

January 2020
Geoff Huston

Addressing 2019

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.

Back around 1992 the IETF gazed into their 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. Doubtless, these device numbers will continue to grow. We continue to increase silicon production volumes and at the same time continue to refine the production process. 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, as 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 whatever volume they might be deployed.

But while the Internet has grown at such amazing rates, the deployment of IPv6 continues at a more measured 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 at this moment. Much of the reason for this apparent contradiction between the designed-size of the IPv4 network and the actual size, which is of course many times larger, 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 20 billion connected devices can be squeezed into some 2 billion active IPv4 addresses. Applications that cannot work behind NATs are no longer useful and no longer used.

However, the pressures of this inexorable growth in the number of deployed devices in the Internet means that the even NATs cannot absorb these growth pressures forever. NATs extend the effective addressable space by up to 16 bits, and permits the time-based sharing of addresses. Both of these are effective, but that do not transform the address space into an infinitely elastic resource. Inevitably, either we will see the fragmenting of the IPv4 Internet into a number of disconnected parts, so that the entire concept of a globally unique and coherent address pool layered over a single coherent packet transmission realm 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, but my (admittedly cynical and perhaps jaded) personal view lies in a future of fragmentation of the network.

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 2019.

IPv4 in 2019

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, although by the end of 2019 the end was in sight.

The allocation of 11.9 million addresses in 2019 on top of a base of 3.67 billion addresses that are already allocated at the start of the year represents a growth rate of 0.5% for the year for the total allocated IPv4 public address pool. This is less than one tenth of the growth rate in 2010 (the last full year before the onset of IPv4 address exhaustion).

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
Allocated (M)	189.4	248.8	201.0	114.9	65.1	63.9	34.8	22.2	15.6	14.5	11.9
Total (B)	2.72	2.90	3.14	3.34	3.43	3.50	3.59	3.62	3.65	3.67	3.68
Relative Growth	7.9%	6.6%	8.3%	6.4%	2.7%	2.0%	2.6%	0.8%	0.8%	0.5%	0.3%

Table 1 - IPv4 Allocated addresses by year

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

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
APNIC	86.9	120.2	105.2	1.0	1.3	3.7	4.1	3.8	1.8	1.8	2.9
RIPE NCC	43.4	56.0	43.1	40.0	2.0	2.5	3.3	3.4	3.6	4.7	6.4
ARIN	41.1	45.2	23.5	45.0	26.5	26.0	8.6	1.6	0.8	0.6	-2.7
LACNIC	10.5	13.0	24.4	21.0	28.5	19.1	1.8	1.6	1.7	1.3	1.2
AFRINIC	5.9	8.5	9.2	7.9	6.8	12.5	16.9	11.8	7.8	6.1	4.1

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

The term “allocations” is used in a somewhat imprecise manner here. What these tables show is the net change in the total assigned address pool for each RIR on a year-by-year basis. This encompasses allocations, returns, reservation actions and inter-RIR transfers.

In terms of the IPv4 Internet there is a considerable diversity in the situation in each region. As of the end of 2019, AFRINIC was the last remaining Regional Internet Registry (RIR) with remaining IPv4 addresses available for general allocation, with some 2.6 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 2.9 million addresses left in this pool, and the RIPE NCC exhausted this pool in the latter part of 2019, and at the start of 2020 recorded a block of 512 addresses as “available”. LACNIC has a residual pool of 48,640 available addresses, while ARIN had 4,096 available addresses as of January 1, 2020.

We can use the address allocation data from 2019 and perform a forward extrapolation on this to predict when the available address pools of the two RIR’s with remaining addresses will be exhausted. This linear projection model is shown in Figure 1.

As shown in Figure 1, the address consumption rate for APNIC has been relatively steady across the period 2017 through 2019, and at this consumption rate the pool will last for a further 19 months, in July 2021. AFRINIC will likely exhaust its remaining addresses in mid-2020, at current consumption rates. The overall picture of exhaustion dates for each RIR is shown in Table 3.

