BGP in 2022

Geoff Huston
APNIC Labs

The Highlights

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

29 Years of Routing the Internet

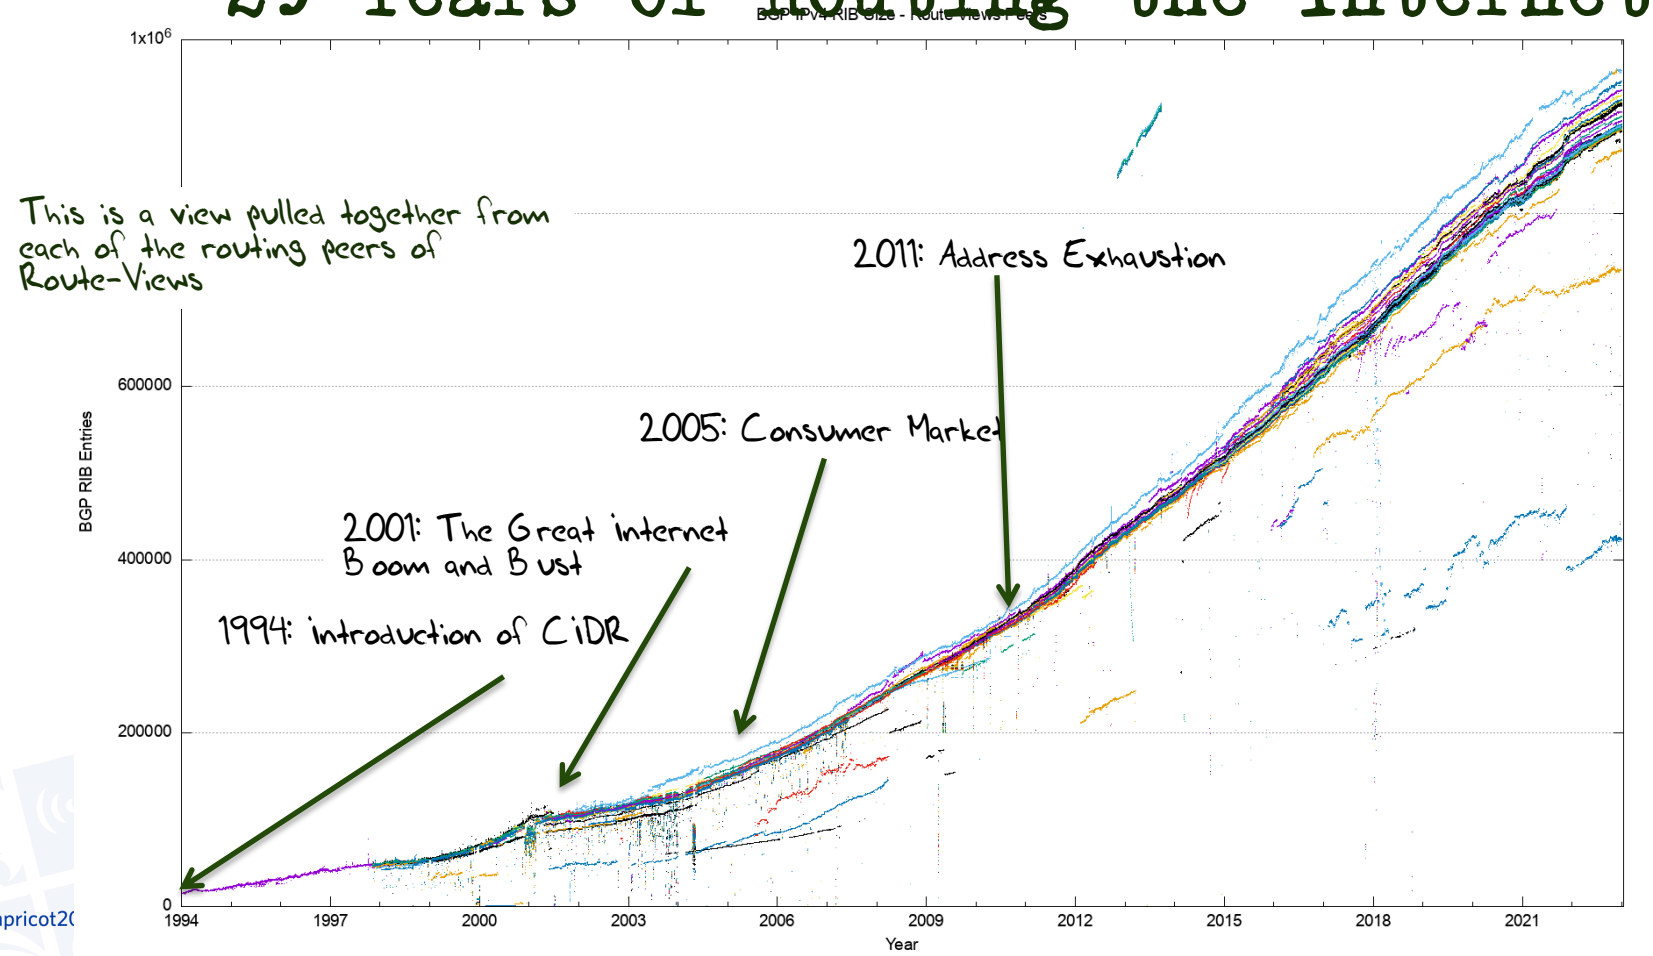

2022 in detail

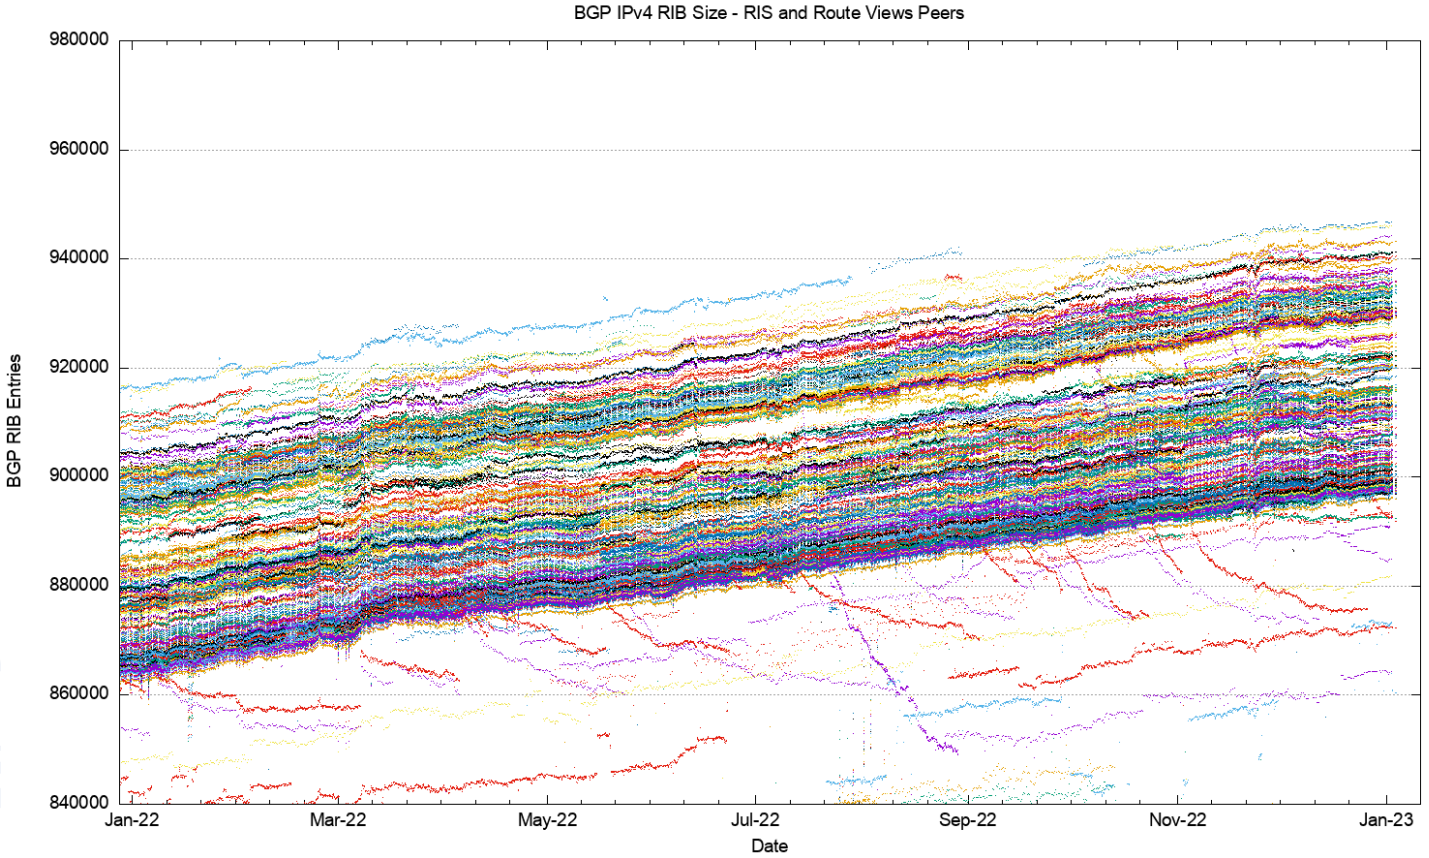

2022 in detail

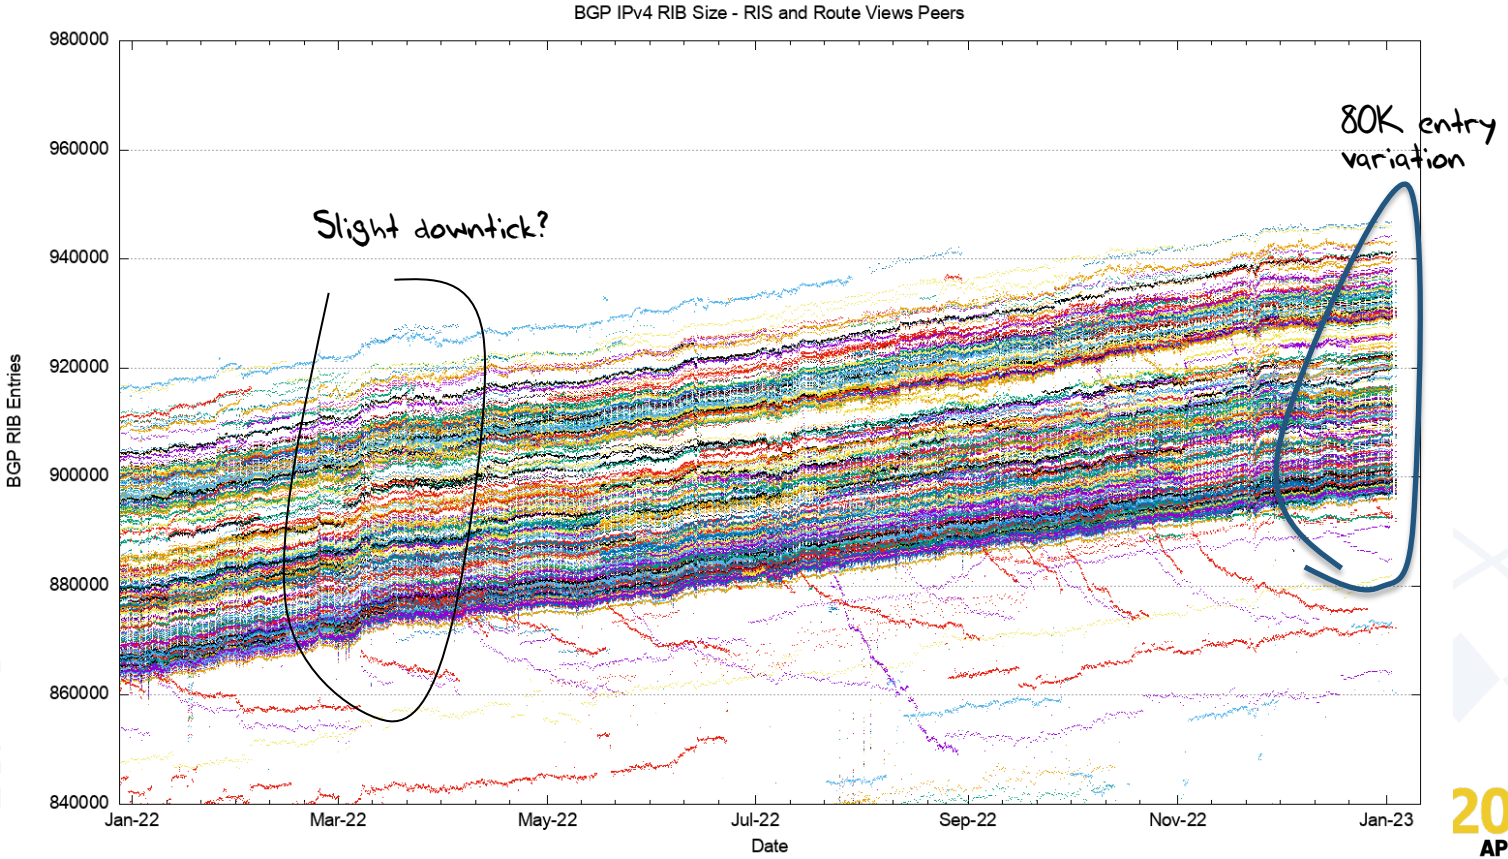

2022: Assigned vs Recovered

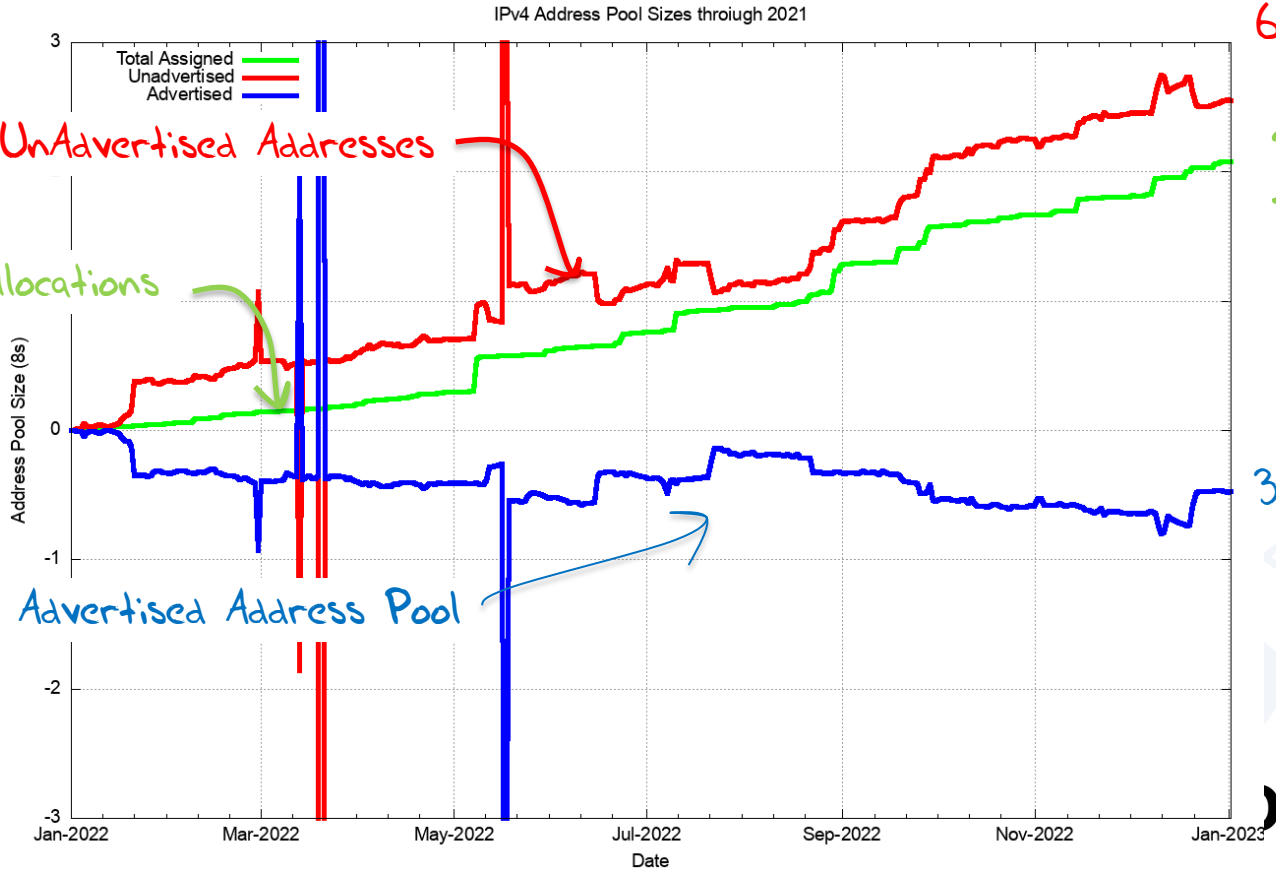

Change in UnAdvertised Addresses

RiR Allocations

Change in the Advertised Address Pool

62M unadvertised
+9.5M delta

3.7B allocated
+1.5M delta

306M advertised
-8M delta

2022: Assigned vs Recovered

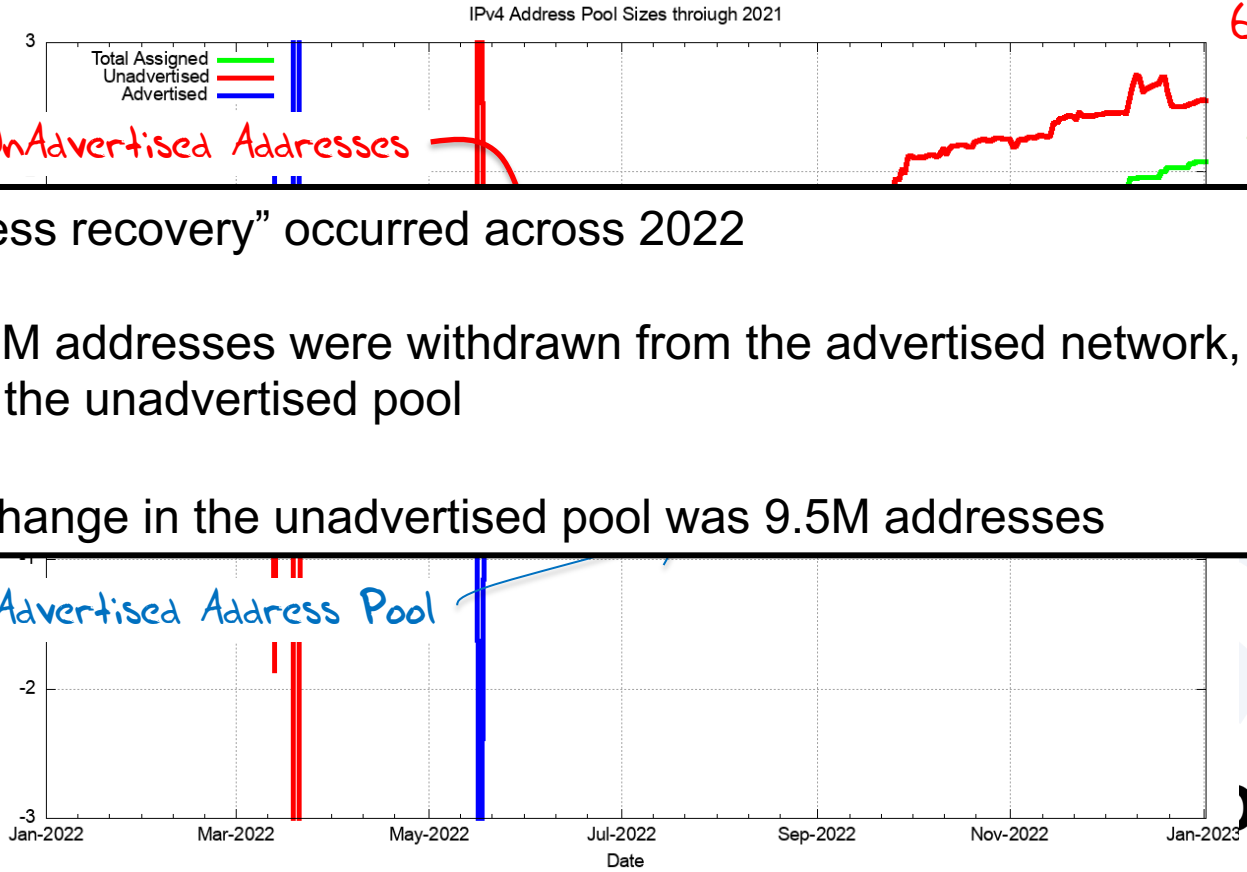

Change in UnAdvertised Addresses

62M unadvertised
+9.5M delta

3.7B allocated

No “address recovery” occurred across 2022

A net of 8M addresses were withdrawn from the advertised network, and shifted to the unadvertised pool

The net change in the unadvertised pool was 9.5M addresses

Change in the Advertised Address Pool

What happened in 2022 in V4?

- From the look of the routing growth plots, the growth of the size of the IPv4 network **is slowing down**
- The number of entries in the IPv4 default-free zone reached 944,000 by the end of 2022
- The pace of growth of the routing table was slightly lower than the rolling 5-year average, with **36,000 new entries in 2022** (was 40,000 in 2021)
- The AS position was slightly lower with **1,400 new AS's advertised in 2022** (was 2,400 in 2021)
- Transit relationships have not changed materially over 2022 for most networks
- **The overall IPv4 routing growth trends slowed down through 2022**

