United States Department of Commerce
National Institute of Standards and Technology
National Telecommunications and Information Administration

Request for Comments on Deployment of Internet Protocol, Version 6

Comments Submitted by

Geoff Huston,

Research Scientist, Asia Pacific Network Information Centre

These comments are intended to address the issues raised in Section II A, regarding estimates of how many IPv4 address have been allocated, how many are still available, and how long the remaining address space will continue to meet the needs of users in the United States, as well as users in other countries around the world.

The Request for Comment cites in a footnote a presentation I made to the APNIC Open Policy Meeting in August 2003. I would like to submit the following commentary as an explanation of the methodology used, as well as providing an update on the projections of IPv4 consumption incorporating data up to the end of February 2004.

Summary

This study examines the available data concerning the consumption of IPv4 addresses, looking at the allocation data published by the Internet Assigned Numbers Authority (IANA), the allocation data published by the Regional Internet Registries (RIRs) and the address span encompassed by advertisements within the Internet global routing system.

The study examines the trends of consumption since and examines whether the current consumption trends are best modeled using an exponential or a linear growth trend. The study concludes that the best model of the current consumption rates of address space can be found in an analysis of the global routing table, and the best fit model is that of linear growth.

It concludes that the IPv4 world, in terms of address availability, could continue up to 2030 without reaching any fixed boundary of address exhaustion. However, in forming this model a number of relatively sweeping assumptions have been made, and to combine them as is 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 If the question is restricted to the next decade (i.e. up to 2014), then we can answer with some level of confidence that there is really no visible evidence of IPv4 exhausting its address pool based on the available address consumption data.

The conclusion drawn here is that if the prime driver for the adoption of IPv6 lies in a looming shortage of public IPv4 addresses, then this is not a near term prospect, and any impetus for the adoption for IPv6 would need to be based on factors other than an imminent exhaustion of IPv4 address space.

Introduction

The IPv4 address pool is of finite size, and, as the IPv4 Internet grows it makes continual demands on previously unallocated address space. It is possible to pose the question of "How long can the IPv4 address pool last in the face of a continual growth in the demand for addresses?"

This submission looks at one approach to attempt to provide some indication of when the IPv4 address pool is likely to be exhausted as a consequence of this continued demand for addresses, and describes the various assumptions that have been made in order to make this prediction.

Predictions of Address Consumption

Predicting the point of IPv4 address exhaustion was first undertaken within the IETF in the early 1990s[1] as part of an exercise in looking at the future requirements of the Internet in terms of both address space and routing capabilities. The initial outcomes of the consequent IETF activities to alleviate the pressure on the address pool were clearly visible by the mid-1990's: the Classless address architecture (Classless Inter-Domain Routing, or "CIDR") was very effective in improving the address utilization efficiency, and the pressures of ever-increasing consumption of a visibly finite address resource were alleviated. The IETF actions to address the longer term requirements of address space was of course IPv6.

Distributing the Address Space

There are three stages in address distribution. The pool of IP addresses is managed by the Internet Assigned Numbers Authority (IANA). Blocks of address space have been defined through standards actions of the IETF for use as global unicast addresses, private-use unicast addresses, multicast addresses, reserved for future actions, and other blocks are marked as reserved, and are not to be used outside of local host contexts.

Currently there are 3,707,764,736 addresses that are managed as global unicast addresses. 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 global unicast address pool is 221 /8s, with a further 16 /8s reserved for multicast use, 16 /8s held in reserve, and 3 /8s are designated as not for use in the public Internet. (Figure 1)

---

1 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." This prediction was documented in RFC 1752. No further predictions have been undertaken within the IETF on this topic since that date.

.

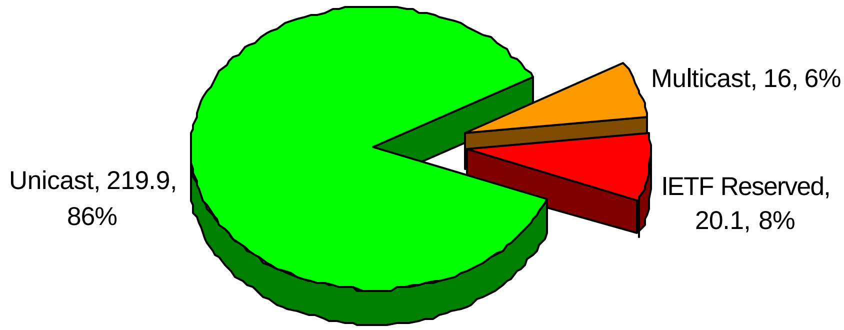

Figure 1 – Breakdown of IPv4 address space

The process of address management is that the IETF defines the global unicast address pool to the IANA. IANA allocates blocks of global unicast addresses to Regional Internet Registries (RIRs). The RIRs allocate address blocks to Local Internet Registries (LIRs) or Internet Service Providers (ISPs). These address blocks are then advertised in the Internet's Border Gateway Protocol (BGP) routing table (Figure 2).

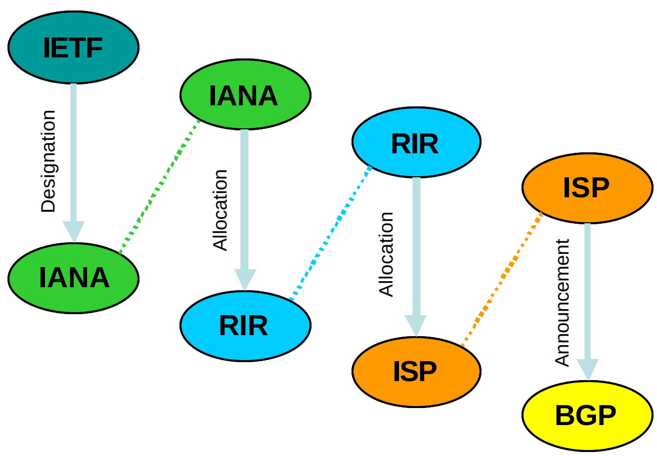

Figure 2 – Address Distribution Framework

In looking at futures there are three distinct questions 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? The related question is when will demand for addresses outstrip the available address pool?

The IANA Registry

The first data set to examine is the IANA registry file[2]. This registry indicates that of these 221 /8 blocks 86 /8 blocks are still held as unallocated by the IANA, 133.9 /8 blocks have been allocated. (Figure 3).

