

May 2009

Geoff Huston

Predicting the End of the World

For some years now I've been running a set of scripts that attempt to model the consumption of IPv4 addresses and then use this model to look forward in time and predict the date at which the pool of unallocated IPv4 addresses will be exhausted. The results of this model, together with a description of the process used in the model can be found at <http://ipv4.potaroo.net>. As the predicted date of exhaustion is getting to be a "right here and right now problem" rather than a "comfortably in the vague future, so its really someone else's problem" its probably worth revisiting the model and the assumptions it makes so that you can choose for yourself whether to believe its predictions - or continue to press on with IPv4 and blithely ignore the entire problem for yet another day.

The complete output of this automated model is shown in Figure 1. This figure shows the pool of IANA's addresses being depleted, the total size of the allocated address pool, and the amount of address space that is visible as advertised address space in the routing system and the amount of unadvertised address space. The figure also shows the amount of unallocated address space held by all the RIRs, and the totals for each of the RIRs. The model generates this data for previous years, using data published by the RIRs and historical BGP routing data, and then uses a combination of statistical techniques and simulation to generate a likely model of the future. In this column I would like to explore in detail the inner working of this model in some detail.




Figure 1 – IPv4 Consumption model

I've been running this model using a daily run of the script since early 2003, and over time the predicted date of exhaustion of the IPv4 address pool has varied greatly. In early 2003 the Internet was still recovering from the boom and bust that had immediately preceded it. The average address consumption rate was some 64 million addresses per year (slightly less than 4 /8s) and if the pace of expansion of the network had proceeded at that speed into the future, then at that rate IPv4 address exhaustion was predicted to occur at some time in 2021 (<http://www.potaroo.net/ispcol/2003-08/ale.html>).

Things change, and the address consumption model has to pick up on changes in address consumption patterns and reflect them in the prediction. The average address consumption rate has tripled in the past 6 years, and we are now averaging some 12 /8s per year in IPv4 address consumption. At this rate the last 25 /8s will be consumed in the next two years, and the first part of address exhaustion, namely the exhaustion of the IANA unallocated address pool, will probably occur in 2011.

Figure 2 shows a plot of the predicted exhaustion date since late 2007. The horizontal axis is the date the prediction was made, and the vertical axis is the date of predicted exhaustion. The red line is the predicted IANA exhaustion date, while the green line charts the data when the first RIR will exhaust all its unallocated addresses.



Figure 2 – Address Exhaustion Prediction Dates

Another way to look at this data is to look at the time remaining to exhaustion from the time the prediction was made, and look at the variation in this period over time. If the model were predicting the same date for exhaustion all the time, then each day the time left until exhaustion would reduce by the same single day. This "constant" prediction can be compared to the actual variation in the predicted dates in the following figure (Figure 3). In May 2009 we now have slightly over 2 years to go before IANA exhausts its address pool and around 2.8 years to go before the first RIR exhausts its pool of addresses.



Figure 3 – Time Remaining until Exhaustion

If we have about two years to go, surely the model should've stabilized, and the predictions from the model should've converged by now, and we should see a pretty stable prediction of the point of exhaustion. So why does this predicted data of exhaustion continue to jump around?

While we are asking critical questions, why has there been these discontinuities in the predictions over the past 18 months? What's been changing to make the model generate those abrupt changes to the predicted end dates? What happened between the middle of 2008 up until the present time to make the model shift such that the data of IANA exhaustion was a constant 2 years out? Why has the predicted date of RIR exhaustion changed dramatically a month ago? It is reasonable to expect that as the date of address exhaustion gets closer that the model would start to converge to a stable point and the day-to-day variations would lessen in size. After all, if we want to place some degree of credibility in this prediction it may be useful to understand what is behind the prediction in terms of the workings of the model itself. So in this month's column I'd like to try and explain the workings of this address consumption model, the changes that have been made to this model in the past 18 months and the impact of these changes on the predicted exhaustion point.

