



IPv4 - How long do we have?

December 2003

Geoff Huston

One of those stories that keeps on appearing from time to time is the claim that somewhere in the world, or even all over the world, we are "running out of IP addresses," referring to the consumption of unallocated IPv4 addresses [1]. In one sense this is a pretty safe claim, in that the IPv4 address pool is indeed finite, and, as the IPv4 Internet grows it makes continual demands on previously unallocated address space. So the claim that the space will be exhausted at some time in the future is a relatively safe prediction. But the critical question is not "if" but "when," because this is a question upon which many of our current technology choices are based.

Given this revived interest in the anticipated longevity of the IPv4 address space, it is timely to revisit a particular piece of analysis that has been a topic of some interest at various times over the past decade or more. The basic question is: "How long can the IPv4 address pool last in the face of a continually growing network?" This article looks at one approach to attempt to provide some indication of "when." Like all predictive exercises, many assumptions have to be made, and the approach described here uses just one of numerous possible predictive models—and, of course, the future is always uncertain.

The IPv4 Address Space

The initial design of IPv4 was extremely radical for its time in the late 1970s. Other contemporary vendor-based computer networking protocols were designed within the constraints of minimizing the packet header overhead in order to improve the data payload efficiency of each packet. At the time address spans were defined within the overall assumption that the networks were deployed as a means of clustering equipment around a central mainframe. In many protocol designs 16 bits of address space in the packet headers was considered to be extravagant. To use a globally unique address framework of 32 bits to address network hosts was, at the time, a major shift in thinking about computer networks from a collection of disparate private facilities into a truly public utility.

To further add to the radical nature of the exercise, the Internet Network Information Center was prepared to hand out unique blocks of this address space to anyone who submitted an application. Address deployment architectures in other contemporary protocols did not have the address space to support such address distribution functions, nor did they even see a need for global uniqueness of computer network addresses. Network administrators numbered their isolated corporate or campus networks starting at the equivalent of "1," and progressed onward from there. Obviously network splits and mergers caused considerable realignment of these private addressing schemes, with consequent disruption to the network service.

By comparison, it seemed, the address architecture of the Internet was explicitly designed for interconnection. But even with 32 bits to use in an address field, getting the right internal structure for addresses is not as straightforward as it may initially seem.

The Evolution of the IPv4 Address Architecture

IP uses the address to express two aspects of a connected device: the identity of this device (endpoint identity) and the location within the network where this device can be reached (location or forwarding identity). The original IP address architecture used the endpoint identity to allow devices to refer to each other in end-to-end application transactions, whereas within the network the address is used to direct packetforwarding decisions. The address was further structured into two fields: a *network* identifier and a *host* identifier within that network. The first incarnation of this address architecture used a division at the first octet: the first 8 bits were the network number and the following 24 bits were the host identifier. The underlying assumption was one of deployment across a small number of very large local networks. This view was subsequently refined, and the concept of a class-based address architecture was devised for the Internet. Half of the address space was left as a 8/24-bit structure, called the *Class A* space (allowing for up to 127 networks each with 16,777,216 host identities). A quarter of the remaining space used a 16/16-bit split (allowing for up to 16,128 networks, each with up to 65,536 hosts), defining the *Class B* space. A further eighth of the remaining space was divided using a 24/8-bit structure (allowing for 2,031,616 networks, each with up to 256 hosts), termed the *Class C* space. The remaining eighth of the space was held in reserve.

This address scheme was devised in the early 1980s, and within a decade it was pretty clear that there was a problem with impending exhaustion. The reason was an evident run on Class B addresses. Although very few entities could see their IP network spanning millions of computers, the personal desktop computer was now a well-established part of the landscape, and networks of just 256 hosts were just too small. So if the Class A space was too big, and the Class C too small, then Class B was the only remaining option. In fact, the Class B blocks were also too large, and most networks that used a Class B address consumed only a few hundred of the 65,535 possible host identities within each network. The addressing efficiency of this arrangement was very low, and a large amount of address space was being consumed in order to number a small set of devices. Achieving even a 1 percent host density (expressed as a ratio of number of addressed hosts to the total number of host addresses available) was better than normal at the time, and 10 percent was considered pretty exceptional.