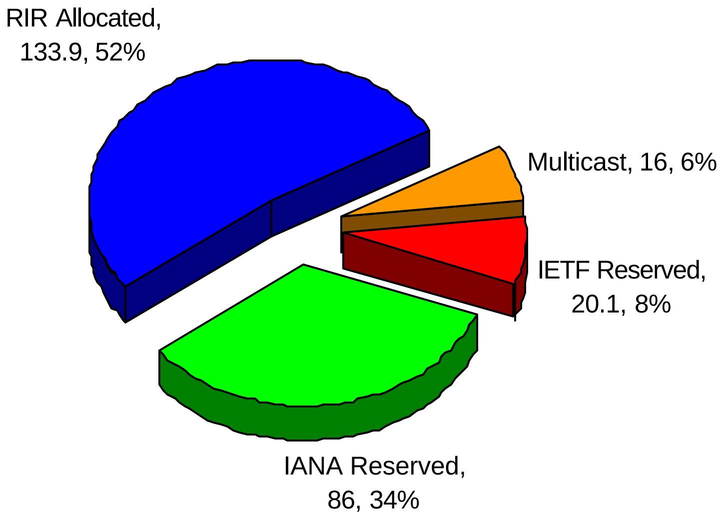

Figure 3 – Address Pool with RIR Allocations

The IANA registry also includes the date of allocation of the address block, so it's possible to construct a time series of IANA allocations. (Figure 4)

[2] This registry is online at http://www.iana.org/assignments/ipv4-address-space

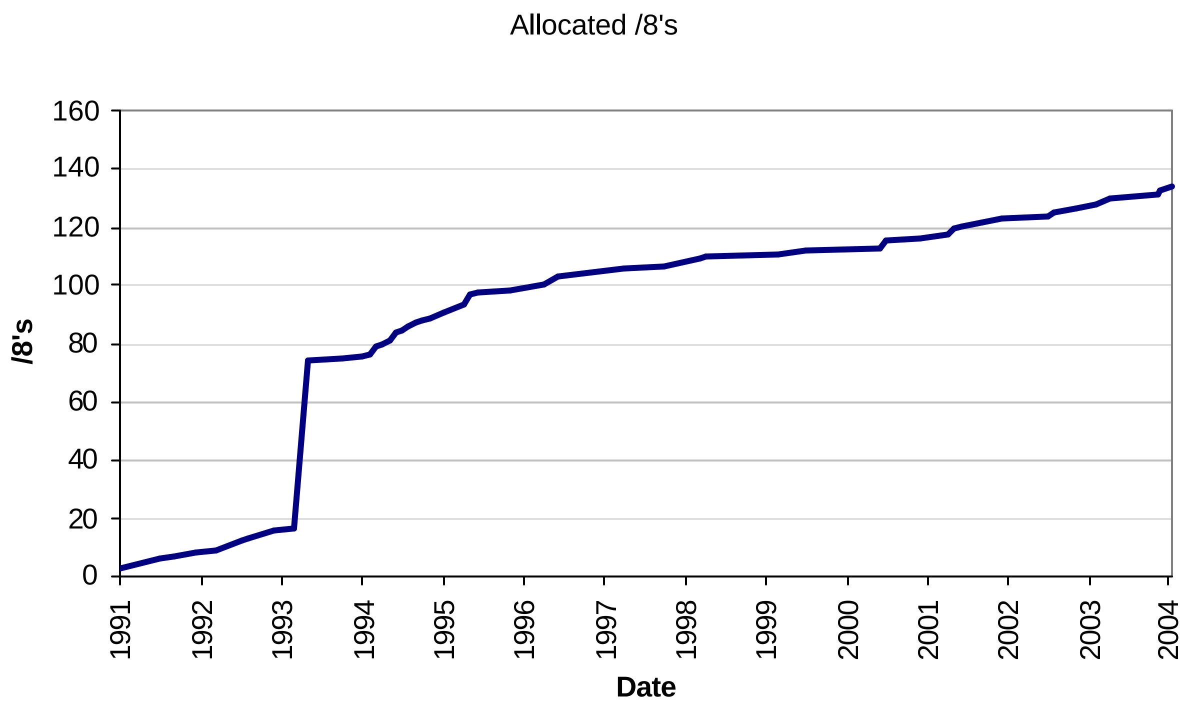

Figure 4 - Time Series of IANA Allocations

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. The log of data is incomplete, and the current status of some address blocks is unclear. It appears that reliable 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, its possible to make some predictions as to the time it will take to exhaust the remaining unallocated 89 /8s. (Figure 5).

