

January 2017  
Geoff Huston

### BGP in 2016

It has become either a tradition, or a habit, each January for me to report on the experience with the inter-domain routing system over the past year, 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.

One reason why we are interested in the behaviour of the routing system is that at its heart the system has no natural bounding function. Our collective unease about routing relates to a potential scenario where every network decided to deaggregate their prefixes and announce only the most specific prefixes, or where every network applied routing configurations that were inherently unstable and rapidly reverted into oscillating states that generated an unending set of routing updates into the BGP realm. In such scenarios, the routing protocol we use, the Border Gateway Protocol, or BGP, will not help us by attempting to damp down the anomalies. Indeed, there is a very real prospect that in such scenarios the protocol behaviour of BGP could well amplify the behaviour!

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 new 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. 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, indicating that it no longer can reach this prefix.

And that's BGP.

What could possibly go wrong?

The first is the sheer 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 may not sound all that challenging until you do some basic calculations and work out that at 100Gbps (which is not uncommon these days) that means that a single such "wire" could present one valid 64 octet IP packet every 5 nanoseconds. Performing a lookup into a data structure of around one million entries for an imprecise match

of a 32-bit value within 5 nanoseconds represents an extremely challenging silicon design problem. The larger the search space, the harder the problem!

Secondly, there is the 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. The router will start to lag in real time, so that the information it is propagating reflects a past local topology, not necessarily the current local topology. If this lag continues then at some point updates may be dropped from the queue. BGP has no inherent periodic refresh capability, so when information is dropped 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 TTL expires. This may cause saturation of the underlying transmission system and trigger further outages which, in turn, may add to the routing load.

So, the critical metrics we are interested in are the size of the routing space and its level of update, or "churn".

### **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 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 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 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 objectively measure this assertion is that there is no "standard" iBGP configuration, and 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.

## **The Data**

### **The IPv4 Routing Table**

Measurements of the size of the routing table have been taken on a regular basis since the start of 1988, although 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. Several events are visible in the plot, such as the introduction of "Classless Inter-Domain Routing" (CIDR) in 1994 dramatically that changed the scaling properties of both the routing system and the IPv4 address space, and appears to have added more than 20 years to the effective lifetime of IPv4. Also, the Internet's boom and bust is visible in the period 1999 – 2003, as is 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 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.




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, and 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 I've combined the BGP RIB counts from both the Route Views peers and the peers of the RIPE NCC's Routing Information Service (RIS). There are two distinct bands in this plot, the upper band is the Route Views peers, and the lower band is generated by the RIS peers. The RIS peers, which are predominately located in Europe, appear to have 30,000 fewer prefixes, and cluster more tightly around their mean as compared to the set of Route Views peers.




Figure 2 – IPv4 routing table 2015-2016, 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 routing table – all views are in fact relative to the perspective of the 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. And thirdly, this illustrates the prevalent use of more specifics to affect traffic engineering. It is often the case that these more specifics are advertised

with a limited scope, and if the changes to the transit arrangements move a BGP speaker in or out of this scope, then one can expect changes in the set of visible routes as a consequence.

The issue of route leaks and the advertisement of more specifics in the routing system could be seen as an instance of a “tragedy of the commons,” ([http://en.wikipedia.org/wiki/Tragedy\\_of\\_the\\_commons](http://en.wikipedia.org/wiki/Tragedy_of_the_commons)) where the self-interest of one actor in attempting to minimise its transit service costs becomes an incremental cost in 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 BGP realm, where there is an extensive reliance on enlightened self-interest to be conservative in one’s own announcements, and the actions by a smaller set of actors are prominent because they fall well outside of the conventional “norm” of inter-domain routing practices.

The next collection of plots (Figures 3 through 8) show some of the vital statistics for IPv4 in BGP since the start of 2011 to the end of 2016.




Figure 3 - IPv4 BGP Routing Table Size (RIB)




Figure 4 - IPv4 More Specific Entries

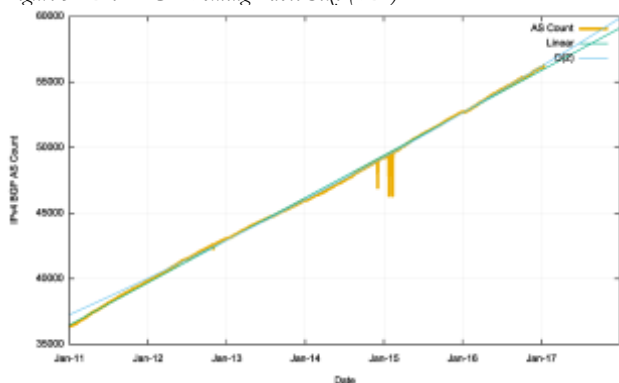


Figure 5 - IPv4 AS Count

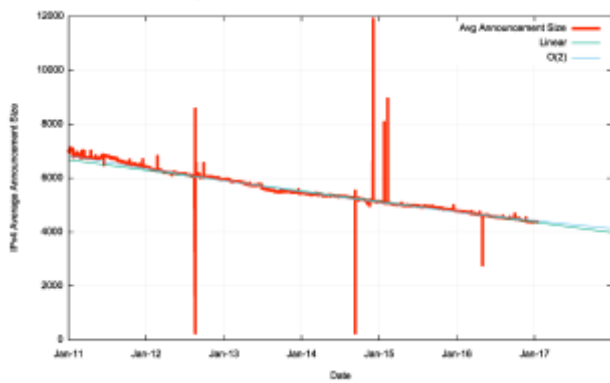


Figure 6 - Average Announcement Size




Figure 7 - IPv4 Advertised Address Space



Figure 8 - IPv4 Average AS Path Length

Figure 3 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 today are more strongly aligned to a far more modest linear growth model than the initial exponential growth models.

