Routing and Addressing in 2017

Geoff Huston

Chief Scientist, APNIC

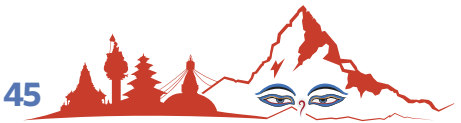

Through the Routing Lens …

There are very few ways to assemble a single view of the entire Internet

The lens of routing is one of the ways in which information relating to the entire reachable Internet is bought together

Even so, its not a perfect lens…

24 Years of Routing the Internet

This is a view pulled together from each of the routing peers of Route-Views

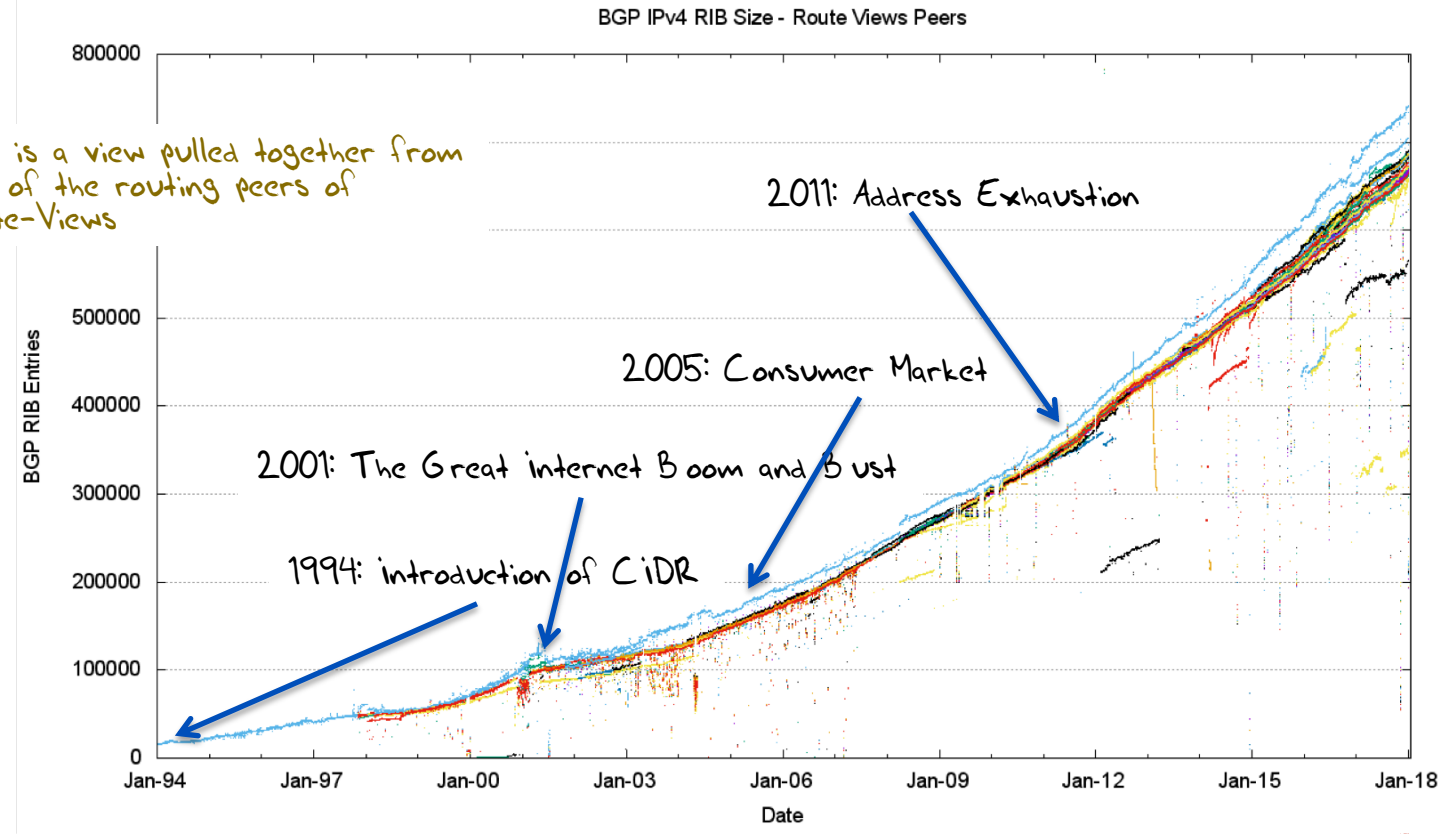

2016-2017 in detail

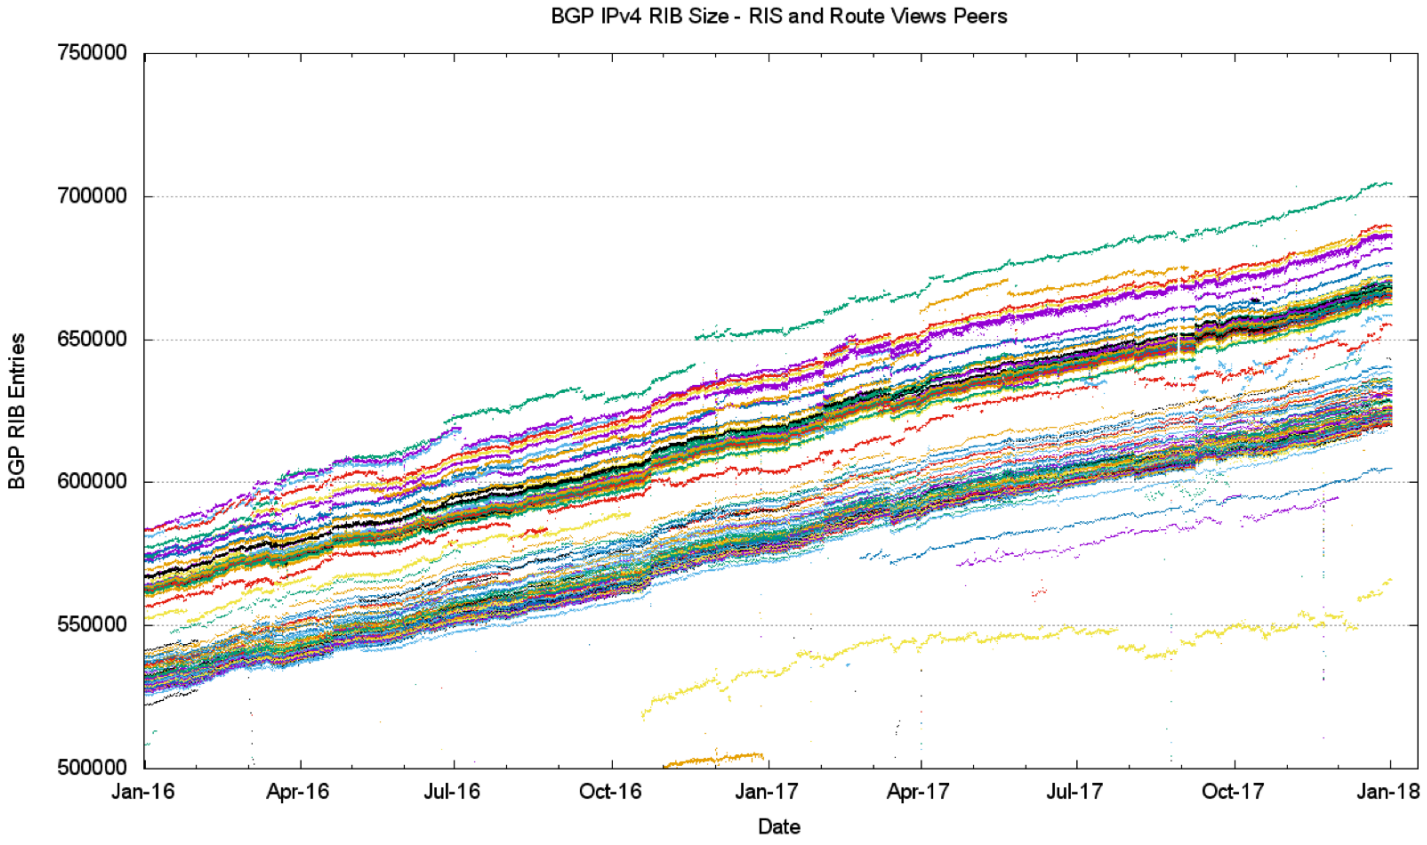

2016-2017 in detail

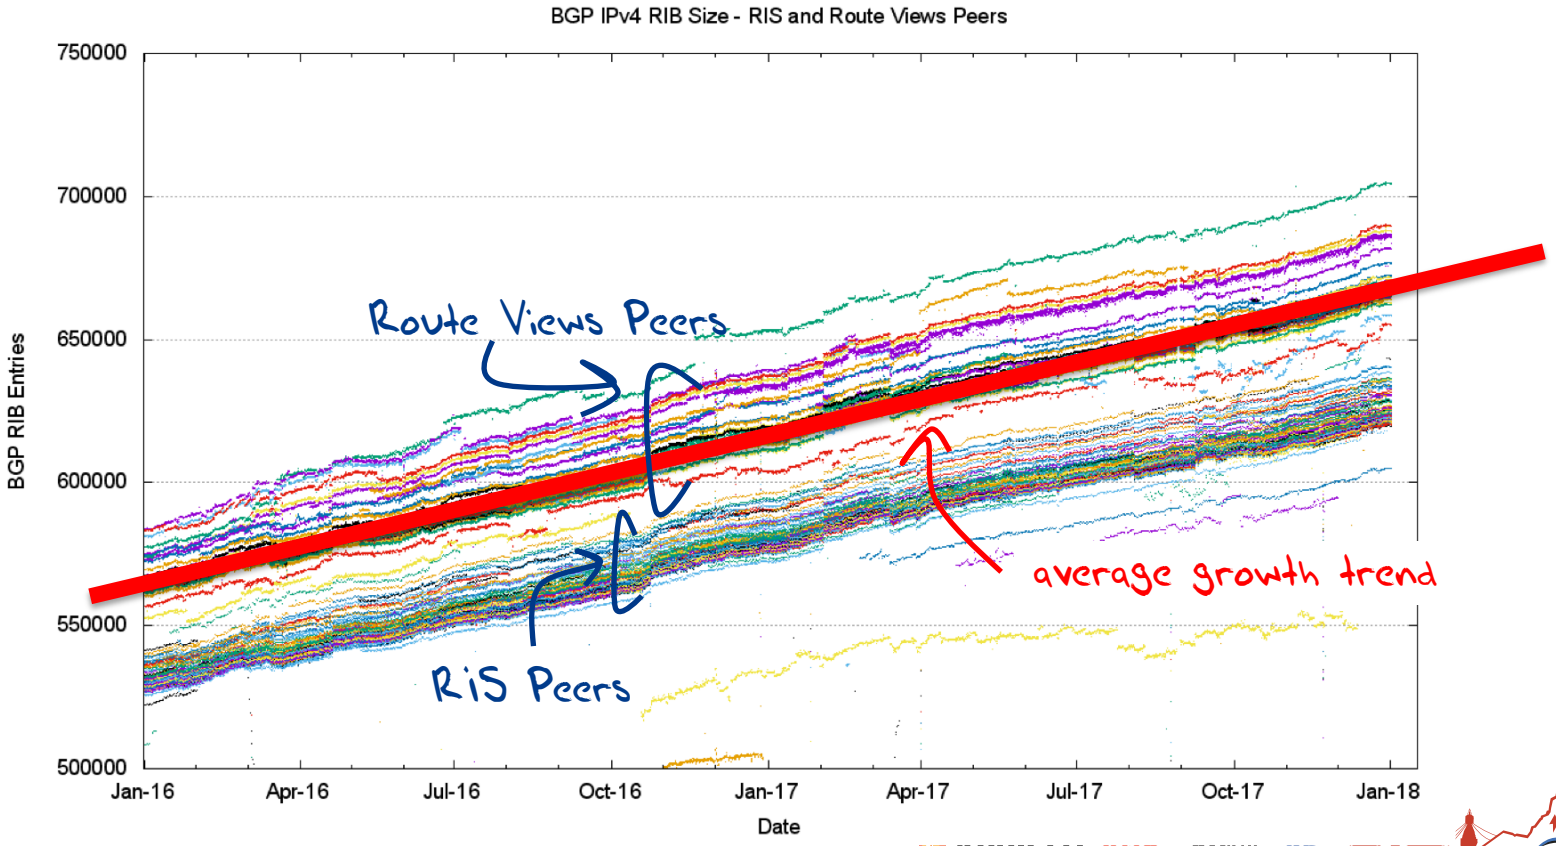

Routing Indicators for IPv4

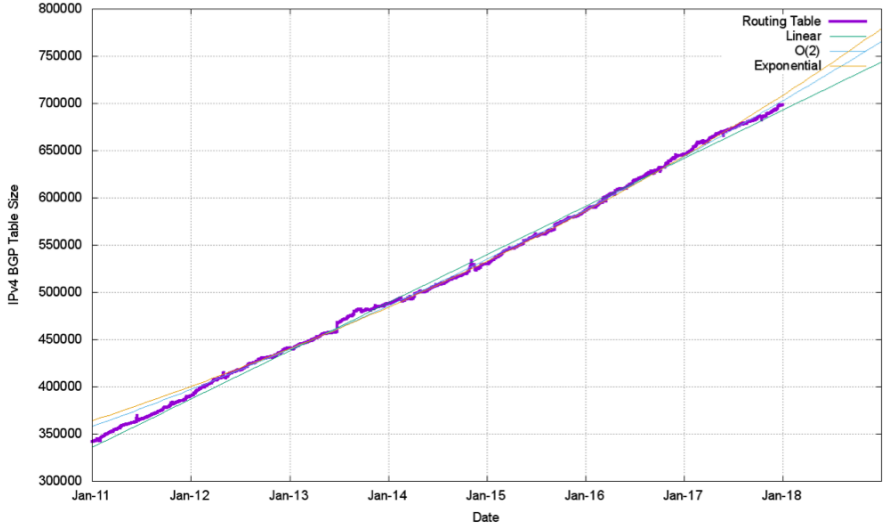

Routing prefixes – growing by some 53,000 prefixes per year

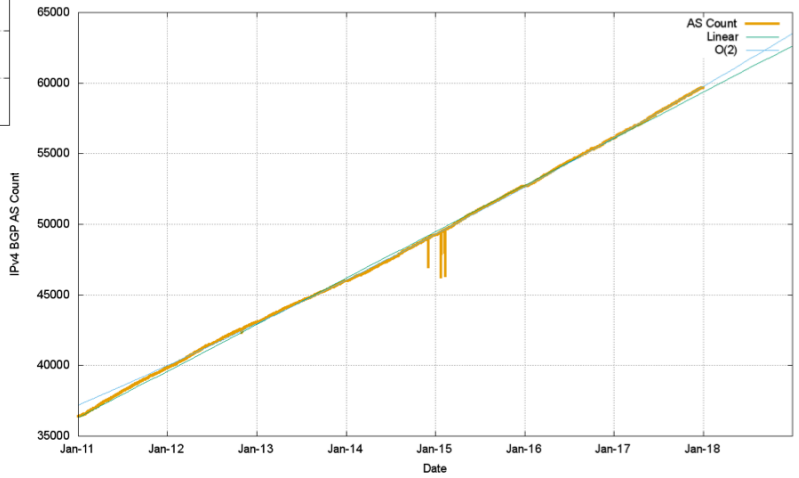

AS Numbers– growing by some 3,400 prefixes per year

Routing Indicators for IPv4

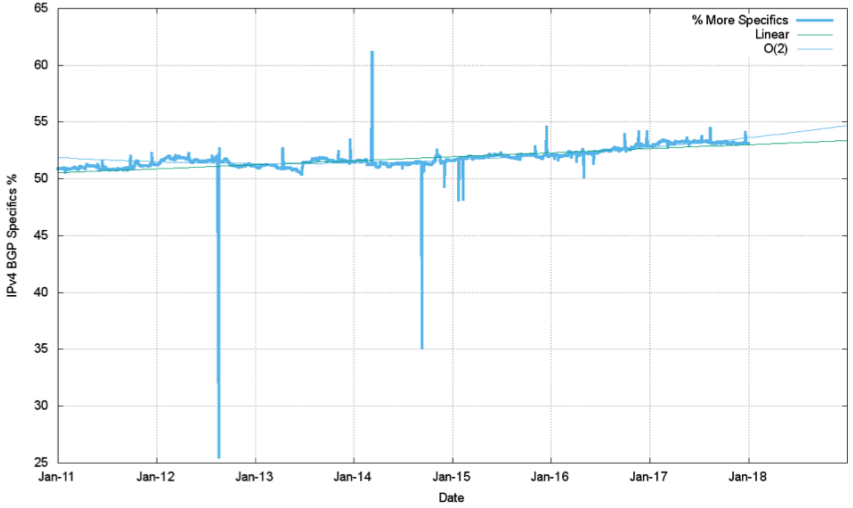