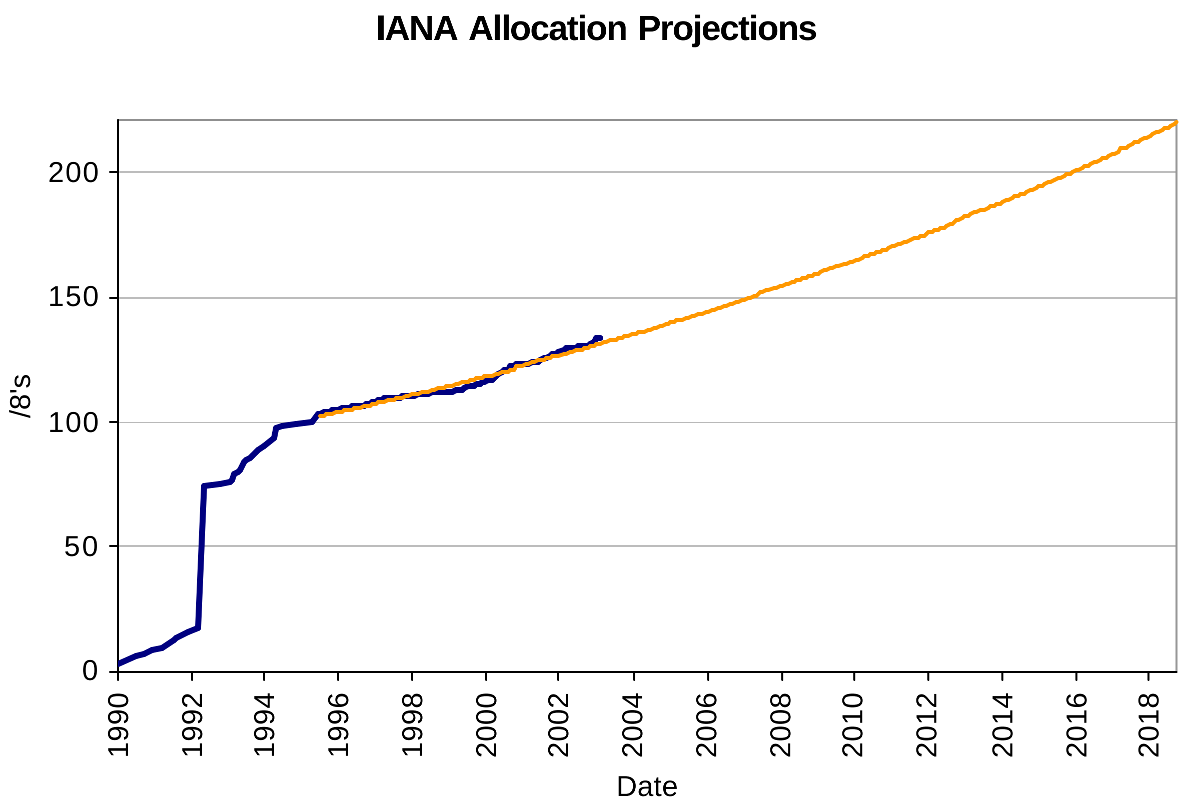

Figure 5 - Projection of IANA Allocations

It is worth a slight digression into the method of projection being used here. The technique is one of using a least squares technique to find a best fit of an exponential growth curve to the data. The underlying assumption behind such exponential projections 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 some constant 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 appear to correlate reasonably to an underlying exponential growth pattern. Whether this growth model will continue into the future is an open issue. It is also possible to fit a linear growth model to the data, in which case the projected exhaustion point shifts from 2019 to 2027.

The projection of 2019 as the date for exhaustion of the IANA unallocated address space pool using this exponential growth model is perhaps surprising, as its 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.

There are a number of reasons why this is the case, and the turning point when these aspects gained traction in the Internet appeared to be around 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 allowed for significantly improved efficiencies to be achieved in using the address space.
- The Regional Internet Registries (RIRs) assumed responsibility for address allocation, and started using conservation-based policies in determining allocations. The RIR process requires each address applicant to demonstrate that they can make efficient and effective use of the address space, and this dampened some of the wilder sets of expectations about an enterprise's address requirements.
- Address Compression technologies became widely deployed. Dynamic Network Address Translation devices (NATs) have, for better or worse, become a common part of the network landscape. NATs allow large 'semi-private' networks to use a very small poll 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 accelerating growth is also an open question. The assumption made in this exercise is that the projections depend in continuity of effectiveness of the RIR policies and their application, continuity of technology approaches and absence of disruptive triggers. While the RIRs have a very well regarded track record and there are strong grounds for confidence that this will continue, the latter two assumptions about technology and disruptive events are not as certain.

The Regional Internet Registries

The RIRs also publish a registry of their transactions in "delegated" files. For each currently allocated address block the RIR registry contains the date of the RIR allocation transaction[3], and the country or region that was the location of the LIR or ISP. Using this data we can break up the 133.9 /8 blocks further, and its evident that the equivalent of 120.8 /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 13.1 /8 blocks. (Figure 6)

[3] There are some anomalies in the consistency of registry data relating to the date of the allocation transaction. Some of these anomalies are found in historical allocations where the registry date is the date when the allocation entry was re-homed to the local RIR. Others are where the original allocation date does not appear to have been recorded in any location.

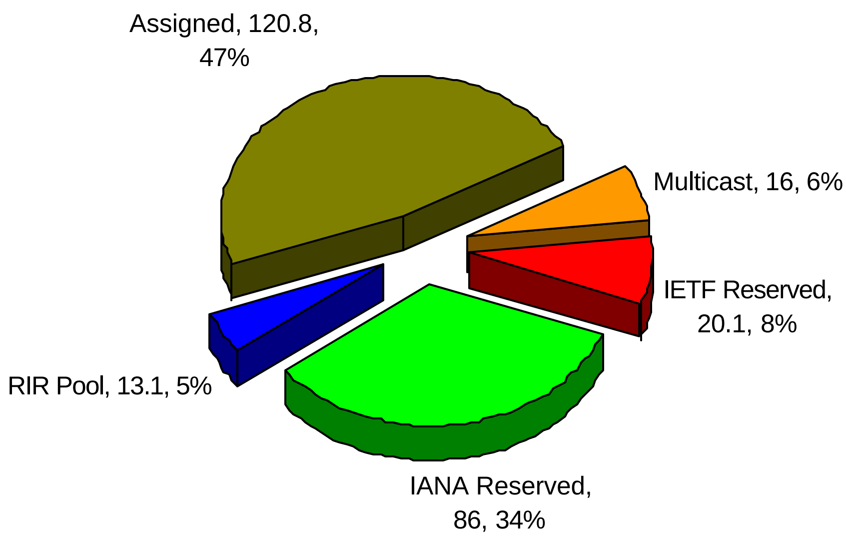

Figure 6 – RIR Assigned Address Pool

These transactions can again be placed in a time series, as shown below (Figure 7).

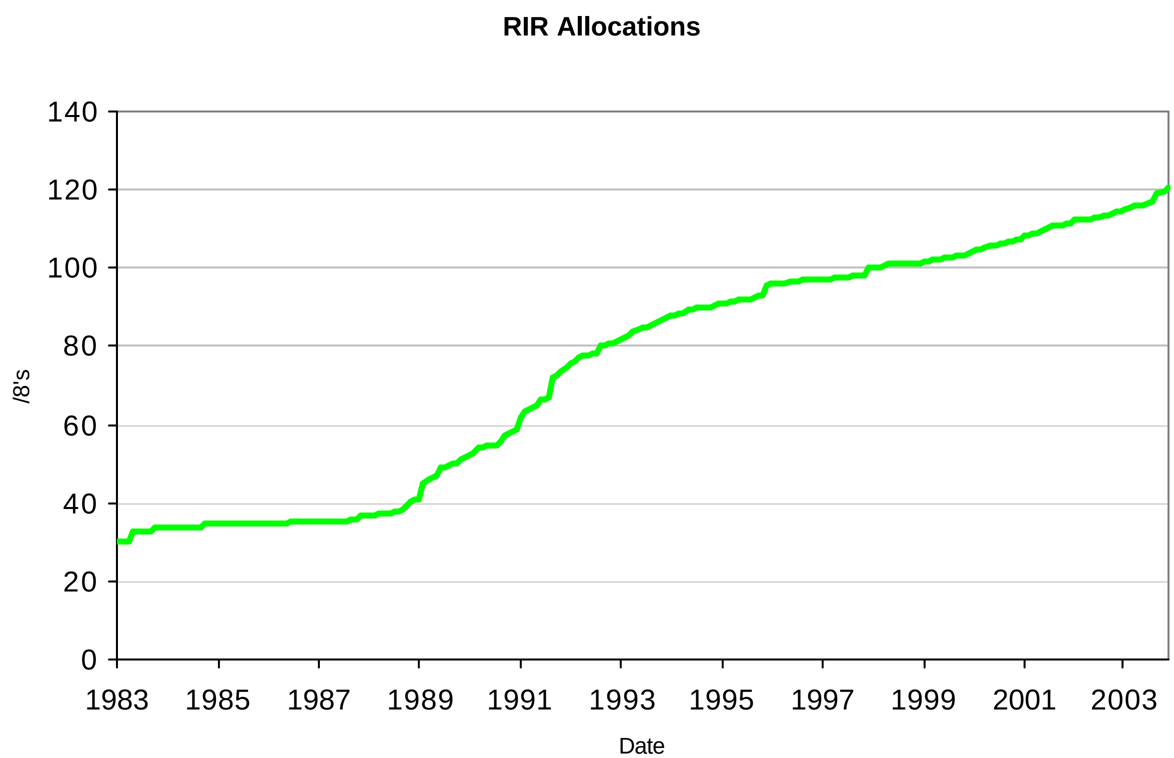

Figure 7 - Time Series of RIR Assignments

The post-1995 data used to extrapolate forward using the same linear regression technique described above, to find a curve of best bit using the same underlying exponential growth model assumptions:

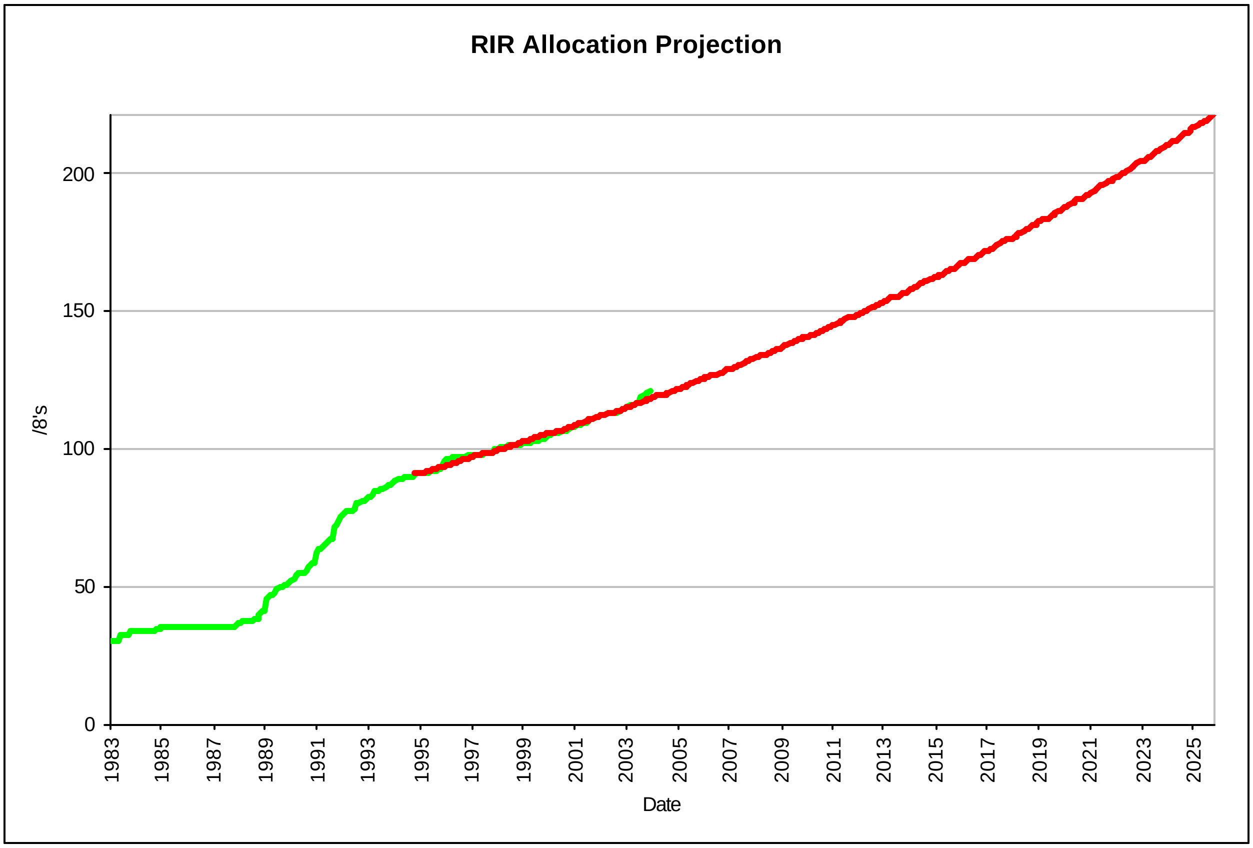

Figure 8 - Projection of RIR Assignments

This form of extrapolation gives a date of 2026 for the time at which the RIRs will exhaust the global unicast address 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.

It is useful to consider why the IANA consumption rate appears to be higher than the RIR consumption rate, and why there is a difference of some 7 years in the predicted exhaustion dates between the two data sets. There are a number of factors that contribute to this difference. The first is that the IANA pool is considered exhausted at the point in time when the last available /8 block is allocated to an RIR, while the RIR pool is considered to be exhausted at the time when the last available address block is allocated to an ISP. This implies that the other RIRs would have exhausted their pool at a slightly earlier point in time. Secondly the RIR practice of using assignment windows to allow subsequent allocations to the same ISP to be aggregated into a single routing advertisement implies that the RIR address pool is fragmented, leading to a slightly higher rate of draw down from the IANA unallocated address pool. Also, as the number of RIRs increases (such as with LACNIC in 2003 and the potential recognition of AFRINIC in the near future) the total amount of address space held in RIR pools increases. However, the most likely reason for the difference is to be found in the data sets themselves. The IANA data set consists of 129 data points with a date granularity no finer than one month, and of this data set, only the most recent 29 points (post 1995) are used in the IANA prediction. The RIR delegated files span some 30,000 entries, with a granularity of one day. The larger data set with a finer level of granularity tends to suggest that a slightly higher level of confidence can be ascribed to the predictive model of the RIR delegation data as distinct to the relatively sparse data set and coarse granularity of the IANA data.

The BGP Routing Table

Once 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's routing table. The next step is to establish the trends of the amount 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[4].

The data as of October 2003 shows that some 29% of the total IPv4 address space is announced in the BGP routing table, while 17% 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% of the address space is held by the RIRs pending assignment or allocation (or at least there is no RIR recorded assignment of the space), while 35% of the total space remains in the IANA unallocated pool. A further 8% of the space is held in reserve (Figure 9).