The Highlights

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

The Route-Views View of IPv6

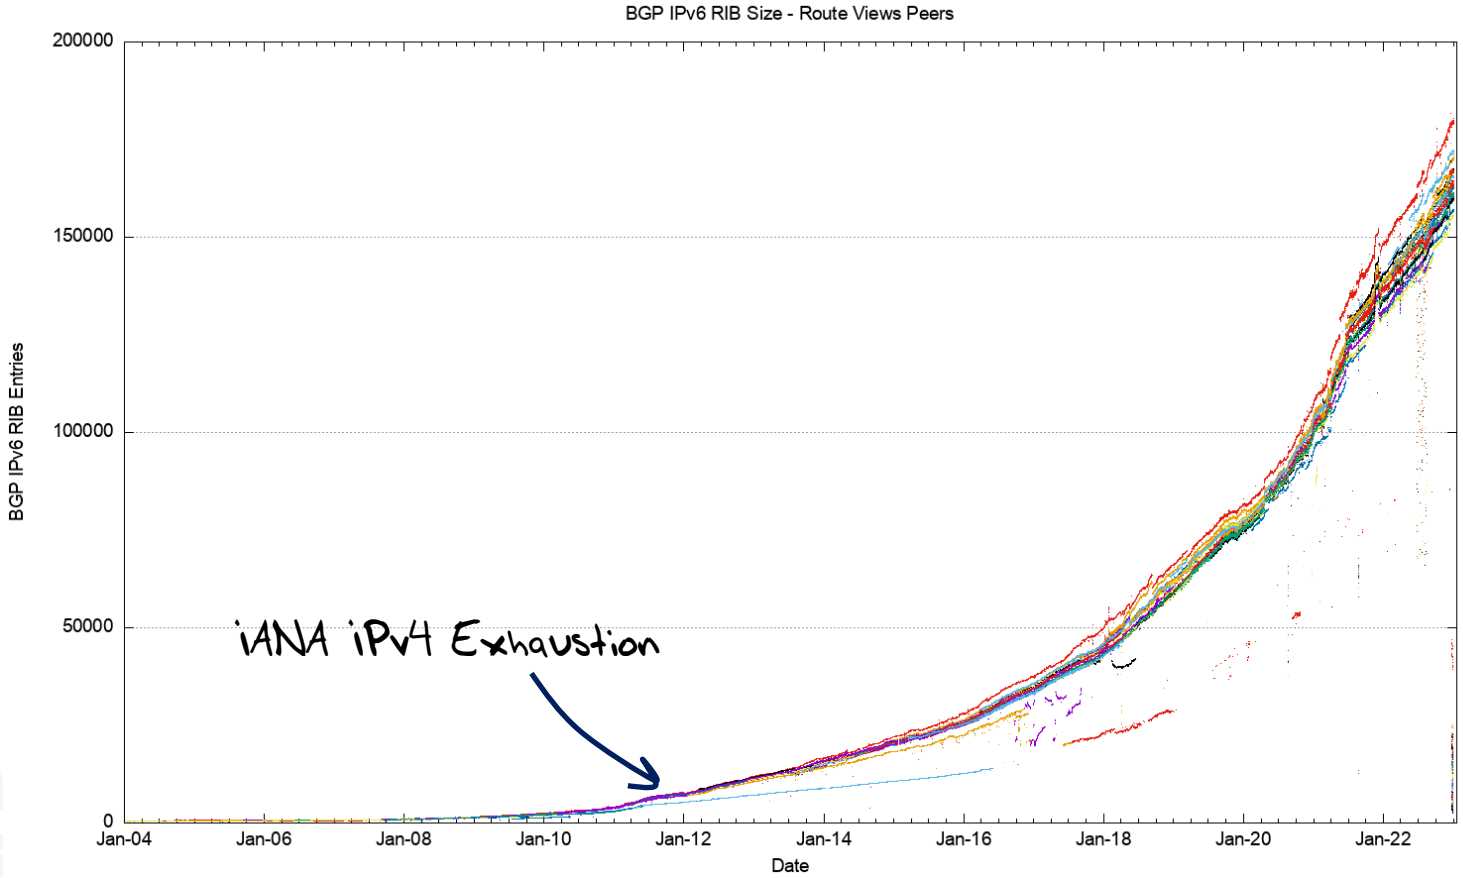

2022 in Detail

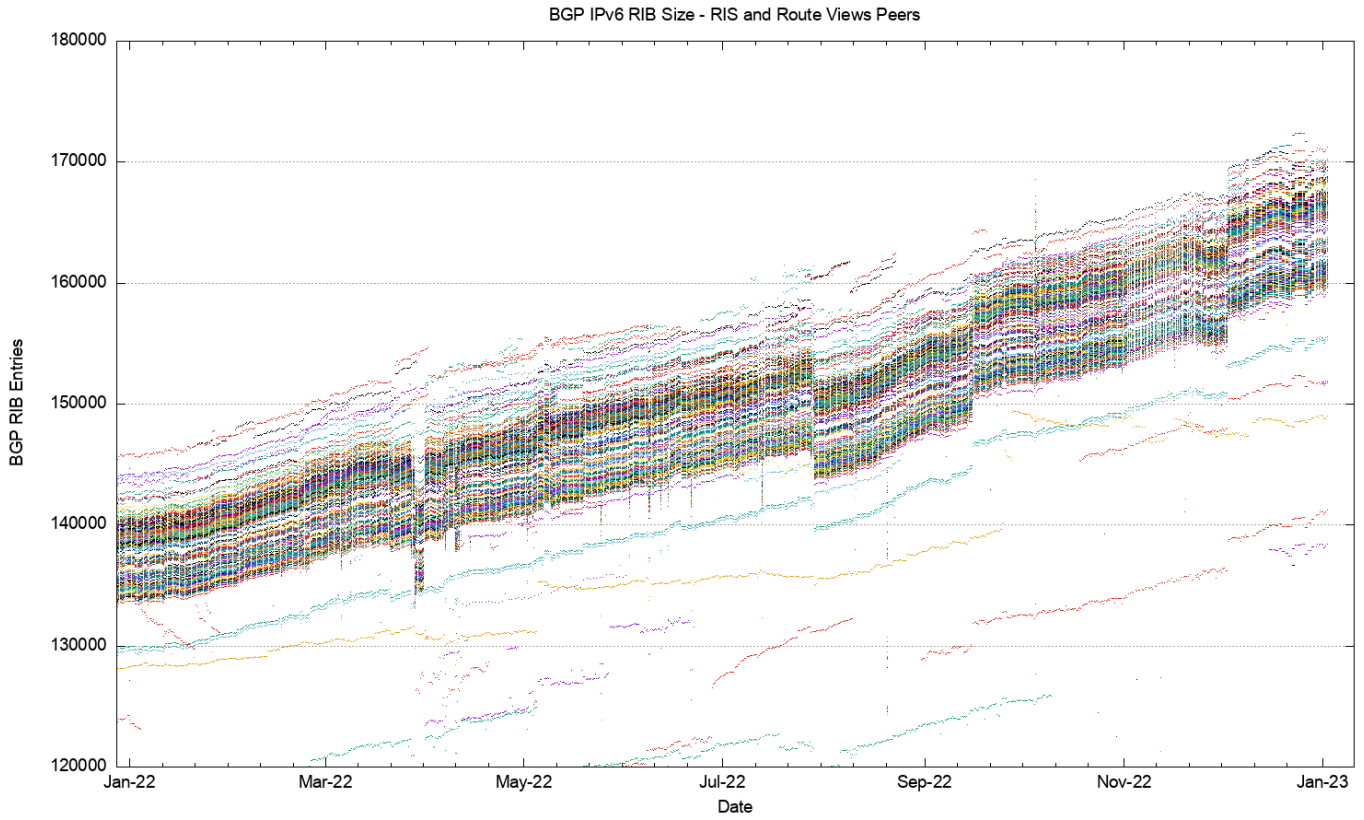

V6 in 2022

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **35,000 route entries p.a.**

  - More use of /48 more specifics
  - More networks advertising IPv6 prefixes

More Specifics in IPv6

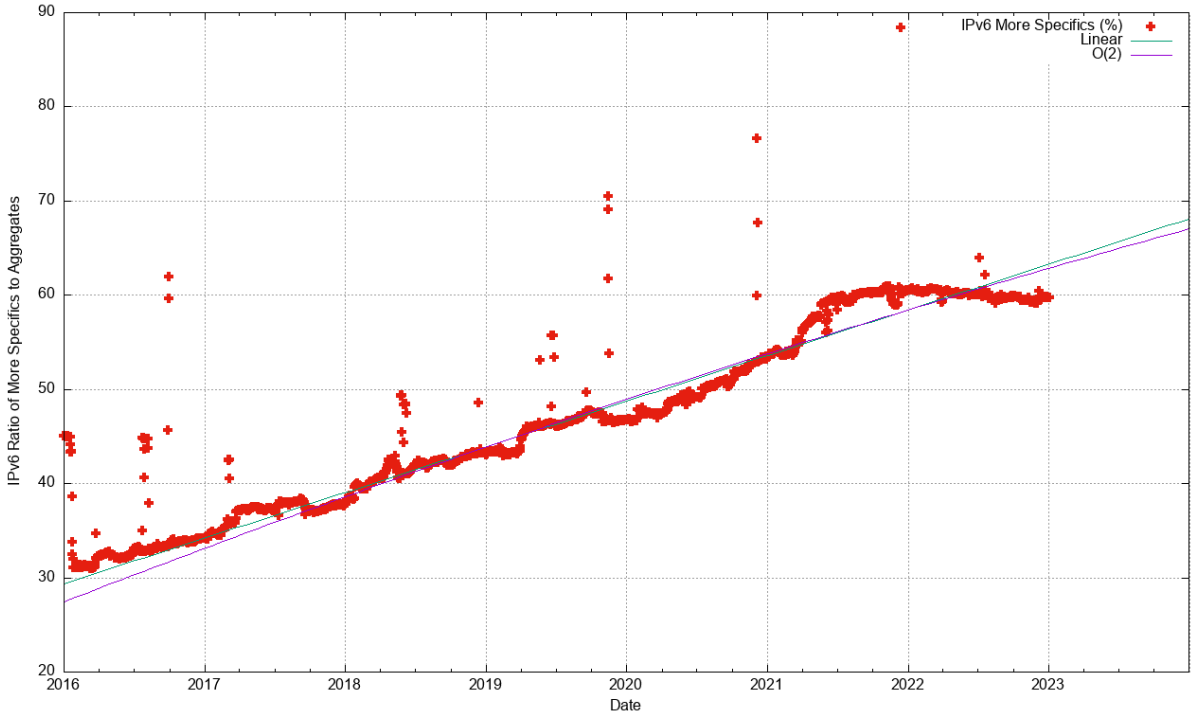

More Specifics in IPv6

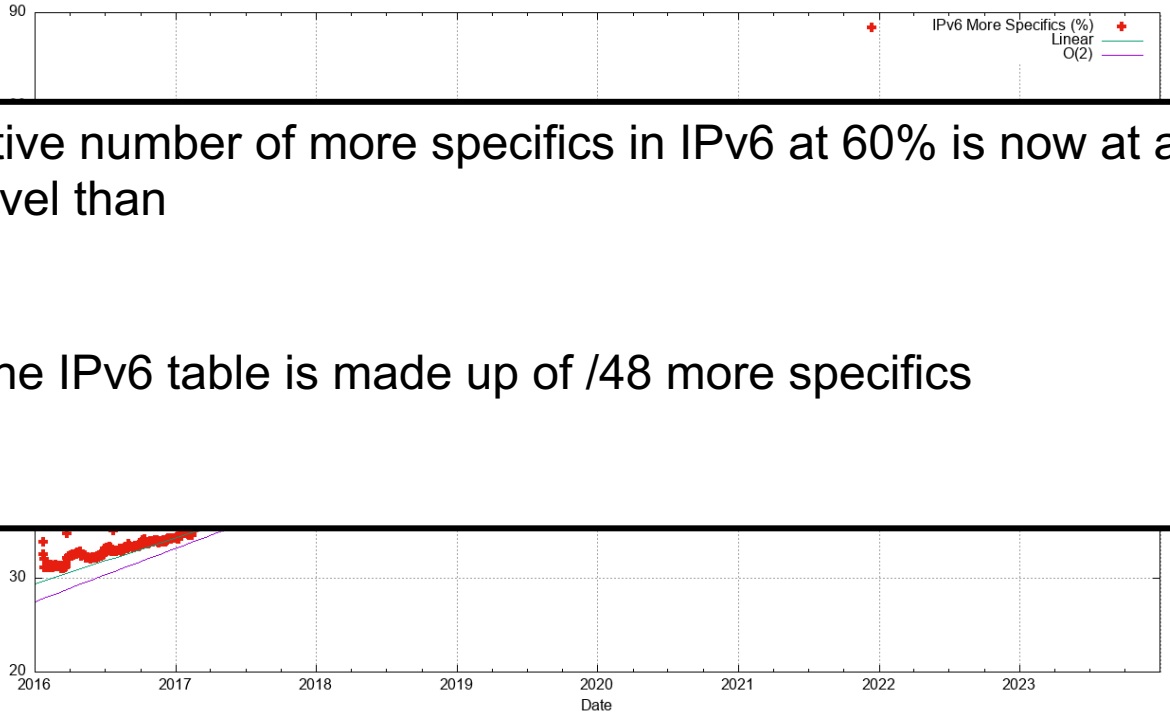

The relative number of more specifics in IPv6 at 60% is now at a higher level than
IPv4