More Specifics are still taking up slightly more than one half of the routing table

But the average size of a routing advertisement continues to shrink

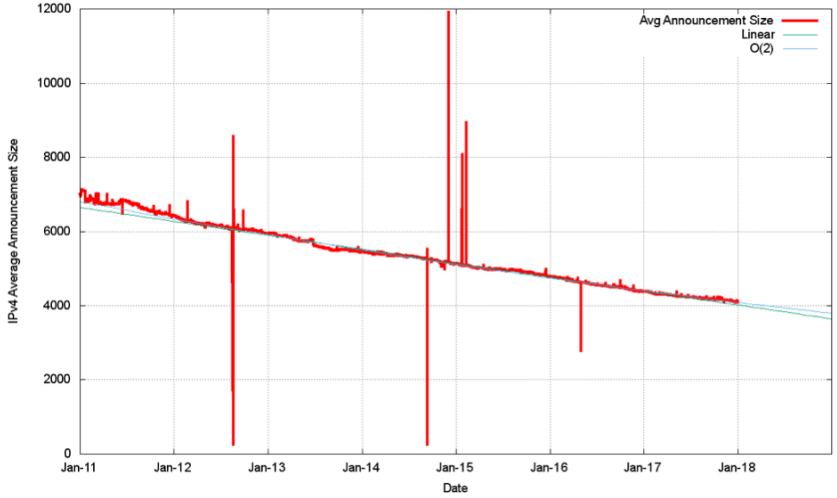

Routing Indicators for IPv4

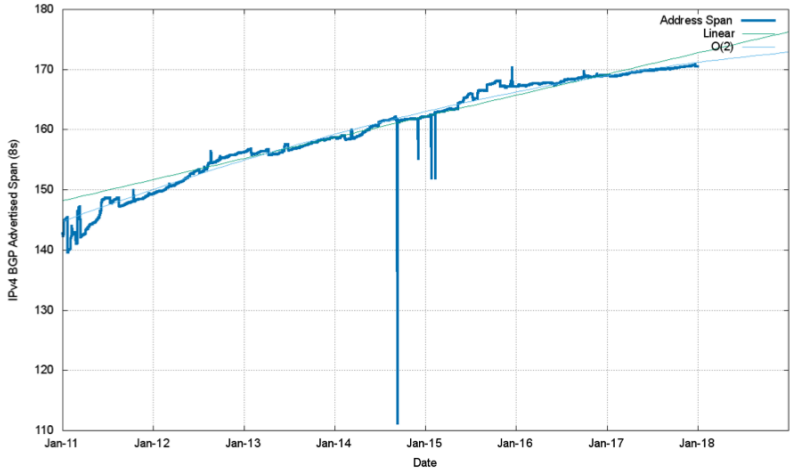

Address Exhaustion is now visible in the extent of advertised address space

The "shape" of inter-AS interconnection appears to be relatively steady

AS Adjacencies (Route-Views)

43,368 out of 60,493 ASNs have 1 or 2 AS Adjacencies (72%)

3,543 ASNs have 10 or more adjacencies

20 ASNs have >1,000 adjacencies

6,934	AS6939	HURRICANE - Hurricane Electric, Inc., US
5,301	AS174	COGENT-174 - Cogent Communications, US
5,037	AS3356	LEVEL3 - Level 3 Communications, Inc., US
2,467	AS3549	LVLT-3549 - Level 3 Communications, Inc., US
2,388	AS7018	ATT-INTERNET4 - AT&T Services, Inc., US
2,068	AS209	OBIT-AS "OBIT" Ltd., RU
2,056	AS57463	NETIX , BG
1,997	AS209	CENTURYLINK, US
1,776	AS37100	SEACOM-AS, MU
1,755	AS6461	ZAYO Bandwidth, US

