

January 2022
Geoff Huston

IP Addressing through 2021

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 still no common sense of urgency about the deployment of this protocol, and still there is no common agreement that the continued reliance on IPv4 is failing us. Much of the reason for this apparent contradiction between the designed population of the IPv4 Internet and the actual device count, 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 can extend the effective addressable space by up to 32 'extra' bits, and they enable the time-based sharing of addresses. Both of these are effective measures in stretching the address space to encompass a larger device pool, but they do not transform the address space into an infinitely elastic resource. The inevitable outcome of this process is that either we will see the fragmenting of the IPv4 Internet into a number of disconnected parts, probably based on the service 'cones' of the various points of presence of the content distribution servers, 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 opinion lies in a future of highly fragmented 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 2021.

IPv4 in 2021

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 2021 the end of the old registry allocation model was in sight with the depletion of the residual pools of unallocated addresses in each of the Regional Internet Registries (RIRs).

It is increasingly difficult to talk about “allocations” in today's Internet. There are still a set of transactions where addresses are drawn from the residual pools of RIR-managed available address space and allocated or assigned to network operators, but at the same time there are also a set of transactions where addresses are traded between network in what is essentially a sale. These address transfers necessarily entail a change of registration details, so the registry records the outcome of a transfer, or sale, in a manner that is similar to an allocation or assignment.

If we want to look at the larger picture of the amount of IPv4 address space that is used or usable by Internet network operators, then perhaps the best metric to use is the total span of allocated and assigned addresses, and the consequent indication of annual change in the change in this total address span.,

What is the difference between "allocated" and "assigned"?

When a network operator or sub-registry has received an *allocation* it can further delegate that IP address space to their customers along with using it for their own internal infrastructure. When a network operator has received an *assignment* this can only be used for their own internal infrastructure.
[<https://www.apnic.net/get-ip/faqs/using-address-space/>]

I personally find the distinction between these two terms somewhat of an artifice these days, so from here on I'll use the term “allocation” to describe both allocations and assignments.

The total IPv4 allocated address pool expanded by some 1.1 million addresses in 2021 on top of a base of 3.682 billion addresses that were already allocated at the start of the year. This represents a growth rate of 0.1% 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) (Table 1).

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Address Span (Billions)	2.985	3.227	3.395	3.483	3.537	3.593	3.624	3.643	3.657	3.657	3.682	3.684	3.685
Annual Change (Millions)	189.4	241.7	168.0	88.4	53.9	55.9	30.6	19.4	13.2	0.6	24.9	2.2	1.1
Relative Growth	7.9%	8.1%	5.2%	2.6%	1.5%	1.6%	0.9%	0.5%	0.4%	0.0%	0.7%	0.1%	0.0%

Table 1 - IPv4 Allocated addresses by year

Where is this supposedly “new” address space coming from? The old model was that unallocated addresses were held in a single pool by the IANA, and blocks of addresses were passed to RIRs who then allocated them to various end entities, either for their own use or for further allocation. But, the IANA exhausted the last of its available address pools some years ago, and these days it holds just 3 /24 address prefixes (<https://www.iana.org/assignments/ipv4-recovered-address-space/ipv4-recovered-address-space.xhtml>), and as the option of dividing this pool into 5 equal chunks of 153.6 individual address is not

viable, then these addresses are likely to sit in the IANA Recovered Address registry for some time (i.e. until one of more of the RIRs return more prefixes recovered from the old “legacy” allocated addresses to the IANA, who would then be able to divide the pool equally and distribute them to each the 5 RIRs. This is unlikely to occur.) There are also addresses that have been marked by the IANA as *reserved* (<https://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>), including blocks of addresses reserved for Multicast use, and the top end of the IPv4 address space, curiously marked as reserved for *future use*. This latter category is a relatively large pool of 268,435,456 addresses (old former “Class E” space) and if ever there was a “future” for IPv4 then it is now. But exactly how to unlock this space and return it to the general use pool is a problem that so far has eluded a generally workable solution.

The topic of releasing the Class E space for use in the public Internet as globally routable unicast address space has been raised from time to time over the past 15 years or so. Some Internet drafts were published for the IETF’s consideration that either directly proposed releasing this space for use (<https://datatracker.ietf.org/doc/html/draft-wilson-class-e-02>), or outlined the impediments in various host and router implementations that were observed to exist in 2008 when these drafts were being developed (<https://datatracker.ietf.org/doc/html/draft-fuller-240space-02>).

The proposals lapsed, probably due to the larger consideration at the time that the available time and resources to work on these issues were limited and the result of effort spent in freeing off this IPv4 space for general use was only going to obtain a small extension in the anticipated date of depletion of the remaining IPv4 address pools, while the same amount of effort spent on working on advancing IPv6 deployment was assumed to have a far larger beneficial outcome.

From time to time this topic reappears on various mailing lists, but the debates tend to circle around the same set of topics one more time, and then lapse.

As the IANA is no longer a source of “new” addresses, then we need to look at the RIR practices to find these 2.2M addresses. When IP address space is returned to the RIR or reclaimed by the RIR according to the RIR’s policies it is normally placed in a RIR-reserved pool for a period of time and marked as “reserved” by the RIR. Marking returned or recovered addresses as “reserved” for a period of time allows various address prefix reputation and related services, including routing records, some time to record the cessation of the previous state of the addresses prefix, prior to any subsequent allocation. Following some period, which has been observed to be between some months and some years, this reserved space is released for re-use. This is the address space we are seeing as expansion of the allocated address pool in 2021.

The record of annual year-on-year change in allocated addresses per RIR over the same twelve-year period is shown in Table 2. There are some years when the per-RIR pool of allocated addresses shrunk in size. This is generally due to inter-RIR movement of addresses, due to administrative changes in some instances and inter-RIR address transfers in others.

