January 2018
Geoff Huston

Addressing 2017

Time for another annual roundup from the world of IP addresses. Let's see what has changed in the past 12 months in addressing the Internet, and look at how IP address allocation information can inform us of the changing nature of the network itself.

There is no doubt that the Internet continues to grow. While the sales volumes of the more traditional forms of personal computers has peaked at some 430 million units per year (according to Gartner's annual roundups of the industry) and sales of handheld smart devices has also peaked at some 1.9 billion units per year, the world of IoT continues to spiral upward. The installed based on IoT units is now at 8.4 billion at the end of 2017.

Back in around 1992 the IETF gazed into the crystal ball and tried to understand how the internet was going to evolve and what demands that would place on the addressing system as part of the "IP Next Generation" study. The staggeringly large numbers of connected devices that we see today were certainly within the range predicted by that exercise. Looking further out, it is doubtless that these numbers will continue to grow. We continue to increase silicon production volumes and at the same time continue to refine the production process to decrease the unit costs of these chips. But, at that time, we also predicted that the only way we could make the Internet work across such a massive pool of connected devices was to deploy a new IP protocol that came with a massively larger address space. It was from that reasoning that IPv6 was designed. This world of abundant silicon was the issue that IPv6 was primarily intended to solve. The copious volumes of address space were intended to allow us to uniquely assign a public IPv6 address to every such device, no matter how small, or in what volume they might be deployed.

But while the Internet has grown at such amazing rates, the deployment of IPv6 continues at a far more leisurely pace. There is no common sense of urgency about the deployment of this protocol, and still there is no hard evidence that the continued reliance on IPv4 is failing us. Much of the reason for this apparent contradiction is that the Internet is now a client/server network. Clients can initiate network transactions with servers, but are incapable of initiating transactions with other clients. Network Address Translators (NATs) are a natural fit to this client/server model, where pools of clients share a smaller pool of public addresses, and only required the use of an address while they have an active session with a remote server. NATs are the reason why in excess of 15 billion connected devices can be squeezed into some 2 billion active IPv4 addresses.

However, the pressures of this inexorable growth in the number of deployed devices means that the even NATs cannot withstand these growth pressures forever. Inevitably, either we will see the fragmenting of an IPv4 Internet into a number of disconnected parts, so that the entire concept of a globally unique and coherent address pool will be foregone, or we will see these growth pressures motivate the further deployment of IPv6, and the emergence of IPv6-only elements of the Internet as it tries to maintain a cohesive and connected whole. There are commercial pressures pulling the network in both of these directions, so it's entirely unclear what path the Internet will follow in the coming years.

Can address allocation data help us to shed some light on what is happening in the larger Internet? Let's look at what happened in 2017.

IPv4 in 2017

It appears that the process of exhausting the remaining pools of unallocated IPv4 addresses is proving to be as protracted as the process of the transition to IPv6.

The allocation of 16 million addresses in 2017 on top of a base of 3,641 million addresses that are already allocated at the start of the year represents a growth rate of 0.43% for the year for the total allocated IPv4 public address pool. This is less that one tenth of the growth in 2010 (the last full year before the onset of IPv4 address exhaustion).

	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
Allocated (M)	203.3	189.4	248.8	201.0	114.9	65.1	63.9	34.8	22.2	15.6
Total (B)	2.52	2.72	2.90	3.14	3.34	3.43	3.50	3.59	3.62	3.65
Relative Growth	7.9%	6.6%	8.3%	6.4%	2.9%	1.9%	1.8%	1.0%	0.6%	0.6%

Table 1 - IPv4 Allocated addresses by year

The record of address allocations per RIR over the past 10 years is shown in Table 2.

	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
APNIC	87.8	86.9	120.2	105.2	1.0	1.3	3.7	4.1	3.8	1.8
RIPE	44.0	43.4	56.0	43.1	40.0	2.0	2.5	3.3	3.4	3.6
ARIN	57.1	41.1	45.2	23.5	45.0	26.5	26.0	8.6	1.6	0.8
LACNIC	12.0	10.5	13.0	24.4	21.0	28.5	19.1	1.8	1.6	1.7
AFRINIC	1.6	5.9	8.5	9.2	7.9	6.8	12.5	16.9	11.8	7.8

Table 2 - IPv4 Allocated addresses (millions) - Distribution by RIR

In terms of the IPv4 Internet there is a considerable diversity in the situation in each region. As of the end of 2017, AFRINIC was the last remaining Regional Internet Registry (RIR) with remaining IPv4 addresses available for general allocation, with some 12.8 million addresses left in its available address pool. APNIC and the RIPE NCC have both adopted "Last /8" policies, where each applicant can receive just a single allocation of up to 1,024 addresses from their respective last /8 address pools. APNIC has 5.6 million addresses left in this pool, and the RIPE NCC has some 9.9 million addresses. LACNIC has a pool of 277,000 remaining addresses, while ARIN has none at all.

We can use the address allocation data from 2017 and perform a forward extrapolation on this to predict when the available address pools of each RIR will filly deplete. This is shown in Figure 1.

The address consumption rate for APNIC reduced slightly in 2017 as compared to previous years, and at this stage the pool will last for a further 3 years at this allocation rate. The RIPE NCC uses a similar address management policy for its remaining pool of addresses, but the consumption rate is slightly higher than that of APNIC, and it increased in 2017 as compared to previous years, so this pool will last for a further 2½ years at its current rate of consumption. LACNIC's remaining address pool will last for a further 1½ years, which is similar to the situation in AFRINIC. ARIN has completely exhausted its available pool. The picture with AFRINIC is not as clear. There have been a small number of relatively large single allocations in recent years. The first half of 2017 also saw a high level of activity, while the allocation rate in the second half of the year was considerably lower. The 2017 average allocation rate of ½ of a /8 per year will see the AFRNIC address pool last for a further 1½ years. This is shown in Table 3.