Consequently, Class B networks were being assigned to networks at an exponentially increasing rate. Projections from the early 1990s forecast exhaustion of the Class B space by the mid-1990s. Obviously there was a problem, and the *Internet Engineering Task Force* (IETF) took on the task of finding some solutions. Numerous responses were devised by the IETF.

As a means of mitigation of the immediate problem, the IETF altered the structure of an IP address. Rather than having a fixed-length network identifier of 8, 16, or 24 bits, the network part of the address could be any length at all, and a network identifier was now the couplet of an IP address field containing a network part and the bit length of the network part. The boundary between the network and host part could change across the network, so rather than having "networks" and "subnetworks" as in the class-based address architecture, there was the concept of a variable length network mask. This was termed the "classless" address architecture (or "CIDR"), and the step was considered to be a short-term expediency to buy some additional time before address exhaustion. The longer-term plan was to develop a new IP architecture that could encompass a much larger connectivity domain than was possible with IPv4.

We now have IPv6 as the longer-term outcome. But what has happened to the short-term expediency of the classless address architecture in IPv4? It appears to have worked very well indeed so far, and now the question is: how long can this supposedly short-term solution last?

Predictions of Address Consumption

Predicting the point of IPv4 address exhaustion has happened from time to time since the early 1990s within the IETF [2]. The initial outcomes of these predictive exercises were clearly visible by the mid-1990s: the classless address architecture was very effective in improving the address utilization efficiency, and the pressures of ever-increasing consumption of a visibly finite address resource were alleviated. But a decade after the introduction of CIDR addressing, it is time to understand where we are heading with the consumption of the underlying network address pool.

Dividing up the Address Space

There are three stages in address allocation. The pool of IP addresses is managed by the *Internet Assigned Numbers Authority* (IANA). Blocks of addresses are allocated to *Regional Internet Registries* (RIRs), who in turn allocate smaller blocks to *Local Internet Registries* (LIRs) or *Internet Service Providers* (ISPs).

Currently 3,707,764,736 addresses are managed in this way. It is probably easier to look at this in terms of the number of "/8 blocks," where each block is the same size as the old Class A network, namely 16,777,216 addresses. The total address pool is 221 /8s, with a further 16 /8s reserved for multicast use, 16 /8s held in reserve, and 3 /8s designated as not for use in the public Internet.


In looking at futures, there are three sources of data concerning address consumption:

- How quickly is the IANA passing address blocks to the RIRs, and when will IANA run out?
- How quickly are the RIRs passing address blocks to LIRs, and when will this run out?
- How much address space is actually used in the global Internet, and how quickly is this growing? When will this run out?

The IANA Registry


So the first place to look is the IANA registry file [3] . This registry reveals that of these 221 /8 blocks, 89 /8 blocks are still held as unallocated by the IANA, 129.9 /8 blocks have been allocated, and the remaining 2.1 / 8 blocks are reserved for other uses. The IANA registry also includes the date of allocation of the address block, so it is possible to construct a time series of IANA allocations, as shown in Figure 1.

Figure 1: IANA Allocated IPv4 /8 Address Blocks



Interestingly, there is nothing older than 1991 in this registry. This exposes one of the problems with analyzing registry data, in that there is a difference between the current status of a registry and a time-stamped log of the transactions that were made to the registry over time. The data published by the IANA is somewhere between the two, and the log data is incomplete; in addition, the current status of some address blocks is unclear. It appears that the usable allocation data starts in 1995. So if we take the data starting from 1995 and perform a linear regression to find a best fit of an exponential projection, it is possible to make some predictions as to the time it will take to exhaust the remaining unallocated 89 /8s. (Figure 2).

