

IPv4 Will Not Be Sufficient For The Next 30 Years

Brian Russell
Department of Computer Science
Rutgers University
morbius@paul.rutgers.edu

Abstract

The Internet has become part of the global infrastructure. The current basis of communication on the Internet is the IPv4 protocol. The future viability of the IPv4 protocol rests on the ability to cope with the expanding number of users and new communication needs that are already coming into existence.

The author takes the position that IPv4 will not suffice for the next 30 years. The increasing demand on the limited address space, the need for more secure communication and quality of service requirements that were not anticipated when IPv4 was created demonstrates that IPv4 will not suffice for the next 30 years.

This paper will provide evidence in support of the position and discuss how the workarounds to the limitations of IPv4 either present new obstacles or fail to meet future needs.

1. Introduction

The power of the Internet has made it a part of the domestic and international infrastructure in less than three decades. The basis of this power is the ability of the machines that exchange information to be able to identify each other by global addresses. The global addressing scheme that has worked from the early 1980's is called IPv4 (Internet Protocol version 4). In IPv4, a 32-bit global IP address serves to uniquely identify a host machine anywhere in the world. In 1977, the number of potential addresses (almost 4.3 billion) was thought to be more than enough to provide for future computer needs [Eklund].

The sad irony is that the very successful role that IPv4 has played will signal its inevitable downfall. Not only has the number of computers using the Internet grown to the point where the remaining global IP addresses are in short supply, there are new devices that use the Internet to communicate and new performance demands from these devices that were not anticipated when IPv4 was first created.

The routers that move data between hosts and their routing and forwarding algorithms are affected by IPv4. The increasing volume of

traffic and the sheer number of hosts and routers affects the efficiency and the reliability of data packet transport between hosts. These pressures will only increase in the future. The practical and theoretical limits of how efficient routers can be are constrained by what data routing requirements can be represented in IPv4.

There are multiple constituencies affected by IPv4. First, small offices and home computer users now have more than one machine connected to the Internet. These users do not have a plug-and-play option when installing new Internet-connected devices; instead they have to deal with the intricacies of Domain Name Servers and routers. These are not things easily managed by the layman and thus they limit access to the Internet and any accompanying financial or recreational benefits.

This is not merely a problem affecting those who do not have deep technical knowledge. Network administrators have to deal with the time and expense of manually configuring and reconfiguring large numbers of machines for Internet access. Clearly a plug-and-play approach would not only make their lives easier, but the needs of the end users of their machines would also be better served.

A final constituency encompasses ISPs and backbone routers. Despite the hierarchical nature of IPv4 addresses, these routers have to deal with changes in IP addresses to create special cases when a customer changes ISP, but keeps the same IP address. These changes become special cases in the backbone routers. There are some 50,000 different routes maintained by the large backbone routers, most of them special cases [Owens 1997]. Every packet that goes through a backbone router involves an examination of an increasingly huge routing table in order to determine the next hop and the overhead of these lookups slows down packet throughput rates for the entire Internet.

The viability of IPv4 is an issue of nothing less than global economic and military consequence. The exhaustion of global IP addresses would mean that existing companies of any size would not be able to expand and new companies would not be able to develop at all. Companies based on new technological demands that could not be supported by IPv4 would strangle. The private individual, who once had available connectivity to the Internet and from that the ability to freely communicate with the rest of the world, would instead have to compete with his fellows for access to the exhausted global IP addresses. The military need for an electronic battlefield creates further demands for global IP addresses and new communication needs.

2. Statement of Position

It is the position of the author that IPv4 will not suffice for the next 30 years. The demands for the remaining unused global IP addresses will exceed the supply in only a few years despite the inventive responses to problems that have been developed. Further, there are new and significant performance demands today and in the near future that IPv4 was never designed to meet and cannot be extended to meet even if there were sufficient addresses available. There are also new communication needs beyond simple unicast that IPv4 cannot support. The projected needs of the electronic battlefield are also making demands on the Internet that must also be met well before the 30 year timeframe.

