BGP in 2020

Geoff Huston
APNIC Labs

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Conclusions

27 Years of Routing the Internet

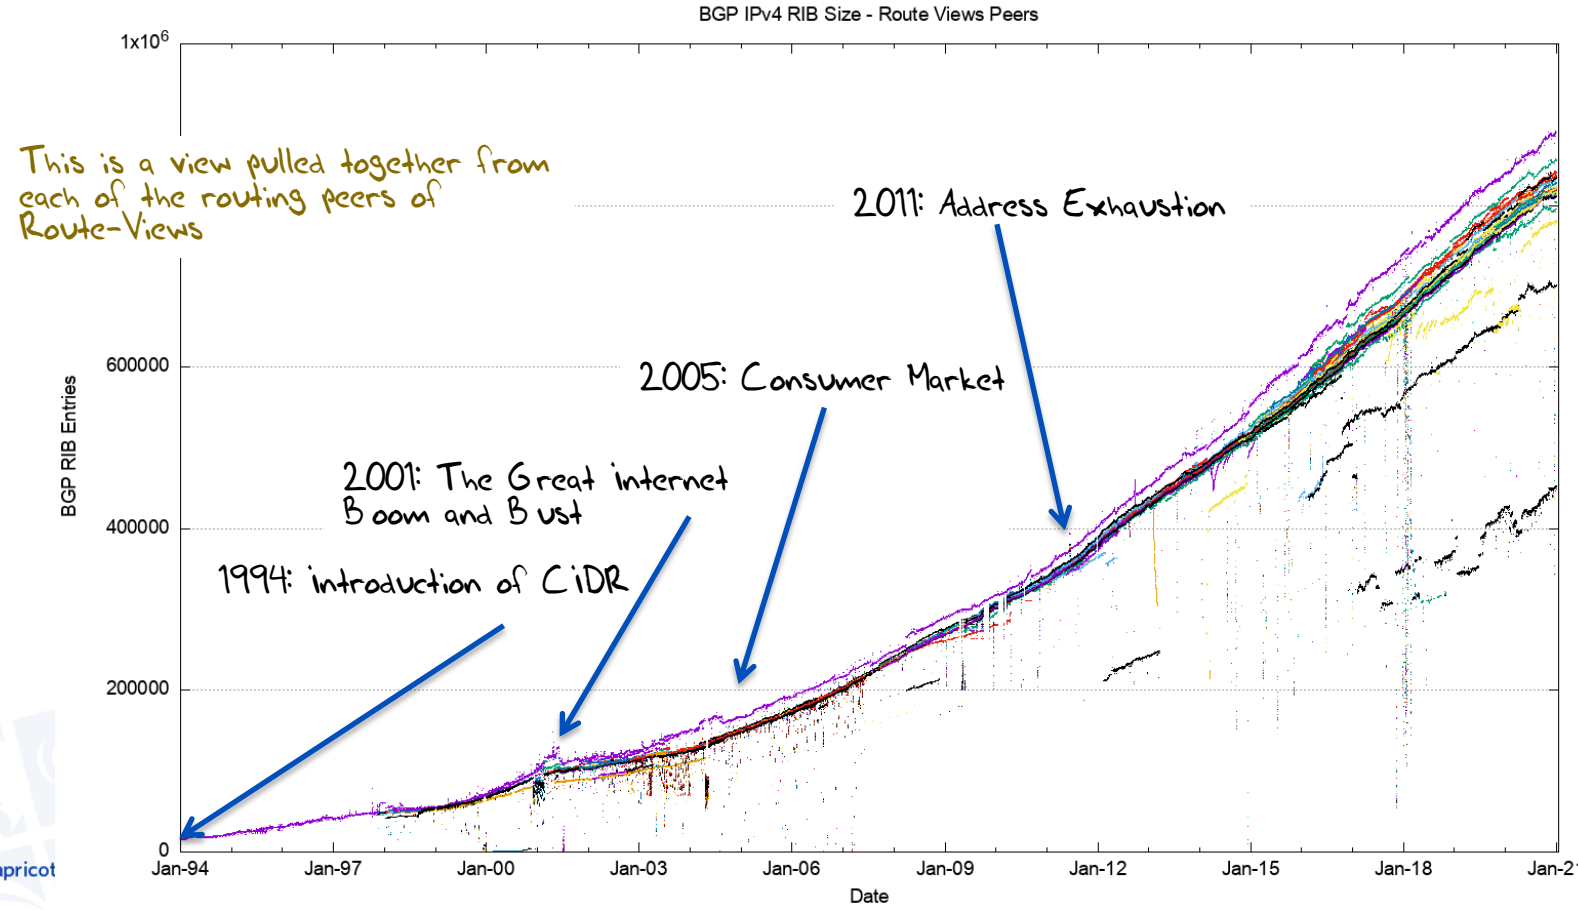

2019-2020 in detail

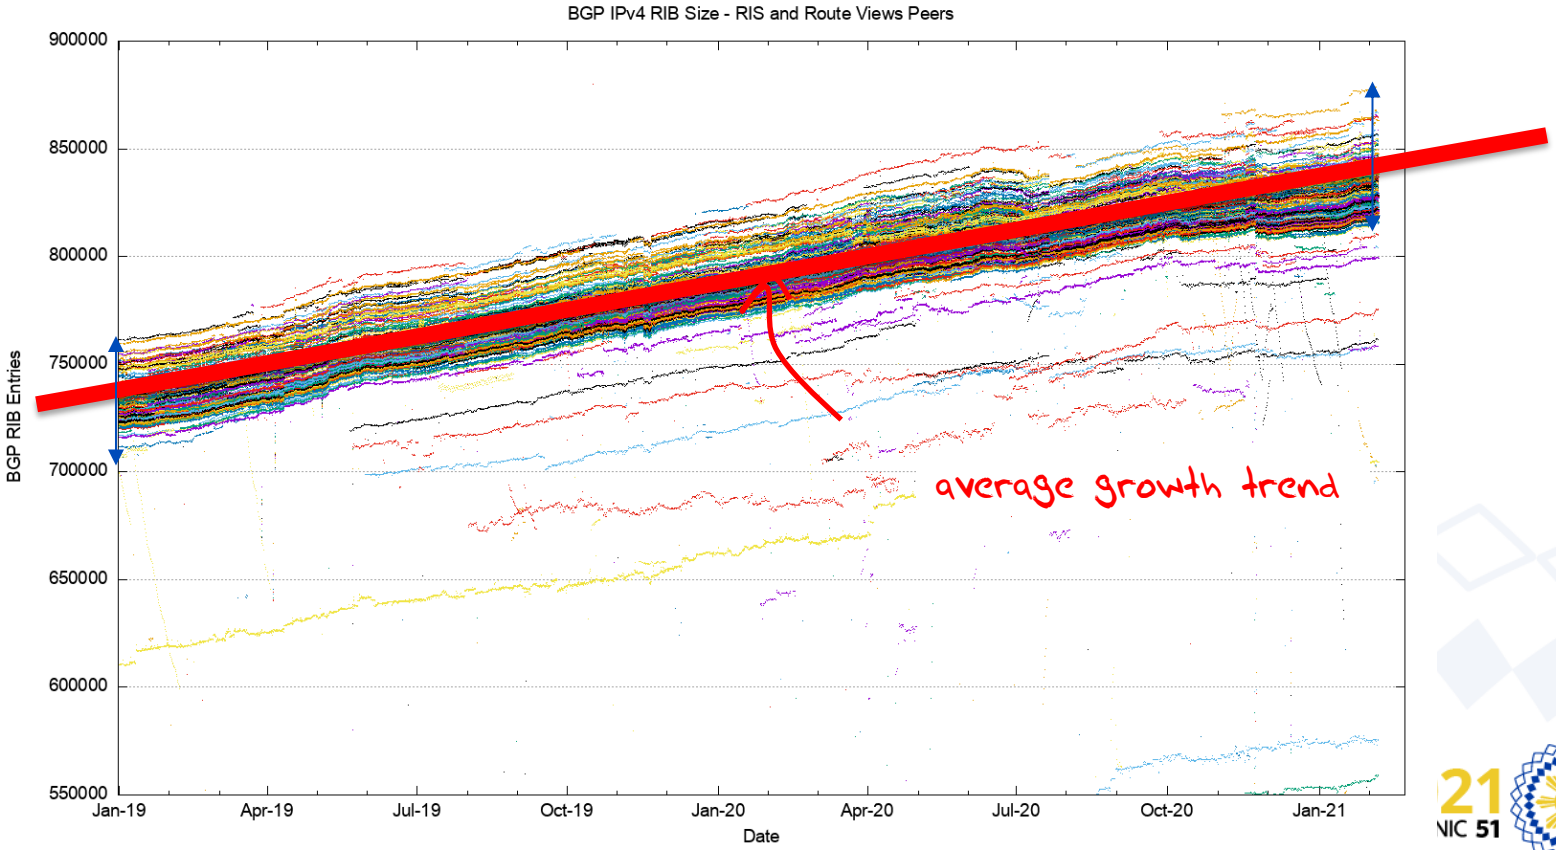

2019: Assigned vs Recovered

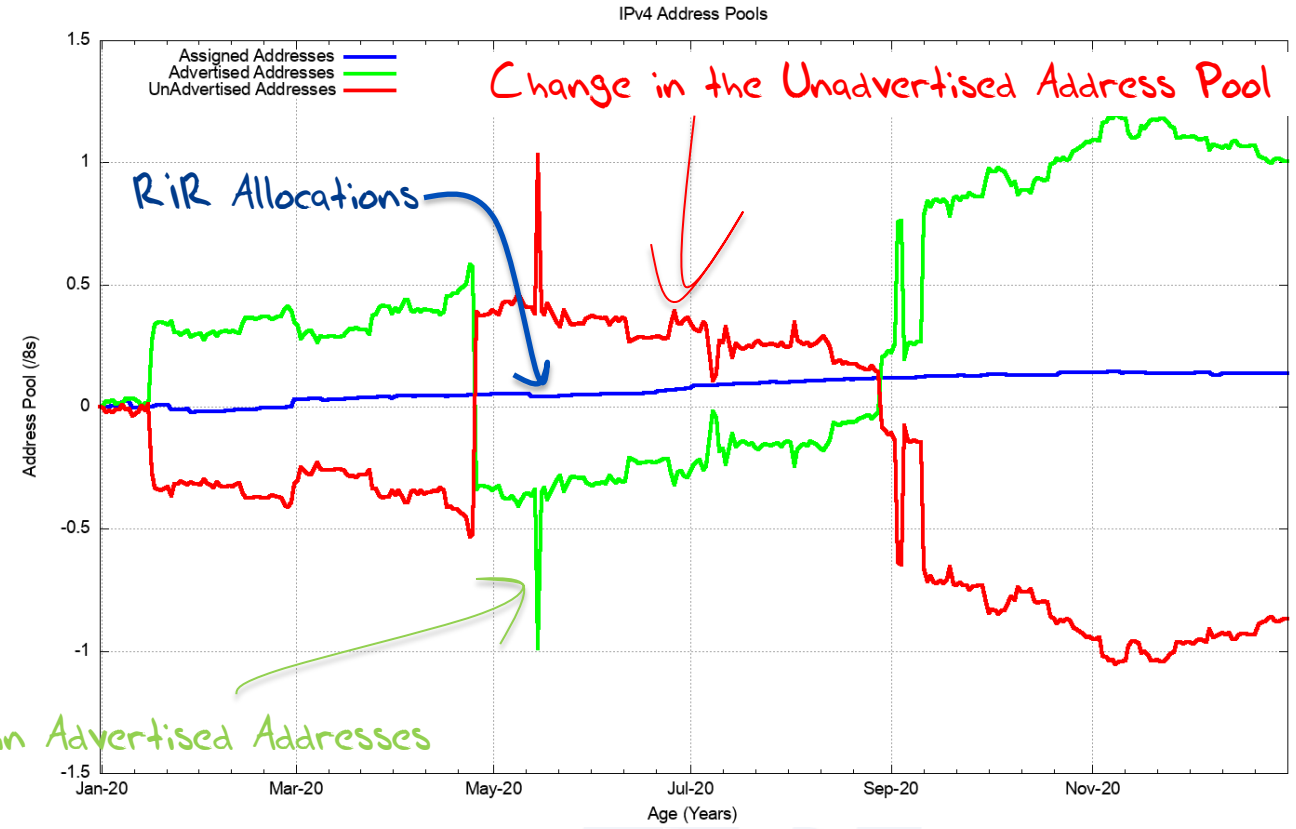

16.5M advertised

2.3M allocated

14.5M recovered

What happened in 2020 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 860,000 by the end of 2020
- The pace of growth of the routing table was slightly lower than the rolling 5-year average, with **52,000 new entries in 2020**
- The AS position was steady with 3,400 new AS's per year
- Transit relationships have not changed materially over 2020 for most networks
- IPv4 address exhaustion is not changing this picture as yet
- Instead, we appear to be advertising shorter prefixes into the routing system

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Conclusions

The Route-Views View of IPv6

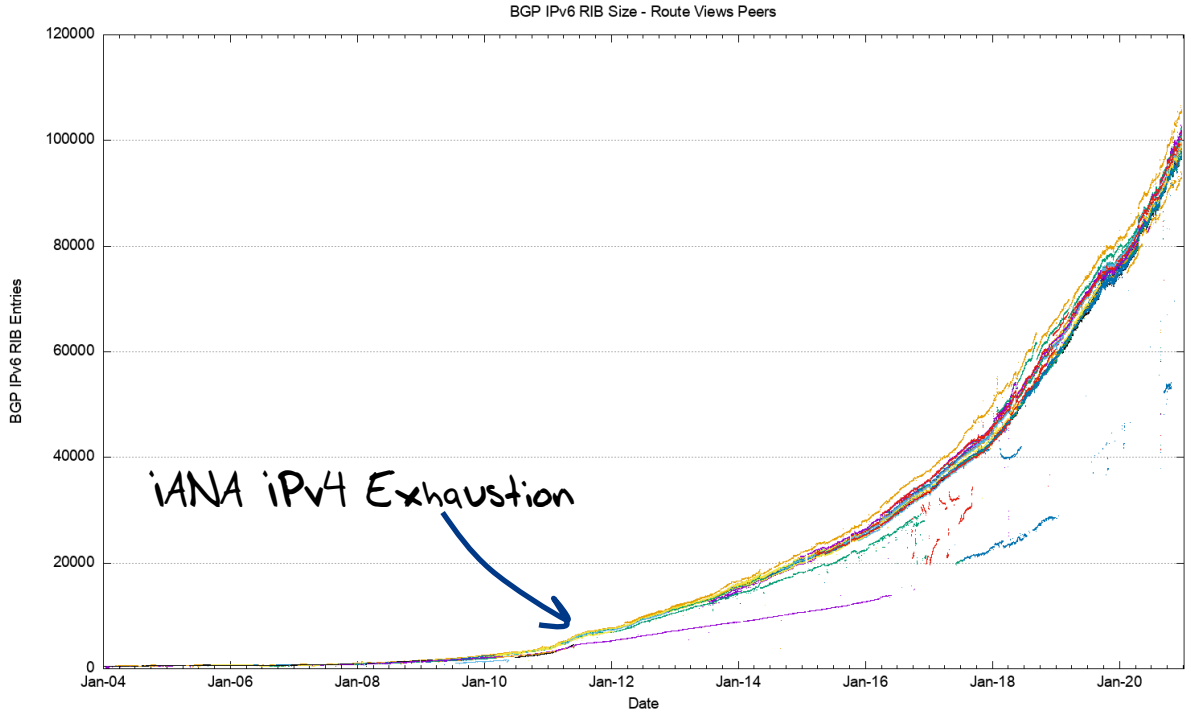

2019-2020 in Detail

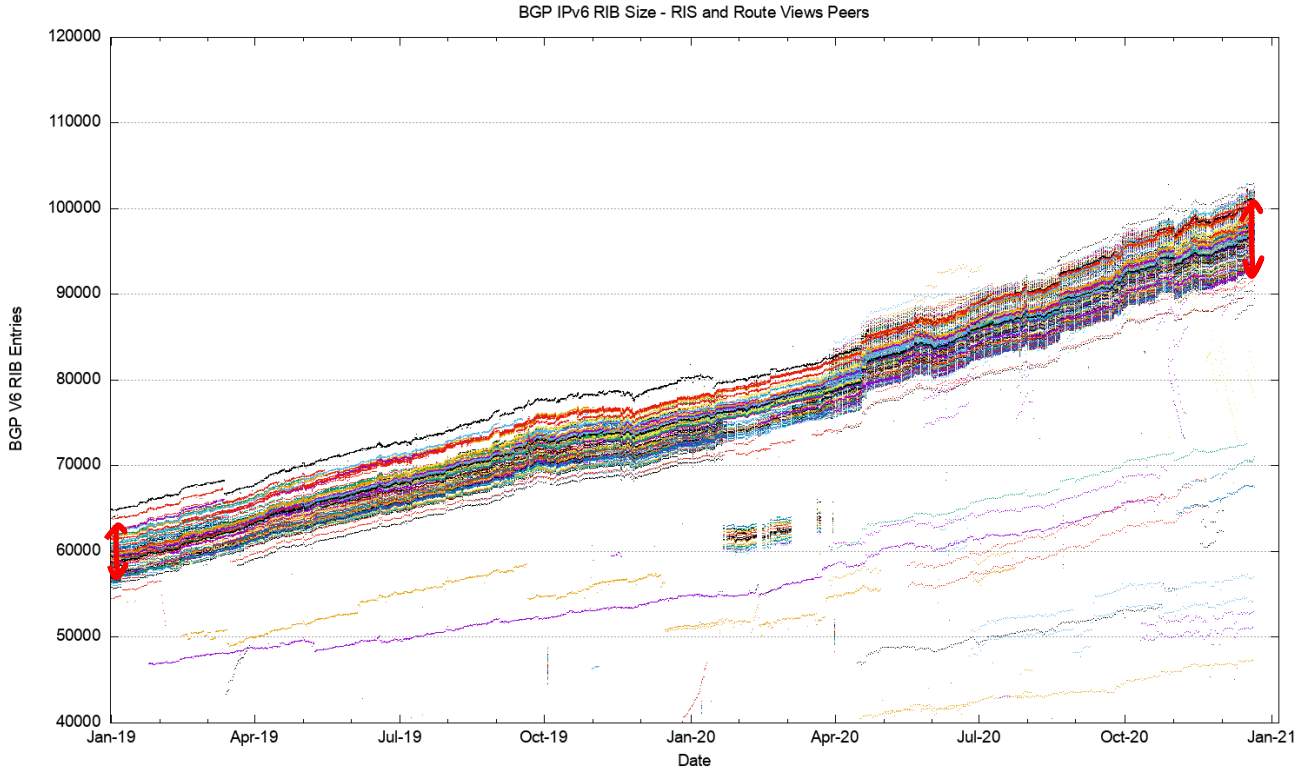

V6 in 2020

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **25,000 route entries p.a.**