Figure 2: IANA Allocated IPv4 /8 Address Blocks



It is worth a slight digression into the method of projection being used here. The technique is one of using a best fit of an exponential growth curve to the data. The underlying

assumption behind such a projection is that the growth rate of the data is proportional to the size of the data, rather than being a constant rate. In network terms, this assumes that the rate of consumption of unallocated addresses is a fixed proportion of the number of allocated addresses, or, in other words, the expansion rate of the network is a proportion of its size, rather than being a constant value. Such exponential growth models may not necessarily be the best fit to a network growth model, although the data since 1995 does indicate an underlying exponential growth pattern. Whether this growth model will continue into the future is an open issue.

The projection of 2019 as the date for consumption of the unallocated address space using this technique is perhaps surprising, because it seems that the network is bigger now than ever, yet the amount of additional address space required to fuel further accelerating growth for a further decade is comparatively small. This is true for many reasons, and the turning point when these aspects gained traction in the Internet appeared to be about 1995. They include:

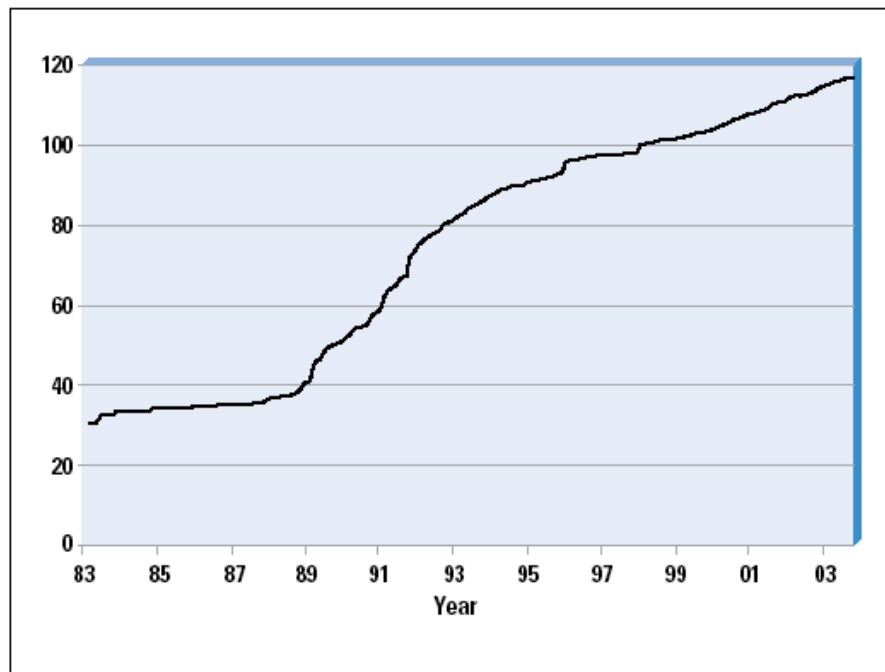
- The first 1.6 billion addresses (equivalent to some 100 /8 blocks) were allocated using the class-based address architecture. Since this date address allocation has used a classless architecture, and this has enabled achievement of significantly improved efficiencies in using the address space.
- The RIRs came into the picture, and started using conservation-based policies in address allocations. The RIR process requires all address applicants to demonstrate that they can make efficient and effective use of the address space, and this has dampened some of the wilder sets of expectations about the address requirements of an enterprise.
- Address compression technologies became widely deployed. Dynamic *Network Address Translation* (NAT) devices have, for better or worse, become a common part of the network landscape. NAT devices allow large "semi-private" networks to use a very small pool of public addresses as the external view of the network, while using private address space within the network. *Dynamic Host Configuration Protocol* (DHCP) has allowed networks to recycle a smaller pool of addresses across a larger set of intermittently connected devices.

