

October 2008

Geoff Huston

## **The Changing Foundation of the Internet: Confronting IPv4 Address Exhaustion**

Throughout its relatively brief history the Internet has continually challenged our preconceptions about networking and communications architectures. For example, the concept that the network itself has no role in management of the network's own resources, and that resource allocation is the result of interaction between competing end-to-end data flows was certainly a novel innovation, and for many it has been a very confronting concept. The approach of designing a network that is unaware of services and service provisioning, and is not attuned to any particular service whatsoever, leaving the role of service support to end-to-end overlays, was again a radical concept in network design. The Internet has never represented the conservative option for this industry, and has managed to define a path that continues to present significant challenges.

From such a perspective it should be entirely unsurprising that the next phase of the Internet's story, that of the transition of the underlying version of the IP protocol from IPv4 to IPv6, refuses to follow the intended script. Where we are now, in mid-2008, with IPv4 unallocated address pool exhaustion looming within the next 18 to 36 months, and IPv6 still largely not deployed in the public Internet, is a situation that was entirely unanticipated and, even in hindsight, entirely surprising to encounter.

The topic examined here is why this situation has arisen, and in examining this question the options available to the Internet to resolve the issue of exhaustion of the supply IPv4 addresses will be analysed. The timing of the exhaustion of IPv4 addresses, and the nature of the intended transition to IPv6 will be examined. The shortfalls in the implementation of this transition will be considered, and the underlying causes of these shortfalls will be identified. The options that are available at this stage will be examined and some likely consequences of such options identified.

### **When?**

This question was first asked on the TCP/IP list in November 1988, and the responses included foreshadowing a new version of IP with longer addresses and undertaking an exercise to reclaim unused addresses [TCP/IP 1988]. The exercise of measuring the rate of consumption of IPv4 addresses has been undertaken many times in the past two decades, with estimates of exhaustion ranging from the late 1990's to beyond 2030. One of the earliest exercises in predicting IPv4 address exhaustion was undertaken by Frank Solensky, presented at IETF 18, August 1990. His findings are reproduced in Figure 1.

At the time the concern was primarily the concern was the rate of consumption of the supply of Class B network addresses (or the supply of /16 prefixes from the address block 128.0.0.0/2, to use current terminology). There were only 16,384 such Class B network addresses within

the Class-based IPv4 address plan, and the rate of consumption was such that the Class B networks would be fully consumed within a further four years, or by 1994. The prediction was strongly influenced by a significant number of international research networks connecting to the Internet in the late 1980's, with the rapid influx of new connections to the Internet creating a surge in demand for Class B networks.

<u>Depletion Dates</u>	
• Assigned Class "B" network numbers	Mar. 11, 1994
• NIC "connected" Class B network numbers	Apr. 26, 1996
• NSFnet address space*	Oct. 19, 1997
• Assigned Class "A-B" network numbers	Feb. 17, 1998
• NIC "connected" Class A-B network numbers	Mar. 27, 2000
• BBN snapshots*	May 4, 2002

\* all types: may be earlier if network class address consumption is not equal.

Figure 1 – Report on IPv4 Address Depletion [Solensky90]

Successive predictions were made in the context of the IETF in the Address Lifetime Expectancy Working Group, where the predictive model was refined from an exponential growth model to a logistical saturation function, attempting to predict the level at which all address demand would be met.

The predictive technique described here is broadly similar, using a statistical fit of historical data concerning address consumption into a mathematical model, then using this model to predict future address consumption rates and thereby predict the exhaustion date of the address pool.

The predictive technique models the IP address distribution framework. Within this framework the pool of unallocated /8 address blocks is distributed by the Internet Assigned Number Authority (IANA) to the five Regional Internet Registries (RIRs). (A "/8 address block" refers to a block of addresses where the first 8 bits of the address values are constant. In IPv4 a /8 address block corresponds to 16,777,216 individual addresses.) Within the framework of the prevailing address distribution policies, each RIR is able to request a further address allocation from IANA when the remaining RIR-managed unallocated address pool falls below a level required to meet the next 9 months allocation activity, and the amount allocated is the number of /8 address blocks required augment the RIR's local address pool to meet the anticipated needs of the regional registry for the next 18 months. However, in practice, the RIRs currently request a maximum of 2 /8 address blocks in any single transactions, and do so when the RIR-managed address pool falls below a threshold of the equivalent of 2 /8 address blocks.

As of August 2008 there are some 39 /8 address blocks left in IANA's unallocated address pool. A predictive exercise has been undertaken using a statistical modelling of historical address consumption rates, using data gathered from the Regional Internet Registries' (RIRs') records of address allocations, and the time series of the total span of address space announced in the

internet's inter-domain default-free routing table as basic inputs to the model. The predictive technique is based on a least squares best fit of a linear function applied to the first order differential of a smoothed copy of the address consumption data series, as applied to the most recent 1,000 days' data (Figure 2).