50% of the IPv6 table is made up of /48 more specifics

The Highlights

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

V4 BGP Table Size Predictions

Date	RIB Size	Prediction
Jan *2018*	*699,000*	
2019	*760,000*	
2020	*814,000*	
2021	*866,000*	
2022	*906,000*	
2023	*942,000*	*944,000*
2024		*977,000*
2025		*1,005,000*
2026		*1,028,000*
2027		*1,045,000*
2028		*1,057,000*

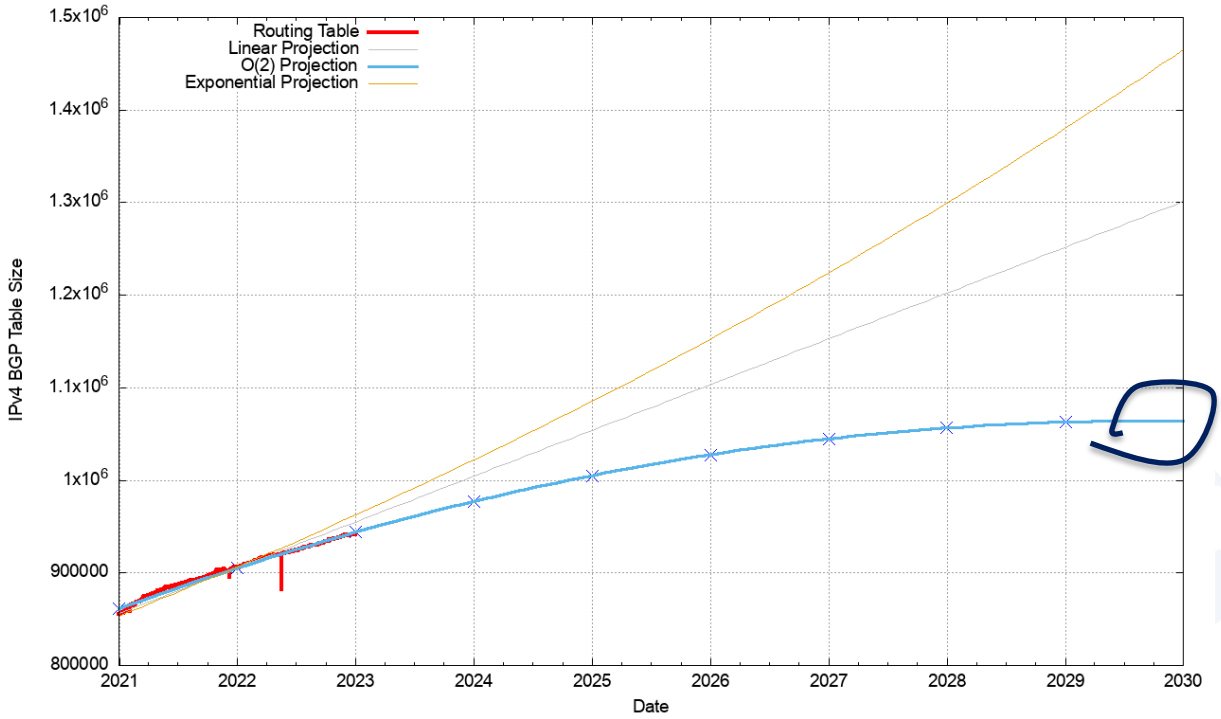

V4 BGP Table Size Predictions

Date	RIB Size	Prediction
Jan 2018	699,000	
2019	760,000	
202		
202		
202		
202		
202		
202		
202		
2027		1,045,000
2028		1,057,000

The best fit projection model is now less than linear

Here I'm using a polynomial function with a decreasing growth rate as a bext fit to the data

This model predicts the IPv4 FIB will peak at some 1.1M entries in 2030 (or thereabouts)

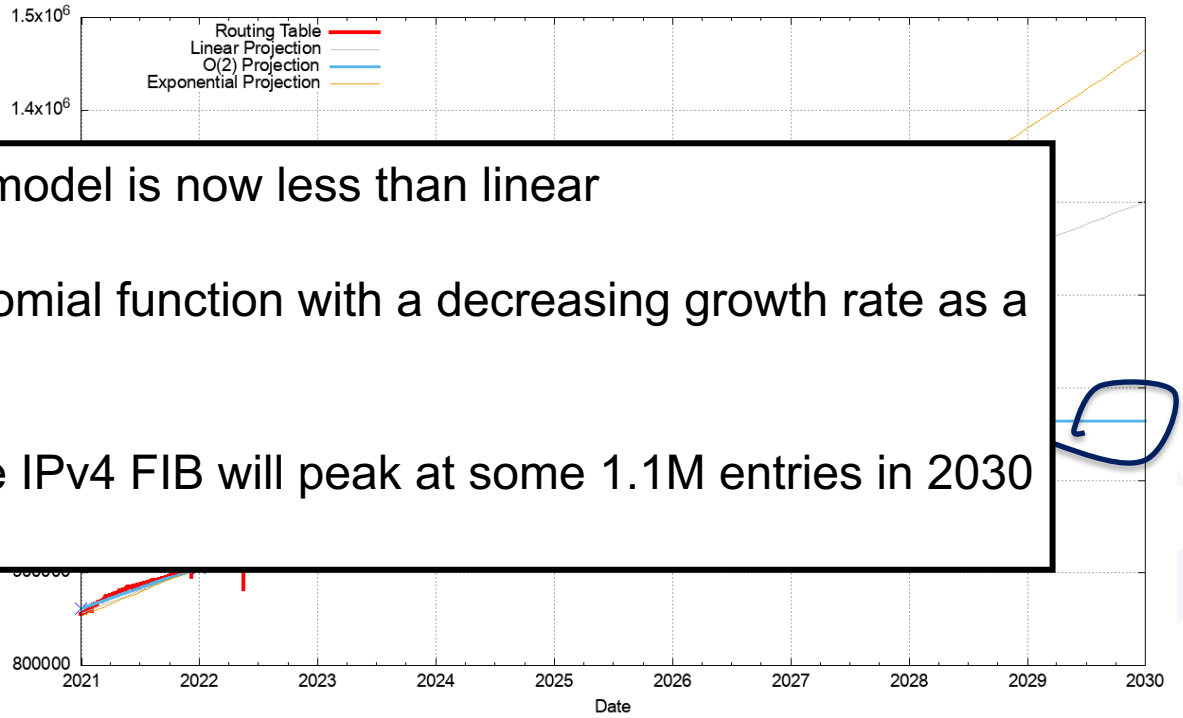

V6 BGP Table Size Predictions

	Linear	Exponential
Jan 2018	45,000	
2019	62,000	
2020	79,000	
2021	104,000	
2022	147,000	
2023	172,000	185,000
2024	199,000	243,000
2025	226,000	320,000
2026	252,000	412,000
2027	279,000	554,000
2028	305,000	723,000

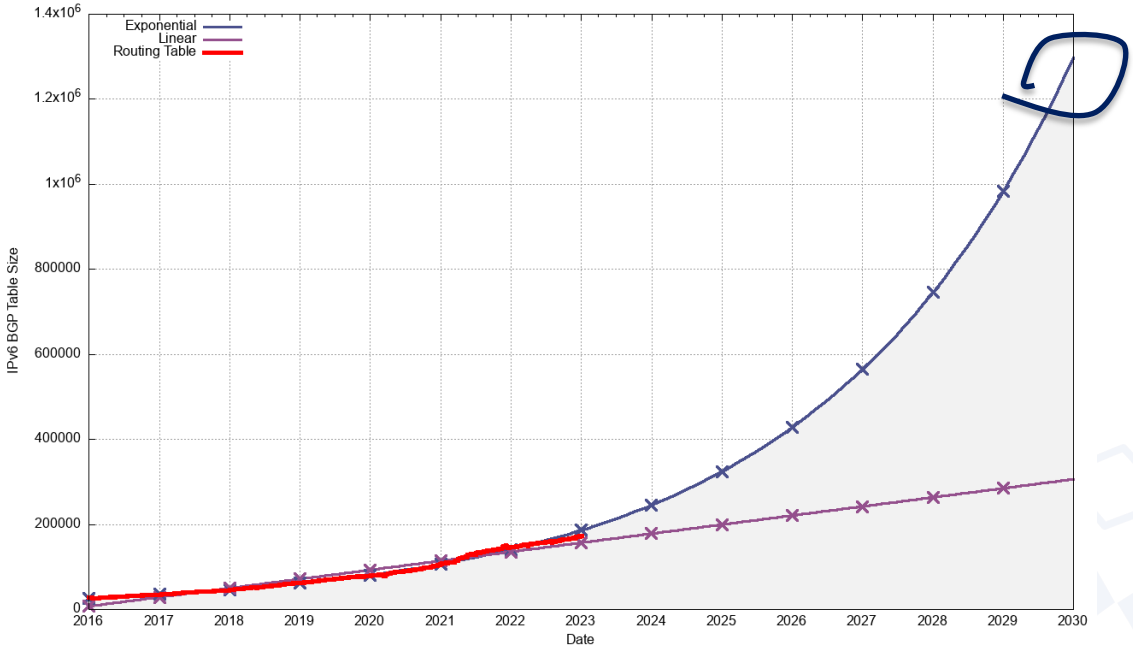

Note that the IPv6 tables are 128bits wide – i.e. 4x the size of the IPv4 tables!