- It's a case of increasing growth, not just constant growth
  - More use of /48 more specifics
  - More networks advertising IPv6 prefixes

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Conclusions

V4 BGP Table Size Predictions

Date	RIB Size	Prediction
Jan *2017*	*646,000*	
2018	*699,000*	
2019	*760,000*	
2020	*814,000*	
2021	*866,000*	
2022		*916,000*
2023		*970,000*
2024		*1,024,000*
2025		*1,078,000*
2026		*1,132,000*

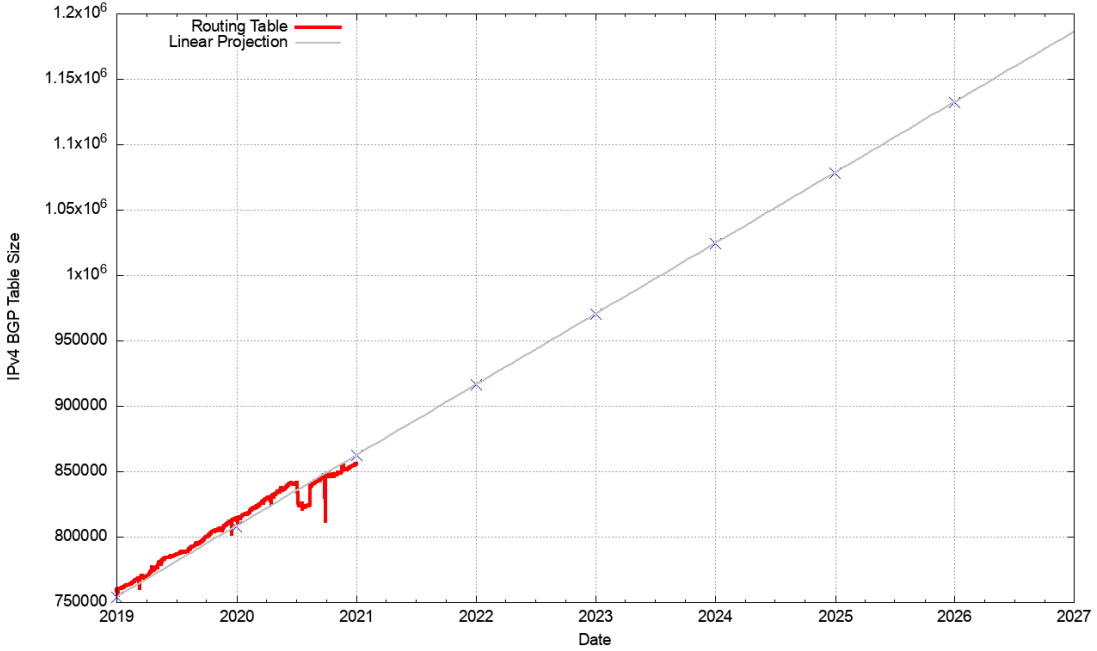

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2017*	*35,000*	
2018	*45,000*	
2019	*62,000*	
2020	*79,000*	
2021	*104,000*	
2022	*118,000*	*135,000*
2023	*136,000*	*175,000*
2024	*155,000*	*228,000*
2025	*174,000*	*296,000*
2026	*192,000*	*384,000*

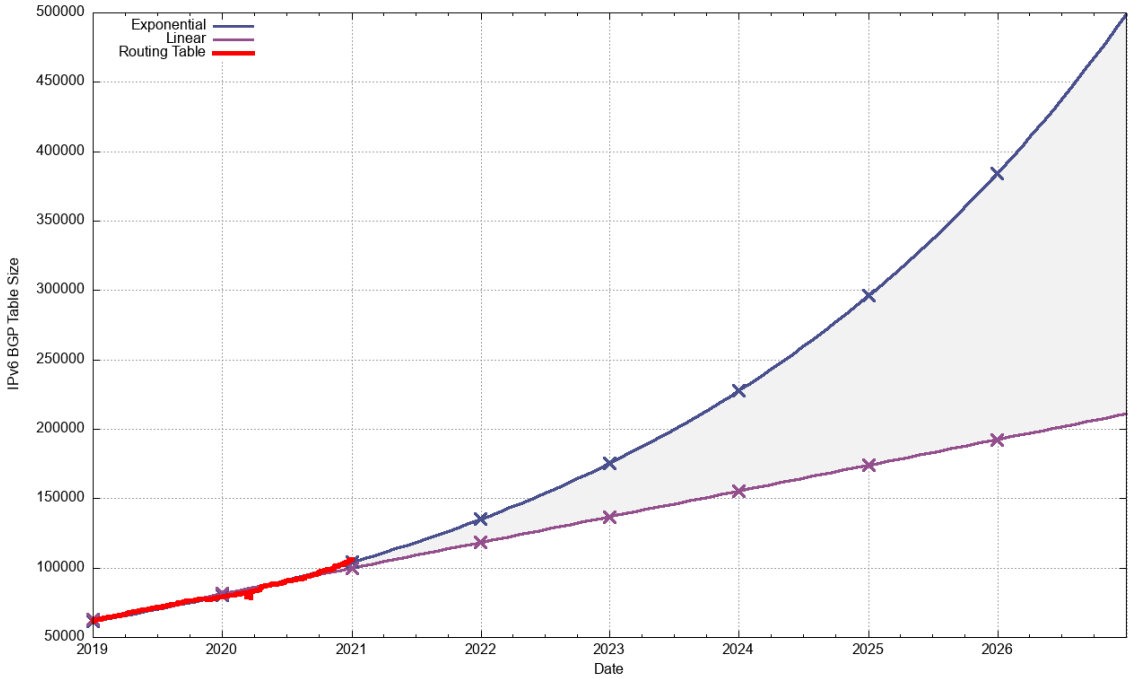

Note that the IPv6 tables are 128bits wide – i.e. 4x the size of the IPv4 tables!

BGP Table Growth

The absolute size of the IPv6 routing table is growing much faster than the IPv4 table

They will require the same memory size in around 4 years time, given that each IPv6 entry is 4 times the memory size of an IPv4 entry

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

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Conclusions

IPv4 BGP Updates

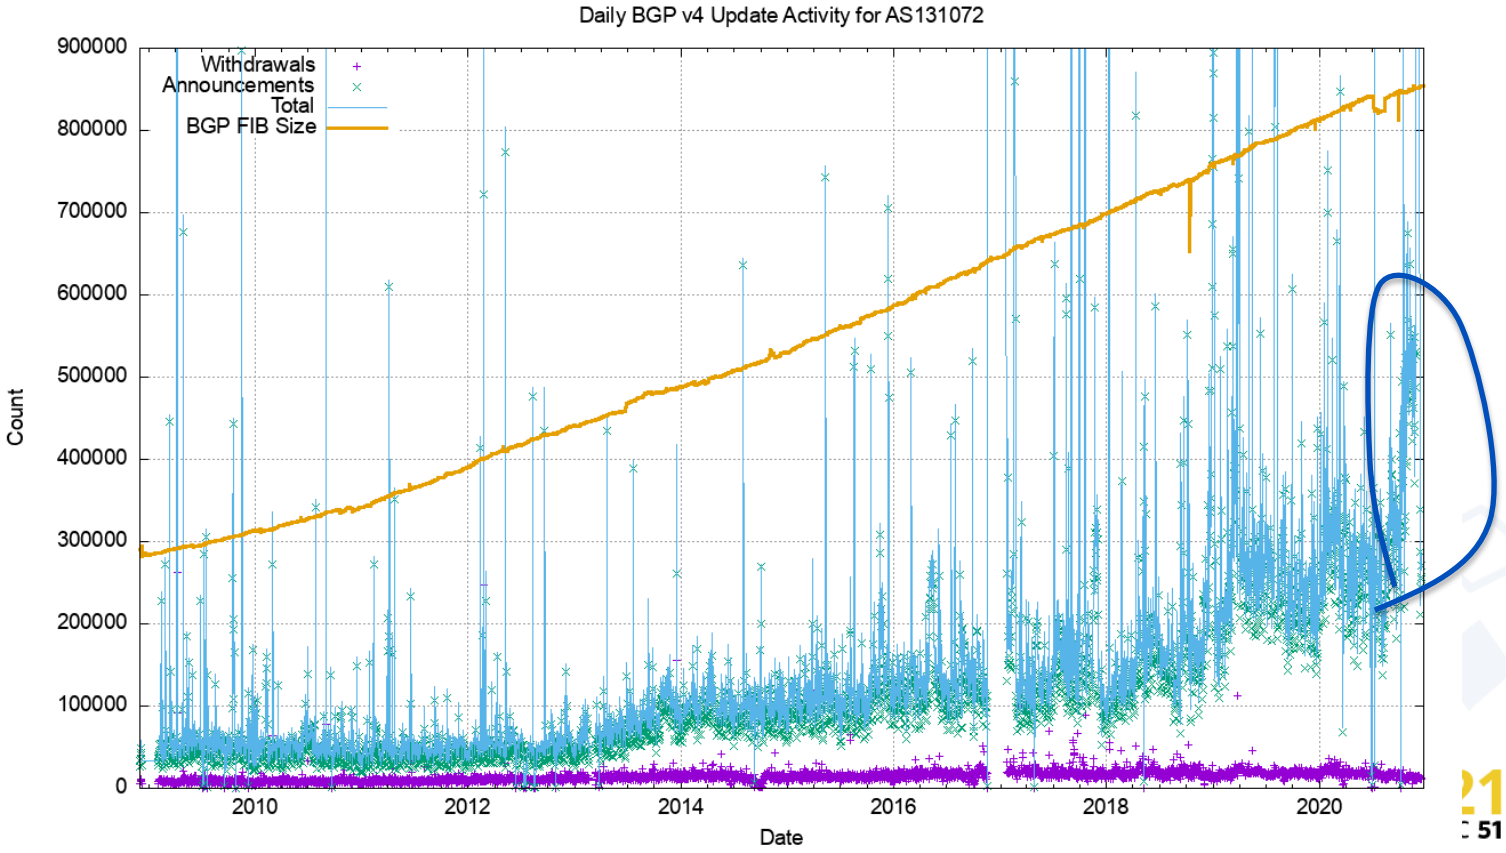

IPv4 BGP Convergence Performance

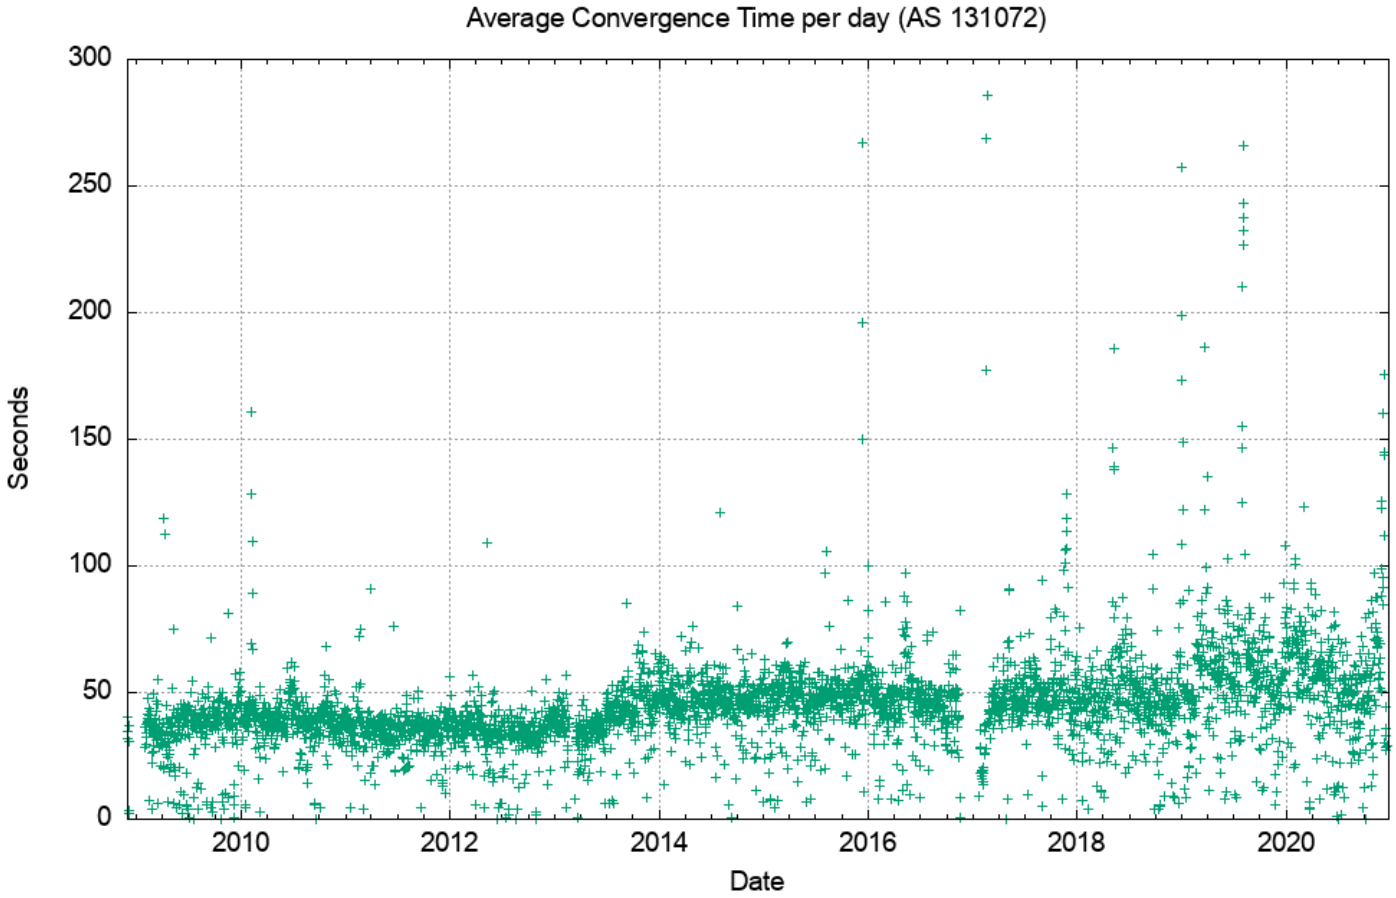

Updates in IPv4 BGP

The beginnings of concern …

- The number of updates per instability event and the time to converge has been relatively constant for many years – its now starting to rise, both in the number of unstable prefixes and the time to converge
- 20% of prefixes generate 80% of all updates. Less than 5% of all origin networks are linked to 80% of all updates. Instability is still concentrated in a small number of pathological cases. Its just that the instability itself from these sources is getting worse.