ARIN	exhausted
LACNIC	exhausted
RIPE NCC	exhausted
AFRINIC	mid-2020
APNIC	late-2021

Table 3 - IPv4 Address Pool Depletion Projections by RIR




Figure 1 – IPv4 RIR pool runout scenarios

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. As of the start of 2020, ARIN has 5.9 million reserved IPv4 addresses, APNIC has 4.4 million, AFRNIC 2.0 million, the RIPE NCC has 780 thousand, and LACNIC 1.2 million. The total pool of reserved IPv4 addresses is some 14.2 million addresses in size, or the equivalent of slightly less than one /8 address block. Finally, the IANA is holding 3,3288 addresses in its recovered address pool in 13 discrete /24 address blocks.

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 was 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 six years is shown in Table 4. This includes both inter-RIR and intra-RIR transfers. Each transfer is treated as a single transaction, and in the case of inter-RIR transfers, this is accounted in the receiving RIR’s totals.

Receiving RIR	2012	2013	2014	2015	2016	2017	2018	2019
APNIC	87	102	175	198	267	280	332	671
RIPE NCC	9	154	919	2,854	2,902	3,004	4,231	4,212
ARIN	30	37	64	310	867	1,283	1,253	1,267
LACNIC					5	13	26	33
AFRINIC							7	7
Total	126	293	1,158	3,362	4,041	4,580	5,849	6,190

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	2018	2019
APNIC	1,822,208	1,962,240	3,863,552	6,300,928	7,843,328	4,377,344	10,082,304	6,494,720
RIPE NCC	65,536	1,977,344	9,644,544	10,557,696	9,453,568	22,962,432	14,431,104	9,47,328
ARIN	4,862,208	5,128,448	4,744,960	38,895,104	17,089,536	37,934,336	37,212,672	28,614,400
LACNIC					59,392	99,328	83,968	61,696
AFRINIC							14,336	64,000
Total	6,749,952	9,068,032	18,253,056	55,753,728	34,445,824	65,373,440	61,824,384	44,982,144