The remainder of this paper has the following outline. Section 3 is concerned with the demands on global IP address space under IPv4, the technological responses to manage the limited address space and how demand will outstrip supply despite the success of those responses. Section 4 is devoted to security needs and how security is undermined by the very technology that manages the limited IPv4 global address space. Section 5 describes the new quality of service needs and how IPv4 fails to support them. Lastly, Section 6 summarizes the evidence and concludes the position.

3. The Supply of IPv4 Global Addresses Cannot Meet the Demand

The initial function of the Internet was to support communication between computers. Every device using the Internet requires a globally unique IP address. Since its inception, the number of host machines and networks joining the Internet has grown at an exponential rate [Hinden 1995][Kumar 2003], which increases the rate of consumption of the finite number of IP addresses. So when will the supply of available IPv4 IP addresses run out? The answer lies in projections based on the history of IPv4 address allocation and consumption.

The allocation of IPv4 addresses is controlled by the Internet Assigned Numbers Authority (IANA), which allocates blocks of addresses to four Regional Internet Registries (RIRs), who the independently allocate smaller blocks of IP addresses to Local Internet Registries (LIRs) and ISPs. Of the 4,294,967,296 possible IP addresses, there are 3,707,764,736 managed by IANA. The remaining 520,093,696 addresses are used for multicast, held in reserve or for purposes other than the public Internet.

Although accurate records of what blocks of IP addresses have been kept since the inception of IANA, accurate information about exactly when blocks of IP addresses have been allocated have only been available since 1995. Fitting an exponential curve to the IANA allocation data from 1995 through the present would be consistent with the exponential increase in the consumption of IPv4 addresses established by

historical information and provide an extrapolation into the future behavior of IPv4 address space consumption. A member of the Internet Architecture Board did exactly that [Huston 2003] and projected that the IANA will have no more IP addresses to allocate around 2019, only 15 years from now.

Of course, a block of IP addresses allocated by the IANA will not be used instantly; it will be held and allocated gradually in smaller blocks by the RIR that received the allocation. The RIRs also have records of how they have allocated IP addresses along with date information. Since the IANA only has reliable date information for 1995 forward, a similar projection can be made based on RIR allocation information to estimate when RIRs will run out of IP addresses. Huston has also fitted an exponential curve to the RIR data and has projected the exhaustion of the RIR address pool around 2026, which is still many years short of the 30 year timeframe.

The final level of IPv4 address consumption is their appearance as announced addresses in Border Gateway Protocol (BGP) routers. In the same paper [Huston 2003], the author fitted an exponential curve to the BGP routing table data and projected an exhaustion of IPv4 address space of around 2029. The projections of 2019 for IANA exhaustion, 2026 for RIR exhaustion and 2029 for announced IP address exhaustion are consistent with the idea that IANA allocates a large block of IP addresses to an RIR, which allocates portions of that block over time to many LIRs or ISPs before making a new request to IANA. The LIRs and ISPs also allocate their blocks to customers over time and those customers eventually advertise their IP addresses which appear in the BGP routing table. What then, is the counterargument to this analysis?

The counterargument is the fitting of an exponential curve to historical data. From 1995 through 2000, historical data showed that an exponential curve was the best fit. Since then, the pattern of IP address consumption suggests that a linear model may be more representative of IPv4 address consumption. In a revision of his earlier work, Huston made projections on IPv4 address exhaustion based on a linear rather than

exponential curve [Huston 2003a]. Using a linear curve, the IANA would exhaust its pool of unallocated blocks around 2030 and the RIRs would run out of unallocated blocks around 2037, the latter date being past the 30 year timeframe.

The linear curve counterargument is made under that same assumptions as the exponential curve projections, namely that the consumption patterns will remain based on existing technologies without the introduction of new technologies that will make new demands on the IPv4 address space. The counterargument is invalidated with the realization that technological advancement will produce new devices that will consume even more IP addresses.

Cellular phones are expected to trend away from circuit switched to packet switched devices, which creates further demand for IP addresses. IP addresses will also be used for automobiles (one is already on the market)[Morton 1997], television sets [Hinden 1995] and refrigerators [Jarvinen 2002]. Robert Hinden also foresees the replacement of analog control systems for heating, ventilation and cooling (HVAC) systems with more energy efficient digital controls as well as large-scale lighting systems and motor controls, where all of the aforementioned devices require IP addresses. He further states that the markets for these control systems are undeveloped (as of 1995) and expected to be large [Hinden1995].