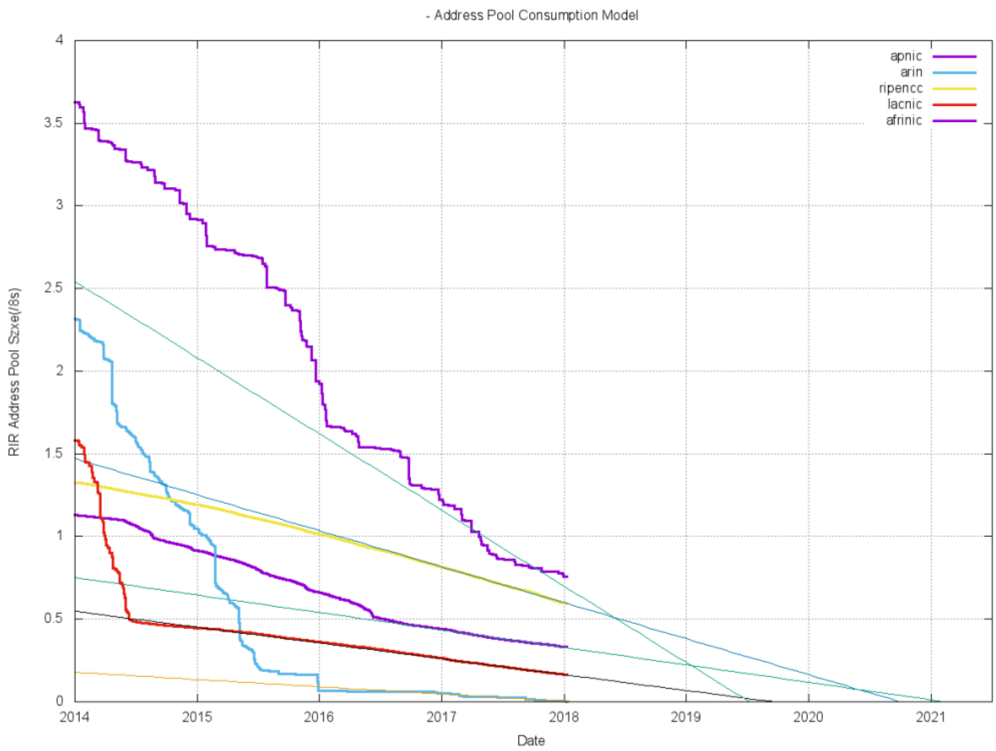

Figure 1 – IPv4 RIR pool runout scenarios

ARIN	exhausted
AFRINIC	mid-2019
LACNIC	mid-2019
RIPE NCC	mid-2020
APNIC	early 2021

Table 3 - IPv4 Address Pool Depletion Projections by RIR

This analysis of the remaining address pools is not quite the complete picture, as each of the RIRs also have reserved some addresses, in accordance with their local policies. There are a variety of reasons for this reservation, including non-contactability of the original address holder, or addresses undergoing a period of 'quarantine' following a forced recovery, or a reservation as prescribed by a local policy. ARIN has 6.0 million reserved IPv4 addresses, APNIC has 4.2 million, AFRINIC 2.0 million, the RIPE NCC has 1.1 million, and LACNIC 1.0 million. The total pool of reserved IPv4 addresses is some 14.4 million addresses in size.

Finally, the IANA is holding 18,688 addresses in its recovered address pool in 64 discrete address blocks. The forthcoming relatively small allocations to each RIR from this address pool will have little in the way of impact on the overall IPv4 picture.

The RIR IPv4 address allocation volumes by year are shown in Figure 2.

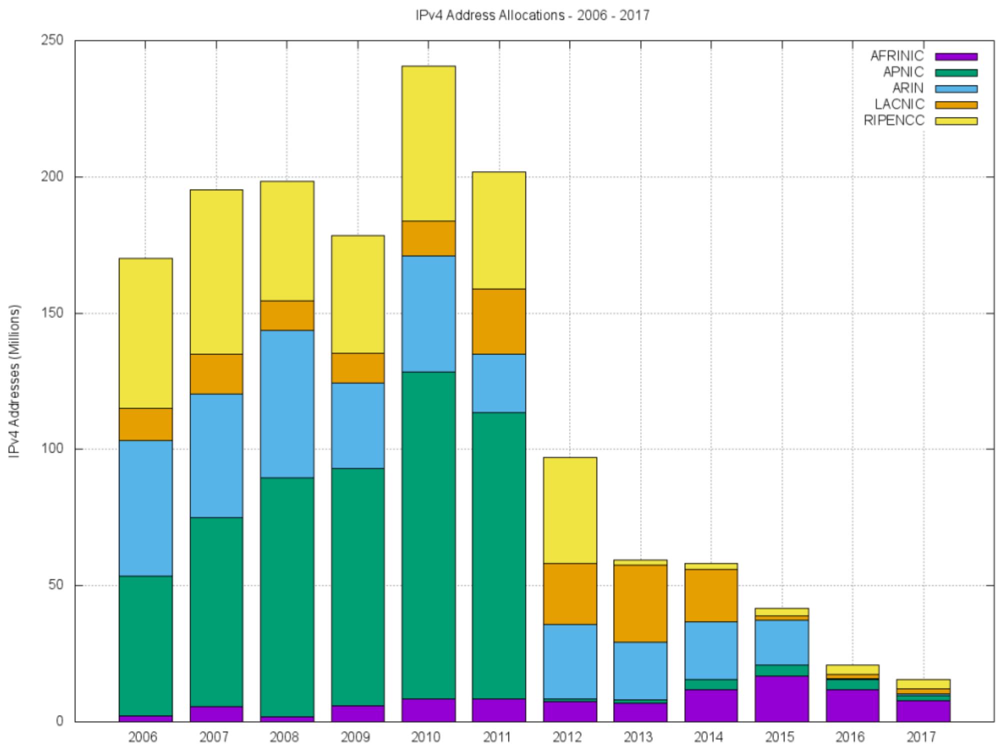

Figure 2 – IPv4 Allocations by RIR by year

IPv4 Address Transfers

In recent years, several RIRs (RIPE NCC, ARIN and APNIC) have included the registration of IPv4 transfers between address holders, as a means of allowing secondary re-distribution of addresses as an alternative to returning unused addresses to the registry. This has been in response to the issues raised by IPv4 address exhaustion, where the underlying motivation as to encourage the reuse of otherwise idle or inefficiently used address blocks through the incentives provided by a market for addresses, and to ensure that such address movement is publically recorded in the registry system.

The numbers of registered transfers in the past four years is shown in Table 4.

Receiving RIR	2012	2013	2014	2015	2016	2017
ARIN	79	31	58	277	727	1,260
APNIC	255	206	437	514	581	466
RIPE NCC	10	171	1,050	2,852	2,411	1,671
Total	**344**	**408**	**1,545**	**3,643**	**3,719**	**3,397**

Table 4 - IPv4 Address Transfers per year

A slightly different view is that of the volume of addresses transferred per year (Table 5).

