

Waiting for IP version 6

January 2003

Geoff Huston

The year has only just started and already I can see the event calendar filling up with a steady stream of IP version 6 summits, workshops and forums, all clamoring for attention. No doubt the claim will be made sooner or later that 2003 will be **the** year for IPv6. Such a claim may have a little more credibility if it was a novel one, but after hearing the same claim made at the start of each of the last five years or so, it's starting to wear a bit thin for me, and with many others I suspect. So will IPv6 ever come, or are we to be left waiting indefinitely?

Some 20 years ago, in January 1983, the ARPANET went through a flag day, and the Network Control Protocol, NCP, was turned off, and TCP/IP was turned on. While there are no doubt some who would like to see a similar flag day where the world turns off its use of IPv4 and switches over to IPv6, such a scenario is a wild-eyed fantasy. Obviously, the Internet is now way too big for coordinated flag days. The transition of IPv6 into a mainstream deployed technology for the global Internet will take some years, and for many there is still a lingering doubt that will happen at all. This seems to be a bit of a contradiction in terms doesn't it? If IPv6 is as good as is claimed, then why aren't we all running it right now? Why hasn't IPv4 run out of available addresses already?

Lets take a closer look at how V6 came about and then look at IPv6 itself in some detail to try and separate the myth from the underlying reality about the timeline for the deployment of IPv6. Maybe then we might be able to venture some answers to these questions.

The Origins of IPv6

The effort that has lead to the specification of IPv6 is by no means a recently started initiative. A workshop hosted by the then Internet Activities Board (IAB) in January 1991 identified the two major scaling issues for the Internet: a sharply increasing rate of consumption of address space and a similar unconstrained growth of the inter-domain routing table. The conclusion reached at the time was that "if we assume that the internet architecture will continue in use indefinitely then we need additional [address] flexibility".

The records of the Internet Architecture Board can be found at www.iab.org/IABmins. While these notes may not represent the most exciting reading matter, what is interesting is the observation that many of the current issues with the Internet were accurately identified over a decade ago in the various IAB workshops.

These issues were taken up later that year by the Internet Engineering Task Force (IETF) with the establishment of the ROAD (ROuting and ADdressing) effort. This effort was intended to examine the

issues associated with the scaling of IP routing and addressing, looking at the rate of consumption of addresses and the rate of growth of the inter- domain routing table. The ultimate objective was to propose some measures to mitigate the worst of the effects of these growth trends. Given the exponential consumption rates then at play, the prospect of exhaustion of the IPv4 Class B space within 2 or 3 years was a very real one at the time. The major outcome of the IETF ROAD effort was the recommendation to deprecate that implicit network / host boundaries that were associated with the Class A, B and C address blocks. In their place the IETF proposed the adoption of an address and routing architecture where the network / host boundary was explicitly configured for each network, and that this boundary could be altered where two or more network address blocks were aggregated into a common single block. This approach was termed “Classless Inter- Domain Routing”, or CIDR. This was a short term measure that was intended to buy some time, and it was acknowledged that it did not address the major issue of defining a longer term scaleable network architecture. But as a short term measure it has been amazingly successful, given that almost ten years and one Internet boom later, the CIDR address and routing architecture for IPv4 is still holding out.

Some would argue that while CIDR was important, it was not the only reason why IPv4 has been able to defy the earlier predictions of its imminent demise. Dynamic Network Address Translation, or NAT, allows a network to use a local private address pool to uniquely number its devices, and then translate these private addresses into public addresses to support transactions involving local and external end points. This way a small pool of public addresses, or even a single address, is used to service a very much larger local private network. It is difficult to estimate the number of devices that are positioned behind NATs, but a highly conservative estimate would see the Internet being at least three times as large as the directly visible part of the Internet.

The IAB, by then renamed the Internet Architecture Board, considered the IETF's ROAD progress in June 1992, still with their eye on the longer term strategy for Internet growth. Their proposal was that the starting point for the development of the next version of IP would be CLNP (Connectionless Network Layer Protocol). This protocol was an element of the Open Systems Interconnection protocol suite (OSI), with CLNP being defined by the ISO 8473 standard. It used a variable- length address architecture, where network-level addresses could be up to 160 bits in length. RFC-1347 contained an initial description of how CLNP could be used for this purpose within the IPv4 TCP/IP architecture and with the existing Internet applications. For the IAB this was a bold step, and considering that the IETF community at the time regarded the OSI protocol suite as a very inferior competitor to their own efforts with IP, it could even be termed a highly courageous step. Predictably, one month later in July 1992, at the IETF meeting in July 1992, this IAB proposal was not well received.

The author recalls a presentation at an IETF plenary session from that time where the OSI protocol suite was termed the road-kill of the Information Superhighway. It was not completely clear that the presenter made the comment in jest!