Table 5 – Volume of Transferred IPv4 Addresses per year

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




Figure 3 – Number of Transfers: 2012 - 2019

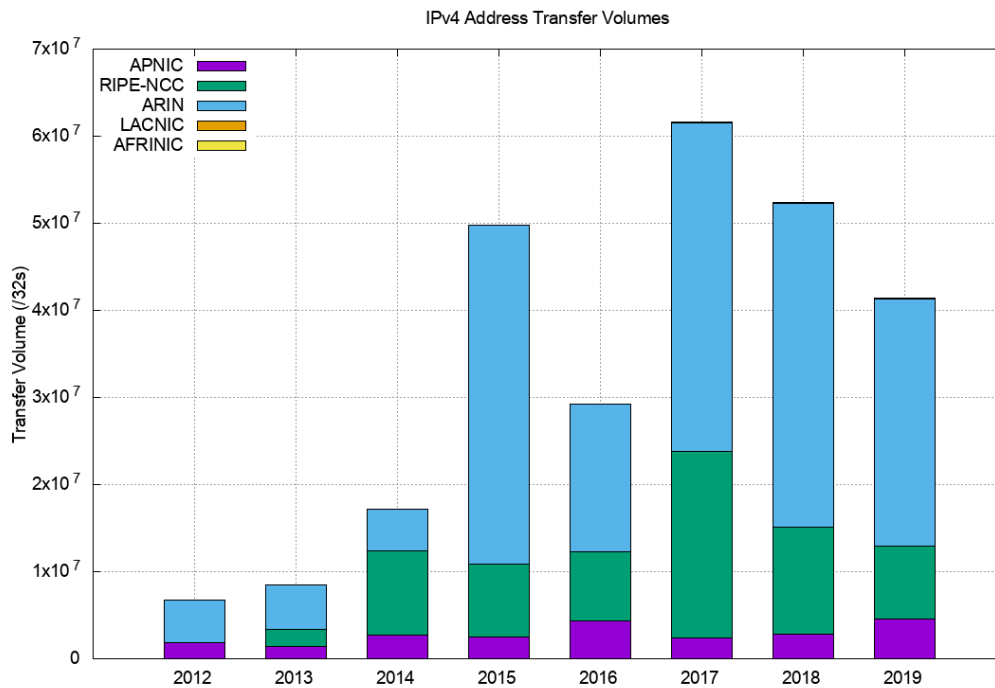


Figure 4 – Volume of Transferred Addresses: 2012 - 2019

The total volume of addresses transferred in this way is four times the volume of allocated addresses across 2019. The aggregate total of addresses in the transfer logs since 2012 is some 300 million addresses, or the equivalent of 17 /8s, which is some 8% of the total delegated IPv4 address space of 3.7 billion addresses.

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. In other words, the opening of a market in addresses would provide incentive for otherwise unproductive address assets to be placed on the market. Providers who had a need for addresses would compete with other providers who had a similar need in bidding to purchase these addresses. In conventional market theory the most efficient user of addresses (here “most efficient” is based on the ability to use addresses to generate the greatest revenue) would be able to set the market price. Otherwise unused addresses would be put to productive use, and as long as demand outstrips supply the most efficient use of addresses is promoted by the actions of the market. In theory.

However, the practical experience with transfer to date is not so clear. The data relating to address re-use is inconclusive, in that between 2011 and late 2017 the pool of unadvertised addresses sat between some 43 and 44 /8s. This pool of unadvertised addresses rose from the start of 2018 and by the end of 2019 there were some 50 /8s that were unadvertised to the public Internet. Rather than reducing the pool of unused (unadvertised) addresses, this pool appears to be expanding in size over the past 24 months.

In relative terms, expressed as a proportion of the total pool of allocated IP addresses, the unadvertised address pool dropped from 28% of the total allocated address pool in 2011 to a low of some 24% at the start of 2016, and subsequently risen to 29% by the end of 2019. This data points to the conclusion that address transfer activity has not made a substantial change in the overall picture of address utilisation efficiency in the past 12 months.




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, or 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 some 20% of the transferred address blocks were originally assigned or allocated by an RIR within

the previous 5 years. In 2019 the trend had reversed, and some 95% of all transferred addresses were more than 5 years old, and 70% of the transferred addresses were more than 25 years old.

By volume, the overall majority of transferred addresses were drawn from legacy address holder, which is similar to the picture in previous years.




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). In 2019 it appears that transfer transactions are uniformly distributed according to age of the addresses. Previous years have seen a high volume of transactions in recently allocated addresses, but this trading of recently allocated addresses has tapered off in the past three years. The comparison of Figures 7 and 8 also leads to the observation that the transfers of older address blocks

have far larger address spans, which corresponds to the very early IPv4 address allocations of /8 and /16 prefixes.

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 25,520 transactions described in the RIRs’ transfer registries from the start of 2012 until the start of 2020, and of these 6,286 entries list transferred address blocks that are the same size as the original allocated block. The remaining 19,234 entries are fragments of the originally allocated address blocks.

These 19,234 transfer entries that have fragmented the original allocation are drawn from 2,104 such original allocations. On average the original allocation is split into 9 smaller address blocks. 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 209,199 distinct address allocations from the RIRs to end entities as of the end of 2019, and the fragmentation reflected in 19,234 of these address blocks is around 9% 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). We will use the RIR-published transfer data for 2019.