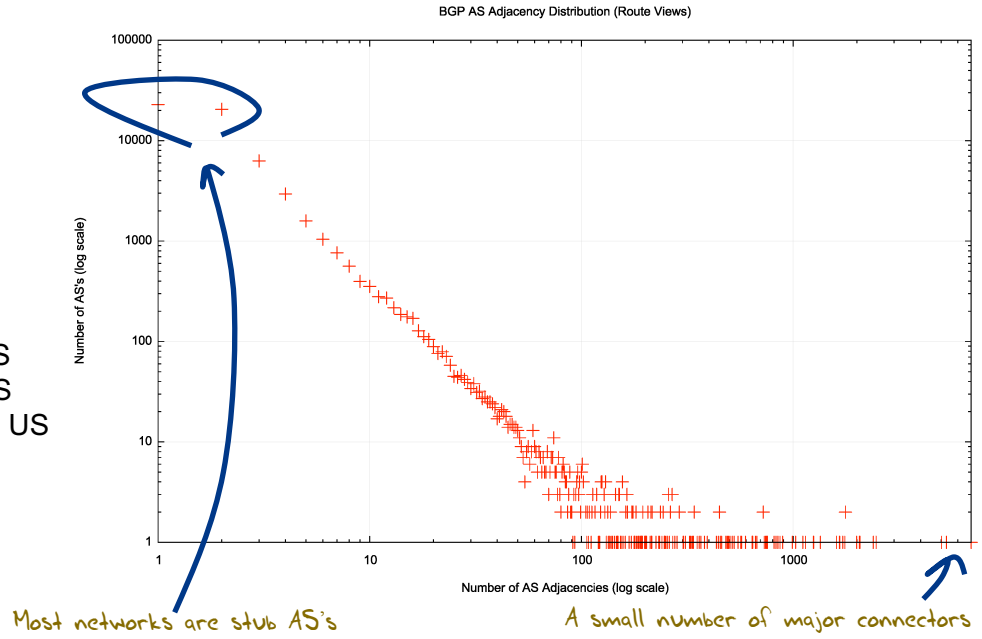

What happened in 2017 in V4?

Routing Business as usual – despite IPv4 address exhaustion!

- From the look of the growth plots, its business as usual, despite the increasing pressures on IPv4 address availability
- The number of entries in the IPv4 default-free zone reached 700,000 by the end of 2018
- The pace of growth of the routing table is still relatively constant at ~53,000 new entries and 3,400 new AS's per year
  - IPv4 address exhaustion is not changing this!
  - Instead, we are advertising shorter prefixes into the routing system

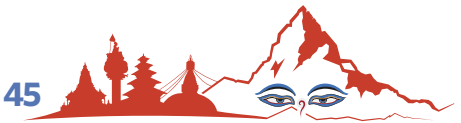

What about IPv4 Address Exhaustion?

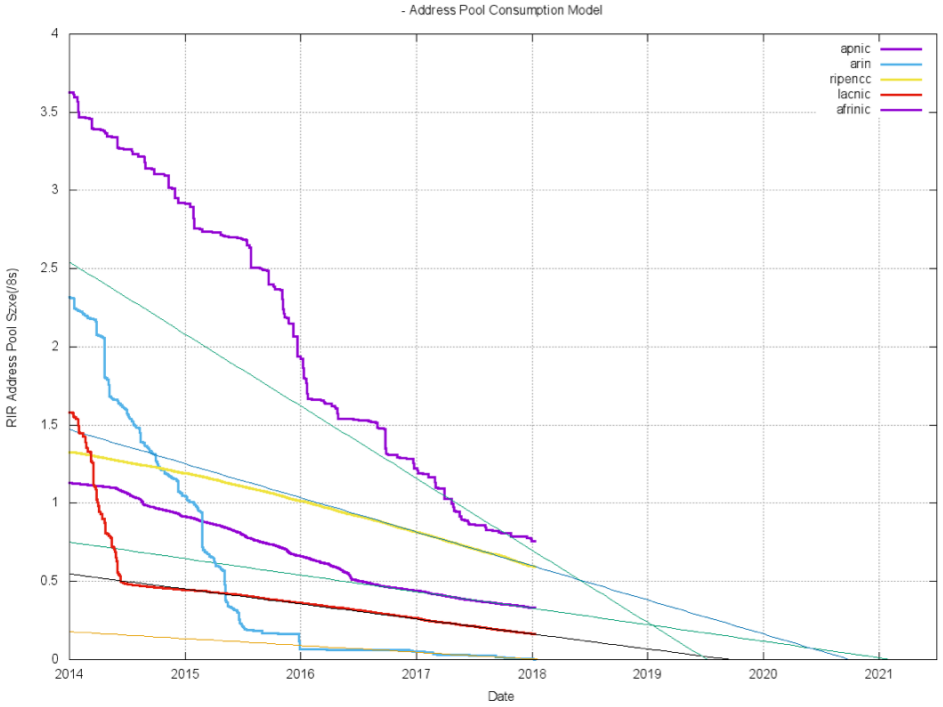

RIR Address Pool runout projections as of the start of 2018:

ARIN	– no free pool left
AFRINIC	– June 2019
LACNIC	– September 2019
RIPE NCC	– September 2020
APNIC	– February 2021

Post-Exhaustion Routing Growth

- What's driving this post-exhaustion growth?
  - Transfers?
  - Last /8 policies in RIPE and APNIC?
  - Leasing and address recovery?

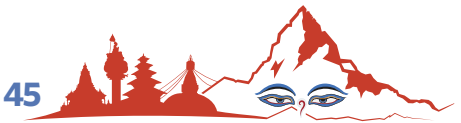

Advertised Address "Age"

2010

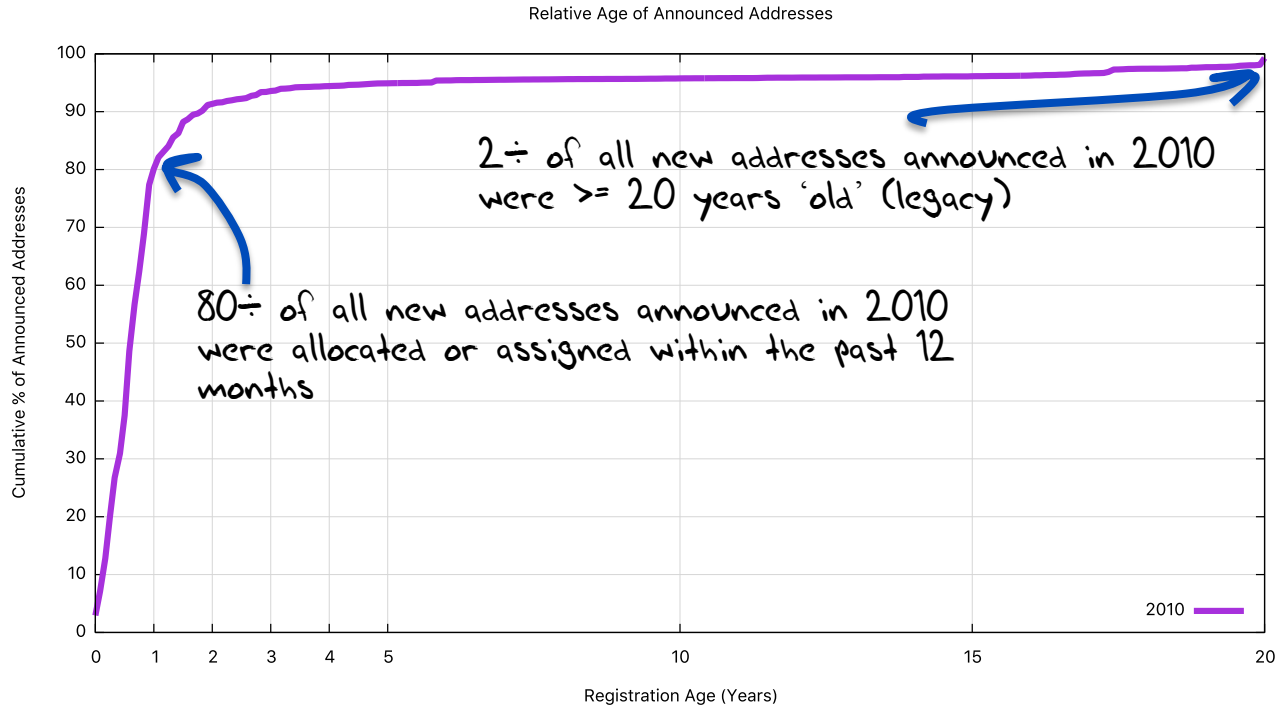

Advertised Address "Age"