Modelling Address Consumption

At the heart of the model are two simple mathematical functions.

The first is a smoothing function which generates a new series where each new value for a given sample point is the average of the old value and a number of values preceding and following the value. For example if a data series contains the sequence ... 3, 5, 4, 11, 4, 5, 4, ... then with a smoothing function using window of 7 sample points the value of 11 would be replaced by the smoothed value 5.14.

The second is a linear least squares line fitting function (http://en.wikipedia.org/wiki/Least_squares). The linear least squares fitting technique is the simplest and most commonly applied form of linear regression and provides a solution to the problem of finding the best fitting straight line through a set of points. Given a set of points (x,y) a linear least squares operation produces a line of the form $y = mx + b$ where the sum of the squares of the distance of the line from each point is minimized. The use of squared offsets is due to the consideration that the use of absolute values requests in discontinuous derivatives that cannot be treated analytically to find a minimum, while the sum of the squares does result in a continuous derivative that can be readily minimized. However the side effect is that this procedure results in outlying points being given disproportionately large weighting. For this reason the smoothing function is used to reduce the outliers in the raw data before applying this least squares procedure.

The address consumption model is based on the allocation data published by the RIRs. The allocation report that each RIR publishes is updated every day with the day's changes.

These daily files are published at:

Afrinic: <ftp://ftp.afrinic.net/pub/stats/afrinic/delegated-afrinic-latest>
APNIC: <ftp://ftp.apnic.net/pub/stats/apnic/delegated-apnic-latest>
ARIN: <ftp://ftp.arin.net/pub/stats/arin/delegated-arin-latest>
RIPE: <ftp://ftp.ripe.net/pub/stats/ripencc/delegated-ripencc-latest>
LACNIC: <ftp://ftp.lacnic.net/pub/stats/lacnic/delegated-lacnic-latest>

For example, on the 21st May 2009 the following allocation actions have been recorded in the RIRs' records:

```
arin|US|ipv4|24.54.96.0|8192|20090521|allocated
arin|US|ipv4|67.231.160.0|4096|20090521|allocated
arin|US|ipv4|68.235.80.0|4096|20090521|allocated
arin|US|ipv4|69.169.64.0|8192|20090521|allocated
arin|US|ipv4|74.114.16.0|2048|20090521|allocated
arin|US|ipv4|74.114.24.0|2048|20090521|assigned

apnic|JP|ipv4|110.233.0.0|65536|20090521|allocated
apnic|JP|ipv4|111.168.0.0|131072|20090521|allocated
apnic|CN|ipv4|202.43.76.0|1024|20090521|allocated
apnic|IN|ipv4|203.55.160.0|256|20090521|assigned

lacnic|AR|ipv4|200.105.120.0|2048|20090521|allocated
```

The stats file format uses 7 columns, separated by a '|' character. The columns are:

- registry name
- country where the recipient of the resources resides
- resource type
- start point of the allocation
- size of the allocation
- date of the registry allocation
- allocation or assignment

The records indicate that on this day ARIN have performed 6 allocations, with a total of 28,672 addresses, APNIC performed 3 allocations, with a total of 197,888 addresses and LACNIC performed a single allocation of 2048 addresses. The total for the 21st May 2009 was an allocation of a total of 228,608 addresses.

From this data a daily total of allocated addresses is generated for each day, as shown in Figure 4.




Figure 4 - Time Series of IPv4 Address Allocations

A smoothing function is used across the resultant time series, using a smoothing window of 71 days. The smoothing function is repeatedly applied to the data using 4 passes. This smoothed data is shown in Figure 5.




Figure 5 - Smoothed Time Series of IPv4 Address Allocations

The first order differential is generated across this smoothed series, giving the average daily number of RIR-allocated addresses per day. The most recent 1095 days (3 years) in the series is used to generate a least squares best fit linear model against this first order differential.




