

January 2023
Geoff Huston

BGP in 2022 – The Routing Table

At the start of each year, I've been reporting 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.

This past year marks a significant point in the evolution of the Internet where the strong growth numbers that were a constant feature of the past thirty years are simply not present in the data. The Internet's growth is slowing down significantly. Have we got to the point of market saturation and there is no more demand capacity to fuel further growth? Or are we reeling from the combinations of a global pandemic, turmoil in energy markets and the signs of increased climate instability so that we are no longer as interested to throw more resources into more network infrastructure investment? Is the Internet losing its investment appeal along with so many other signals of investor disillusion over outlandish growth potential in technology-based services? Let's take a look at the BGP view of 2022 and see how these larger economic and social 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. This is a lot like a social rumour network, where every individual who hears a new rumour immediately informs all their friends. BGP works in a very similar fashion: each time a neighbour informs a BGP speaker about reachability to 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 information provides a better path to the prefix, then the local speaker moves this prefix and associated next-hop forwarding decision to the local forwarding table and informs all its immediate neighbours of a new path to a prefix, implicitly citing itself as the next hop. In addition, there is a withdrawal mechanism, where a BGP speaker determines that it no longer has a viable path to a given prefix, in which case it announces a "withdrawal" to all its neighbours. When a BGP speaker receives a withdrawal, it stores the withdrawal against this neighbour. If the withdrawn neighbour happened to be the currently preferred next hop 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.

What could possibly go wrong?

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 a large 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 tree with 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 *content-addressable memory* to bypass this serial decision limitation. *Ternary content-addressable memory* (TCAM) can search its entire contents in a single memory cycle. It's fast, but it's also a very expensive component 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 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.

Secondly, there is 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 we are interested in are the size of the routing space and its level of updates, or "churn". Here we will look at the first of these metrics, the size of the routing space, and the changes that occurred through 2021, 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 measurement point being used here is a BGP speaker configured within AS131072. This AS generates no traffic and originates no routes in BGP. It's a passive measurement point that has been logging all received BGP updates since 2007. The router is fed with a default-free eBGP feed from AS4608, which is the APNIC network located in Australia, and AS4777, 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 that iBGP is a major contributor to BGP scalability concerns in BGP, the consideration here in trying to measure 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, its 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.

The IPv4 Routing Table

Measurements of the size of the routing table have been taken regularly since the start of 1988, although regular highly detailed snapshots of the routing system only date back to early 1994. Figure 1 shows a rather unique picture of the size of the routing table, as seen by all the peers of the Route Views route collector on an hourly basis.

I should take a moment to mention the Route Views Project (<http://www.routeviews.org/routeviews/>). It 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 Route Views so unique is that it archives these routing tables every two hours and has done so for more than two decades. It also archives every BGP update message. This vast collection of data is a valuable research data source in its own right, and here we are just taking a tiny slice of this data set to look at longer term routing growth trends.

The folk at the Route Views Project, with the support from the University of Oregon and the US National Science Foundation should be commended for their efforts here. This is a very unique data set if you are interested in the evolution of the Internet over the years.

Several broader events are visible in the history of the routing table, such as the busting of the Internet bubble in 2001, and if one looks closely, the effects of the global financial crisis in 2009. 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 free 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.