Receiving RIR	2012	2013	2014	2015	2016	2017
ARIN	6,728,448	5,136,640	4,737,280	37,637,888	15,613,952	37,942,528
APNIC	3,434,496	2,504,960	4,953,088	9,836,288	7,842,816	4,283,136
RIPE NCC	65,536	1,977,344	9,635,328	10,835,712	9,220,864	20,615,168
Total	**10,228,480**	**9,618,944**	**19,325,696**	**58,309,888**	**32,677,632**	**62,840,832**

Table 5 – Volume of Transferred IPv4 Addresses per year (millions of addresses)

A plot of these numbers is shown in Figures 3 and 4.

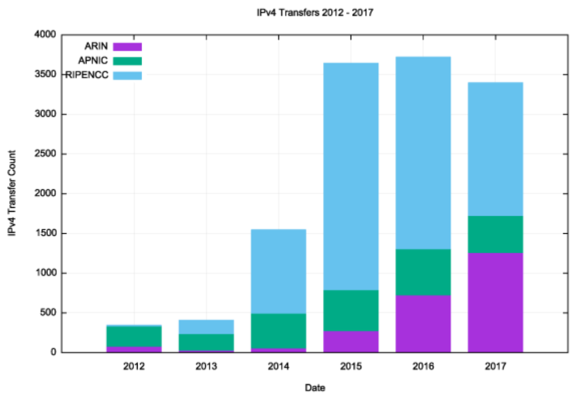

Figure 3 – Number of Transfers: 2012 - 2017

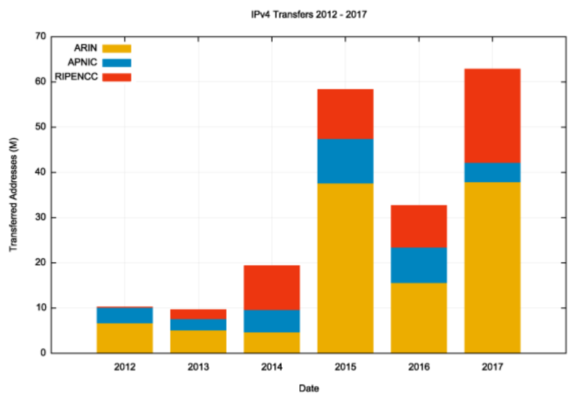

Figure 4 – Volume of Transferred Addresses: 2012 - 2017

The total volume of addresses transferred in this way is four times the volume of allocated addresses across 2017. The aggregate total of addresses in the transfer logs is some 193 million addresses, or the equivalent of 11.5 /8s.

This data raises some questions about the nature of transfers.

The first question is whether address transfers have managed to be effective in dredging the pool of allocated but unadvertised public IPv4 addresses. It was thought that by being able to monetize these addresses, holders of such addresses may have been motivated to convert their networks to use private addresses and resell their holding of public addresses. The numbers appear to show that this has happened, although progress has been slow. At the onset of IPv4 address exhaustion in early 2011 the unadvertised pool was at the equivalent of 54 /8s and it was down to 45 /8s at the end of 2017 (Figure 5). Some 150 million IPv4 addresses have been recirculated back into the advertised address space over this 7 year period. In relative terms the pool dropped from 27% of the total allocated address pool to 21% in the same period (Figure 6). This data appears to support the observation that address scarcity has managed to increase the efficiency of the IPv4 address pool, by bringing into the routed space addresses that were either previously idle or were used in private contexts.

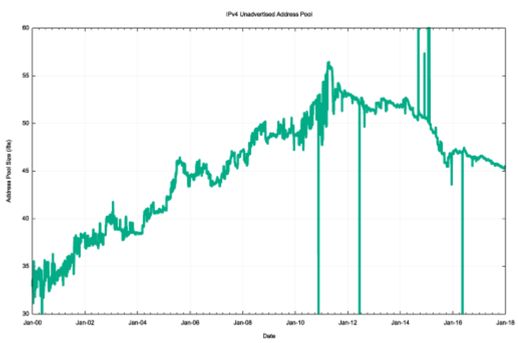

Figure 5 – IPv4 Unadvertised Address Pool size

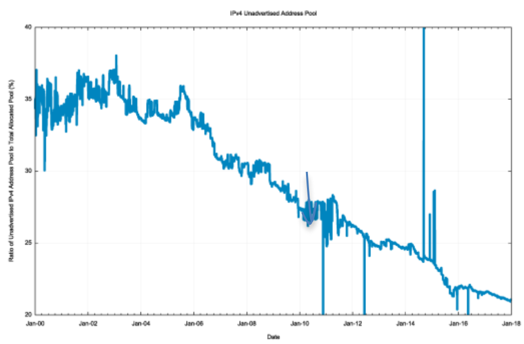

Figure 6 – Ratio of Unadvertised Pool to Total Pool

There is a slightly different aspect to this question, concerning whether the transferred addresses are predominately recently allocated addresses, where there may be the potential for arbitrage between the costs of receiving an address allocation from an RIR and the potential returns from selling these address holdings on the transfer market, and longer held address addresses where the holder is wanting to realise otherwise unused assets. The basic question concerns the "age" distribution of transferred addresses where the "age" of an address reflects the period since it was first allocated or assigned by the RIR system.

The cumulative age distribution of transferred addresses is shown on a year-by-year basis in Figure 7. In 2012 more than 80% of the transferred address blocks were originally assigned or allocated by an RIR within the previous 10 years. In 2016 this has dropped to around to 10% of transferred addresses, but in 2017 the trend reversed, and some 45% of all transferred addresses were less than 10 years old.

The high volumes of transfer activity associated with legacy addresses appears to have peaked in 2016, and the transfer activity has a slightly more even distribution across the age range.

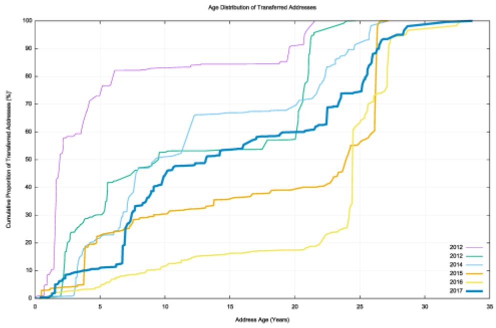

Figure 7 – Age distribution of transferred addresses

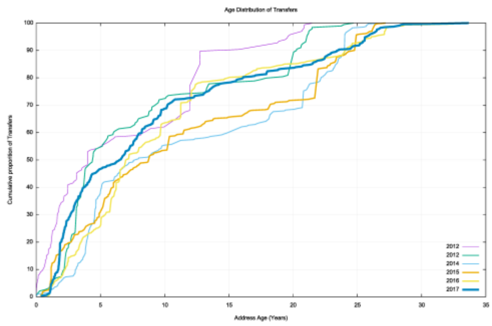

Figure 8 – Age distribution of transfer transactions

If we look at transfer transactions irrespective of the size of each transfer, we get a slightly different picture (Figure 8). One half of all transfer transactions deal with addresses that are no older than 7 years, and this has been the case in each of the past six years. This data indicates that a large number of transaction activity deals with small address blocks that have been allocated relatively recently, while the legacy address blocks tend to be transferred as larger address blocks.