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
APNIC	86.9	119.5	101.0	0.6	1.2	4.6	7.4	6.7	3.2	0.4	10.5	1.7	1.5
RIPE NCC	43.4	52.3	40.5	37.8	1.0	33.8	4.7	4.1	3.7	0.3	12.0	0.4	2.5
ARIN	41.1	27.2	53.8	24.3	19.0	-14.1	2.3	-4.8	-2.3	-0.3	-10.1	-0.9	-1.7
LACNIC	10.5	17.1	13.6	17.3	26.3	18.7	1.2	1.5	1.4	0.1	2.4	1.2	-0.2
AFRINIC	5.9	8.8	9.4	8.5	6.3	12.8	15.0	11.9	7.1	0.2	10.1	-0.2	-0.9
TOTAL	187.8	224.9	218.3	88.5	53.8	55.8	30.6	19.4	13.1	0.7	24.9	2.2	1.2

Table 2 – Annual change in IPv4 Allocated addresses (millions) - Distribution by RIR

Each of the RIRs are running through their final pools of IPv4 addresses. Some of the RIRs have undertaken address reclamation efforts during 2021, particularly in the area of re-designating previously “reserved” addresses as available for allocation as noted above, notably in APNIC and LACNIC.

At the end of 2021, across the RIR system there are some 5.2 million addresses are in the Available pool, held mainly in APNIC (3.5 million) and AFRINIC (1.6 million). Some 12 million addresses are marked as reserved, with 5.5 million held by ARIN, 4 million addresses held by AFRINIC, and 1.8 million addresses held by APNIC. It is evident from this table that there has been a major effort at address reclamation from the “quarantine” pools marked as reserved during 2021, As seen in Table 3, there has been some reduction in the reserved pool in APNIC (700K), ARIN (300K) and RIPE NCC (200K) while the reserved pool in AFRINIC has risen by some 1.2M addresses through 2021.

RIR	Available			Reserved		
	2019	2020	2021	2019	2020	2021
APNIC	2,937,088	4,003,072	3,533,056	4,398,848	2,483,968	1,787,904
RIPE NCC	1,536	328,448	-	1,072,608	965,728	762,104
ARIN	4,096	4,352	4,608	6,137,600	5,509,888	5,244,160
LACNIC	50,688	-	7,168	1,416,448	266,240	224,768
AFRINIC	2,638,848	1,925,888	1,652,480	1,920,256	2,853,888	4,065,024
TOTAL	5,632,256	6,261,760	5,197,312	14,945,760	12,079,712	12,083,960

Table 3 – IPv4 Available and Reserved Pools December 2021

The RIR IPv4 address allocation volumes by year are shown in Figure 1, but it is challenging to understand what is meant by an “allocation” across the entire RIR system as there are some subtle but important differences between RIRs., particularly as they relate to the handling of transfers of IPv4 addresses.

In the case of ARIN, a transfer between two ARIN-serviced entities is conceptually treated as two distinct transactions: a return of the addresses to the ARIN registry and a new allocation from ARIN. The date of the transfer is recorded as the new “allocation date” in the records published by the RIR. Other RIRs treat an address transfer in a manner analogous to a change of the nominated holder of the already-allocated addresses, and when processing a transfer, the RIR’s records preserve the original allocation date for the transferred addresses. When we look at the individual transaction records in the published RIR data, and collect then by year, then in the case of ARIN the collected data includes the volume of transferred addresses that were processed in that year, while the other RIRs only include the allocations performed in that year.

In order to provide a view across the entire system its necessary to use an analysis approach that can compensate for these differences in the ways RIRs record address transactions. In this study, an “allocation” is defined here as a state transition in the registry records from *reserved* or *available* to an *allocated* state. This is intended to separate out the various actions associated with processing address transfers, which generally involve no visible state change, as the transferred address block remains allocated across the transfer, from allocations. This is how the data used to generate Figure 1 has been generated from the RIR published data.




Figure 1 – IPv4 Address Allocations by RIR by year

The number of RIR IPv4 allocations by year, once again generated by using the same data analysis technique as used for Figure 1, are shown in Figure 2.




Figure 2 – IPv4 Allocations by RIR by year

It is clear from these two figures that the average size of an IPv4 address allocation has shrunk considerably in recent years, corresponding to the various IPv4 address exhaustion policies in each of the RIRs.

IPv4 Address Transfers

In recent years, the RIRs have permitted 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 is 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 nine years is shown in Table 4. This includes both inter-RIR and intra-RIR transfers. It also includes both the merger and acquisition-based transfers and the other grounds for of address 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	2020	2021
APNIC	159	185	302	451	841	834	487	528	781	786
RIPE NCC	10	171	1,054	2,836	2,373	2,451	3,774	4,221	4,696	5,743
ARIN				3	22	26	26	68	94	150
LACNIC							2		3	9
AFRINIC							17	27	26	80
Total	169	356	1,356	3,290	3,236	3,311	4,306	4,844	5,600	6,768

Table 4 - IPv4 Address Transfers per year

The differences between RIRs reported numbers are interesting. The policies relating to address transfers do not appear to have been adopted to any significant extent by address holders in AFRINIC and LACNIC serviced regions, while uptake in the RIPE NCC service region appears to be very enthusiastic!

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	2020	2021
APNIC	1.7	2.5	3.9	6.6	8.2	4.9	10.0	4.3	16.6	6.5
RIPE NCC	0.1	2.0	9.6	11.6	9.2	24.6	19.5	26.9	18.2	16.2
ARIN				0.1	0.3	0.2	-	0.3	0.2	0.2
LACNIC							-		-	0
AFRINIC							0.2	0.5	1.2	3.4
Total	1.7	4.5	13.6	18.2	17.6	29.7	29.7	31.9	36.2	26.4

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 - 2021




Figure 4 – Volume of Transferred Addresses: 2012 - 2021

The aggregate total of addresses that have been listed in these transfer logs since 2012 is some 210 million addresses, or the equivalent of 12.5 /8s, which is some 5.7% of the total delegated IPv4 address space of 3.7 billion addresses.

