Global Information Assurance Certification Paper

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at http://www.giac.org/registration/gsec

Internet Protocol v6: Who Cares?

Jeper Benson
GIAC Security Essentials Certification (GSEC)
Practical, version 1.4c, option 1
September 23, 2004

Abstract

Internet Protocol v4 (IPv4, RFC 791) has been around for 25 years. With the advent of Classless Inter-Domain Routing (CIDR, RFC 1517) and Network Address Translation (NAT, RFC 1631), a passing glance might reveal that the current state of affairs in the IP addressing realm is 'good enough'. This paper brings into question the address shortage and takes a look at the evolutionary Internet Protocol 6 (IPv6, RFC 2460). It might seem that IPv4 is enough and some could conceivably question why we should even care about IPv6. This discourse will cover some of the benefits an organization might garner by adopting IPv6 and reasons for end users to do the same. I will conclude with remarks on current efforts to integrate the IPv4 address space with that of IPv6.

Are we really running out of IP addresses?

It is the role of the Internet Assigned Numbers Authority (IANA) to assign IP addresses from unallocated address space.[1] IPv4 uses a 32-bit address space; that is just under 4.5 billion possible unique addresses. IPv4 was designed at a time when mainframe computer systems were connected for timesharing purposes with hundreds of devices connected to the network, as opposed to hundreds of millions. Although IPv4 has been around since 1981 it has not been substantially changed since then. It is robust, easily implemented and scalable. Despite these features the initial design did not take into account the exponential growth the Internet would have, the depletion of 32 bit addresses or the extensive entries in Internet routing tables.

The Address Shortage Problem

When IP addresses are assigned there are three steps involved in the allocation. IANA allocates blocks of addresses to Regional Internet Registries who in turn allocate blocks to Local Internet Registries and the Local Internet Registries hand out addresses to the end user. Early on in the history of the Internet, the notion of classes was used when assigning IP addresses; the most common were Class A, Class B and Class C. Each address had two parts, the first identified a unique network and the second part identified a unique host in that network.[2]

Address Class	# Network Bits	# Hosts Bits	Decimal Address Range
Class A	8 bits	24 bits	1-126
Class B	16 bits	16 bits	128-191
Class C	24 bits	8 bits	192-223

Because addresses were assigned in these three sizes, a lot of addresses were wasted. Despite the fact that there are only 126 Class A networks available, this accounts for half of the total available addresses. These addresses were supposed to be for very large networks. Because only the first octet was fixed in this type of network this yielded many possible hosts, for a total of 2^{24}-2 (16,777,214) unique IP addresses. Class B networks, in turn were to be used for medium sized networks. The first two octets were fixed in this network, allowing for 2^{14} (16,384) Class B networks that could each have 2^{16}-2 (65,534) unique IP addresses. Class C networks were used for mid sized businesses; the first three octets are part of the network identifier. The last octet is used to identify each host, this allows for a total of 2^{21} (2,097,152) networks each with 2^8-2 (254) possible hosts on it.

Because of the way the three network addresses were divided up, the most commonly sought address type was a Class B address. Many organizations had

4

between 254 and 16 million hosts on their network, making a Class B network the logical choice. If an organization had a couple thousand hosts on their network that was still an enormous waste of IP addresses, since a Class B IP address would allow over 65 thousand hosts on the network. This situation became such a problem that the Internet was running out of unassigned addresses even though only 3% of the assigned addresses were actually being used.[3] It was this combination of rapid growth and waste that spurred the development of IPv6 in the early 1990's. It was approved in 1994 and eventually made a standard in 1998.[4] Although IPv6 was proposed early, something had to be done until it could be implemented.

Classless Inter-Domain Routing

IP addresses were in short demand and routing tables were reaching capacity, a restructuring was required. A new method of address assignment was developed that would increase the efficiency of IP dispersal as well as minimize route table entries. It is known as Classless Inter-Domain routing. CIDR accomplished these two goals: it allowed for more efficient allocation of IP addresses by creating network identifiers that ranged from 13 to 27 bits. This range of network identifiers created the possibility for networks to exist that were as small as 32 hosts or over 500,000 hosts.[5] CIDR also allows for route aggregation, which decreased the size of Internet global routing tables.