Figure 6 - First Order Differential and Linear Fit

The function of the linear best fit to the first order differential ($ax + b$) is integrated to a quadratic ($(ax^2)/2 + bx + c$) and a value for c is calculated that minimizes the sum of the error of this quadratic to the original data series.

While there is strong level of variation in allocation rates through the year, the quadratic function is entirely uniform. In order to create a better approximation of the variations in the data, for each day of the year over the past three years the average deviation of the original data to the model is calculated, and this is applied as a correction to the quadratic function in order to include seasonal variations into the model.




Figure 7 - Daily Correction Factor

The model is then projected forward, applying the average variation to the model according to the day of the year of the projection, carrying this forward until the IPv4 address space is fully allocated.




Figure 8 - Projection of Allocated Address Pool

There are a number of additional aspects to the model, as it attempts to model the actions of the IANA in allocating address blocks to the RIRs. The way this is done is to calculate the relative proportion of total allocations made by each RIR for each day, then using a least squares best fit to this data to generate a linear trend of this relative proportion. This is used to then calculate the relative allocations performed by each RIR in the overall projected allocation.




Figure 9 - Calculation of relative allocation rate of each RIR

This data can be used in conjunction with the allocation data to model the consumption of the unallocated address pool held by each RIR. When the free pool is less than a threshold level IANA then allocates additional address blocks to the RIR, depleting the IANA pool while replenishing the RIR's address pool. While it is possible to simulate the individual allocations performed by each RIR, this is not performed in this model.

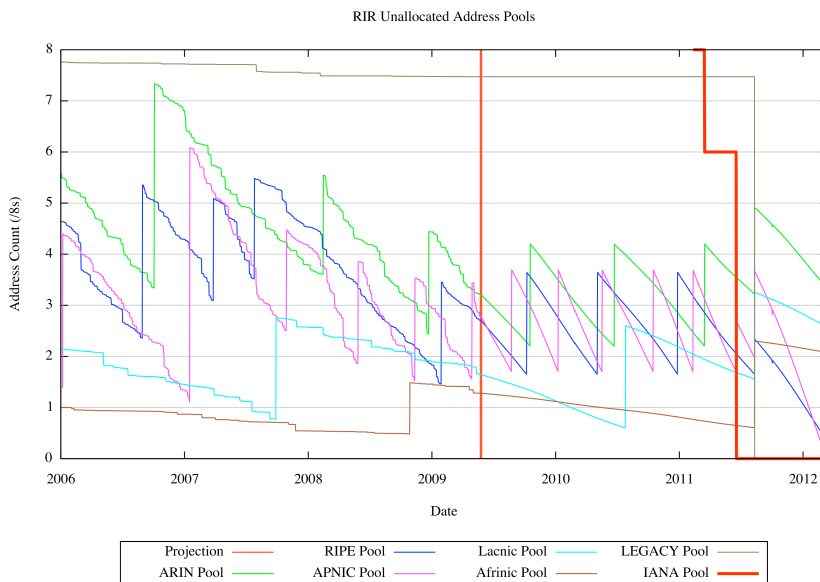


Figure 10 - RIR Unallocated Address Pools

So from this process the model calculates the date at which IANA will allocate its last remaining unallocated /8 address block, and the data at which the first RIR will exhaust its unallocated address pool.

However, the question I would like to address is why does the projected date of exhaustion of the IPv4 address pool keep changing?

Changes to Address Management Policies and Practices

The first answer to this question is that we are changing the policies and practices associated with the management of the IPv4 address pool.

The last 5 /8s

In 2008 a policy proposal was considered by each of the RIRs calling for IANA to "reserve" the last 5 /8 unallocated address blocks. When IANA is set to allocate the 6th last /8 address block it will also then allocate the remaining 5 /8 address blocks to each of the RIRs, thereby exhausting its remaining available address space. Each of the RIRs is to determine how it will manage this "last /8" and it appears that the allocation regime to be used for this last address block will differ from the current practice and policies. As the aim of this prediction exercise is to measure the point at which the current policy regime that relates to the distribution of IPv4 address comes to a halt through exhaustion of the address pool, the model assumes that these last 5 /8s will not be managed by the RIRs in the same way as current unallocated space is managed. These last 5 /8s are not included in the model of each RIR's normal allocation activity. The effective result of adoption of this policy is to bring the effective date of exhaustion closer by some 4 months or so. I made this change to the model in October 2008, hence the jump "inward" of the prediction by almost 6 months at that time (Figure 3).

