv6. But unfortunately, the requirement to move to IPv6 from the public sector did not have any ripple effect on the private sector to adopt IPv6.

So until we are running out of IPv4 addresses, there will not be large commercial IPv6 deployments. This is unfortunate since the Internet will significantly gain in stability if step-by-step migrations from IPv4 to IPv6 will occur. Moving to IPv6 is not a small change to the Internet infrastructure!

Fundamentals of IPv4 Classful Addressing

Initially, Internet Protocol version 4 (IPv4) addresses were defined in IETF RFC 791 as a two-part object combining: a "network identifier" and a "host identifier". The network identifiers are assigned by IANA. The host identifiers are assigned by the network manager. There are five classes of addresses: A (7 bits for the network number/24 bits for the host number) B (14 for the net/16 for the hosts) C (21 for the net/8 for the hosts), D (for multicast) and E (experimental). Each address is normally represented as four decimal numbers separated by dots such as 128.88.12.7. Class A/B/C are called classful addressing.

Preserving Internet Addresses

One of the first problems encountered with IPv4 addressing was the market demand for class B. To cope with that challenge, Internet routing with classful addressing was replaced in 1993 by Classless Inter-Domain Routing (CIDR). CIDR introduces an extension called supernetting or network prefix to the classful address: 128.88.12.7/8.

Classless Addressing and Supernetting:
Internet routers use this network prefix or classless addressing instead of the full IP address to route IP packets on the Internet. The supernet can be seen as the country or area code for the telephone network. Supernets collapse a number of contiguous addresses into one by providing the pair: Network address/Network Prefix Length, where the network address is the first address in the contiguous block and the network prefix length allows interfering the full network address, through a mask comparison, in the supernet.

CIDR:
With CIDR, a single route advertisement can therefore cover a block or old-style addresses and so addresses can be assigned hierarchically. This means that large blocks are delegated to large service providers, which then can break up their allocation, keep some, and delegate smaller blocks to smaller providers.

Internet routing with classless addressing is performed with the Border Gateway Protocol (BGP) version 4. BGP is a "path vector" protocol. BGP handles routing communication based on configured policies between Autonomous System (AS), which defines "a set of networks under the same management organization". BGP routers exchange network reachability information which lists the complete path for each and every possible AS destination on the Internet between service providers.

Preserving and Managing Addresses on the LAN

Like service providers on the Internet, enterprise networks have adopted for a long time various protocols and tools to better manage the growth of their IP addresses in particular with subnets, DHCP and NAT.

Subnetting:

Subnetting for the LAN is the equivalent of supernetting for the Internet. The subnet address format is:
Network number- Subnet-Host number
The subnet can have any length and is specified by a "mask" determined by a "comparison-under-mask" operation. Enterprise networks have designed subnets to sub divide their IP networks into unequal pieces, each having their own subnet mask. The network administrator "right sizes" the addressing of each subnet with variable-length subnet mask (VLSM). VLSM like CIDR makes efficient use of an organization's assigned IP address space. It reduces as well the amount of routing information by leveraging the subnet to route packets to their final destination.

DHCP:
When an IP address has been given to a host, it can be found using the Address Resolution Protocol (ARP) that map the IP address to the equipment hardware Ethernet MAC address. To simplify adding, changing and moving IP addresses in an enterprise network, Dynamic Host Configuration Protocol (DHCP) was created. IP addresses can be assigned for a limited time or until the end station relinquishes it. DHCP supplies IP host address, subnet mask, and local gateway information in response to end-system broadcast requests.

NAT:
Network Address Translation (NAT) was designed for enterprise networks that are not fully connected to the Internet and accept having all their outgoing traffic "rewritten" by an address translation gateway. With NAT, the "partially connected" network uses "private addresses", that is addresses that are not routable in the Internet and are not "public". This enables NAT to be used for securing the entrance to private networks.

IPv6 Features & Benefits

Between the Internet Protocol version 4 (IPv4) and version 6 (IPv6), IPv5 defined in RFC 1819 was intended for streaming traffic but has never widely been used. IPv6 was designed mainly to address space exhaustion and Internet backbone routing issues. IPv6 can be seen as a conservative extension to IPv4. Most application-layer protocols do not need major changes to support IPv6 except if they embed network addresses.