Network Address Translation

Once CIDR was implemented the number of IP addresses that were wasted decreased drastically. However even with the innovation of CIDR the IP shortage was still a problem. That is, until a technique known as Network Address Translation was developed. NAT allows a single IP address to grant Internet access to an entire network of computers. This is done by mapping IP addresses from non-routable private addresses to public IP addresses as defined in RFC 1918. NAT routers can be used to translate between any two address realms; typically they sit at the border between public and private networks and work by creating bindings between addresses.[6] Without NAT we would most likely be using IPv6 already. There are several different types of NAT including Static NAT, Dynamic NAT, NAPT and bidirectional NAT but those are outside the scope of this paper.

The combination of CIDR and NAT is a powerful address saving technique. In 2003 BBC Online reported that we would run out of IP addresses as early as 2005,[7] implying that some places would do so before others. This turned out to be false.[8] We will have IP addresses for quite some time, as many as 20 years or more.[9] IANA keeps time stamped log files that are publicly accessible on transactions that are made to the registry over time. Detailed analysis of

5

Regional Internet Registries and BGP routing tables make it evident that despite CIDR and NAT we will eventually run out of addresses.[10] But we will do so globally, not regionally.

If it is true that we aren't running out of IPv4 addresses in the near future, then what is the big deal with IPv6? What's the rush?
Knowing that IPv4 addresses will run out means we need a replacement protocol. Fortunately IPv6 is exactly that. IPv6 was designed with the future of the Internet in mind, paying particular attention to the oversights of IPv4.

A look at Internet Protocol 6

IPv6 features include:
- New header format
- Large address space
- Improved efficiency in routing and packet handling
- Auto configuration and plug and play
- Built-in security
- Better support for QoS
- Extensibility

The IPv6 header has been streamlined, now it has a fixed length of 40 bytes. Non-essential fields and optional fields have been moved behind the IPv6 header. In addition to streamlining the header, fields have been added that define how certain traffic is handled. This allows for Quality of Service support even when the payload is encrypted with IPSec.[11] The headers are 64–bit aligned, meaning they take advantage of the new generation of 64 bit processors resulting in lower overhead than IPv4 options. If the optional extension headers are present they occur in this order:

- Hop-byHop
- Destination
- Routing
- Fragmentation
- Authentication and Encapsulating Security Payload

The Hop-by-Hop header carries information that needs to be examined by all the nodes along the destination path. This header replaces the Time To Live header in IPv4.

The Destination header carries information that can only be examined by the destination node.

The Routing header is used by the source node to list the path that the packet must take to reach its destination.

6

The Fragmentation header is used by the source to indicate that the packet has been fragmented to fit within the maximum transmission unit. If this is the case the packet is assembled by the end nodes instead of the routers as is done in IPv4.

Authentication and Encapsulating Security Payload headers are used in IPSec to provide security services to ensure authentication, integrity and confidentiality of a packet.[12]

Large address space

IPv6 uses a 128 bit address space and allows for an extremely large number of addresses, over 3.4×10^{38}. The large addressing space will allow the allocation of large address blocks to ISPs and other organizations. This will allow more efficient and scalable routing, also reducing the size of routing tables.