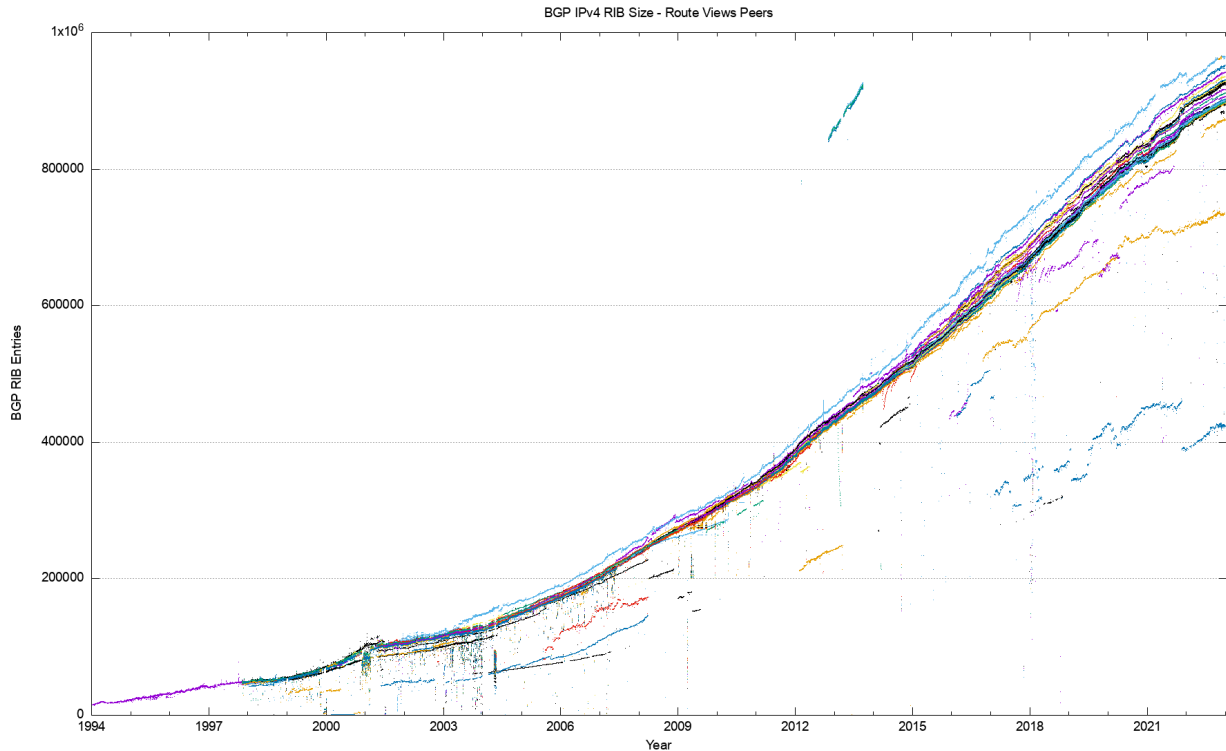


Figure 1 – IPv4 routing table since 1994 as seen by Route Views peers

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 2, where the BGP routing table counts from both the Route Views peers and the peers of the RIPE NCC's Routing Information Service (RIS) are combined. There is not a single plot in this figure where each BGP speak sees essentially the same network. There is a variance across the various peers of these route collectors that is around 50,000 routes.

