

January 2025  
Geoff Huston

### BGP in 2024

At the start of each year, it's been my habit to report on the behaviour of the Internet's inter-domain routing system over the previous 12 months, looking in some detail at some metrics from the routing system that can show the essential shape and behaviour of the underlying interconnection fabric of the Internet.

The strong growth numbers that were a constant feature of the past thirty years of the Internet's story are simply not present in the data in recent times. Is there is no more demand capacity to fuel further growth? Is the Internet losing its investment appeal along with so many other signals of investor disillusion over outlandish growth predictions in technology-based services? Or has the massive transition into content distribution networks for digital services meant that there is a declining demand for the traditional form of content distribution over on-demand network access?

Let's take a look at the BGP view of 2024 and see how these larger technical and economic considerations are reflected in the behaviour of the Internet's inter-domain routing system.

One reason why we are interested in the behaviour of the routing system is that at its heart the routing system has no natural self-constraint. Our collective unease about routing relates to a potential scenario where every network decides to disaggregate their prefixes and announce only the most specific prefixes, or where every network applies routing configurations that are inherently unstable, and the routing system rapidly reverts into oscillating between unstable states that generate an overwhelming stream of routing updates into the inter-domain routing space. In such scenarios, the routing protocol we use, the Border Gateway Protocol, or BGP, will not help us by attempting to damp down such behaviour. Indeed, there is a very real prospect that in such scenarios the protocol behaviour of BGP could well amplify the situation!

BGP is an instance of a Bellman-Ford distance vector routing algorithm. This algorithm allows a collection of connected devices (BGP speakers) to each learn the relative topology of the connecting network. The basic approach of this algorithm is very simple: each BGP speaker tells all its other neighbours about what it has learned if the newly learned information alters the local view of the network. Each time an adjacent BGP neighbour informs a BGP speaker about a change of reachability for an IP address prefix, the BGP speaker compares this new reachability information against its stored knowledge that was gained from previous announcements from other neighbours. If this new reachability information describes a better network path to the prefix, then the local speaker stores this prefix and associated next-hop forwarding decision into the local forwarding table and then informs all its adjacent neighbours of this new path to a prefix, implicitly citing itself as the next hop. In addition, there is a withdrawal mechanism. When a BGP speaker determines that it no longer has a viable path to a given prefix it announces a withdrawal of this prefix to all its adjacent neighbours. When a BGP speaker receives a withdrawal, it stores the withdrawal against this neighbour. If this withdrawn route happened to be the currently preferred next hop (neighbour) for this prefix, then the BGP speaker will examine its per-neighbour data sets to determine which stored announcement represents the best path from those that are still extant. If it can find such an alternative path, it will copy this into its local forwarding table and announce this new

preferred path to all its BGP neighbours. If there is no such alternative path, it will announce a withdrawal to its neighbors, indicating that it no longer can reach this prefix.

And that's the one paragraph summary of BGP.

The first metric of interest is the size of the routing tables. Each router needs to store a local database of all prefixes announced by each routing peer. In addition, conventional routing design places a complete set of "best" paths into each line card and performs a lookup into this forwarding data structure for each packet. This represents an extremely challenging silicon design problem. The larger the routing search space, the more challenging the problem!

Why does memory size matter for a router?

If you look at the internals of a high-speed Internet router operating the default-free zone of the Internet one of the more critical performance aspects of the unit is to make a forwarding decision for each packet within the mean inter-packet arrival time, and preferably within the inter-arrival time of minimum-sized IP packets.

A router line card with an aggregate line rate across all of its serial interfaces of some 10Tbps (which is probably not that large an aggregate capacity by today's standards) needs to process each packet within 70 nanoseconds, assuming that the average packet size is 900 octets). If the average memory access cycle time is 10ns then this implies that the router line card processor needs to scan the entire decision space within just 7 memory access operations just to keep pace with the anticipated peak packet rate. A densely packed binary search across 1M entries will require an average of 20 decisions when using conventional serial binary decision logic, so it's clear that some other decision approach is needed here. These very high-speed decision tables are often implemented using high speed *content-addressable memory* to bypass this limitation. *Ternary content-addressable memory* (TCAM) can search its entire contents in a single memory cycle. It's fast, but it's also one of the most expensive components of a high-speed router line card.

TCAM size is what you purchase when you buy the router, so you need to pay attention to not only what you need today, but what you may need over the operational lifetime of the unit. If the router is to be useful in, say, 5 years from now, then you need to deploy units that can maintain their switching performance levels five years from now. That often implies configuring your units with sufficient TCAM memory to contain the combination of IPv4 and IPv6 routing tables that are not only adequate for today but are adequate to meet the routing table requirements some years into the future. Getting it wrong means that you've spent too much on your switching equipment if you over-provision or are forced to retire the equipment prematurely if you under-provision. What this means is this size question is an important question both to network operators and to designers and vendors of network switching equipment.

There is also the consideration of the overall stability of the system. Processing a routing update requires several lookups into local data structures as well as local processing steps. Each router has a finite capacity to process updates, and once the update rate exceeds this local processing capability, then the router will start to queue up unprocessed updates. In the worst case, the router will start to lag in real-time, so that the information a BGP speaker is propagating reflects a past local topology, not necessarily the current local topology. If this lag continues, then at some point unprocessed updates may be dropped from the queue. BGP has no inherent periodic refresh capability, so when information is dropped, then the router and its neighbours fall out of sync with the network topology. At its most benign, the router will advertise "ghost" routes where the prefix is no longer reachable, yet the out-of-sync router will continue to advertise reachability. At its worst, the router will set up a loop condition and as traffic enters the loop it will continue to circulate through the loop until the



packet's TTL expires. This may cause saturation of the underlying transmission system and trigger further outages which, in turn, may add to the routing load.

The two critical metrics of the Internet's routing table that we are most interested in are the *size* of the routing space and its level of updates, or *churn*. Here we will concentrate on the first of these metrics, the size of the routing space, and the changes that occurred through 2024, and use this data to extrapolate forward and look at 5-year projections for the size of the routing table in both IPv4 and IPv6.