Over this period, we had the exhaustion of the IPv4 address space pools in IANA in January 2011, APNIC in April 2011 (serving the Asia Pacific region), in the RIPE NCC in September 2012 (serving Europe and the Middle East), LACNIC in May 2014 (serving Latin America and the Caribbean), and ARIN in September 2015 (serving North America). The five year period since the start of 2011 has seen the span of addresses advertised in the routing system slowing down (Figure 7). However, at the same time there has been a consistent level of growth in the number of entries in the routing table over the same period. The result of these two factors is that the average announcement in the IPv4 routing table is spanning fewer addresses, or, to put it another way, the granularity of the IPv4 routing space is getting finer. As Figure 6 shows, the average BGP announcement size has dropped from 7,000 host addresses at the start of 2010 to 4,200 addresses at the end of 2016. While /24 announcements are steady at a little over 50%, the relative number of /22 announcements is increasing, while the relative number of larger announcements, including up to /21s, are decreasing. The topology of the network has remained relatively consistent, with the growth in the Internet being seen as increasing density of interconnectivity, rather than through extending transit paths, so the average AS path length has remained relatively constant at 5.8 for this period for this observation AS (Figure 8).

The summary of the IPv4 BGP network over the 2014-2016 period is shown in Table 1.

	Jan-14	Jan-15	Jan-16	Jan-17	2014 growth	2015 growth	2016 growth
<b>Prefix Count</b>	488,000	530,000	587,000	646,000	9%	11%	10%
<b>Root Prefixes</b>	237,000	257,000	281,000	304,000	8%	9%	8%
<b>More Specs</b>	251,000	287,000	306,000	342,000	14%	7%	12%
<b>Address Span (/8s)</b>	158.6	162.1	167.2	169.0	2%	3%	1%
<b>AS Count</b>	46,000	49,000	52,700	56,100	7%	8%	6%
<b>Transit AS</b>	6,400	7,000	7,600	7,800	9%	9%	3%
<b>Stub AS</b>	39,000	42,000	45,100	48,300	8%	7%	7%

Table 1 – IPv4 BGP Table Growth Profile

In terms of prefix count the size of the routing table continues to grow at a rate of some 10% p.a. The number of routed stub AS numbers (new edge networks) grew by 7% in 2016, similar to the growth rate of the previous two years. The effects of increasing scarcity of IPv4 addresses is evident, with the span of advertised networks growing by just 1% through 2016. It appears that the drivers for growth in the IPv4 network in 2016 continued at a pace that is similar to that of the previous 24 months. However, as the address supply is slowing down the compensatory move is that the address space being divided up into smaller units, and presumably this routing change is accompanied by the increasing use of IPv4 Network Address Translation to accommodate the network's growth pressures.

The overall conclusions from this collection of observations is that the IPv4 network continues to grow, but as the supply of new addresses is slowing down, what is now becoming evident is more efficient use of addresses, which results in the granularity of the IPv4 inter-domain routing system becoming finer. The density of inter-AS interconnection continues to increase. The growth of the Internet is not "growth from the edge" as the network is not getting any larger in terms of average AS path change. Instead, the growth is happening by increasing the density of the network by attaching new networks into the existing transit structure and peering at established exchange points. This makes for a network whose diameter, measured in AS hops, is essentially static, yet whose density, measured in terms of prefix count, AS interconnectivity and AS Path diversity, continues to increase. This denser mesh of interconnectivity could be potentially problematical in terms of convergence times if the BGP routing system used a dense mesh of peer connectivity, but the topology of the network continues along a clustered hub and spoke model, where a small number of transit ASs directly

service a large number of stub edge networks. This implies that the performance of BGP in terms of time and updates required to reach convergence continues to be relatively static.

### The IPv6 BGP Table Data

A similar exercise has been undertaken for IPv6 routing data. There is a considerable diversity in the number of routes seen at various vantage points in the Internet, as shown when looking at the prefix counts advertised by all the peers of Route Views (Figure 9).