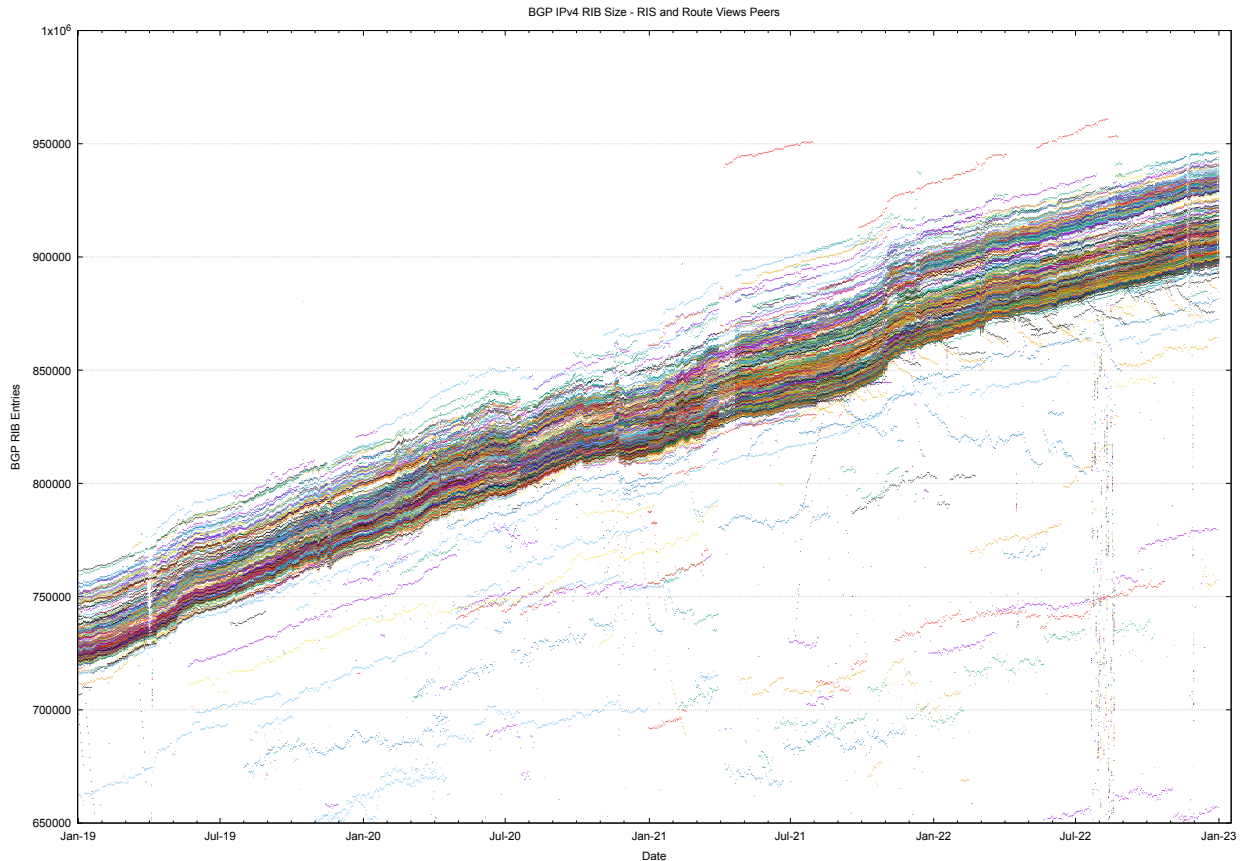


Figure 2 – IPv4 routing table 2019-2022, as seen by Route Views and RIS peers

Figure 2 illustrates an important principle in BGP, that there is no single authoritative view of the Internet’s inter-domain routing table – all views are relative to the perspective of each BGP speaker. It also illustrates 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.

And finally, in this overview of the BGP system Figure 3 looks at the year 2022 in terms of the number of BGP routing entries, using a 6-hour sample set for the year. The picture for the year is in many ways unremarkable. The average number of routing entries in the IPv4 network rose from some 880,000 at the start of the year to 915,000, a rise of some 35,000 entries. Across the peers of the Route Views and RIS peers the variance in the size of the routing table is some 50,000 entries, and this level of variable has been generally steady through the year.

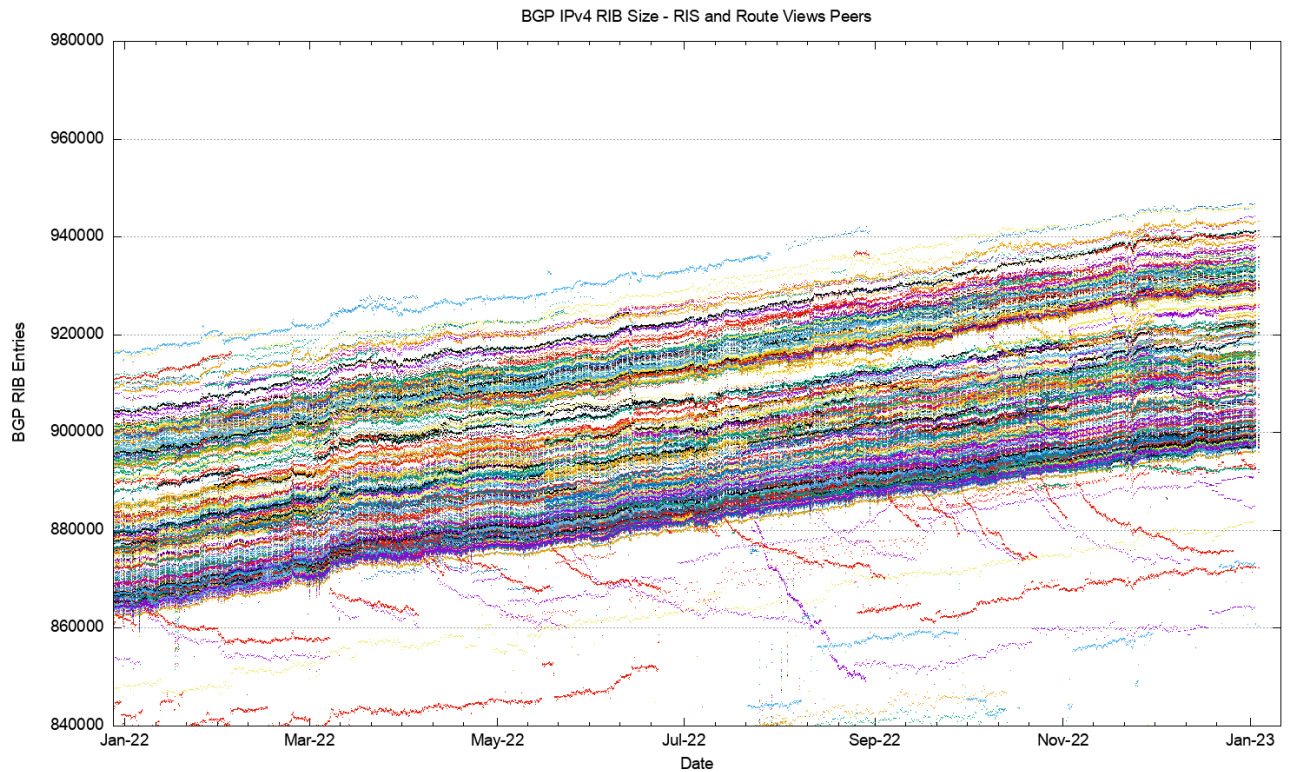


Figure 3 – IPv4 routing table across 2022, as seen by Route Views and RIS peers

The next collection of plots (Figures 4 through 13) contains some of the vital characteristics for the IPv4 BGP network since the start of 2016 to the end of 2022.