V6 BGP Updates

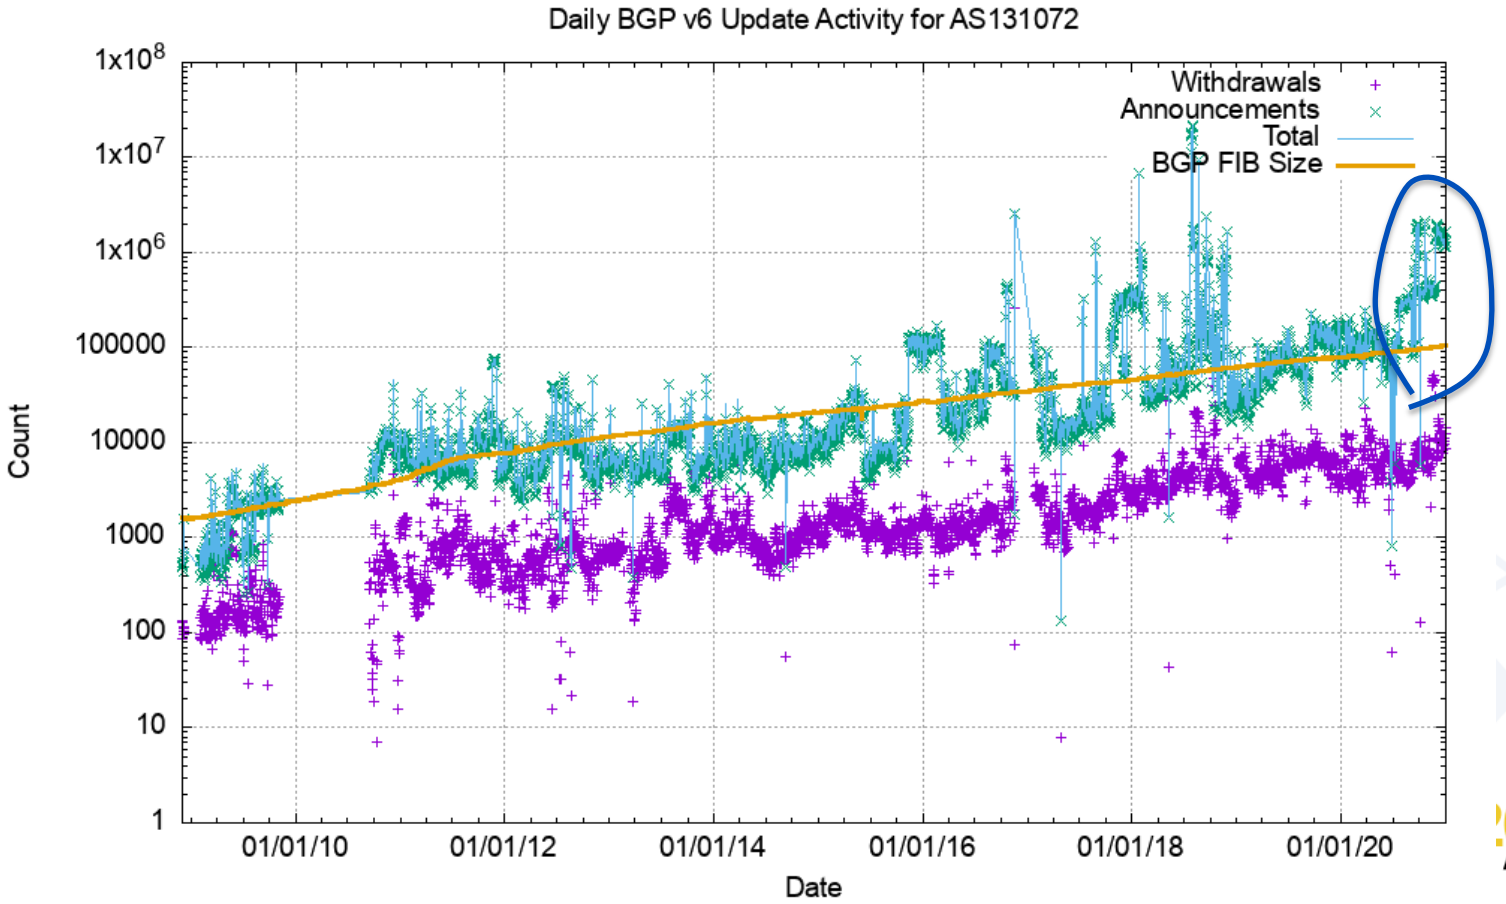

V6 Convergence Performance

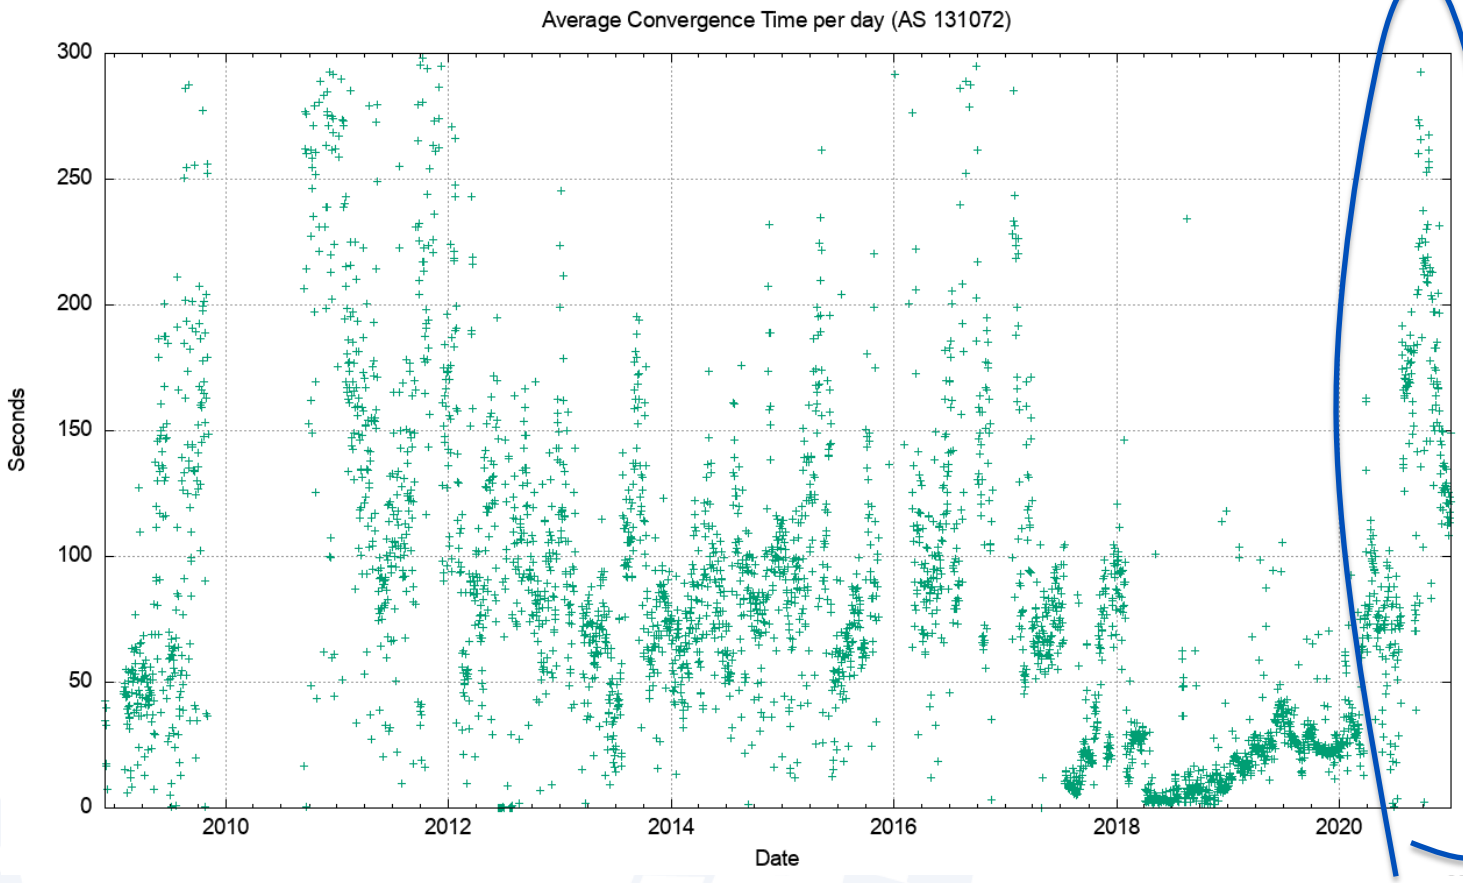

Updates in IPv6 BGP

Nobody is looking…

- Compared to IPv4, the IPv6 network exhibits a high level of routing instability, which is unexpected as the old overlay approaches are disappearing and the topology of IPv6 is now converging to the same topology as IPv4
- Just 5 AS's generate 30% of the BGP update load in the past 2 weeks. Instability is still concentrated in a small number of pathological cases. The instability itself from these sources is getting worse.

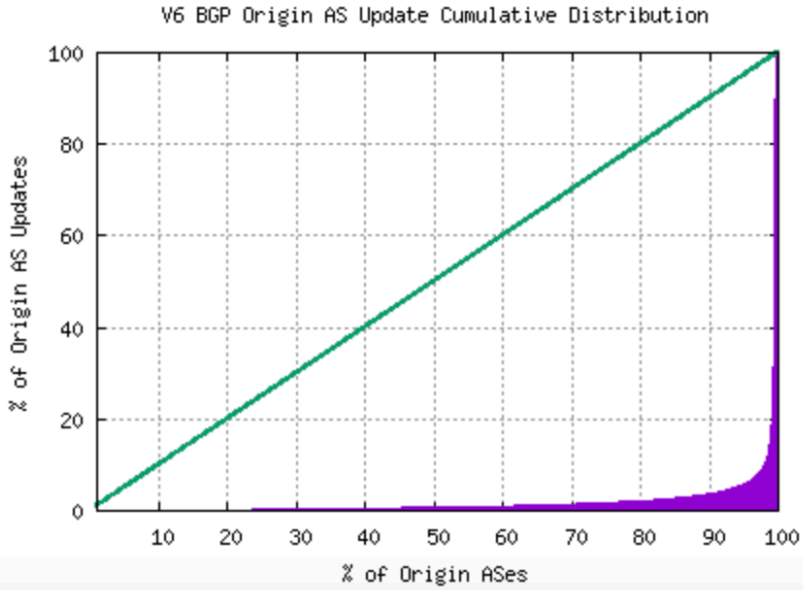

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Conclusions

Routing Futures

- There is still 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. Instability levels are rising, generally driven by a small set of highly unstable “super generators”

- 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-capacity line cards

Some Practical Suggestions

- Understand your hardware's high speed FIB capacity in the default-free parts of your network
- Review your IPv4 / IPv6 portioning - a dual-stack eBGP router will need 970,000 32-bit IPv4 slots and 175,000 128-bit IPv6 slots for a full eBGP routing table in line cards over the coming 24 months if they are using a full eBGP FIB load (plus internal routes of course)
- Judicious use of default routes in your internal network may allow you drop this requirement significantly
- Using a hot cache for line card FIB cache would reduce the memory requirement significantly without visible performance cost

That's it!