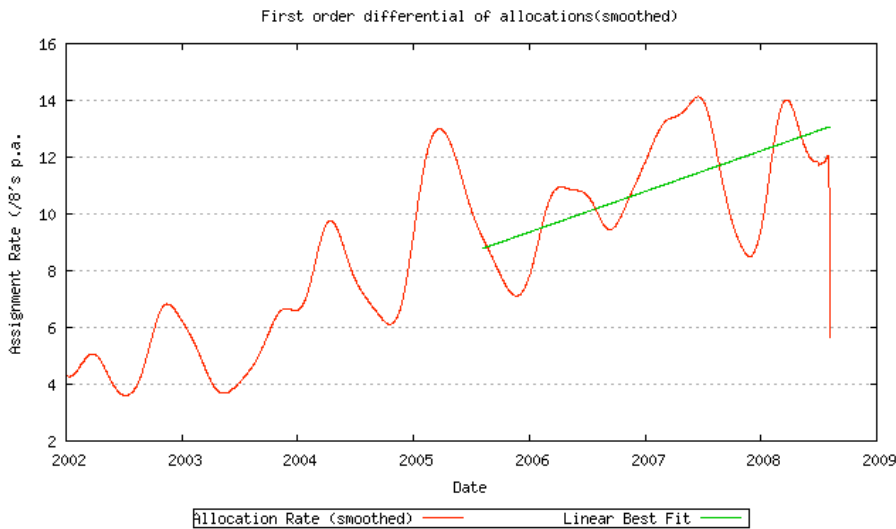


Figure 2 – First Order Differential of Daily IPv4 Address Allocation Rates [Huston08]

The linear function, which is a best fit to the first order differential of the data series is integrated to provide a quadratic time series function to match the original data series. The projection model is further modified by analysing the day-of-year variations from the smoothed data model, averaged across the past three years, and applying this daily variation to the projection data, in order to take into account the level of seasonal variations in the total address consumption rate that has been observed in the historical data (Figure 3).




Figure 3 – Address Consumption Model [Huston08]

The anticipated rate of consumption of addresses from this central pool of unallocated IPv4 addresses is expected to average 15 /8's per year over 2009, and slightly more in 2010.

RIR behaviours are modelled using the current RIR operational practices and associated address policies, which is used to predict the times when each RIR will be allocated a further 2

/8's from IANA. This RIR consumption model, in turn, allows the IANA address pool to be modelled.

This anticipated rate of increasing address consumption will see the remaining unallocated addresses that are held by IANA reach the point of exhaustion in February 2011. The most active RIRs are anticipated to exhaust their locally managed unallocated address pools in the months following the time of IANA exhaustion (Figure 4).




Figure 4 – IANA Depletion Prediction [Huston08]

The assumption behind this form of prediction is that the current policy framework relating to the distribution of addresses will continue to apply without any further alteration through to complete exhaustion of the unallocated address pool, and that the demand curves will remain consistent, in so far as there will be no forms of disruption to demand, such as a panic rush on the remaining addresses or some introduced externality that impacts on total address demand, and that the level of return of addresses to the unallocated address pool will not vary significantly from existing levels of address return.

While the statistical model is based on a complete data set of address allocations and a detailed hourly snapshot of the address span advertised in the Internet's routing table, there is still a considerable level of uncertainty associated with this prediction.

Firstly, the behaviour of the Internet Service Provider (ISP) industry and the other entities who are the direct recipients of RIR address allocations and assignments are not ignorant of the impending exhaustion condition, and there is some level of expectation of some form of last minute rush or panic on the part of such address applicants when exhaustion of this address pool is imminent. The predictive model described here does not include such a last minute acceleration of demand.

The second factor is the skewed distribution of addresses in this model. In the period from 1 January 2007 until 20 July 2008 there were 10,402 allocation or assignments transactions recorded in the RIRs' daily statistics files. These transactions accounted for a total 324,022,704 individual IPv4 addresses, or the equivalent of 19.3 /8's. Precisely one half of this address space was allocated or assigned in just 107 such transactions. In other words, some 1% of the recipients of address space in the past 18 months have received some 50% of all the allocated address space. The distribution of allocation sizes is indicated in Figure 5. The reason why this distribution is relevant here is that this predictive exercise assumes that while individual actions are hard to predict with any certainty, the aggregate outcome of many individuals' actions assumes a much greater level of predictability. This observation about

aggregate behaviour is not the case in this situation, and the predictive exercise is very sensitive to the individual actions of a very small number of recipients of address space because of this skewed distribution of allocations. Any change in the motivations of these larger-sized actors that results in an acceleration of demand for IPv4 will have a very significant impact on the predictions of the longevity of the remaining unallocated IPv4 address pool.