2 /8s at a time

The RIRs also recently changed their practices with respect to requesting /8 address blocks from the IANA. The prevailing policy allows each RIR to request additional address space when its pool of available space falls below either one half of a /8 or is less than 9 months of supply at the current address allocation rate. The amount of space that will be allocated by IANA will be a number of /8s which will be sufficient to satisfy the requirements of the RIR to meet its needs for the next 18 months. For example, if an RIR is allocating address at a rate of 4 /8s per year, then once the RIR's unallocated pool falls below 3 /8s then the RIR can request a further 4 /8s from IANA, bringing its total pool of unallocated address space to more than 6 /8s, which will meet its projected 18 month needs. The RIR's have agreed between themselves to abide by further self-imposed constraints such that they will not request any address space from IANA until their pool of available addresses is approximately 2 /8s and the maximum request in any single transaction with IANA will be 2 /8s. This was a

change I made along with a number of other changes in May 2008. The cumulative effect of these changes was to draw the IANA exhaustion date inward by 6 months and the first RIR's exhaustion date inward by almost 12 months (Figure 3).

As of May 2009 there are 30 /8s left in the IANA pool. With the last 5 /8s being held over, then the remaining 25 /8s will be distributed between the RIRs according to the practices described above. The outcome of this process is a current projection that of the pool of 25 /8s remaining for 'normal' allocations, APNIC will receive a further 11 /8s, ARIN and the RIPE NCC a further 6 /8s each and LACNIC will receive a further 2 /8s.

Legacy Address Space Management

The second answer to this question concerns some refinements to the model over the treatment of the so-called "legacy" address space. In the IANA IPv4 address registry there are some 49 /8 address blocks where IANA has registered the address space as "LEGACY". These address blocks correspond to the old Class B address space and the old Class C address space, and allocations from these blocks pre-date the RIR system. While the legacy address space is, on the whole, heavily fragmented, it is still useable address space, and can be used to meet current demand for IP addresses. Of interest to the address consumption model is the observation that these address blocks are not fully allocated or assigned. It appears that there are the equivalent of 9.31 /8s left in this pool of addresses that are, at present, unallocated.

When I originally modelled the RIR's use of this space I used a model similar to that of the IANA address management function. Following IANA exhaustion, when an RIR was close to exhausting all of its unallocated address pool it would withdraw the equivalent of a /8 from the unallocated space in the legacy address pool. In this way an RIR with a higher allocation rate would consume a higher number of addresses from this legacy address pool than an RIR with a lower address allocation rate.

Distribution of unallocated legacy address space

This model does not correspond to my understanding of the intentions in managing this pool of addresses, and the model now assumes that the RIRs will distribute the unallocated "holes" in the legacy address space equally between all five RIRs. The implication is that each RIR has a fixed amount of additional unallocated address space over that allocated by IANA, and the RIR's allocation rate will not impact this amount of 'legacy' space that is available to the RIR. When this change was applied to the address consumption model it brought the RIR exhaustion time in by approximately 6 months. I made this change to the model in May 2009, and the net result was to draw the RIR exhaustion date inward by some 6 months (Figure 3).

188/8

The model assumes that each RIR will assign from the IANA-allocated address space and fully assign from that space before moving on to assign from the legacy address space pools. However, in February 2009 the RIPE NCC has commenced assignments from the legacy address block 188/8. The model has been further adjusted to incorporate 188/8 into the RIPE NCC's current address pool, and 188/8 will be treated in the same way as all other current RIPE NCC address blocks in the calculation of the threshold of available space for IANA allocations. However in terms of the allocation of unallocated legacy address space the RIPE NCC is assumed to have already drawn its allocation of space down by 1 /8. I also made this change in May 2009, one week after the preceding change.