Whether these factors will continue to operate in the same fashion in the future is an open question. Whether future growth in the use of public address space operates from a basis of a steadily accelerated growth is also an open question. The assumption made in this exercise is that the projections depend on continuity of effectiveness of the RIR policies and their application, continuity of technology approaches, and absence of disruptive triggers. Although the RIRs have a very well-regarded track record and there are strong grounds for confidence that this will continue, obviously the latter two assumptions about technology and disruptive events are not all that comfortable. With that in mind, the next step is to look at the RIR assignment data.

The RIR Registries


The RIRs also publish a registry of their transactions in "stats" files. For each currently allocated or assigned address block the RIRs have recorded, among other items, the date of the RIR assignment transaction that assigned an address block to a LIR or ISP. Using this data we can break up the 129.9 /8 blocks further, and it is evident that the equivalent of 116.7 /8 blocks have been allocated or assigned by the RIRs, and the remaining space, where there is no RIR allocation or assignment record, is the equivalent of 13.2 /8 blocks. These transactions can again be placed in a time series, as shown in Figure 3.

Figure 3: RIR Assigned IPv4 /8 Address Blocks



The post-1995 data used to extrapolate forward using the same linear regression technique described previously to find a curve of best fit using the same underlying growth model assumptions yields:

Figure 4: RIR Assigned IPv4 /8 Address Blocks—Projection




This form of extrapolation gives a date of 2026 for the time at which the RIRs will exhaust the number pool. Again the same caveats about the use of this approach as a reliable predictor apply here, and the view forward is based on the absence of large-scale disruptions, or some externally induced change in the underlying growth models for address demand.

The BGP Routing Table

When addresses are assigned to end networks, the expectation is that these addresses will be announced to the network in the form of routing advertisements. So some proportion of these addresses is announced in the Internet routing table. The next task is to establish the trends of the amount of address space covered by the routing table. The approach used has been to take a single view of the address span of the Internet. This is the view from one point, inside the AS1221 network operated by Telstra.

The data as of October 2003 shows that some 29 percent of the total IPv4 address space is announced in the Border Gateway Protocol (BGP) routing table, whereas 17 percent has been allocated to an end user or LIR but is not announced on the public Internet as being connected and reachable. A total of 5 percent of the address space is held by the RIR's pending assignment or allocation (or at least there is no RIR recorded assignment of the space), while 35 percent of the total space remains in the IANA unallocated pool. A further 8 percent of the space is held in reserve (Figure 5).

Figure 5: IPv4 /8 Address Space



This BGP data is based on an hourly inspection of the amount of address space advertised within the Internet routing table. The data collection commenced in late 1999, and the data gathered so far is shown in Figure 6. The problem with this data is that there is some considerable amount of fluctuation in the amount of address space advertised over time. The major step changes are due to a small number of /8 advertisements that periodically are announced and withdrawn in BGP. In order to obtain reasonable data for generating projections, some noise reduction on this data needs to be undertaken. The approach used has been to first filter the data using a constant value of 18 /8 prefix announcements, and then use a sliding average function to create a smoothed time series. This is indicated in Figure 7.

The critical issue when using this data for projection is to determine what form of function can provide a best fit to the data. A good indication of the underlying trends in the data can be found by analyzing the first-order differential of the data. An underlying increasing growth model would have an increasing first-order differential, whereas a decreasing growth model would have a negatively inclined differential. A least-squares best-fit analysis of the data shows that the growth rates have not been consistent over the past three years. A reasonable fit for this data appears to be a constant growth model, or a linear growth projection, with a consumption rate of some 3 /8 blocks per year.

Figure 6: Advertised IPv4 /8 Address Space (/8 Blocks)





Figure 7: Smoothed IPv4 /8 Advertised Address Blocks



Combining the Three Views


One question remains before we complete the projections for IPv4 address space. There are 43.3 /8 blocks, or some 17 percent of the total IPv4 address space that has been allocated for use, but is not visible in the Internet routing table. This is a very significant amount of address space, and if it is growing at the same rate as the advertised space, then this will have a significant impact on any overall model of consumption of the use of address space.