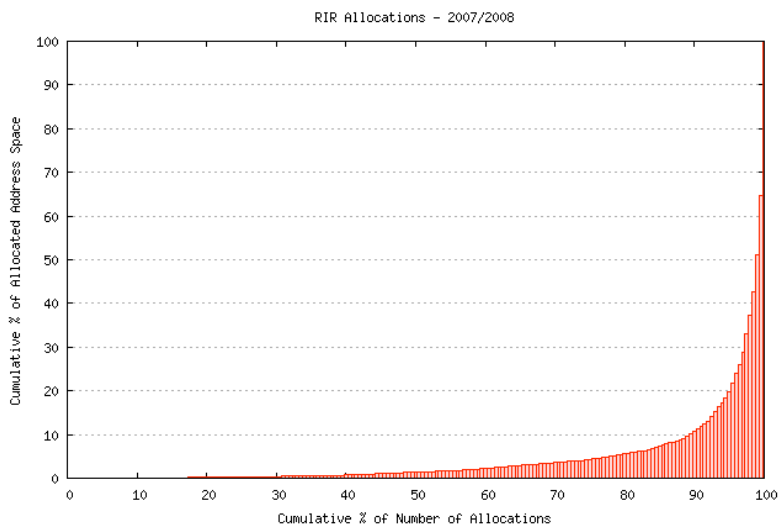


Figure 5. *Distribution of Sizes of Allocated Address Blocks*

The third factor is that this model assumes that the policy framework remains unaltered, and that all unallocated addresses are allocated or assigned under the current policy framework, rather than under a policy regime that is substantially different from today's framework. The related assumption here is that the cost of obtaining and holding addresses remains unchanged, and that the perceptions of future scarcity of addresses have no impact on the policy framework of address distribution of the remaining unallocated IPv4 addresses.

Given this potential for variation within this set of assumptions, a more accurate summary of the current expectations of address consumption would be that the exhaustion of the IANA unallocated IPv4 address pool will occur sometime between July 2009 and July 2011, and that the first RIR will exhaust all its useable address space within three to 12 months from that date, or between October 2009 and July 2012.

## What Next?

Apart from the exact date of exhaustion that is predicted by this modelling exercise, none of the information relating to exhaustion of the unallocated IPv4 address pool should be viewed as particularly novel information. From the IETF's Routing and Addressing (ROAD) study of 1991, it was recognised that the IPv4 address space was always going to be completely consumed at some point in the Internet's future [RFC1380].

Such predictions of the potential for exhaustion of the IPv4 address space were the primary motivation for the adoption of Class-less Inter-Domain Routing (CIDR) in BGP, and the corresponding revision of the address allocation policies to craft a more exact match between planned network size and the allocated address block. These predictions also motivated the protracted design exercise of what was to become the IPv6 protocol across the 1990's within the IETF. The prospect of address scarcity engendered a conservative attitude to address management that, in turn, was a contributory factor in driving the widespread deployment of Network Address Translators (NATs) [RFC1631] in the Internet the past decade. By any

reasonable metric this industry has had ample time to study this issue, ample time to devise various strategies, and ample time to make plans and execute them.

And this has been the case for adopting classless address allocations, the adoption of CIDR in BGP and the extremely widespread deployment of NATs. But all of these were short term measures, while the longer term measure, that of the transition to IPv6, was what was intended to come after IPv4. But IPv6 has not been the subject of widespread adoption so far, while the time of anticipated exhaustion of IPv4 has been drawing closer. Given almost two decades of advance warning of IPv4 address exhaustion, and a decade since the first stable implementations of IPv6 were released, it would be reasonable to expect that this industry, and each individual actor within this industry, is aware of the problem here and the need for a stable and scalable long term solution as represented by IPv6. It would be reasonable to anticipate that with respect to IPv6 transition, the industry has already planned what actions will be taken, and is aware of the triggers that will invoke such actions, and approximately when this will occur.

However, such an expectation appears to be ill-founded when considering the broad extent of the actors in this industry, and there is little in the way of a common commitment as to what will happen after IPv4 address exhaustion, nor even any coherent view of plans that industry actors are making in this area.

This makes the exercise of predicting the actions within this industry following address exhaustion somewhat challenging, so instead of immediately describing future scenarios, it may be useful to first describe the original plan for the Internet's response to IPv4 address exhaustion.

### **What Was Intended?**

The original plan, devised in the early 1990's by the Internet Engineering Task Force (IETF), to address the IPv4 address shortfall was the adoption of classless inter-domain routing (CIDR) as a short term measure to slow down the consumption of IPv4 addresses by reducing the inefficiency of the address plan, and the longer term plan of the specification of a new version of the Internet Protocol that would allow for adoption well before the time of exhaustion of the IPv4 address pool.

The industry also adopted the use of NATs as an additional measure of increasing the efficiency of address use, although the IETF did not take a strongly supportive position with respect to NATs. For many years the IETF did not undertake the standardization of NAT behaviours, presumably on the basis that NATs were not consistent with the IETF's advocacy of end-to-end coherence of the Internet at the IP level of the protocol stack.

Over the 1990's the exercise of the specification of a successor IP protocol to version 4 was undertaken by the IETF, and the IETF's view of the longer term response was refined to be advocacy of the adoption of the IPv6 protocol and the use of this protocol as the replacement for IPv4 across all parts of the network.