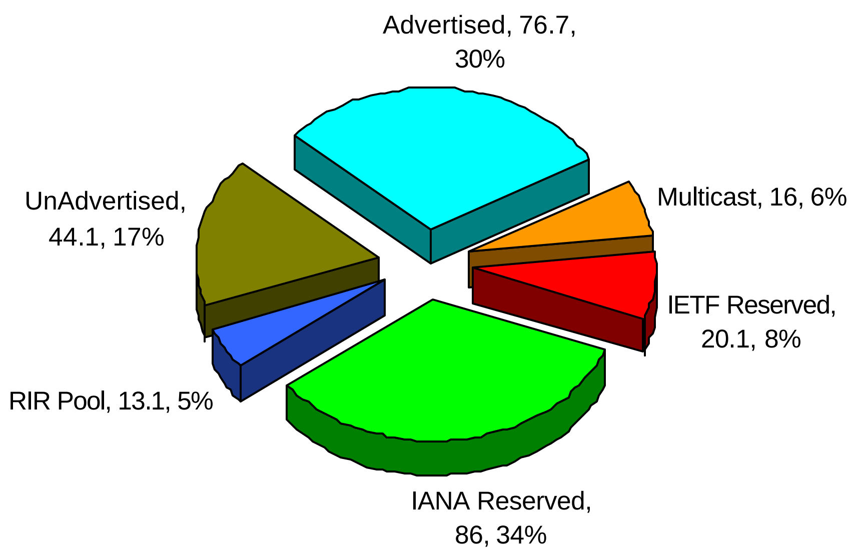

Figure 9 - IPv4 Announced Address Space

---

[4] It would normally be expected in a uniformly connected network that all addresses would be seen from all locations, and that the point at which the observations are being made as to the total span of addresses being advertised in the Internet would be immaterial. Analysis of routing views using “route collectors” suggests that this is not entirely the case, and that there are differences between various viewpoints. The scale of difference is not considered to be so large that it creates any significant uncertainty in the related predictive analysis based on the data.

This BGP data is based on an hourly inspection of the amount of address space advertised within the Internet's routing table. The data collection commenced in late 1999, and the data gathered so far is shown in Figure 10.

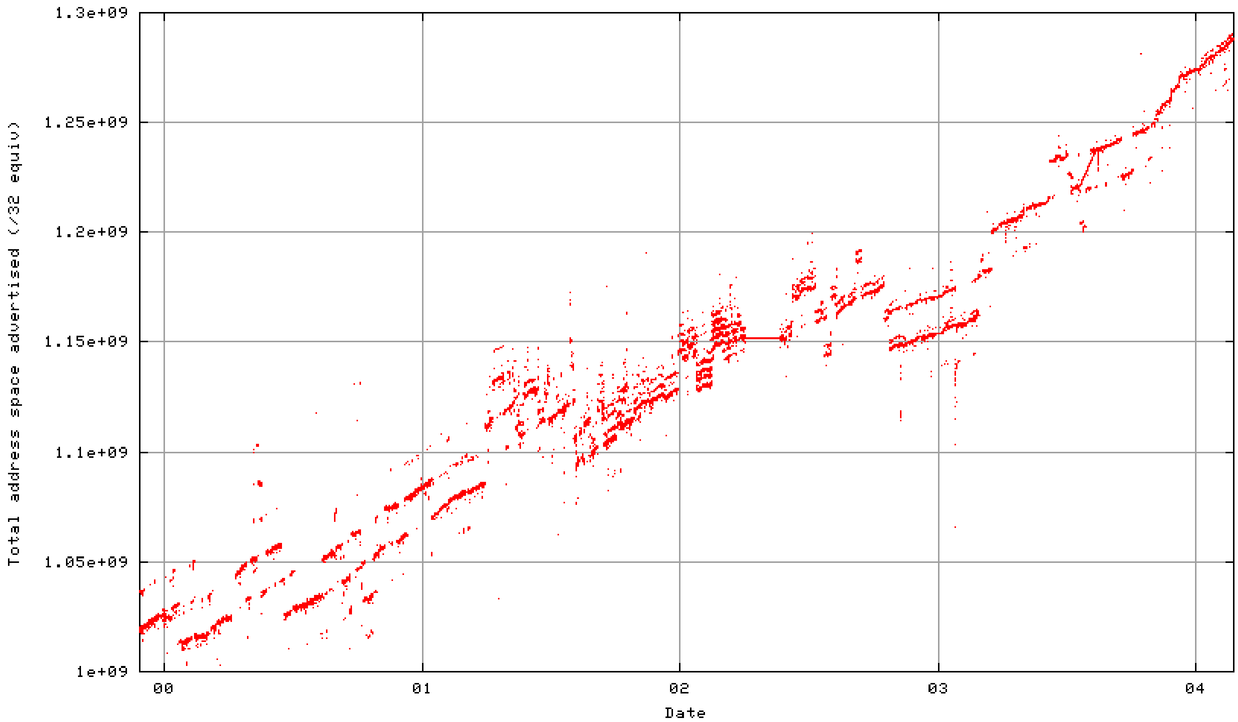

Figure 10 - Advertised Address Space as seen at AS1221

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 are periodically 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 approached 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 11.