At this point the address consumption model assumes that no other legacy address blocks will be used for allocations prior to the exhaustion of the IANA address pool by the RIPE NCC, or by any of the other RIRs.

Distribution of Address Assignments

There are two major assumptions behind this form of modelling.

Firstly, it assumes that tomorrow will be much the same as today. This assumption includes the consideration that the influences that shape today, or in other words the changes between today and yesterday will continue tomorrow, and that the second order change, or the difference between the change in the day before yesterday and yesterday and the change between yesterday and today will also be reflected in tomorrow's reading. This

model is relatively blind to discontinuities, disruptions, or other forms of externalities that may occur in the future, and at the same time can be overly sensitive to isolated events of a disruptive nature that have occurred in the past, as the impact of such events will be factored into the predictive model as a continuing impact. There are other forms of modelling, such as demand modelling, discrete event simulations and such, that have different properties in terms of their strengths and weaknesses as a predictive technique. This assumption tends to produce predictions that are relatively stable.

The Law of Large Numbers and Heavy Tail Skew

Secondly, it assumes that the population is sufficiently large that the actions of individuals do not have any appreciable impact on the model. This is perhaps a paraphrasing of the “law of large numbers” (http://en.wikipedia.org/wiki/Law_of_large_numbers) which predicts stable long term results for random events. Or, to quote Wikipedia: “Given a random variable with a finite expected value, if its values are repeatedly sampled, as the number of these observations increases, the sample mean will tend to approach and stay close to the expected value (the average for the population).” So the assumption we are making here is that even though the actions of an individual entity in requesting an individual allocation from an RIR is, to some extent random, or unpredictable, the actions of many such individual entities over time has a stable mean. How valid is this assumption? How many allocation actions are we using in this model?

In the 1095 day period there have been 19,577 recorded allocations made by the RIRs, and these have accounted for the equivalent of 35.257 /8s, or a little under an average of 12 /8s per year. However, only 221 of these allocations, or a little over 1% of the total number of allocations, account for one half of the total allocated address space. Some 90% of the address space, or some 32 /8s were allocated in some 10% of the total number of allocation transactions, or around 2,000 transactions. The data does not provide a clear indication of the number of distinct entities involved here, but given that the prevailing address allocation policy is to allocate for the needs over the forthcoming 12 months, the number of distinct entities involved here may be considerably lower than 2,000. A view of the distribution of address allocations is shown in the following figure.




Figure 11 - Distribution of Allocations since May 2006

I'm not sure if it is a classic example of a heavy tail distribution (http://en.wikipedia.org/wiki/Heavy-tailed_distribution) or not, however it is certainly highly skewed. If this distribution of address continues in the coming months then of the remaining 25 /8s, some 12 /8s will be allocated in 140 individual allocations in the coming 24 months. Precisely when these 140 allocations is not as easy to model given that the model would be attempting to predict the actions of some 140 individuals over 24 months, or the equivalent of some 6 allocations per month. If this collection entities determined to accelerate their applications for further addresses then the endpoint of the address pool could be brought forward by more than one year, and, similarly, if they delayed their application for additional allocations of address space the exhaustion point could be deferred for a considerable period.

The conclusion here is that due to the heavily skewed nature of the distribution of allocations the prediction technique used here has a very high degree of uncertainty associated with it, and this uncertainty is not reduced as the pool of remaining addresses declines further in size.

Consistency and Quality of the Data

The model used here deliberately uses published RIR data collected from the five RIRs, as well as data collected from the IANA and data gathered from dumps of the inter-domain routing table. How good is this data? How consistent is this data?

There is one notable inconsistency in the statistics files that are published by the RIRs. The general interpretation of these daily statistics files is that each RIR allocation action is recorded in the file in its own entry, describing the allocated resource, the date when the allocation was made, the size of the allocation and the country of record of the entity to whom the allocation has been made.