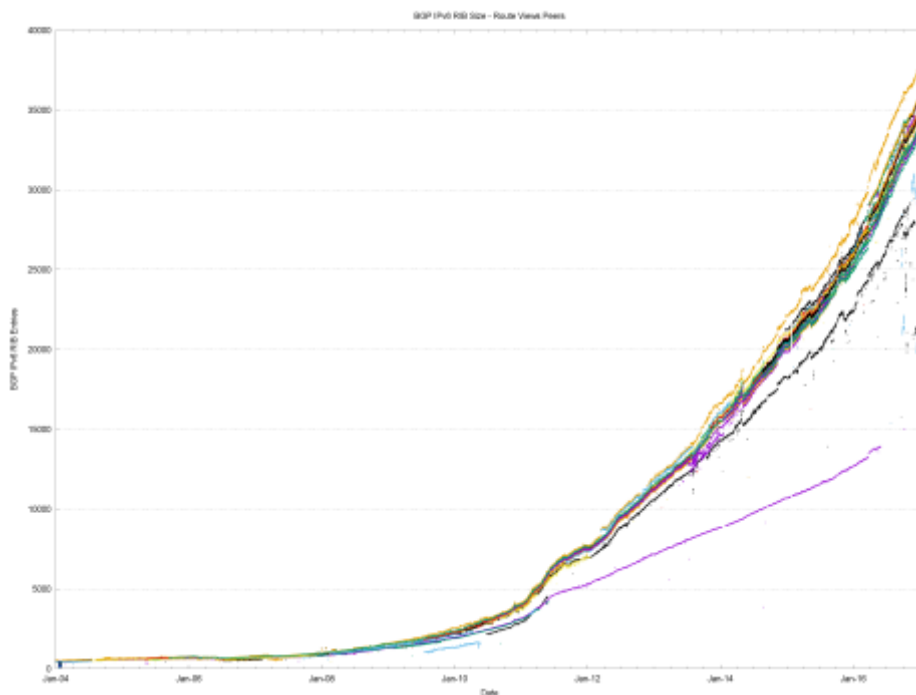


Figure 9 – IPv6 routing table since 2004 as seen by Route Views Peers

A more detailed look at 2015 and 2016 incorporating both Route Views and RIS (Figure 10) shows that in IPv6 there is no visible disparity in the route sets announced by RIA peers as compared to Route Views peers.




Figure 10 – IPv6 routing table 2015 - 2016 as seen by Route Views and RIS Peers

The comparable figures for the IPv6 Internet are shown in Figures 11 through 16.




Figure 11 - IPv6 BGP Routing Table Size (RIB)




Figure 12 - IPv6 More Specific Entries




Figure 13 - IPv6 AS Count

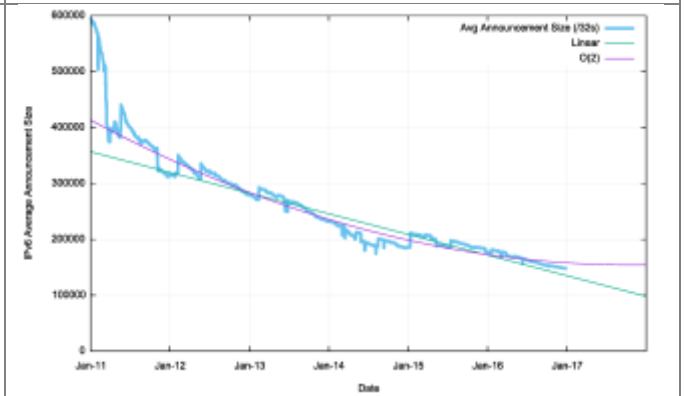


Figure 14 - Average Announcement Size




Figure 15 - IPv6 Advertised Address Space




Figure 16 - IPv6 Average AS Path Length

	Jan-14	Jan-15	Jan-16	Jan-17	2014 growth	2015 growth	2016 growth
<b>Prefix Count</b>	16,700	21,000	27,200	34,800	26%	30%	28%
<b>Root Prefixes</b>	11,400	14,600	17,800	22,900	28%	22%	29%
<b>More Specifics</b>	5,300	6,400	9,400	11,900	21%	47%	27%
<b>Address Span (/32s)</b>	56,000	58,200	71,000	76,600	4%	22%	8%
<b>AS Count</b>	7,900	9,100	10,700	12,700	15%	18%	19%
<b>Transit AS Count</b>	1,600	1,700	2,000	2,400	6%	18%	20%
<b>Stub AS Count</b>	6,300	7,400	8,700	10,300	17%	18%	18%

Table 2 - IPv6 BGP Table Growth Profile

The summary of the IPv6 BGP profile for period 2014 through to the start of 2017 is shown in Table 2. It appears that the last quarter of 2016 saw the growth in the IPv6 routing table fall below the growth levels of the first three quarters.

There are some interesting aspects to this growth where the characteristics of IPv4 look to be appearing in IPv6. The number of more specific advertisements in IPv6 is high, and is currently approximately 1/3 of the total number of advertised prefixes. This is less than the level in IPv4, but it is still high.




Figure 17 – IPv6 Table Growth in 2016

## BGP Update Stability

We will now look at the profile of BGP updates across 2016 to assess whether the stability of the routing system, as measured by the level of BGP update activity is changing.

### IPv4 Stability

Figure 18 shows the daily update activity for the eBGP session at AS131072, since mid-2007. There are some aspects of BGP behaviour that are visible in this plot that do not have an obvious cause.

The first of these is the number of withdrawals. Since 2007 the number of advertised prefixes has risen from 230,000 to 650,000, yet the number of observed withdrawals has remained relatively constant at some 15,000 withdrawals per day. There is no particular reason why the withdrawal count would be held steady while the number of announced prefixes has more than doubled.

The second is the number of update messages per day. This was steady at some 50,000 updates per day from 2007 until 2013. During 2013 the volume of updates doubled to some 100,000 updates per day, which it maintained for most of the ensuing 24 months. During 2016 the number of updates per day rose once more, and approached 150,000 updates per day by the end of the year. This is not exactly reassuring news. It has been fortuitous that the BGP update rate has been held steady for so many years, as this has implied that the capability of BGP systems has not required constant upgrades in terms of processing capability. In the same way that there was no clear understanding of why the BGP update rate was steady for so many years, it's also unclear why the rate has started to increase in 2016. This makes efforts to attempt to stem this upward trend rather uncertain.