The military has its own plans that would increase demand for IP addresses. The Department of Defense wants to assign a unique IP address to every piece of equipment and every soldier on the electronic battlefield [Robinson 2004].

In light of increasing demand for IP addresses caused by technological advancement and trends toward widespread use of digital control systems, the optimistic projections of a viable supply of IPv4 addresses lasting for more than 30 years are clearly unrealistic.

In addition to increasing demand, the way in which IP addresses are allocated has further reduced the number of available addresses. To improve routing, blocks of addresses are allocated

hierarchically to organizations that needed many Internet connections. The most significant bits of the addresses in the hierarchical allocation are the network portion of the address and have the same unique value for the network portion of the address. The remaining least significant bits of the address are used by the organization to construct IP addresses as needed. The initial method of address allocation divided address allocation up along byte boundaries and is called classful addressing. Under classful addressing, class A addresses have a network portion that is only one byte long, but allow the owner to construct over 16 million IP addresses with the remaining three bytes of address. The network portion of a class B address is two bytes long and provides 65536 IP addresses to the owner. The class C addresses have a network portion that is 3 bytes long and provides up to 254 IP addresses. The obvious powers of 2 suggested by this addressing scheme were not met exactly because some of the address values were reserved.

The problem with the classful address allocation scheme is that it does not match the needs of the requesting organizations closely enough and is wasteful of IP addresses. An organization that needed a few hundred or a few thousand IP addresses would be issued a class B address that allocated 65536 addresses. Any of those addresses not used by the organization were not available to anyone else, which served to deplete the number of available IP addresses faster than they were actually used.

The response to the wastefulness of classful addressing was to divide the network portion of the IP address along bit boundaries rather than on byte boundaries. This scheme was called classless interdomain routing (CIDR). Dividing network portions of IP addresses along bit boundaries is less wasteful than the classful addressing scheme, but CIDR does not eliminate unused IP addresses. Division on bit boundaries implies that the number of addresses allocated is still a power of 2. An organization that needs a number of IP addresses that is not an exact power of 2 is allocated a block of addresses that is the next higher power of 2. The IP addresses beyond the organization's needs up to that next higher power of two remain unused and unavailable.

Since the IP needs of an organization can be reasonably expected to be any integer value rather than an exact power of 2, the most efficient aggregate allocation scheme currently in use can be expected to waste half the addresses allocated on average. This waste on top of the already limited space inherent to IPv4 supports the argument that IPv4 is not sufficient to meet current and future demands of Internet users for global IP addresses.

A second response to the depletion of IP addresses is Network Address Translation (NAT). NAT is used primarily by small businesses, teleworkers and residential users with multiple hosts [Phifer 2000]. NAT manages the intersection of private networks and the public Internet. In a private network, the device addresses are unique within the private network, but there is no guarantee that these same addresses are unique in the public Internet. NAT is the mechanism that translates nonunique private device addresses to public Internet addresses and vice versa.

This bidirectional private-public translation is handled by NAT routers on the gateway between the private network and the public Internet. A NAT router has a single IP address and looks like a single device to the rest of the Internet no matter how many devices it services on the private network. Packets originating within the private network have a source address taken from the private network and a port number. The NAT router changes the private source address and port number to its own public Internet address and selects a port number from its own pool of available ports. The NAT router maintains a table of the private address and port mappings. The IP and TCP checksums of the outgoing packet are also recalculated by the NAT router. The outgoing packet is then sent to its destination and appears to the receiver as if it had been sent from the NAT router. Any responses sent from the receiver have the IP address of the NAT router as the destination address along with the port number selected by the originating NAT router.

Upon receiving a response packet from the Internet, the NAT router must use its public-private mapping table to translate the destination

address from the IP address of the NAT router to the private address of the originating device within the private network and the port number selected by the NAT router to the port number selected by the device in the private network. The NAT router then forwards the packet to the device within the private network.