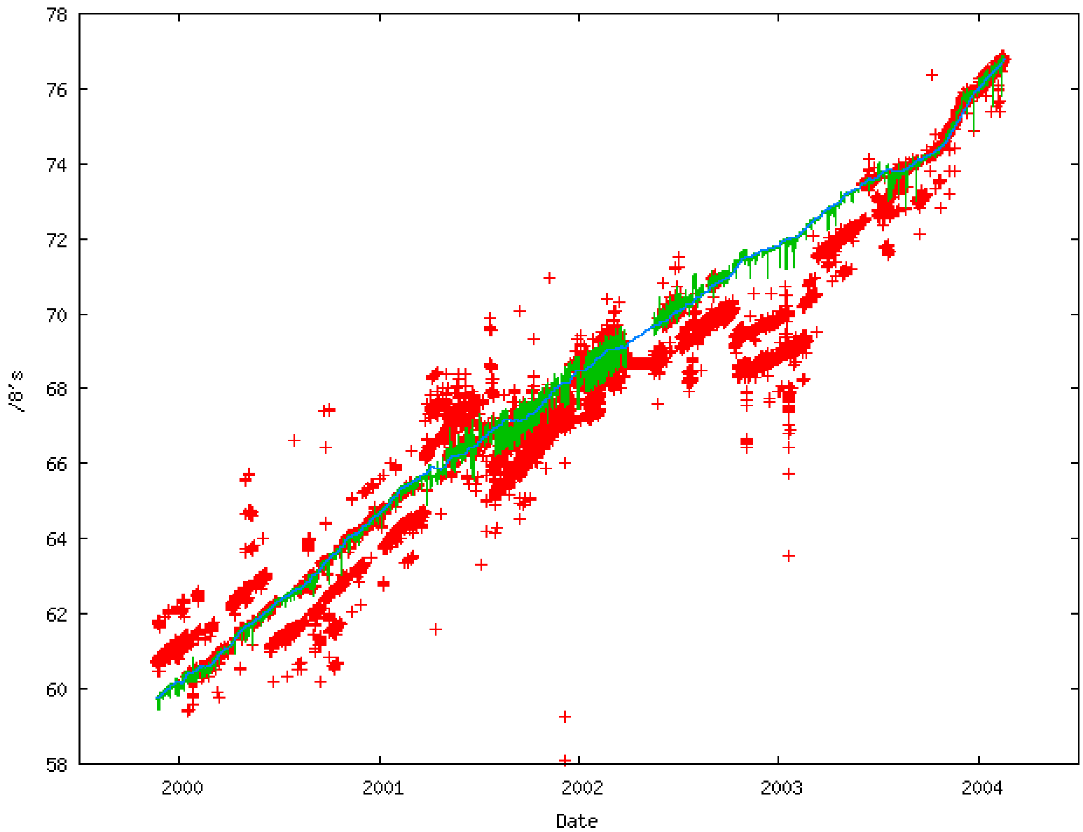

Figure 11 Smoothed BGP Data

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 analysing the first order differential of the data. An underlying increasing growth model would have an increasing first order differential, while a decreasing growth model would have a negatively inclined differential. Figure 12 shows the first order differential of the BGP data and the fit of constant growth and exponential growth models to the data. It appears that the most consistent fit to the data set is that of a constant growth model.

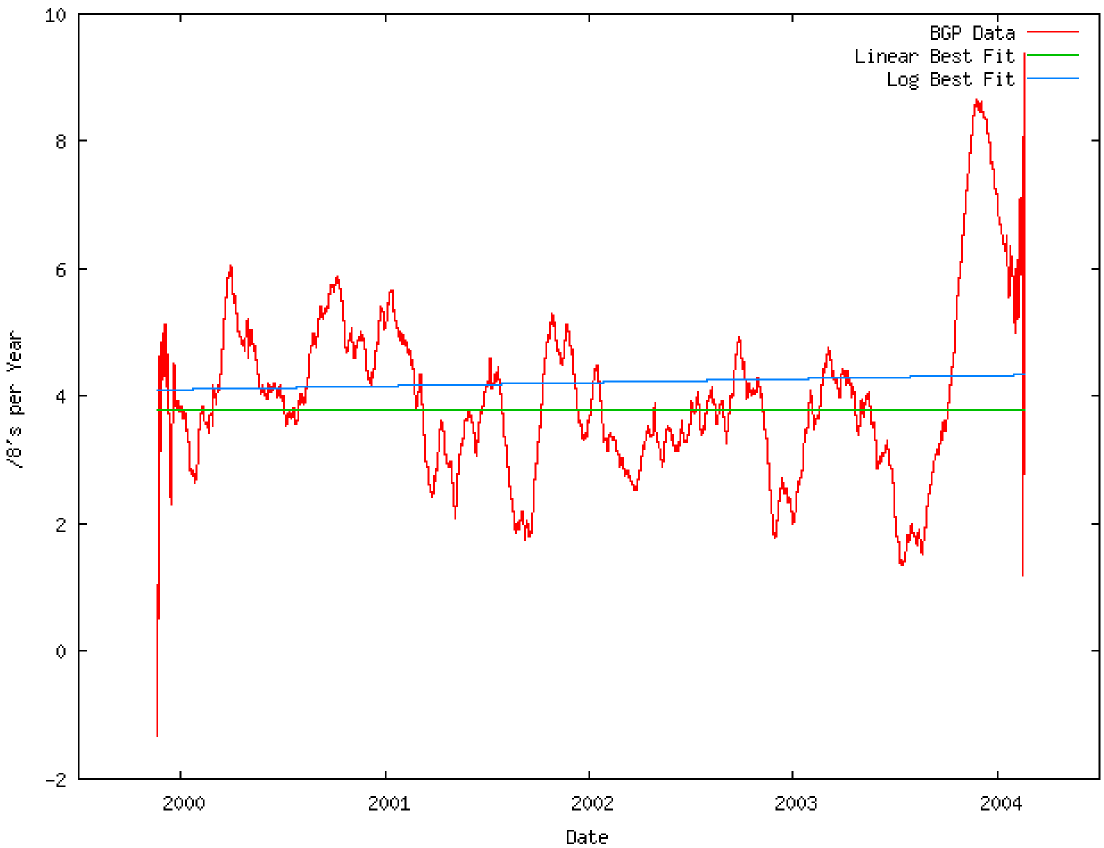

Figure 12. First order differential of BGP Data

It is now possible to produce two models, using constant growth and increasing growth, that extrapolate forward to complete exhaustion of the IPv4 address space. These models are shown in Figure 13.