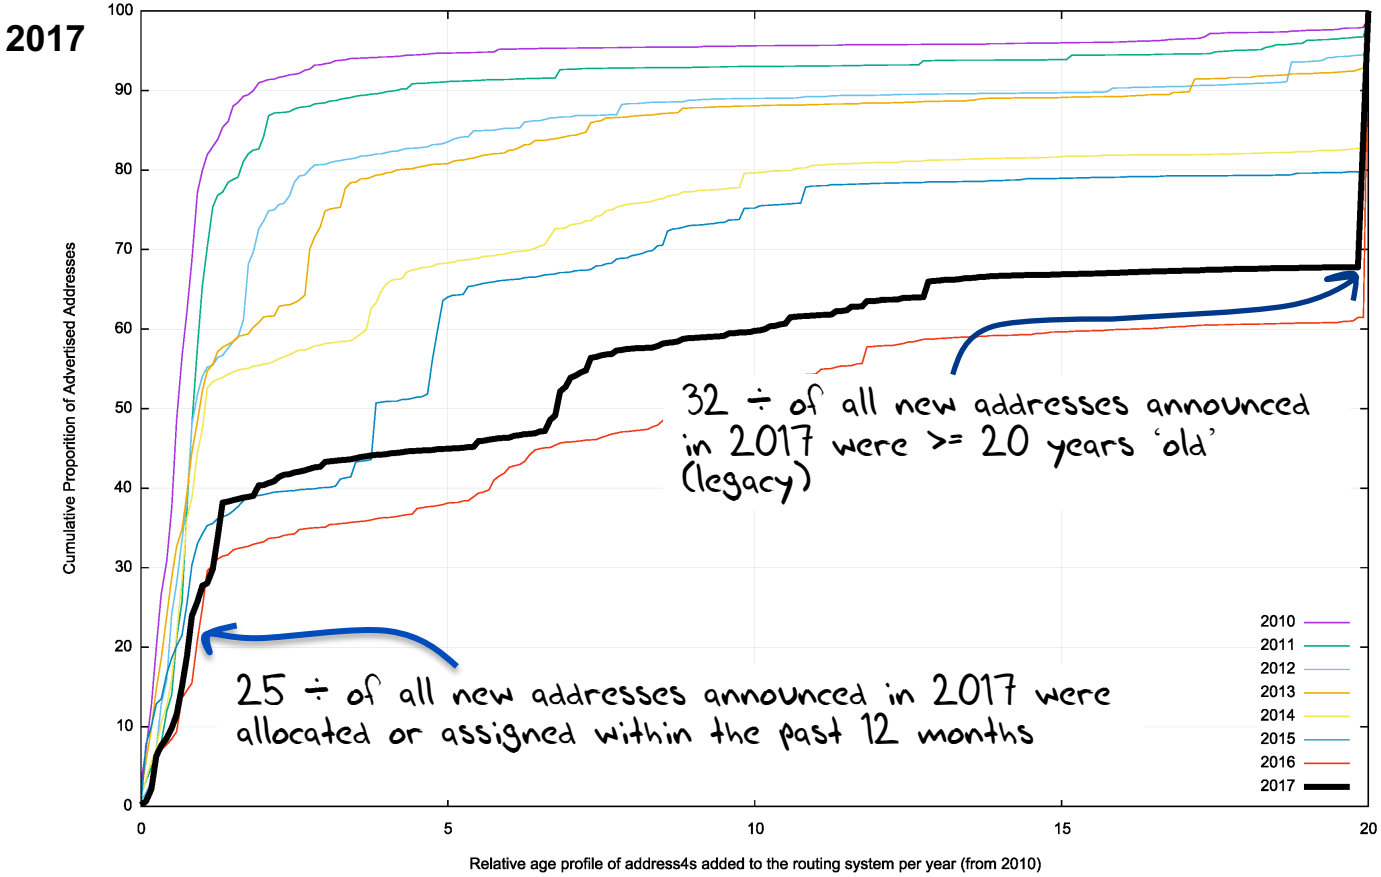

2000 - 2017: Advertised vs Unadvertised

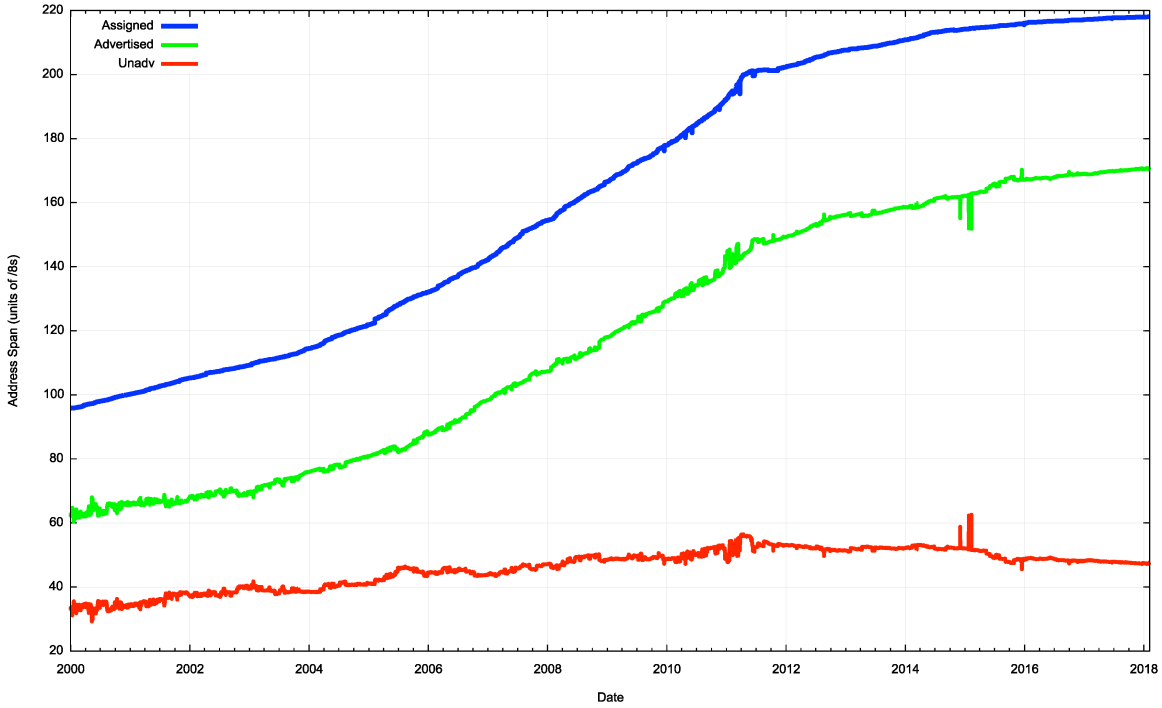

2000 - 2017: Unadvertised Addresses

2017: Assigned vs Recovered

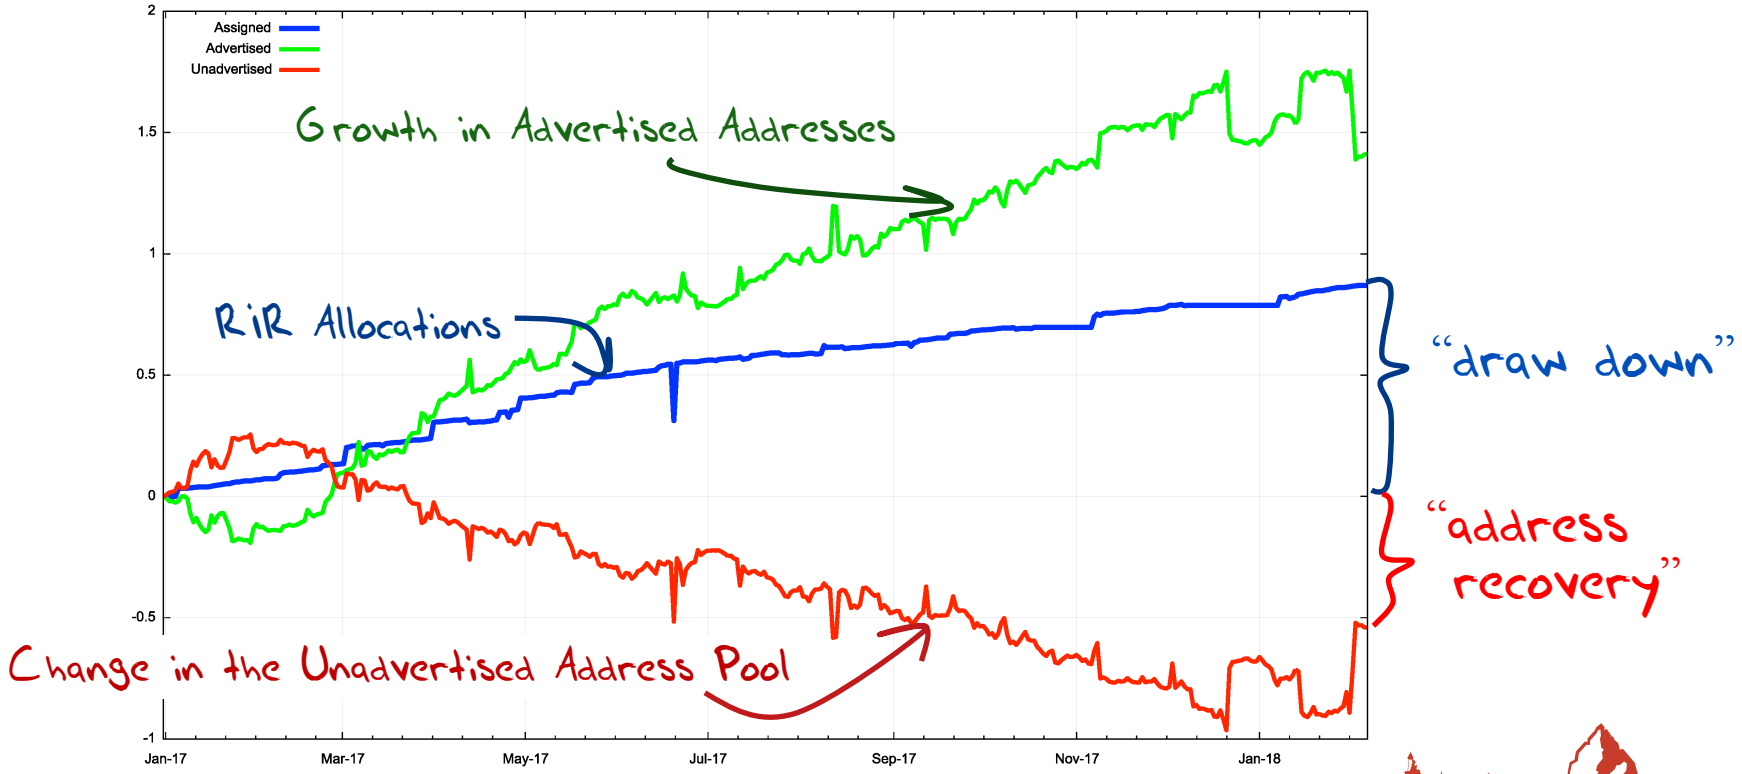

V4 in 2017

The equivalent of 1.4 /8s was added to the routing table across 2017

- Approximately 0.8 /8s were assigned by RIRs in 2015
  - 0.5 /8's assigned by Afrinic
  - 0.3 /8s were assigned by RIPE NCC and APNIC (Last /8 allocations)

- And a net of 0.6 /8's were recovered from the pool of previously Unadvertised addresses