Are Transfers Performing Unused Address Recovery?

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 and recycling these addresses back into active use.

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 transfers is not so clear. The data relating to address re-cycling is inconclusive, in that between 2011 and late 2017 the pool of unadvertised addresses sat between some 38 and 40 /8s. This pool of unadvertised addresses rose from the start of 2018 and by early 2020 there were just under 50 /8s that were unadvertised in the public Internet. This 2-year period of increase in the unadvertised address pool appeared to be a period where IPv4 addresses were being hoarded, though such a conclusion from just this high-level aggregate data is highly speculative and probably unjustified.




Figure 5 – IPv4 Unadvertised Address Pool Size

There has been a substantial reduction in the size of this unadvertised address pool across 2021, and an expansion of the data across 2021 is shown in Figure 6. The major change in 2021 was the announcement of seven /8s from the address space originally allocated to the US Department of Defence in the early days of the then ARPANET. On the 20th of January 2021, AS 8003, an AS registered to “Global Resource Systems”, a contractor to the US Department of Defence, advertised 11.0.0.0/8 and 22.0.0.0/8. On the 7th of April another block, 33.0.0.0/8 was advertised, again by AS 8003. On the 19th and 20th of April a further 4 /8 address blocks, 21.0.0.0/8, 26.0.0.0/8, 28.0.0.0/8 and 30.0.0.0/8, were advertised by AS8003. On the 7th of September these addresses were shifted over to AS 749, a network registered to the US Defence Data Network.

In April 2021 this set of address advertisements was picked up in the media, and a number of motives were ascribed to this action. The likely explanation is far more prosaic, and probably sits in category of defensive measures, as the advertisement of these addresses by a legitimate source makes efforts to co-opt these addresses as part of an attack a little more obvious. At the end of 2021 AS749 originates more IPv4 addresses than any other network, namely some 211,581,184 addresses, or the equivalent of a /4.34 in prefix length notation, or some 5% of the total IPv4 address pool.




Figure 6 – IPv4 Address Pool changes through 2021

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 2020. During 2021, this figure has dropped to 20%, largely due to the advertisement of the legacy US DoD address space, rather than the activation of previously unadvertised address space. This 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 7).




Figure 7 – Ratio of Unadvertised Pool Size to Total Pool Size

This data appears to point to a somewhat sluggish transfer market. The number of transfer transactions is rising, but the total volume of transferred addresses is falling. The address market has not been all that effective in flushing out otherwise idle addresses and re-deploying them into the routed network.

However, as with all other commodity markets, the market price of the commodity reflects the balancing of supply and demand and the future expectations of supply and demand. What can be seen in the price of traded IPv4 addresses over the past 8 years? One of the address brokers, Hilco Streambank, publish the historical price information of transactions (if only all the address brokers did the same, as a market with open price information for transactions can operate more efficiently and fairly than markets where price information is occluded). Figure 8 uses the Hilco Streambank transaction data to produce a time series of address price.

There are a number of distinct behaviour modes in this data. The initial data prior to 2016 reflected a relatively low volume of transactions with stable pricing just below \$10 per address. Over the ensuing 4 years, up to the start of 2019 the price doubled, with small blocks (/24s and /23as) attracting a price premium. The price stabilised for the next 18 months at between \$20 to \$25 per address, with large and small blocks trading as a similar unit price. The most recent 18 months has seen a new dynamic which is reflective of an exponential rise in prices, and the price has lifted to between \$45 and \$60 per address by the end of 2021. There is a greater variance of prices at present, particularly compared to the situation at the start of 2021.

If prices are reflective of supply and demand it appears that demand has increased at a far greater level than supply, and the higher price reflects some form of scarcity premium being applied to addresses in recent times (Figure 8).




Figure 8 – IPv4 Price Time Series (data from Hilco Streambank)

Is supply of tradable IPv4 address declining? One way to provide some insight into answering this question is to look at the registration age of transferred addresses. Are such addresses predominately recently allocated addresses, or are they longer held address addresses where the holder is wanting to realise the inherent value in 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 by transaction is shown on a year-by-year basis in Figures 9 and 10. Some 10% of all transferred addresses in 2021 were drawn from legacy address holders, as shown in Figure 9. It appears that the effort to recycle the legacy address pool has all but run its course and the volume of transferred legacy addresses has declined sharply.

Address holders appear to hold recently allocated addresses for the policy-mandated minimum holding period of some 2 years, but then a visible proportion of these holders transfer these addresses on the market. Some 8% of addresses that were transferred in 2021 were allocated in 2019, and 15% of transferred addresses were allocated since 2016. As shown in Figure 10, these transactions account for some 40% of the transfer transactions in 2021, which is somewhat larger than previous years. It's also useful to consider the situation described in Figure 10 in the light of the data in Table 4. In 2021, Some 84% of all transfers recorded in the registry system are transfers in the RIPE region, so this behaviour of receiving an allocated IP address, holding it for the policy-mandated minimum holding time, and then trading the block on the address market is behaviour that is predominately seen in the RIPE region.

In 2012 some 90% of the transferred address blocks were originally assigned or allocated by an RIR within the previous 10 years. In 2021 some 20% of these transferred addresses relate to addresses that were originally allocated by an RIR in the previous 10 years. Given that the onset of IPv4 address exhaustion occurred some 10 years ago, and the address allocations performed since then were performed under the constraints of exhaustion allocation policies, then in 2021 some 20% of transferred addresses, and 60% of all transfers transactions involve addresses originally allocated under address exhaustion policy frameworks.