In terms of what has happened in the past 15 years, the adoption of CIDR was extremely effective, and most parts of the network were transitioned to use CIDR within a little under two years, with the transition declared to be complete by the IETF in June 1996. And, as noted already, NATs have been adopted across many, if not most, parts of the network. The most common point of deployment of NATs has been not been at an internal point of demarcation between provider networks, but at the administrative boundary between the local customer network and the ISP, so that the common configuration of customer premises equipment includes NAT functionality. Customers effectively own and operate NAT devices as a commonplace aspect of today's deployed Internet.

CIDR and NAT have been around for more than a decade now, and the address consumption rates have been held at very conservative levels in that period, particularly so when considering that the bulk of the Internet's population was added well after the advent of CIDR and NATs.

The longer term measure, that of the transition to IPv6, has not proved to be as effective in terms of adoption in the Internet.

There was never going to have been a "flag day" transition where, in a single day, simultaneously across all parts of every network the IP protocol was flipped over to use IPv6 instead of IPv4. The Internet is too decentralized, too large, too disparate and too critical for such actions to be orchestrated, let alone be completed with any chance of success. A flag day, or any such form of coordinated switchover, was never a realistic option for the Internet.

If there was no possibility of a single coordinated switch over the IPv6, the problem is that there was never going to be an effective piecemeal switch over either. By this it is meant that there was never going to be a switch over where host by host, and network by network, IPv6 is substituted for IPv4 on a piecemeal and essentially uncoordinated basis. The problem here is that IPv6 is not "backward compatible" with IPv4. When host uses IPv6 exclusively then that host has no direct connectivity to any part of the IPv4 network. If an IPv6-only host is connected to an IPv4-only network, then the host is effectively isolated. This situation does not bode well for a piecemeal switchover, where individual components of the network are switched over from IPv4 to IPv6 on a piecemeal basis. Each host that switches over to IPv6 essentially disconnects itself from the IPv4 Internet at that point.

Given this inability to support backward compatibility, what was planned for the transition to IPv6 was a "dual stack" transition. Rather than switching over from IPv4 to IPv6 in one operation on both hosts and networks, a two-step process has been proposed: firstly switching from IPv4 only to a "dual stack" mode of operation that supports both IPv4 and IPv6 simultaneously, and secondly, and at a much later date, switching from dual stack IPv4 and IPv6 to IPv6 only. During the transition more and more hosts are configured with dual stack. The idea is that dual-stack hosts prefer to use IPv6 to communicate with other dual stack hosts, and revert to use IPv4 only when an IPv6-based end-to-end conversation is not possible. As more and more of the Internet converted to dual stack there was anticipated to be a declining use of IPv4, until the point was reached that there was no further value in continuing to support IPv4. In this dual stack transition scenario no single flag day is required and the dual stack deployment can be undertaken in a piecemeal fashion. There is no requirement to coordinate hosts with networks, and as dual stack capability is supported in networks the attached dual stack hosts can make use of IPv6. This scenario still makes some optimistic assumptions, particularly relating to the achievement of universal deployment of dual stack, at which point all IPv4 functionality is no longer being used, and support for IPv4 can be terminated. Knowing when this point has been reached is unclear, of course, but, in principle there is no particular timetable for the duration of the dual stack phase of operation.




Figure 6. A view of Dual Stack diffusion of IPv6.

There are always variations, and in this case it is not necessarily that each host must operate in dual stack mode for such a transition. A variant of the NAT approach can perform a

rudimentary form of protocol translation, where a protocol translating NAT (or NAT-PT [RFC2766]) essentially transforms an incoming IPv4 packet to an outgoing IPv6 packet, and vice versa, using algorithmic binding patterns to map between IPv4 and IPv6 addresses. While this relieves the IPv6-only host of some additional complexity of operation at the expense of some added complexity in DNS transformations and service fragility, the essential property still remains that in order to speak to an IPv4-only remote host the combination of the local IPv6 host and the NAT-PT have to generate an equivalent IPv4 packet. In this case the complexity of the dual stack is now replaced by complexity in a shared state across the IPv6 host and the NAT-PT unit. Of course this solution does not necessarily operate correctly in the context of all potential application interactions, and there are significant concerns with the integrity of operation of NAT-PT devices, a factor which motivated the IETF to deprecate the existing NAT-PT specification [RFC4966]. On the other hand, the lack of any practical alternatives has led to the IETF to subsequently reopen this work, and once again look at specifying the standard behaviour of such devices [Bagnulo 2008].

The detailed progress of a dual stack transition is somewhat uncertain, as it involves the individual judgment of many actors as to when it may be appropriate to drop all support for IPv4 and rely solely on IPv6 for all connectivity requirements. However, one factor is constant in this envisaged transition scenario, and whether its dual stack in hosts, or dual stack via NAT-PT, or various combinations thereof, there is the consistent requirement that there are sufficient IPv4 addresses to span the entire Internet's addressing needs across the complete duration of the dual stack transition process. Under this dual stack regime every new host on the Internet was envisaged to need access to both IPv6 and IPv4 addresses, in order to converse with any other host using IPv6 or IPv4. Of course this approach works for as long as there is a continuing supply of IPv4 addresses, implying that the timing of the transition was such that it was meant to have completed by the time that IPv4 address exhaustion was going to happen, as indicated in Figure 7.