## The BGP Measurement Environment

In trying to analyse long baseline data series the ideal approach is to keep as much of the local data gathering environment as stable as possible. In this way, the changes that occur in the collected data reflect changes in the larger environment, as distinct from changes in the local configuration of the data collection equipment.

The major measurement point being used here is a BGP speaker configured within AS 131072. This network generates no traffic and originates no routes in BGP. It's a passive network that has a single BGP speaker that has been logging all received BGP updates since 2007. The router is fed with a default-free BGP feed from AS 4608, which is the APNIC network located in Australia, and AS 4777, which is the APNIC network located in Japan, for both IPv4 and IPv6 routes.

There is also no internal routing (iBGP) component in this measurement setup. While it has been asserted at various times over the years that iBGP is a major contributor to BGP scalability concerns in BGP, the consideration here in trying to quantify this assertion is that there is no "standard" iBGP configuration, as each network has its own rather unique configuration of Route Reflectors and iBGP peers. This makes it hard to generate a "typical" iBGP load profile, let alone analyse the general trends in iBGP update loads over time.

In this study, the scope of attention is limited to a simple eBGP configuration that is likely to be found as a "stub" AS at the edge of the Internet. This AS is not an upstream for any third party, it has no transit role, and does not have a large set of BGP peers. It's a simple view of the routing world that I see when I sit at an edge of the Internet. Like all BGP views, it is totally unique to this network, and every other network will see a slightly different Internet with different metrics. However, the behaviour seen by this stub network at the edge of the Internet is probably similar to most other stub networks at the edge of the Internet. While the fine details may differ, the overall picture is probably much the same. This BGP view is both unique and typical at the same time.

To complement this single view of the BGP network we will use the recourses of two large *route collector* systems. We use *RouteViews*, a project supported by the Network Startup Resource Center (NSRC) at the University of Oregon. RouteViews has also enjoyed significant levels of support from the National Science Foundation and numerous industry entities (<https://www.routeviews.org/routeviews/supporters/>). This route collector has been in operation continuously since 1997 and holds much of the routing history of the Internet in its archives.

The **RouteViews Project** was originally intended to offer a multi-perspective real-time view of the inter-domain routing system, allowing network operators to examine the current visibility of route objects from various points in the inter-domain topology. What makes RouteViews so unique is that it archives these routing tables every two hours and has done so for more than two decades. Their system also archives every BGP update message to complement the regular snapshot of the routing information base (RIB). This vast collection of data is a valuable research data trove.

The folk at the RouteViews Project, operated with the support from the Network Startup Resource Center (NSRC) at the University of Oregon and the US National Science Foundation and a number of other corporate supporters should be commended for their efforts here. This is a very unique data set if you are interested in understanding the evolution of the Internet over the years.

We also use the *Routing Information Service* (RIS), operated by the RIPE NCC. This service is operated by the RIPE Network Coordination Centre, the Regional Internet Registry for Europe, Middle East and Central Asia. Similarly to RouteViews, RIS operates 27 route collectors located across a diverse collection of locations (<https://www.ris.ripe.net/peerlist/all.shtml>).

There are currently some 777 BGP sessions that are being collected across all routing peers of these two route collector systems, with 398 route sessions providing an IPv4 routing table, and 379 IPv6 sessions.

### The IPv4 Routing Table

Measurements of the size of the routing table have been taken regularly since the start of 1988, roughly coinciding with the early days of the NSFNET in the United States.

Figure 1 shows a rather unique picture of the size of the routing table, as seen by all the peers of the RouteViews route collector system since its inception in 1997 and incorporating earlier measurements conducted on the NSFNET.