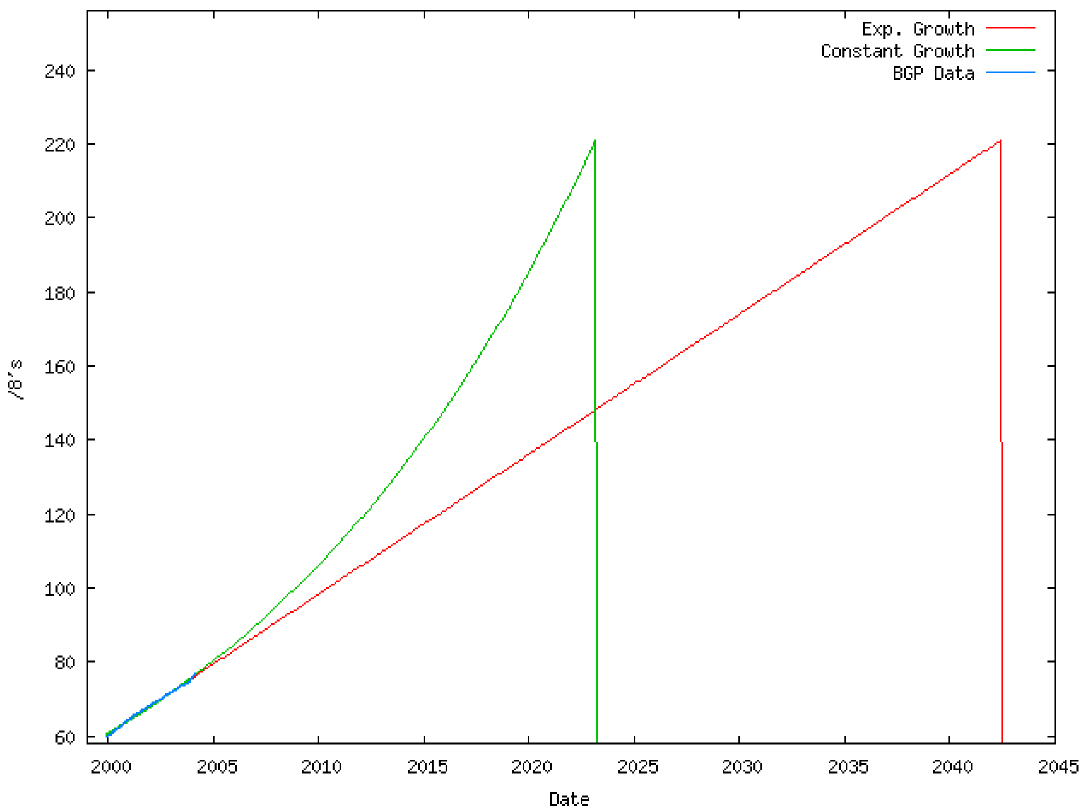

Figure 13 IP Address Consumption models based on Routing Advertisements

The constant growth model predicts complete exhaustion in mid 2042, while an exponential growth model has a prediction of complete exhaustion of 2023.

Combining the Three Views

Before completing the projections for IPv4 address space there is one remaining question. There are 43.3 /8 blocks, or some 17% of the total IPv4 address space which has been allocated for use, but is not visible in the Internet's routing table. This is a very significant amount of address space, and if it's growing at the same rate as the advertised space, then this will have a significant impact on any overall model of consumption of 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 can into operation in the mid 1990's, or some intrinsic inefficiency in the current system. If its the latter, then its likely that this pool of unannounced addresses will grow in direct proportion to the growth in the announced address space, while if its the former then the pool will remain relatively constant in size in the future.

The RIR allocation data and look at the allocation dates of unannounced address space (Figure 14). 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.

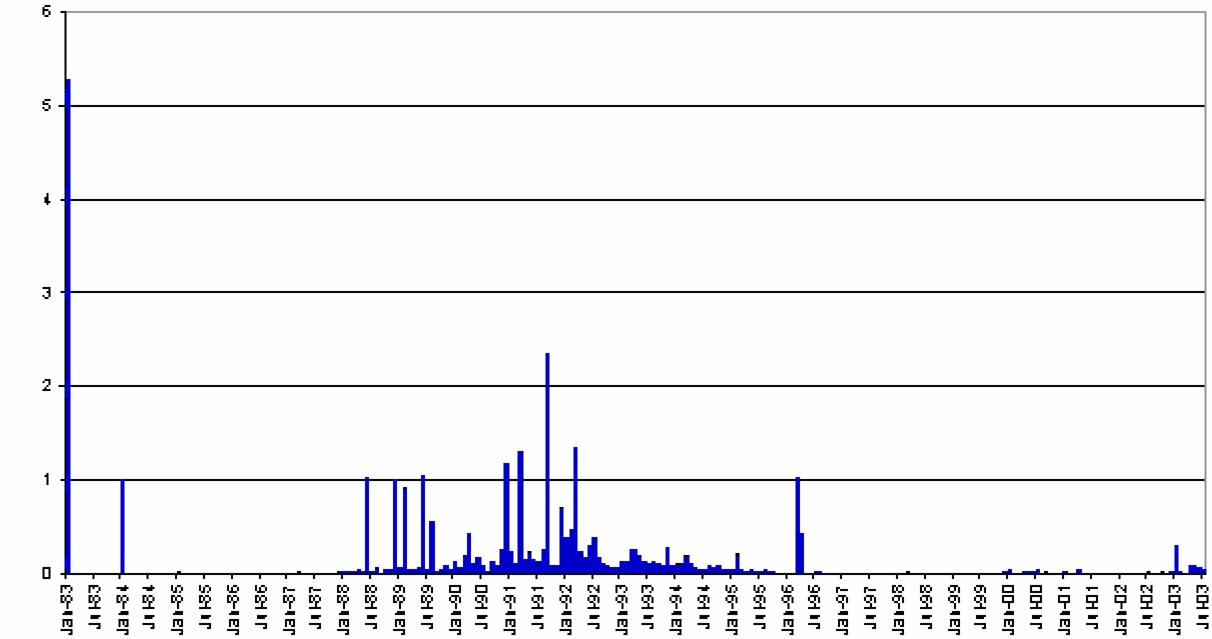

Fig 14 - Age Distribution of Unadvertised Address Space

This confirms that in recent years the overall majority of the address space that has been assigned by the RIRs appears in the Internet's routing table. This in turn implies that projections of the amount of address space advertised in the routing table provide a reasonable correlation to underlying mechanics 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 pool of allocated but unadvertised addresses it is possible to 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.