The second question is whether the transfer process is further fragmenting the address space by splitting up larger address blocks into successively smaller address blocks. There are 11,607 address blocks described in the transfer registries up to the end of 2017, and of these 4,558 entries list transferred address blocks that are the same size as the original allocated block. The remaining 7,049 entries are fragments of the originally allocated address blocks.

One third of all original address blocks that are transferred (1,921 out of 6,278) are split into smaller fragments with multiple holders, and on average this results in approximately slightly less than four different holders of transferred address fragments.

This data implies that the answer to the second question is that address blocks are being fragmented as a result of address transfers, but in absolute terms this is not a major issue. There are some 182,026 distinct address allocations from the RIRs to end entities, and the fragmentation of 1,921 of these address blocks is around 1% of the total pool of allocated address prefixes.

The third question concerns the inter-country flow of transferred addresses. Let's look at the ten countries that sourced the greatest volume of transferred addresses, irrespective of their destination (i.e. including 'domestic' transfers within the same country) (Table 6), and the ten largest recipients of transfers (Table 7), and the ten largest country-to-country address transfers (Table 8). In this case we will use the published transfer data in all years up to the end of 2017.

Rank	CC	Addresses	Country Name
1	US	110,978,304	USA
2	CA	19,074,048	Canada
3	FR	12,230,912	France
4	RO	7,321,856	Romania
5	GB	7,081,728	UK
6	RU	4,895,232	Russian Federation
7	JP	3,455,744	Japan
8	DE	3,221,504	Germany
9	CN	3,122,432	China
10	HK	2,089,984	Hong Kong SAR

Table 6 – Top 10 Countries Sourcing Transferred IPv4 addresses

Rank	CC	Addresses	Country Name
1	US	113,937,408	USA
2	FR	13,216,000	France
3	CN	10,270,208	China
4	IN	6,789,632	India
5	GB	6,669,824	UK
6	JP	5,580,800	Japan
7	IR	4,161,792	Iran
8	RU	3,353,600	Russian Federation
9	SA	2,930,176	Saudi Arabia
10	DE	2,800,896	Germany

Table 7 – Top 10 Countries Receiving Transferred IPv4 addresses

Rank	From-CC	To-CC	Addresses	From	To
1	US	US	98,388,224	USA	USA
2	CA	US	14,513,408	Canada	USA
3	FR	FR	12,069,120	France	France
4	GB	GB	5,093,376	UK	UK
5	US	IN	4,158,976	USA	India
6	JP	JP	3,324,416	Japan	Japan
7	RU	RU	3,068,160	Russia	Russia
8	CN	CN	2,605,824	China	China
9	CA	CN	2,359,296	Canada	China
10	US	CN	2,245,632	USA	China

Table 8 – Top 10 Country-to-Country IPv4 address transfers

The transfer logs contain 7,436 domestic address transfers, with a total of 139,465,472 addresses, while 4,808 transfers appear to result in a movement of addresses between countries, involving a total of 54,947,072 addresses.

The total volume of addresses reassigned in this manner, some 194 million IPv4 addresses over eight years, is far less than the underlying pre-exhaustion address demand levels that peaked at some 250 million addresses in a single year. It appears that the address supply hiatus has motivated most Internet service providers to use address sharing technologies, and, in particular, Carrier Grade NAT (CGN), on the access side and server pooling on the content side as a means of increasing the level of sharing of addresses. This has been accompanied by a universal shift of the architecture of the Internet to a client/server model, where clients are dynamically assigned IP addresses in order to communicate with servers (via NATs) and many servers use common IP addresses via name-sharing constructs. The result is that the pressure of the IP address space has been relieved to a considerable extent, and the sense of urgency to migrate to an all-IPv6 network has been largely, but not completely, mitigated in recent years.

The outstanding question about this transfer data is whether all address transfers that have occurred have been duly recorded in the registry system. This question is raised because registered transfers require conformance to various registry policies, and it may be the case that only a subset of transfers are being recorded in the registry as a result. This can be somewhat challenging to detect, particularly if such a transfer is expressed as a lease or other form of temporary arrangement, and if the parties agree to keep the details of the transfer confidential.

It might be possible to place an upper bound on the volume of address movements that have occurred in any period is to look at the Internet's routing system. One way to shed some further light on what this upper bound on transfers might be is through a simple examination of the routing system, looking at addresses that were announced in 2017 by comparing the routing stable state at the start of the year with the table state at the end of the year (Table 9).

	Jan-17	Jan-18	Delta	Unchanged	Re-Home	Removed	Added
Announcements	646,059	698,680	52,621	557,812	17,366	70,881	123,502
Root Prefixes:	309,093	332,487	23,394	275,278	11,041	23,166	46,168
Address Span (/8s)	158.34	160.86	2.52	149.06	2.64	6.70	9.91
More Specifics:	336,966	366,193	29,227	282,534	6,325	47,716	77,334
Address Span (/8s)	56.04	57.37	1.33	49.72	0.75	6.32	6.90

Table 9 – IPv4 BGP changes over 2017

While the routing table grew by 52,621 entries over the year, the nature of the change is slightly more involved. Some 70,881 prefixes that were announced at the start of the year were removed from the routing system through the year, and 123,502 prefixes were announced by the end of the year that were not announced at the start of the year. (Without the scope of this study I have not tracked the progress of announcements through the year, and it is likely that more prefixes were announced and removed on a transient basis through the course of the year.) A further 17,366 prefixes had changed their originating Autonomous System number, indicating some form of change in the prefix's network location in some manner (Table 9).

We can compare these changed prefixes against the transfer logs for the two year period 2016 and 2017. Table 10 shows the comparison of these routing numbers against the set of transfers that were logged in these two years.