For example the following entry from the RIPE NCC's daily stats file

```
ripenncc|NL|ipv4|213.247.64.0|16384|20090317|allocated
```

describes an allocation of the address block 213.247.64.0/18 to an entity whose country is the Netherlands, or "NL".

When the same entity receives two or more allocations of addresses then there are multiple entries in the stats file to record these actions. However, ARIN uses a slightly different approach in the case where an allocation is made that 'expands' a previous allocation.

For example ARIN has the following entry in its stats file report:

```
arin|US|ipv4|24.49.96.0|8192|20081104|allocated
```

which reports that on the 4th November 2008 ARIN allocated 24.49.96.0/19 to an entity in the US. However prior to the 8th April 2009 the entry used to contain:

```
arin|US|ipv4|24.49.96.0|4096|20081104|allocated
```

The interpretation here is that ARIN appears to have originally allocated 24.49.96.0/20 on the 4th of November 2008 and on the 8th April 2009 ARIN has allocated 204.49.112.0/20 to the same entity, and revised the details of the original allocation to encompass the larger address range. The other RIRs use a different practice and in this same situation the other RIRs would report two allocations, and would have the following entries in their stats files:

```
arin|US|ipv4|24.49.96.0|4096|20081104|allocated  
arin|US|ipv4|24.49.112.0|4096|20090408|allocated
```

In order to ensure that all the data sets reflect that same underlying reporting procedures I've had to reprocess the ARIN data report and replace any instance of these amalgamated entries with the individual allocation actions, as determined by using daily differences in the ARIN report file. Fortunately each RIR not only publishes the current daily stats file, but also keeps online all the previous daily status files, allowing this day-by-day comparison to be made. This change was made in May 2008, and, together with some other changes made to the model at the time, had a significant impact on the predictive model. The ARIN recording practice has the side effect of pushing some recent large allocations back into the past, reducing ARIN's apparent current allocation rates, and thereby reducing the overall growth drivers in the model. Bringing these allocation back to their actual date corrects this anomaly.

Given five separate RIRs each performing the same reporting function there is the potential for two or more RIRs to publish mutually inconsistent entries, particular if it relates to some form of movement of the registration of a resource allocation from one registry to another.

Another issue is a relatively minor occurrence, and it relates to the recording of the date of allocation. When an entity requests the movement of an allocation record from one RIR's registry to another under the terms of the early registration transfer procedures the original data of resource allocation is not always maintained in the transfer of the registration record in the stats file. Instead the accepting RIR records the data of the transfer itself.

Another potential source of inconsistency relates to the completeness of reporting of the legacy allocations. While some RIRs have endeavoured to reproduce all the early IANA allocations from the legacy Class B and C

space in the daily stats file, other RIRs have not adopted this approach and the total pool of allocated address space managed by the RIR is not fully described in the RIR's stats file.

In order to assess the extent of these inconsistencies in the reported data, I have compared the data in the daily stats file against the RIR's database reports that the RIRs make available under a research use agreement. A report of the inconsistencies in the RIR's data can be found at <http://bgp.potaroo.net/bogons/rir-data.html>.

The reason why this is important in the context of the operation of this model is in the calculation of the total size of the unallocated space in the legacy address blocks. There are some 23.9 million addresses (23,987,465 to be exact) which are marked in the databases as being allocated or assigned, but are not visible in the daily stats files.

Business Cycles

So far the reasons we've examined for changes in the predictive outcomes of this model of Ipv4 address consumption have been somewhat introspective, as they focus on the model itself and the data it uses. But the reasons for change are not entirely internal.

The last three years has seen the **Dow Jones stock index** hit its high in October 2007 and then head to almost half that level by early 2009




Figure 12 - Dow Jones Industrial Average Index - May 2006 to May 2009
(From moneycentral.msn.com)

Obviously this has an impact on the pace of product deployment in the Internet which, in turn, has impacted address consumption rates.