Figure 7. The Planned Dual Stack Transition Model

If this transition were to commence in earnest at the present time, in late 2008, and take an optimistic five years to complete, then at the current address consumption rate we will require a further 90 to 100 /8 address blocks to span this five year period. A more conservative estimate of a 10 year transition will require a further 200 to 250 /8 address blocks, or the entire IPv4 address space again, assuming that we will use IPv4 addresses in the future in the precisely the same manner as we've been using them in the past and with precisely the same level of utilization efficiency as we've managed to date.






Figure 8. The Current Outlook for Dual Stack Transition

Clearly, waiting for the time of IPv4 unallocated address pool exhaustion to act as the signal to industry to commence the deployment of IPv6 in a dual stack transition framework is a totally flawed implementation of the original dual stack transition plan (Figure 8)

Either the entire process of dual stack transition will need to be undertaken across a far faster time span than has been envisaged, or the manner of use of IPv4 addresses, and, in particular their utilization efficiency, in the context of dual stack transition support will need to differ markedly from the current manner of use of addresses. It may be the case that a number of forms of response are required, which poses some challenging questions as there is no agreed precise picture of what markedly different and significantly more efficient form of address use is required here. To paraphrase the situation, its clear that we need to do "something" differently, and do so as a matter of some urgency, but we have no clear agreement on what that "something" is that we are meant to be doing differently. This is not an optimal situation.

What was intended as a transition mechanism for IPv6 is still the only feasible approach that we are aware of, but the forthcoming exhaustion of the unallocated IPv4 address pool now calls for novel forms of use of IPv4 addresses within this transitional framework, and this may well entail the deployment of various forms of address translation technologies that we have not yet defined, let alone standardized. The transition may also call for scaling capabilities from the inter-domain routing system that also head into unknown areas of technology and deployment feasibility.

## Why?

At this point it may be useful to consider how and why this situation has arisen.

If the industry needed an abundant supply of IPv4 addresses to be on hand to underpin the entire duration of the dual stack transition to IPv6, then why didn't they follow the lead of the IETF and commence this transition while there was still an abundant supply of IPv4 addresses on hand? If network operators, service providers, equipment vendors, component suppliers, application developers, and every other part of the Internet's supply chain were aware of the need to commence a transition to IPv6 well before effective exhaustion the remaining pool of IPv4 addresses, then why didn't the industry make a move earlier? Why was the only clear signal for a change in the operation of the Internet to commence a dual stack transition to IPv6 one that has been activated too late to be useful for the industry to act on in an efficient manner?

One possible reason may lie in a perception of technical immaturity of IPv6 as compared to IPv4. It is certainly the case that many network operators in the Internet are highly risk adverse and tend to operate their networks in a mainstream path of technologies rather than constant use of leading edge advance releases of hardware and software solutions. Does IPv6 represent some form of unacceptable technical risk of failure that has prevented its adoption? This does not appear to be the case, either in terms of observed testing, or in terms of observation of perceptions about the technical capability of IPv6. The IPv6 protocol is functionally complete, internally consistent, and capable of being used in almost all contexts where IPv4 is used today. IPv6 works as a platform for all forms of transport protocols, and is fully function as an internetwork layer protocol that is functionally equivalent to IPv4. IPv6 NATs exist, DHCPv6 provides dynamic host configuration for IPv6 nodes, the DNS can be completely equipped with IPv6 resource records and operate using IPv6 transport for queries and responses. Perhaps the only notable difference between the two protocols in the ability to perform host scans in IPv6, where probe packets are sent to successive addresses. In IPv6 the address density is extremely low as each host's low order 64 bit interface address is more or less unique and within a single network the various interface addresses are not clustered sequentially in the number space. The only known use of address probing to date has been in various forms of hostile attack tools, so the lack of such a capability in IPv6 is generally seen as a feature rather than an impediment. IPv6 deployment has been undertaken in a small scale for many years, and while the size of the deployed IPv6 base remains small, the level of experience gained with the functionality of the technology has been significant. It is possible to draw the conclusion that IPv6 is technically capable and this capability has been broadly tested in almost every scenario except that of universal use across the Internet.

It also does not appear that the reason was a lack of information or awareness of IPv6. The efforts of promotion of IPv6 adoption have been underway for almost a decade now in earnest. All regions and many of the larger economies have instigated programs to promote the adoption of IPv6 and provide information to local industry actors of the need to commence a dual stack transition to IPv6 at an early date. In many cases these promotional programs have enjoyed broad support from both public and industry funding sources. The coverage of these promotional efforts has been widespread in industry press reports. Indeed, perhaps the only criticism of this effort is a case of too much promotion, with a possible result that the effectiveness of the message to industry of a need to commence IPv6 adoption has been dulled through constant repetition.