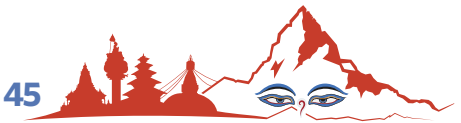

The Route-Views View of IPv6

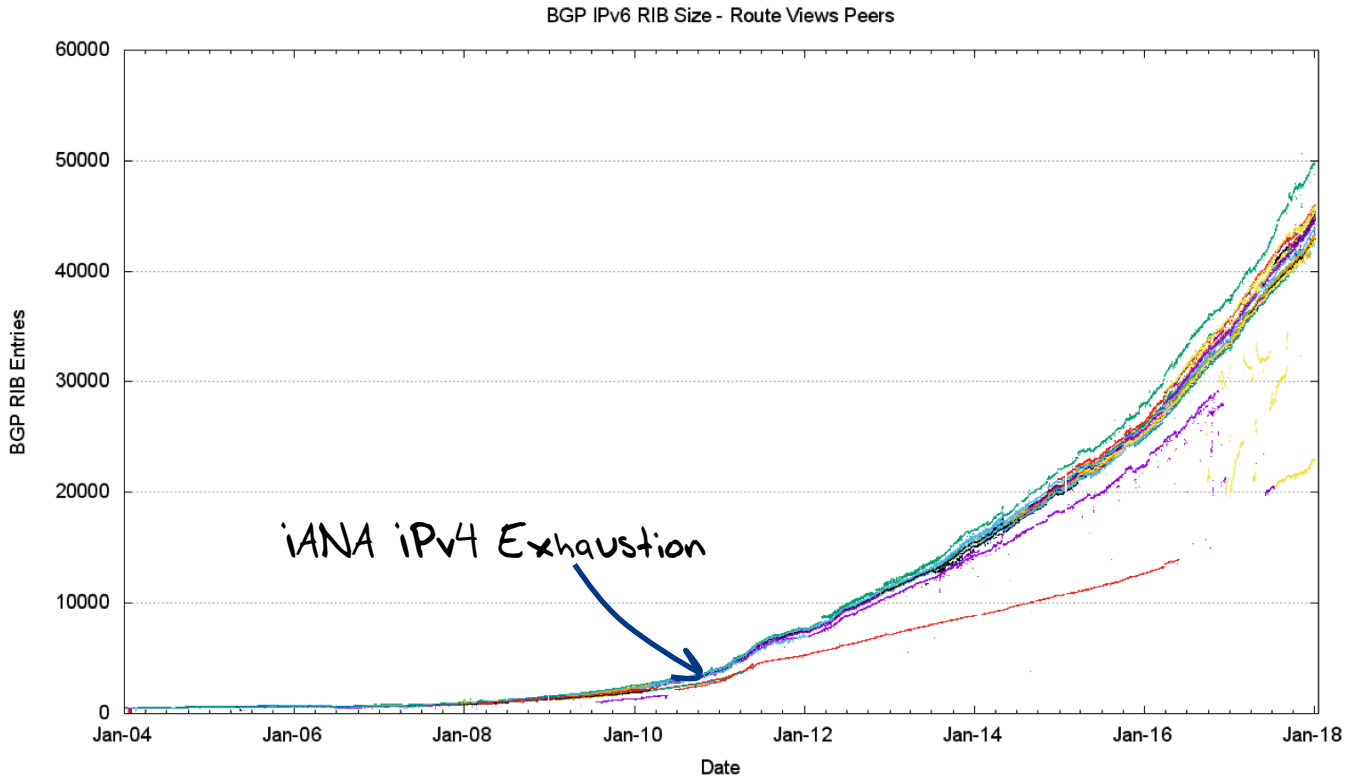

2016-2017 in Detail

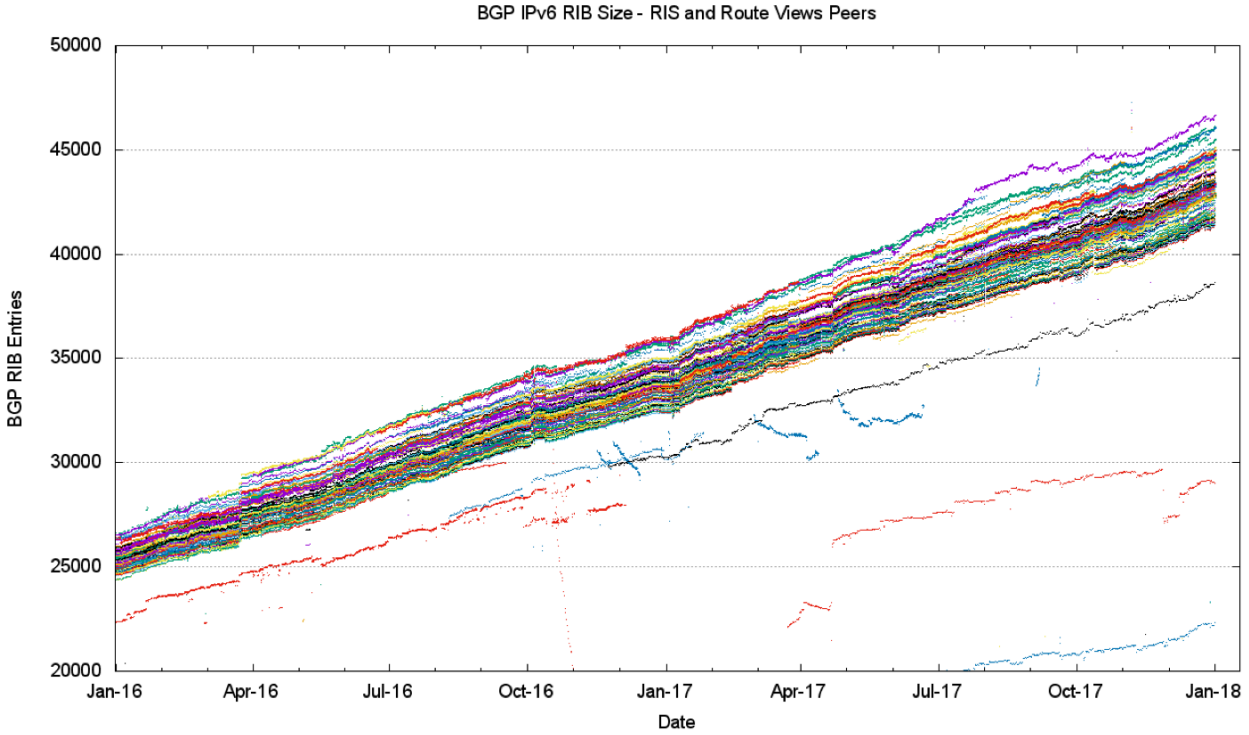

Routing Indicators for IPv6

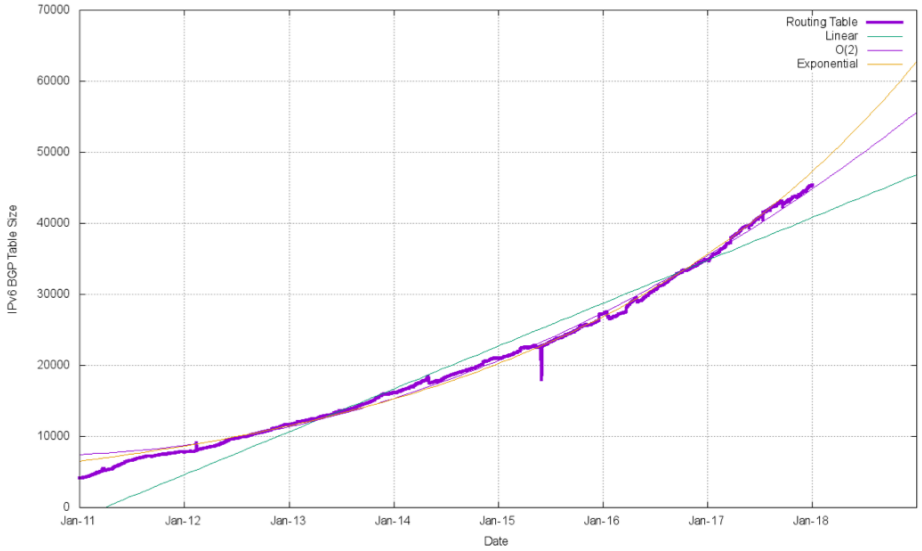

Routing prefixes – growing by some 10,000 prefixes per year

AS Numbers– growing by some 1,700 prefixes per year (which is half the V4 growth)

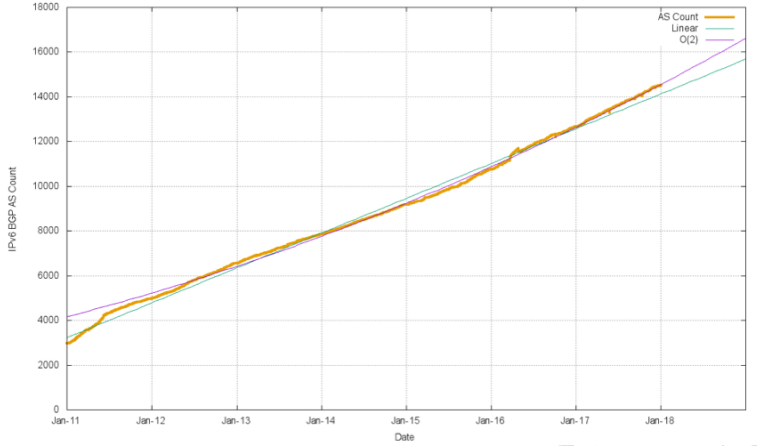

Routing Indicators for IPv6

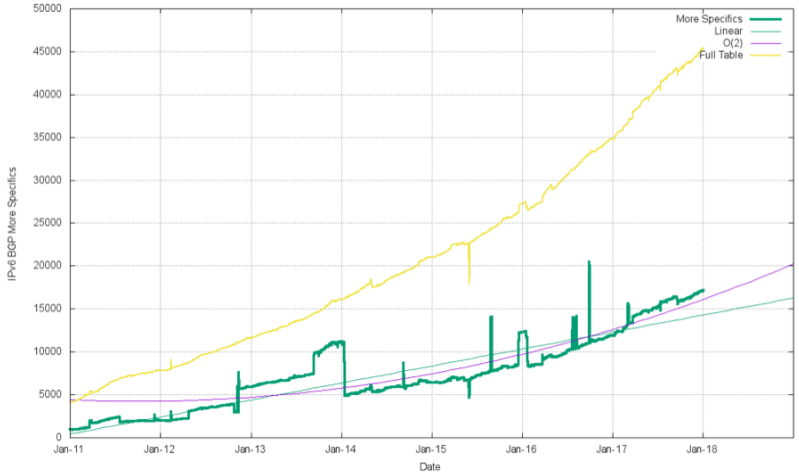

More Specifics now take up more than one third of the routing table

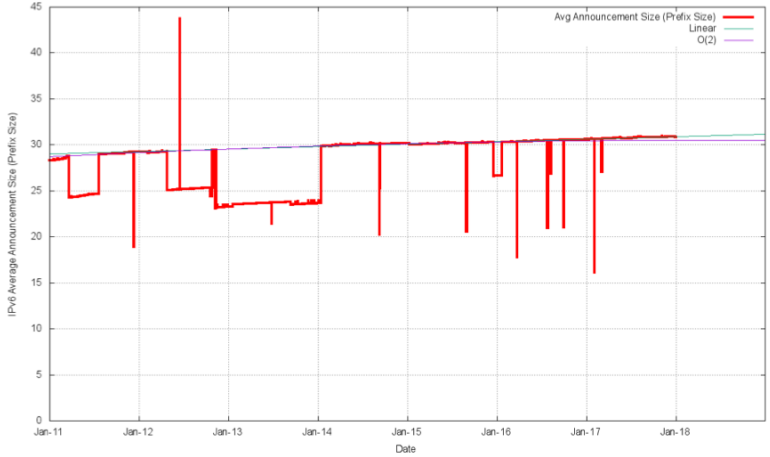

The average size of a routing advertisement is getting smaller

Routing Indicators for IPv6

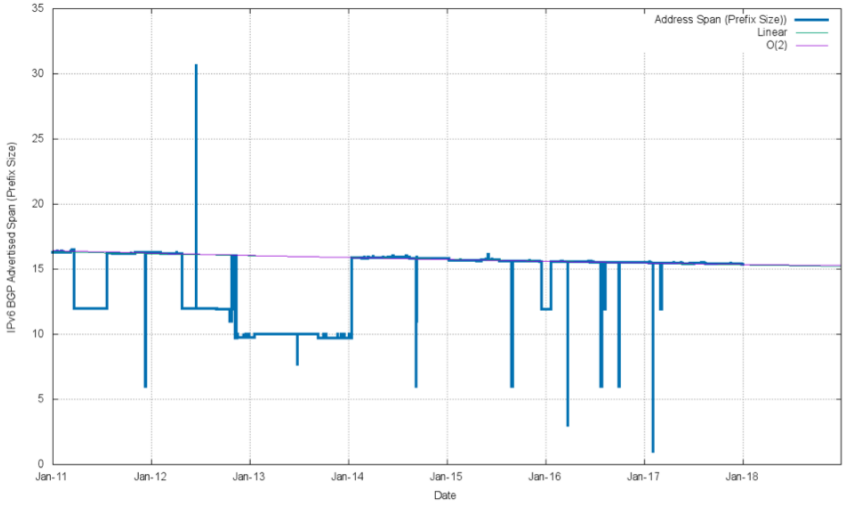

Advertised Address span is growing at a linear rate

The "shape" of inter-AS interconnection in IPv6 appears to be steady, as the Average AS Path length has been held steady through the year