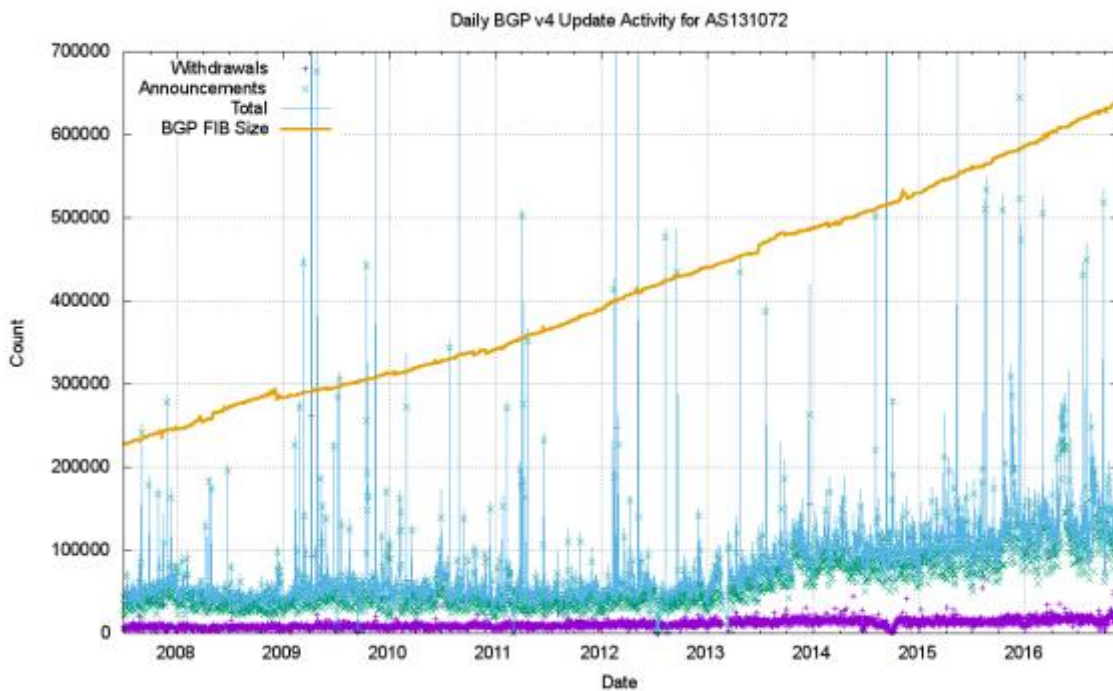


Figure 18 – IPv4 BGP update counts

What is also intriguing is that most of these 150,000 updates per day are generated from a pool of between 20,000 to 40,000 prefixes. A plot of the daily number of different prefixes that are the subject of BGP updates is shown in Figure 19. It is evident that this count has not substantially changes in 2016, so the heightened instability is due to more updates to reach convergence, rather than more unstable prefixes.




Figure 19 – IPv4 unstable prefixes per day

Where this higher instability level is evident is in the average time to reach convergence. Up until 2013 the average convergence time was around 70 – 75 seconds. Across 2013 this rose to 100 seconds, and has remained at this level through to the end of 2016 (Figure 20).

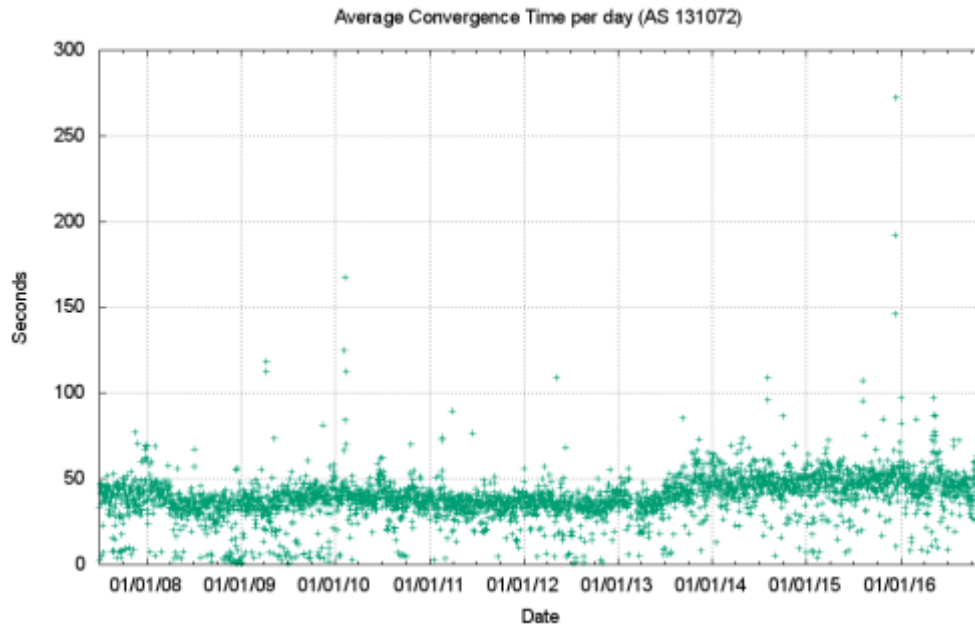


Figure 20 – IPv4 Average routing convergence time per day

However, the instability in BGP is not widespread. Half of all the BGP updates over a period of a week are attributed to less than 1% of the unstable prefixes, and just 50 origin ASNs accounted for 40% of these updates. It appears that the network is generally highly stable, and that a very small number of prefixes appear to be wedged into highly unstable configurations over periods that extend for weeks rather than hours. The cumulative distribution of BGP updates by prefix and by origin AS, shown in Figures 21 and 22 for the final week of 2016 shows the highly skewed nature of unstable prefixes in the routing system.