The question here is whether this "invisible" address pool is a legacy of the address allocations policies in place before the RIR system came into operation in the mid 1990s, or

some intrinsic inefficiency in the current system. If it is the latter, then it is likely that this pool of unannounced addresses will grow in direct proportion to the growth in the announced address space, whereas if it is the former, then the size of the pool will remain relatively constant in the future.


We can look back through the RIR allocation data and look at the allocation dates of unannounced address space (Figure 8). This view indicates that the bulk of the space is a legacy of earlier address allocation practices, and that since 1997, when the RIR operation was fully established, there is an almost complete mapping of RIR allocated address space to BGP routing announcements. The recent 2003 data indicates that there is some lag between recent allocations and BGP announcements, most probably due to the time lag between an LIR receiving an allocation and subsequent assignments to end users and advertisement in the routing table.

Figure 8: Age Distribution of Unadvertised Address Blocks (/8 Address Blocks)



This confirms that in recent years all the address space that has been assigned by the RIRs appears in the Internet routing table, implying that projections of the amount of address space advertised in the routing table is a good correlation to projections of address space consumption. With this in mind it is now possible to construct a model of the address distribution process, working backward from the BGP routing table address size. From the sum of the BGP table size and the LIR holding pool, we can derive the total RIR-managed address pool. To this number is added the RIR holding pool low size and its low threshold where a further IANA allocation is required. This allows a view of the entire system, projected forward over time, where the central driver for the projection is the growth in the network itself, as described by the size of the announced IPv4 address space. This is shown in Figure 9.

Figure 9: IPv4 Projections of Address Consumption



It would appear that the point of effective exhaustion is the point where the RIRs exhaust available address space to assign. In this model, RIR exhaustion of the unallocated address pool would occur in 2037.

Uncertainties

Of course such projections are based on the underlying assumption that tomorrow will be much like today, and the visible changes that have occurred in the past will smoothly translate to continued change the future. This assumption obviously has some weaknesses, and many events could disrupt this prediction.

Some disruptions could be found in technology evolution. An upward shift in address take-up rates could occur because of an inability of NAT devices to support emerging popular applications. Widespread deployment of peer-to-peer applications implies the need for persistent address presentation, which may imply greater levels of requirement for public address space. The use of personal mobile IP devices (such as PDAs in their various formats) using public IPv4 addresses would place a massive load on the address space, simply because of the very large volumes associated with deployment of this technology [4].

Other disruptions have a social origin, such as the boom and bust cycle of Internet expansion in recent years. Another form of disruption in this category could be the adoption of a change in the distribution function. The current RIR and LIR distribution model has been very effective in limiting the amount of accumulation of address space in holding pools, and allocating addresses based on efficiency of utilization and conformance to the routing topology of the network.

Many other forms of global resource distribution use a geopolitical framework, where number blocks are passed to national entities, and further distribution is a matter of local policy [5]. The disruptive nature of such a change would be to immediately increase the number of "holding" points in the distribution system, locking away larger pools of address space from being deployed and advertised and generating a significant upward change in

the overall address consumption rates due to an increase in the inefficiency of the altered distribution function.

The other factor to be aware of is the steadily decreasing "buffer" of unallocated addresses that can be used to absorb the impacts of a disruptive change in address consumption rates. Although at present some 60 percent of the address space—or some 2.6 billion addresses—are available in the unallocated address pools or held in reserve, this pool will reduce over time. If a disruptive event is, for example, a requirement to directly address some 500 million devices, then such an event would reduce the expectancy of address space availability by some years, assuming it occurred within the period when sufficient address space remains to meet such a surge of demand.

The other source of uncertainty is that this form of predictive modeling assumes that the ratios of actual connected devices and the amount of address space deployed to service this device pool remain relatively constant.