The Complete Pack

27 Years of Routing the Internet

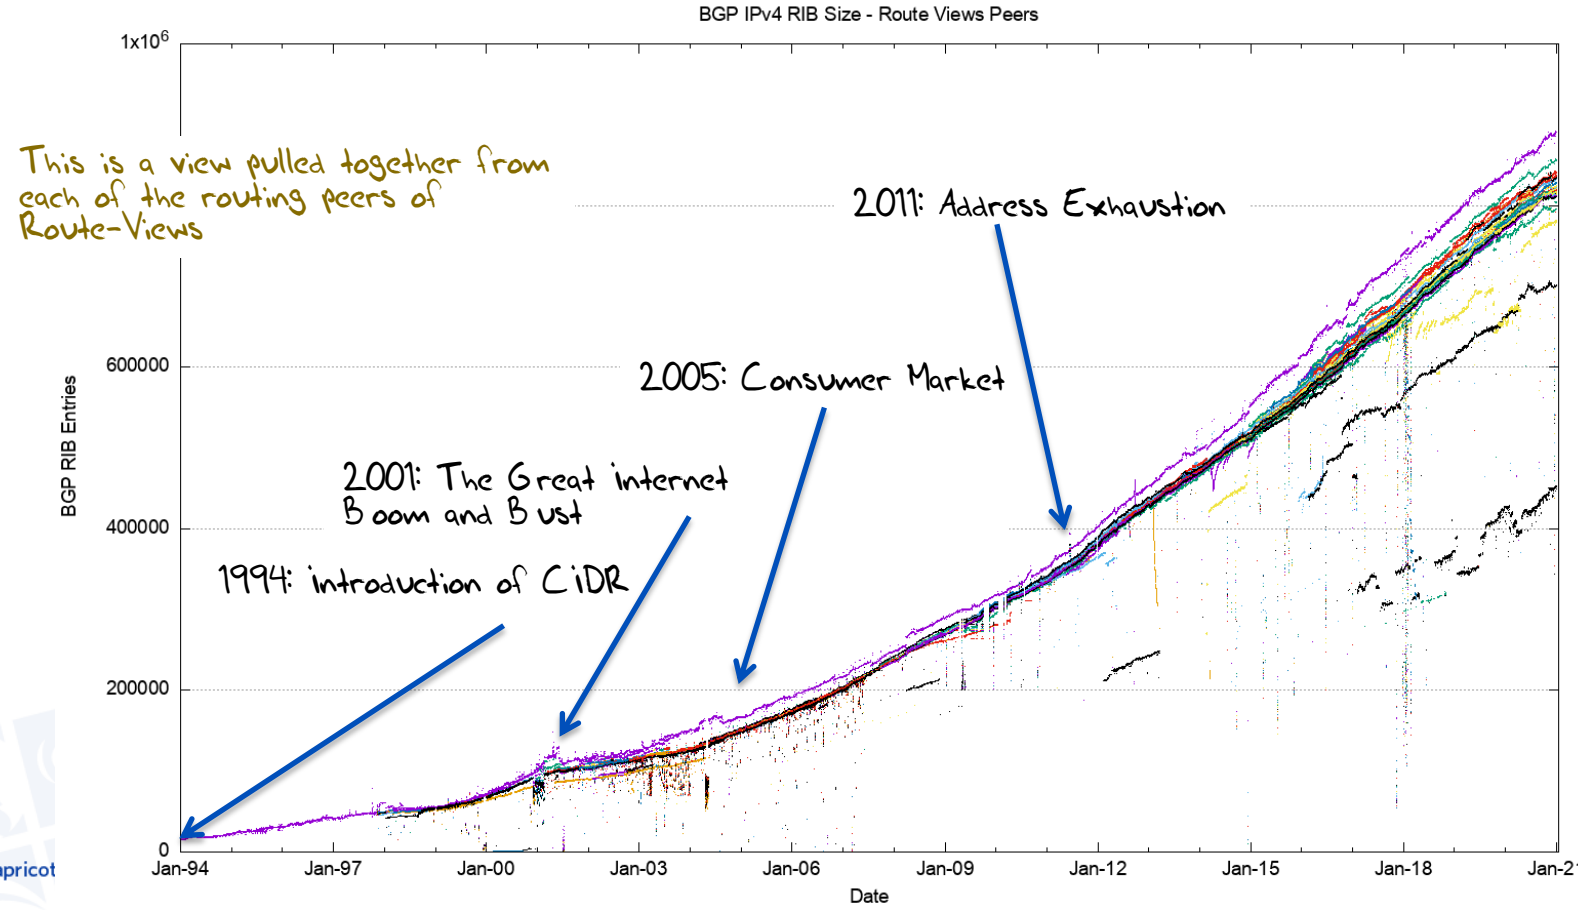

2017-2020 in detail

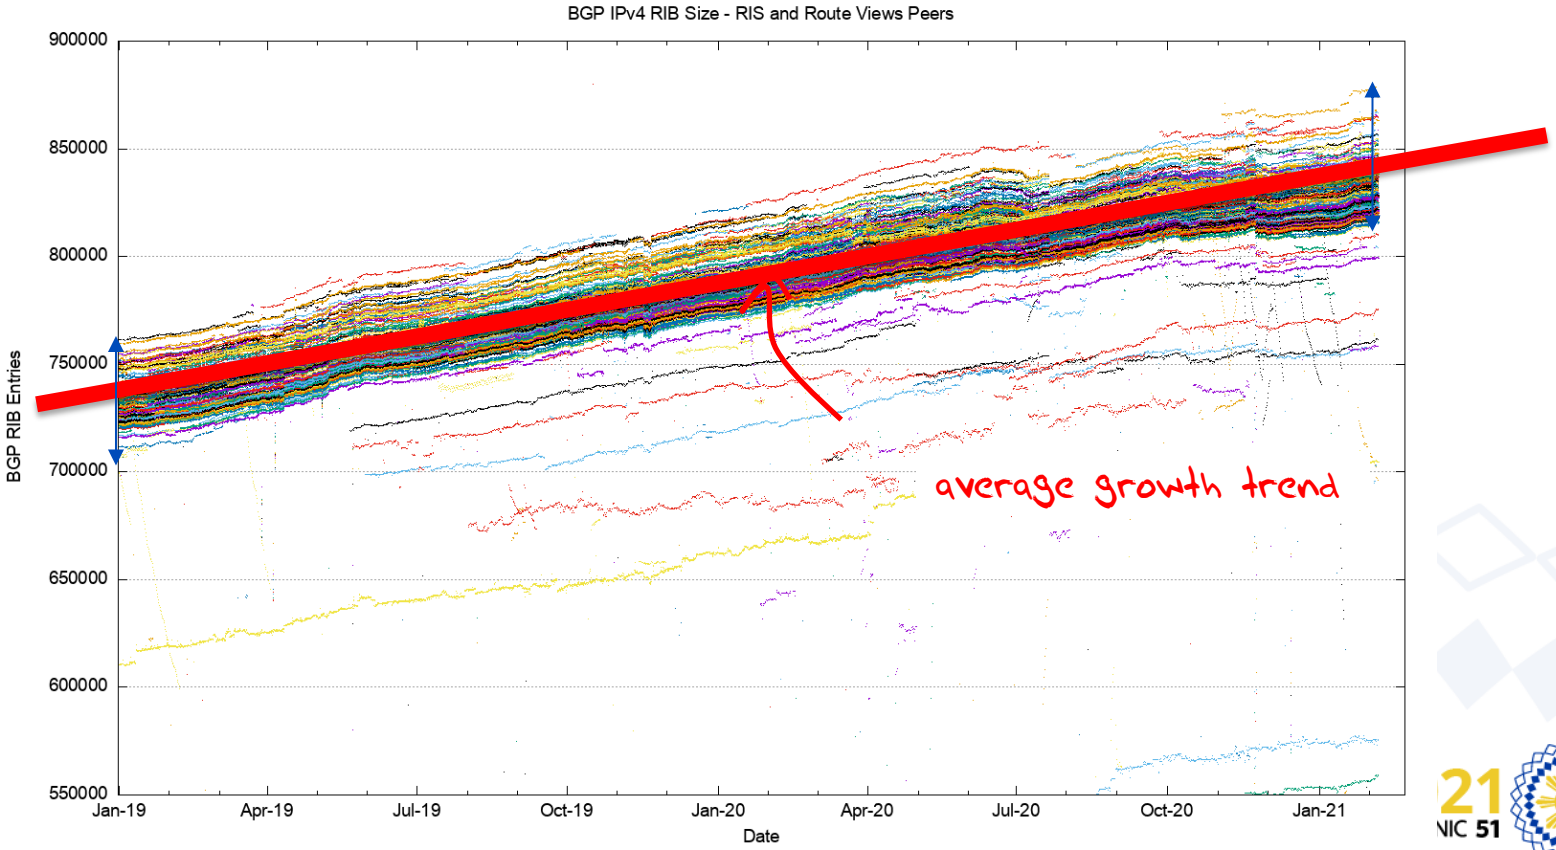

Routing Indicators for IPv4

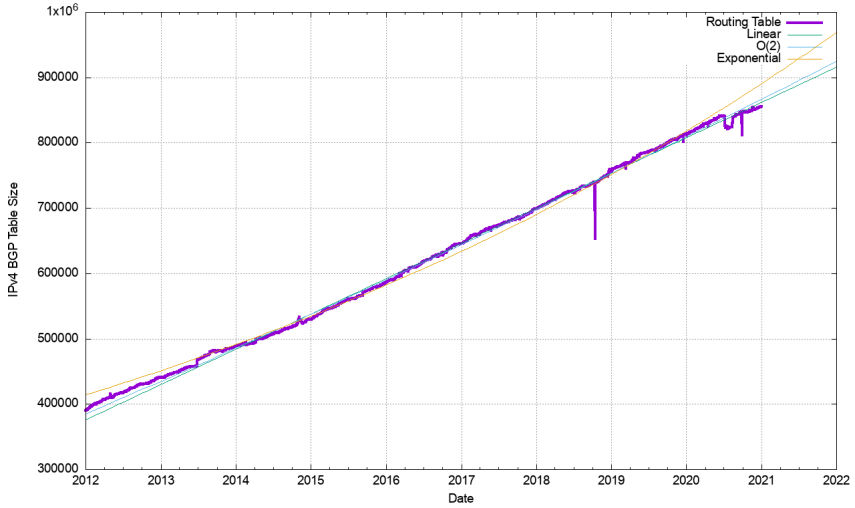

Routing prefixes – growing by some 51,000 prefixes per year

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

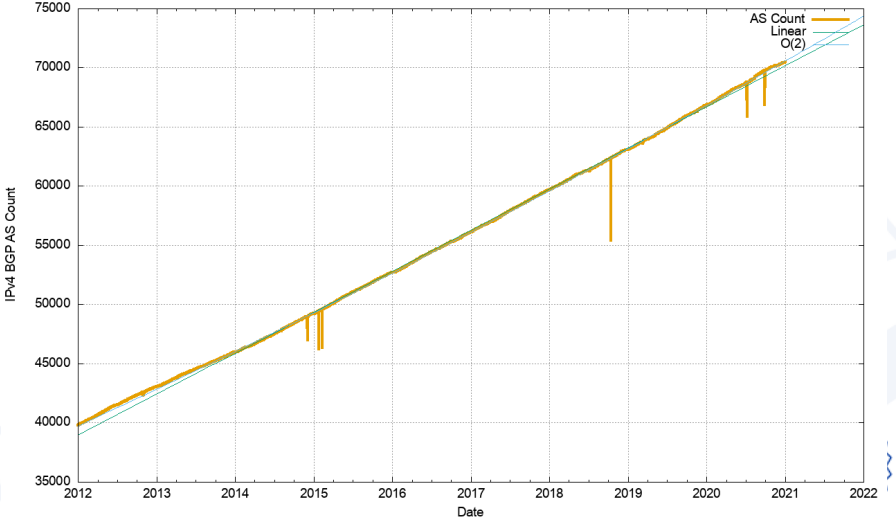

Routing Indicators for IPv4

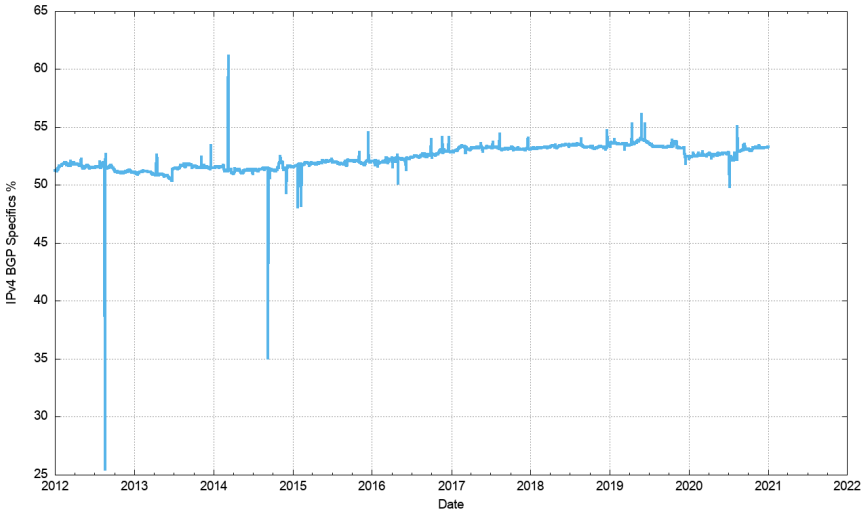

More Specifics are still taking up slightly more than one half of the routing table

But the average size of a routing advertisement continues to shrink

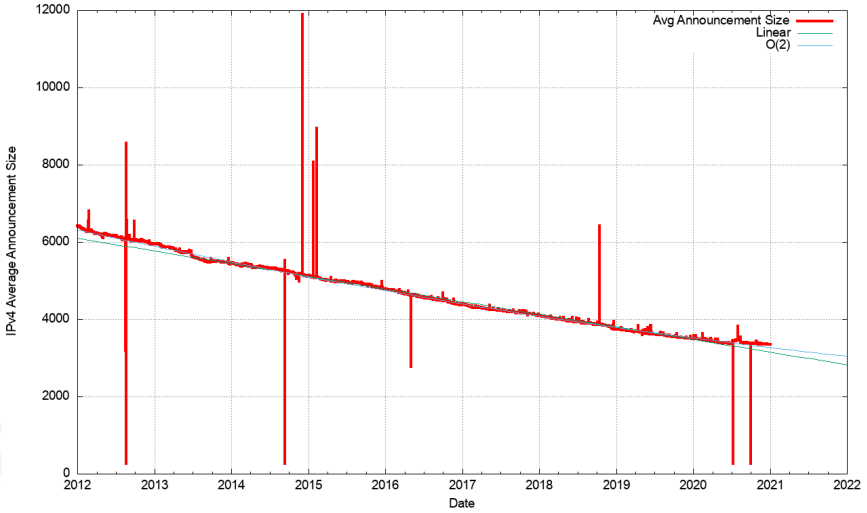

Routing Indicators for IPv4

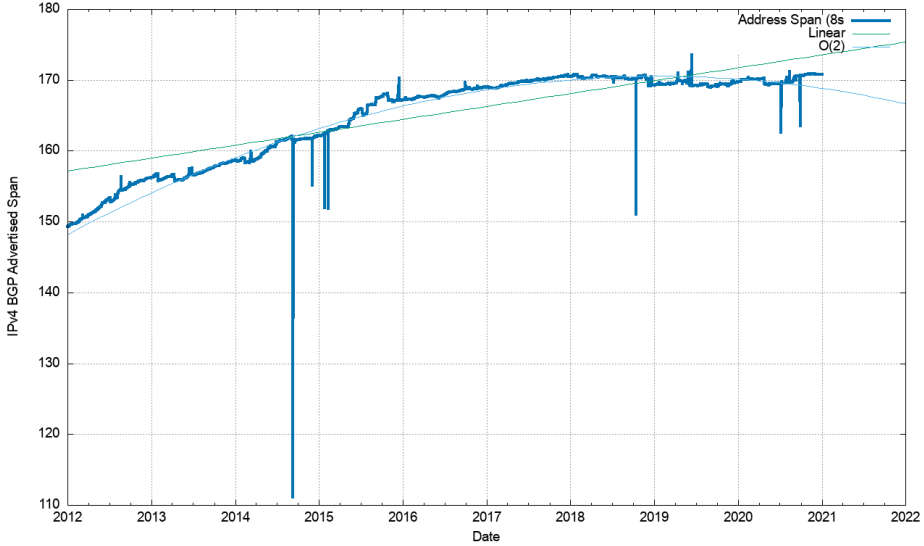

Address Exhaustion is now visible in the extent of advertised address space

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

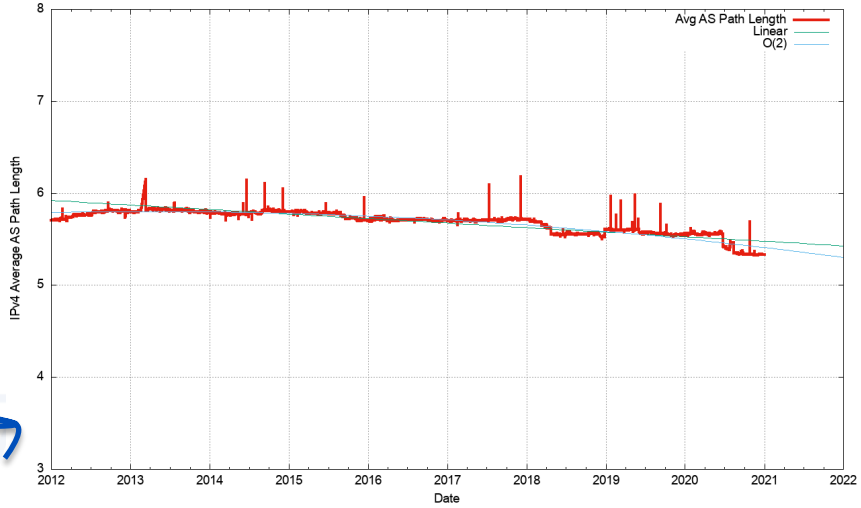

AS Adjacencies (AS131072)

57,343 out of 70,532 ASNs have 1 or 2 AS Adjacencies (82%)
2,342 ASNs have 10 or more adjacencies
9 ASNs have >1,000 adjacencies

5,727	AS6939	HURRICANE - Hurricane Electric, Inc., US
4,433	AS3356	LEVEL3 - Level 3 Communications, Inc., US
3,707	AS174	COGENT-174 - Cogent Communications, US
1,896	AS6461	ZAYO Bandwidth, US
1,711	AS7018	ATT-INTERNET4 - AT&T Services, Inc., US
1,701	AS3257	GTT-Backbone, DE
1,349	AS2914	NTT America, US
1293	AS1299	TELIANET Telia Carrier, SE
1,148	AS3549	LVLT-3549, US

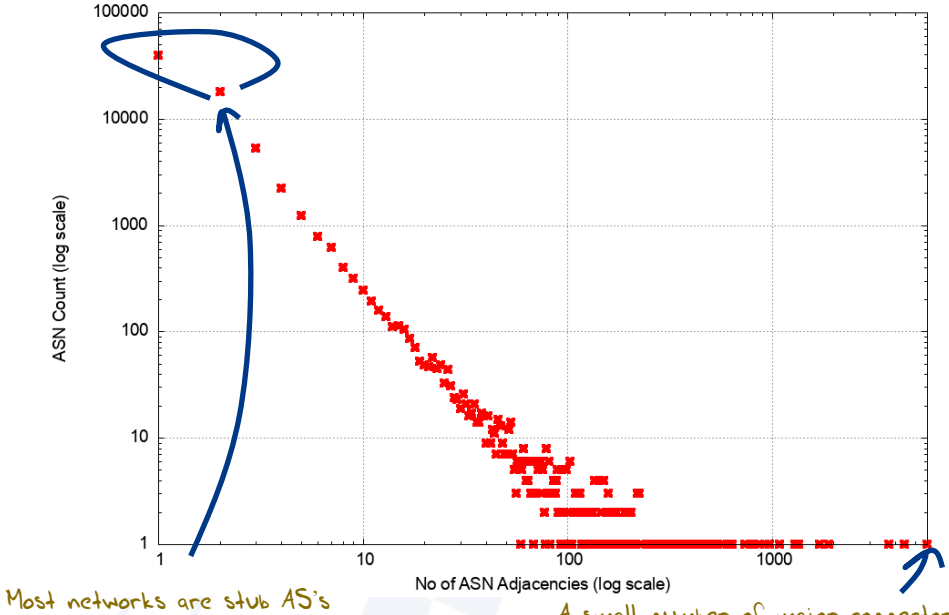

What happened in 2020 in V4?