Figure 4 - IPv4 BGP Routing Table Size (RIB)




Figure 5 - IPv4 Announced Address Span




Figure 6 - IPv4 More Specific Announcements




Figure 7 - IPv4 Relative Proportion of More Specific Announcements

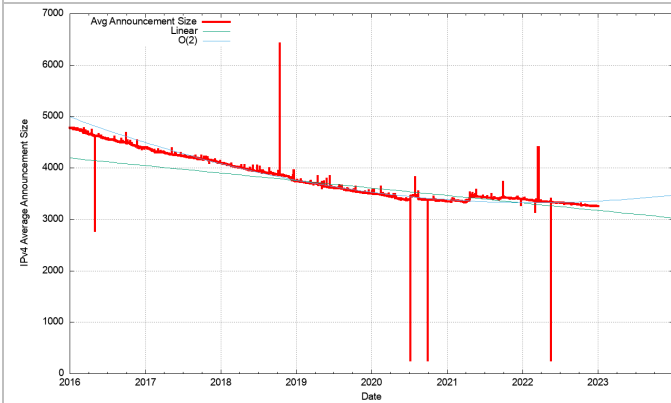


Figure 8 - IPv4 Average Announcement Size

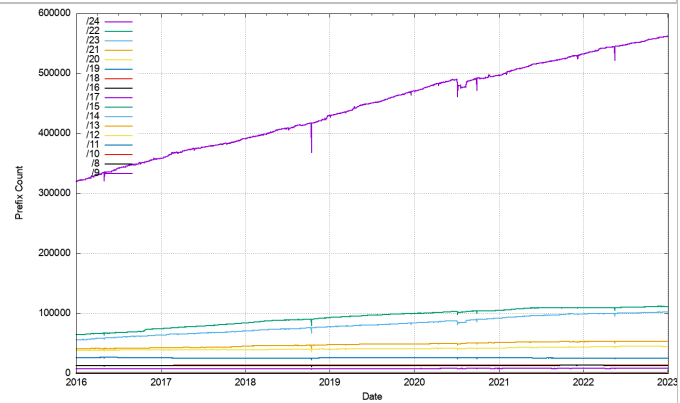


Figure 9 - IPv4 Prefix Size Counts

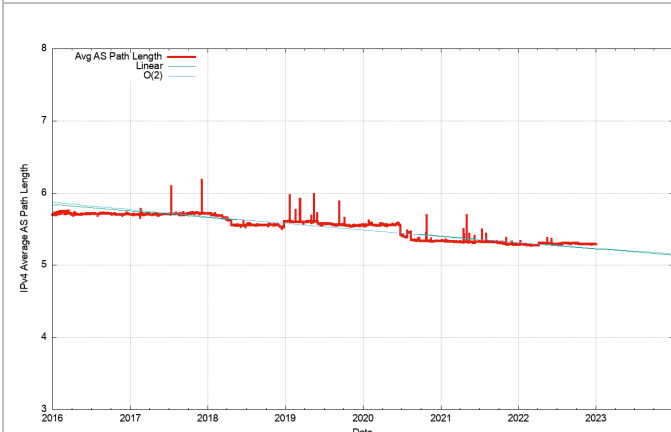


Figure 10 - IPv4 Average AS Path Length




Figure 11 - IPv4 AS Count

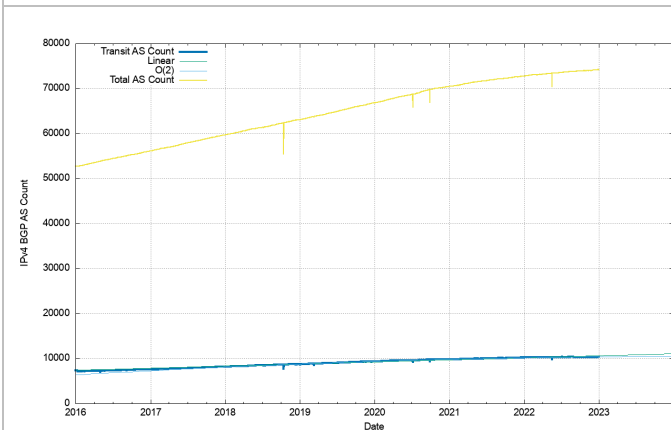


Figure 12 - Transit vs Total AS Count

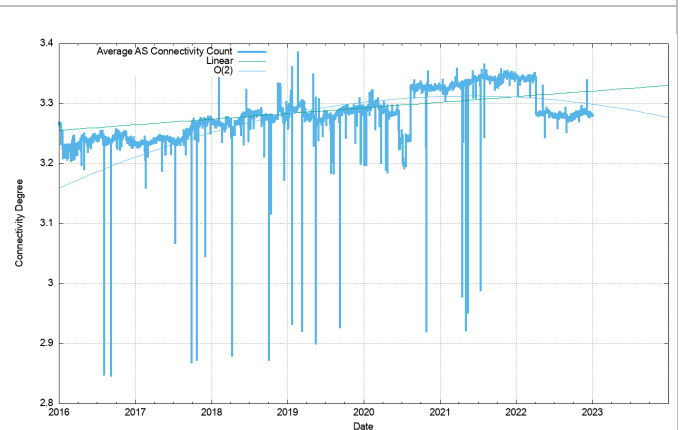


Figure 13 - AS Connectivity Degree

Figure 4 shows the total number of routes in the routing table over this period. This is a classic "up and to the right" Internet trajectory, but it should be noted that growth trends in the Internet over 2022 are much lower

than previous years. It seems that the growth of the Internet, in terms of the size of the IPv4 routing table is slowing down. We'll return to this data when we look at projections for the size of the BGP FIB table later in this article.

The seven-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 5). 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 the total span of advertised IPv4 addresses declined.

In terms of more specific advertisements and covering aggregate advertisements, both counts increased across 2022 (Figure 6), but as shown when looking at the ratio of the two counts, the count of more specifics grew more slowly than the count of covering aggregates, as the ratio declined over 2022 (Figure 7). As the number of routing entries increased, while the total advertised address span decreased the average address span of an IPv4 route continued to decrease across 2022, and the average prefix size is now slightly smaller than a /20. (Figure 8). Prefixes sizes of /24, /23 and /22 now account for 80% of the entire IPv4 routing table (Figure 9).

The topology of the network has remained relatively consistent, with the growth of the Internet being seen as increasing density of interconnectivity, rather than through extending transit paths. This trend was not evident in 2022, and the average AS Path length was constant through the year (Figure 10).

The clearest evidence of the slowing of the growth in the IPv4 internet can be seen in the AS count. The growth of the AS count started to decline in late 2020 and continued to decline in 2021 and 2022. This is a likely signal of network saturation in many markets (Figure 11). The number of transit networks was held constant across 2022, which appears to be a related signal of market saturation (Figure 12).

The year-by-year summary of the IPv4 BGP network over the 2019-2022 period is shown in Table 1.

	Jan-19	Jan-20	Jan-21	Jan-22	Jan-23	Annual Growth			
						2019	2020	2021	2022
