

August 2013
Geoff Huston

When?

At the April 2013 ARIN meeting, the inevitable question came up once more: "Exactly when is ARIN going to run out of IPv4 addresses?" Various dates have been proposed as an answer to this question, based on various methods of prediction. As the date is indeed getting closer, it may well be worth the time to review ARIN's situation, and make a few predictions here about the likely date when ARIN's exhausts its remaining pool of IPv4 addresses.

I have been publishing such a date for some years now at <http://ipv4.potaroo.net> (which, by the way, currently predicts that ARIN will reach its last /8 block of addresses in mid-April 2014), but in fact this web site is not answering the right question in ARIN's case. Not only is it not answering the right question, it also somewhat misleading. The prediction of an exact date gives the impression that this is all highly predictable, with a high degree of both accuracy and confidence. Paradoxically, the closer we get to exhaustion the less reliable this forecast becomes, and we probably need to consider confidence intervals rather than try to foretell a particular date.

Surely, as we get closer to exhaustion of ARIN's remaining inventory of IPv4 addresses, it is easier to predict the time when that last address block heads out the door? What's so hard about generating a prediction on a date when the addresses will run out, and why would it get more uncertain the closer we get to the date? In this article I'd like to spend a little time explaining why the remaining life for IPv4 in ARIN is becoming less certain rather than more certain as ARIN's IPv4 address pool drains out to the bottom, and then look at the question of predicting the effective exhaustion of the ARIN IPv4 address pool.

Of Elephants and Mice

The main factor here that introduces uncertainty onto the predictive model is that we are operating in a bi-modal world of elephants and mice. I'd like to explore this factor in a little more detail before heading into the statistics.

Statistical predictive techniques all start with the premise that: "Tomorrow will be a lot like today." Even though this will not be the case for everyone, it will be the case for so many folk that we can refine this a little to say: "On average, tomorrow will be a lot like today."

But this assumes a steady state world, which does not reflect the changes we see over time. So we can refine this premise: "On average, the differences that we observed between yesterday and today will be present as differences between today and tomorrow". Again, if we take a single individual case this will probably not be true. Individuals do not typically behave in such a precise regular pattern (or perhaps, despite thinking otherwise, we really are creatures of regular and highly predictable habitual actions, but that's a different area of behavioural study!). However, whether it applies in individual cases, when taken in large enough numbers, in other words looking at the collective actions of masses of individuals, the individual variations are less important, and the average of all the individuals becomes a more accurate description of the group as a whole. Large relatively homogenous groups tend to behave with a higher degree of statistical alignment than individuals.

We can apply this form of statistical analysis to the group of entities who draw addresses from the address registries, and look at the average rate of address consumption, and also look at the variation over time. The first figure is a plot of the amount of addresses allocated by ARIN each day since 1 January 2005.

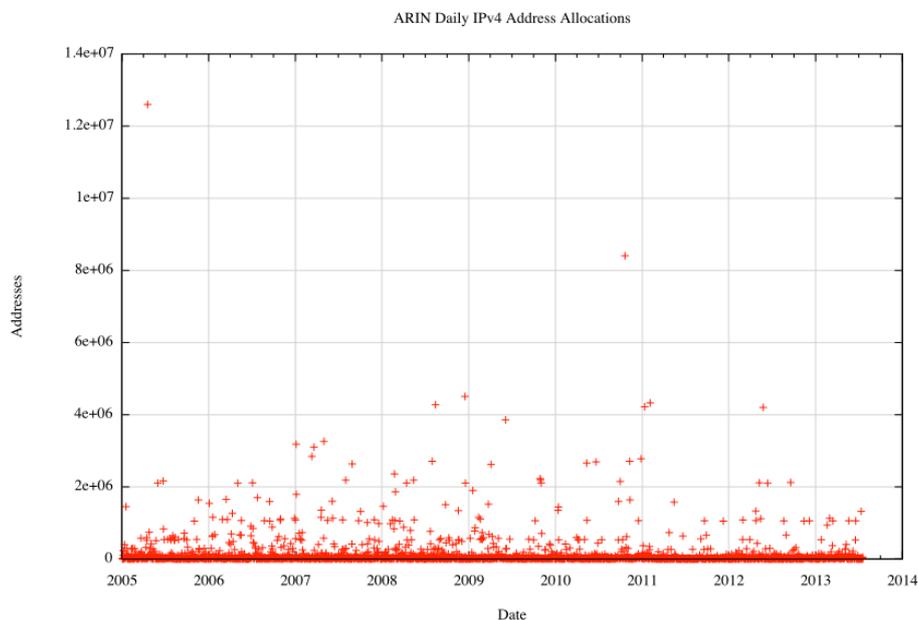


Figure 1 - Daily IPv4 Address Allocations: 2005 to the present

This is a very "noisy" sequence, with a small number of high value outliers that correspond to those days when large allocations of a /9 or /10 in size were made. Any underlying trend in the data sequence is drowned out by these variants. Even using a logarithmic scale does not show any clear trend within the daily allocation data, as shown in Figure 2.