A more likely area to examine in terms of possible reasons why industry has not engaged in dual stack transition deployment is that of the business landscape of the Internet. The Internet can be viewed as a product of the wave of progressive deregulation in the telecommunications sector in the 1980's and early 1990's. The search undertaken by new players in the deregulated industry for a competitive edge that was capable of unseating the dominant position of the legacy incumbents found the Internet as their competitive lever. The result was perhaps unexpected, as it was not one that replaced one vertically integrated operator with a collection of similarly structured operators whose primary means of competitive was in terms of price efficiency across an otherwise undifferentiated service market as we saw in the mobile telephony industry. In the case of the Internet the result was not one that attempted to impose convergence on this industry, but one that stressed divergence at all levels, accompanied by branching role specialization at every level in the protocol stack and at every point in the supply chain process. In the framework of the Internet consumers are exposed to all parts of the supply process, and were not reliant on an integrator to package and supply a single all-embracing solution. Consumers make independent purchases of their platform technology, their software, their applications, their access provider and their means of advertising their own capabilities to provide goods and services to others, all as independent decisions, all as a result of this direct exposure to the consumer of every element in the supply chain.

What we have today is an industry structure that is highly diverse, broadly distributed, strongly competitive and intensely focussed on meeting specific customer needs in a price

sensitive market, operating on a quarter-by-quarter basis. Bundling and vertical integration of services has been placed under intense competitive pressure, and each part of the network has been exposed to specialized competition in its right. For consumers this has generated significant benefits. For the same benchmark price of around USD 15 - 30 per month, or its effective equivalent in a local currency's purchasing power, today's Internet user enjoys multi-megabit per second access to a richly populated world of goods and services. The price of this industry restructure has been a certain loss of breadth and depth of the supply side of the market. If consumers do not value a service, or even a particular element of a service, then there is no benefit in incurring marginal additional cost in providing the service. In other words, If the need for a service is not immediate, then it is not provided. For all service providers right through the supply side the focus is on current customer needs, and this focus on current needs, as distinct from continued support of old products or anticipatory support of possible new products, excludes all other considerations.

Why is this change in the form of operation of communications industry an important factor in the adoption of IPv6? The relevant question in this context is that of placing IPv6 deployment and dual stack transition into a viable business model. IPv6 was never intended to be an end user visible technology. It offers no additional functionality to the end user, nor any direct cost savings to the customer or the supplier. Current customers of Internet service providers do not need IPv6 today, and neither current nor future customers are aware that they may need it tomorrow. For end users of Internet services email is email and web-based delivery of services is just the web. Nothing will change that perspective in an IPv6 world, so from that respect customers do not have a particular requirement for IPv6, as opposed to a generic requirement for IP access, and will not value such an IPv6-based access service today in addition to an existing IPv4 service. For an existing customer IPv6 and dual stack simply offers no visible value. So if the existing customer base places no value on the deployment of IPv6 and dual stack then the industry has little incentive to commit to the expenditure to provide it. Any IPv6 deployment across an existing network is essentially an unfunded expenditure exercise that erodes the revenue margins of the existing IPv4-based product. And as long as there is sufficient IPv4 address space remaining to cover the immediate future needs, looking at this on the basis of a quarter-by-quarter business cycle, then the decision to commit to additional expenditure and lower product margins to meet the needs of future customers using IPv6 and dual stack deployments is a decision that can comfortably be deferred for another quarter. This business structure of today's Internet appears to represent the major reason why the industry has been incapable of making moves on dual stack transition within a reasonable timeframe as it relates to the timeframe of IPv4 address pool exhaustion.

What of the strident calls for IPv6 deployment? Surely there is substance to the arguments to deploy IPv6 as a contingency plan for the established service providers in the face of impending IPv4 address exhaustion, and if that is the case why have service providers discounted the value of such contingency motivations? The problem to date is that IPv4 address exhaustion is now not a novel message, and, so far, NATs have neutralized the urgency of the message. NATs are well understood, they appear to work reliably, applications work across NATs, and have influenced the application environment to such an extent that now no popular application can be fielded unless it can operate across NATs. For conventional client-server applications this represents no particular problem. For peer-to-peer based applications the rendezvous problem with NATs has been addressed through applications gateways and rendezvous servers. Even the variability of NAT behaviour is not a service provider liability, and it is left to applications to load additional functionality to detect specific NAT behaviour and make appropriate adjustments to the behaviour of the application. The conventional industry understanding to date is that NATs can work acceptably well within the application and service environment. In addition, for an ISP NATs are an externalized cost, as they are essentially funded and operated by the customer and not the ISP. The service provider's perspective is that considering that NATs have been so effective in externalizing the costs of IPv4 address scarcity from the ISP for the past five years, then surely they will continue to be effective for the next quarter. To date the costs of IPv4 address scarcity have

been passed across to the customer in the form of NAT-equipped CPE devices and to the application in the form of higher complexity in certain forms of application rendezvous. The ISP has not had to absorb these costs into its own costs of operation. From this perspective, the benefit of IPv6 that there is no further need for edge NATs, and the costs and associated complexities and vulnerabilities of NAT operation can be eliminated in an IPv6 environment hold no traction with the ISP in terms of marginal cost improvement from their current situation, given that they are not exposed to these costs and risks in any case at present.