Improved efficiency in routing and packet handling

The simplified IPv6 header means there will be less overhead for routers and the elimination of NAT means applications will not have to deal with address translation. The sheer size of IPv6 address space and the multilevel address hierarchy naturally lend themselves to efficient and scalable routing. By allocating large blocks of addresses to ISPs and other organizations it will allow them to aggregate their internal users and announce a single prefix to the Internet. This will significantly reduce the size of routing table entries.[13]

Auto configuration and plug and play

IPv6 supports automatic address configuration, with or without a DHCP server. In the absence of a DHCP server addresses are derived from prefixes advertised by local routers. In the absence of a router, hosts on the same link can automatically configure themselves. To illustrate how this is done I will refer to the sending computer as PC1 and the receiving computer as PC2. If PC1 wants to determine the address of PC2, PC1 sends a neighbor solicitation message to PC2. After PC2 receives the solicitation message, PC2 replies with a neighbor advertisement message containing its address. After PC1 receives the neighbor advertisement, both PC1 and PC2 can communicate with each other.[14]

Built-in security

IPv6 has built in, mandatory IP Security (IPSec, RFC 2401). IPSec is a standard for Virtual Private Networks and has become an industry standard. IPSec provides data integrity, confidentiality and authentication. It has two main modes: the Authentication Header (AH) protocol and the Encapsulated Security Payload (ESP) protocol. The AH provides message integrity, anti-replay, and source authentication, thereby eliminating source spoofing. ESP is a companion protocol to AH and offers the same features in addition to confidentiality. ESP provides the ability to encrypt the contents of the message, adding another layer to the defense in depth strategy. Both AH and ESP work independent of any particular encryption algorithm, although there are some algorithms that they must support in order to meet with IETF standards.[15]

A disadvantage of IPSec is that the AH combined with NAT just doesn't work.[16] Part of what NAT does is to modify IP packets. This creates a problem since IPSec is intended to prevent, among other things, unauthorized modification of the IP packet. As long as we are using IPv4 we don't have to worry about IPSec breaking NAT, because IPSec is optional in IPv4. If we want to implement secure data exchange we can use IPSec ESP instead of the AH method. The hash created by ESP does not include the outer packet header fields.

The elimination of NAT isn't as bad as one might think. Under IPv6 the hosts will have Internet routable IP addresses, but there should still be a firewall performing traffic filtering between the host and the Internet.

Better support for QoS

QoS allows applications to request and receive predictable service levels in terms of data throughput capacity (bandwidth), latency variations and delay.[17] A new field in the IPv6 header called Flow Label will allow routers to identify a set of particular packets as belonging together and provide special handling between the source and destination for them. Because the traffic is identifiable in the header, the payload can be encrypted.[18] This will be ideal for Voice over IP (VoIP), which is basically the transporting of packets containing digitized voice over the network. VoIP allows businesses to talk to other branches using a PC phone, over their corporate intranet, instead of using public phone lines. This eliminates the need for paying long distance phone bills. It is especially useful for international calls and can be advantageous for large corporations, particularly with employees that are spread across a large geographical region.[19]

The combination of QoS and VoIP is very tempting not only for corporations, but for home users as well. If it isn't plug and play it will be much more difficult to sell to the home user market. However, technology seems to be paving the way to do that as Vonage, a consumer VoIP service, has just struck a deal with Cisco Systems' Linksys division to enter the consumer VoIP market. Linksys aims to deliver a wired and wireless router that will be an all in one solution, allowing

8

consumers to use phone, fax and broadband. Linksys promises to incorporate QoS for prioritizing voice packets, achieving clear telephone reception and Universal Plug-and-Play (UpnP) for easy setup and configuration.[20]

Extensibility

IPv6 is not restricted to 40 bytes of options as is IPv4. New features can be added by inserting extension headers after the IPv6 header, and the overall size of the IPv6 header is only constrained by the size of the IPv6 packet.[21] IPv6 provides mandatory IPSec extension headers, making encryption, authentication and virtual private networks easier to implement. It also provides confidentiality with less impact on network performance. Fields have been added to the header to define how traffic is handled, this happens even when the payload is encrypted through IPSec.

Benefits of Adopting Internet Protocol 6

One might ask why a corporation would move to IPv6. Even ignoring the undeniable truth that we will eventually deplete IPv4 addresses, the clear advantages of IPv6 present a stunning argument to make the move.