Type	Listed as Transferred	Unlisted	Ratio
Re-Homed			
All	1,123	16,243	6.9%
Root Prefixes	891	9,745	9.1%
Removed			
All	2,746	68,135	4.0%
Root Prefixes	1,655	21,510	7.7%
Added			
All	6,602	116,900	5.6%
Root Prefixes	4,021	42,147	9.5%

Table 10 – Routing changes across 2017 compared to the Transfer Logs

These figures show that some 4-10% of changes in advertised addresses are reflected as changes as recorded in the RIRs' transfer logs. This should not imply that the remaining 90-96% of changes in advertised prefixes reflect unrecorded address transfers. There are many reasons for changes in the advertisement of an address prefix and a change in the administrative controller of the address is only one potential cause. However, it does establish some notional upper ceiling on the number of movements of addresses in 2017, some of which relate to transfer of operational control of an address block, that have not been captured in the transfer logs.

Finally, we can perform an age profile of the addresses that were Added, Removed and Re-Homed during 2017, and compare it to the overall age profile of IPv4 addresses in the routing table. This is shown in Figure 9. In terms of addresses that were added in 2017, they differ from the average profile due to a skew in favour of "recent" addresses, and 20% of all announced addresses were allocated or assigned in the past 18 months. In terms of addresses that were removed from the routing system, there is a disproportionate volume of removed addresses that are between 2 and 10 years old. 20% of removed addresses are more than 20 years old, where almost 70% of all advertised addresses are more than 20 years old. Addresses that Re-Home appear to be disproportionally represented in the age bracket of between 7 to 15 years old.

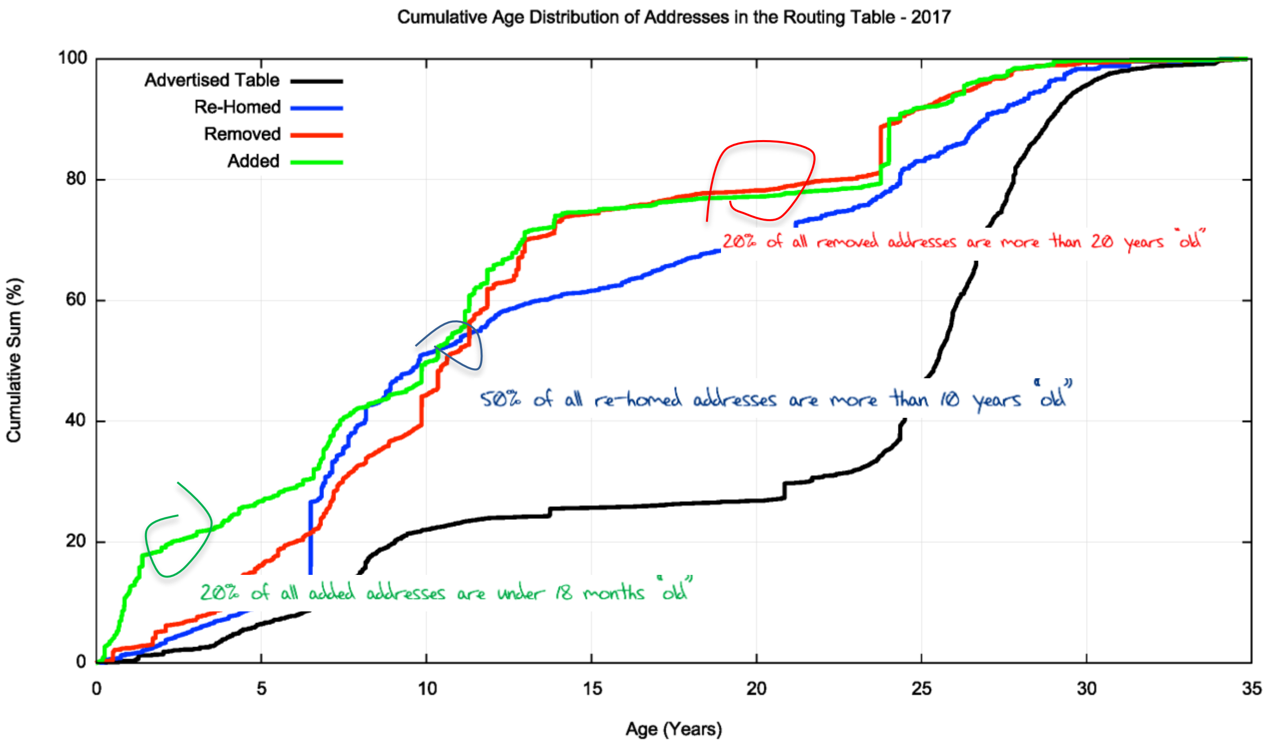

Figure 9 – Change in the size of the BGP routing table across 2016

However, as IPv4 moves into its final stages we are perhaps now in a position to take stok of the overall distribution of IPv4 addresses and look at where the addresses landed up. Table 11 shows the ten countries that have the largest pools of allocated IPv4 addresses.

Rank	CC	IPv4 Pool	% of Total	Addrs per capita	CC Name
1	US	1,613,129,216	37.56%	4.972	USA
2	CN	339,378,176	7.90%	0.241	China
3	JP	203,720,192	4.74%	1.598	Japan
4	GB	123,396,376	2.87%	1.865	UK
5	DE	121,341,312	2.83%	1.478	Germany
6	KR	112,435,456	2.62%	2.205	Korea
7	BR	84,118,528	1.96%	0.402	Brazil
8	FR	81,492,528	1.90%	1.254	France
9	CA	70,089,216	1.63%	1.914	Canada
10	IT	54,182,464	1.26%	0.913	Italy

Table 11 – IPv4 Allocated Address pools per National Economy

Slightly more than one third of all IPv4 addresses are allocated to entities that are registered as US entities. If we divide this address pool by the current population of each national entity, then we can generate an address per capita index. For the curious, the value of just under 5 addresses per capita for the United States is not the highest values. The numbers for the Seychelles and the Holy See are far higher! The global totals of 3.7 billion addresses with an estimated global population of 7.5 billion people gives an overall value of 0.49 addresses per capita. The full table of IPv4 allocations per national economy can be found at http://resources.potaroo.net/iso3166/v4cc.html.

IPv6 in 2017

Obviously, the story of IPv4 address allocations is only half of the story, and to complete the picture it's necessary to look at how IPv6 has fared over 2017.