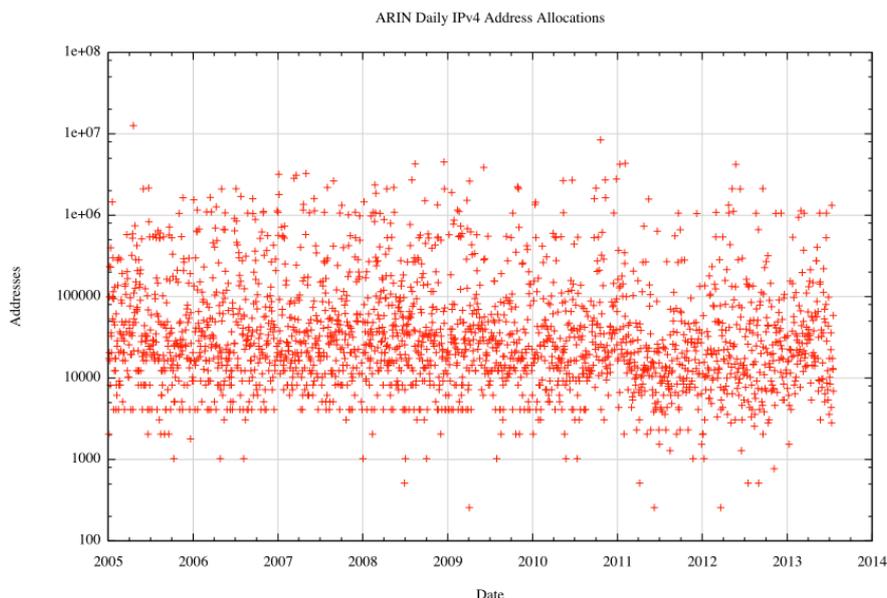


Figure 2 - Daily IPv4 Address Allocations: 2005 to the present - log scale

Perhaps a more critical question to ask is: What is causing this statistical "noise" and is it significant?

Behind this question is a question about the "uniformity" of the population that draws addresses from the pool. Are we seeing the result of a large number of individual transactions, or is this high level of variance caused by a skew in the distribution of individual allocations?

Since address allocations are a reflection of the nature of the Internet industry itself we can rephrase this question somewhat: How "balanced" is this industry? Do we see a diverse collection of actors between large and small? Or has this industry bifurcated into a small number of very large entities and a larger pool of small entities? To address this question, the address allocation profile can reveal a lot about the "profile" of this industry.

Lets looks at the set of recent ARIN allocations, performed from 1 January 2012 to the last week of July 2013, and look at the range of sizes of address allocations:

Allocation Size	No. of Allocations	Total Address Count	% Addresses	Cumulative %
/10	1	4,194,304	9.81%	9.81%
/11	3	6,291,456	14.72%	24.53%
/12	12	12,582,912	29.43%	53.96%
/13	5	2,621,440	6.13%	60.09%
/14	11	2,883,584	6.74%	66.84%
/15	20	2,621,440	6.13%	72.97%
/16	48	3,145,728	7.36%	80.33%
/17	68	2,228,224	5.21%	85.54%
/18	86	1,409,024	3.30%	88.83%
/19	129	1,056,768	2.47%	91.30%
/20	380	1,556,480	3.64%	94.95%
/21	449	919,552	2.15%	97.10%
/22	991	1,014,784	2.37%	99.47%
/23	200	102,400	0.24%	99.71%
/24	485	124,160	0.29%	100.00%
Total	2,888	42,752,256 (2.54 /8s)		

Table 1 – ARIN IPv4 Address Allocations Jan 2012 - July 2013

Of the 2,888 individual allocations performed in the 19 months from January 2012 until late July 2013, some 16 allocations, or 0.5% of all allocations in this period, accounted for 54% of all allocated addresses. This distribution makes any statistical predictive technique challenging, as the most appropriate predictive exercise is no longer some form of averaging function applied to a relatively large homogenous population, but instead we see the outcome being heavily influenced by the actions of a relatively small number of entities, each of whom are requesting a /12 or larger when they request an IPv4 address allocation from the address registry.

Who are these folk who are making these large address requests? Sorting through the allocation data published by ARIN it is possible to list those entities who received a /12 or larger in the period January 2012 - July 2013:

CC	Address	Size	Date	Entity
US	23.192.0.0	1,310,720	2013-07-12	Akamai Technologies Inc
US	40.128.0.0	1,048,576	2012-11-13	Windstream Communications Inc
US	47.58.0.0	1,048,576	2012-05-15	Vodafone Americas Inc.
US	54.192.0.0	1,048,576	2013-06-19	Amazon Technologies Inc.
US	54.224.0.0	1,048,576	2012-03-01	Amazon Technologies Inc.
US	100.140.0.0	1,310,720	2012-04-24	T-Mobile USA, Inc.
US	100.160.0.0	2,097,152	2012-05-09	T-Mobile USA, Inc.
US	100.192.0.0	4,194,304	2012-05-25	T-Mobile USA, Inc.
US	107.192.0.0	1,048,576	2012-04-25	AT&T Internet Services
US	107.208.0.0	1,048,576	2012-12-06	AT&T Internet Services
US	162.160.0.0	2,097,152	2012-06-13	T-Mobile USA, Inc.
US	162.192.0.0	1,048,576	2013-03-01	AT&T Internet Services

US	162.224.0.0	1,048,576	2013-05-21	AT&T Internet Services
US	172.0.0.0	1,048,576	2012-08-20	AT&T Internet Services
US	172.32.0.0	2,097,152	2012-09-18	T-Mobile USA, Inc.
US	172.224.0.0	1,048,576	2013-03-15	Akamai Technologies Inc

Table 2 – ARIN's Large IPv4 Address Allocations Jan 2012 – July 2013

These large address blocks addresses have been assigned to 6 entities, 4 of which are service providers and 2 are content providers.