In a day and age where corporations do not even realize that they have been hacked[22] and suffer large financial losses,[23] it is extremely important to encrypt data. IPSec offers end-to-end secure communications between users and devices. Clearly the corporate world can benefit from the security advantages of IPv6. Not only does IPv6 offer secure communications, it eases the burden of administration by providing automatic address configuration. Given that IT security budget spending is on the rise[24] and the already low cost incentive of having IPv6 built in to most of the popular Operating Systems[25] there should be even fewer reasons to stay with IPv4.

Due to the growth of wireless devices, wireless networking is becoming more and more prevalent in the corporate world. In these environments users can find themselves moving from access point to access point. Mobile IPv6 allows a client node to remain reachable regardless of its location on an IPv6 network.[26]

Current Integration Efforts

The Internet Corporation for Assigned Names and Numbers has started handing out IPv6 addresses.[27] Initially support will be seen on Japan and Korea's country codes (.jp and .kr respectively), France will be next. In a meeting in Kuala Lumpur, Malaysia ICANN has said that it has added IPv6 to the Internet's DNS root server system. This means that businesses and individuals who want to sign

9

up for an IPv6 service will be able to communicate with people using IPv4 addresses.[28]

The Department of Defense has committed to Ipv6 compliance by 2008. Traditionally the DOD has used proprietary infrastructures, however according to John Osterholz, the director of architecture and interoperability for the Department of Defense, the necessity for real-time information has moved the department from an infrastructure of data links between proprietary systems to a secure, global enterprise built on the next generation of open systems.[29]

IPv4 and IPv6: together at last

It is clear that, as more users and devices join the Internet IPv6, will be required. We will also need a smooth transition into the pre-existing IPv4 network. One way to do this is to use tunnels, tunneling encapsulates IPv6 traffic within IPv4 packets so they can be sent over an IPv4 backbone. This allows IPv6 networks to communicate with an IPv4 infrastructure between them. The tunnel is not tied to a specific protocol; it is designed to implement a point-to-point encapsulation scheme. Each protocol has to be setup separately for each link.[30] If you want to communicate with IPv4 and IPv6 networks you will need to use dual-stack routers.

In preparation for this move to IPv6, organizations are beginning to run IPv6 internally and using tunnel brokerage services like Freenet6 to bridge their IPv6 networks to the IPv4 Internet. The way it works is to set up a tunnel using the Tunnel Setup Protocol (RFC 3053) this negotiates and automatically creates configured tunnels on dual-stack hosts or routers. The client connects to a tunnel broker, through client software, exchanges the protocol version and then authenticates to the tunnel broker. The client software is supported on most operating systems.

In a dual-stack backbone deployment, all routers in the network maintain both IPv4 and IPv6 protocol stacks. The key requirement for a dual-stack enabled site is that it has an IPv6 unicast global prefix and appropriate entries in a DNS that map between host names and IP addresses for both IPv4 and IPv6.[31] Applications then have to choose between IPv4 and IPv6, the application chooses the correct address based on the type of IP traffic and requirements of the communication. This approach requires not only that all routers support IPv6 but also that they have enough memory for both IPv4 and IPv6 routing tables.

Tunneling over a dual stack environment can be accomplished several ways:
- IPv6 over IPv4 tunneling
- Manually configured tunnels
- Generic Routing Encapsulation (GRE)
- 6 over 4 tunnels

10

- 6to4 tunnels
- ISATAP
- MPLS

IPv6 over IPv4 tunneling encapsulates IPv6 traffic within IPv4 packets and sends the data over an IPv4 backbone. This can be done in a number of ways:

Manually configured tunnels (RFC 2893) In this tunnel type both end points need to be configured with IPv6 and IPv4 addresses, usually a dual stack router will forward tunnel traffic based on the configuration.[32]

GRE tunnels encapsulate packets within GRE packets and transport the data over IPv4 networks.[33]

6 over 4 tunnels automatically set up tunnels based on the IPv4-compatible IPv6 addresses.[34]

6to4 tunneling uses an IPv4 address embedded in the IPv6 address to identify the end point of the tunnel and setup the tunnel automatically.[35]

ISATAP is similar to 6to4 tunneling but is designed for use in a local site.[36]

Multi-Protocol Label Switching (MPLS) is a packet forwarding technology that uses labels to make data forwarding decisions. The label is a four byte, fixed length identifier that is placed between the data link layer header and the network layer header.[37]

In June of 2003 Freenet6 furthered the integration efforts by offering a NAT traversal solution on operating systems supporting UDP encapsulation (Windows XP, FreeBSD, and Linux)[38] In order to run the TSP (Tunnel Setup Protocol) client the host must run dual-stack networking environments and must have a valid IPv4 address as well as have UDP port 3653 open.