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 860,000 by the end of 2020
- The pace of growth of the routing table was slightly lower than the rolling 5-year average, with 52,000 new entries in 2020
- The AS position was steady with 3,400 new AS's per year
- Transit relationships have not changed materially over 2020 for most networks
- IPv4 address exhaustion is not changing this picture as yet
- Instead, we appear to be advertising shorter prefixes into the routing system

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

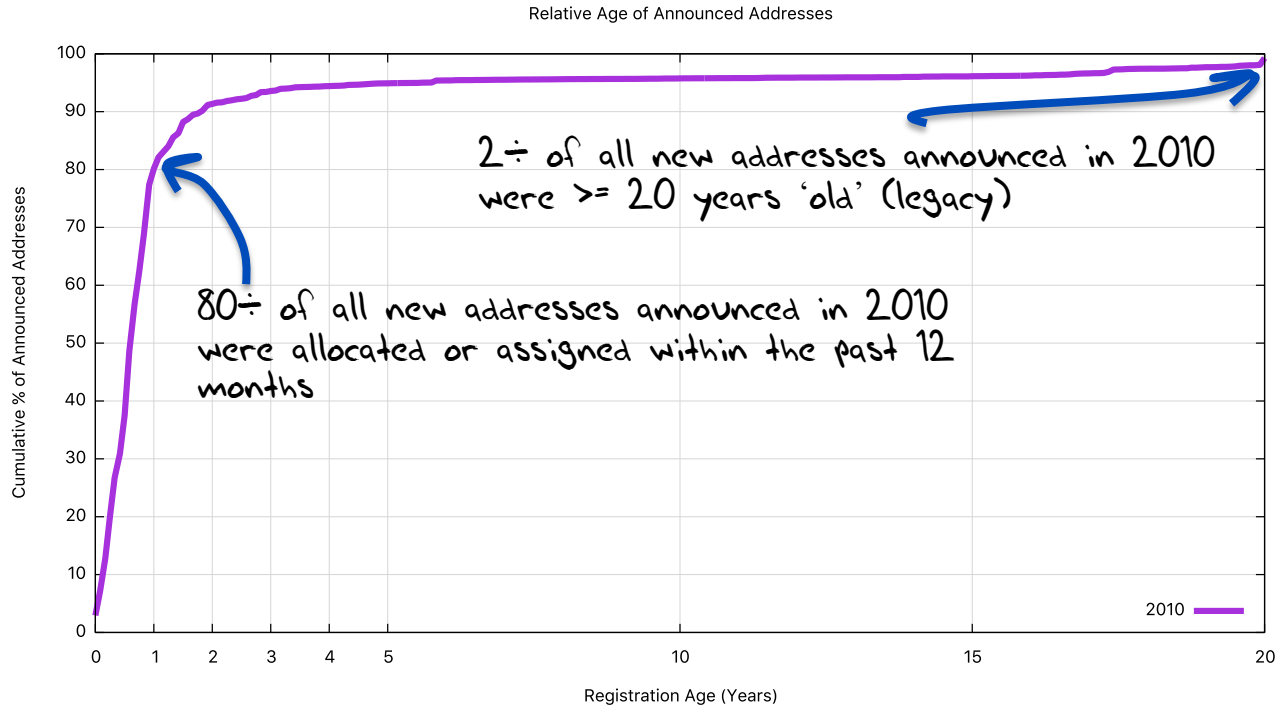

Advertised Address "Age"

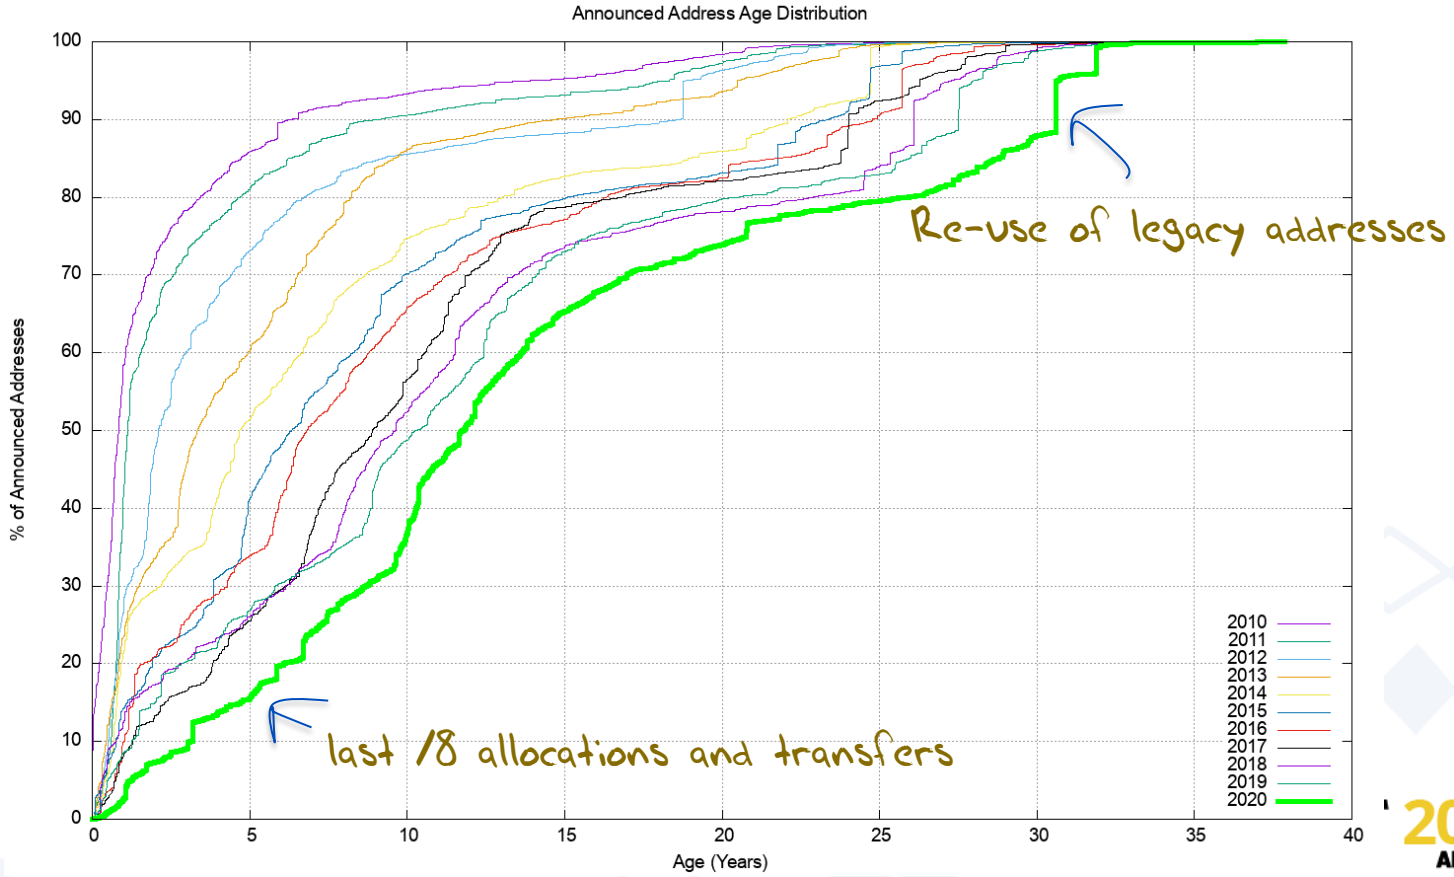

IPv4 Advertised vs Unadvertised

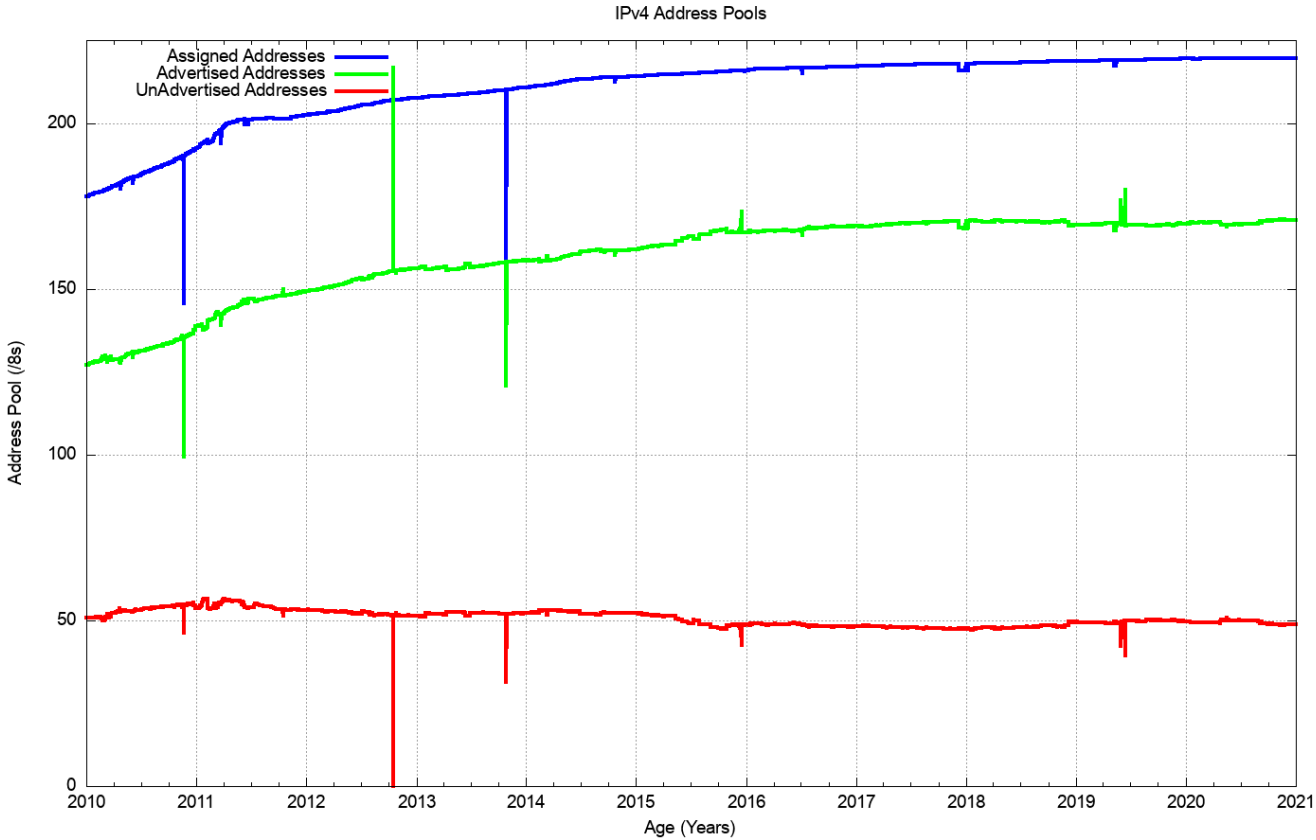

2005 - 2020: Unadvertised Addresses

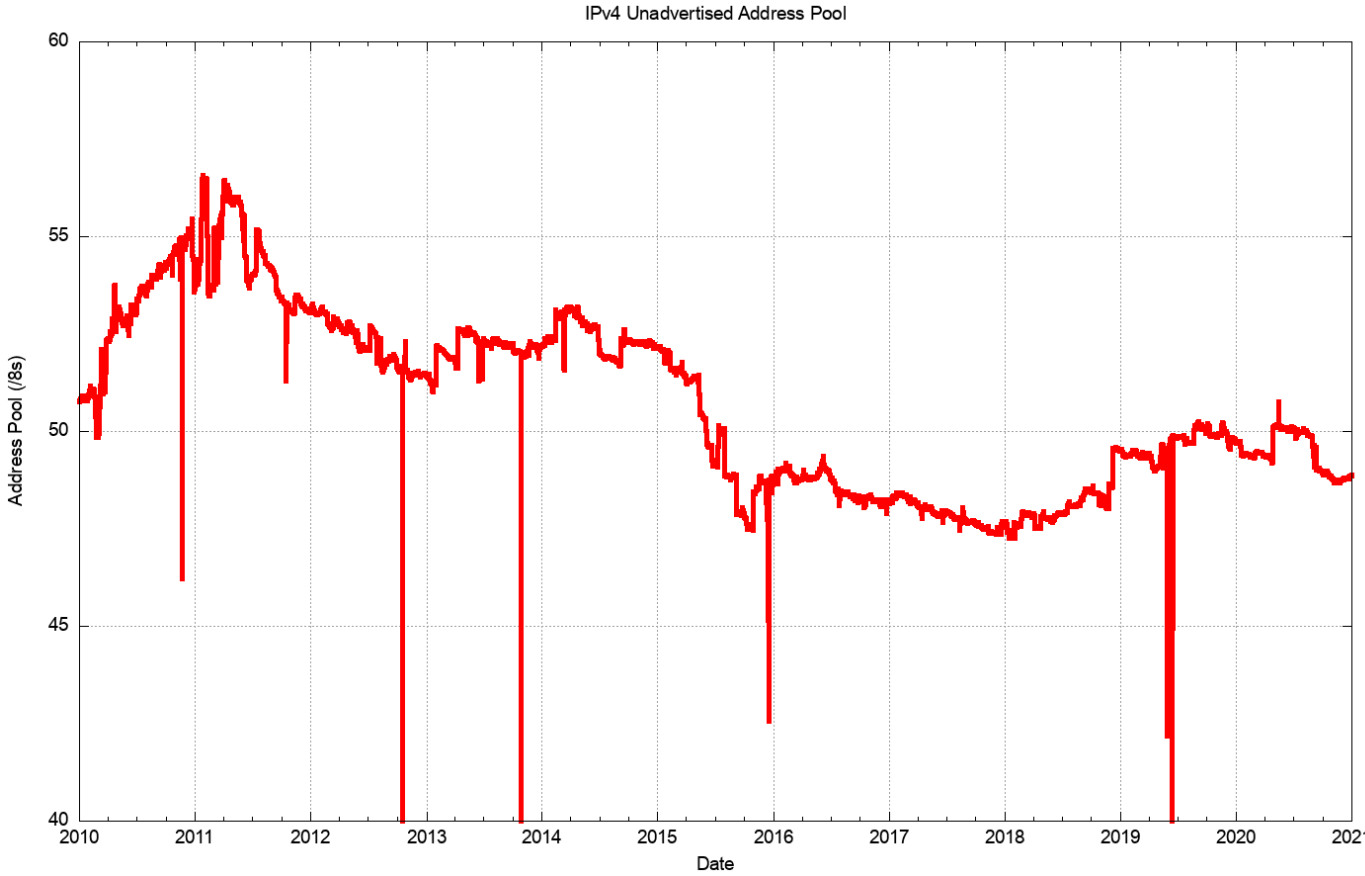

2019: Assigned vs Recovered

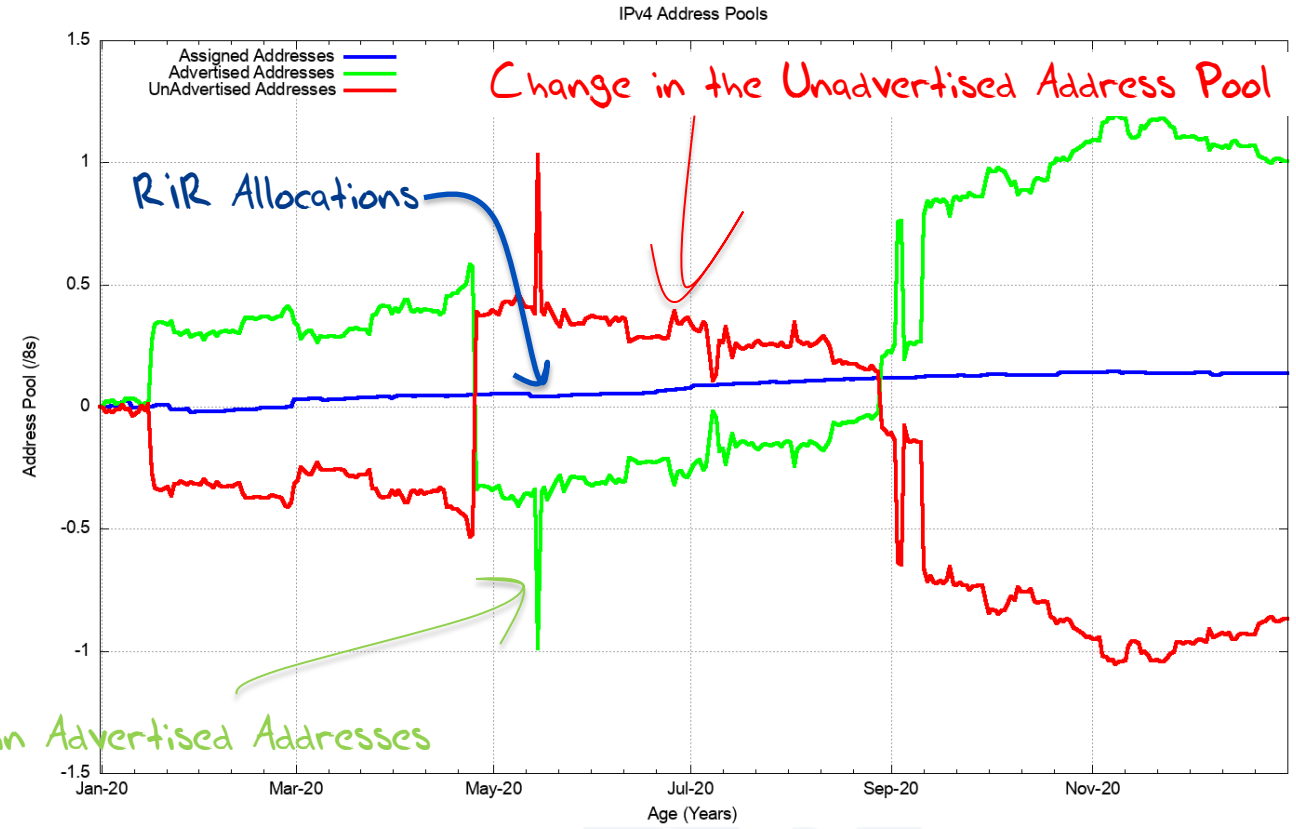

V4 in 2020

- 16,855,124 addresses were **added** to the routing table across 2020
- 2,306,432 addresses were assigned by RIRs in 2019
- And a net of 14,548,692 addresses were drawn from the pool of unadvertised addresses

In 2020 we saw legacy blocks transferring away from ISPs / end user sites and heading towards cloud SPs.