Rank	CC	Addresses (Source)	Country
1	US	29,126,400	United States of America
2	CN	2,263,808	China
3	SG	1,633,024	Singapore
4	IL	1,378,816	Israel
5	DE	1,229,312	Germany
6	JP	1,116,672	Japan
7	IR	1,112,832	Iran
8	RU	1,092,608	Russian Federation
9	CA	713,984	Canada
10	SA	703,744	Saudi Arabia

Table 6 – Top 10 Countries Sourcing Transferred IPv4 addresses in 2019

Rank	CC	Addresses (Recipient)	Country
1	US	31,458,304	United States of America
2	CN	2,761,216	China
3	IL	1,373,952	Israel
4	RU	1,286,656	Russian Federation
5	IR	1,224,192	Iran
6	CA	727,296	Canada
7	NL	645,376	Netherlands
8	JP	593,408	Japan
9	GB	583,680	United Kingdom
10	DE	495,360	Germany

Table 7 – Top 10 Countries Receiving Transferred IPv4 addresses in 2019

Rank	From	To	Addresses	From	To
1	US	US	28,781,568	United States of America	United States of America
2	CN	CN	2,259,712	China	China
3	IL	IL	1,360,896	Israel	Israel
4	US	SG	1,097,728	United States of America	Singapore
5	RU	RU	1,046,528	Russian Federation	Russian Federation
6	IR	IR	1,022,720	Iran	Iran
7	JP	JP	589,312	Japan	Japan
8	CA	CA	455,680	Canada	Canada
9	CN	SG	452,608	China	Singapore
10	PH	PH	409,600	Philippines	Philippines
11	ES	ES	408,320	Spain	Spain
12	US	SA	376,832	United States of America	Saudi Arabia
13	HU	HU	345,856	Hungary	Hungary
14	NL	NL	302,080	Netherlands	Netherlands
15	DE	DE	271,360	Germany	Germany
16	US	DE	271,360	United States of America	Germany
17	GB	GB	250,368	United Kingdom	United Kingdom
18	AU	AU	250,112	Australia	Australia
19	NL	DE	241,920	Netherlands	Germany
20	US	JP	201,472	United States of America	Japan

Table 8 – Top 20 Country-to-Country IPv4 address transfers in 2019

The 2019 transfer logs contain 4,414 domestic address transfers, with a total of 39,377,792 addresses (of which 1,056 are US domestic transfers that encompass 28,781,568 addresses) while 1,417 transfers appear to result in a movement of addresses between countries, involving a total of 6,707,968 addresses.

It appears that the IPv4 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 name-based 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. The result is that the pressure of the IPv4 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 over this period.

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 2019 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-19	Jan-20	Delta	Unchanged	Re-Home	Removed	Added
Announcements	760,025	814,253	54,228	651,625	27,101	81,299	135,527
Root Prefixes:	352,449	387,073	34,624	321,753	13,732	25,049	51,588
Address Span (/8s)	169.27	169.73	0.46	160.11	2.06	8.84	7.56
More Specifics:	407,576	427,180	19,604	329,872	13,369	56,250	83,939
Address Span (/8s)	64.27	64.02	-0.25	56.08	1.32	5.60	6.68

Table 9 – IPv4 BGP changes over 2019

While the routing table grew by 54,228 entries over the year, the nature of the change is slightly more involved. Some 81,299 prefixes that were announced at the start of the year were removed from the routing system through the year, and 135,527 prefixes were announced by the end of the year that were not announced at the start of the year. (Within the scope of this study I have not tracked the progress of announcements through the year, and it is likely that many more prefixes were announced and removed on a transient basis through the course of the year.) A further 27,101 prefixes had changed their originating Autonomous System number, indicating some form of change in the prefix’s network location in some manner.

We can compare these changed prefixes against the transfer logs for the two-year period 2018 and 2019. Table 10 shows the comparison of these routing numbers against the set of transfers that were logged in these two years.

Type	Listed	Unlisted	Ratio
Re-Homed			
All	1,651	25,450	6.09%
Root Prefixes	1387	12027	10.34%
Removed			
All	4,135	77,164	5.09%
Root Prefixes	2,465	22,584	9.84%
Added			
All	8,739	126,788	6.45%
Root Prefixes	5,987	45,601	11.61%

Table 10 – Routing changes across 2018 compared to the Transfer Logs

These figures show that some 5%-12% of changes in advertised addresses are reflected as changes as recorded in the RIRs’ transfer logs. This should not imply that the remaining 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 2019, 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 2019, 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 2019, 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 30 months. 20% of all added and removed addresses are more than 20 years old. Addresses that re-home appear to be disproportionately represented in the age bracket of between 5 to 10 years old.