If we want to understand how long the remaining 31,230,976 unallocated IPv4 addresses (as of 22 August 2013) will last for ARIN, then it's a case of working through a combination of the projections of the average of the 151 smaller allocations per month, and the far less predictable average of slightly under a single large allocation each month. It's the sum of these two factors that will determine how and when ARIN will exhaust its address pool.

However, its not the only factor here that influences the outcome. There are a couple of policy issues as well that are relevant to this consideration.

ARIN's Policy Framework

ARIN's policy framework relating to the run down of the IPv4 address pool is slightly different to that of the RIPE community in Europe and the Middle East and APNIC in Asia Pacific. The RIPE and APNIC regional address policy communities had developed so-called "Last /8" policies, where the last 16,777,216 addresses (or the equivalent of a /8 address block) has been withheld from conventional allocation nby the regional registry, and, instead, once all others IPv4 address pools have been exhausted, each applicant for IPv4 address space can be assigned a total of up to 1,024 addresses from the last /8. Ever. This policy effectively holds a limited pool of IPv4 addresses for "latecomers", which, while the very limited quantity per application implies that such allocations are not useful for conventional address purposes in service provider scenarios, may be adequate for some forms of IPv4 NAT based deployment, such as, for example, 5-tuple NATs, that deliver a very high level of address compression.

ARIN does not have a "last /8" policy. ARIN's analogous policy reserves the last 4,194,304 addresses (or a /10) for assisting in IPv6 deployment. Here's the extract from ARIN's Policy Manual that describes this:

4.10 Dedicated IPv4 block to facilitate IPv6 Deployment

When ARIN receives its last /8 IPv4 allocation from IANA, a contiguous /10 IPv4 block will be set aside and dedicated to facilitate IPv6 deployment. Allocations and assignments from this block must be justified by immediate IPv6 deployment requirements. Examples of such needs include: IPv4 addresses for key dual stack DNS servers, and NAT-PT or NAT464 translators. ARIN staff will use their discretion when evaluating justifications.

This block will be subject to a minimum size allocation of /28 and a maximum size allocation of /24. ARIN should use sparse allocation when possible within that /10 block.

In order to receive an allocation or assignment under this policy:

1. the applicant may not have received resources under this policy in the preceding six months;
2. previous allocations/assignments under this policy must continue to meet the justification requirements of this policy;
3. previous allocations/assignments under this policy must meet the utilization requirements of end user assignments;

4. the applicant must demonstrate that no other allocations or assignments will meet this need;
5. on subsequent allocation under this policy, ARIN staff may require applicants to renumber out of previously allocated / assigned space under this policy in order to minimize non-contiguous allocations.

[<https://www.arin.net/policy/nrpm.html#four10>]

For ARIN, the effective time of exhaustion of “general use” IPv4 addresses is when the available pool of IPv4 addresses in ARIN has dropped to a /10 in size, not a /8.

However, in their reporting of IPv4 address pools, ARIN have already marked an IPv4 address block (23.128.0.0/10) as *reserved* for this purpose, so in our exercise we will run the simulation until ARIN's available pool of IPv4 addresses is fully exhausted.

The IANA IPv4 Runout Global Address Policy

The next policy factor in modelling ARIN's runout is the residual pool of IPv4 addresses held by the IANA. Currently, IANA holds 20,466,432 IPv4 addresses, which are returned addresses from the original IPv4 address allocations made before the RIR framework was out in place.

The policy that describes the management of this IANA address pool reads:

When one of the RIRs declares it has less than a total of a /9 in its inventory, the Recovered IPv4 pool will be declared active, and IP addresses from the Recovered IPv4 Pool will be allocated as follows:

- a. Allocations from the IANA may begin once the pool is declared active.
- b. In each "IPv4 allocation period", each RIR will receive a single "IPv4 allocation unit" from the IANA.
- c. An "IPv4 allocation period" is defined as a 6-month period following 1 March or 1 September in each year.

[<http://www.icann.org/en/resources/policy/global-addressing/allocation-ipv4-post-exhaustion>]

It appears that if the IANA returned address pool remains at its current size of size of 20,466,432 addresses, then each RIR will receive a /11 from the pool in the first allocation period, then a /12 at the start of the following March or September, then a /13, and so on in six monthly intervals for the next 7 years, once any of the RIR's IPv4 address inventory falls below a /9 in size.

Previous Experience with IPv4 Runout

In April 2011 APNIC reached its last /8, and in September 2012 the RIPE NCC also ran down to its last /8. The two experiences in run out were quite dissimilar.