The Route-Views View of IPv6

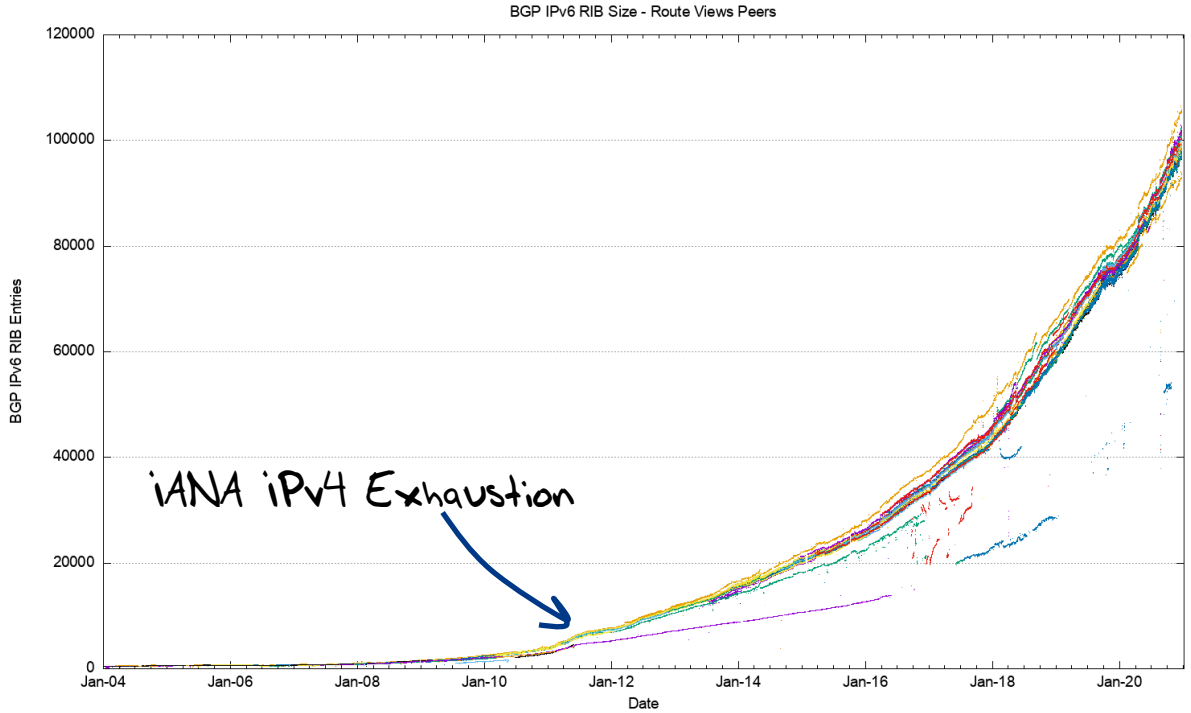

2018-2019 in Detail

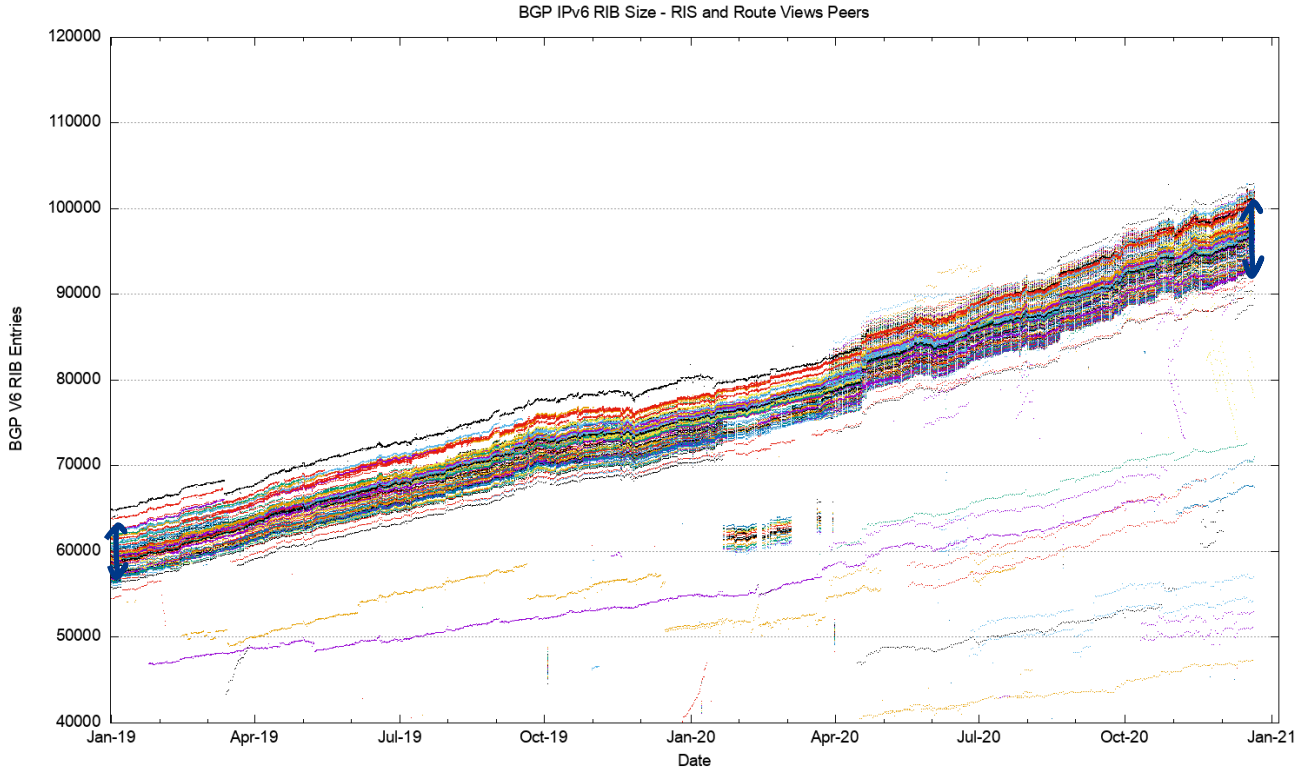

Routing Indicators for IPv6

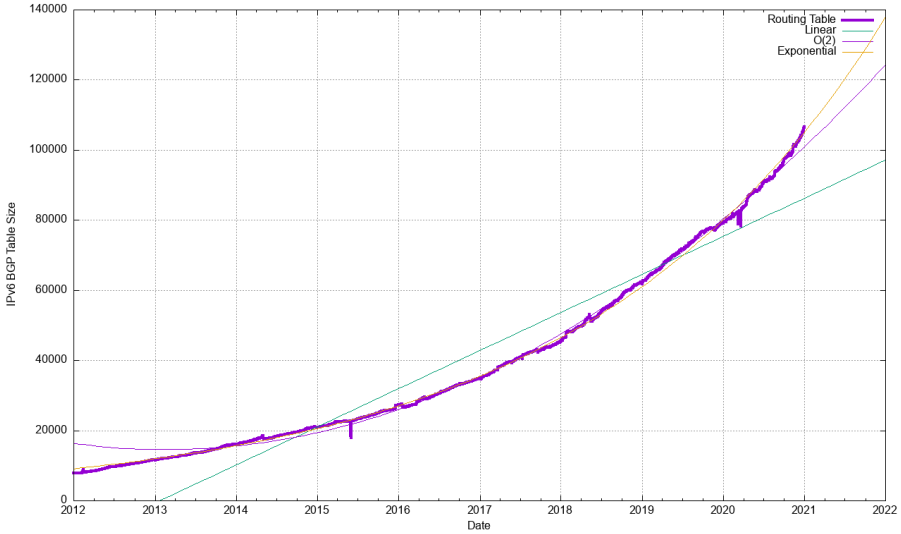

Routing prefixes – growing by some 25,000 prefixes per year

AS Numbers – growing by some 2,200 ASNs per year (which is 60÷ the V4 growth)

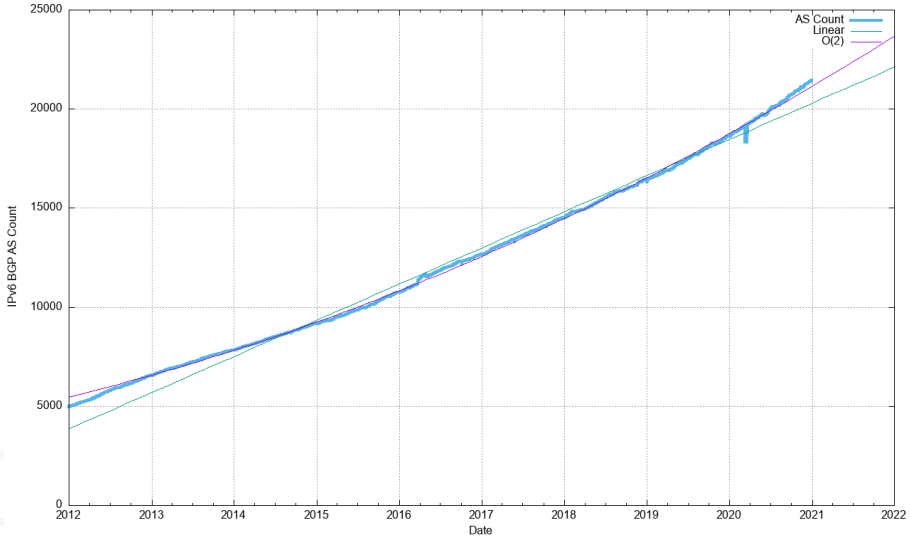

Routing Indicators for IPv6

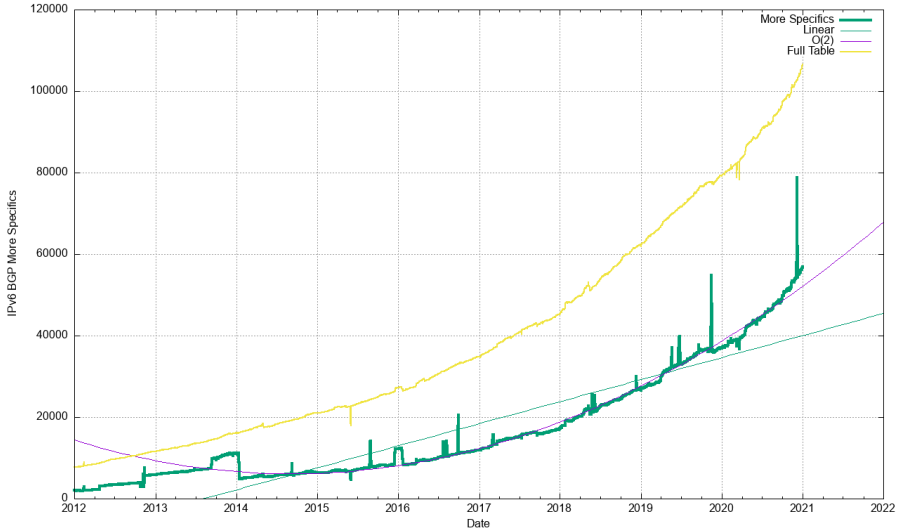

More Specifics now take up one half of the routing table

The average size of a routing advertisement is getting smaller

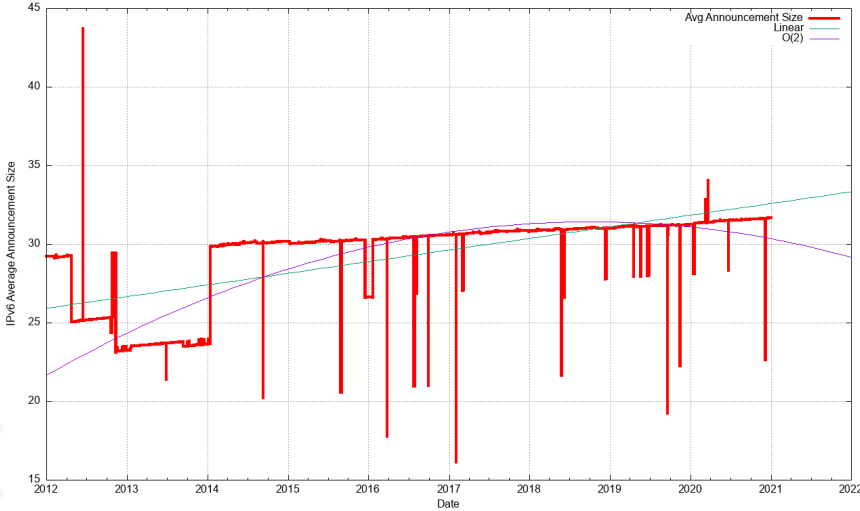

Routing Indicators for IPv6

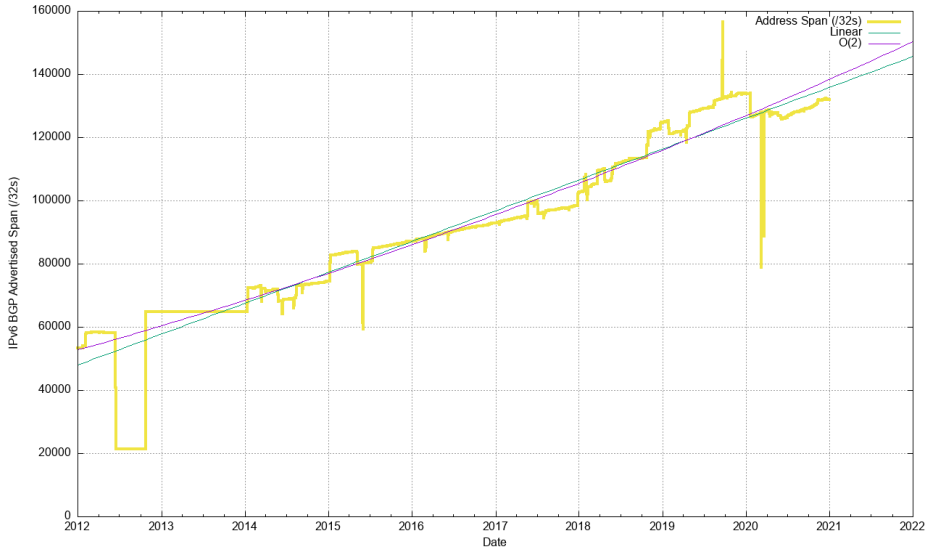

Advertised Address span is growing at an exponential rate

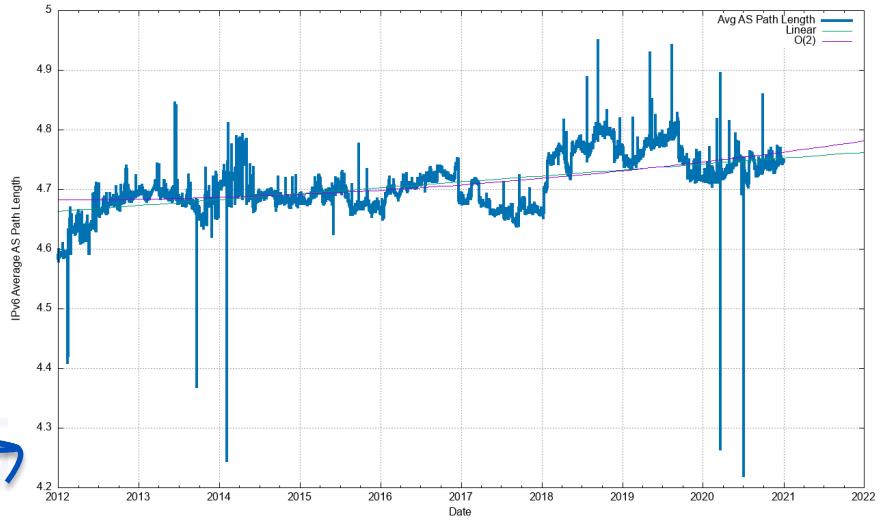

The "shape" of inter-AS interconnection in IPv6 is rising slightly. Local connections appear to be replacing overlay trunk transits

AS Adjacencies (AS131072)

16,961 out of 21,489 ASNs have 1 or 2 AS Adjacencies (80%)
798 ASNs have 10 or more adjacencies
4 ASNs have >1,000 adjacencies

5,177	AS6939	HURRICANE - Hurricane Electric, Inc., US
1,315	AS3356	LEVEL3 - Level 3 Communications, Inc., US
1,196	AS174	COGENT-174 - Cogent Communications, US
1,191	AS1299	Telia Carrier, SE

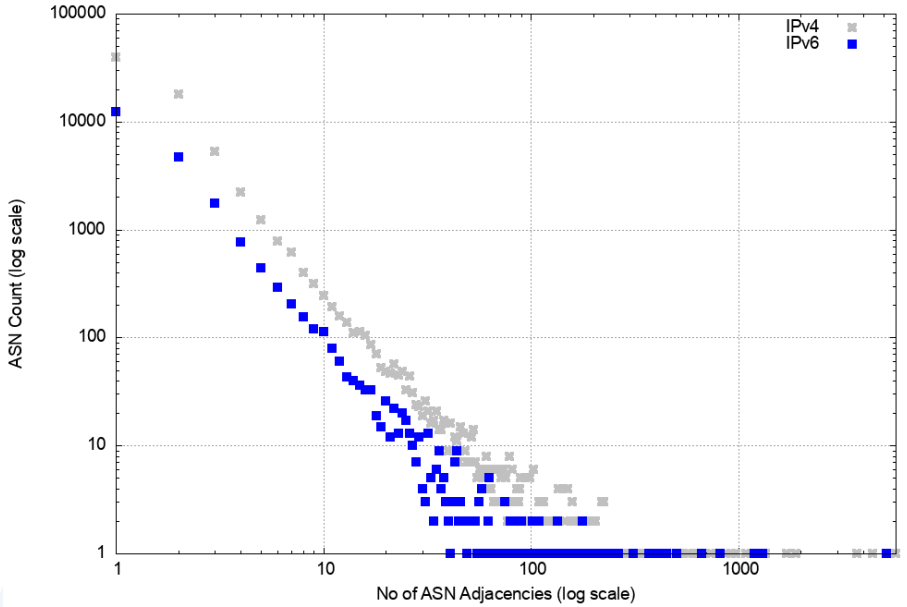

V6 in 2020

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **25,000 route entries p.a.**

- It's a case of increasing growth, not just constant growth
  - More use of /48 more specifics
  - More networks advertising IPv6 prefixes

What to expect

BGP Size Projections

How quickly is the routing space growing?