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2018*	*45,000*	
2019	*62,000*	
2020	*79,000*	
2021		
2022		
2023		
2024		
2025		
2026	*252,000*	*412,000*
2027	*279,000*	*554,000*
2028	*305,000*	*723,000*

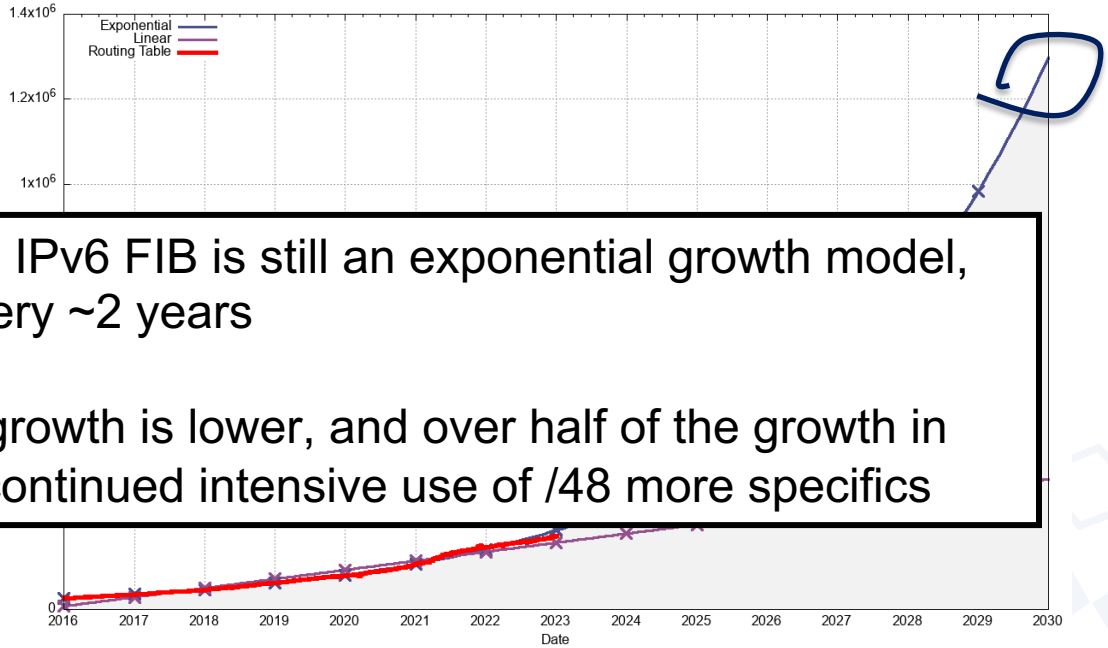

The best fit model for the IPv6 FIB is still an exponential growth model, with a doubling factor every ~2 years

The underlying network growth is lower, and over half of the growth in this model is due to the continued intensive use of /48 more specifics

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

These two tables will require the same storage/lookup size in around 1 year from now, given that each IPv6 entry is 4 times the bit size of an IPv4 entry