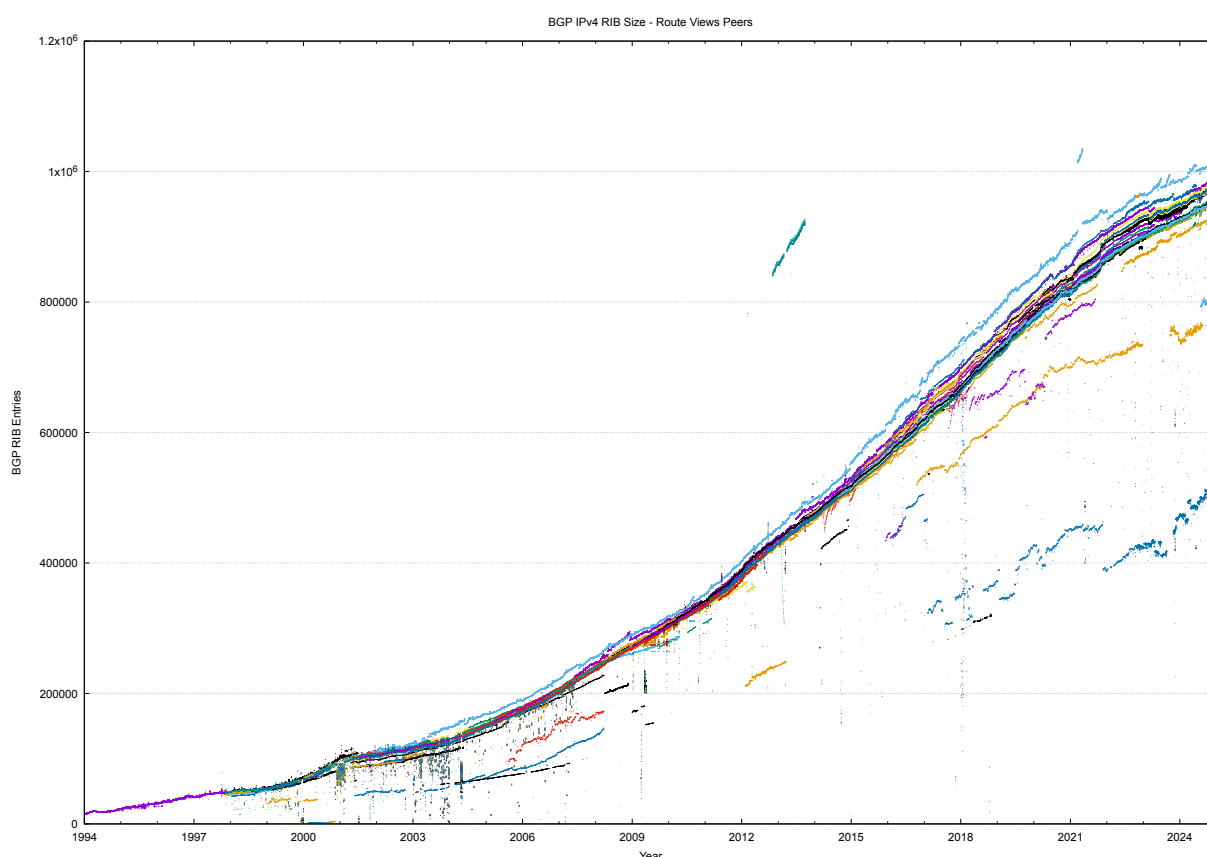


Figure 1 – IPv4 routing table since 1994 as seen by RouteViews peers

Several events are visible in the history of the routing table, such as the bursting of the Internet bubble in 2001, the impact of the global financial crisis in 2009 and the lingering effects of the Covid-related shutdowns in late 2020. What is perhaps surprising is one ongoing event that is not visible in this plot. Since 2011 the supply of IPv4 addresses has been progressively constrained as the unassigned address pools of the various Regional Internet Registries have been exhausted. Yet there is no visible impact on the rate of growth of the number of announced prefixes in the global routing system since 2011. In terms of the size of the routing table it's as if the exhaustion of IPv4 addresses has not happened at all for the ensuing decade. It is only by 2021 that we see some tapering of the growth of the size of the IPv4 routing table.

BGP is not just a reachability protocol. Network operators can manipulate traffic paths using selective advertisement of more specific addresses, allowing BGP to be used as a traffic engineering tool. These more specific advertisements often have a restricted propagation. This is evident in Figure 1, where there is no single

plot in this figure (in a entirely homogenous routing environment each BGP peer would see precisely the same routing table). There is a variance across the various peers of these route collectors that is around 50,000 routes.

We can zoom in to look at just 2024, taking the collection of RouteViews and RIS BGP peers (Figure 2).




Figure 2 – IPv4 routing table 2024, as seen by all RouteViews and RIS peers

Both Figures 1 and 2 illustrate an important principle in BGP, that there is no single authoritative view of the Internet’s inter-domain routing table, as all views are relative to the perspective of each BGP speaker. These figures also illustrate that at times the cause of changes in routing is not necessarily a change at the point of origination of the route which would be visible to all BGP speakers across the entire Internet, but it may well be a change in transit arrangements within the interior of the network that may expose, or hide, collections of routes.

The issue of the collective management of the routing system can be seen as an example of the condition of the *tragedy of the commons* where the self-interest of one actor in attempting to minimise its transit service costs becomes an incremental cost in the total routing load that is borne by other actors. To quote the Wikipedia article on this topic “In absence of enlightened self-interest, some form of authority or federation is needed to solve the collective action problem.” This appears to be the case in the behaviour of the routing system, where there is an extensive reliance on enlightened self-interest to be conservative in one’s own announcements.

The next collection of plots (Figures 3 through 12) contains some of the vital characteristics for the IPv4 BGP network since the start of 2017 to the end of 2024, using the routing data for a single BGP session, collected from AS 131072.




Figure 3 - IPv4 BGP Routing Table Size (RIB)

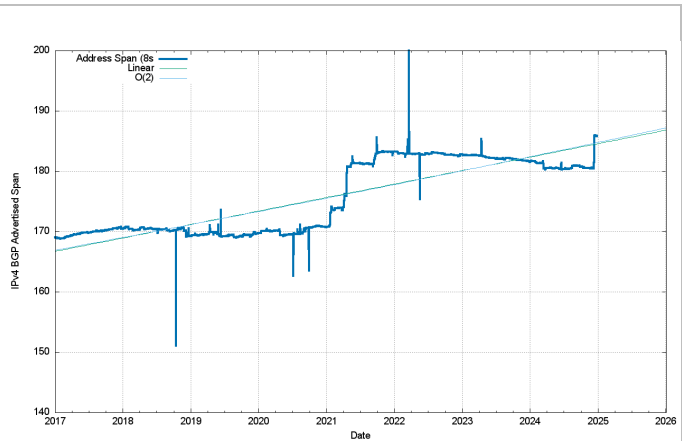


Figure 4 - IPv4 Announced Address Span




Figure 5 - IPv4 More Specific Announcements




Figure 6 - IPv4 Relative Proportion of More Specific Announcements




Figure 7 - IPv4 Average Announcement Size




Figure 8 - IPv4 Prefix Size Counts




Figure 9 - IPv4 Average AS Path Length

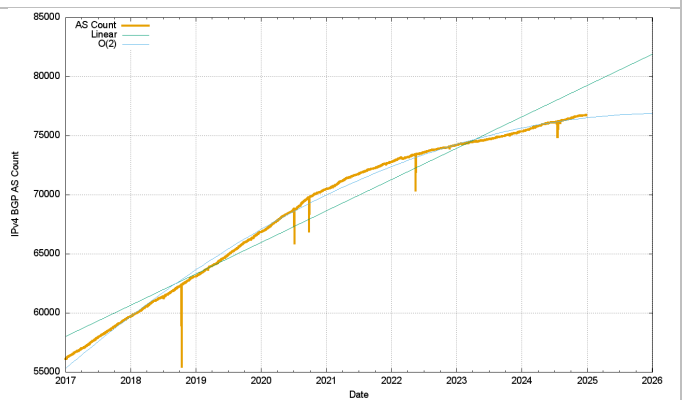


Figure 10 - IPv4 AS Count




Figure 11 - Transit vs Total AS Count

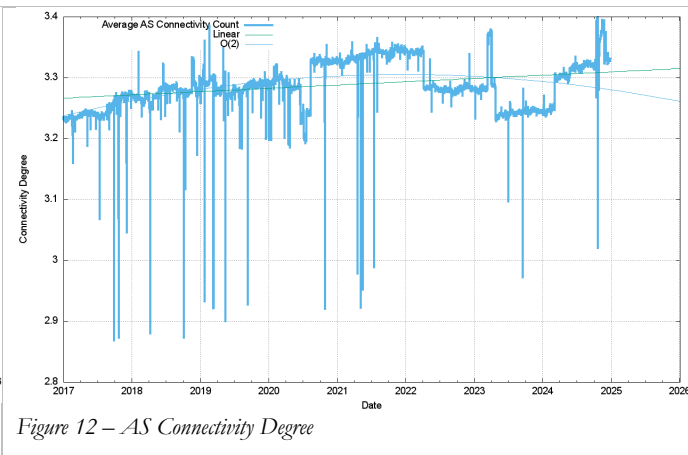


Figure 12 - AS Connectivity Degree

Figure 3 shows the total number of routes in the routing table over this period. The routing table has continued to grow, largely as a result of the advertisement of more specifics.