What are the projections of future BGP FIB size?

V4 - Daily Growth Rates

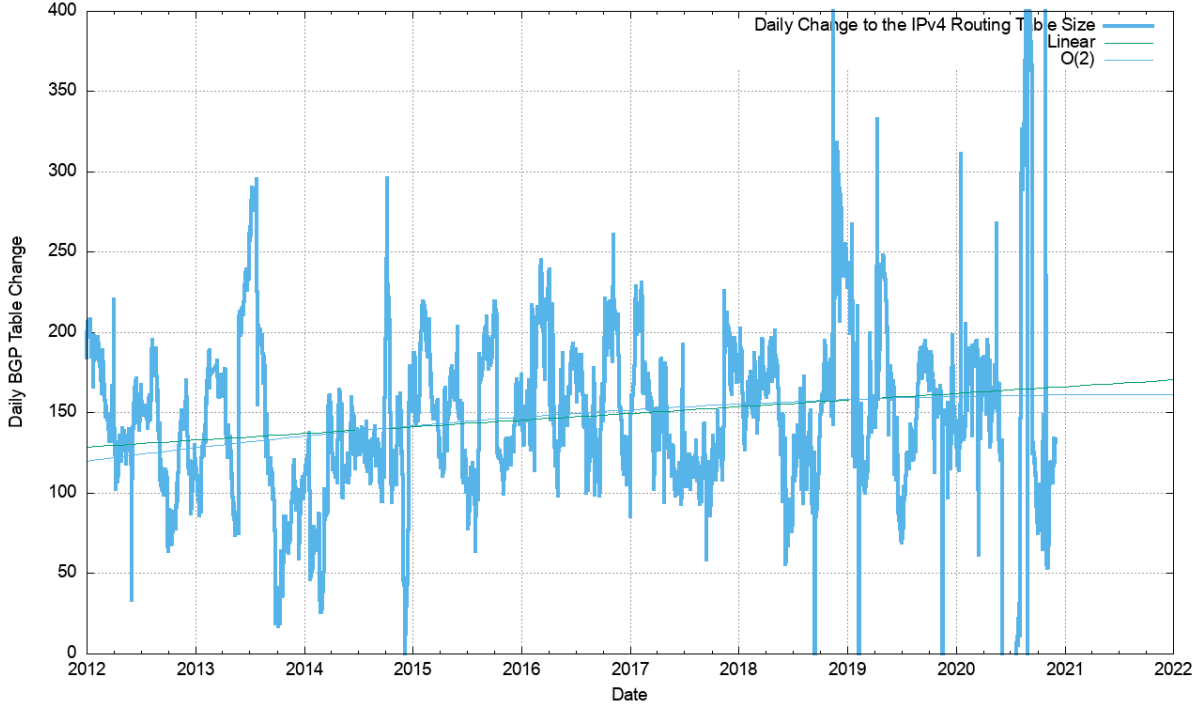

Growth in the V4 network appears to be constant at a long-term average of 150 additional routes per day, or some 54,000 additional routes per year

V4 BGP Table Size Predictions

Date	RIB Size	Prediction
Jan *2017*	*646,000*	
2018	*699,000*	
2019	*760,000*	
2020	*814,000*	
2021	*866,000*	
2022		*916,000*
2023		*970,000*
2024		*1,024,000*
2025		*1,078,000*
2026		*1,132,000*

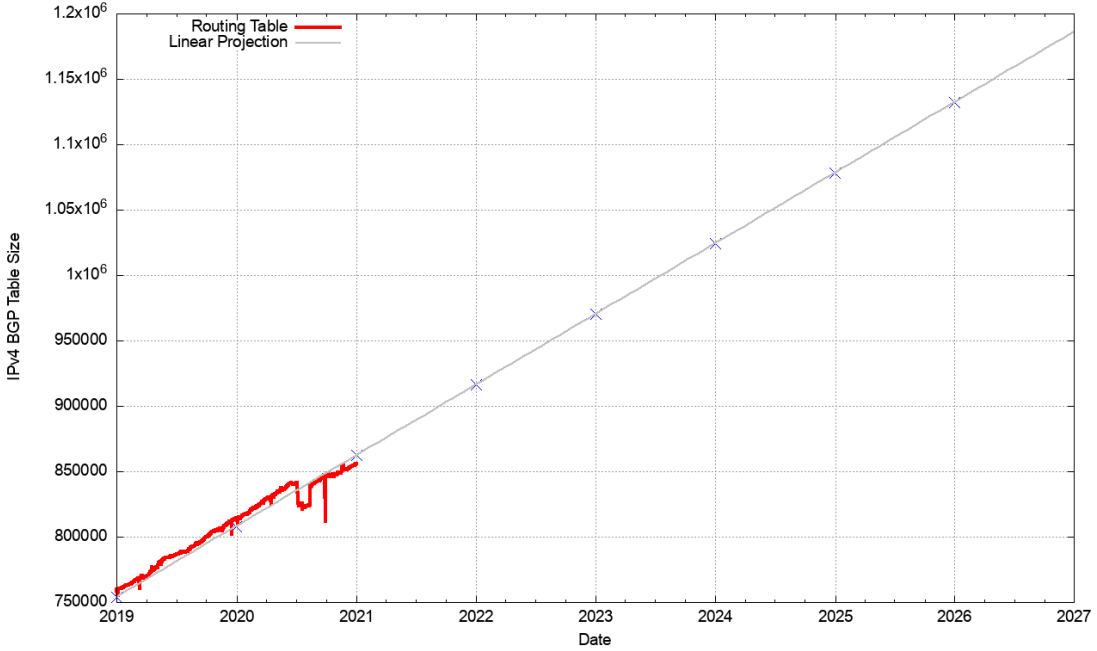

V6 - Daily Growth Rates

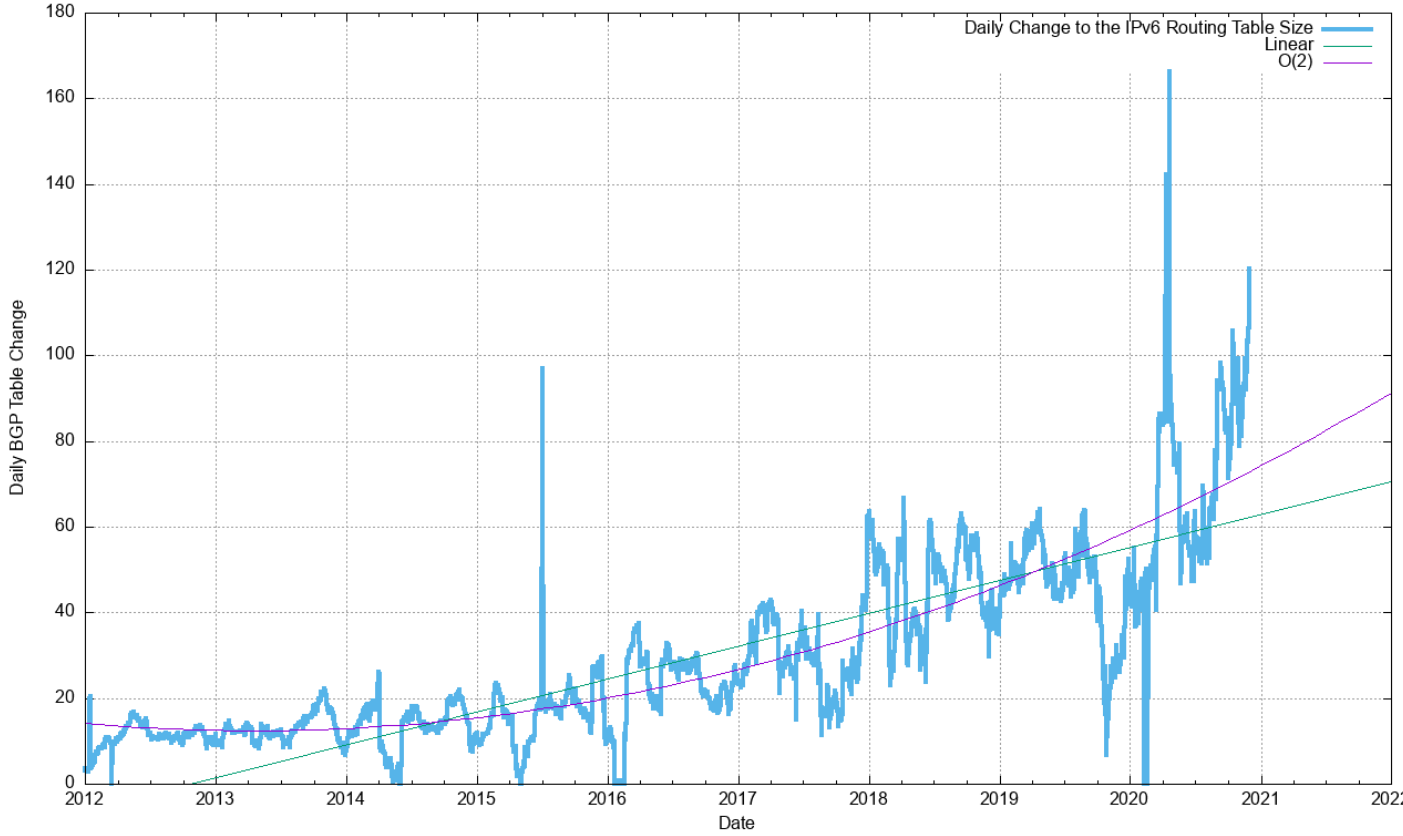

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2017*	*35,000*	
2018	*45,000*	
2019	*62,000*	
2020	*79,000*	
2021	*104,000*	
2022	*118,000*	*135,000*
2023	*136,000*	*175,000*
2024	*155,000*	*228,000*
2025	*174,000*	*296,000*
2026	*192,000*	*384,000*

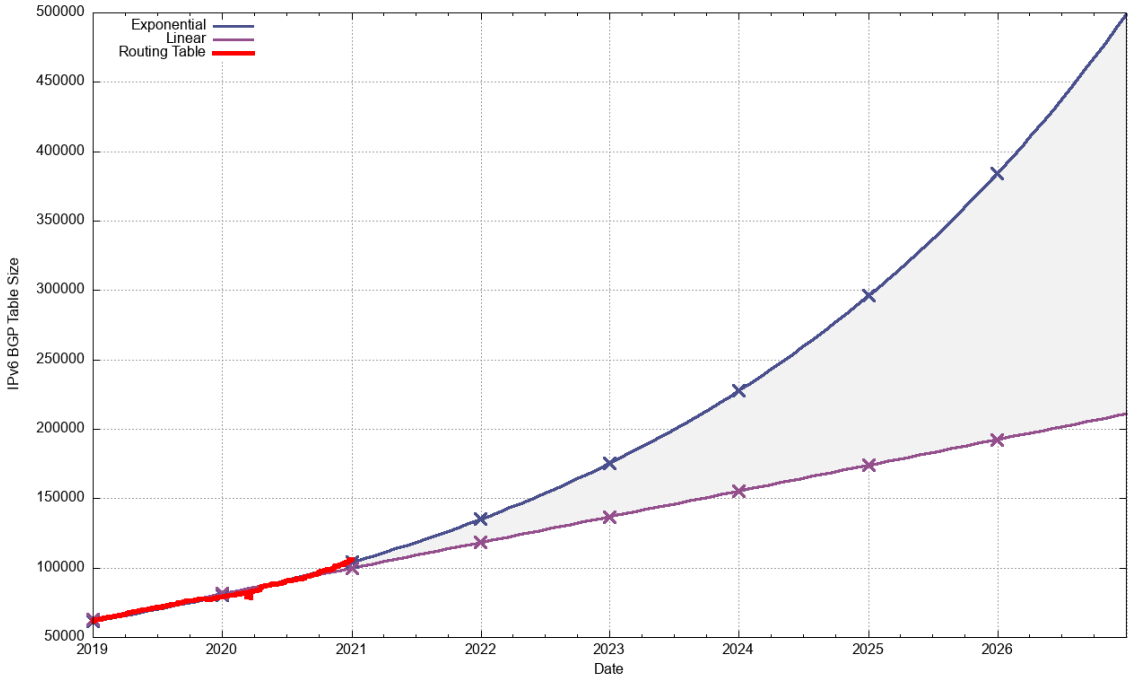

BGP Table Growth

The absolute size of the IPv6 routing table is growing much faster than the IPv4 table

They will require the same memory size in around 4 years time, given that each IPv6 entry is 4 times the memory size of an IPv4 entry