IPv6 uses a somewhat different address allocation methodology than IPv4, and it is a matter of choice for a service provider as to how large an IPv6 address prefix is assigned to each customer. The original recommendations published by the IAB and IESG in 2001, documented in RFC3177, envisaged the general use of a /48 as an end site prefix. Subsequent consideration of long term address conservation saw a more flexible approach being taken with the choice of the end site prefix size being left to the service provider. Today's IPv6 environment has some providers using a /60 end site allocation unit, many use a /56, and other providers use a /48. This variation makes a comparison of the count of allocated IPv6 addresses somewhat misleading, as an ISP using /48's for end sites will require 256 times more address space to accommodate a similarly sized same customer base as a provider who uses a /56 end site prefix, and 4,096 times more address space than an ISP using a /60 end site allocation!

For IPv6 let's use both the number of discrete IPv6 allocations and the total amount of space that was allocated to see how IPv6 fared in 2016.

Comparing 2015 to 2016 the number of individual allocations of IPv6 address space has risen by some 20%. By contrast, the number of IPv4 allocations has fallen by 16% in this same period (Table 12).

Allocations	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
IPv6	473	841	1,243	2,477	3,700	3,403	3,840	4,407	4,733	5,594	5,765
IPv4	6,312	6,969	6,701	7,758	10,061	8,619	7,110	10,853	11,732	9,787	9,440

Table 12 - Number of individual Address Allocations, 2007 - 2017

The amount of IPv6 address space distributed in 2016 had risen by some 25% over 2015 levels, but in 2017 the total volume of allocated addresses fell by the same amount, back to the same total volume of addresses as in 2015. The number of allocations increased, however, indicating that in 2017 there were no anomalous extremely large allocations of IPv6 address space (Table 13).

Addresses	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
IPv6 (/32s)	6,916	15,634	1,555	4,754	20,009	18,136	23,935	17,513	20,225	25,301	19,986
IPv4 (/32s)(M)	203.9	203.3	189.4	248.8	201.0	114.9	65.1	63.9	34.8	22.2	51.9

Table 13 – Volume of Address Allocations, 2007 - 2017

Regionally, each of the RIRs saw IPv6 allocation activity in 2017 that was on a par with those seen in the previous year, with the exception of LACNIC, which saw a 50% increase in allocations and APNIC, which saw a 20% decline in allocations (Table 14).

Allocations	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
RIPE NCC	150	413	595	1,012	1,565	1,661	2,057	2,143	2,206	2,141	2,051
ARIN	196	213	357	567	959	545	523	505	602	646	684
APNIC	61	158	185	637	610	561	505	503	778	1,681	1,369
LACNIC	38	43	93	212	447	560	683	1,196	1,061	1,010	1,549
AFRINIC	18	14	13	49	119	76	72	60	86	116	112
	473	841	1,243	2,477	3,700	3,403	3,840	4,407	4,733	5,594	5,765

Table 14 - IPv6 allocations by RIR

The address assignment data tells a slightly different story. Table 15 shows the number of allocated IPv6 /32's per year. It appears that 2016 was an anomalous year for the RIPE NCC, in that the allocation totals for 2015 and 2017 are roughly the same. APNIC allocated a larger total in 2017, thanks to three large allocations: a /24 into Japan, a /22 into India and a /21 into China.

IPv6 (/32s)	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
RIPE NCC	1,468	964	1,052	2,406	3,174	3,892	6,286	8,217	12,031	21,707	12,844
ARIN	148	14,486	236	780	6,344	1,660	12,558	5,241	641	1,088	1,372
APNIC	5,236	139	170	1,335	9,486	3,783	4,442	2,644	2,109	1,236	4,228
LACNIC	51	35	87	197	948	4,605	597	1,359	974	1,182	1,429
AFRINIC	13	10	9	36	147	4,196	51	51	4,471	78	113
	6,916	15,634	1,555	4,754	20,099	18,136	23,935	17,513	20,225	25,301	19,986

Table 15 - IPv6 address allocation volumes by RIR

Dividing addresses by allocations gives the average IPv6 allocation size in each region (Table 16). APNIC average allocations increase in size due to the large allocations already noted. Overall, the average IPv6 allocation size remains a /30.

Average Allocation	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
RIPE NCC	/28.7	/30.8	/31.2	/30.8	/31.0	/30.8	/30.4	/30.1	/29.6	/28.7	/29.4
ARIN	/32.4	/25.9	/32.6	/31.5	/29.3	/30.4	/27.4	/28.6	/31.9	/31.2	/31.0
APNIC	/25.6	/32.2	/32.1	/30.9	/28.0	/29.2	/28.9	/29.6	/30.6	/32.4	/30.4
LACNIC	/31.6	/32.3	/32.1	/32.1	/30.9	/29.0	/32.2	/31.8	/32.1	/31.8	/32.1
AFRINIC	/32.5	/32.5	/32.5	/32.4	/31.7	/26.2	/32.5	/32.2	/26.3	/32.6	/32.0
All	/28.1	/27.8	/31.7	/31.1	/29.6	/29.6	/29.4	/30.0	/29.9	/29.8	/30.2

Table 16 – Average IPv6 address allocation size by RIR

The number and volume of IPv6 allocations per RIR per year is shown in Figures 10 and 11.