Concluding Remarks

As IPv6 becomes more widespread ISPs will start handing out IPv6 IP addresses. Some reports claim that as many as 50% of ISPs will have IPv6 IP addresses by 2006. As IPv6 addresses begin horning in to the IPv4 space, home users as well as organizations will begin to move to IPv6. Regardless of whether or not you believe that IPv4 address space is going to run out in the near or distant future, the fact remains that IPv6 adoption has been set in motion. It is just a matter of time before IPv6 permeates the Internet and eventually crowds out IPv4. There is no doubt that IPv4 and IPv6 will have to coexist for some time.

11

Corporations will move to IPv6 as VoIP takes hold, streaming media becomes even more popular and QoS becomes more important. ISPs will start handing out more IPv6 addresses; many hardware manufacturers already support IPv6 and IPv4 dual stacks. Eventually users will move to IPv6 not for the advantages that it provides, but because its plug and play features make it easier to do than not to.

[1] Internet Assigned Numbers Authority. "IP Address Services." 29-Apr-2003. URL: http://www.iana.org/ipaddress/ip-addresses.htm (September 23, 2004).

[2] Pacific Bell Internet. "Classless Inter-Domain Routing (CIDR) Overview." 1999. URL: http://public.pacbell.net/dedicated/cidr.html (September 23, 2004).

[3] Pacific Bell Internet. "Classless Inter-Domain Routing (CIDR) Overview." 1999. URL: http://public.pacbell.net/dedicated/cidr.html (September 23, 2004).

[4] Hagen, Silvia. "Learn it, love it." December 19, 2002. URL: http://searchnetworking.techtarget.com/originalContent/0,289142,sid7_gci870277,00.html (September 23, 2004).

[5] Pacific Bell Internet. "Classless Inter-Domain Routing (CIDR) Overview." 1999. URL: http://public.pacbell.net/dedicated/cidr.html (September 23, 2004).

[6] Phifer, Lisa. "The trouble with NAT." July, 2003. URL: http://www.cisco.com/warp/public/759/ipj_3-4/ipj_3-4_nat.html (September 23, 2004).

[7] BBC News. "Tackling the net's numbers shortage." October 26, 2003. URL: http://news.bbc.co.uk/2/hi/technology/3211035.stm (September 23, 2004).

[8] RIPE NCC. "IPv4 Address Space: October 2003." Ocotober, 2003 URL: http://www.ripe.net/rs/ipv4-ncc-20031030.html (September 23, 2004).

[9] Huston, Geoff. "How long have we got." July, 2003. URL: http://www.potaroo.net/ispcolumn/2003-07-v4-address-lifetime/ale.html (September 23, 2004).

[10] Huston, Geoff. "How long have we got." July, 2003. URL: http://www.potaroo.net/ispcolumn/2003-07-v4-address-lifetime/ale.html (September 23, 2004).

[11] Microsoft Corporation. "Tackling the net's numbers shortage." March, 2004. URL: http://download.microsoft.com/download/5/2/5/525343cc-7ba4-4e3b-a96a-c7a040d98d2d/IPv6.doc (September 23, 2004).

[12] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[13] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[14] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[15] SANS Institute. "IPSec Overview, SANS Security essentials book 4". 2004

[16] Phifer, Lisa. "The trouble with NAT." July, 2003. URL: http://www.cisco.com/warp/public/759/ipj_3-4/ipj_3-4_nat.html (September 23, 2004).