The good news …

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 - Daily Updates

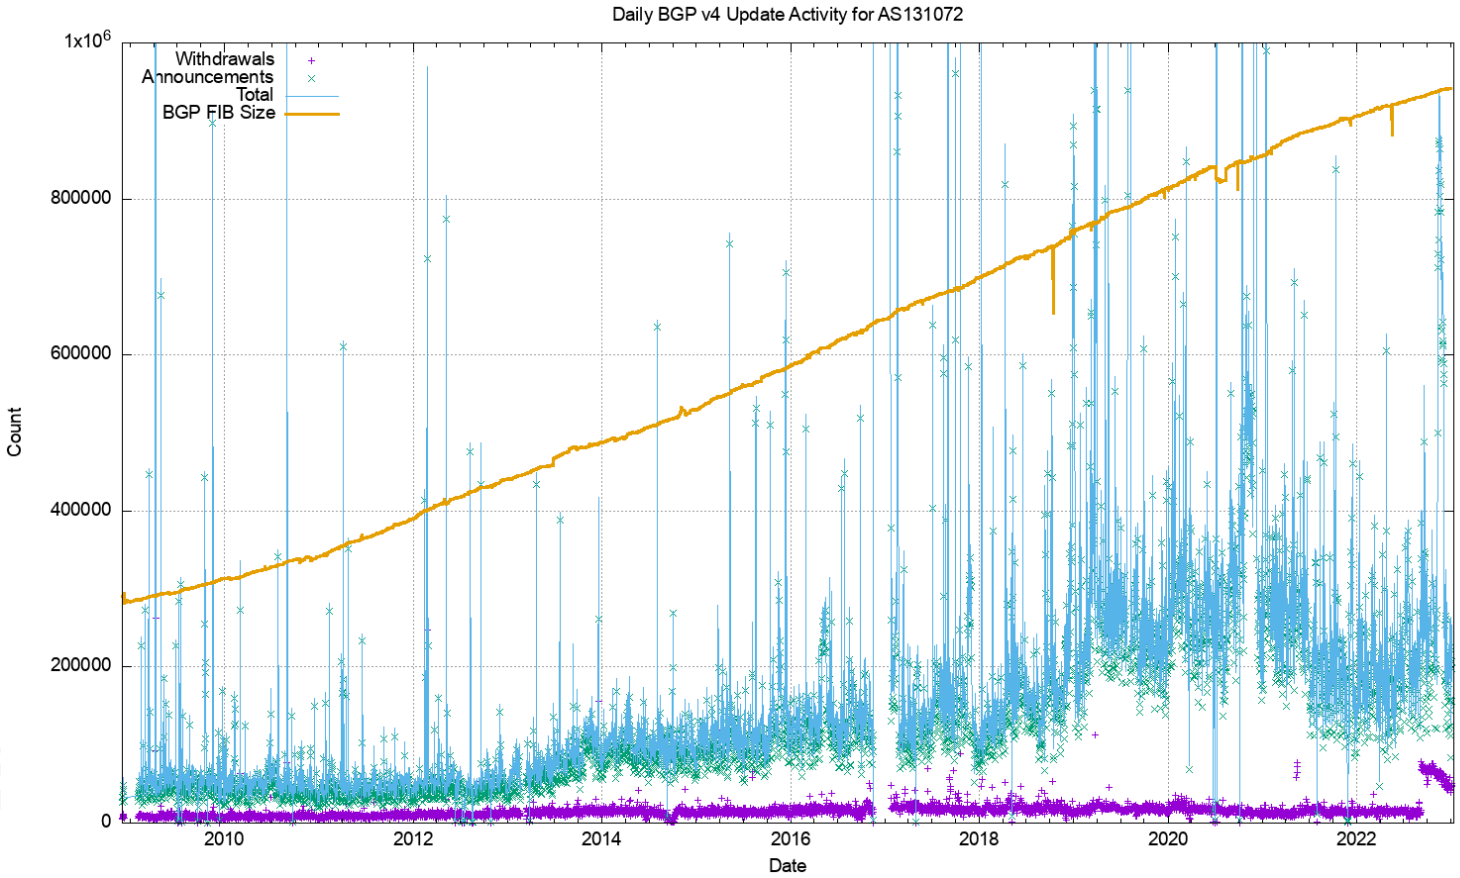

IPv4 BGP Updates - Daily Updates

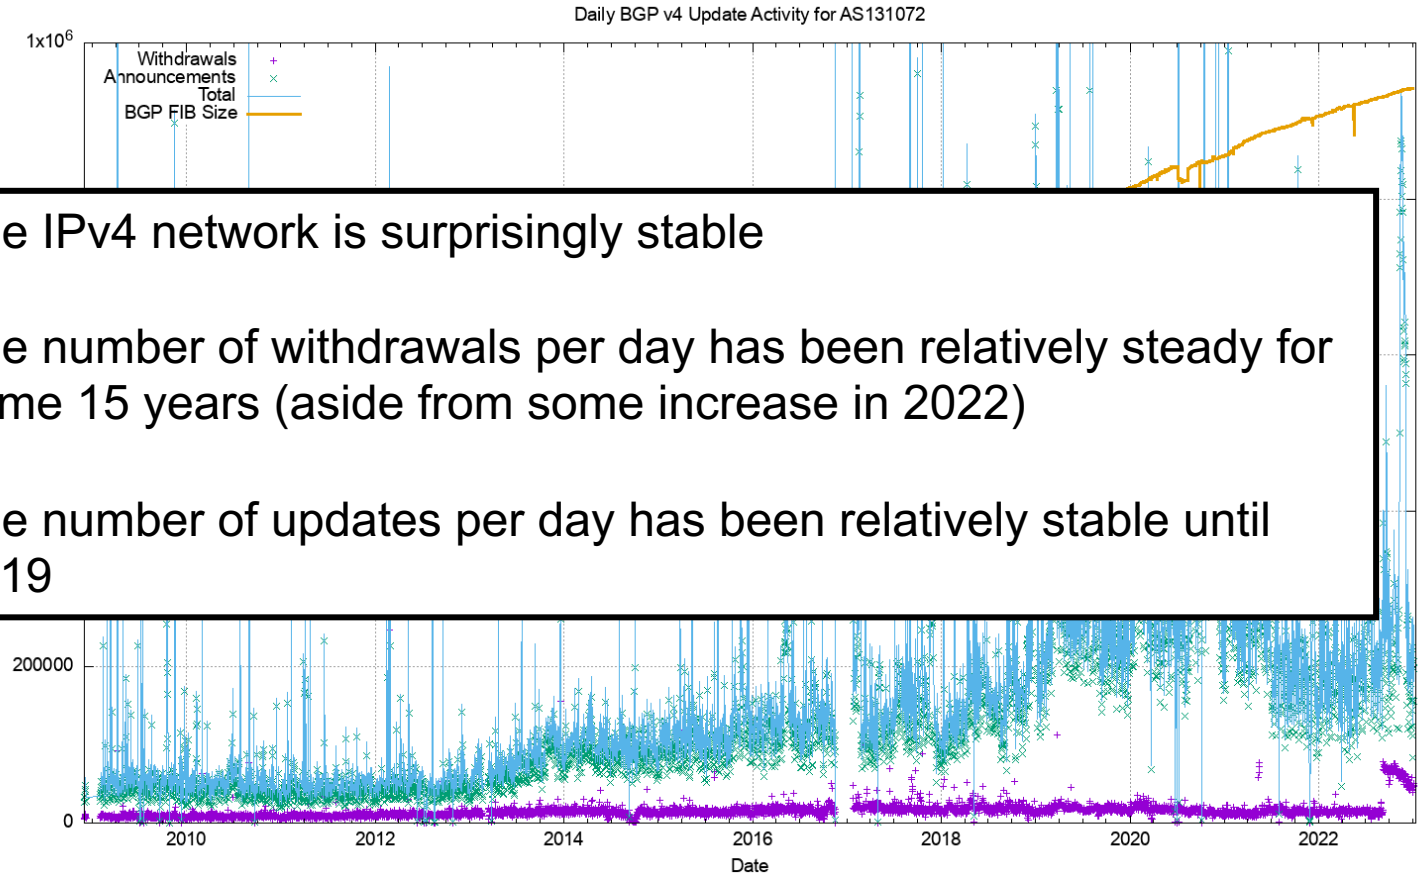

The IPv4 network is surprisingly stable

The number of withdrawals per day has been relatively steady for some 15 years (aside from some increase in 2022)