The real beauty of this approach is that there can be many machines in the private network that all share the single public IP address of the NAT router. Since the 16-bit port value in the IP packet header suggests as many as 65536 devices on a private network can share a single NAT router, this implies that up to 65536 times 4.2 billion devices could all share the 32-bit public address space if everyone used NAT, thus greatly reducing the demand on the limited size of the IPv4 address space and extending the viability of IPv4. Unfortunately, there are problems inherent to NAT that make it unsuitable for the needs of Internet users.

The first problem is the fact that only the NAT router has an IP address, not the devices in the private network shielded by the NAT router. Each private-public mapping maintained by the NAT router is evanescent and ends after overt termination by the private host or a timeout of the unused mapping by the NAT router itself. Since the devices in the private network do not have permanent IP addresses, incoming packets cannot be directed to them. This means that the devices in the private network cannot act as servers. Similarly, peer-to-peer applications inherently require the IP addresses of the participating peer devices, but devices have no permanent IP address and therefore P2P applications do not work with NAT shielded devices [Kumar 2003].

A second problem concerns the translation of addresses. Many popular applications protocols carry source and destination addresses in the actual packet data and are not accessible to a NAT router [Phifer 2004]. Internet Control Message Protocol (ICMP) and Simple Network Management Protocol (SNMP) packets contain 32-bit IP addresses. Worse, File Transfer Protocol (FTP) packets contain IP addresses as character strings. Application specific algorithms (called Application-Level Gateways) can translate

source and destination addresses in packet data for compatibility with NAT, but the necessity of ALGs for IP address translation in data implies the inability to create new, NAT-compatible applications that carry IP addresses in data, especially when the data may be of arbitrary complexity.

The inability of devices in a NAT shielded private network to act as servers, or to participate in P2P applications, or the problems of translating IP addresses in application data are not the only problems with NAT that make it an unsuitable approach to dealing with the supply and demand problem of the limited IPv4 address space. NAT contributes significantly to the problem of secure communication with IPv4. The security problems are described in the next section.

4. Internet Security Needs Are Not Met By IPv4

IPv4 was conceived in a time when the Internet was used by a small and mostly known community. As the number of Internet users grew, the world proved to be a much less friendly place. The IP security protocol (IPsec) was developed as a network-layer response to threats to communication security on the Internet. The two main IPsec protocols are the Authentication Header protocol (AH) and the Encapsulation Security Payload protocol (ESP). IPsec-AH provides authentication of the sender of messages to prevent impostors from masquerading as someone else, and integrity checks to guarantee that the data that is received is the same data that was sent. IPsec-ESP provides authentication and integrity checks as well as encryption to guarantee that no one without proper decryption keys can examine or alter data exchanged on the Internet. The security issues with IPv4 lie not in the soundness of these security measures, but rather in how these security measures are undermined by the measures taken to compensate for the insufficient IPv4 address space. It is the attempts to compensate for the insufficient address space that actually serve to make secure communication impossible in IPv4.

IPsec-AH was designed to prevent spoofing, man-in-the-middle attacks and unauthorized

modification of data between the source and the destination. These goals are accomplished by the generation of a hash value of each outgoing packet, including the source and destination IP addresses. The hash is inserted into an AH header added to the outgoing packet before it is sent. The receiver uses IPsec-AH to recompute the hash value and any discrepancies between the recomputed hash value and the hash value in the AH header will cause IPsec-AH to consider the received packet to be invalid and drop it. However, if the sender or receiver of packets is a device on a private network using NAT, then no communication is possible. NAT changes the source address of packets crossing the border from private network to public Internet and the destination address of packets crossing from public Internet to private network. The unfortunate consequence of this address modification is that the receiver using IPsec-AH will recompute the hash for the packet including the changed addresses and get a hash value different from the hash value in the AH packet header. The recipient will then treat the packet as having been modified without authorization and drop it and every other packet sent from a device inside a NAT shielded private network. An author summed up the inevitable conclusion most succinctly: "AH + NAT simply cannot work" [Phifer 2000].