In the case of APNIC the last 4 months saw an increased rate of allocation that was markedly higher than that seen in the previous months. The rundown of APNIC's IPv4 address pool is shown in Figure 3.

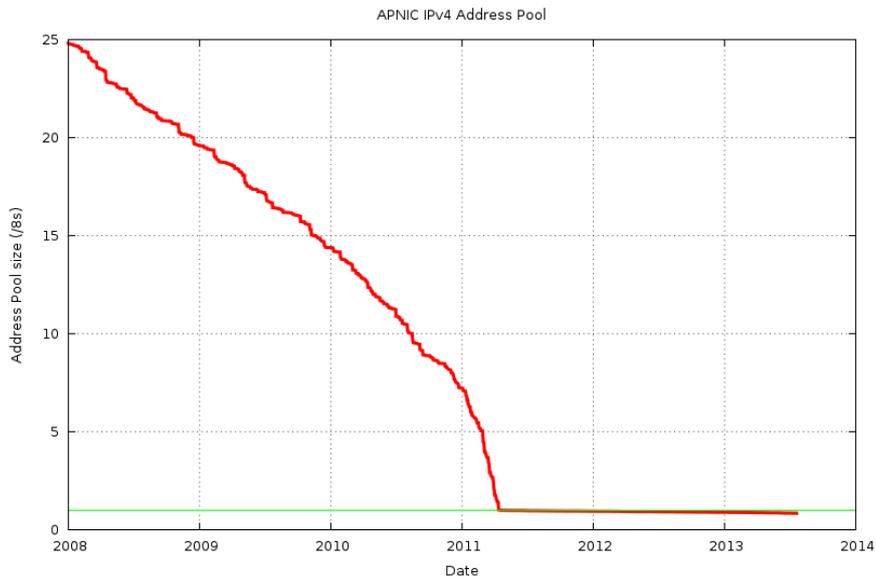


Figure 3 – APNIC IPv4 Address Pool: 2008 - present

APNIC’s address allocation rate was rising steadily from 2006, from an average of 4 /8s p.a. to just under 8 /8s p.a. by later 2010. However in December 2010 the allocation rate jumped to a rate of up to 18 /8s p.a., and the last 6 /8s were allocated in a period of 4 months. This acceleration of the address consumption rate is clearly shown in Figure 4, which shows the smoothed daily allocation rate for APNIC across this period.

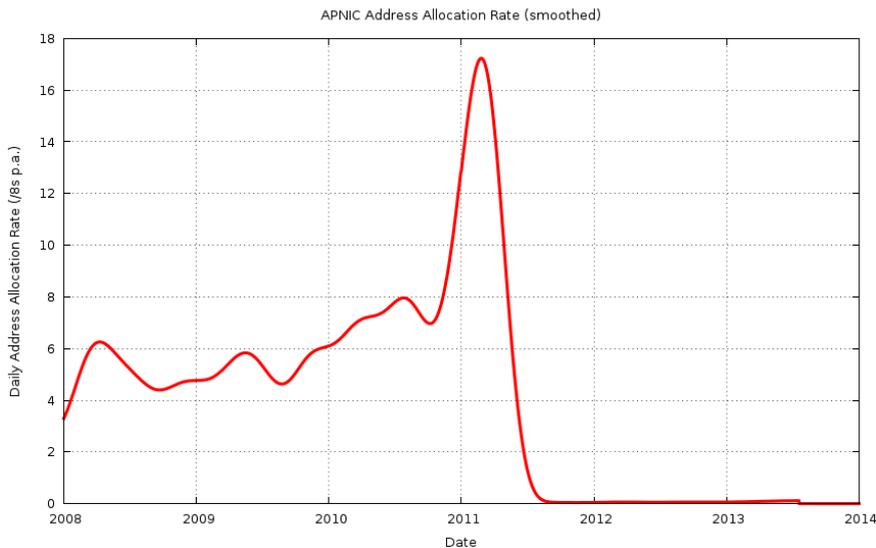


Figure 4 – APNIC IPv4 Address Consumption Rate: 2008 - present

In the case of the RIPE NCC no such last minute allocation acceleration was visible in the data, and the allocation rate remained steady as an average of some 3 /8s p.a. through 2012, right to the point where the last /8 was reached in September of that year (Figures 5 and 6).