Prefix Count	760,000	814,000	860,000	906,000	940,000	7%	6%	5%	4%
Root Prefixes	353,000	387,000	400,000	420,000	445,000	10%	3%	5%	6%
More Specs	407,000	427,000	460,000	486,000	495,000	5%	8%	6%	2%
Address Span (/8s)	169.3	169.8	171.4	183.3	182.8	0%	1%	7%	0%
AS Count	63,100	66,800	70,400	72,800	74,200	6%	5%	3%	2%
Transit AS	9,000	9,600	10,200	10,800	10,800	7%	6%	6%	0%
Stub AS	54,100	57,200	60,200	62,000	63,400	6%	5%	3%	2%

Table 1 – IPv4 BGP Table Growth Profile

In terms of advertised prefixes, the size of the routing table continues to grow, but the 4% recorded through 2022 is lower than the numbers seen for the previous three years, and this declining growth is seen on the other metrics of the IPv4 routing space. This observation supports a sub-linear growth model of the routing table size, with a growth rate of on average of 93 prefixes per day for the year. The number of routed Stub AS numbers (new edge networks) grew by 2% in 2022, which is one third of the growth rate of 2 – 3 years earlier.

These metrics show that the drivers for growth in the IPv4 BGP network are declining. It's likely that we are seeing a number of factors at play. The first is the saturation of many Internet markets, 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. 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. Finally, the continued deployment of IPv6 cannot be ignored. Within the 10 economies with the largest span of advertised addresses (collectively, these 10 economies advertise 74% of the span of advertised IPv4

addresses) 6 of these economies are also in the 10 economies with the largest span of advertised IPv6 addresses (collectively, these same 10 economies advertise 70% of the span of advertised IPv6 addresses). Looking at just these six economies, namely the United States, China, Japan, Germany, Brazil, and the United Kingdom, they advertise 69% of the entire advertised IPv6 address span and 65% of the advertised IPv4 address span. The level of IPv6 use in 5 of these 6 economies have a greater use of IPv6 deployment than the total global deployment level of 33% of end users, while the deployment levels in China is slightly below this global average.

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 Route Views (Figure 14).




Figure 14 – IPv6 routing table since 2004 as seen by Route Views 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. The period from mid-2019 to mid-2021 saw an increase in the growth function, where the doubling interval fell to slightly under two years. However, from mid-2021 the growth pattern has changed again and is now best modelled by a linear growth model with a growth rate of some 25,000 additional route entries per year, or an average of 68 new routing entries per day. This is still lower than the IPv4 routing table growth of some 35,000 IPv4 routes per year.