The comparison of Figures 9 and 10 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.




Figure 9 – Age distribution of transferred addresses




Figure 10 – Age distribution of transfer transactions

Are transfers effectively recovering and recycling address blocks that have fallen into disuse, and passing them back into active use? The available data indicates that this is not a strong theme of address transfers. The more prevalent current behaviour is where transfers obtain allocations from the RIPE NCC registry, hold it for the policy-mandated minimum holding time, and then monetise the addresses by trading them on the transfer market.

Do Transfers Fragment the Address Space?

The next 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 33,193 transactions described in the RIRs' transfer registries from the start of 2012 until the start of 2022, and of these 11,245 entries list transferred address blocks that are smaller than the original allocated block. Some 34% of transfers implicitly perform fragmentation of the original allocation.

These 11,245 transfer entries that have fragmented the original allocation are drawn from 4,229 such original allocations. On average the original allocation is split into 2.6 smaller address blocks. This data implies that the answer to this 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 229,019 distinct address allocations from the RIRs to end entities as of the end of 2021, and the fragmentation reflected in the addition of 7,016 more specific entries of these original address blocks is around 3% of the total pool of allocated address prefixes.

Imports and Exports of Addresses

The next question concerns the international flow of transferred addresses. Let's look at the ten economies that sourced the greatest volume of transferred addresses, irrespective of their destination (i.e. including 'domestic' transfers within the same economy) (Table 6), and the ten largest recipients of transfers (Table 7), and the ten largest international address transfers (Table 8). We will use the RIR-published transfer data for 2021 as basis for these tables.

Rank	CC	Addresses	Source Economy
1	GB	3,697,152	United Kingdom
2	ZA	3,143,936	South Africa
3	AU	2,387,968	Australia
4	US	2,044,000	USA
5	RU	1,409,536	Russia
6	CN	1,332,224	China
7	DE	1,081,920	Germany
8	CH	1,029,632	Switzerland
9	TH	1,017,856	Thailand
10	HU	785,664	Hungary

Table 6 – Top 10 Countries Sourcing Transferred IPv4 addresses in 2021

Rank	CC	Addresses	Destination Economy
1	US	3,625,312	USA
2	GB	3,367,168	United Kingdom
3	ZA	3,339,520	South Africa
4	RU	1,301,248	Russia
5	DE	1,242,944	Germany
6	TH	1,082,624	Thailand
7	CH	1,080,320	Switzerland
8	CN	890,880	China
9	HU	781,056	Hungary
10	BR	742,912	Brazil

Table 7 – Top 10 Countries Receiving Transferred IPv4 addresses in 2021

There are many caveats about this data collection, particularly relating to the source and meaning of this economy-based geolocation. Even if we use only the country-code entry in the RIR’s registry records, then we get a variety of meanings. Some RIRs use the principle that the recorded country code entry corresponds to the physical location of the headquarters of nominated entity that is the holder of the addresses, irrespective of the locale where the addresses are used on the Internet. Other RIRs allow the holder to update this geolocation entry to match the holder’s intended locale where the addresses will be used. It is generally not possible to confirm the holder’s assertion of location, so whether these self-managed records reflect the actual location of the addresses or reflect a location of convenience is not always possible to determine. When we look at the various geolocation services, or which Maxmind is a commonly used service, where are similar challenges of location. These services generally intend to associate an address with a location that relates to where the address is physically located. At times this is not easy to establish, such as with tunnels used in VPNs. Is the “correct” location the location of the tunnel ingress or tunnel egress? Many of the fine-grained differences in geolocation services reflect the challenges in dealing with VPNs and the various ways these location services have responded. There is also the issue of cloud-based services. Where the cloud service uses anycast the address is located in many locations at once. In the case where the cloud uses conventional unicast, the addresses use may be fluid across the cloud service’s points of presence based on distributing addresses to meet the demands for the service. The bottom line is that these location listings are a “fuzzy” approximation rather than a precise indication of location.

With that in mind let’s now look at imports and exports of addresses of 2021 transfers where the source and destination of the transfers are in different economies.

Rank	From	To	Addresses	Source	Destination
1	AU	US	2,089,984	Australia	USA
2	CN	BR	524,288	China	Brazil
3	US	GB	303,360	USA	United Kingdom
4	US	FR	268,288	United Kingdom	France
5	GB	US	238,592	United Kingdom	USA
6	DE	US	231,680	Germany	USA
7	US	DE	149,248	USA	Germany
8	AU	JP	133,376	Australia	Japan
9	GB	DE	132,608	United Kingdom	Germany
10	US	FI	132,096	USA	Finland
11	SG	CN	131,072	Singapore	China
12	US	ES	128,000	USA	Spain
13	NL	US	120,576	Netherlands	USA
14	NL	DE	89,600	Netherlands	Germany
15	CN	SG	87,040	China	Singapore
16	LV	US	80,896	Latvia	USA
17	GB	ES	77,824	United Kingdom	Spain
18	GB	IT	75,776	United Kingdom	Italy
19	US	HK	73,472	USA	Hong Kong
20	US	PH	66,816	USA	Philippines