Figure 21 – Distribution of BGP Updates by Prefix

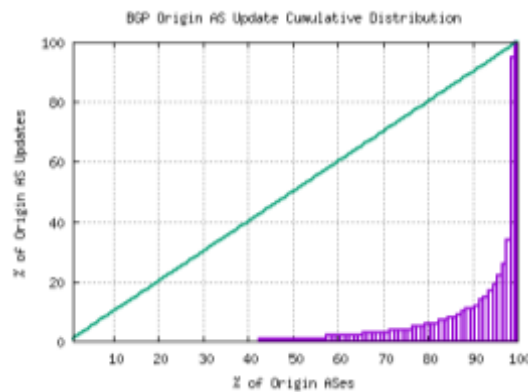


Figure 22 – Distribution of BGP Updates by Origin AS

### IPv6 Stability

The theory says that the IPv6 routing network should be behaving in a very similar manner to the IPv4 environment. It is a smaller network, but essentially similar in terms of topology. Is the stability profile of IPv6 the same as IPv4 or not?

Figure 23 shows the profile of Ipv6 updates since 2008. In common with IPv4, the number of withdrawals per day appears to be relatively constant at some 3,500 withdrawals per day. The number of updates per day is quite different from IPv4 in that the number has been rising since 2015, and appears to be clustered around a point of 40,000 updated per day with outlier days as high as 120,000 updates per day.

This is due to an increase in the number of updates to reach a converged state rather than any notable increase in the number of unstable prefixes (Figure 24).