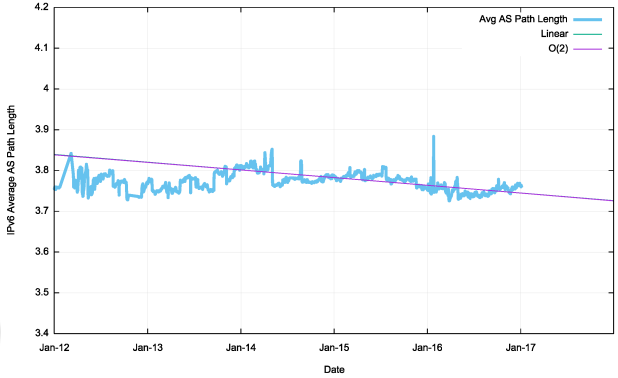

AS Adjacencies (Route Views)

10,415 out of 14,975 ASNs have 1 or 2 AS Adjacencies (69%)
1,020 ASNs have 10 or more adjacencies
5 ASNs have >1,000 adjacencies

4,190 AS6939 HURRICANE - Hurricane Electric, Inc., US
1,711 AS174 COGENT-174 - Cogent Communications, US
1,436 AS3356 LEVEL3 - Level 3 Communications, Inc., US
1,179 AS37100 SEACOM-AS, MU
1,041 AS1299 Telia Carrier, SE

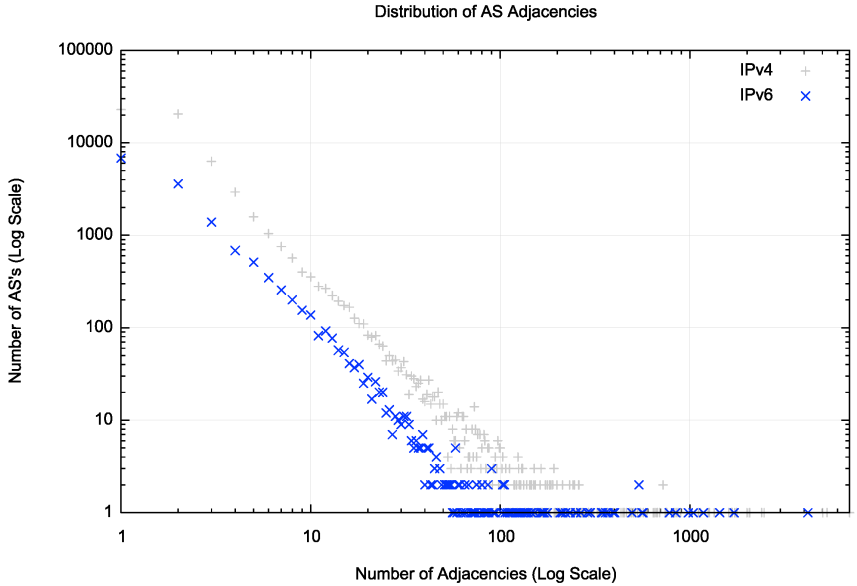

V6 in 2017

- Overall IPv6 Internet growth in terms of BGP is steady at some **10,000 route entries p.a.**

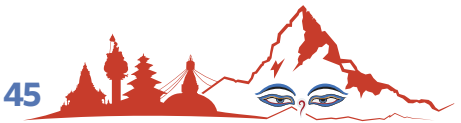

What to expect

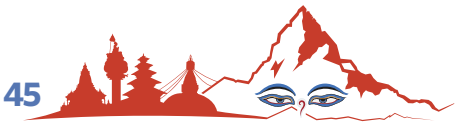

BGP Size Projections

For the Internet this is a time of some uncertainty

- Registry IPv4 address run out
- Uncertainty over the impacts of market-mediated movements of IPv4 on the routing table
- Uncertainty over the timing of IPv6 takeup leads to a mixed response to IPv6 so far, and no clear indicator of trigger points for change for those remaining IPv4-only networks

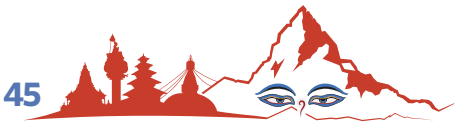

V4 - Daily Growth Rates

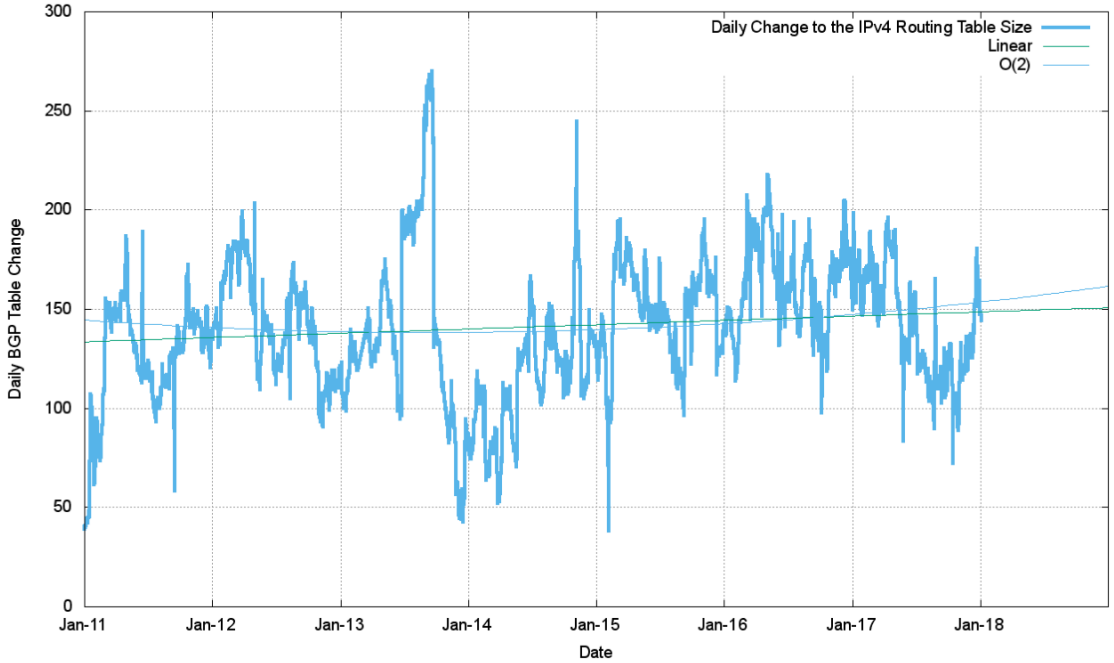

Growth in the V4 network appears to be constant at a long term average of 140 additional routes per day, or some 51,000 additional routes per year

V4 BGP Table Size Predictions

Jan *2016*	*587,000*
2017	*646,000*
2018	***699,000***
2019	*744,000*
2020	*795,000*
2021	*846,000*
2022	*897,000*
2023	*948,000*