The eight-year period since the start of 2016 has seen the span of addresses advertised in the routing system slowing down until the start of 2021 (Figure 4). In 2021 we saw a number of large address blocks being advertised in the routing system by agencies associated with the US Department of Defence. Across 2022 and 2023 the total span of advertised IPv4 addresses declined. There was a sharp rise in the total span of address space on the 12<sup>th</sup> December 2024, coinciding with the advertisement of a total span of 81,224,704 addresses (the equivalent of 4.8 /8s) by ASes operated by Amazon. Their major network, AS16509, announced a span of 154,961,152 addresses at the end of 2024, or some 4.97% of the total IPv4 announced address span, second only to the 224,851,712 addresses announced by AS749 for the US Department of Defence. The collection of Amazon ASes collectively announces a span of 157,809,408 addresses at the end of 2024.

In terms of more specific advertisements and covering aggregate advertisements (Figure 5), the number of covering aggregate announcements increased across 2024, as did the number of more specifics. Looking at the ratio of these two counts, the ratio has remained steady over 2024 (Figure 6). The average prefix size is now somewhat smaller than a /20 (Figure 7). Prefixes sizes of /24, /23 and /22 in total now account for 83% of the entire IPv4 routing table (Figure 8).

The topology of the network has remained relatively consistent, with the growth of the Internet being seen as relatively static density of interconnectivity. The average AS Path length has been constant since mid-2020. (Figure 9).

Further evidence of the slowing of the growth in the IPv4 internet can be seen in the AS count (Figure 10). The growth of the AS count started to decline in late 2020 and continued to decline in the ensuing years. This is a likely signal of network saturation in many markets. The number of transit networks was held constant across 2024, which appears to be a related signal of market saturation (Figure 11).

The year-by-year summary of the IPv4 BGP network (as seen by AS 131072) over the 2021-2025 period is shown in Table 1.

	Jan-21	Jan-22	Jan-23	Jan-24	Jan-25	Annual Growth			
						2021	2022	2023	2024
<b>Prefix Count</b>	860,000	906,000	940,000	943,000	996,000	5%	4%	0%	6%
<b>Root Prefixes</b>	400,000	420,000	445,000	457,000	470,000	5%	6%	3%	3%
<b>More Specs</b>	460,000	486,000	495,000	486,000	526,000	6%	2%	-2%	8%
<b>Address Span (/8s)</b>	171.4	183.3	182.8	181.7	185.8	7%	0%	-1%	2%
<b>AS Count</b>	70,400	72,800	74,200	75,300	76,700	3%	2%	1%	2%
<b>Transit Ases</b>	10,200	10,800	10,800	11,000	11,300	6%	0%	2%	3%
<b>Stub Ases</b>	60,200	62,000	63,400	64,300	65,400	3%	2%	1%	2%

Table 1 - IPv4 BGP Table Growth Profile



In terms of advertised prefixes, the size of the routing table grew by some 53,000 entries, or 6%. The number of root prefixes increased by 13,000 entries, while the number of more specific routes increased by 40,000 entries. This year (2024) saw the IPv4 routing table resume the growth metrics of 2021 and 2022.

The total span of advertised addresses increased by some 69M IPv4 /32s across the year, due to the advertising of a large pool of addresses by Amazon in December 2024.

The number of routed Stub AS numbers (edge networks) grew by 1.6% in 2024, and the total number of ASes visible in the IPv4 network grew by 351 ASes across 2024, or 3.2%.

Let's look in a little more detail at the 10 networks that had the highest route object net growth, and the highest net route withdrawal count for the year (Table 2).

Net Growth						
AS Num	Change	Jan-24	more specs	Dec-24	more specs	Name, CC
AS16509	<b>3,007</b>	9,471	1,601	12,478	7,523	AMAZON-02, US
AS174	<b>1,922</b>	2,313	1,627	4,235	3,210	Cogent, US
AS4155	<b>1,440</b>	848	-	2,288	-	USDA, US
AS984	<b>1,308</b>	82	28	1,390	687	OCTOPUS WEB, US
AS28202	<b>1,115</b>	10	-	1,125	1,115	Rede Brasileira, BR
AS17561	<b>1,113</b>	-	-	1,113	-	LARUS, HK
AS20011	<b>970</b>	31	18	1,001	56	INTERNET-SOLUTIONS, ZA
AS7018	<b>745</b>	2,502	245	3,247	193	ATT-INTERNET, US
AS37069	<b>655</b>	177	83	832	801	MOBINIL, EG
AS152194	<b>655</b>	-	-	655	261	CTG Server, HK
Net Withdrawal						
AS Num	Change	Jan-24	more specs	Dec-24	more specs	Name, CC
AS8551	<b>-1,867</b>	3,327	3,292	1,460	1,427	Bezeqint, IL
AS18566	<b>-1,855</b>	1,855	426	-	-	Megapaths, US
AS5650	<b>-1,349</b>	2,010	287	661	239	Frontier, US
AS46573	<b>-766</b>	898	-	132	-	Layerhost, US
AS8452	<b>-764</b>	2,394	2,380	1,630	1,616	Telecom Egypt, EG
AS61317	<b>-723</b>	1,393	18	670	7	Heficed, US
AS64050	<b>-671</b>	811	373	140	23	BGP Network, HK
AS8151	<b>-621</b>	12,173	9,891	11,552	9,340	Uninet, MX
AS6849	<b>-508</b>	1,039	1,021	531	511	UKRTELNET, UA
AS9394	<b>-482</b>	483	379	1	-	TieTong Telecom, CN

Table 2 – IPv4 Advertised Prefix Changes – Top 10 ASes

The majority of additional route objects are more specifics. The highest growth in route objects over 2024 was AS16509, where the number of more specifics grew by 5,922 routes. One of the more notable actions was AS 28202, Rede Brasileria, which advertised the same 10 aggregates and added a further 1,115 more specifics.