This model also assumes some form of continuity of current address allocation policies. This is not a likely scenario, because it is likely that address policies will reflect some notion of balance between the level of current demand against future demands. As the unallocated address pool shrinks it is possible that policies will alter to express the increased level of competitive demand for the remaining resource. Consumption rates would be moderated by such a change in allocation policy. The commonly cited intended evolutionary path for the Internet is to a transition to ubiquitous use of IPv6, and at some point in that transition process it is reasonable to assume that further demands for IPv4 space will dwindle. It may be that at such a "crossover" time allocation policies may then be altered to reflect a drop in both current and future demands for IPv4 address space.

In attempting to assess the possible future path of address allocation policies, it is also evident that, from a market rationalist perspective, there is a certain contrivedness about the current address allocation process. The current address management system assumes a steady influx of new addresses to meet emerging demands, and the overall address utilization efficiency is not set by any form of market force, but by the outcomes of the application of RIR address allocation policies to new requests for address space. A market rationalist could well point to the use of market price as a means of determining the most economically efficient form of utilization of a commodity product. Such a position is based on the observation that the way that the consumer chooses between alternative substitutable services is by a market choice that is generally price sensitive.

If price is removed from an IPv4 address market, the choices made by market players are not necessarily the most efficient choices, and some would argue that the current situation underprices IPv4 at the expense of IPv6.

However, in venturing into these areas we are perhaps straying a little too far from exploring the degree of uncertainty in these predictive exercises. A discussion of the interaction between various forms of distribution frameworks and likely technology outcomes is perhaps a topic for another time.

So just how long does IPv4 have?

The assumptions used here include assuming that the trends in the growth in the advertised space are directly proportional to the future consumption rates for IP addresses, and that the constant growth model remains a best fit for this time series of data. It also assumes a continuation of the current utilization efficiency levels in the Internet, a continuing balance between public address utilization and the use of various forms of address compression, and continuity of current address allocation policies, as well as the absence of highly disruptive events. With all this in mind, then it would appear that the

IPv4 world, in terms of address availability, could continue for up to another three decades or so without reaching any fixed boundary of exhaustion.

But it must be remembered that each of these assumptions is relatively sweeping, and to combine them as we have done here is pushing the predictive exercise to its limits, or possibly beyond them. Three decades out is way over the event horizon for any form of useful prediction for the Internet, so if we restrict the question to at most the next five to eight years, then we can answer with some level of confidence that, in the absence of any significant disruptions to the current deployment model of the Internet, there is really no visible evidence that IPv4 will exhaust its address pool by 2010, based on the available address consumption data.

Data Sources

IANA IPv4 Address Registry: <http://www.iana.org/assignments/ipv4-address-space>

Registry Stats report files:

APNIC:	ftp://ftp.apnic.net/pub/apnic/stats
ARIN:	ftp://ftp.arin.net/pub/stats
LACNIC:	ftp://ftp.lacnic.net/pub/stats
RIPE NCC:	ftp://ftp.ripe.net/ripe/stats

BGP Address Data: <http://bgp.potaroo.net>

Notes

- [1] "Tackling the net's number shortage." BBC News, World Edition, 26 October 2003. The item starts with the claim: "BBC ClickOnline's Ian Hardy investigates what is going to happen when the number of net addresses—Internet Protocol numbers— runs out sometime in 2005."
<http://news.bbc.co.uk/2/hi/technology/3211035.stm>
- [2] The work was undertaken in the *Address Lifetime Expectations* (ALE) Working Group of the IETF in 1993 - 1994. The final outcome from this effort was reported from the December 1994 meeting of this group: "Both models currently suggest that IPv4 addresses would be depleted around 2008, give or take three years."
- [3] This registry is online at:
<http://www.iana.org/assignments/ipv4-address-space>
- [4] On the other hand, it is evident that the growth of the Internet in recent years has been fueled by the increasing prevalence of NAT devices. In order for applications to be accepted into common use in today's Internet, they need to be able to function through various NAT-based constraints, and increasing sophistication of applications in operating across NAT devices is certainly evident today.
- [5] Such a geopolitical distribution system is used in the E.164 number space for telephony ("ENUM").

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