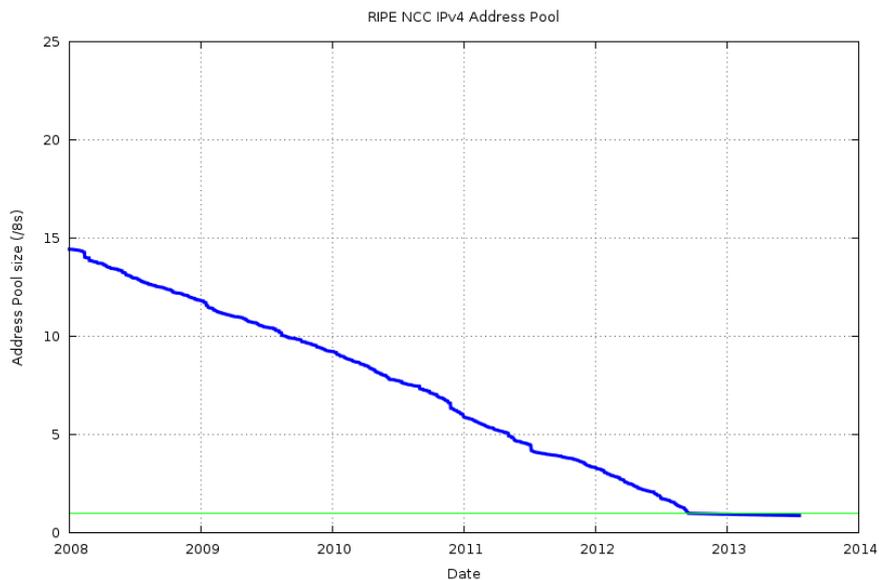


Figure 5 – RIPE NCC IPv4 Address Pool: 2008 - present

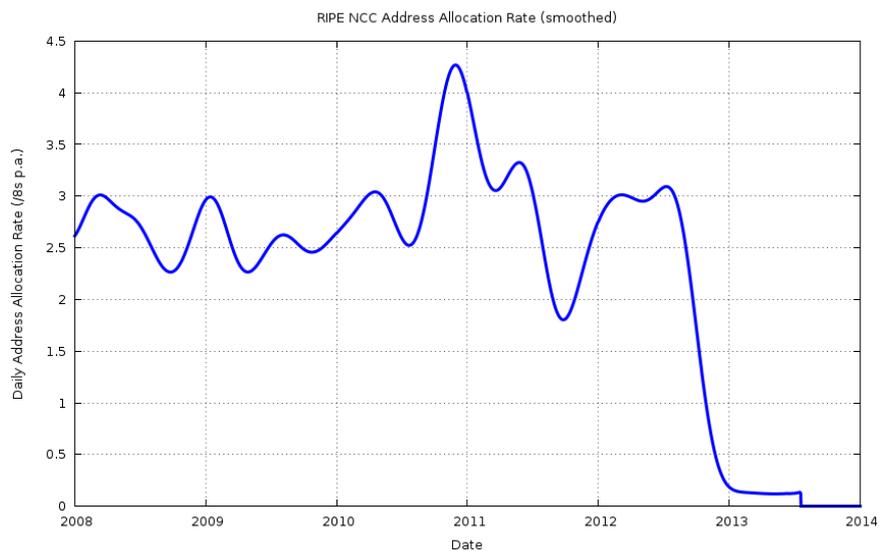


Figure 6 – RIPE NCC IPv4 Address Consumption Rate: 2008 - present

ARIN's Runout

A model of IPv4 address exhaustion for ARIN starts with the data so far. Instead of looking at the daily address allocation rates, as in Figures 1 and 2, we can instead look at the model of the RIR as a reservoir of addresses. Instead of looking at the amounts of addresses withdrawn from the reservoir each day, let's start with the quantity of addresses held in the reservoir each day. This is shown in Figures 3 and 5 for APNIC and the RIPE NCC respectively, and in Figure 7 for ARIN.

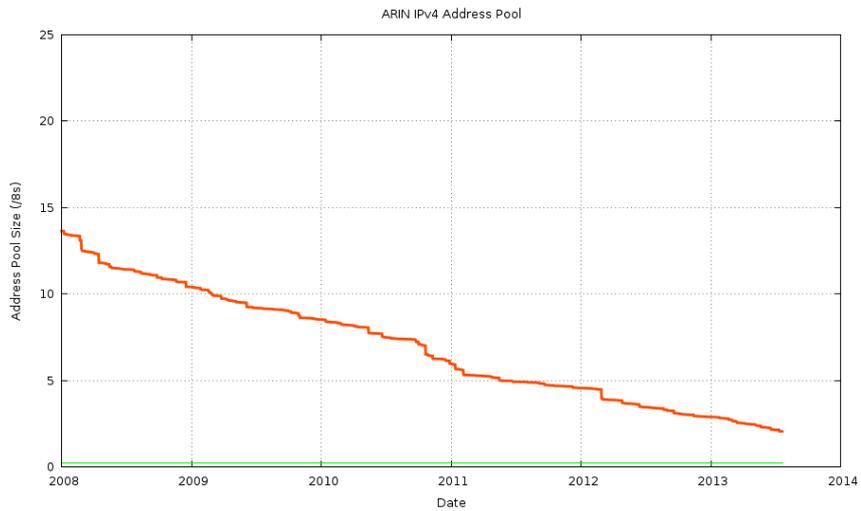


Figure 7 – ARIN IPv4 Address Pool: 2008 - present

Looking at the data series in Figure 7, starting from 1 January 2011, we can apply a least squares algorithm to the data series, to get a best fit straight line. This corresponds to a “best fit” fixed allocation rate of the equivalent of 1.43 /8s p.a. This best fit model can be compared to the smoothed daily allocation rate, as shown in Figure 8.