We can also look at the total span of advertised addressers for each AS, comparing the advertised address space at the start of 2024 with that at the end of the year. This data is shown in Table 3.

Net Growth (M Addresses)					Net Withdrawal (M Addresses)				
AS Num	Change	Jan-24	Dec-24	Name, CC	AS Num	Change	Jan-24	Dec-24	Name, CC
AS16509	<b>109.08</b>	45.88	154.96	AMAZON-02, US	AS9394	<b>-17.17</b>	17.18	0.01	China TieTong, CN
AS8075	<b>13.27</b>	52.40	65.67	Microsoft, US	AS1239	<b>-9.46</b>	12.04	2.58	SPRINTLINK, US
AS174	<b>8.63</b>	27.27	35.90	Cogent, US	AS367	<b>-3.40</b>	15.73	12.33	DNIC, US
AS37963	<b>3.98</b>	8.08	12.05	AliBaba, CN	AS12576	<b>-3.28</b>	6.84	3.56	EE, GB
AS56	<b>3.69</b>	0.13	3.82	DNIC, US	AS17676	<b>-3.05</b>	43.48	40.43	SoftBank Corp., JP
AS3257	<b>1.84</b>	6.69	8.53	GTT, US	AS7029	<b>-2.78</b>	13.94	11.16	Windstream, US
AS14618	<b>1.69</b>	16.51	18.20	AMAZON-AES, US	AS4565	<b>-1.79</b>	1.79	0.00	Megapath, US
AS38091	<b>1.38</b>	0.11	1.48	HelloVision LG, KR	AS33771	<b>-1.57</b>	3.35	1.77	SafariCom, KE
AS6167	<b>1.28</b>	12.66	13.94	CELLCO, US	AS10455	<b>-1.57</b>	2.43	0.85	LUCENT, US
AS6730	<b>1.17</b>	0.96	2.12	SUNRISE, CH	AS18566	<b>-1.53</b>	1.53	0.00	Megapath, US

Table 3 – IPv4 Advertised Address span – Top 10 ASes

The picture of IPv4 growth for the year 2024 is dominated by Amazon. AS16509 expanded from 9,471 routes spanning 45.88M addresses to 12,478 routes, spanning 154.96M addresses. The next largest was Microsoft, AS8075, growing from 52.40M addresses to 65.76M addresses. It is interesting that 7 of the top 10 entities whose advertised address span grew in 2024 are US-based. There are some drops in the advertised address count from the ISP sector, including Sprintlink, EE, Softbank and Windstream.

It's likely that we are seeing a number of factors at play behind these changes in the IPv4 network:

- The first is the saturation of many Internet markets across the globe, so that the amount of “green field” expansion is far lower than, say, a decade ago.
- Secondly, we are seeing considerable concentration on the service market, where the level of utilization of addresses is vastly greater by both content and service publishers and by end clients. The service and client numbers may be growing, but that does not necessarily imply the use of more IPv4 addresses or more routing table entries. NATs are fully entrenched in the IPv4 world, and in the service provider market volume economics implies that few larger providers are more efficient than a greater number of smaller providers.
- Thirdly, this concentration in the service market has been accompanied by further consolidation in the access market, particularly in mobile access networks. This consolidation of client access networks creates greater efficiencies in shared address solutions.
- The continued deployment of IPv6 cannot be ignored. Within the 10 economies with the largest span of advertised addresses (collectively, these 10 economies advertise 75% of the span of advertised IPv4 addresses) 7 of these economies are also in the 10 economies with the largest span of advertised IPv6 addresses (collectively, these same 10 economies advertise 76.5% of the span of advertised IPv6 addresses). Looking at just these 7 economies, namely the United States, China, Japan, Germany, the United Kingdom, Brazil, and Canada, they advertise 76.4% of the entire advertised IPv6 address span and 71% of the advertised IPv4 address span. The level of IPv6 use in all 7 economies have a greater use of IPv6 deployment than the total global deployment level of 39.7% of end users.
- Finally, I suspect that the most significant dissipating factor has been the inexorable rise of private content distribution networks (CDNs). The conventional drivers for growth in the routing table were based on the product of an increasing number of networks with discrete routing policies and the need for these networks to balance incoming traffic across multiple network paths (traffic engineering using more specifics). The rapid uptake of CDNs by service platforms has meant a significantly reduced level of dependence on BGP transit services to deliver content to users. The public part of the Internet is rapidly shrinking to the last mile access network, which connects directly to a number of CDNs. The transit part of the Internet is just not so critical these days. It is unlikely that we will move away from the extensive use of CDNs in the foreseeable future, and if anything, their level of use within the Internet's service portfolio will increase further.

In various parts of the network the number of IPv4 entries in the default-free zone is between 950,000 and 1,000,000 entries. The net routing table growth in 2024 was some 53,000 new entries, a significant increase over the 3,000 new entries seen in 2023. A net increase of 1,400 new AS numbers were seen in the IPv4 network across the year, compared to 1,100 in 2023. While 2023 saw a slowing of the IPv4 network growth, 2024 has seen the growth numbers resume to the levels last seen in 2021 and 2022.

### **The IPv6 BGP Table Data**

A similar exercise has been undertaken for IPv6 routing data. As with the IPv4 network, there is diversity in the number of IPv6 routes seen at various vantage points, as shown when looking at the prefix counts advertised by all the peers of RouteViews (Figure 13).