Figure 9 – Change in the size of the BGP routing table across 2019

However, as IPv4 moves into its final stages we are perhaps now in a position to take stock 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	Per-Capita	Economy
1	US	1,611,080,448	43.76%	4.90	United States of America
2	CN	340,246,784	9.24%	0.24	China
3	JP	204,528,640	5.56%	1.61	Japan
4	DE	124,218,752	3.37%	1.51	Germany
5	GB	115,310,872	3.13%	1.72	United Kingdom
6	KR	112,476,160	3.06%	2.19	Korea
7	BR	86,113,536	2.34%	0.41	Brazil
8	FR	83,199,504	2.26%	1.27	France
9	CA	69,951,744	1.90%	1.88	Canada
10	IT	55,135,296	1.50%	0.93	Italy

Table 11 – IPv4 Allocated Address Pools per National Economy

If we divide this address pool by the current population of each national entity, then we can derive an address per capita index. The global total of 3.70 billion allocated addresses with an estimated global population of 7.7 billion people gives an overall value of 0.37 IPv4 addresses per capita. It appears that early adopter communities tend to have high than average per capita values, while later adopters tend to fall below the global average. The full table of IPv4 allocations per national economy can be found at <http://resources.potaroo.net/iso3166/v4cc.html>.

IPv6 in 2018

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 2019.

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 2019.

Comparing 2018 to 2019 the number of individual allocations of IPv6 address space has risen, as has IPv4 activity (Table 12).

Allocations	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
IPv6	1,243	2,477	3,700	3,403	3,840	4,407	4,733	5,594	5,765	6,309	6,927
IPv4	6,701	7,758	10,061	8,619	7,110	10,853	11,732	9,787	9,440	10,199	14,033

Table 12 - Number of individual Address Allocations, 2007 - 2018

The amount of IPv6 address space distributed in 2019 is slightly less than the amount seen in 2018 (Table 13).

Addresses	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
IPv6 (/32s)	1,555	4,754	20,009	18,136	23,935	17,513	20,225	25,301	19,986	39,256	36,084
IPv4 (/32s)(M)	189.4	248.8	201	114.9	65.1	63.9	34.8	22.2	15.6	14.5	11.9

Table 13 - Volume of Address Allocations, 2007 - 2018

Regionally, each of the RIRs saw IPv6 allocation activity in 2019 that was on a par with those seen in the previous year, with the exception of the RIPE NCC, which saw a 18% increase in allocations and LACNIC with a 11% increase in allocation activity. (Table 14).

Allocations	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
RIPE NCC	595	1,012	1,565	1,661	2,057	2,143	2,206	2,141	2,051	2,617	3,105
ARIN	357	567	959	545	523	505	602	646	684	658	605
APNIC	185	637	610	561	505	503	778	1,681	1,369	1,474	1,484
LACNIC	93	212	447	560	683	1,196	1,061	1,010	1,549	1,450	1,618
AFRINIC	13	49	119	76	72	60	86	116	112	110	115
	1,243	2,477	3,700	3,403	3,840	4,407	4,733	5,594	5,765	6,309	6,927

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. There were two /20 allocations in the year, one to CERNET in China and a second to Cable One in the United States. A second large allocation, a /21, was also made to CERNET in China in 2019.