Table 8 – Top 20 Economy-to-Economy IPv4 address transfers in 2021

The 2021 transfer logs contain 4,033 domestic address transfers, with a total of 17,723,296 addresses, with the largest activity by address volume in domestic transfers in South Africa, the United Kingdom and Russia. Some 2,735 transfers appear to result in a movement of addresses between countries, involving a total of 8,663,040 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 2021 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-21	Jan-22		Delta	Unchanged	Re-Home	Removed	Added
Announcements	855,767	906,456		50,689	697,928	29,226	82,098	132,787
Address Span (/8s)	236.44	249.61		13.17	208.18	5.35	11.45	24.62
Root Prefixes:	399,812	423,948		24,136	336,842	15,080	27,593	49,000
Address Span (/8s)	170.75	183.29		12.54	160.20	3.20	6.19	17.89
More Specifics:	455,955	482,508		26,553	361,086	14,146	54,505	83,787
Address Span (/8s)	50.54	50.62		0.08	58.28	2.15	5.26	6.73

Table 9 – IPv4 BGP changes over 2021

While the routing table grew by 50,689 entries over the year, the nature of the change is slightly more involved. Some 92,098 prefixes that were announced at the start of the year were removed from the routing system at some time through the year, and 132,787 prefixes were announced by the end of the year that were not announced at the start of the year. More transient prefixes may have appeared and been withdrawn throughout the year of course, but here we are comparing two snapshots rather than looking at every update message. A further 24,226 prefixes had changed their originating Autonomous System number, indicating some form of change in the prefix's network location in some manner.

If we look at the entirety of all updates through 2021 then we see a larger collection of transient address prefixes. A total of 1,128,986 distinct prefixes were observed in the total of all BGP updates through 2021, or 273,219 additional prefixes from the initial set at the start of the year.

We can compare these prefixes that changed in 2021 against the transfer logs for the two-year period 2020 and 2021. Table 10 shows the comparison of these routing numbers against the set of transfers that were logged in these two years.

Type	Listed Transfer	Not a Listed Transfer	Ratio
Re-Homed			
All	2,382	26,844	8.2%
Root Prefixes	1,991	12,773	13.5%
Removed			
All	3,966	78,132	4.8%
Root Prefixes	2,748	24,845	10.0%
Added			
All	7,026	125,761	5.3%
Root Prefixes	5,070	43,930	10.3%

Table 10 – Routing changes across 2021 compared to the Transfer Log Entries for 2020 - 2021

These figures show that some 5%-10% of changes in advertised addresses from the beginning to the end of the year are reflected as changes as recorded in the RIRs' transfer logs. This shouldn't 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 2021, 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 2021 and compare it to the overall age profile of IPv4 addresses in the routing table. This is shown in Figure 11. In terms of addresses that were added in 2021, they differ from the average profile due to a skew in favour of "older" addresses, and 20% of all announced addresses were allocated or assigned more than 30 years ago. This is due to the advertisement of the legacy US DoD /8 allocations, as already noted. Addresses that re-home appear to be disproportionately represented in the age bracket of between 15 to 30 years old.




Figure 11 – Changes to the BGP routing table across 2021 by Address Prefix Age

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. However, I have to note that the assignation of a country code in an address registration reflects the country where address holder is located (the corporate location), and not necessarily the country where the addresses will be deployed.

Rank	CC	IPv4 Pool	% Total	Per-Capita	Economy
1	US	1,619,530,848	43.9%	4.89	United States of America
2	CN	343,938,048	9.3%	0.24	China
3	JP	190,204,416	5.2%	1.50	Japan
4	DE	123,997,312	3.4%	1.48	Germany
5	KR	112,486,912	3.1%	2.19	South Korea
6	GB	112,046,184	3.0%	1.65	United Kingdom
7	BR	87,127,552	2.4%	0.41	Brazil
8	FR	82,836,112	2.2%	1.27	France
9	CA	69,632,000	1.9%	1.84	Canada
10	IT	55,131,712	1.5%	0.91	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 7.8 billion people gives an overall value of 0.47 IPv4 addresses per capita.

Rank	CC	IPv4 Pool	% Total	Per-Capita	Economy
1	SC	7,244,288	0.2%	73.66	Seychelles
2	VA	10,752	0.0%	13.42	Holy See
3	GI	254,464	0.0%	7.55	Gibraltar
4	US	1,619,530,848	43.9%	4.89	United States of America
5	SG	24,064,512	0.7%	4.11	Singapore
6	MU	4,777,472	0.1%	3.76	Mauritius
7	LI	116,256	0.0%	3.05	Liechtenstein
8	SE	29,993,576	0.8%	2.97	Sweden
9	NL	50,198,304	1.4%	2.93	Netherlands
10	NO	15,680,528	0.4%	2.89	Norway
-		3,684,980,232	100.0%	0.47	World

Table 12 – IPv4 Allocated Address Pools ranked by per-Capita holdings

The full table of IPv4 allocations per national economy can be found at <http://resources.potaroo.net/iso3166/v4cc.html>.

IPv4 Address Leasing

It is worth noting that the address market includes leasing as well as sales. Should an entity who requires IPv4 addresses enter the market and perform an outright purchase of the addresses from an existing address holder, or should they execute a timed lease to have the use of these addresses for a specified period and presumably return these addresses at the end of the lease? This lease versus buy question is a very conventional question in market economics and there are various well-rehearsed answers to the question. They tend to relate to the factoring of market information and scenario planning.

If a buyer believes that the situation that led to the formation of a market will endure for a long time, and the goods being traded on the market are in finite supply while the level of demand for these goods is increasing, then the market will add an escalating scarcity premium to the price goods being traded. The balancing of demand and supply becomes a function of this scarcity premium imposed on the good. Goods in short supply tend to become more expensive to buy over time. A holder of these goods will see an increase in the value of the goods that they hold. A lessee will not.

If a buyer believes that the market only has a short lifespan, and that demand for the good will rapidly dissipate at the end of this lifespan, then leasing the good makes sense, in so far as the lessee is not left with a valueless asset when the market collapses.