The IETF outcome was not just a restatement of direction, but a sweeping redefinition of the respective roles and membership of the various IETF bodies, including that of the IAB.

Of course such a structural change in the composition, roles and responsibilities of the bodies that collectively make up the IETF could be regarded as upheaval without definite progress. But perhaps

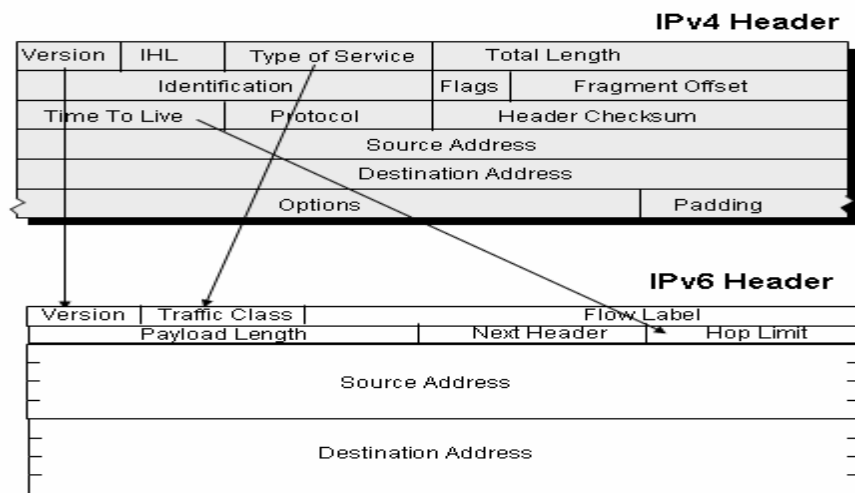
this is an unkind view, as the IAB position also pushed the IETF into a strenuous burst of technical activity. The IETF immediately embarked on an effort to undertake a fundamental revision of the internet protocol that was intended to result in a protocol that had highly efficient scaling properties in both addressing and routing. There was no shortage of protocols offered to the IETF during 1992 and 1993, including the fancifully named TUBA, PIP, SIPP and NAT to mention but a few, as part of a process intended to understand the necessary attributes of such a next generation protocol. The IETF also canvassed various industry sectors to understand the broad dimensions of the requirements of such a protocol.

In 1994 the IETF Next Generation protocol design team defined the core IPv6 protocol. The essential characteristic of the protocol was that of an evolutionary refinement of the version 4 protocol, rather than a revolutionary departure from V4 to an entirely different architectural approach.

IPv6 has had a variety of names - the original IAB documents refer to IP version 7, working on the assumption that the protocol numbers 5 and 6 were already in use in research networks. When some doubt was cast on the use of protocol 6, the effort was renamed IPng, for "next generation". The final word from the IANA was that protocol number 6 was unused, and the final specification was named version 6 of the IP protocol.

IPv6 Changes

The major strength of the IPv6 protocol is the use of fixed length 128 bit address fields. Other packet header changes include the dropping of the fragmentation control fields from the IP header, dropping the header checksum and length, and altering the structure of packet options within the header and adding a flow label. But it is the extended address length that is the critical change with IPv6. A 128 bit address field allows an addressable range of 2 to the 128th power, and 2 to the power of 128 is an exceptionally large number. On the other hand if we are talking about a world that is currently capable of manufacturing more than a billion silicon chips every year, and recognizing that even a 10⁻³ density ration would be a real achievement, then maybe its not all that large a number after all. There is not doubt that such a protocol has the ability to encompass a network that spans billions of devices, which is a network attribute that is looking more and more necessary in the coming years.



Look, No NATS!

It's not just the larger address fields per se, but also the ability for IPv6 to offer an answer to the address scarcity work-arounds being used in IPv4 that is of value here. The side-effect of these larger address fields is that there is then no forced need to use network address translators (NATs) as a means of increasing the address scaling factor. NATs have always presented operational issues to both the network and the application. NATs distort the implicit binding of IP address and IP identity and allow only certain types of application interaction to occur across the NAT boundary. Because the "interior" to "exterior" address binding is dynamic, the only forms of applications that can traverse a NAT are those that are initiated on the "inside" of the NAT boundary. The exterior cannot initiate a transaction with an interior end point simply because it has no way of addressing this remote device. IPv6 allows all devices to be uniquely addressed from a single address pool, allowing for coherent end-to-end packet delivery by the network. This in turn allows for the deployment of end-to-end security tools for authentication and encryption and also allows for true peer-to-peer applications.

In such a light IPv6 can be seen as an attempt to regain the leverage of the original IP network architecture: that of a simple and uniform network service that allows maximal flexibility for the operation of the end-to-end application. It is often the case the complex architectures scale very poorly, and from this perspective IPv6 appears to be a readily scalable architecture.