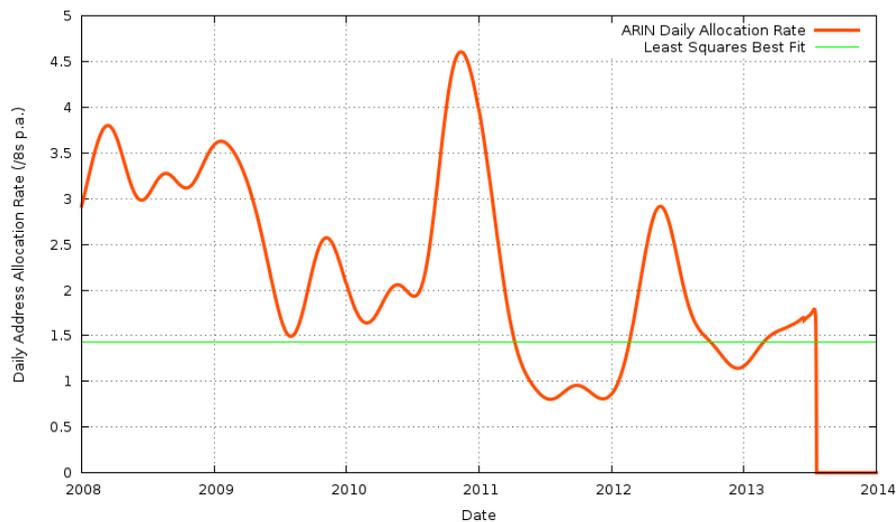


Figure 8 – ARIN IPv4 Address Consumption Rate: 2008 - present

If we take this simple linear model, and extrapolate it forward to the point when the ARIN pool reaches a /10 in size, we get an exhaustion model shown in Figure 9.

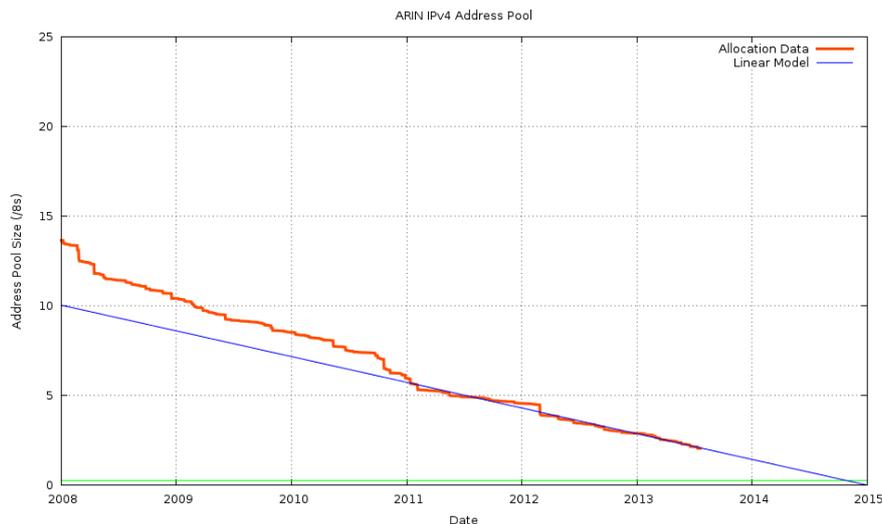


Figure 9 – ARIN IPv4 Address Consumption Model

This linear model exhausts the ARIN address pool on the 28th December 2014, at which time the only address pool left in ARIN’s registry is that which has been set aside to facilitate IPv6 deployment, and conventional address allocations would cease at that time.

A conventional model of a simple linear projection based on a least squares best fit to address consumption data since 1 January 2011 provides a prediction of address runout for ARIN for sometime in the third quarter of 2014. But this simple extrapolation of demand does not take in to account the IANA IPv4 Runout Global Address Policy, which will provide a small increment to ARIN's remaining address pool once ARIN's remaining address pool shrinks below the initial /9 threshold and every 6 months thereafter.

How realistic is this model?

This linear “steady consumption” model of address consumption is a reasonable match against the experience of the RIPE NCC under runout, when the exhaustion of the IPv4 address pool did not trigger any notable last minute panic or any major acceleration of the address consumption rate.

However, APNIC’s experience was markedly different. Some four months prior to exhaustion the address consumption rate jumped from a rate equivalent of 8 /8s p.a. at the start of 2011 to a final rate of the equivalent of 18 /8s p.a. at the time of address pool exhaustion. Even in the year prior to this, 2010, the address consumption rate for APNIC climbed from 6 /8s p.a. at the start of 2010 to a final rate of 8 /8s p.a. at the end of the year. The challenge in examining the APNIC data is that there is no clear correlation between the 2010 address consumption data and the 2011 data. Any curve fitting technique, whether it's a polynomial function or an exponential function, could be shaped to fit the APNIC 2011 data, but such a curve fitting exercise would be just that, a mathematical technique that in and of itself would offer no particular insight into the underlying demand models in the AP region that lead to the accelerated demands in 2011.

However there is one policy difference between APNIC and the RIPE NCC that may be relevant here. In APNIC, the “horizon” for the justification of need for further address allocations was held steady at 12 months thought the address depletion period, while in the RIPE NCC, and in ARIN, the “horizon” for demonstrated need for addresses was reduced to 3 months following the exhaustion of the IANA /8 registry in February 2011.