The more general observation is that, at the current point in time, for the service provider industry, IPv6 has all the negative properties of revenue margin erosion with no immediate positive offsets. This observation lies at the heart of why the service provider industry has been so resistant to the call for widespread deployment of IPv6 services to date.

It appears that the current situation is not the outcome of a lack of information about IPv6, nor a lack of information about the forthcoming exhaustion of the IPv4 unallocated address pool. Nor is it the outcome of concerns over technical shortfalls or uncertainties in IPv6, as there is no evidence of any such technical shortcomings in IPv6 that prevent its deployment in any meaningful fashion. A more likely explanation for the current situation is an inability of a highly competitive deregulated industry to be in a position to factor in longer term requirements into short term business logistics.

## What Next?

It is now possible to return to considering some questions relating to IPv4 address exhaustion. Will the exhaustion of the current framework that supplies IP addresses to service providers cause all further demand for addresses to cease at that point? Or will exhaustion increase the demand pressure for addresses in response to various forms of panic and hoarding behaviours in addition to continued demand from growth?

The size and value of the installed base of the Internet using IPv4 is now very much larger than the size and value of incremental growth of the network. In address terms the routed Internet currently (as of 14 August 2008) spans 1,893,725,831 IPv4 addresses, or the equivalent of 112.2 /8 address blocks. Some 12 months ago the routed Internet spanned 1,741,837,080 IPv4 addresses or the equivalent of 103.8 /8 address blocks, representing a net annual growth of 10% in terms of advertised address space.

This leads to the observation that even in the hypothetical scenario where all further growth of the internet was forced to use IPv6 exclusively, while the installed base remains using IPv4, then it is highly unlikely that the core value of the Internet will shift away from its predominate IPv4 installed base in the short term.

Moving away from the hypothetical scenario, this implies that the relative size and value of new Internet deployments will be such that these new deployments may not have sufficient critical mass by virtue of their volume and value so as to be in a position to force the installed base to underwrite the incremental cost to deploy IPv6 and convert the existing network assets to dual stack operation in this timeframe. The corollary of this observation is that new Internet network deployments will need to communicate with a significantly larger and valuable IPv4-only network, at least initially. As IPv6 is not backward compatible with IPv4, this further implies that hosts in these new deployments will need to cause IPv4 packets with public addresses in their packet headers to be sent and received, either by direct deployment of dual stack or by proxies in the form of protocol-translating NATs. In either case the new network will require some form of access to public IPv4 addresses. In other words after exhaustion of the unallocated address pools, new network deployments will continue to need to use IPv4 addresses.

From this observation it appears highly likely that the demand for IPv4 addresses will continue at rates comparable to current rates across and after the IPv4 unallocated address pool exhaustion event. The exhaustion of the current framework of supply of IPv4 addresses will not trigger an abrupt cessation of demand for IPv4 addresses and this event will not cause the deployment of IPv6-only networks, at least in the short term of the initial years following IPv4 address pool exhaustion. It is therefore possible to indicate that immediately following this exhaustion event there will be a continuing market need for IPv4 addresses for deployment in new networks. While a conventional view is that this is likely to occur in a scenario of dual stacked environments, where the hosts are configured with both IPv4 and IPv6, and the networks are configured to also support the host operation of both protocols, it is also conceivable to envisage the use of deployments where hosts are configured in an IPv6 only mode and network equipment undertakes a protocol-translating NAT function. In either case the common observation is that we appear to have a continuing need for IPv4 addresses well after the event of IPv4 unallocated pool exhaustion, and IPv6 alone is no longer a sufficient response to this issue.

## How?

If there is continuing demand, then what is the source of supply in an environment where the current supply channel, namely the unallocated pool of addresses, is exhausted? The options for the supply of such IPv4 addresses are limited.

In the case of established network operators some IPv4 addresses may be recovered through the more intensive use of NATs in existing networks. A typical scenario of current deployment for ISPs involves the use of private address space in the customer's network and a NAT at the interface between the customer network and the service provider infrastructure (the Customer Premises Equipment, or CPE). One option for increasing the IPv4 address utilization efficiency could involve the use of a second level of NAT within the service provider's network, or the so-called "carrier-grade" NAT option [Nishitani 2008]. This option has some attraction in terms of increasing the port density utilization of public IPv4 addresses, by effectively sharing the port address space of the public IPv4 address across multiple CPE NATs, allowing the same number of public IPv4 addresses to be utilized across a larger number of end customer networks.