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

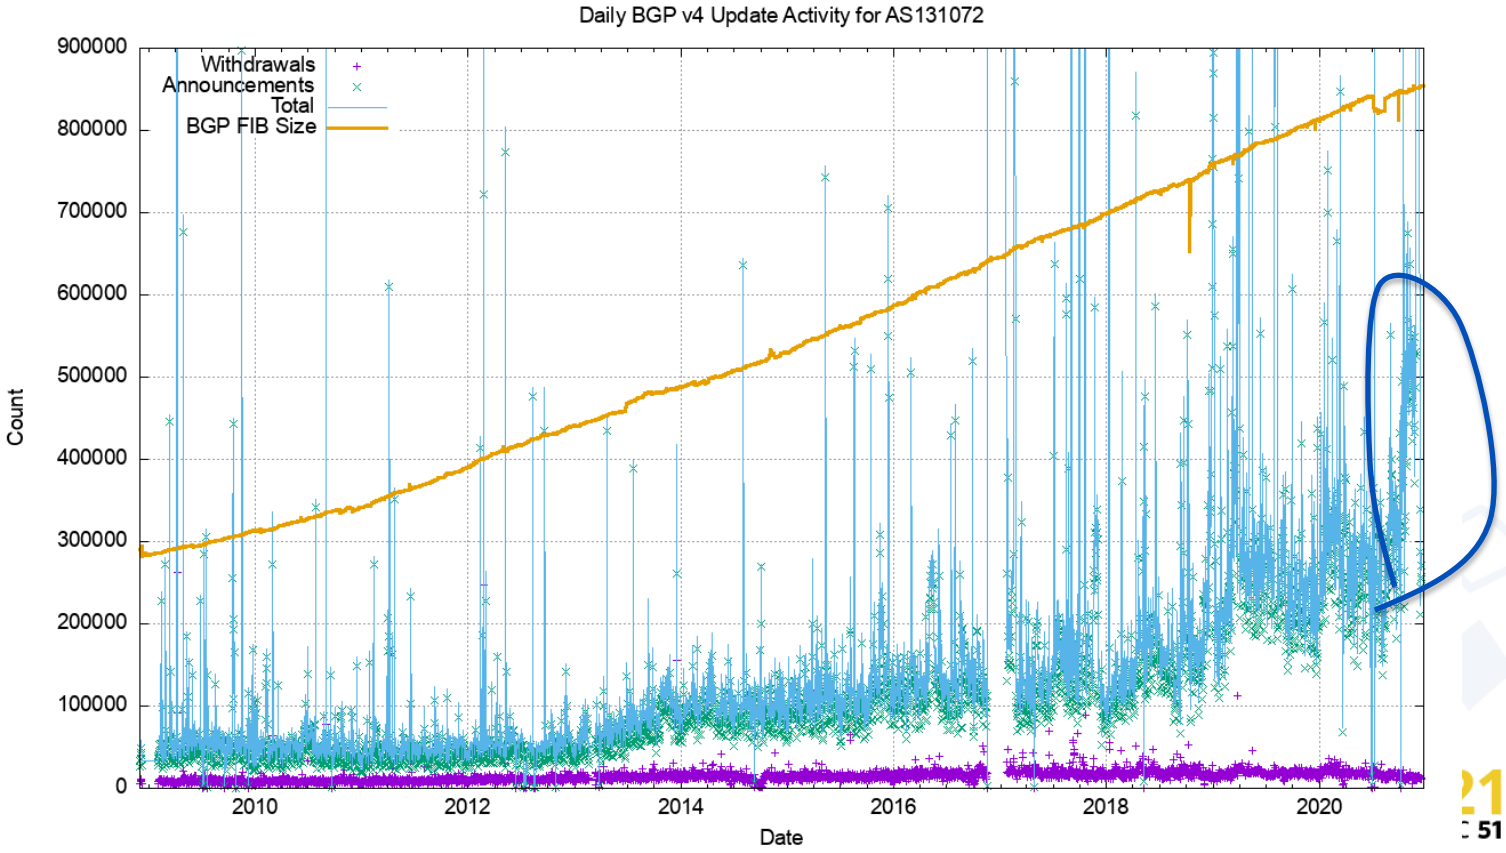

IPv4 BGP Convergence Performance

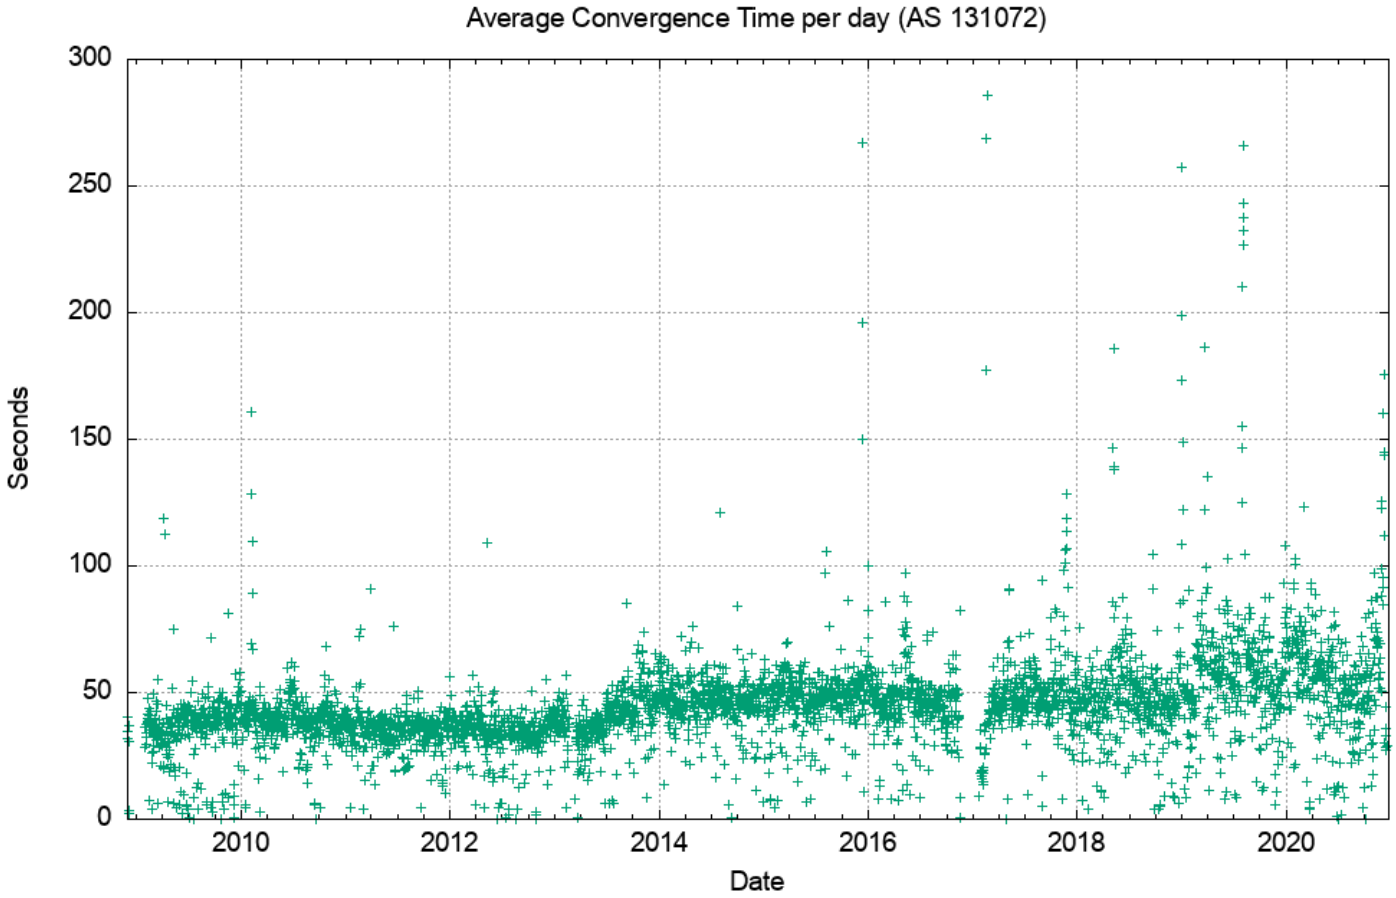

Updates in IPv4 BGP

The beginnings of concern …

- The number of updates per instability event and the time to converge has been relatively constant for many years – its now starting to rise, both in the number of unstable prefixes and the time to converge
- 20% of prefixes generate 80% of all updates. Less than 5% of all origin networks are linked to 80% of all updates. Instability is still concentrated in a small number of pathological cases. Its just that the instability itself from these sources is getting worse.

V6 BGP Updates

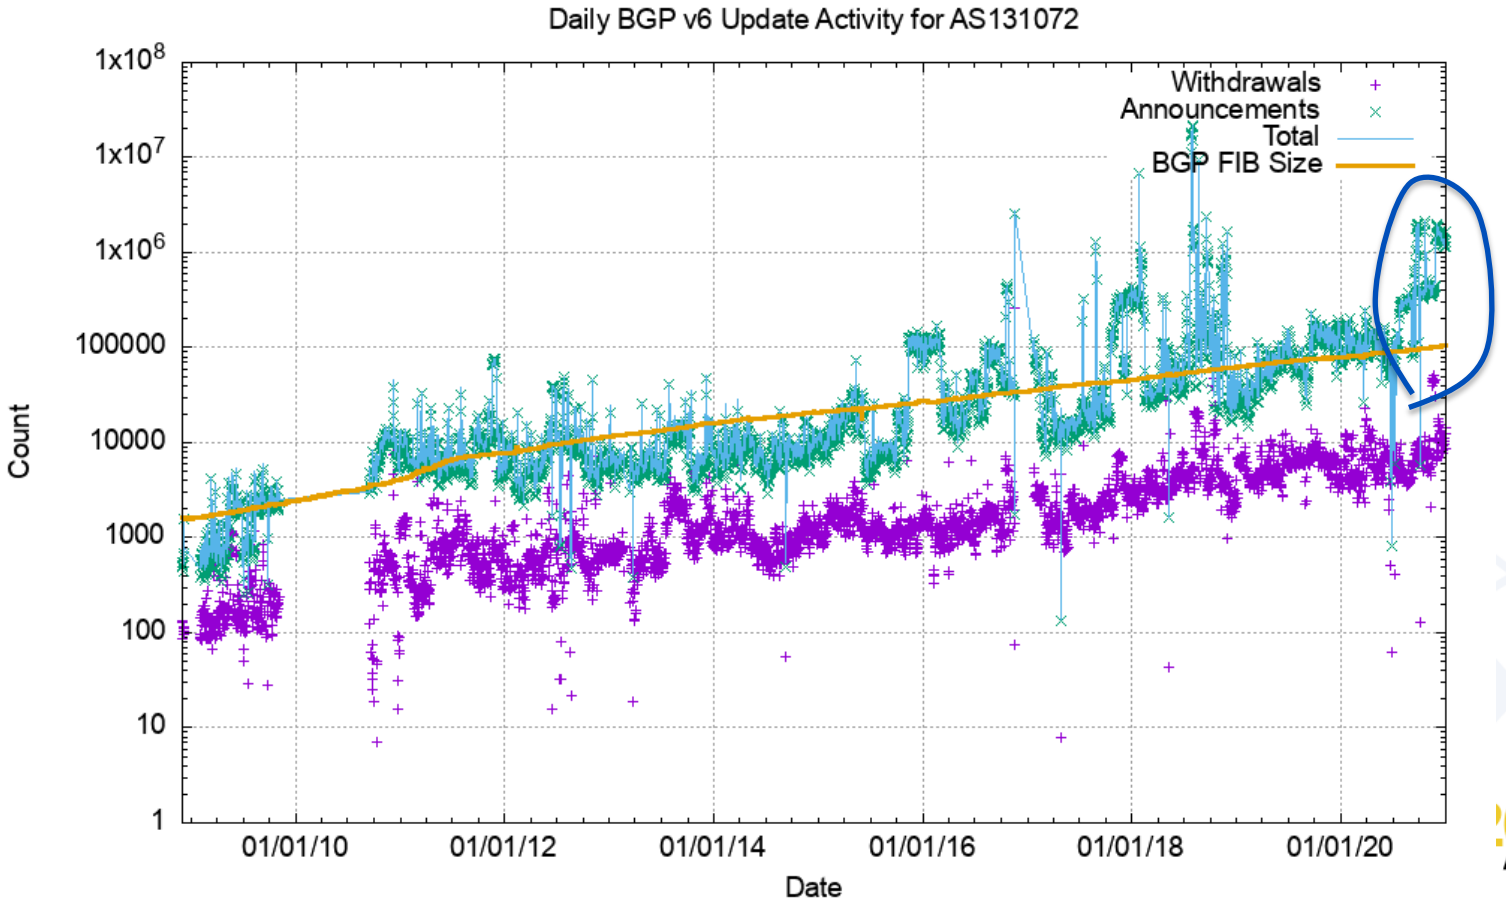

V6 Convergence Performance

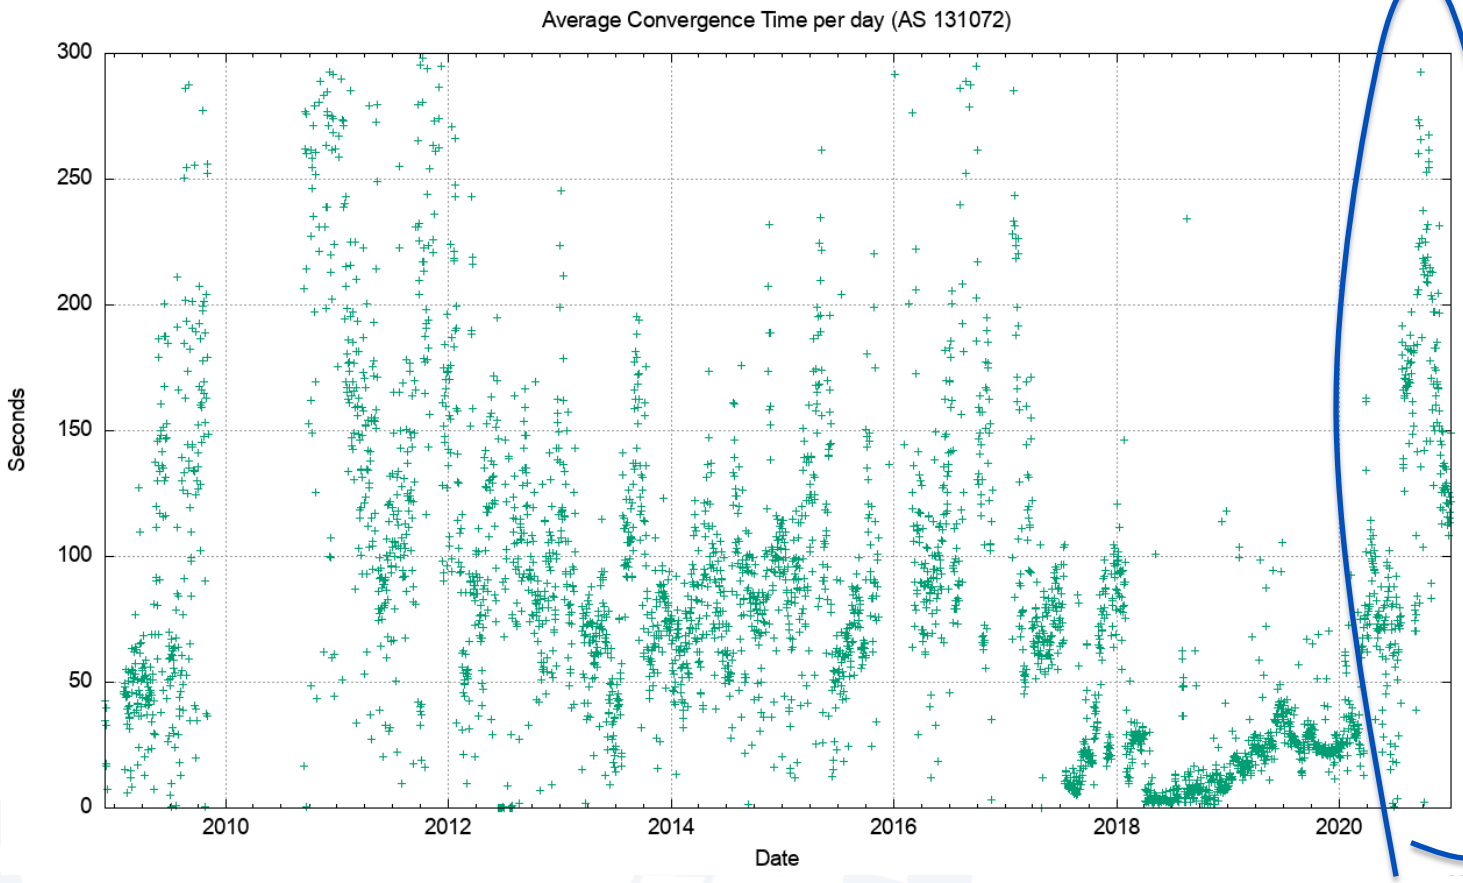

Updates in IPv6 BGP

Nobody is looking…

- Compared to IPv4, the IPv6 network exhibits a high level of routing instability, which is unexpected as the old overlay approaches are disappearing and the topology of IPv6 is now converging to the same topology as IPv4
- Just 5 AS's generate 30% of the BGP update load in the past 2 weeks. Instability is still concentrated in a small number of pathological cases. The instability itself from these sources is getting worse.

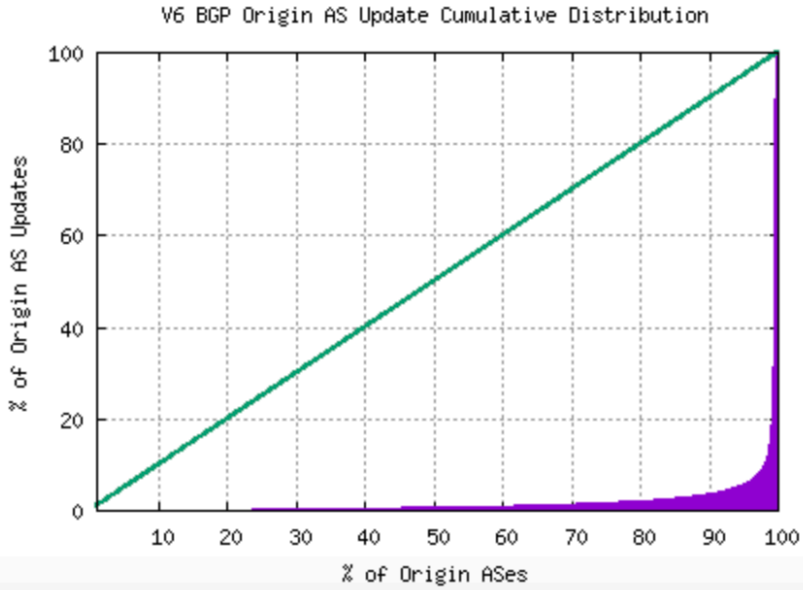

Routing Futures

- There is still 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. Instability leves are rising, generally driven by a small seet of highly unstable “super generators”

- 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-capacity line cards

Some Practical Suggestions

- Understand your hardware's high speed FIB capacity in the default-free parts of your network
- Review your IPv4 / IPv6 portioning - a dual-stack eBGP router will need 970,000 32-bit IPv4 slots and 175,000 128-bit IPv6 slots for a full eBGP routing table in line cards over the coming 24 months if they are using a full eBGP FIB load (plus internal routes of course)
- Judicious use of default routes in your internal network may allow you drop this requirement significantly
- Using a hot cache for line card FIB cache would reduce the memory requirement significantly without visible performance cost

That's it!


2021
APRICOT
APNIC 51
ONLINE
22 February – 4 March 2021
#apricot2021