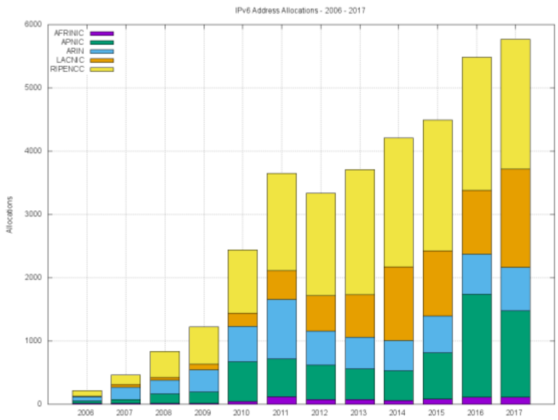

Figure 10 – Number of IPv6 Allocations per year

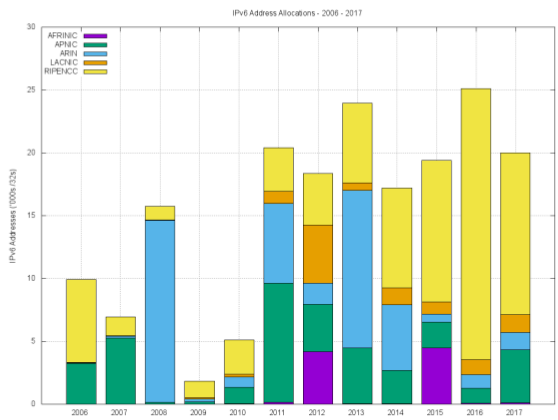

Figure 11 – Volume of IPv6 Allocations per year

Rank	2013		2014		2015		2016		2017	
1	USA	485	Brazil	946	Brazil	815	Brazil	774	Brazil	1,114
2	Brazil	473	USA	457	USA	540	USA	603	USA	635
3	UK	248	UK	239	China	267	China	509	Germany	270
4	Russia	246	Germany	215	Germany	249	Germany	266	Russia	220
5	Germany	195	Russia	201	UK	216	Australia	219	Australia	211
6	Netherlands	134	Netherlands	181	Russia	183	UK	211	China	208
7	France	132	France	122	Netherlands	170	Netherlands	198	Netherlands	194
8	Sweden	112	Switzerland	103	Australia	123	Russia	173	UK	190
9	Australia	102	Italy	103	Spain	119	India	161	Indonesia	187
10	Italy	98	Australia	101	France	116	Indonesia	159	Argentina	178

Table 17 - IPv6 allocations by Economy

Table 17 shows the countries who received the largest number of IPv6 allocations, while Table 18 shows the amount of IPv6 address space assigned on a per economy basis for the past 5 years (using units of /32s).

Rank	2013		2014		2015		2016		2017	
1	USA	12,537	USA	4,930	South Africa	4,440	UK	9,571	China	2,245
2	China	4,135	China	2,127	China	1,797	Germany	1,525	USA	1,498
3	UK	782	UK	1,090	UK	1,297	Netherlands	1,312	Germany	1,364
4	Germany	651	Brazil	863	Germany	1,269	USA	1,137	Russia	1,358
5	Russia	523	Germany	749	Netherlands	1,010	Russia	1,005	Netherlands	1,296
6	Netherlands	463	Netherlands	719	Russia	864	France	926	Spain	1,170
7	Brazil	450	Russia	716	Brazil	755	Brazil	727	India	1,087
8	France	435	France	436	Spain	708	Spain	702	UK	1,072
9	Italy	339	Italy	410	Italy	707	Italy	679	Brazil	1,049
10	Switzerland	265	Switzerland	369	USA	662	China	596	France	714

Table 18 - IPv6 Address Allocation Volumes by Economy (/32s)

Three of the countries in Table 17 listed as having received the highest volumes of allocated addresses in 2016, namely China, Russia and Spain have IPv6 deployments that are under 5% of their total user population. To what extent are allocated IPv6 addresses visible as advertised prefixes in the Internet's routing table?

Figure 12 shows the overall counts of advertised, unadvertised and total allocated address volume for IPv6 since mid 2009. Aside from the obvious discontinuity in early 2013, when a registration of a single /18 national address allocation for the Brazil National Registry of a /18 was replaced by the actual end user allocations, it's clear that the pool of unadvertised IPv6 addresses appears to the growing at a faster rate than the pool of advertised addresses in IPv6..

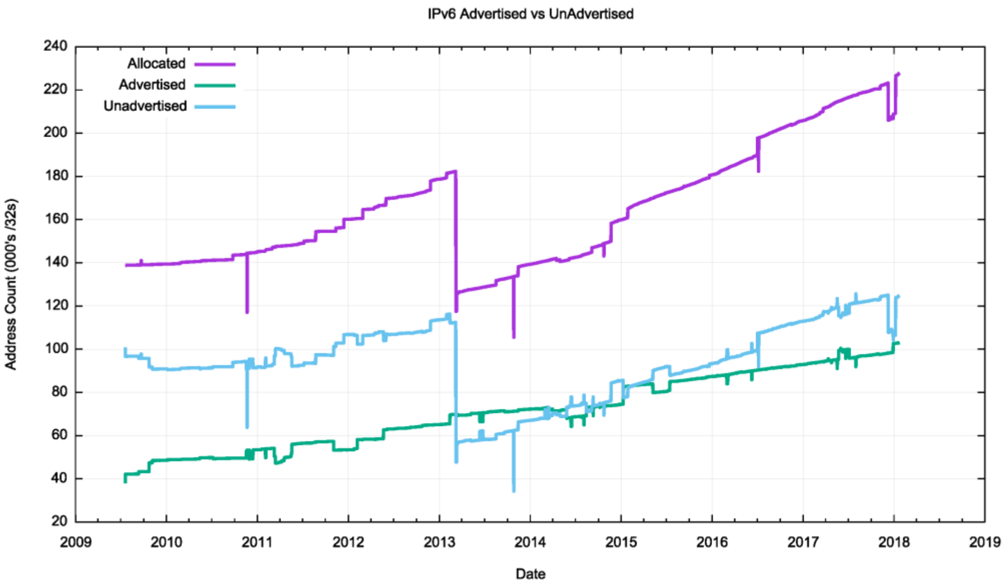

Figure 12 – Allocated, Unadvertised and Advertised IPv6 addresses

It is probably clearer to see the ratio of advertised to unadvertised addresses expressed as a percent, as shown in Figure 13. By the end of 2017 slightly less than half of the total pool of allocated IPv6 addresses was advertised in BGP.

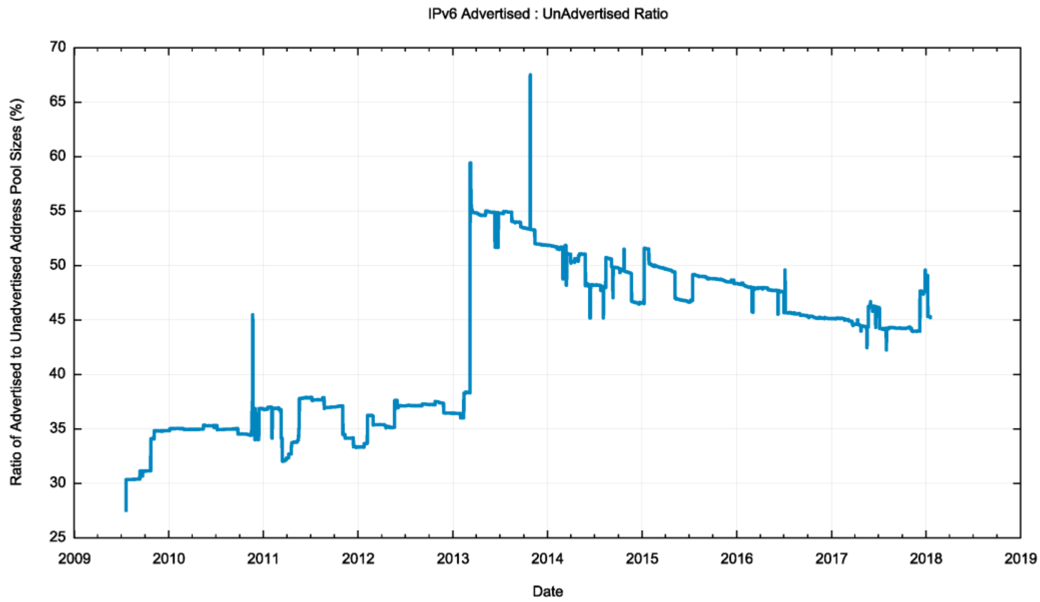

Figure 13 –Advertised IPv6 Addresses as a percentage of the Allocated Address Pool