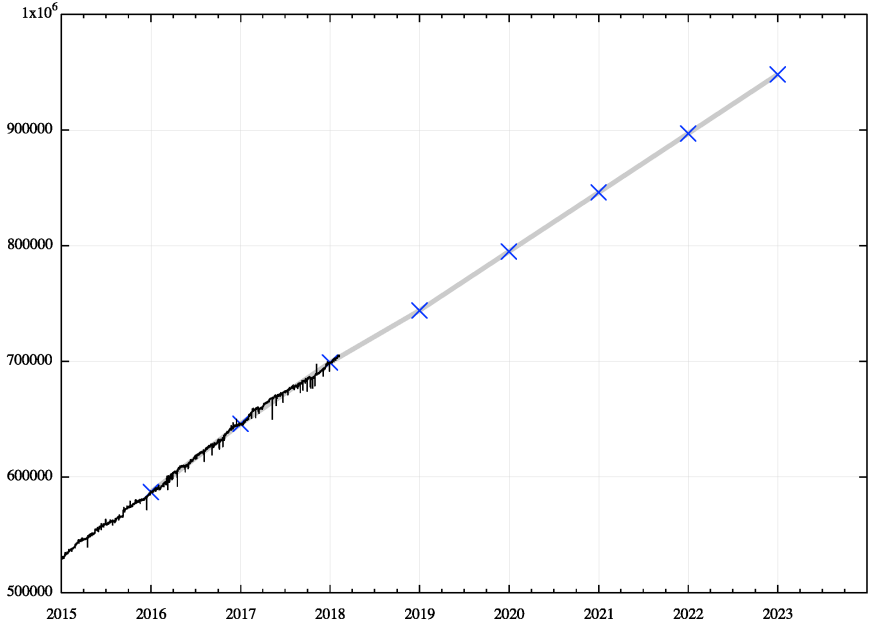

V6 - Daily Growth Rates

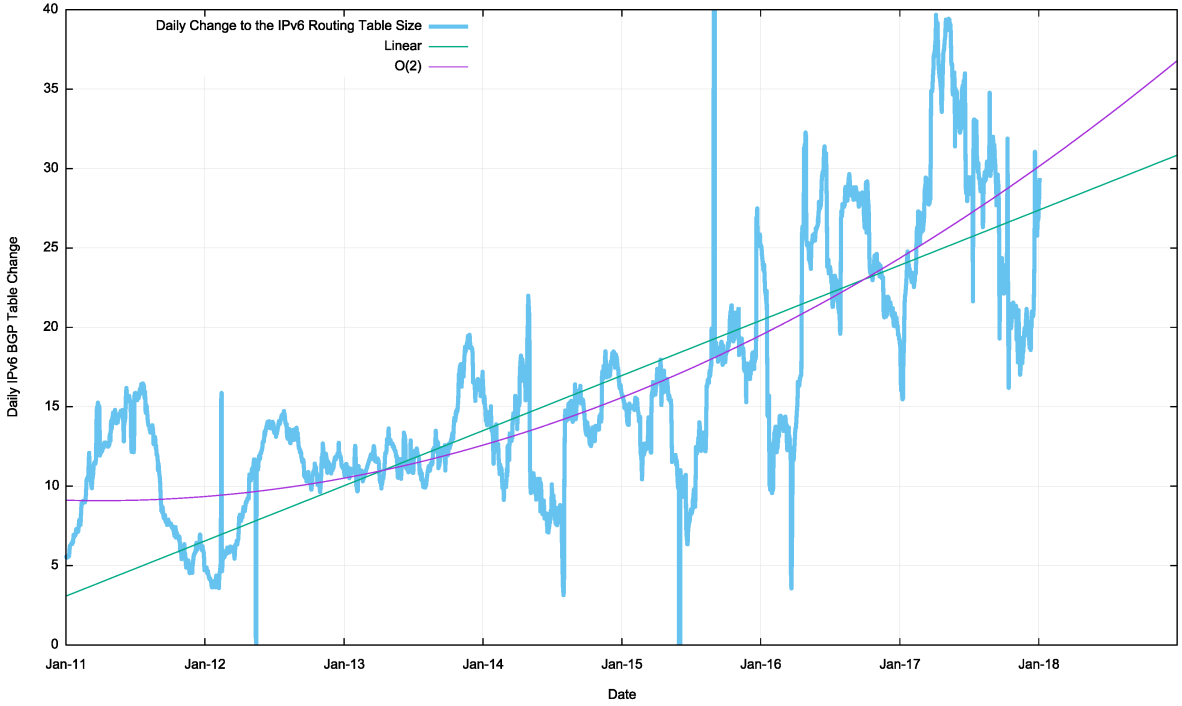

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2016*	*27,000*	*27,000*
2017	*37,000*	*37,000*
2018	***45,000***	***45,000***
2019	*47,000*	*63,000*
2020	*53,000*	*83,000*
2021	*59,000*	*110,000*
2022	*65,000*	*146,000*
2023	*71,000*	*194,000*

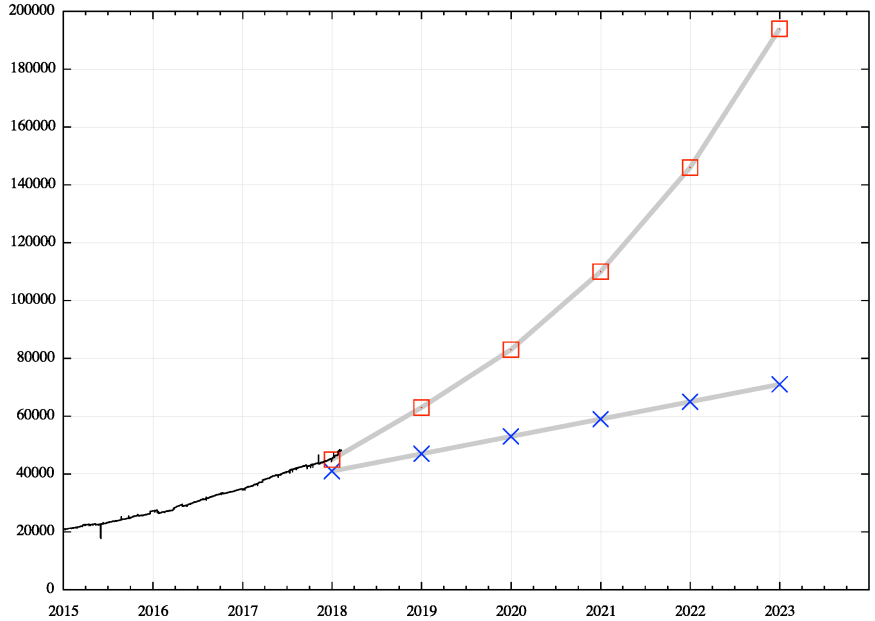

BGP Table Growth

Nothing in these figures suggests that there is cause for urgent alarm -- at present

- The overall eBGP growth rates for IPv4 are holding at a modest level, and the IPv6 table, although it is growing at a faster relative rate, is still small in size in absolute terms
- As long as we are prepared to live within the technical constraints of the current routing paradigm, the Internet's use of BGP will continue to be viable for some time yet

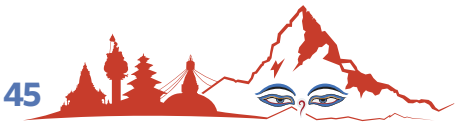

BGP Updates

- What about the level of updates in BGP?

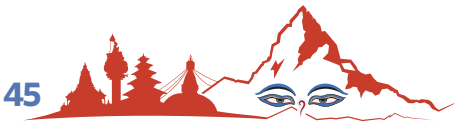

IPv4 BGP Updates

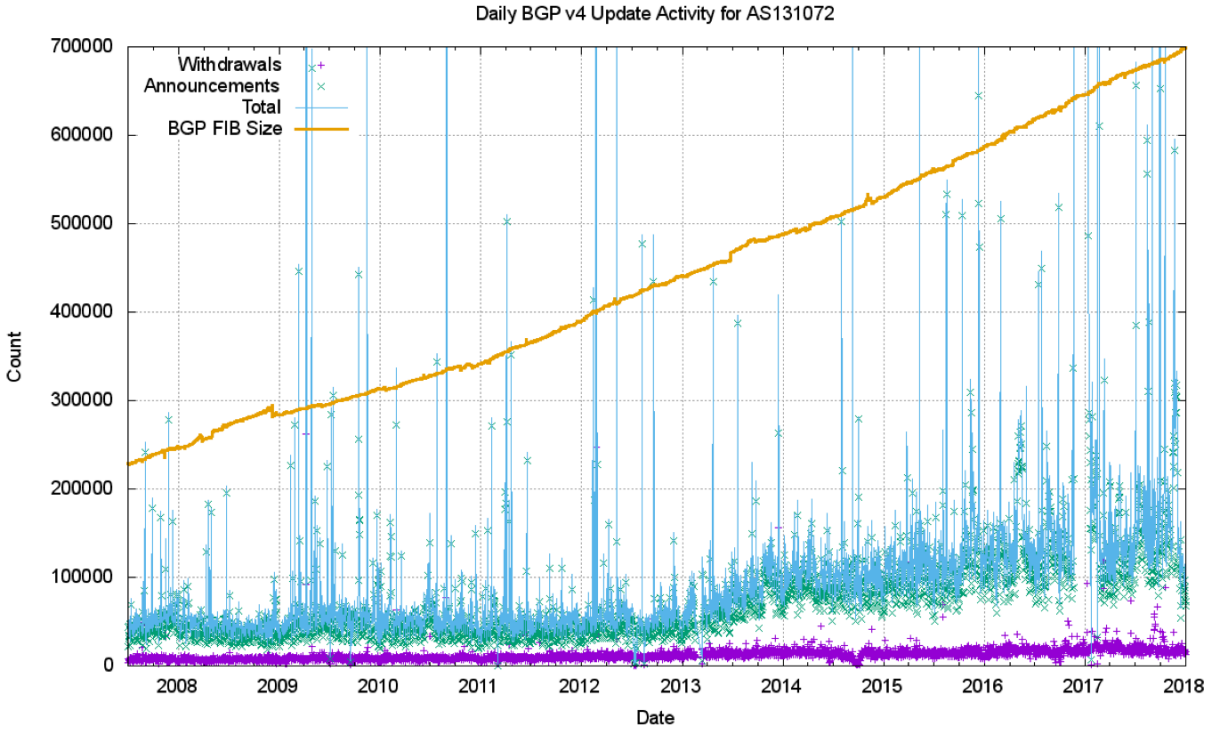

IPv4 BGP Convergence Performance

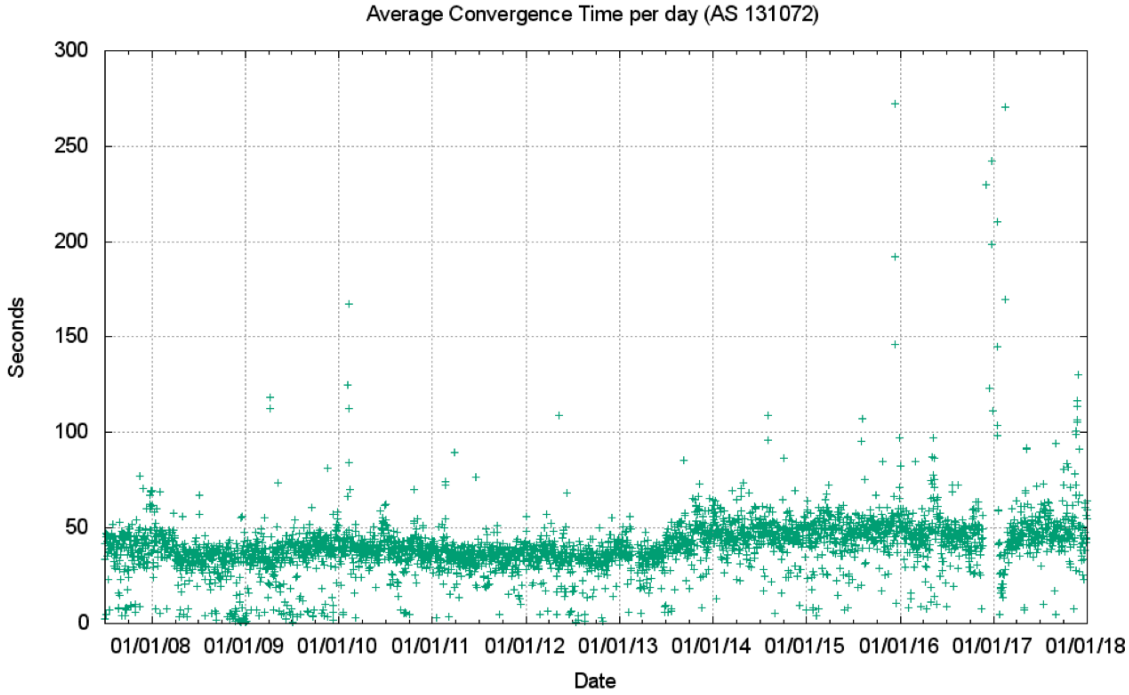

Updates in IPv4 BGP

Still no great level of concern …

- The number of updates per instability event has been relatively constant, which unanticipated.
- Likely contributors to this outcome are the damping effect of widespread use of the MRAI interval by eBGP speakers, and the compressed topology factor, as seen in the relatively constant V4 AS Path Length