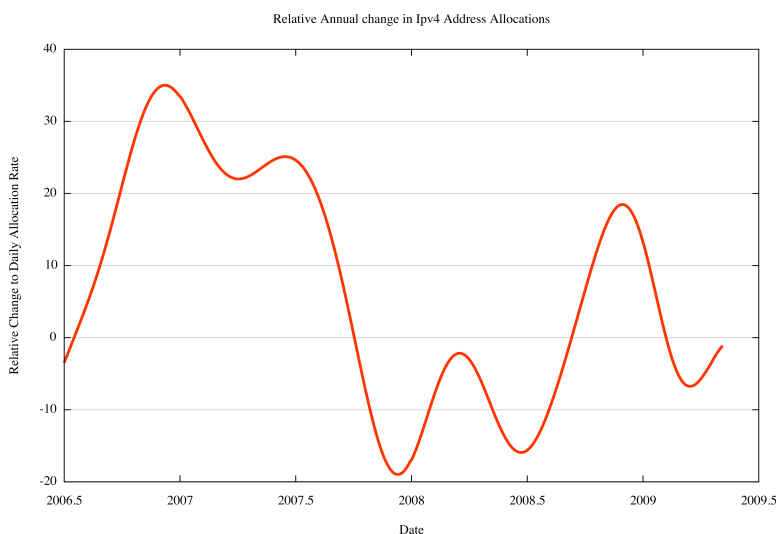


Figure 13 - Relative Annual Change in Daily IPv4 Allocation Rates

The correlation is not precise against the stock market indices but the allocation data shows a strong downswing starting in July 2007 and continuing across 2008, with some upward movement in more recent months.

Another perspective on the same events can be seen in Figure 3, looking at the number of days until exhaustion of the address pool. However, given that this data uses a 3 year sliding window, it can take between 6 to 9 months for a new trend in the data to have a visible impact on the prediction. The impact of the slowdown in the predicted end date on the prediction of address exhaustion was visible in this prediction by mid 2008, and for a 1 year period the rate of consumption of addresses has been slowing to an extent that the prediction of IANA exhaustion has been at a constant 2 years in the future. It is only in the past month that the projected exhaustion date has recommenced drawing inward, which appears to correlate to the slow signs of a recovery in the stock indices in most developed markets around the world.

More Refinements to the Model?

It is certainly possible to tweak this model in various ways. The RIR allocation algorithm could be refined to take into account weekdays and weekends, and also simulate the range of individual allocation sizes for each RIR. I could also use [Fourier analysis](#) to extract out some of the strong periodic elements in the data series and use these functions to form the predictive model. Alternatively, the least squares best fit algorithm could use a weighting on the data points to give a higher weighting to more recent data over older data. The RIR data could be assessed in further detail to resolve some of the lingering inconsistencies.

But, ultimately the real question here is what would all this refinement to the address consumption model achieve over and above what we already know about this situation in terms of its predictive capability?

Its clear to say that the exhaustion of the unallocated IPv4 address pools within the current distribution regime is as close to a certainty as anything is in this world. Sometime soon, or maybe a little sooner, or maybe just a little later, the IPv4 unallocated address pools will run out for each RIR, and at that stage changes will necessarily come into play. Within two or maybe three years from today the current lines of supply of IPv4 addresses will probably dry up. Given the skewed nature of the distribution of allocations it is difficult to be any more precise than this and although the mathematical model may claim today that exhaustion will occur at 10:32 am on the 14th of June 2011, the range of uncertainty in such a prediction spans years rather than seconds. No matter when it does occur, after that time the life of an IPv4-only network will start to get quite messy and certainly very expensive, and it will only get worse from then over time.

If this looming exhaustion of the current supply of IPv4 addresses isn't a sufficiently loud and clear message to each and every player in this industry to **DEPLOY IPV6 PRODUCTS AND SERVICES NOW!** then I'm very much afraid that no other message is going to work either.

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>