Scarcity also has several additional consequences, one of which is the pricing of substitute goods. At some point the price of the original good rises to the point that substitution looks economically attractive, even if the substitute good has a higher cost of production or use. In fact, this substitution price effectively sets a price ceiling for the original scarce good.

Some commentators have advanced the view that an escalating price for IPv4 increases the economic incentive for IPv6 adoption, and this may indeed be the case. However, there are other potential substitutes that have been used, most notably NATs (Network Address Translators). While NATs do not eliminate the demand pressure for IPv4, they can go a long way to increase the address utilisation efficiency if IPv4 addresses. NATs allow the same address to be used by multiple customers at different times. The larger the pool of customers that share a common pool of NAT addresses the greater the achievable multiplexing capability.

The estimate as to how long the market in IPv4 addresses will persist is effectively a judgement as to how long IPv4 and NATs can last and how long it will take IPv6 to sufficiently deployed to be viable as an IPv6-only service. At that point in time there is likely to be a tipping point where the pressure for all hosts and networks to support access to services over IPv4 collapses. At that point, the early IPv6-only adopters can dump all their remaining IPv4 resources onto the market as they have no further need for them, which would presumably trigger a level of market panic to emerge as existing holders are faced with the prospect of holding a worthless asset and are therefore under pressure to sell off their IPv4 assets while there are still buyers in the market.

While a significant population of IPv4-only hosts and networks can stall this transition and increase scarcity pressure, if the scarcity pressure becomes too great the impetus of IPv6-only adoption increases to the level that the IPv6-connected base achieves market dominance. When this condition is achieved the IPv4 address market will quickly collapse.

IPv6 in 2021

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 2021.

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 prefix 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 using a /56, and many other providers using 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 2021.

Comparing 2020 to 2021 the number of individual allocations of IPv6 address space has remained stable while the number of IPv4 allocations has increased somewhat (Table 13).

Allocations	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
IPv6	1,279	2,461	3,588	3,302	4,020	4,537	4,736	5,593	5,765	6,311	6,922	5,455	5,471
IPv4	6,698	7,739	10,062	8,608	7,073	11,114	11,644	9,911	8,091	9,707	13,112	6,263	7,829

Table 13 - Number of individual Address Allocations, 2009 - 2021

The amount of IPv6 address space distributed in 2021 is 30% more than the amount that was allocated in 2020 (Table 14).

Addresses	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
IPv6 (/32s)	1,091	5,844	15,018	17,740	24,103	17,917	20,234	25,300	19,986	39,278	36,507	21,835	28,690
IPv4 (/32s)(M)	190.1	248.0	201.0	114.9	64.9	68.5	45.3	22.5	21.2	16.1	15.4	5.6	4.3

Table 14 – Volume of Address Allocations, 2009 – 2021

Regionally, each of the RIRs saw IPv6 allocation activity in 2021 that was on a par with those seen in the previous year, with the exceptions of LACNIC, which saw a 50% decrease in allocations in 2021, and the RIPE NCC where the number of allocations lifted back to levels last seen in 2018 (Table 15).

Allocations	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
ARIN	394	638	1,039	611	560	512	604	646	684	659	605	644	671
APNIC	194	670	641	599	540	528	778	1,681	1,369	1,474	1,484	1,498	1,392
RIPE NCC	642	1,048	1,647	1,756	2,151	2,227	2,207	2,141	2,051	2,620	3,104	1,403	2,542
LACNIC	35	50	132	253	696	1,208	1,061	1,009	1,549	1,448	1,614	1,801	730
AFRINIC	14	55	129	83	73	62	86	116	112	110	115	109	136
	1,279	2,461	3,588	3,302	4,020	4,537	4,736	5,593	5,765	6,311	6,922	5,455	5,471

Table 15 - IPv6 allocations by RIR

The address assignment data tells a slightly different story. Table 16 shows the number of allocated IPv6 /32's per year. There were two /20 allocations in the year, one to CERNET in China (an earlier allocation of an IPv6 /20 was also made to CERNET in 2019) and a second to the National Knowledge Network in India. A third large allocation, a /22, was made to China's State Oil and Gas Pipeline Group.

Addresses (/32s)	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
ARIN	257	584	2,280	1,676	12,581	5,232	642	1,088	1,372	845	5,522	4,975	374
APNIC	175	3,241	9,506	3,807	4,462	2,663	2,109	1,236	4,228	19,690	7,945	7,365	10,185
RIPENCC	617	1,869	2,425	3,729	6,385	8,608	12,039	21,717	12,844	17,335	21,180	7,737	17,329
LACNIC	32	46	652	4,325	608	1,363	973	1,181	1,429	1,336	1,496	1,669	660
AFRINIC	9	104	155	4,203	67	51	4,471	78	113	72	364	89	142
	1,091	5,844	15,018	17,740	24,103	17,917	20,234	25,300	19,986	39,278	36,507	21,835	28,690

Table 16 - IPv6 address allocation volumes by RIR

Dividing addresses by allocations gives the average IPv6 allocation size in each region (Table 16). Overall, the average IPv6 allocation size remains around a /30, with the RIPE NCC and APNIC averaging larger individual IPv6 allocations than the other RIRs.

	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
ARIN	/32.6	/32.1	/30.9	/30.5	/27.5	/28.6	/31.9	/31.2	/31.0	/31.6	/28.8	/29.1	/32.8
APNIC	/32.1	/29.7	/28.1	/29.3	/29.0	/29.7	/30.6	/32.4	/30.4	/28.3	/29.6	/29.7	/29.1
RIPENCC	/32.1	/31.2	/31.4	/30.9	/30.4	/30.0	/29.6	/28.7	/29.4	/29.3	/29.2	/29.5	/29.2
LACNIC	/32.1	/32.1	/29.7	/27.9	/32.2	/31.8	/32.1	/31.8	/32.1	/32.1	/32.1	/32.1	/32.1
AFRINIC	/32.6	/31.1	/31.7	/26.3	/32.1	/32.3	/26.3	/32.6	/32.0	/32.6	/30.3	/32.3	/31.9
	/32.2	/30.8	/29.9	/29.6	/29.4	/30.0	/29.9	/29.8	/30.2	/29.4	/29.6	/30.0	/29.6

Table 17 – Average IPv6 address allocation size by RIR

The number and volume of IPv6 allocations per RIR per year is shown in Figures 12 and 13.




