Opinion: The Mythology of IP Version 6

Geoff Huston
July 2003

Disclaimer: This is an opinion piece and, therefore, the author takes some liberties in making his points. I hope you as the reader take this in the spirit in which it is intended—a gentle poke at ourselves that sometimes we oversell ourselves and our technology.

In January 1983, the *Advanced Research Projects Agency Network* (ARPANET) experienced a "flag day," and the Network Control Protocol, NCP, was turned off, and TCP/IP was turned on. Although 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.

Let's look more closely at how IPv6 came about, and then look at IPv6 itself in some detail to try to separate the myth from the underlying reality about the timeline for the deployment of IPv6. Maybe then we can suggest some answers to these questions.

IPv6

The effort that has led 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."

These issues were considered 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 two or three years was a very real one at the time. The major outcome of the IETF ROAD effort was the recommendation

to deprecate the 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 proposed that this boundary could be altered such that two or more network address blocks may be aggregated into a common, single block.

This approach was termed *Classless Interdomain 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, scalable 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 although 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 ROAD progress in June 1992, still with its eye on the longer-term strategy for Internet growth. The board's proposal was that the starting point for the development of the next version of IP would be *Connectionless Network Layer Protocol* (CLNP). This protocol was an element of the *Open System Interconnection* (OSI) protocol suite, 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 long. 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 its own efforts with IP, it could even be termed a highly courageous step. Predictably, one month later in July 1992, at the IETF meeting this IAB proposal was not well received.

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

> At an IETF plenary session from that time, 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!

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, because 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, as well as PIP, SIPP and NAT.

This effort was part of a process intended to understand the necessary attributes of such a next-generation protocol.

The IETF formed an *Internet Protocol Next Generation (IPng) Directorate* in 1994, and canvassed various industry sectors to understand the broad dimensions of the requirements of such a protocol. This group selected the IPv6 Protocol from a set of proposals, largely basing its selection on the so-called "Simple Internet Protocol," or SIP proposal. The essential characteristic of the protocol was that of an evolutionary refinement of the Version 4 protocol, rather than a revolutionary departure from Version 4 to an entirely different architectural approach.

The final word from the *Internet Assigned Numbers Authority* (IANA) was that protocol number 6 was unused, and the final specification was named Version 6 of the Internet Protocol.

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. It was renamed IPng, for "next generation."

The major strength of IPv6 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 one in one thousand address utilization rate would be a real achievement, then maybe it is not all that large a number after all. There is no 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.

Its not just the larger address fields per se, but also the ability for IPv6 to offer an answer to the address scarcity workarounds 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 NAT as a means of increasing the address scaling factor. NAT has always presented operational issues to both the network and the application. NAT distorts the implicit binding of IP address and IP identity and allows 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.

IPv6, as a protocol architecture, is not a radical departure from the architecture of IPv4. The same datagram delivery model is used, with the same minimal set of assumptions about the underlying network capabilities, and the same decoupling of the routing and forwarding capabilities. The use of an address field in the IP header to contain the semantics of both location and identity was not altered in any fundamental way. The changes made by IPv6 could be seen as conservative set of decisions, based on falling back to the IPv4 protocol model for guidance, on the principle that IPv4 is an operating proof of concept for this architectural approach.

In such a light, IPv6 can be seen as an attempt to regain the advantage 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 that complex architectures scale very poorly, and from this perspective the core of 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 consequently 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 IPv6 it is 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 is 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 highly capable in attempting to undertake highly secure transactions. 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 IP Security (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 IPv6 is somehow implicitly superior to IPv4 is an overly enthusiastic claim that falls into the category of "IPv6 myth."

Now I should qualify this, because there is a distinction between the protocol and its environment of deployment. In the case of IPv4, this protocol capability is compromised in many environments in the face of various forms of deployed active middleware such as NAT. It's too early to tell with IPv6, but the line of argument is that NAT-based active middleware has been deployed as a means of address extension, and in a IPv6 world such devices are no longer necessary, and will not be deployed. So perhaps one could say that IPv6 enables a path toward widespread peer-to-peer authentication and transport security at the protocol level, but whether the deployment models faithfully follow along such a path remains an open question.

"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 IPv6 address fields are entirely appropriate for such large-scale deployments. IPv6 includes a developing concept of stateless auto configuration and *Neighbour Discovery* mechanisms.

But if the claim is more about the technology to support mobility 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. IPv4 overloaded the semantics of an address to include both identity and locality within an address, and IPv6 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 current location of the network. This topic remains the subject of activity within the IETF in both IPv4 and IPv6.

"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 numerous 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 signalling 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 IPv6 offers "bundled" support for differentiated *Quality of Service* (QoS), whereas IPv4 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 claim conveniently omits to note that the flow identification field in the IPv6 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 is a case of nothing has changed.

"Only IPv6 Supports Auto-Configuration"

Another common claim is that only IPv6 offers "plug-and-play" auto configuration. Again this is an overenthusiastic 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 IPv4 and IPv6.

"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 they do to 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 - 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 IPv6 does not make IP routing any easier, nor any more scalable.

"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 multihomed 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 IPv6 performs no differently from IPv4 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 time frame for "rapid," such as a few seconds or even a few milliseconds, this is not really the case.

"IPv6 Provides Better Support for Multihomed 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. Multihoming is relatively easy if you are allowed to globally announce the network 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 multihoming 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 percent is reserved and another 6 percent is used for multicast. Forty-one percent of the space has already been allocated, and the remaining 37 percent (or some 1.5 billion addresses) is yet to be allocated. Prior to 1994, some 36 percent of the address space had been allocated. Since that time, and this includes the entire Internet boom period, a further 15 percent 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 IPv6 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 IEEE 802.11 "hot-spot" networking to the expectations arising from various forms of *third-generation* (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 advantage through sheer volume of production, and it is in 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.

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 the Executive Director of the Internet Architecture Board, and is a member of the APNIC Executive Committee. 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