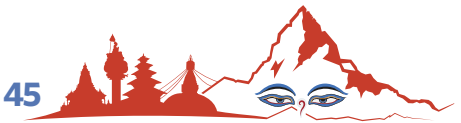

V6 BGP Updates

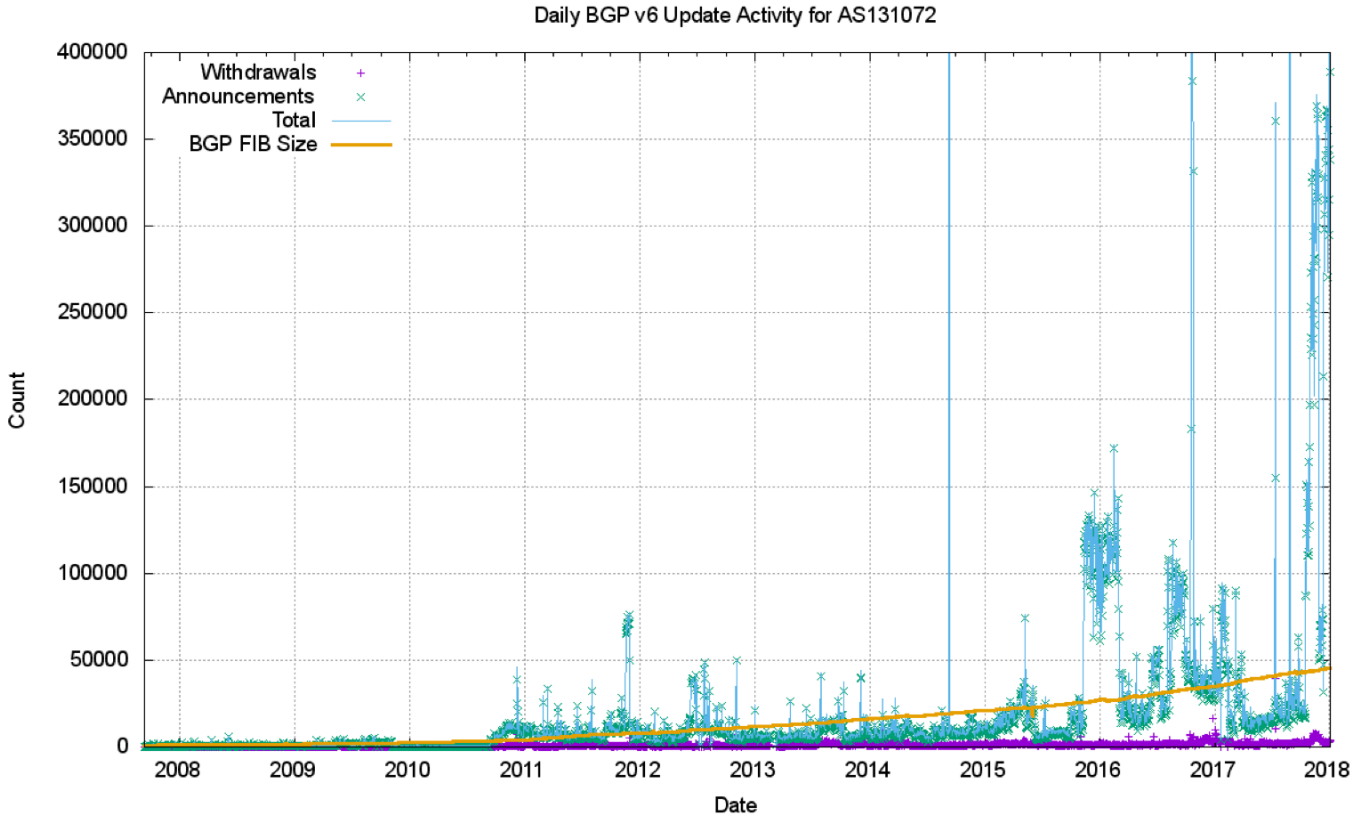

V6 Convergence Performance

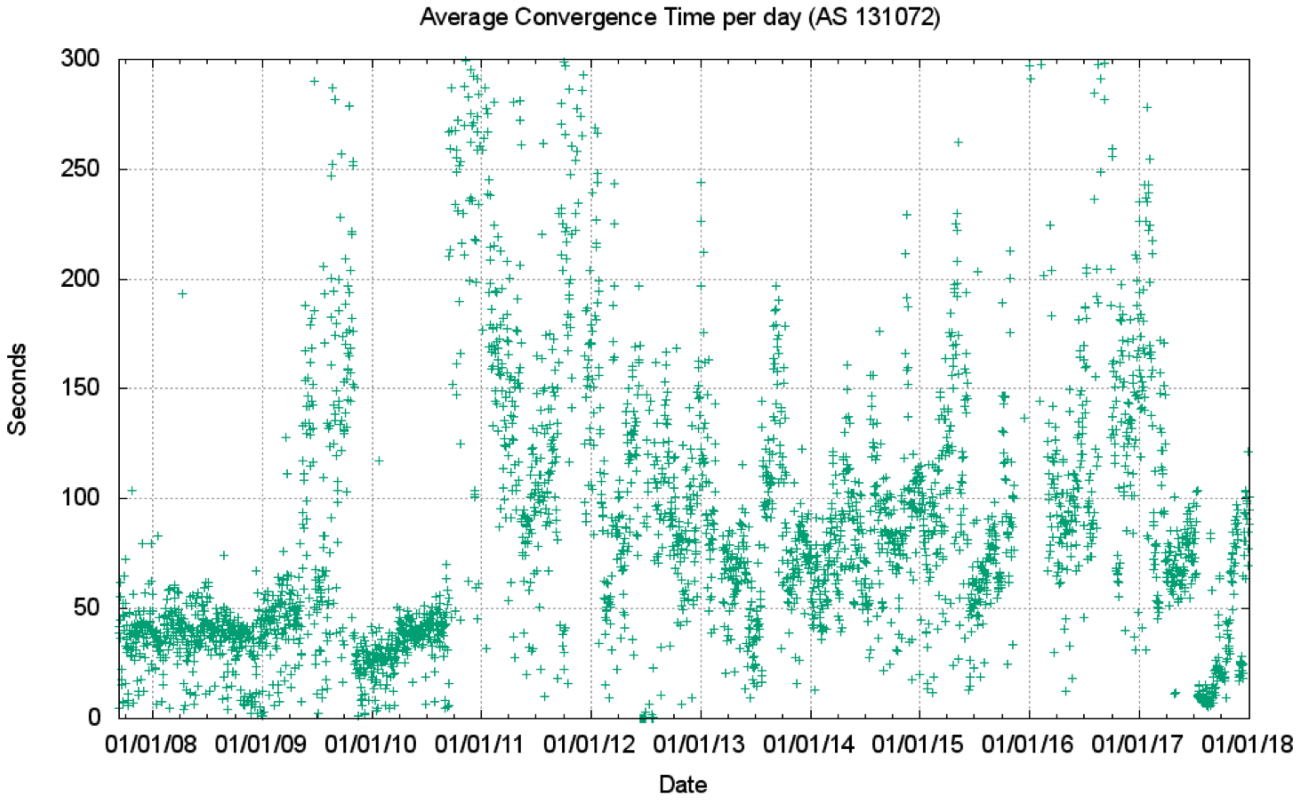

Updates in IPv6

BGP Route Updates are very unequally distributed across the prefix set – they appear to affect a very small number of prefixes which stand out well above the average

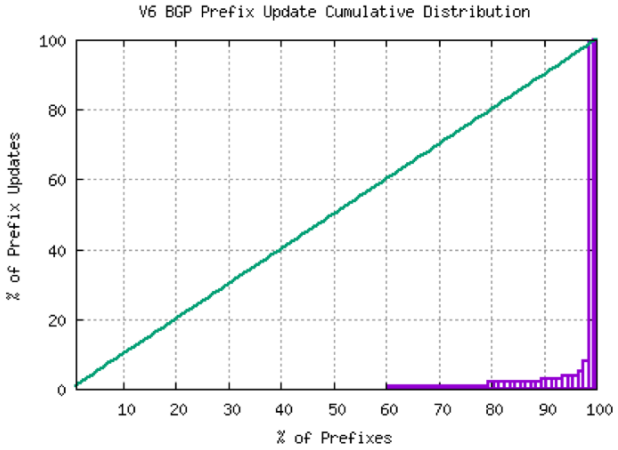

Updates in IPv6

50 Most active ASes for the past 31 days

RANK	ASN	UPDs	%	Prefixes	UPDs/Prefix	AS NAME
1	133481	3459540	42.64%	85	40700.47	AIS-FIBRE-AS-AP AIS Fibre, TH
2	131445	1219687	15.03%	67	18204.28	AIS3G-2100-AS-AP Advance Wireless Network, TH
3	10226	566196	6.98%	16	35387.25	ETL-IX-AS-AP Enterprise of Telecommunications Lao, LA
4	133193	237089	2.92%	21	11289.95	PEATHAILAND-AS-AP Provincial Electricity Authority (PEA), TH
5	17552	231485	2.85%	85	2723.35	TRUE-AS-AP True Internet Co.,Ltd., TH
6	55430	204560	2.52%	290	705.38	STARHUBINTERNET-AS-NGNBN Starhub Internet Pte Ltd, SG

The busiest 6 origin AS's prefixes accounted for 70% of all BGP IPv6 prefix updates

Routing Futures

- There is little in the way of scaling pressure from BGP as a routing protocol – the relatively compressed topology and stability of the infrastructure links tend to ensure that BGP remains effective in routing the internet

- The issues of FIB size, line speeds and equipment cost of line cards represent a more significant issue for hardware suppliers – we can expect cheaper line cards to to use far smaller LRU cache local FIBs in the high speed switches and push less used routes to a slower / cheaper lookup path. This approach may also become common in very high speed line cards

Some Practical Suggestions

- Understand your hardware's highspeed FIB capacity in the default-free parts of your network
- Review your IPv4 / IPv6 portioning - a dual-stack eBGP router will need 850,000 IPv4 slots and 85,000 IPv6 slots for a full eBGP routing table in line cards over the coming 24 months
- Judicious use of default routes in your internal network may allow you drop this requirement significantly

That's it!

Questions?