Figure 12 – Number of IPv6 Allocations per year




Figure 13 – Volume of IPv6 Allocations per year

It might be tempting to ascribe the decline in 2020 of IPv6 allocations from the RIPE NCC to the year where many European countries were hit hard by COVID-19 measures. Arguing against that is the observation that countries all over the world have been similarly affected, yet the decline in IPv6 allocation activity in 2020 is only seen in the data from the RIPE NCC. However, it's an interesting question to ask as to why the IPv6 address allocation activity has slumped in the European economies, but not in China, the US and Brazil (Table 19).

Rank	2017		2018		2019		2020		2021	
1	Brazil	1,114	Brazil	1,049	Brazil	1,112	Brazil	1,394	USA	619
2	USA	634	Russia	638	USA	538	USA	588	Russia	576
3	Germany	270	USA	595	Russia	502	Indonesia	389	Brazil	508
4	Russia	220	Germany	308	Germany	407	India	226	Netherlands	448
5	Australia	211	China	253	Indonesia	366	Netherlands	199	India	390
6	China	208	Indonesia	213	Netherlands	342	Germany	192	UK	304
7	Netherlands	194	UK	184	UK	223	Bangladesh	182	Bangladesh	213
8	UK	190	Bangladesh	183	Bangladesh	202	Russia	128	Germany	196
9	Indonesia	187	India	168	France	179	Australia	118	Indonesia	110
10	Argentina	178	Netherlands	162	China	165	China	115	Hong Kong	108

Table 18 - IPv6 allocations by Year by Economy

Table 18 shows the countries who received the largest number of individual IPv6 allocations, while Table 19 shows the amount of IPv6 address space assigned on a per economy basis for the past 5 years (using units of /32s).

Rank	2017		2018		2019		2020		2021	
1	China	2,245	China	17,647	China	6,787	China	6,765	China	5,424
2	USA	1,497	Russia	4,675	USA	5,510	USA	5,051	Russia	4,409
3	Germany	1,364	Germany	1,932	Russia	3,716	Brazil	1,358	India	4,281
4	Russia	1,358	UK	1,209	Germany	2,522	Netherlands	1,331	Netherlands	3,390
5	Netherlands	1,296	Singapore	1,055	Netherlands	2,516	Germany	716	UK	2,249
6	Spain	1,170	Netherlands	1,025	UK	1,355	Russia	715	Germany	896
7	India	1,087	Brazil	1,007	France	1,182	UK	552	Ukraine	651
8	UK	1,072	USA	874	Italy	1,052	Italy	391	Lithuania	633
9	Brazil	1,049	Spain	851	Brazil	1,049	France	390	Brazil	502
10	France	714	France	722	Spain	854	Turkey	290	USA	491

Table 19 - IPv6 Address Allocation Volumes by Year by Economy (/32s)

We can also look at the allocated address pools for the 25 national economies with the largest allocated address pools in IPv6, and the current picture is shown in Table 20.

While the United States also tops this list in terms of the total pool of allocated IPv6 addresses, with some 20% of the total span of allocated IPv6 addresses, the per capita number is lower than many others in this list. Sweden has a surprisingly high number. The large address pools are likely due to early IPv6 allocations, made under a somewhat different allocation policy regime that that used today.

Some twenty years ago it was common practice to point out the inequities in the state of IPv4 address deployment. Some US universities had more IPv4 addresses at their disposal than some highly populated developing economies, and the disparity was a part of the criticism of the address management practices of the time. The RIR system was intended to address this issue of predisposition to a biased outcome. The concept behind the system that within the regional community each community had the ability to develop their own address distribution policies and could determine for themselves what they meant by such terms as “fairness” and “equity” and then direct their regional address registry to implement these policies. While IPv4 had a very evident early adopter reward, in that the address allocations in the IPv4 class-based address plan could be quite extravagant, the idea was that in IPv6, where the address allocations were developed from the outset through local bottom-up policy determinate frameworks, such evident inequities in the outcome would be avoided. It was also envisaged that with such a vast address plan provided by 128 bits of address space, the entire concept of scarcity and inequity would be largely irrelevant. 2^{128} is a vast number and the entire concept of comparison between two vast pools of addresses is somewhat irrelevant. So when we look at the metric of /48s per head of population don’t forget that a /48 is actually 80 bits of address space, which is massively larger than the entire IPv4 address space. Even India’s average of 0.1 /48s per capita is still a truly massive number of IPv6 addresses!

However, before we go too far down this path it is also useful to bear in mind that the 128 bits of address space in IPv6 has become largely a myth. We sliced off 64 bits in the address span for no particularly good reason, as it turns out. We then sliced off a further 48 bits for again no particularly good reason. So the vastness of the address space represented by 128 bits in IPv6 is in fact not so vast. The usable address prefix space in IPv4 roughly equates a /32 end address in IPv4 with around a /48 prefix in IPv6. So perhaps this comparison of /48s per capita is not entirely fanciful, and there is some substance to the observation that there are inequities in the address distribution in IPv6 so far. However, unlike IPv4, the exhaustion of the IPv6 address space is still quite some time off, and we still believe that there are sufficient IPv6 addresses to support a uniform address utilisation model across the entire world of silicon over time.