The Mythology of IPv6

Good as all this is, these attributes alone have not been enough so far to propel IPv6 into broad scale deployment, and there has been considerable enthusiasm to discover additional reasons to deploy IPv6. Unfortunately most of these reasons fall into the category of myth, and in looking at V6 its probably a good idea, as well as fair sport, to expose some of these myths as well.

IPv6 is More Secure

A common claim is that IPv6 is more "secure" than IPv4. It's more accurate to indicate that IPv6 is no more or less secure than IPv4. Both IPv4 and IPv6 offer the potential to undertake secure transactions across the network, and both protocols are potentially superior than attempting to undertake highly secure transactions in the face of various forms of active middleware such as NATs. Yes, the IPv6 specification includes as mandatory support for Authentication and Encapsulating Security Payload extension headers, but no, there is no 'mandatory to use' sticker associated with these extension headers, and, like IPv4 IPSEC, it is left to the application and the user to determine whether to deploy security measures at the network transport level. So, to claim that V6 is somehow implicitly superior to V4 is an overly enthusiastic claim that falls into the category of "IPv6 myth".

IPv6 is Required for Mobility

It is also claimed that only IPv6 supports mobility. If one is talking about a world of tens of billions of mobile devices, then the larger V6 address fields are entirely appropriate for such large scale deployments. But if the claim is more about the technology to support mobility rather than the number of mobile devices, then this claim also falls short. The key issue with mobility is that mobility at a network layer requires the network to separate the functions of providing a unique identity for each connected device, and identifying the location within the network for each device. As a device "moves" within the network its identity remains constant while its location is changing.

V4 overloaded the semantics of an address to include both identity and locality within an address, and V6 did not alter this architectural decision. In this respect IPv4 and IPv6 offer the same levels of support for mobility. Both protocols require an additional header field to support a decoupled network identity, commonly referred to as the “home address”, and then concentrate on the manner of the way in which the home agent maintains a trustable and accurate copy of the mobile node or network’s current location. This topic remains the subject of activity within the IETF in both V4 and V6.

IPv6 is Better for Wireless Networks

Mobility is often associated with wireless, and again there has been the claim that somehow IPv6 is better suited for wireless environments than IPv4. Again this is well in the realm of myth. Wireless environments differ from wireline environments in a number of ways. One of the more critical differences is that a wireless environment may experience bursts of significant levels of bit error corruption, which in turn will lead to periods of non-congestion-based packet loss within the network. A TCP transport session is prone to interpreting such packet loss as being the outcome of network-level congestion. The TCP response is not only retransmission of the corrupted packets, but also an unnecessary reduction of the sending rate at the same time. Neither IPv4 nor IPv6 have explicit signaling mechanisms to detect corruption-based packet loss, and in this respect the protocols are similarly equipped, or ill-equipped as in this case, to optimize the carriage efficiency and performance of a wireless communications subnet.

IPv6 offers better QoS

Another consistent assertion is that V6 offers “bundled” support for differentiated Quality of Service (QoS), whereas V4 does not. The justification for this claim often points to the 20-bit flow label in the IPv6 header as some kind of instant solution to QoS. This conveniently omits to note that the flow identification field in the V6 header still has no practical application in large scale network environments. Both IPv4 and IPv6 support an 8 bit traffic class field, which includes the same 6 bit field for differentiated service code points, and both protocols offer the same fields to an Integrated Services packet classifier. From this perspective QoS deployment issues are neither helped nor hindered by the use of IPv4 or IPv6. Here, again, it’s a case of nothing has changed.

Only IPv6 supports Auto-Configuration

Only IPv6 offer plug and play auto-configuration is another common claim. Again this is an over-enthusiastic statement given the widespread use of the Dynamic Host Configuration Protocol (DHCP) in IPv4 networks these days. Both protocol environments support some level of “plug and play” auto-configuration capability and in this respect the situation is pretty much the same for both V4 and V6.

IPv6 Solves Routing Scaling

It would be good if IPv6 included some novel approach that solved, or even mitigated to some extent, the routing scaling issues. Unfortunately, this is simply not the case, and the same techniques of address aggregation using provider hierarchies apply as much to IPv6 as IPv4. The complexity of routing is an expression of the product of the topology of the network, the policies used by routing entities and the dynamic behaviour of the network, and not the protocol being routed. The larger address space does little to improve on capability to structure the address space in order to decrease the routing load. In this respect V6 does not make IP routing any easier, nor any more scaleable.