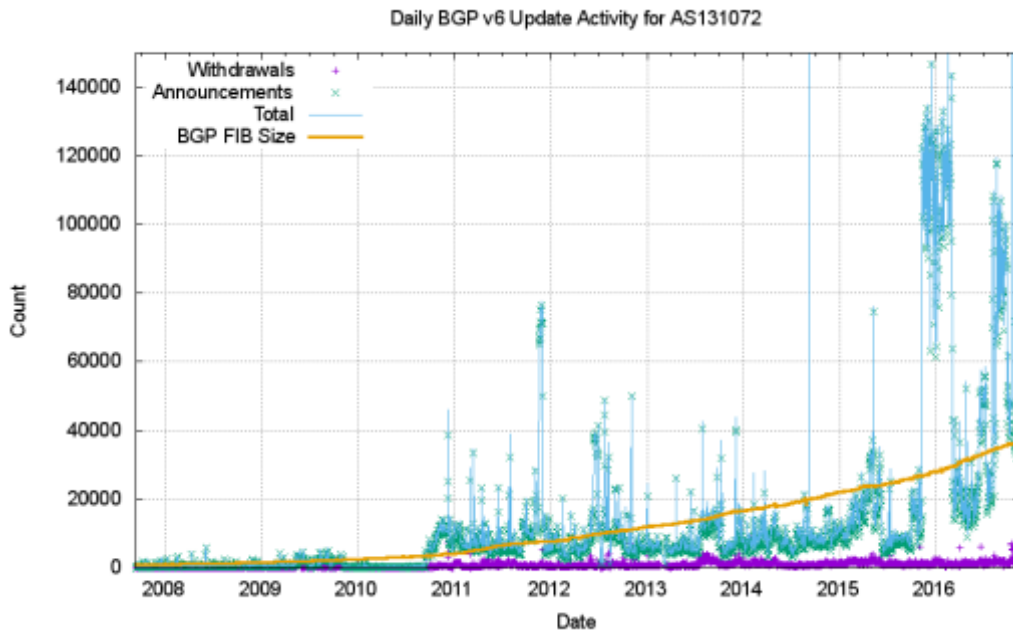


Figure 23 – IPv6 BGP update counts

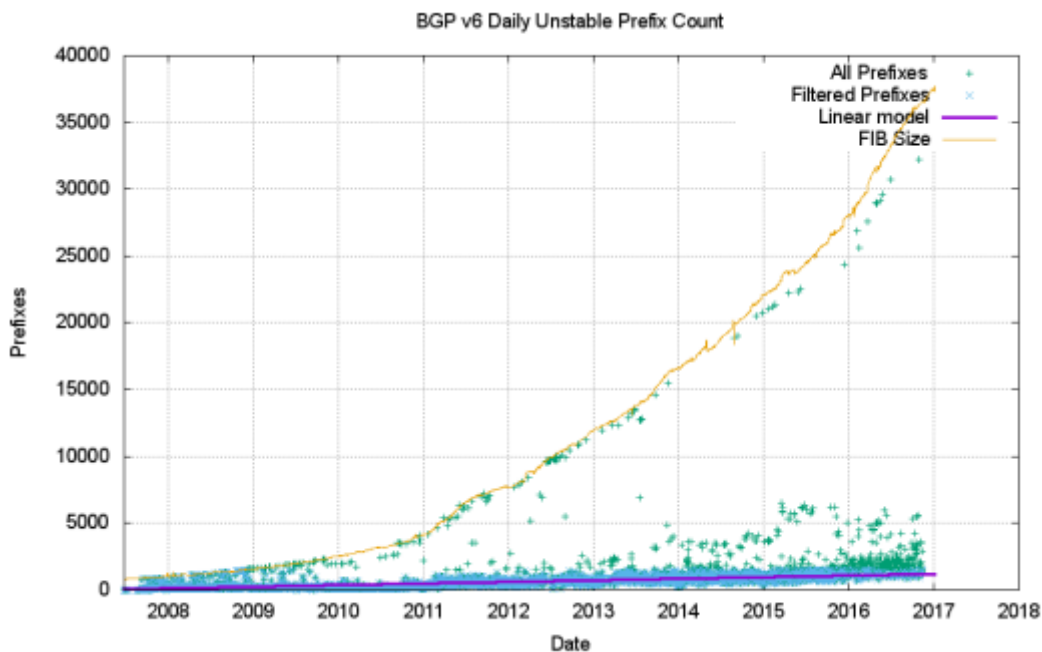


Figure 24 – Unstable IPv6 prefix count

This greater instability in the IPv6 network is also visible in the daily average time to reach convergence. Prior to 2009 the average convergence time was less than 50 seconds, which was comparable to the IPv4 network at the time. However, this metric has increased significantly since then, and average daily convergence times now range from a low point to some 40 seconds to some 300 seconds (Figure 25).

It is also evident that the distribution of updates across the set of announced prefixes and originating ASNs is more skewed than IPv4. The most unstable 50 IPv6 prefixes accounted for slightly more than one half of the total update volume, and the most unstable 50 Origin ASNs account for some 80% of the updates. The distribution of updates is shown in Figure 25 and 26.

It is not immediately obvious why IPv6 has this higher instability component than IPv4. A concern is that this instability remains a persistent condition as the IPv6 network continues to grow, which would create a routing environment that would impose a higher processing overhead than we had anticipated, with its attendant pressures on BGP processing capabilities in the network.

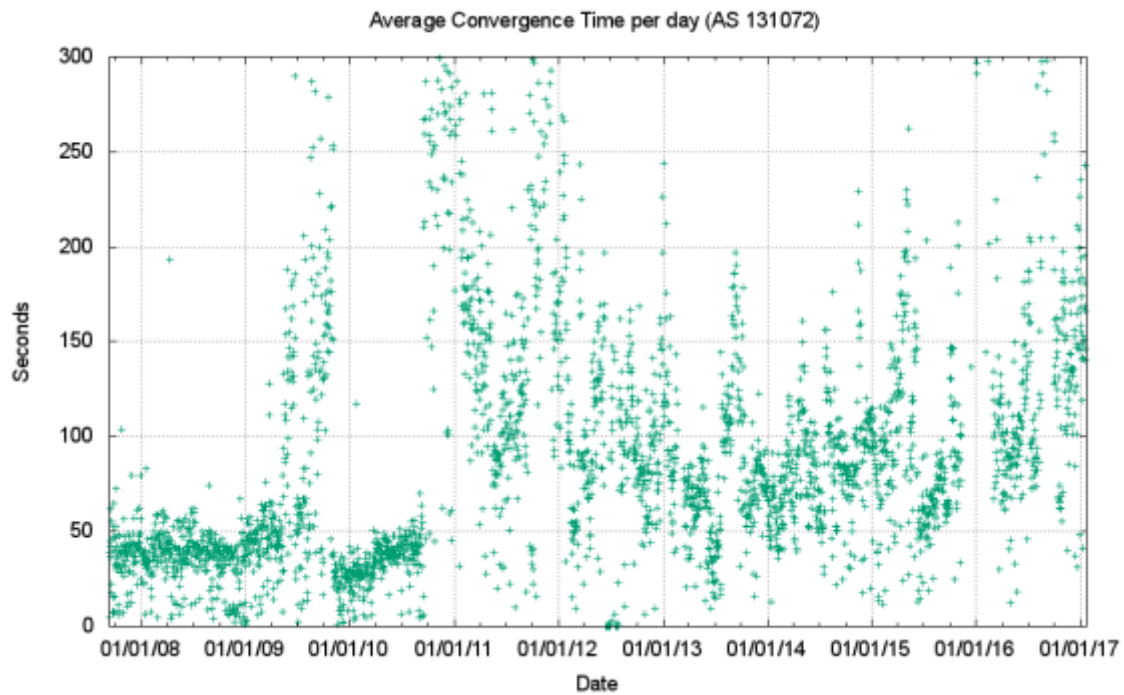


Figure 25 – IPv6 Average routing convergence time per day

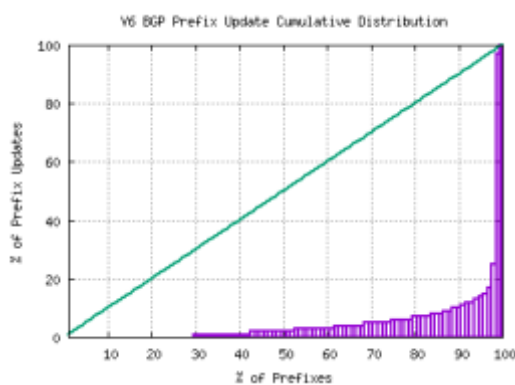


Figure 26 – Distribution of BGP IPv6 Updates by Prefix

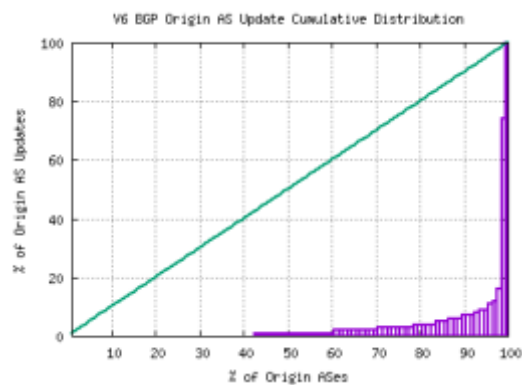


Figure 27 – Distribution of BGP IPv6 Updates by Origin AS

## The Predictions

What can this data from 2016 tell us in terms of projections of the future of BGP in terms of BGP table size?