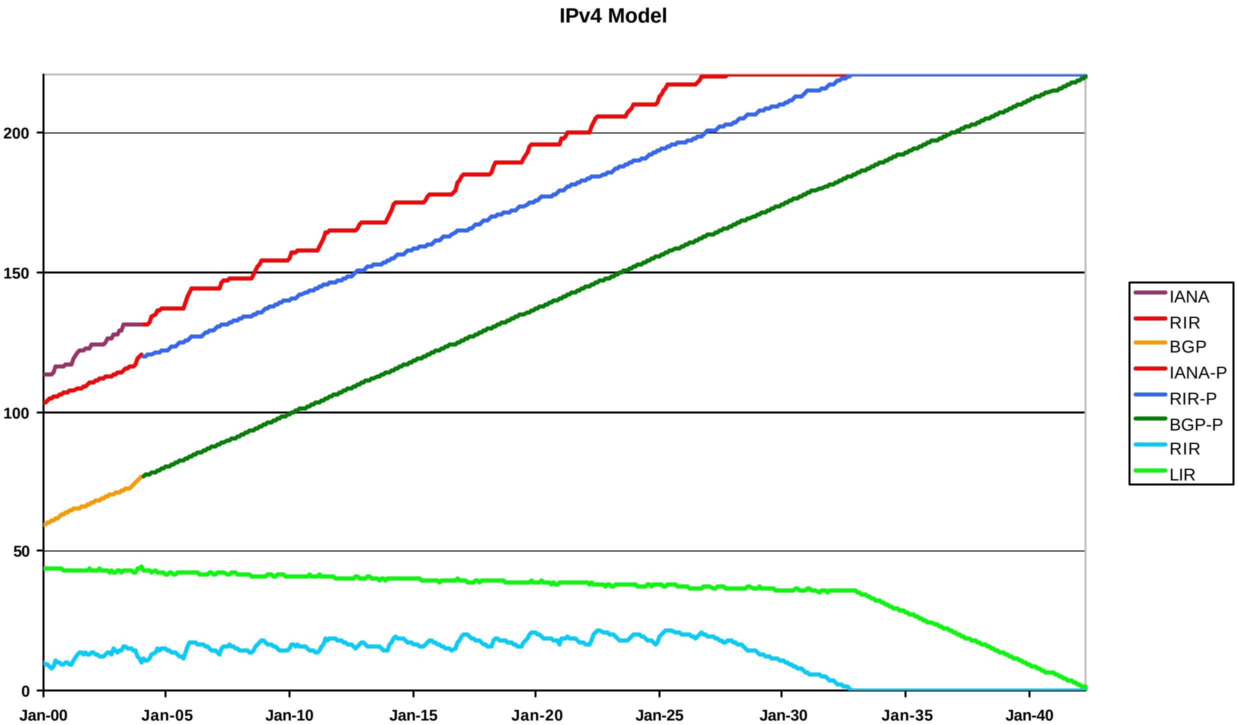

Figure 15 - Address Consumption Model

This model tracks the process of IANA allocation to the RIRs, the RIRs operating allocation windows to their LIRs and ISPs, and LIRs and ISPs advertising this address space to the routing system. The underlying assumption of the growth model is that of a continuation of a constant growth model, with a constant consumption rate of some 3.4 /8 blocks per year. Within this mode the IANA pool will be exhausted in mid 2027, and the RIR pool will cease to operate using allocation windows early 2028. While there would be sufficient addresses within the allocation system to sustain a further 5 years of RIRs drawing from their allocation pools, it is reasonable to predict that the current address allocation practices would be forced to change in early 2028. Accordingly, this linear growth model predicts an effective exhaustion date of current allocation IPv4 allocations in 2028,

An accelerating growth model of Internet address consumption, assuming that the trends over the past 3 years in address consumption rates continue to grow at the same accelerating rate would see effective exhaustion occur a decade earlier, in 2018.

Uncertainties and Assumptions

Of course such projections are based on the underlying assumption that the visible changes that have occurred in the past will smoothly translate to continued change the future. There are some obvious weaknesses in this assumption, 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 NATs to support some new popular class of 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 due to the very large volumes associated with deployment of this particular technology[5].

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 geo-political framework, where number blocks are passed to national entities, and further distribution is a matter of local policy[6]. 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 address that can be used to absorb the impacts of a disruptive change in address consumption rates. While, at present, some 60% 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 there remains sufficient address space 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.

The assumption is also made that market behaviours will remain relatively constant over the period under examination. This is a relatively weak assumption, in that as the unallocated pool of addresses diminishes in size it is reasonable to anticipate both an increase in market demand for alternative technologies that have a reduced (or no) requirement for IPv4 address space (such as IPv6), and also an emergence of secondary address trading markets where address space distribution is realigned to reflect a distribution that maximizes the utility value of the address space. In this predictive investigation the externalities associated with modified market behaviours in the face of a visibly diminishing address resource will be assumed to have no impact on the rate of consumption of address resources.

This particular model also assumes some form of continuity of current address allocation polices. This is not a likely scenario, as it is likely that address policies will reflect some notion of balance between

[5] On the other hand it is evident that the growth of the Internet in recent years has been fuelled by the increasing prevalence of NATs. 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 NATs is certainly evident today.

[6] Such a geo-political distribution system is used in the E.164 number space for telephony.

the level of current demand and 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, and it may be that at such a "cross over" 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 the current address allocation process does not operate as a pure market. 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. By removing price 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 a smooth market-led transition into IPv6.

IPv4 Consumption Prediction

Three factors have substantially altered the IPv4 Internet in intervening decade since the original IETF predictions of address space exhaustion. They are:

Dynamic Host Configuration Protocol (DHCP)

> The original IP architecture assumed that every device would be assigned a unique address on a persistent basis. DHCP allowed network administrators and service providers to maintain a pool of addresses and 'lend' them to devices for the period during which they were attached to the network. This has realized considerable savings in terms of address consumption. For example, in a dial-up service environment, the size of the required address pool is not the number of contracted customers of the service, but the maximal number of simultaneously connected customers, or, in other words, the size of the address pool.

Network Address Translation (NAT)

> Various host census estimates indicate that the number of devices that are in some fashion "attached" to the Internet is somewhere between 2 and 3.5 billion. The total amount of address space being advertised within the global routing tables encompasses some 1.3 billion addresses. This would imply that, if the advertised address space were fully utilized, then the number of devices behind NATs is either equal to, or up to double the number of directly addresses hosts. NATs are very common in today's Internet. The reasons for their popularity are as much, if not more, to do with the level of protection they are perceived to offer against various forms of malicious attack as they have to do with their capability to undertake address compression. There are three commonly cited drawbacks to NAT deployment, those of its inability to provide protection of the packet header in the context of the IPSEC end-to-end security protocol, its inability to allow persistent services to be located within the "private

“ realm of the NAT and the inability to support application protocols that carry IP addresses within their payload. None of these reasons has observed to be an impediment to the widespread deployment of NATs.

Address Allocation Policies

The third factor is the policies associated with the current address allocation function undertaken by the RIRs. The common use of an 80% utilization threshold as the necessary precondition for a Service Provider to obtain additional address resources, after the initial allocation, and careful management of the initial allocation requests to ensure that the initial allocation is capable of supporting the documented rollout of the network have combined to ensure that the policy objectives of the RIR system, those of conservation and responsible use of the address resource, appear to be achieved.

The assumptions used here include a continuation of the current utilization efficiency levels in the Internet, and a continuing balance between public address utilization and the use various forms of address compression, continuity of current address allocation policies, as well as the absence of highly disruptive events.

With all this in mind, it would appear that the IPv4 Internet, in terms of address availability, could continue up to around 2030 without reaching any fixed boundary of address exhaustion.

But it must be remembered that each of the assumptions made here are relatively sweeping, and to combine them as is done here is pushing the predictive exercise to its limits. Three decades out is well over the event horizon for any form of useful prediction for the Internet. If the question is restricted to the next decade (i.e. up to 2014), then we can answer with some level of confidence that there is really no visible evidence of IPv4 exhausting its address pool based on the available address consumption data.

The conclusion drawn here is that if the prime driver for the adoption of IPv6 lies in a looming shortage of public IPv4 addresses, then this is not a near term prospect, and any impetus for the adoption for IPv6 would need to be based on factors other than an imminent exhaustion of IPv4 address space.

Data Sources

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

Registry "delegated" report files:

APNIC: ftp://ftp.apnic.net/pub/apnic/delegated

ARIN: ftp://ftp.arin.net/pub/delegated

LACNIC: ftp://ftp.lacnic.net/pub/delegated

RIPE NCC: ftp://ftp.ripe.net/ripe/delegated

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