IPv6 (/32s)	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
RIPE NCC	1,052	2,406	3,174	3,892	6,286	8,217	12,031	21,707	12,844	17,311	21,188
ARIN	236	780	6,344	1,660	12,558	5,241	641	1,088	1,372	845	5,522
APNIC	170	1,335	9,486	3,783	4,442	2,644	2,109	1,236	4,228	19,690	7,495
LACNIC	87	197	948	4,605	597	1,359	974	1,182	1,429	1,338	1,515
AFRINIC	9	36	147	4,196	51	51	4,471	78	113	72	364
	1,555	4,754	20,099	18,136	23,935	17,513	20,225	25,301	19,986	39,256	36,084

Table 15 - IPv6 address allocation volumes by RIR

Dividing addresses by allocations gives the average IPv6 allocation size in each region (Table 16). ARIN average allocations increased in size due to some larger allocations in 2019. Overall, the average IPv6 allocation size remains around a /30.

Average	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019
RIPE NCC	/31.2	/30.8	/31.0	/30.8	/30.4	/30.1	/29.6	/28.7	/29.4	/29.3	/29.3
ARIN	/32.6	/31.5	/29.3	/30.4	/27.4	/28.6	/31.9	/31.2	/31.0	/31.6	/28.8
APNIC	/32.1	/30.9	/28.0	/29.2	/28.9	/29.6	/30.6	/32.4	/30.4	/28.2	/29.7
LACNIC	/32.1	/32.1	/30.9	/29.0	/32.2	/31.8	/32.1	/31.8	/32.1	/32.1	/32.1
AFRINIC	/32.5	/32.4	/31.7	/26.2	/32.5	/32.2	/26.3	/32.6	/32.0	/32.6	/30.3
	/31.7	/31.1	/29.6	/29.6	/29.4	/30.0	/29.9	/29.8	/30.2	/29.3	/29.6

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	2015		2016		2017		2018		2019	
1	Brazil	815	Brazil	774	Brazil	1,114	Brazil	1,049	Brazil	1,112
2	USA	540	USA	603	USA	635	Russia	635	USA	541
3	China	267	China	509	Germany	270	USA	594	Russia	502
4	Germany	249	Germany	266	Russia	220	Germany	308	Germany	407
5	UK	216	Australia	219	Australia	211	China	253	Indonesia	366
6	Russia	183	UK	211	China	208	Indonesia	213	Netherlands	342
7	Netherlands	170	Netherlands	198	Netherlands	194	UK	184	UK	223
8	Australia	123	Russia	173	UK	190	Bangladesh	183	Bangladesh	202
9	Spain	119	India	161	Indonesia	187	India	168	France	179
10	France	116	Indonesia	159	Argentina	178	Netherlands	162	China	165

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	2015		2016		2017		2018		2019	
1	South Africa	4,440	UK	9,571	China	2,245	China	17,647	China	6,787
2	China	1,797	Germany	1,525	USA	1,498	Russia	4,651	USA	5,111
3	UK	1,297	Netherlands	1,312	Germany	1,364	Germany	1,932	Russia	3,716
4	Germany	1,269	USA	1,137	Russia	1,358	UK	1,209	Germany	2,522
5	Netherlands	1,010	Russia	1,005	Netherlands	1,296	Singapore	1,055	Netherlands	2,516
6	Russia	864	France	926	Spain	1,170	Netherlands	1,025	UK	1,355
7	Brazil	755	Brazil	727	India	1,087	Brazil	1,007	France	1,182
8	Spain	708	Spain	702	UK	1,072	USA	874	Italy	1,052
9	Italy	707	Italy	679	Brazil	1,049	Spain	851	Brazil	1,049
10	USA	662	China	596	France	714	France	722	Singapore	854

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

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 2010, while Figure 13 shows the advertised address span as a percentage of the total span of allocated and assigned IPv6 addresses.




Figure 12 – Allocated, Unadvertised and Advertised IPv6 addresses




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

The drop in the allocated address span in 2013 is the result of a change in LACNIC where a single large allocation into Brazil was replaced by the recording of direct allocation and assignments to ISPs and similar end entities.