At the outset, it's necessary to observe that this is a time of extreme uncertainty in the BGP prediction business! With this caveat that we are now heading deep into highly speculative areas, and the associated warning that the predictions being made here come with a very high level of uncertainty, let's look at the predictions for the Internet's routing system for the coming few years.

### Forecasting the IPv4 BGP Table

Figure 28 shows the data set for BGP from January 2014 until January 2017. This plot also shows the fit of these most recent 3 years of data to various models. The first order differential, or the rate of growth, of the BGP routing table is shown in Figure 29. The rate of growth of the routing table appears to be steady between 130 and 160 additional entries per day. If the first order differential matches a flat line, then the data set matches a linear increase model. The consequent prediction of IPv4 BGP table size using this constant growth model of 148 additional routing entries per day is shown in Table 3.




Figure 28 – IPv4 BGP Table 2014 - 2017




Figure 29 -First Order Differential of Smoothed IPv4 BGP Table Size

	IPv4 Table	IPv4 Prediction
Jan 2013	441,172	
Jan 2014	488,011	
Jan 2015	529,806	
Jan 2016	586,879	
Jan 2017	645,974	
Jan 2018		697,000
Jan 2019		751,000
Jan 2020		806,000
Jan 2021		861,000
Jan 2022		915,000

Table 3 – IPv4 BGP Table Size Prediction

With the caveat that this prediction assumes that tomorrow will be a lot like today and that the influences that shape tomorrow have already shaped today, then it's reasonable to predict that the IPv4 routing table five years from now, at the start of 2022, will contain an additional 270,000 entries, making a total for IPv4 of some 915,000 entries in the BGP routing table at that time.

It's difficult to portray this prediction as reasonable under the current circumstances. Given that that last 'normal' year of supply of available IPv4 address to fuel continued growth in the IPv4 Internet was 2010, why has the growth of the IPv4 routing table occurred with such regularity?

### A Second Look at IPv4 Routing Advertisements

These predictions for the size of the IPv4 BGP network growing by a continued 54,000 new routing entries per year assume that the near-term future will continue to play out much the same as the recent past. But, as we've noted, the issues related to IPv4 address exhaustion make this assumption somewhat implausible. So, let's take a more detailed look at IPv4 across 2016, and to do this I'll take a comparison of a snapshot of the routing table as of the start of 2016 to that at the end of the year.

At the start of the year the BGP routing table in AS131072 had 586,917 entries, and at the end of the year it had 646,059 entries. The routing table grew by 59,141 entries through the year – right? Mathematically that correct, but it's not the entire story. There were only 502,846 routing entries that were in both routing snapshots. Some 67,504 routes were present at the start of the year, but not at the end, and 126,645 entries that were in the end of year snapshot, but not at the beginning of the year.

The issue here is that BGP is not just used to glue the Internet together in a reachability sense. Routing is also the only tool we have to adjust the path taken by incoming traffic, so in a sense we could say that the routing table contains the cross product of reachability and routing policies. At any point in time there are a collection of "traffic engineering" prefixes, and a set of reachability prefixes. So can we differentiate between what appears to be the background of traffic engineering route changes from the routes that appear to be announcing reachability to previously unreachable addresses? One approach is to divide the routing table into "root" prefixes, that announce reachability, and "more specific" prefixes that refine this reachability for parts of this announced space. At the start of the year there were 286,249 root prefixes and 300,669 more specifics. At the end of the year there were 309,092 root prefixes and 336,947 more specifics. That's a net growth of 22,843 roots and 36,298 more specifics. The comparison table is shown in Table 4.

	Jan-16	Jan-17	Delta	Unchanged	Re-Home	Removed	Added
<b>Announcements</b>	586,918	646,059	59,141	502,846	16,928	67,504	126,645
<b>Root Prefixes</b>	286,249	309,092	22,843	252,411	10,803	22,080	46,238
<b>Address Span (/8s)</b>	156.35	158.40	2.04	147.31	2.52	5.58	8.57
<b>More Specifics</b>	300,669	336,967	36,298	250,435	6,125	45,424	80,407
<b>Address Count (/8s)</b>	51.86	56.04	4.18	47.06	0.81	4.94	8.17

Table 4 – IPv4 BGP changes over 2016

Figure 30 shows the relative age (as determined by the date of registration of the address) for addresses that were announced each year since 2010.