A more detailed look at the most recent four years incorporating both Route Views and RIS data (Figure 15) shows some increasing diversity between various BGP views as to what constitutes the “complete” IPv6 route set, and the variance at the end of 2022 now spans some 10,000 prefix advertisements.




Figure 15 – IPv6 routing table 2019 - 2022 as seen by Route Views and RIS peers

Figure 16 shows a detailed view of the routing table as seen by each of the Route Views and RIS peers across 2022.




Figure 16 – IPv6 routing table across 2021, as seen by Route Views and RIS peers

The comparable profile figures for the IPv6 Internet are shown in Figures 17 through 26.




Figure 17 - IPv6 BGP Routing Table Size (RIB)




Figure 18 - Announced Address Span (/32s)




Figure 19 - IPv6 More Specific Prefix Advertisements




Figure 20 - IPv6 Relative Proportion of More Specific Announcements

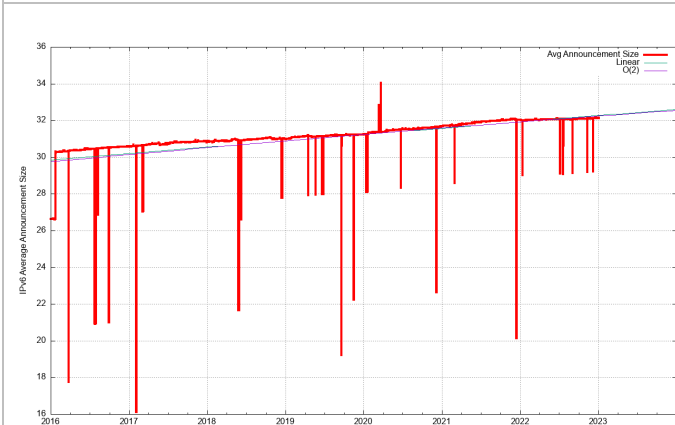


Figure 21 - IPv6 Average Prefix Size (Prefix Size)

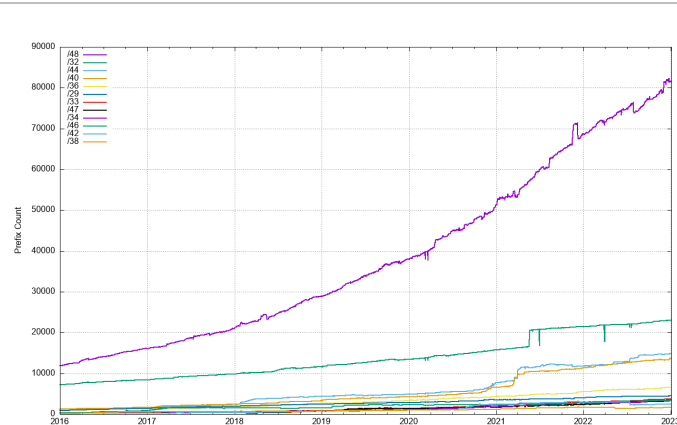


Figure 22 - IPv6 Prefix Size Counts




Figure 23 IPv6 Average AS Path Length


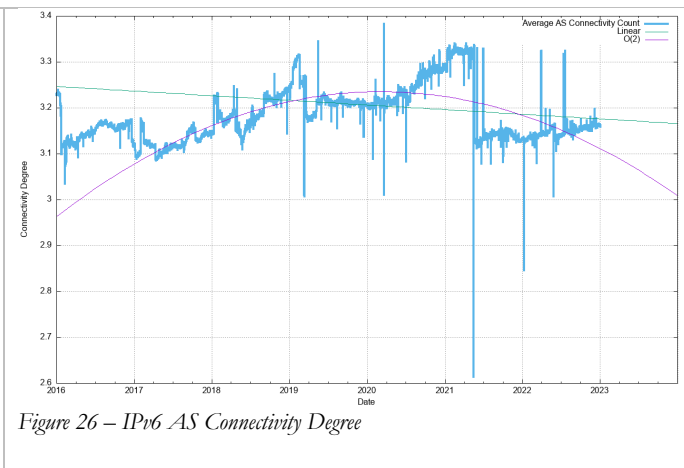
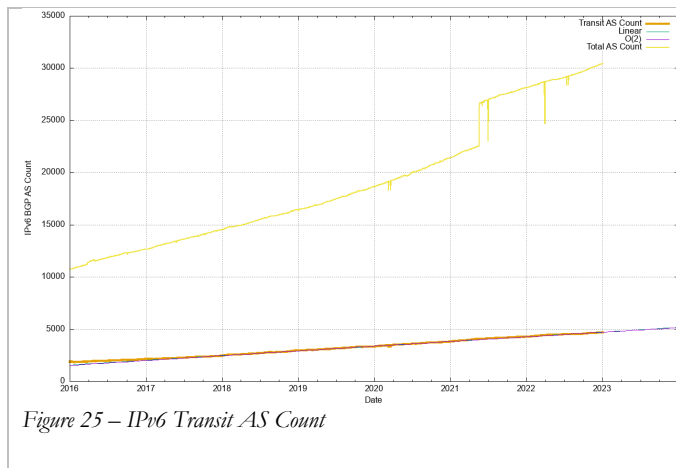