From a history of careful conservation of IPv4 addresses, where some 77% of allocated or assigned IPv4 addresses are advertised in the BGP routing table, a comparable IPv6 figure of 45% does not look all that impressive. But that's not the point. We chose the 128-bit address size in IPv6 to allow addresses to be used without overriding concerns about conservation. We are allowed to be inefficient in address utilisation. At the start of 2020 we have advertised an IPv6 address span which is the equivalent of some 140,000 /32s, or some 9 billion end-site /48 prefixes. That is just 0.003% of the total number of /48 prefixes in IPv6.

We can also look at the allocated address pools for the top twenty five national economies in IPv6, and the current picture is shown in Table 19.

Rank	CC	Allocated /48s	% Total	/48s p.c.	Advertised /48s	Deployment	Name
1	US	3,329,586,952	9%	10.1	773,656,081	55.4%	USA
2	CN	3,135,766,596	8%	2.2	352,841,264	16.2%	China
3	DE	1,473,446,444	4%	17.9	540,410,664	46.8%	Germany
4	GB	1,323,434,197	3%	19.8	135,659,537	33.3%	Ukraine
5	FR	921,895,024	2%	14.1	1,175,557	38.4%	France
6	RU	790,298,901	2%	5.5	65,555	7.7%	Russia
7	JP	658,120,883	2%	5.2	427,466,771	36.1%	Japan
8	NL	615,448,886	2%	35.9	4,259,858	24.8%	Netherlands
9	AU	611,320,934	2%	24.4	274,129,217	25.4%	Australia
10	IT	610,668,567	2%	10.3	786,433	4.0%	Italy
11	BR	430,661,203	1%	2.0	68,323,447	31.4%	Brazil
12	SE	422,904,132	1%	42.1	269,025,292	9.0%	Sweden
13	ES	363,855,899	1%	7.8	196,610	3.8%	Spain
14	PL	352,780,492	1%	9.3	598,025	12.4%	Poland
15	KR	344,653,833	1%	6.7	971,306	14.4%	Korea

16	AR	340,920,912	1%	7.6	273,219,645	9.1%	Argentina
17	ZA	311,824,529	1%	5.4	197,632	0.3%	South Africa
18	EG	269,090,818	1%	2.7	131,072	11.6%	Egypt
19	CH	212,271,251	1%	24.7	4,259,849	38.6%	Switzerland
20	TR	170,983,447	0%	2.1	65,536	0.0%	Turkey
21	CZ	169,541,740	0%	15.9	393,224	12.8%	Czechia
22	TW	166,395,918	0%	7.0	16,081,382	43.5%	Taiwan
23	UA	165,413,014	0%	3.8	16,396	0.2%	Ukraine
24	IR	156,041,222	0%	1.9	65,538	0.0%	Iran
25	NO	141,754,711	0%	26.2	17,432,580	25.0%	Norway

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 9%, 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.

The year 2017 saw a sharp rise in IPv6 deployment, influenced to a major extent by the deployment of IPv6 services in India, notably by the Jio service. 2018 was a quieter year, although the rise in the second half of the year is due to the initial stirrings of mass scale IPv6 deployment in the major Chinese service providers. This movement accelerated in 2019 and the overall move of some 5% in the network-wide IPv6 deployment levels has a lot to do with the very rapid rise of the deployment of IPv6 in China (At the end of 2019 this IPv6 deployment level was approaching 20% of the entire China user population, which is a massive achievement considering the size of that country’s user population.)




Figure 14 – IPv6 Deployment measurement 2010 - 2020

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 ultimate 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 that particular service provider's considerable user population. The deployment in India has been accompanied by a conservative use of IPv6 addresses. The 112 million/48s allocated to date to India equates to 0.1 /48's per capita, a number far lower than any other economy with significant IPv6 deployment.

On the other hand, the other massive user population, that of China, has only stirred with visible deployment of IPv6 in the latter part of 2018.

It has also been assumed that we should see 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 so far such deployments are taking place using private IPv4 addresses, and they rely on NATs and 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. 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 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