[17] Cisco Systems. "QoS Frequently Asked Questions." URL: http://www.cisco.com/en/US/tech/tk543/tk545/technologies_q_and_a_item09186a00800cdfab.shtml (August 23, 2004).

[18] Microsoft Corporation. "Introduction to IP Version 6." March, 2004. URL: http://download.microsoft.com/download/5/2/5/525343cc-7ba4-4e3b-a96a-c7a040d98d2d/IPv6.doc (September 23, 2004).

[19] Hermes Group. "Voice over IP Technologies Ready." September 14, 2000 URL: http://techlibrary.commweb.com/detail/RES/968953727_329.html (September 23, 2004).

[20] Linksys. "New Linksys Voice Over IP (VoIP) Solutions Help Consumer and Small Offices Save Money on Phone Calls." August 24, 2004. URL: http://www.linksys.com/press/press.asp?prid=171 (September 24, 2004).

[21] Microsoft Corporation. "Introduction to IP Version 6." March, 2004. URL: http://download.microsoft.com/download/5/2/5/525343cc-7ba4-4e3b-a96a-c7a040d98d2d/IPv6.doc (September 23, 2004).

[22] SiliconValley.com. "Hacker accessed customer information, Acxiom reports." August 07, 2003 URL: http://www.siliconvalley.com/mld/siliconvalley/news/6484554.htm (September 23, 2004).

[23] Computer Security Institute. "Cyber crime bleeds U.S. corporations, survey shows; financial losses from attacks climb for third year in a row." April 07, 2002 URL: http://www.gocsi.com/press/20020407.jhtml?_requestid=73439 (September 23, 2004).

[24] Sharma, Dinesh. "IT security budgets expected to rise." June 07, 2004 URL: http://news.com.com/IT+security+budgets+expected+to+rise/2100-1009_3-5227840.html (September 23, 2004).

[25] IPv6fourm. "IPv6 Host Operating System Implementation." URL: http://www.ipv6forum.com/navbar/links/v6oslist.htm (September 23, 2004).

[26] Davies, Joseph. "Introduction to Mobile IPv6." September 1, 2004. URL: http://www.microsoft.com/technet/community/columns/cableguy/cg0904.mspx#EHAA (September 23, 2004).

[27] ICANN. "What is ICANN." June 09, 2004 URL: http://www.icann.org/faq/#WhatisICANN (September 23, 2004).

[28] Reardon, Marguerite. "IPv6 domains primed for launch." July 26, 2004 URL: http://news.zdnet.co.uk/internet/0,39020369,39161602,00.htm (September 23, 2004).

[29] CNN Money. "Defense Department Will Require IPv6 Compliance, Says DoD's John Osterholz." June 26, 2003 URL: http://money.cnn.com/services/tickerheadlines/mw/054991.htm (September 23, 2004).

[30] Cisco Systems. "Intranet and Extranet VPN Business Scenarios." URL: http://www.cisco.com/en/US/products/hw/vpndevc/ps333/products_configuration_guide_chapter09186a008009b349.html#1057710 (August 23, 2004).

[31] Cisco Systems. "IPv6 Deployment Strategies." URL: http://www.cisco.com/univercd/cc/td/doc/cisintwk/intsolns/ipv6_sol/ipv6dswp.pdf (August 23, 2004).

[32] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[33] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[34] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[35] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[36] Lee, Dean. Stewart, Elliot. "Internet Protocol version 6 (IPv6) Conformance and Performance Testing." URL: http://www.ixiacom.com/library/white_papers/wp_display.php?skey=ipv6 (September 23, 2004).

[37] Cisco Systems. "MPLS FAQ For Beginners." URL: http://www.cisco.com/en/US/tech/tk436/tk428/technologies_q_and_a_item09186a00800949e5.shtml#qa1 (September 23, 2004).

[38] Hexago. "TSP Client Frequently Asked Questions." 2004 URL: http://www.hexago.com/index.php?pgID=39 (September 23, 2004).