IPsec-ESP uses encryption to provide confidentiality. A packet under IPsec-ESP travels in one of two modes: Transport mode or Tunnel mode. Both modes are undermined by NAT routers. In Transport mode, NAT changes the source address and port information in the outgoing packet and must recalculate the checksum in the packet header. When the packet arrives at its destination, IPsec-ESP will detect an altered checksum and drop the packet as invalid. In Tunnel mode, the sender copies the packet header information to a separate packet header, encrypts the entire packet including the original packet header and then makes the encrypted packet as the payload of a larger packet. NAT modifies the source address and port information of the outgoing packet and the checksum in the TCP/UDP header. IPsec-ESP in the receiver decrypts the payload and discovers that the source and port information in the decrypted packet do

not match the source and port information in the enclosing packet header and drops the packet as invalid. Thus, NAT completely undermines the communication of confidential information under IPsec-ESP. This is just another way that NAT undermines security under IPv4.

NAT is also the root of another security problem. A host in a large corporate or academic private network can launch spam or attacks on other hosts outside the private network and delay or evade subsequent identification because any evidence of the attacks contains the IP address of the NAT router and not the attacking host [Srisur 2001]. This means that there is no evidence incriminating any specific device or user in the private network, stifling efforts to locate and capture an attacker before other attacks can be launched.

The combination of NAT and IPv4 has been shown to undermine the security measures taken to protect users of the Internet. IPsec-AH and NAT together do not even allow data to successfully travel from sender to receiver under IPv4. NAT allows attackers to resist investigation into spam and attacks on other hosts by hiding within private networks. IPsec-ESP and NAT cannot be combined in IPsec Transport mode under IPv4.

5. IPv4 Is Unable To Meet New Quality of Service Needs

The Internet was initially conceived to move data between computers. The data consisted of the bits, bytes, character strings and numbers that the computers were designed to use. When there was a large amount of data moving between computers, there was no concern if some data packets got lost or there were interruptions in the flow of data during network congestion as long as all the data got to its destination "fast enough." IPv4 was designed with the assumption that all traffic could be congestion controlled and that data got to its destination as long as the routers did their best efforts to move the data packets. However, the nature of Internet communication has not remained limited to data transfers between computers and the new quality of service (QoS)

requirements imposed by these new data transfers cannot be met by IPv4.

A significant portion if not the majority of Internet use in the immediate future will be the transmission of real-time multimedia information that will consist of audio/video conferencing and multicasting [Morton 1997]. This is high-volume, high-bandwidth information transfer that requires consistent throughput that can only be achieved by imposing fixed constraints on delay and jitter. Treating multimedia traffic as if it were congestion controlled would allow random interruptions in the audio/video presentation.

Given that router bandwidth is finite, the only way to guarantee sufficient bandwidth is to reserve either specific routers in their entirety or reserve bandwidth within specific routers [Hinden 1995]. This prevents other traffic from reducing the bandwidth during the multimedia communication. One problem inherent to IPv4 is that there is no way to define a reservation request protocol and therefore no way to reserve a bandwidth channel. It is also not possible in IPv4 to identify a packet as belonging to a reserved bandwidth channel.

Another IPv4 problem is the inability to define different priorities of real-time communication in response to router congestion. One author offered the example of an HDTV multicast that would run at a lower priority than the comparable low-resolution broadcast [Morton 1997]. Should the relevant routers become congested, the HDTV portion of the broadcast could be dropped in favor of the low-resolution image, which would be more desirable than losing the audio portion of the broadcast. This would place decisions about the relative priorities of different parts of the broadcast in the hands of the sender. Unfortunately, no such provisions exist in IPv4, so routers may respond to congestion with less than desirable results.

Real-time multimedia communication through the Internet is not something that will happen many years in the future, it is already happening now [Morton 1997]. IPv4 does not have the provisions to reserve bandwidth through routers to give multimedia broadcasts the quality they

require in terms of guaranteed throughput and worse, IPv4 does not have the provisions for the quality of service requirements that will continue to arise in the near future.

6. Conclusion

The question of whether or not IPv4 will be sufficient for the next 30 years is an issue that affects many different constituencies, ranging from small offices and home users who are faced with the technologically overwhelming task of managing multiple Internet devices to network administrators and backbone routers that have to deal with manual configuration and decreasing routing efficiency. Economic growth hinges on the ability of devices to communicate with each other in ways beyond those anticipated by the designers of IPv4.