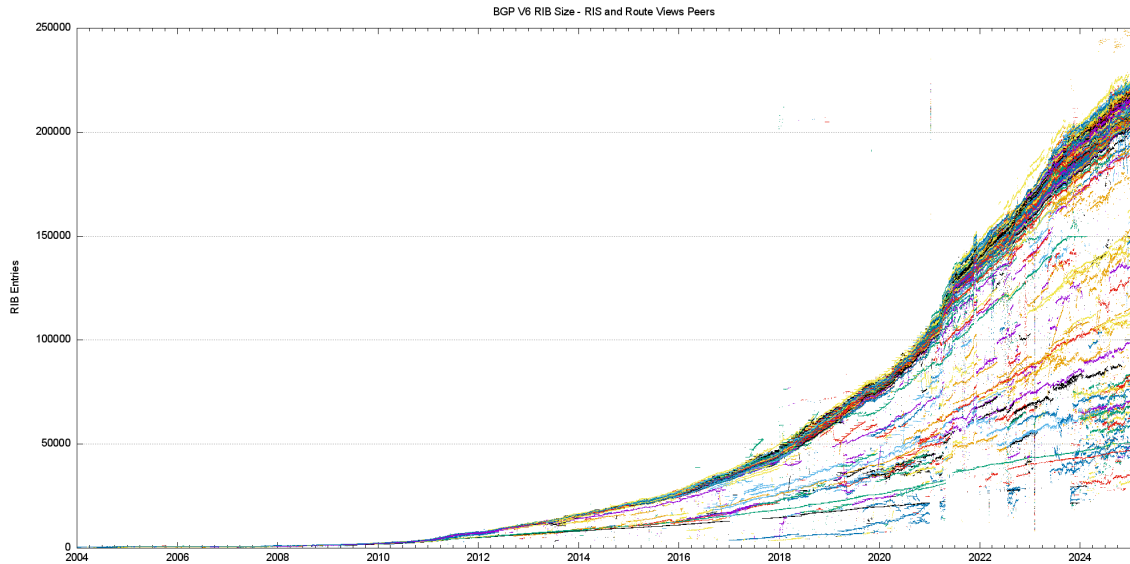


Figure 13 – IPv6 routing table since 2004 as seen by RIS and RouteViews peers

There are a number of distinct phases in the growth trends that are visible in this history of the IPv6 routing table. The period between 2004 and mid 2018 could be modelled by an exponential growth function with a doubling interval of three years. There was a sharp increase across late 2021, and again in mid 2023.

A more detailed look at the most recent four years incorporating both RouteViews and RIS data (Figure 14) shows some increasing diversity between various BGP views as to what constitutes the “complete” IPv6 route set, and the variance at the end of 2024 now spans some 25,000 prefix advertisements.




Figure 14 – IPv6 routing table 2021 - 2024 as seen by RouteViews and RIS peers

Figure 15 shows a detailed view of the routing table as seen by each of the RouteViews and RIS peers across 2024.




Figure 15 – IPv6 routing table across 2024, as seen by RouteViews and RIS peers

The comparable profile figures for the IPv6 Internet are shown in Figures 16 through 25.




Figure 16 - IPv6 BGP Routing Table Size (RIB)




Figure 17 – Announced Address Span (/32s)




Figure 18 - IPv6 More Specific Prefix Advertisements

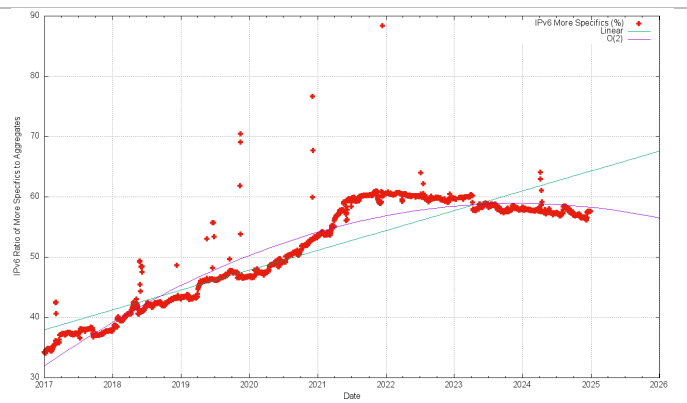


Figure 19 – IPv6 Relative Proportion of More Specific Announcements




Figure 20 - IPv6 Average Prefix Size (Prefix Size)




Figure 21 - IPv6 Prefix Size Counts




Figure 22 IPv6 Average AS Path Length




Figure 23 - IPv6 AS Count




Figure 24 - IPv6 Transit AS Count




Figure 25 - IPv6 AS Connectivity Degree

The growth of the IPv6 routing table has been relative steady across 2024, growing by some 20,000 routes across the year (Figure 16). Routing advertisements of /48s are by far the most prevalent prefix size in the IPv6 routing table, and some 46% of all prefixes are /48's. A total of 75% of the IPv6 table entries are composed of /48, /32, /44, and /40 prefixes (Figure 21).

RIR allocations of IPv6 addresses show a different pattern, with 74% of the 66,7854 IPv6 address allocations recorded in the RIRs' registries are either a /32 (47%) or a /29 (24%). Only 20% of allocations are a /48. What is evident is that there is no clear correlation between an IPv6 address allocation prefix size (as used by the address registries in the address allocation process) and the advertised address prefix size. Many IPv6 address holders do not advertise their entire allocated IPv6 address prefix in a single routing advertisement.

Why is the IPv6 routing table being fragmented so extensively? The conventional response is that this is due to the use of more specific route entries to perform traffic engineering. Another possible reason is the use of



more specifics to counter efforts of route hijacking. This latter rationale also has some credibility issues, given that it appears that most networks appear to accept a /64 prefix, and the disaggregated prefix is typically a /48, so as a countermeasure for more specific route hijacks, advertising /48's may not be all that effective.