Figure 24 - IPv6 AS Count



The growth of the IPv6 routing table has been relative steady across 2022, growing by some 25,000 routes across 2022 (Figure 12). Routing advertisements of /48s are the most prevalent prefix size in the IPv6 routing table at the end of 2022, and some 47% of all prefixes are /48's. A total of 78% of the IPv6 table entries are composed of /48, /32, /44, and /40 prefixes (Figure 22).

RIR allocations of IPv6 addresses show a different pattern, with 74% of address allocations are either a /32 (49%) or a /29 (25%). Only 19% 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 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 be slicing 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 is to use no minimum size filter. In theory that would imply that a /128 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 belief appears to be the case. 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. Slightly less than 1% of all advertised prefixes are more specific than a /48.

The summary of the IPv6 BGP routing table profile for period 2018 through to the start of 2023 is shown in Table2. As is the case in IPv4, the IPv6 network growth rate has dropped significantly in 2022, and the overall growth rate is between a half and a third of the growth rates seen in 2021. Given that the overall extent of the IPv6 network encompasses approximately one third of the IPv4 network it is hard to argue that the IPv6 network is now a saturated market. What is more likely is that the appetite for further deployment of IPv6 has lessened in the past 12 months.

	Jan-19	Jan-20	Jan-21	Jan-22	Jan-23	Annual Growth			
						2019	2020	2021	2022
Prefix Count	62,400	79,400	105,500	146,500	172,400	27%	33%	39%	18%
Root Prefixes	35,400	42,300	49,200	57,800	69,400	19%	16%	17%	20%
More Specifics	27,000	37,100	56,300	88,700	103,000	37%	52%	58%	16%
Address Span (/32s)	124,900	133,800	132,000	142,300	157,000	7%	-1%	8%	10%
AS Count	16,470	18,650	21,400	28,140	30,430	13%	15%	31%	8%
Transit AS Count	3,190	3,590	4,100	4,580	4,990	13%	14%	12%	9%
Stub AS Count	13,280	15,060	17,300	23,560	25,440	13%	15%	36%	8%

Table 2 – IPv6 BGP Table Growth Profile

The Predictions

What can this data from 2022 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 2016 until December 2022. 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 to around 100. The current best fit to the data is an order(2) polynomial function that predicts that the growth in the IPv4 routed entry count will stop by 2028.




Figure 27 – IPv4 BGP Table 2016 – 2022




Figure 28 - First Order Differential of Smoothed IPv4 BGP Table Size – 2016 - 2022

However, the data since 2016 in Figure 28 suggests a different trend, where the growth rate is slowly declining to a rate at the end of 2022 of 100 new entries per day. What is the change to these projections if we use just the most recent 24 months to generate a prediction model for the IPv4 routing table size? This is shown in Figure 29.