The 32-bit global IP address space, which was once thought to be more than sufficient, is rapidly running out as more computers and other devices join the Internet. Expecting the remaining 32-bit addresses to suffice for the next 30 years in the face of such increasing demand would be unrealistic to say the least.

NAT has been a good response in any ways to the problem of limited IPv4 address space, but it has also caused many problems. Devices in private networks cannot act as servers or participate in P2P applications when NAT changes packets. NAT has also served to undermine the security provisions that have been created to protect the users of the Internet. The combined use of NAT and IPsec network level security outright prevents successful communication between devices under IPv4. Secure communication has proven to be a necessity in a world that has demonstrated remarkable malfeasance.

Lastly, new demands for real-time multimedia conferencing and multicasting have imposed performance requirements that cannot even be communicated in IPv4. IPv4 also provides no means for specifying appropriate forms of graceful degradation in the face of potential router congestion.

Insufficient address space to meet increasing demand, workarounds that prevent successful secure communication and the inability to provide appropriate quality of service for real-time applications are problems that are inherent to IPv4 and cannot be fixed by a retrofit of IPv4.

The inevitable conclusion is that IPv4 will certainly not suffice for the next 30 years. The creators of IPv4 deserve tremendous credit for the creation of an Internet protocol that was successful for many years, but it is clear that its time has passed.

REFERENCES

- [Eklund] Thomas Eklund. Standards for IPv6 based Mobility - Is IPv6 the Answer? Slide Presentation. Includes a quote attributed to Vinton Cerf: "32 bits should be enough address space for Internet." <http://www.ipv6forum.org/navbar/presentations/mobile-isp-eklund/mobile-isp-eklund.pdf>
- [Emigh 2002] Jacqueline Emigh. IPv6: Workarounds to IPv4 Stand in the Way, <http://networking.earthweb.com/netsp/article.php/1402161>. July 11, 2002.
- [Fritsche 2004] Wolfgang Fritsche. Euro Wireless. <http://www.sys-con.com/wireless/article.cfm?id=92>, 2004.
- [Hinden 1995] Robert Hinden. IP Next Generation Overview. <http://playground.sun.com/pub/ipng/html/INET-IPng-paper.html>. 1995.
- [Holdrege 2001] Matt Holdrege, Pyda Srisuresh. RFC 3027, January 2001.
- [Huston 2003] Geoff Huston. IPv4 - How long have we got? The ISP column, July 2003.
- [Huston 2003a] Geoff Huston. IPv4 Address Lifetime Expectancy Revisited. Slide presentation to RIPE Plenary. September 2003.
- [Jarvinen 2002] Antti Jarvinen. Comparing IPv4 and IPv6 Mobility and autoconfiguration for residential networks. Helsinki University of Technology, 2002.
- [JSI 2004] JSI FAQ 6522. Using Internet Protocol Security and Network Address Translation and Internet Security Acceleration Server. <http://www.jsiinc/SUBN/tip6500/rh6522.htm>.
- [Kumar 2003] Santosh Kumar. Lecture 33 IP Next Generation, IPv6, IP Next Layer (IPNL). <http://ww.cse.iitk.ac.in/users/braman/courses/cs625-fall2003/lec-notes/lec-notes33-4.html>, 2003.
- [Morton 1997] David Morton. Understanding IPv6, PC Network Advisor, May 1997.
- [Owens 1997] Keith Owens. IPv6 - Why does the Internet need a new protocol? http://www.luv.asn.au/overheads/ipv6/ipv6_why.html.
- [Phifer 2000] Lisa Phifer. The Trouble with NAT, Internet Protocol Journal, December 2000.
- [Posey 2003] Brian M. Posey MCSE. IPv4 versus IPv6. Tech Republic, April 16, 2003.
- [Robinson 2004] Brian Robinson. IPv6: Built for speed. <http://www.fcw.com/fcw/articles/2004/0830/tec-ipv6-08-30-04.asp>. August 30, 2004.
- [Srisur 2001] Pyda Srisuresh, Kjeld Borch Egevang. RFC 3022, 2001.