This brings up the related topic of the minimum accepted route object size. The common convention in IPv4 is that a /24 prefix advertisement is the smallest address block will propagate across the entire IPv4 default-free zone. More complex minimum size rules have largely fallen into disuse as address trading appears to have sliced up many of the larger address blocks into smaller sizes. If a /24 is the minimum accepted route prefix size in IPv4, what is the comparable size in IPv6? There appears to be no common consensus position here, and the default action many network operators appears have no minimum size filter at all. In theory, that would imply that a /128 route object would be accepted across the entire IPv6 default-free zone, but a more pragmatic observation is that a /32 would be assuredly accepted by all networks, and it appears that many network operators believe that a /48 is also generally accepted. Given that a /48 is the most common prefix size in today's IPv6 network this view appears to be widespread. However, we also see prefixes smaller in size than a /48 in the routing table with /49, /52, /56 and /64 prefixes present in the IPv6 BGP routing table. Some 0.7% all advertised prefixes are more specific than a /48.

The summary of the IPv6 BGP routing table profile for period 2020 through to the start of 2025 is shown in Table 4. The IPv6 network growth rate is somewhat lower than previous years, with a 10% growth in routing entries, and a 4% growth in the advertised address span.

	Jan-21	Jan-22	Jan-23	Jan-24	Jan-25	Annual Growth			
						2021	2022	2023	2024
<b>Prefix Count</b>	105,500	146,500	172,400	201,200	221,500	39%	18%	17%	10%
<b>Root Prefixes</b>	49,200	57,800	69,400	84,000	94,000	17%	20%	21%	12%
<b>More Specs</b>	56,300	88,700	103,000	117,200	127,500	58%	16%	14%	9%
<b>Address Span (/32s)</b>	132,000	142,300	157,000	155,000	161,000	8%	10%	-1%	4%
<b>AS Count</b>	21,400	28,140	30,430	32,500	34,360	31%	8%	7%	6%
<b>Transit Ases</b>	4,100	4,640	4,990	5,400	5,800	13%	8%	8%	7%
<b>Stub Ases</b>	17,300	23,500	25,440	27,100	28,560	36%	8%	7%	5%

Table 4 – IPv6 BGP Table Growth Profile

The pressures for further expansion on the IPv6 network appear to be more idiosyncratic for each market sector and region, rather than being expressed as a more general imperative. Where there is at present scant IPv6 adoption, as is the case in most of Africa, the Middle East, Eastern and Southern Europe, and the western part of Latin America, there is no apparent sense of urgency to make the shift (Figure 26). It would appear that the Internet market is largely a saturated one and the smaller pace of network growth in those regions appears, for the moment, be adequately accommodated in the continued use of IPv4 NATs.

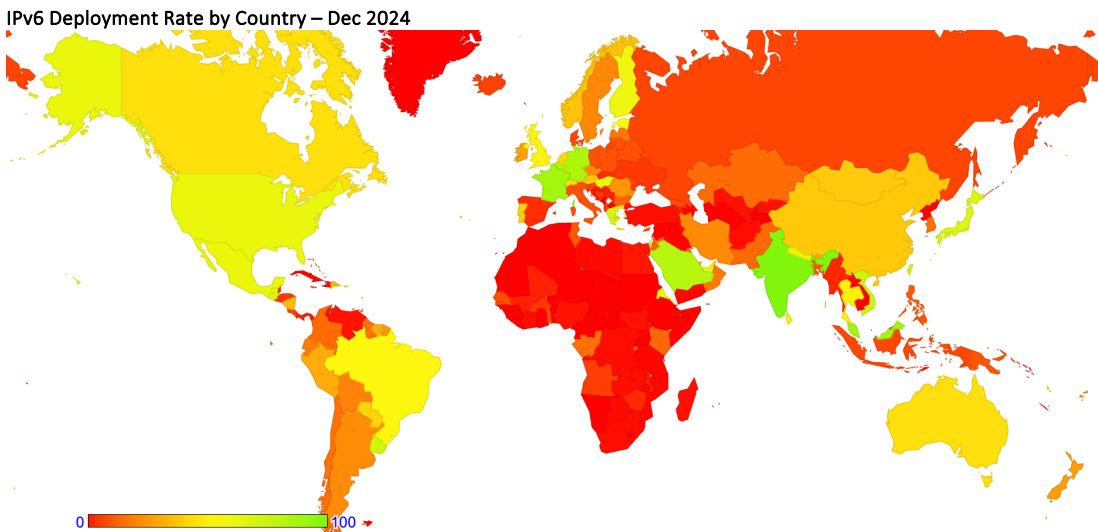


Figure 26 – IPv6 Adoption as of December 2024 (<https://stats.labs.apnic/net/ipv6>)

# The Predictions

What can this data from 2024 tell us in terms of projections of the future of BGP in terms of BGP table size?

## Forecasting the IPv4 BGP Table

Figure 27 shows the data set for BGP from January 2017 until December 2024. This plot also shows the fit of these most recent 5 years of data to various growth models. The first-order differential, or the rate of growth, of the BGP routing table is shown in Figure 28. The linear average rate of growth of the routing table appears to be falling slowly from 140 to 160 additional entries per day in 2016 to around 100 per day at the start of 2024.

There are a number of potential models to match this data. One model is to take the five-year average daily rate of change and apply this as a continuous model for the next five years. This is a “linear” model and takes the current dynamics of the IPv4 Internet, making the assumption that these dynamics will operate largely unchanged over the projection period. The second model is to look at the trend in the changes of rate of change and match this to a linear model. If the first order differential of a data series is a linear function, then the original data can be represented as a second order polynomial. The final model used here is to model the log of the data series as a linear model and therefore derive an exponential model for the data series. The application of these three projection models to the original data series is shown in Figure 27, and the first order differential of the data (the daily rate of change) is shown in Figure 28.




Figure 27 – IPv4 BGP Table 2017 – 2024




Figure 28 - First Order Differential of Smoothed IPv4 BGP Table Size – 2017 - 2024

Figure 28 appears to support the supposition of the primary data series fitting a second order polynomial model. The rate of change of the daily change to the count of IPv4 routing objects has been declining since 2020, where the range of daily change has fallen from a band of values between 100 to 200 new entries per day to a new range that sits in a band between 50 to 100 new entries per day for the first half of 2024. This rate has risen above 100 new entries per day for the second half of 2024.

The projections of the linear and polynomial best fit models are shown in Table 4 (the resumption of an exponential growth model appears to be highly unlikely in this late phase of the IPv4 network).