The potential drawback of this approach is that of added complexity in NAT behaviour for applications, given that an application may have to traverse multiple NATs, and the behaviour of the compound NAT scenario becomes in effect the behaviour of the most conservative of the NATs in the path in terms of binding times and access. Another potential drawback is that some applications have started to take up the option of use of multiple simultaneous transport sessions in order to improve the performance of the download of multi-part objects. For single-level CPE NATs with more than 60,000 ports to be used for the customer network this application behaviour had little impact, but the presence of a carrier NAT servicing a large number of CPE NATs may well restrict the number of available ports per connection, which, in turn impacts on the utility of various forms of applications that operate in this highly parallel mode. Allowing for a peak simultaneous demand level of 500 ports per customer provides a potential use factor of some 100 customers per IP address. Given a large enough common address pool this may be further improved by statistical multiplexing by a factor of 2 or 3, allowing for between 200 and 300 customers per NAT address. Of course such approximations are very coarse, and the engineering requirement to achieve such a high level of NAT utilization would be significant. Variations on this engineering approach are possible in terms of the internal engineering of the ISP network and the control interface between the CPE NATs and the ISP equipment, but the maximal ratio of 200 - 300 customers per public IP address appears to be a reasonable upper bound without unduly impacting on application behaviours.

Another option is based on the observation that of the currently allocated addresses some 42% of these addresses, or the equivalent of some 49 /8 address blocks, are not advertised in the

inter-domain routing table, and are presumed to be either used in purely private contexts, or are currently unused. This pool of addresses could also be used as a supply stream for future address requirements, and while it may be overly optimistic to assume that the entirety of this unadvertised address space could be used in the public Internet, it is possible to speculate that a significant amount of this address pool could be used in such a manner, given the appropriate incentives. Speculating even further, if this were used in the context of intensive carrier-grade NATs with an achieved average deployment level of, say, 10 customers per address, an address pool of 40 /8s would be capable of sustaining some 7 billion customer attachments.

Of course no such recovery option exists for new entrants, and in the absence of any other supply option this situation will act as an effective barrier to entry into the Internet Service Provider market. In cases where the barriers to entry effectively shut out new entrants there is a strong trend for the incumbents to form cartels or monopolies and extract monopoly rentals from their clients. However it is unlikely that the lack of supply will be absolute, and a more likely scenario is that addresses will change hands in exchange for money. Or, in other words, it is likely that such a situation will encourage the emergence of markets in addresses. Existing holders of addresses have the option to monetise all or part of their held assets, and new entrants, and others, have the option to bid against each other for the right to use these addresses. In such an open market the most efficient utilization application would tend to be able to offer the highest bid, which, in an environment dominated by scarcity would tend to provide strong incentives for deployment scenarios that offer high levels of address utilization efficiency.

It would therefore appear that there are options available to this industry to increase the utilization efficiency of deployed address space, and thereby generate pools of available addresses for new network deployments. However, the motive for so doing will probably not be phrased in terms of altruism or alignment to some perception of common good. Such motives sit uncomfortably within the commercial world of the deregulated communications sector. Nor will it be phrased in terms of regulatory impositions. Not only will it take many years to halt and reverse the ponderous process of public policy and its expression in terms of regulatory measures, it is also the case that the "common good" objective here transcends the borders of regulatory regimes. This consideration tends to leave this argument with one remaining mechanism that will motivate the industry to significantly increase the address utilization efficiency is that of monetizing addresses and exposing the costs of scarcity of addresses to the users of addresses. The corollary of this approach is the use of markets to perform the address distribution function, creating a natural pricing function based on levels of address supply and demand.

In the second part of this article I'd like to indulge in a little speculation on addresses and look at the various aspects of using a market to undertake the address distribution function post exhaustion.

## References

- [Bagnulo 2008] NAT64/DNS64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, M. Bagnulo, P. Matthews, I. van Beijnum, work in progress, Internet Draft, draft-bagnulo-behave-nat64-00.txt, June 2008.
- [Huston 2008] The IPv4 Internet Report, G. Huston, August 2008.  
<http://ipv4.potaroo.net>
- [Nishitani 2008] Carrier Grade Network Address Translator (NAT) Behavioral Requirements for Unicast UDP, TCP and ICMP, T. Nishitani, S. Miyakawa, work in progress, Internet draft, draft-nishitani-cgn-00.txt, July 2008.
- [Solensky 1990] Internet Growth, F. Solensky, Steering Group Report, p61, Proceedings of the 18<sup>th</sup> IETF Meeting, August 1990.  
<http://www.ietf.org/proceedings/prior29/IETF18.pdf>
- [TCP/IP 1988] TCP/IP Mailing List, Message Thread: "Running out of Internet Addresses", November 1988.  
[http://www-mice.cs.ucl.ac.uk/multimedia/misc/tcp\\_ip/8813.mm.www/index.html#121](http://www-mice.cs.ucl.ac.uk/multimedia/misc/tcp_ip/8813.mm.www/index.html#121)
- [RFC1380] IESG Deliberations on Routing and Addressing, P. Gross, P. Almquist, RFC1380, November 1992.
- [RFC1631] The IP Network Address Translator (NAT), K. Egevang, P. Francis, RFC1631, May 1994.
- [RFC2766] Network Address Translation - Protocol Translation (NAT-PT), G. Tsirtsis, P. Srisuresh, RFC2766, February 2000.
- [RFC4966] Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status, C. Aoun, E. Davies, RFC4966, July 2007.

---

## Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

---

## About the Author

GEOFF HUSTON is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He graduated from the Australian National University with a B.Sc, and M.Sc. in Computer Science. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

<http://www.potaroo.net>