The major specifications of IPv6 can be found at the IETF IPv6 working group and include:

IPv6 Header Improvements:
The IPv6 header format has been simplified for faster processing. It has a fixed size of 40 bytes and each field is 64-bits taking therefore advantage of the current generation of 64-bits processors. Support for encoding options has been improved by daisy-chaining an extension to the header.

IPv6 Addressing:
As described in detail earlier, IPv6 introduces 128-bits addresses; half of the bits are used to identify the network and the other half to identify the host interface. Addresses are written as

eight 16-bit integers. Each integer is represented by four hexadecimal digits as in:

FEDC:BA98:7654:3210:FEDC:BA98:7654:3210

IPv6 introduces a new type of address: anycast address which gives an address to a group of hosts but for which data is sent to only one of them. Broadcast addresses are replaced by multicast addresses and a single host interface can be assigned multiple addresses of any type (unicast, multicast, anycast). Addresses can be manually configured, assigned via DHCP or auto-configured/auto-generated.

IPv6 Auto-Configuration of Addresses:
IPv6 enables devices to have their addresses auto
configured. Hosts can get their network identification auto-configured, or expanded from a 48-bit MAC address (e.g., Ethernet address) auto-generated pseudo-random number. This auto-configuration enables "plug-and-play" deployment in particular of new consumer devices such as cell phones and home appliances and do not require any manual configuration or DHCP server.

Mobile IPv6:
With mobile IPv6, even though the mobile node changes locations and addresses, the existing connections through which the mobile node is communicating are maintained. To accomplish this, connections to mobile nodes are made with a specific address that is always assigned to the mobile node, and through which it is always reachable. Mobile IPv6 provides transport layer survivability when a node moves from one link to another by performing address maintenance for mobile nodes at the Internet layer.

IPv6 Routing Efficiency:
Like CIDR, the larger IPv6 address space enables the use of multiple levels of hierarchy inside the address space. Each level helps to aggregate the traffic at that level. So large address blocks are allocated to ISPs so that they can aggregate the prefixes of all their customers into a single prefix and announce that one prefix to the IPv6 Internet. Similarly, an enterprise network can use only one prefix for the entire network of the organization.

IPv6 QoS:
QoS in IPv6 is handled in the same way it is currently handled in IPv4 through the Traffic Class field implementing the DiffServ model. But IPv6 header has a new field named Flow label which can contain a label identifying a specific flow such as video stream or videoconference. The source node generates this flow label for the QoS devices in the path to take appropriate actions based on this label.

IPv6 Security:
While the use of IPSec is optional in IPv4, IPSec is mandatory in IPv6 and is part of the IPv6 protocol suite. Network implementers can enable IPSec in every IPv6 node and to provide authentication of the node, data integrity and data privacy.

Transitioning from IPv4 to IPv6

Before a transition to an end-to-end native IPv6 network, both service providers and enterprises have a number of migration paths from IPv4 to IPv6 in particular dual IPv4/IPv6 stack, IPv4 tunnels for IPv6 traffic and IPv6 NAT PT (protocol translation). The key is to deploy IPv6 at the edge where the applications and the hosts reside, and then toward the core where the cost to moving to IPv6 are higher and the network operations are more challenging.

IPv4/IPv6 Dual stack backbone:
A core router supporting dual stack can forward packets to IPv4 and IPv6 nodes or to dual stack nodes. A host supporting dual stack can have applications that are not upgraded to run with IPv6 to coexist with IPv6 enabled applications.

IPv6 over IPv4 Tunnels:
IPv4 tunnels encapsulate IPv6 packets to connect isolated IPv6 sites or remote IPv6 networks over an IPv4 backbone.

IPv6 NAT PT:
IPv6 NAT PT provides bi-directional connectivity between IPv4 and IPv6 domains. A dual-stack router with interfaces in both IPv4 and IPv6 networks is capable of performing this function. The difference between IPv6 NAT PT and classic IPv4 NAT is that translations should be done both ways. Packets routed towards IPv4 hosts should have their source/destination addresses changed to IPv4 address equivalents and vice versa.