This implies that at any time the quantity of the “future demand” for the RIPE NCC was one quarter of that exposed in APNIC, by virtue of the difference in the time period relating to demonstrated need. And the same holds for ARIN. Relating this to table 2, the schedule of address consumption from the

6 largest address consuming entities in ARIN over the 2012 – 2013 period, this would imply that AT&T Internet Services is expressing a demand level of some 300,000 addresses per month, and is due to return to ARIN for a further allocation of 1M addresses in September 2013. By the same reasoning it would be reasonable to anticipate that Akami Technologies will return in October 2013 for a further allocation of some 1.3M addresses. This shorter horizon in the address allocation policy framework tends to reduce the volume of unmet demand at any point in time, and tends to point to a model of runout that is more aligned to a steady state model as experienced by the RIPE NCC, rather than one of accelerating demand close to the exhaustion time.

However, a close examination of Figure 8 shows that over 2013 the demand levels in ARIN have increased from a rate of around 1.2 /8s p.a. at the start of the year to a rate of around 1.8 /8s p.a. in July 2013. What if this accelerated demand rate were to persist in the coming months?

One way to answer this is to adjust the time interval used to generate the predictive model. If we just consider the allocation data from the 1st January 2013 through to the middle of August 2013, then a linear best fit model for this 2013 data raises the model's address consumption rate, from 1.43 /8s p.a. which is the average address consumption rate since January 2011 to 1.62 /8s p.a. An O(2) polynomial fit gives a very similar outcome, as shown in Figure 10. Why do the two functions show such a high degree of similarity? The 2013 allocation data is itself very regular.

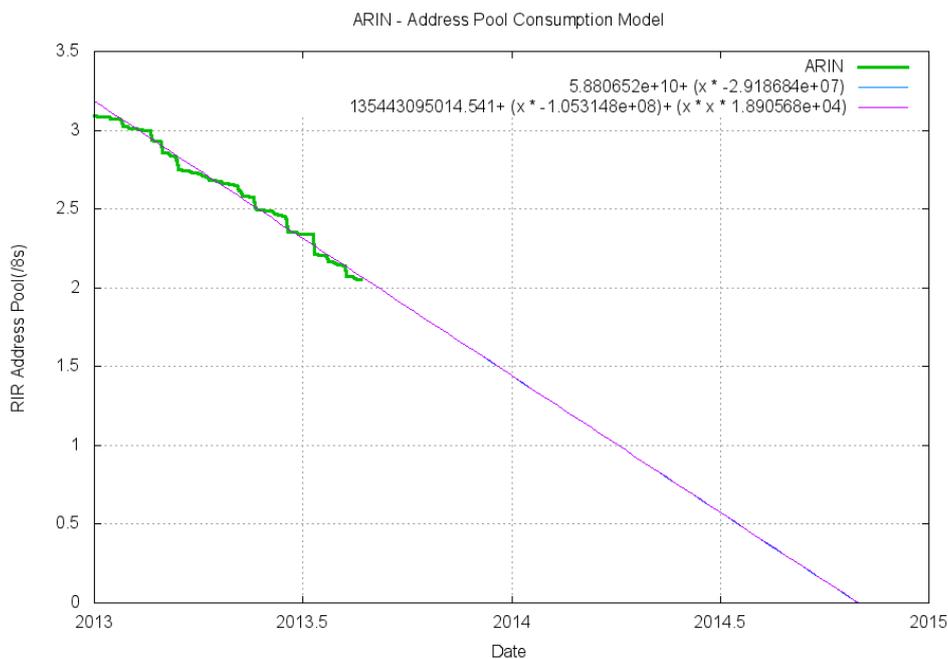


Figure 10 – ARIN IPv4 Address Consumption Model: 2013 data

In this case, when using just the 2013 address allocation data as the input to the projection model, the projected date of the exhaustion of ARIN's IPv4 address pool is the 29th October 2014, some two months earlier than the original prediction.

What this leads to is a prediction that, in terms of the application of a statistical projection technique, we are looking at a likely exhaustion date for ARIN in the last quarter of 2014, most likely in October or November of 2014.

The same caveat applies to these estimates as well: This simple extrapolation of demand does not take in to account the IANA IPv4 Runout Global Address Policy, which will provide a small increment to ARIN's remaining address pool once ARIN's remaining address pool shrinks below the initial /9 threshold and every 6 months thereafter.

Theme and Variations

What about the issue of the distribution of allocation sizes? To what extent does the small pool of large allocations add uncertainty to these predictions? Can we assign a confidence level to this prediction?

How can the variation on allocation sizes shown in Table 1 be used to give us a confidence interval for these predictions? The issue is, as outlined above, that it is difficult to predict the actions of a relatively small number of entities who are recipients of large allocations from ARIN, and this makes the statistical techniques used so far to have an associated high level of uncertainty.

While we are asking questions, can we include the impacts of the application of the IANA IPv4 Global Runout Address Policy?

One way to quantify this uncertainty is with an experimental technique of discrete event simulation. The data shown in Table 1 can be used to assign probabilities for the allocation of individual prefix sizes. The data spans 568 days of allocations, so the recorded allocation of 11 /14 prefixes over this period can be interpreted as saying that on any single day in this period the probability of assigning a /14 is 11/568, or 1.9%.