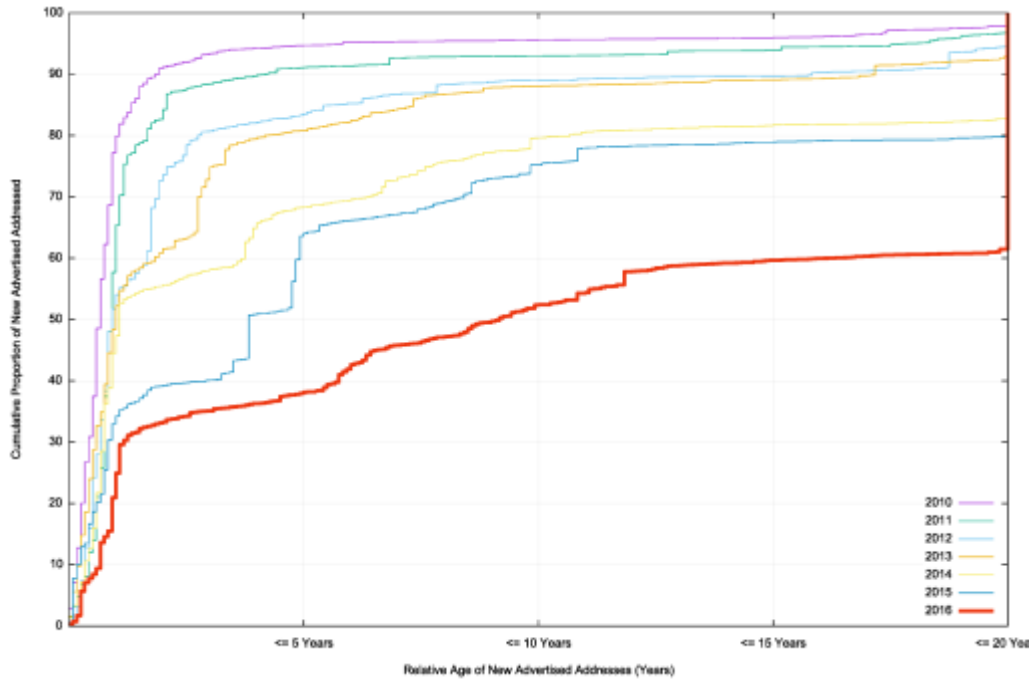


Figure 30 – Relative age profile of advertised addresses per year (from 2010)

In 2010, 80% of all newly advertised addresses in the year were allocated in the same year. Most of the growth in advertised addresses in 2010 came of new allocations of address space. But with the onset of address exhaustion from 2011 onward this level has dropped. In 2014 less than half the newly advertised addresses were allocated in 2014. In 2016 this figure has dropped to 30% of addresses that were added to the routing system were drawn from the remaining unallocated address pools in 2016. Another 30% were originally allocated since 1996, or within the last 20 years. The remaining 40% were part of the early allocations that form the “legacy” address pool.

This data tends to support a useful objective of address trading in a post-address exhaustion environment, namely to allow holders of otherwise idle or non-publically used IPv4 addresses to release them for use by current network operators.

Can we reasonably expect the IPv4 address table to reach 915,000 entries in five years from now? It is one possible scenario, but in so saying that, it also would imply that this protracted transition to IPv6 will extend over the next five years. Indeed, that scenario could be interpreted as a clear signal of failure in the overall transition to IPv6 by then. Are the prospects for IPv6’s adoption so poor? What can we see in the Internet’s routing data about the prospects for IPv6?

### Forecasting the IPv6 BGP Table

The same technique can be used for the IPv6 routing table. Figure 31 shows the data set for BGP from January 2014 until January 2017.

The first order differential, or the rate of growth of the IPv6 BGP routing table is shown in Figure 32. The number of additional routing entries has grown from 10 new entries per day at the start of 2014 to a peak of over 30 in early 2017, although this has declined to around 20 per day by the end of 2016. Obviously, this is far lower than the equivalent figure in the IPv4 domain, which is growing by some 150 new entries per day, but it does show a consistent level of increasing growth.

This implies that a linear growth model is inappropriate for modelling growth in IPv6. A better fit to the data is a compound growth model, with a doubling factor of some 22 months. It is possible to fit a linear model to the first order differential of the data, which can be used to derive an  $O(2)$  polynomial fit to the original data. The fit of a linear,  $O(2)$  polynomial and an exponential model against the data is also shown in Figure 31.




Figure 31 – IPv6 BGP Table Size from January 2014



Figure 32 – First Order Differential of IPv6 BGP Table Size

The projections for the IPv6 table size are shown in Table 5.

	IPv6 Table	IPv6 Prediction Linear	IPv6 Prediction Exponential
Jan 2013	11,600		
Jan 2014	16,158		
Jan 2015	20,976		
Jan 2016	27,241		
Jan 2017	37,469		
Jan 2018		43,244	49,619
Jan 2019		50,290	65,376
Jan 2020		57,335	86,136
Jan 2021		64,399	113,594
Jan 2022		71,445	149,639

Table 5 – 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. At this point these figures are not a cause for any significant concern.

It appears from these projections that for the next five years, the significantly larger size of the IPv4 network will continue to drive the overall costs of BGP routing, and the IPv6 BGP network will operate, in effect, in the margins of oversupply in meeting the demands of IPv4. It will be some time before there is significant change in the relativities of the two protocols from a routing perspective.

## Conclusions

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. In addition, there is still a set of confused signals relating to IPv6 adoption.

While a small number of providers have made significant progress in public IPv6 deployments for their respective customer base, the overall majority of the Internet is still exclusively using IPv4. This is despite the fact that among that small set of networks that have deployed IPv6 are some of the largest ISPs in the Internet! 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. Despite this uncertainty, nothing in this routing data indicates any serious cause for alarm in the current trends of growth in the routing system. There is no evidence of the imminent collapse of BGP.

None of the metrics indicate that we are seeing such an explosive level of growth in the routing system that it will fundamentally alter the viability of carrying a full BGP routing table anytime soon. In terms of the projections of table size in the IPv4 and IPv6 networks, the BGP sky is firmly well above us, and it's not about to fall on our heads just yet!

---

## Disclaimer

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

---

## Author

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

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