	IPv4 Table	Projection	
		Linear	O(2) Poly
Jan-19	760,000		
Jan-20	814,000		
Jan-21	857,000		
Jan-22	906,000		
Jan-23	942,000		
Jan-24	944,000		
Jan-25	996,000	1,014,313	980,293
Jan-26		1,057,234	994,518
Jan-27		1,100,154	1,002,370
Jan-28		1,143,075	1,003,849
Jan-29		1,186,113	998,933
Jan-30		1,229,033	987,649
Jan-31		1,271,954	969,991

Table 5 – IPv4 BGP Table Size Prediction

The linear projection model is somewhat unlikely, in my opinion. The drivers for continued growth of the IPv4 network do not appear to be clearly evident, so the projection of continued growth of the number of IPv4 FIB entries with an annual net gain of 43,000 entries is somewhat unrealistic. The O(2) polynomial projection model predicts that this first order differential will reach the zero point by early 2028 and then decline. It must be stressed it's just a mathematical model that fits the recent data, and nothing more.

Given that that last “normal” year of supply of available IPv4 address to fuel continued growth in the IPv4 Internet was now some fifteen years ago in 2010, perhaps the more relevant question is: Why has the growth of the IPv4 routing table persisted with such regularity in the ensuing fifteen years?

It should be remembered that a dual-stack Internet is not the objective in this time of transitioning the Internet to IPv6. The ultimate objective of the entire transition process is to support an IPv6-only network. An important part of the process is the protocol negotiation strategy used by dual-stack applications, where IPv6 is the preferred protocol wherever reasonably possible. In a world of ubiquitous dual-stack deployment all applications will prefer to use IPv6, and the expectation is that in such a world the use of IPv4 would rapidly plummet.

The challenge for the past decade or more has been in attempting to predict when in time that tipping point that causes demand for IPv4 to plummet may occur.

### Forecasting the IPv6 BGP Table

The same technique can be used for the IPv6 routing table. Figure 29 shows the data set for BGP from January 2017 until December 2024.




Figure 29 – IPv6 BGP Table Size from January 2017

The first order differential, or the rate of growth of the IPv6 BGP routing table is shown in Figure 30. The number of additional routing entries has grown from 10 new entries per day at the start of 2012 to a peak of some 230 new entries per day in May 2021.

The first half of both 2023 and 2024 had a steady period of increase in the growth of the daily count of new entries, and this declined in the second half of the year. Current growth rates are some 50-100 new entries per day, which is more than half the rate of growth in the IPv4 network (Figure 30).




Figure 30 - First Order Differential of IPv6 BGP Table Size

It appears that the period of exponential growth of the IPv6 network is over, and the best fit of a model to the recent data is either a linear model or an  $O(2)$  polynomial, as shown in Figure 31.




Figure 31 - Projections of IPv6 BGP Table Size

The projection for the IPv6 table size is shown in Table 6.

	IPv6 Table	Projection		
		Linear	O(2) Poly	Exp.
Jan-19	62,000			
Jan-20	79,000			
Jan-21	107,000			
Jan-22	147,000			
Jan-23	172,000			
Jan-24	201,000			
Jan-25	222,000	222,000	236,000	271,000
Jan-26		248,000	273,000	347,000
Jan-27		274,000	314,000	443,000
Jan-28		299,000	357,000	566,000
Jan-29		325,000	402,000	724,000
Jan-30		351,000	451,000	925,000
Jan-31		377,000	501,000	1,182,000

Table 6 – IPv6 BGP Table Size Prediction

The linear and polynomial projections in Table 5 provide a reasonable estimate of the low and high bounds of the growth of the IPv6 BGP routing table in the coming years.

If IPv6 continues to grow exponentially over the next five years, doubling every 30 months or so, then the size of the IPv6 routing table will be a little over one million entries at the start of 2030. The data from the previous two years suggests that such a level of growth is extremely unlikely, and a linear growth model is a closer fit to the recent past, and an average growth rate of 26,000 new entries per year is a better fit to this recent data.

## Conclusion

These predictions for the routing system are highly uncertain. The correlation between network deployments and routing advertisements has been disrupted by the hiatus in supply of IPv4 addresses, causing more recent deployments to make extensive use of various forms of address sharing technologies, and making fundamental alterations to the architecture of the service model of the Internet.

While a number of access providers and service platforms have made significant progress in public IPv6 deployments for their respective customers, the majority of the Internet user base (some 60% the Internet's user base) is still exclusively using IPv4 as of the end of 2024 (Figure 32).





Figure 32 - IPv6 Deployment 2012 - 2024

These predictions as to the future profile of the routing environment for IPv4 and IPv6, using extrapolation from historical data, can only go so far in providing a coherent picture for the near-term future. As well as the technical issues relating to the evolution of IP technology and the IPv6 transition there are also broader factors such as the state of the global communications economy and the larger global economy. Investment in communications infrastructure, as with most other forms of infrastructure investment is not generally a short-term proposition. The major benefits tend to be realised in increased efficiency of economic production, rather than short-term windfall gains from infrastructure investment. This means that short term expedient measures, such as a response to a global pandemic or a rapid escalation of energy prices due to regional conflict, can interrupt infrastructure investment programs. The question behind the recent slowing of the growth in both the IPv4 and IPv6 aspects of the Internet's routing space is whether this slowdown is due to market saturation in the case of IPv4 or a dissipation of collective market impetus in the case of IPv6, or an interruption due to these short-term exogenous market factors. In the latter case we would expect growth to resume once more when the current global market conditions dissipate, while an underlying condition of market saturation is a more permanent state.

If the concern is that the routing system is growing at a rate that is faster than our collective ability to throw available technology at it, then there is no serious cause for alarm in the current trends of growth in the routing system. There is no evidence of the imminent collapse of BGP. Far from it!

However, the size of the inter-domain routing table is only one half of the story. The stability of the routing system is also very important, and to complete this look at the routing system in 2024 we will also need to look at the dynamic behaviour of the routing system. The profile of BGP update churn in 2024 is a topic we'll look at in detail in the next article that looks at the larger picture of addressing and routing across 2024.

---

## 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](http://www.potaroo.net)*