If we regard the allocation of each individual prefix as an independent event, then for each day we can generate a particular allocation profile by adding an allocation of each prefix size according to a random draw. In other words, the allocation profile of a particular day will include a /14 allocation only if a random value drawn from the range 0 – 1,000 is 19 or less, which should occur 1.9% of the time.

Applying this method over 500,000 separate randomized sequences of simulating ARIN’s profile of allocations from the current date to the point of exhaustion of the ARIN address pool gives a distribution of exhaustion dates shown in Figure 11. The useful aspect of this technique is that each potential exhaustion date now has a probability, and we can use a cumulative distribution of these individual probability values to derive an interval when exhaustion is likely to occur, and an associated level of confidence of that prediction (Figure 12).

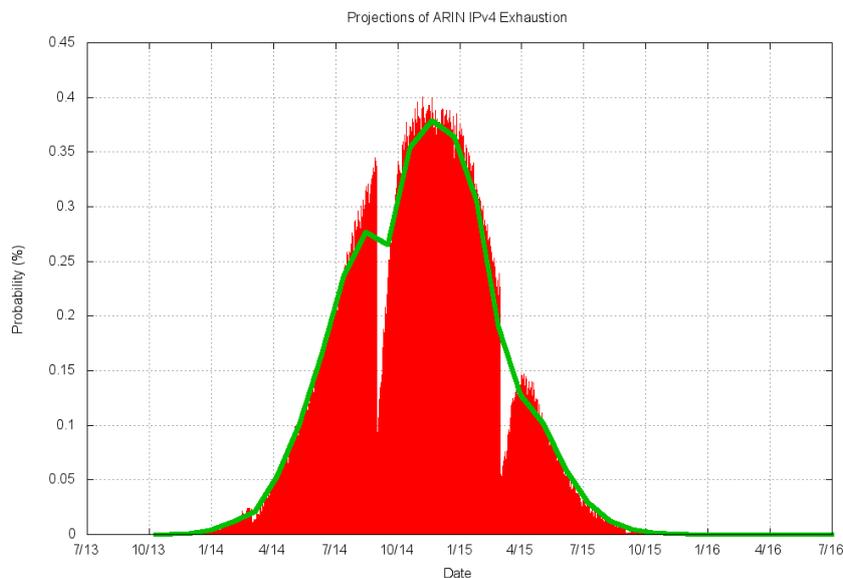


Figure 11 – Variance of Exhaustion Dates

The extreme outliers using this technique are an early date of mid-October 2013, and a late date of January 2016. But of course these outlier exhaustion dates are extremely unlikely. This technique predicts that, with a 90% level of confidence, ARIN will exhaust its remaining IPv4 addresses between the 16th May 2014 and the 10th May 2015. At a 50% confidence level, the exhaustion interval for ARIN is from the 19th August 2014 to the 15th January 2015, with a median date of the 19th November 2014..

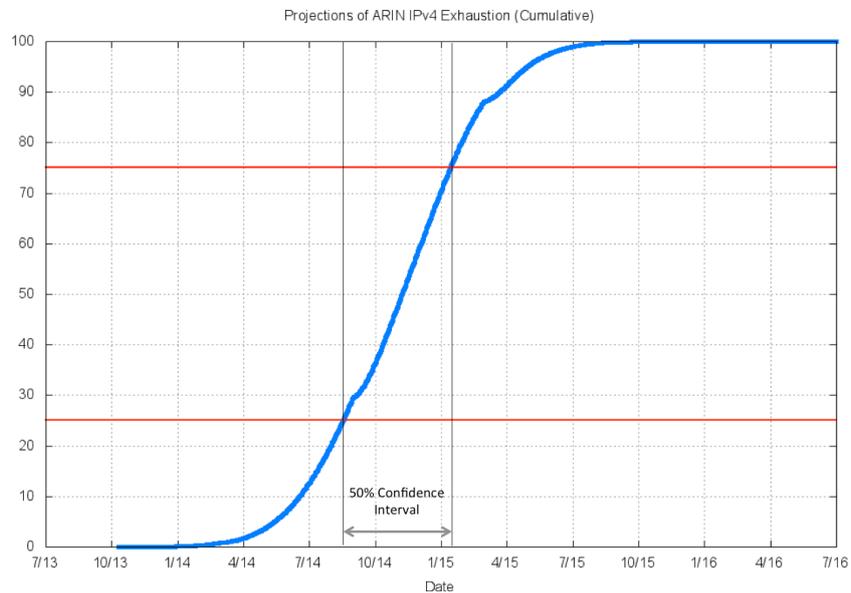


Figure 12 – Confidence Intervals of Exhaustion Dates

Which gives us some basis to be able to answer the original question, assuming, as we noted at the outset, that the profile of address consumption for ARIN for tomorrow will continue to be a lot like today.

When?

Sometime in the latter half of 2014 or possibly early in 2015. Probably.

In the last quarter of 2014. Perhaps.

Disclaimer

The views expressed are the authors' and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

About the Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. 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.

www.potaroo.net