Where is this ending up? We can take a comparable look at the allocated address pools for the top ten national economies in IPv6, and the current picture is shown in Table 19

Rank	CC	Allocated (/48s)	%_of_Allocated	/48s per capita	IPv6 Use	Country Name
1	US	2,920,401,219	10%	9.0	38%	USA
2	CN	1,400,111,130	5%	1.0	0%	China
3	GB	1,179,713,714	4%	17.8	26%	UK
4	DE	1,144,717,750	4%	13.9	43%	Germany
5	FR	793,182,300	3%	12.2	21%	France
6	JP	636,231,848	2%	5.0	24%	Japan
7	AU	587,006,962	2%	24.0	15%	Australia
8	IT	497,942,546	2%	8.4	2%	Italy
9	NL	402,194,722	1%	23.6	12%	Netherlands
10	SE	391,774,513	1%	39.5	5%	Sweden

Table 19 – IPv6 Allocated Address pools per National Economy

While the United States also tops this list in terms of the total pool of allocated IPv6 addresses, the per capita number is lower than many others in this list. Sweden has a surprisingly high number yet estimates of the population of IPv6-capable users in that country point to a deployment rate of just 5%, considerably lower than many other countries listed here. But for IPv6 its still relatively early days and no doubt the picture will change as deployment of IPv6 matures.

The Outlook for the Internet

Once more the set of uncertainties that surround the immediate future of the Internet are considerably greater than the set of predictions that we can be reasonably certain about.

There has been much in the way of progress in the transition to IPv6 in 2017, but that does not necessarily mean that other providers will quickly follow this lead. While a number of service operators have reached the decision point that the anticipated future costs of NAT deployment are unsustainable for their service platform, there remains a considerable school of thought that says that NATs will cost effectively absorb some further years of Internet device population growth. At least that's the only rationale I can ascribe to a very large number of service providers who are making no visible moves to push out Dual-Stack services at this point in time. Given that the objective of this transition is not to turn on Dual-Stack everywhere, but to turn off IPv4, there is still some time to go, and the uncertainty lies in trying to quantify what that time might be.

The period of the past few years has been dominated by the mass marketing of mobile internet services, and the growth rates for 2014 through to 2016 perhaps might have been the highest so far recorded were it not for the exhaustion of the IPv4 address pool. In address terms this growth in the IPv4 Internet is being almost completely masked by the use of Carrier Grade NATs in the mobile service provider environment, so that the resultant demands for public addresses in IPv4 are quite low and the real underlying growth rates in the network are occluded by these NATs.

In theory, there is no strict requirement for IPv6 to use NATs, and if the mobile world were deploying dual stack ubiquitously then this would be evident in the IPv6 address allocation data. And we see this in India, where the rollout of the Jio mobile service through 2016 and into 2017 has now encompassed some 90% of their considerable user population. On the other hand, the other massive user population, that of China, still shows no visible signs of deploying IPv6 as yet.

We should also be seeing IPv6 address demands for deployments of large scale sensor networks and other forms of deployments that are encompassed under the broad umbrella of the Internet of Things. This does not necessarily imply that the deployment is merely a product of an over-hyped industry, although that is always a possibility. It is more likely to assume that such deployments take place using private IPv4 (or IPv6 ULA addresses) addresses, and once more rely on NATs or application level gateways to interface to the public network. Time and time again we are lectured that NATs are not a good security device, but in practice NATs offer a reasonable front-line defence against network scanning malware, so there may be a larger story behind the use of NATs and device based networks than just a simple conservative preference to continue to use an IPv4 protocol stack.

We are witnessing an industry that is no longer using technical innovation, openness and diversification as its primary means of propulsion. The widespread use of NATs in IPv4 limit the technical substrate of the Internet to a very restricted model of simple client/server interactions using TCP and UDP. The use of NATs force the interactions into client-initiated transactions, and the model of an open network with considerable flexibility in the way in which communications take place is no longer being sustained in today's network. Incumbents are entrenching their position and innovation and entrepreneurialism are taking a back seat while we sit out this protracted IPv4/IPv6 transition.

What is happening is that today's internet carriage service is provided by a smaller number of very large players, each of whom appear to be assuming a very strong position within their respective markets. The drivers for such larger players tend towards risk aversion, conservatism and increased levels of control across their scope of operation. The same trends of market aggregation are now appearing in content provision, where a small number of content providers are exerting a completely dominant position across the entire Internet.

The evolving makeup of the Internet industry has quite profound implications in terms of network neutrality, the separation of functions of carriage and service provision, investment profiles and expectations of risk and returns on infrastructure investments, and on the openness of the Internet itself. The focus now is turning to the regulatory agenda. Given the economies of volume in this industry, it was always going to be challenging to sustain an efficient, fully open and competitive industry, but the degree of challenge in this agenda is multiplied many-fold when the underlying platform has run out of the basic currency of IP addresses. The pressures on the larger players within these markets to leverage their incumbency into overarching control gains traction when the stream of new entrants with competitive offerings dries up, and the solutions in such scenarios typically involve some form of public sector intervention directed to restore effective competition and revive the impetus for more efficient and effective offerings in the market.

As the Internet continues to evolve, it is no longer the technically innovative challenger pitted against venerable incumbents in the forms of the traditional industries of telephony, print newspapers, television entertainment and social interaction. The Internet is now the established norm. The days when the Internet was touted as a poster child of disruption in a deregulated space are long since over, and these days we appear to be increasingly looking further afield for a regulatory and governance framework that can continue to challenge the increasing complacency of the newly-established incumbents.

It is unclear how successful we will be in this search. We can but wait and see.

Disclaimer

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

Author

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

www.potaroo.net