Rank	CC	Allocated (/48s)	% Total	/48s p.c.	Advertised (/48s)	Deployed	Name
1	CN	3,936,092,249	17.4%	2.7	1,655,222,454	17.8%	China
2	US	3,807,856,200	16.8%	11.5	1,022,265,534	11.0%	USA
3	DE	1,490,944,680	6.6%	17.8	1,031,211,619	11.1%	Germany
4	GB	1,458,438,378	6.4%	21.5	404,885,857	4.3%	UK
5	RU	1,073,414,454	4.7%	7.4	177,953,050	1.9%	Russia
6	FR	949,227,921	4.2%	14.5	160,147,671	1.7%	France
7	NL	802,750,772	3.5%	46.8	373,539,852	4.0%	Netherlands
8	JP	663,101,642	2.9%	5.2	508,415,880	5.5%	Japan
9	IT	648,744,989	2.9%	10.7	408,634,242	4.4%	Italy
10	AU	618,398,908	2.7%	24.3	307,830,331	3.3%	Australia
11	BR	536,371,604	2.4%	2.5	353,980,127	3.8%	Brazil
12	SE	434,635,090	1.9%	43.0	354,176,919	3.8%	Sweden
13	IN	425,591,614	1.9%	0.3	87,989,109	0.9%	India
14	PL	386,859,241	1.7%	10.2	206,078,405	2.2%	Poland
15	ES	381,812,772	1.7%	8.2	94,643,292	1.0%	Spain
16	AR	349,046,886	1.5%	7.7	283,176,806	3.0%	Argentina
17	KR	344,915,979	1.5%	6.7	4,321,983	0.0%	Korea
18	ZA	317,986,504	1.4%	5.4	289,465,297	3.1%	South Africa
19	EG	270,139,394	1.2%	2.6	269,942,784	2.9%	Egypt
20	CH	236,388,527	1.0%	27.3	113,156,004	1.2%	Switzerland
21	TR	216,006,682	1.0%	2.6	45,105,048	0.5%	Turkey
22	CZ	188,022,895	0.8%	17.6	109,987,216	1.2%	Czech Republic
23	UA	179,044,525	0.8%	4.1	69,595,577	0.7%	Ukraine
24	IR	176,357,383	0.8%	2.1	28,537,229	0.3%	Iran
25	TW	168,165,391	0.7%	7.1	154,817,108	1.7%	Taiwan

Table 20 – IPv6 Allocated Address pools per National Economy – December 2021

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 14 – Allocated, Unadvertised and Advertised IPv6 addresses




Figure 15 – 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 40% 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 2022 we have advertised an IPv6 address span which is the equivalent of some 150,000 /32s, or some 9 billion end-site /48 prefixes. That is just 0.003% of the total number of /48 prefixes in IPv6.

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. The next year, 2018, was a quieter year, although the rise in the second half of the year is due to the initial efforts of mass scale IPv6 deployment in the major Chinese service providers. This movement accelerated in 2019 and the overall move of some 5% in IPv6 deployment levels had a lot to do with the very rapid rise of the deployment of IPv6 in China.




Figure 16 – IPv6 Deployment measurement 2010 – 2021

In 2021 the growth patterns for IPv6 are more diffuse around the world with a 3% overall growth rate, although there has been steady growth in IPv6 deployment in Brazil, Colombia, Ecuador and Argentina. The regions where IPv6 deployment is low compared to this 30% Internet-wide average includes Africa, Southern and Eastern Europe, the Middle East and Central Asia (Figure 17).

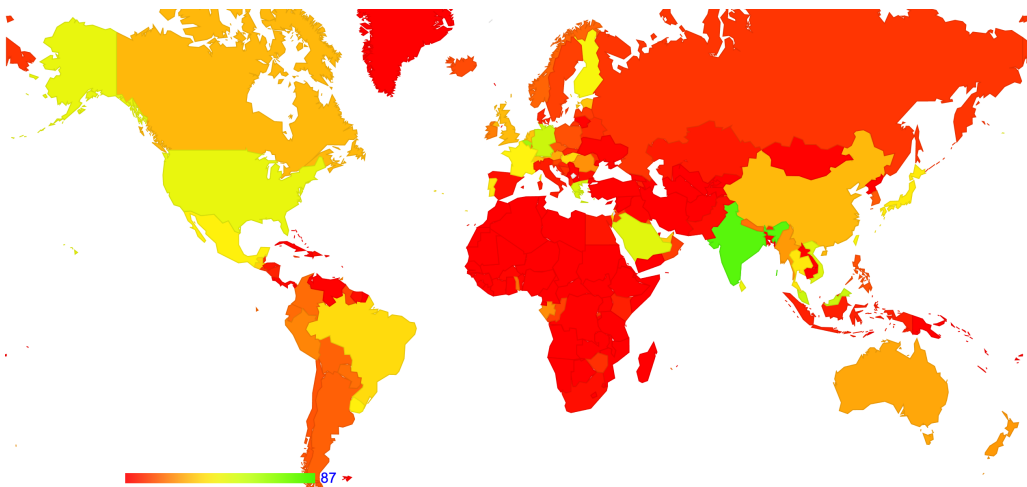


Figure 17 – IPv6 Deployment measurement - December 2021

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 decade has been dominated by the mass marketing of mobile internet services, and the Internet's growth rates for 2014 through to 2016 perhaps might have been the highest so far recorded. This would've been visible in the IP address deployment data 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 IPv6 the extremely large size of the address space masks out much of this volume. A single IPv6 /20 allocation to an ISP allows for 268 million /48 allocations, or 68 billion /56 allocations, so much of the growth in IPv6-using networks is simply hidden behind the massive address plan that lies behind IPv6.

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.

More generally, 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 AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net