IPv6 provides better support for Rapid Prefix Renumbering

If provider-based addressing is to remain an aspect of the deployed IPv6 network, then one way to undertake provider switching for multi-homed end networks is to allow rapid renumbering of a network common prefix. Again, it has been claimed that IPv6 offers the capability to undertake rapid renumbering within a network to switch to a new common address prefix. Again V6 performs no differently from V4 in this regard. As long as “rapid” refers to a period of hours or days then, yes, IPv4 and IPv6 both support “rapid” local renumbering. For a shorter timeframe for “rapid”, such as a few seconds or even a few milliseconds, this is not really the case.

IPv6 provides better support for Multi-Homed sites

This leads on to the more general claim that IPv6 supports multi-homing and dynamic provider selection. Again this is an optimistic claim, and the reality is a little more tempered. Multi-homing is relatively easy if you are allowed to globally announce the network’s address prefix without recourse to any form of provider-based address aggregation. But this is a case of achieving a local objective at a common cost of the scalability of the entire global routing system, and this is not a supportable cost. The objective here is to support some form of multi-homing of local networks where any incremental routing load is strictly limited in its radius of propagation. This remains an active area of consideration for the IETF and clear answers, in IPv4 or IPv6, are not available at present. So at best this claim is premature and more likely the claim will again fall into the category of myth rather than firm reality.

IPv4 has run out of addresses

Again, this is in the category of myth rather than reality. Of the total IPv4 space, some 6% is reserved and another 6% is used for multicast. 51% of the space has already been allocated, and the remaining 37% (or some 1.5 billion addresses) is yet to be allocated. Prior to 1994 some 36% of the address space had been allocated. Since that time, and this includes the entire Internet boom period, a further 15% of the available address space was allocated. With a continuation of current policies it would appear that IPv4 address space will be available for many years yet.

So Why IPv6 Anyway?

The general observation is that V6 is not a “feature-based” revision of IPv4 – there is no outstanding capability of IPv6 that does not have a fully functional counterpart in IPv4. Nor is there a pressing urgency to deploy IPv6 because we are about to run out of available IPv4 address space in the next few months or even years within what we regard as the “conventional” Internet. It would appear that the real drivers for network evolution lurk in the device world. We are seeing the various wireless technologies, ranging from Bluetooth for personal networking through the increasingly pervasive 802.11 hot-spot networking to the expectations arising from various forms of 3G large radius services being combined with consumer devices, control systems, identification systems and various other forms of embedded dedicated function devices. The silicon industry achieves its greatest leverage through sheer volume of production, and it is the combination of Internet utility with the production volumes of the silicon industry that we will see demands for networking that encompasses tens, if not hundreds, of billions of devices. This is the world where IPv6 can and will come into its own, and I suspect that it is in this device and utility mode of communications that we will see the fundamental drivers that will lead to widespread deployment of IPv6 support networks.

Of course predicting the future is easy – the tough bit is getting it right! And there are a very diverse set of views on this topic. But for me I confidently expect my wait for IPv6 to be a mainstream global network service to come to a successful conclusion sometime soon.



Samuel Beckett's *Waiting for Godot* was first performed some 50 years ago on the 5th January, 1953 in Paris. It was an adventurous play, and provoked a mixed reception from audiences. Kenneth Tynan, the London theatre critic, wrote at the time it was produced in London that “it has no plot, no climax, no denouement, no beginning, no middle and no end.” With such a notable lack of these traditional theatre attributes, it should come as no surprise to learn that *Waiting for Godot* is credited as being the play that irrevocably changed theatre. A useful reference for Beckett's work online is at <http://samuel-beckett.net/>.

An alternate presentation of "Waiting for Godot" can be found online at <http://www.musearts.com/cartoons/pigs/godot.html>.

As a postscript to this article, the IPv6 Forum has prepared a [response](#) that offers a rebuttal to these points. You are encouraged to review both perspectives and judge for yourself.

Disclaimer

The above views do not represent the views of the Internet Society, nor do they represent the views of the author's employer, the Telstra Corporation. They were possibly the opinions of the author at the time of writing this article, but things always change, including the author's opinions!

About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. Huston is currently the Chief Scientist in the Internet area for Telstra. He is also a member of the Internet Architecture Board, and is the Secretary of the APNIC Executive Committee. He was an inaugural Trustee of the Internet Society, and served as Secretary of the Board of Trustees from 1993 until 2001, with a term of service as chair of the Board of Trustees in 1999 – 2000. He is author of *The ISP Survival Guide*, ISBN 0-471-31499-4, *Internet Performance Survival Guide: QoS Strategies for Multiservice Networks*, ISBN 0471-378089, and coauthor of *Quality of Service: Delivering QoS on the Internet and in Corporate Networks*, ISBN 0-471-24358-2, a collaboration with Paul Ferguson. All three books are published by John Wiley & Sons.

E-mail: gih@telstra.net