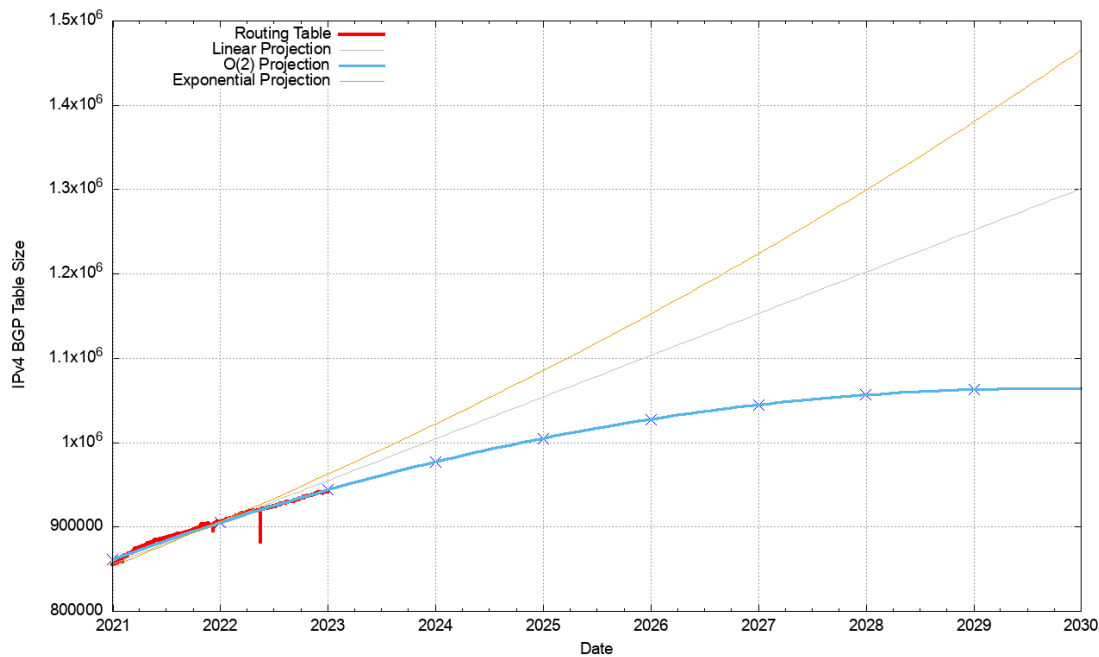


Figure 29 – IPv4 BGP Table Predictions

The dominant trend in the data series is the continuous decline in growth, and the best fit of a simple model to the data is an order(2) polynomial function. This model predicts a peak size of the IPv4 BGP table of some 1,070,000 entries in mid 2029. However, it must be stressed it’s just a mathematical model that fits the data, and nothing more.

This model can be used to predict that the IPv4 routing table in the coming years. Some five years from now, at the start of 2028, the model predicts that the IPv4 BGP routing table will contain an additional 115,000 entries, making a total for IPv4 of some 1,057,000 entries in the BGP IPv4 routing table at that time.

	IPv4 Table	IPv4 Prediction
Jan-18	699,000	
Jan-19	760,000	
Jan-20	814,000	
Jan-21	866,000	
Jan-22	906,000	
Jan-23	942,000	944,000
Jan-24		977,000
Jan-25		1,005,000
Jan-26		1,028,000
Jan-27		1,045,000
Jan-28		1,057,000
Jan-29		1,062,000

Table 3 – IPv4 BGP Table Size Prediction

It’s challenging to ascribe a degree of certainty into this prediction. Given that that last ‘normal’ year of supply of available IPv4 address to fuel continued growth in the IPv4 Internet was now over a decade 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 decade?

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. The assumption behind these predictions, and similar predictions that have been made over the past twenty years, is that such a tipping point is at least five more years in the future from the time of the prediction. This may not be a reasonable assumption, but it’s been our informal working mode of operation through this period.

Forecasting the IPv6 BGP Table

The same technique can be used for the IPv6 routing table. Figure 30 shows the data set for BGP from January 2016 until December 2022.




Figure 30 – IPv6 BGP Table Size from January 2016

The first order differential, or the rate of growth of the IPv6 BGP routing table is shown in Figure 31. 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 2022 had a steady period of growth of between 50 and 70 new entries per day, while the data for the second half of 2022 shows an increase to a little over 100 new entries per day by the end of the year. (Figure 31).




Figure 31 -First Order Differential of IPv6 BGP Table Size

There are two visible trends in the IPv6 data. The 5-year trend can be fitted against an exponential growth model with a doubling interval of a little under three years. The more recent 2-year trend can be fitted against

a linear growth model The fit of a linear projection and an exponential model of projected IPv6 table size is shown in Figure 32.




Figure 32 -Projections of IPv6 BGP Table Size

The projections for the IPv6 table size are shown in Table 4.

	IPv6 Table	IPv6 Prediction	IPv6 Prediction
		Linear	Exponential
Jan-18	45,000		
Jan-19	62,000		
Jan-20	79,000		
Jan-21	106,000		
Jan-22	147,000		
Jan-23	172,000	172,000	185,000
Jan-24		199,000	243,000
Jan-25		226,000	320,000
Jan-26		252,000	421,000
Jan-27		279,000	554,000
Jan-28		305,000	729,999
Jan-29		332,000	959,000

Table 4 – IPv6 BGP Table Size Prediction

The linear and exponential projections in Table 4 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 under one million entries at the start of 2029. The data from the previous two years suggests that such a level of growth is 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.

What is also unknown is the extent to which the slowing growth rate for IPv6 in the past 24 months is related to the COVID-19 pandemic, and potentially the economic side effects of the turmoil in global energy markets due to the Russian invasion of the Ukraine. The communications market is not immune from such global effects.

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 two thirds of the Internet's user base) is still exclusively using IPv4 as of the end of 2022 (Figure 33).




Figure 33 - IPv6 Deployment 2012 - 2022

The predictions as to the future profile of the routing environment for IPv4 and IPv6 that use 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. 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.

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 2022 we will also need to look at the dynamic behaviour of the routing system. The profile of BGP update churn in 2022 is a topic we'll look at in detail in the next article that looks at addressing and routing across 2022.

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