The number of updates per day has been relatively stable until 2019

IPv4 Unstable Prefixes per Day

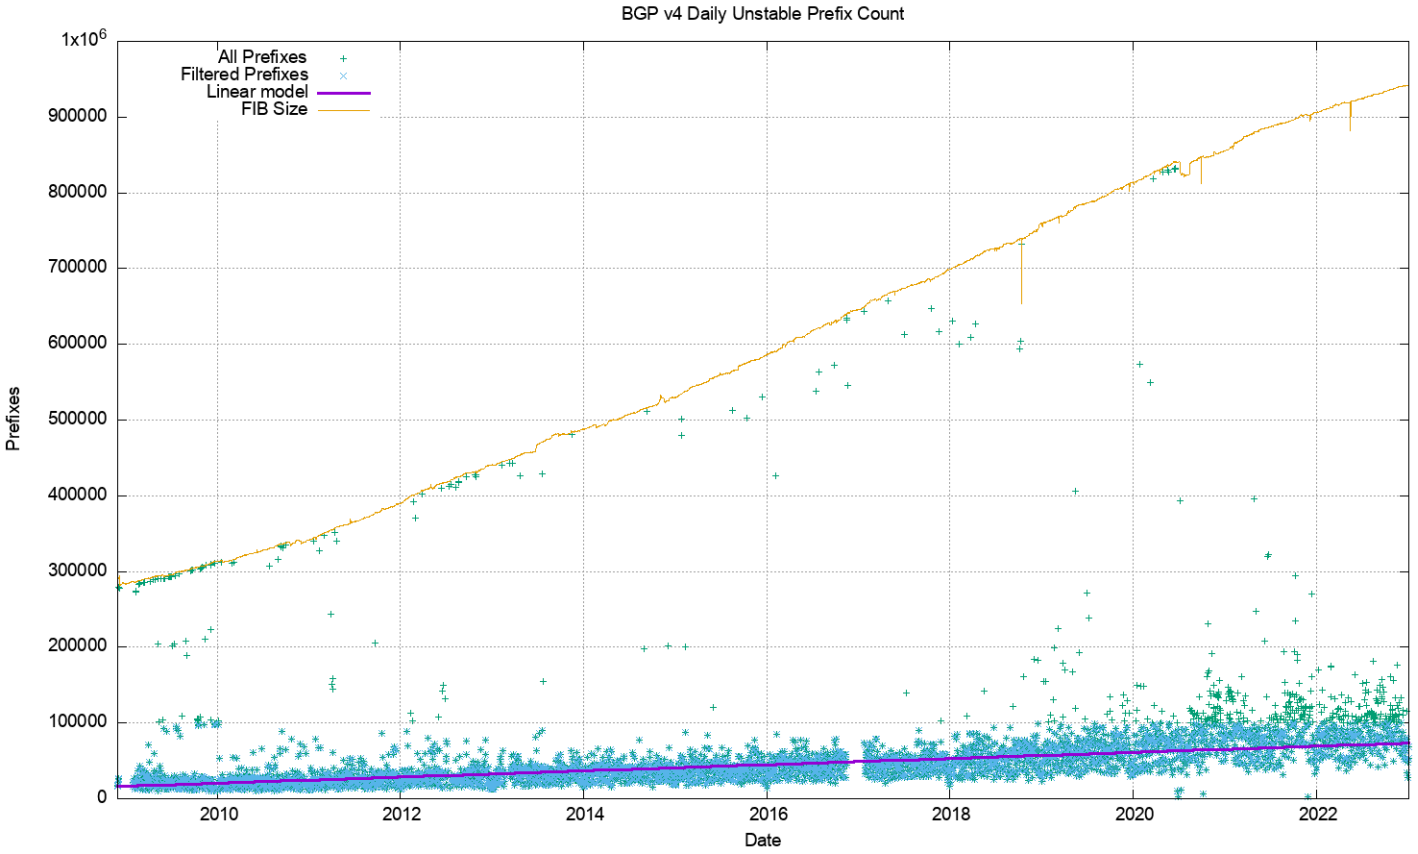

IPv4 Unstable Prefixes per Day

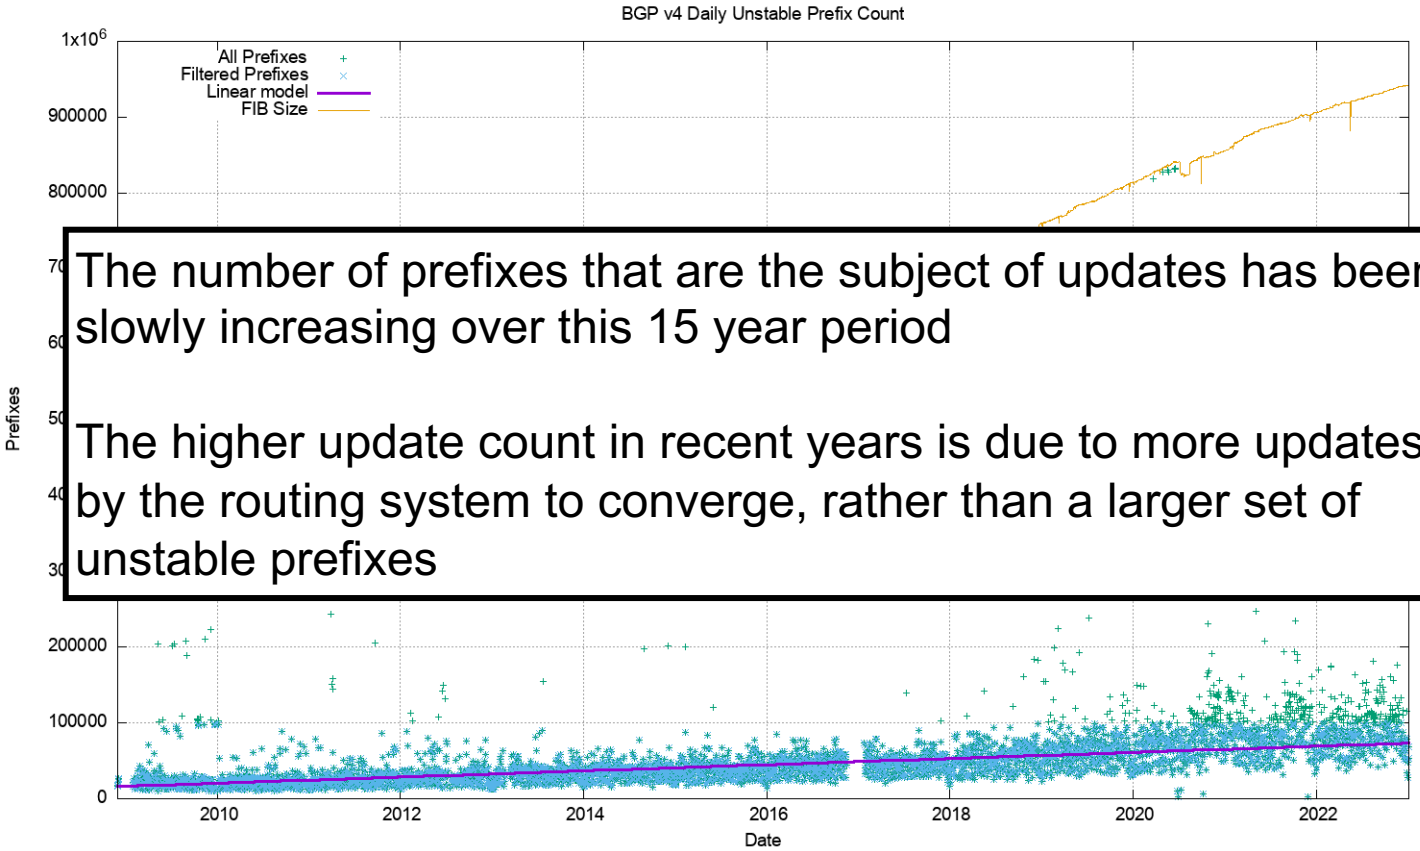

The number of prefixes that are the subject of updates has been slowly increasing over this 15 year period

The higher update count in recent years is due to more updates used by the routing system to converge, rather than a larger set of unstable prefixes

IPv4 BGP Convergence Performance

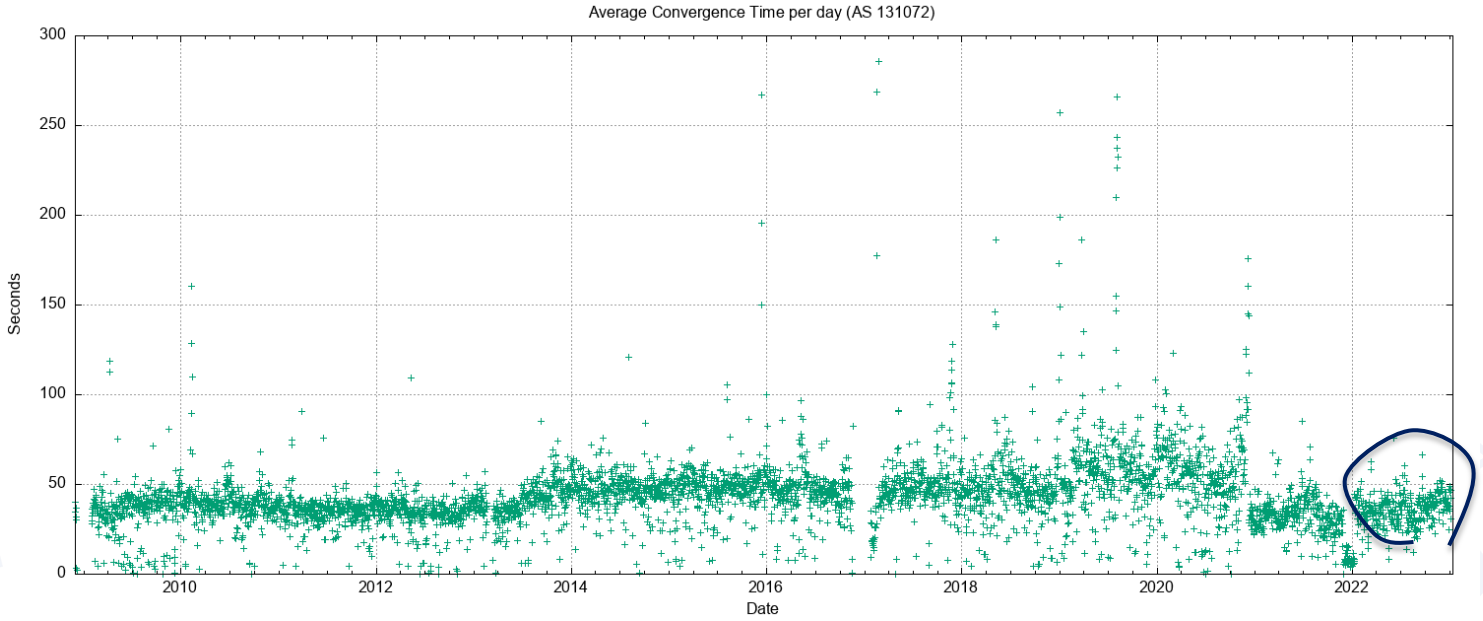

IPv4 BGP Convergence Performance

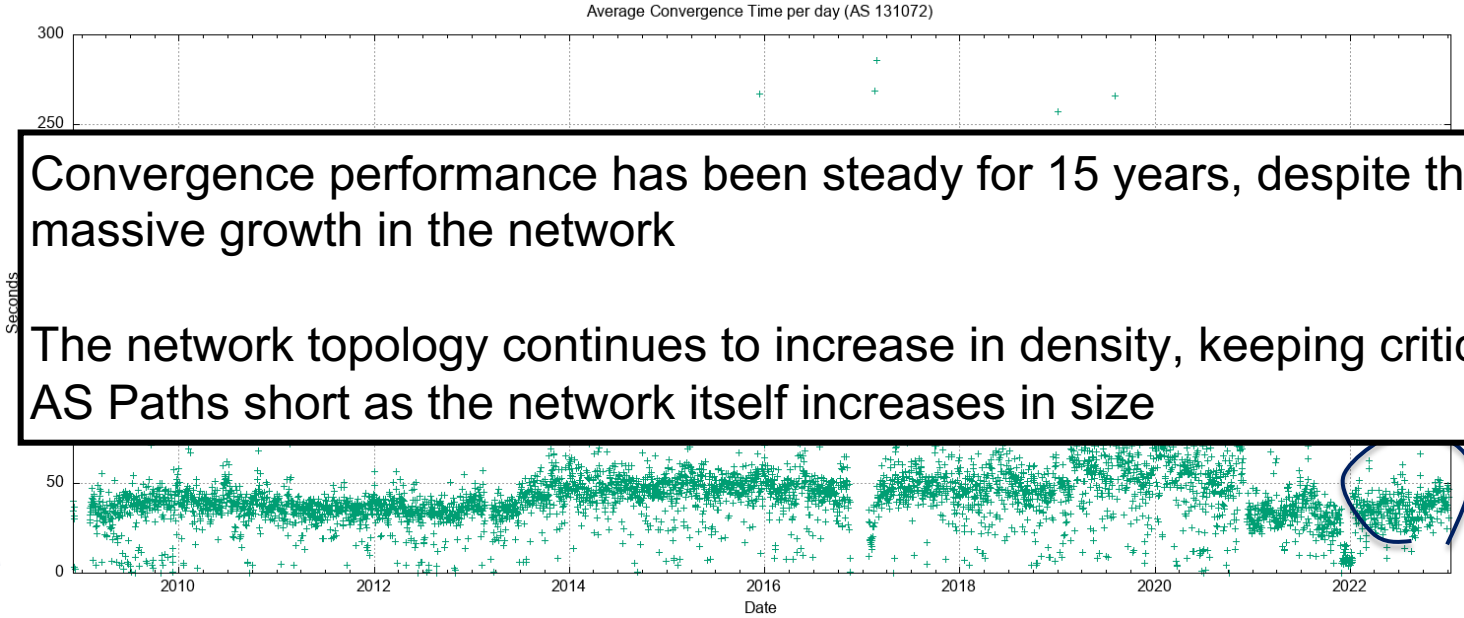

Convergence performance has been steady for 15 years, despite the massive growth in the network

The network topology continues to increase in density, keeping critical AS Paths short as the network itself increases in size

Updates in IPv4 BGP

The IPv4 inter-domain routing system is still highly stable …

- The number of updates per instability event and the time to converge to a stable forwarding state has been relatively constant for many years - it rose in 2019 - 2020 and has declined again in 2021, and stabilized in 2022
- 20% of prefixes generate 80% of all updates. Less than 5% of all origin networks are linked to 80% of all updates. **Instability is concentrated in a small number of highly unstable cases.**

V6 BGP Updates

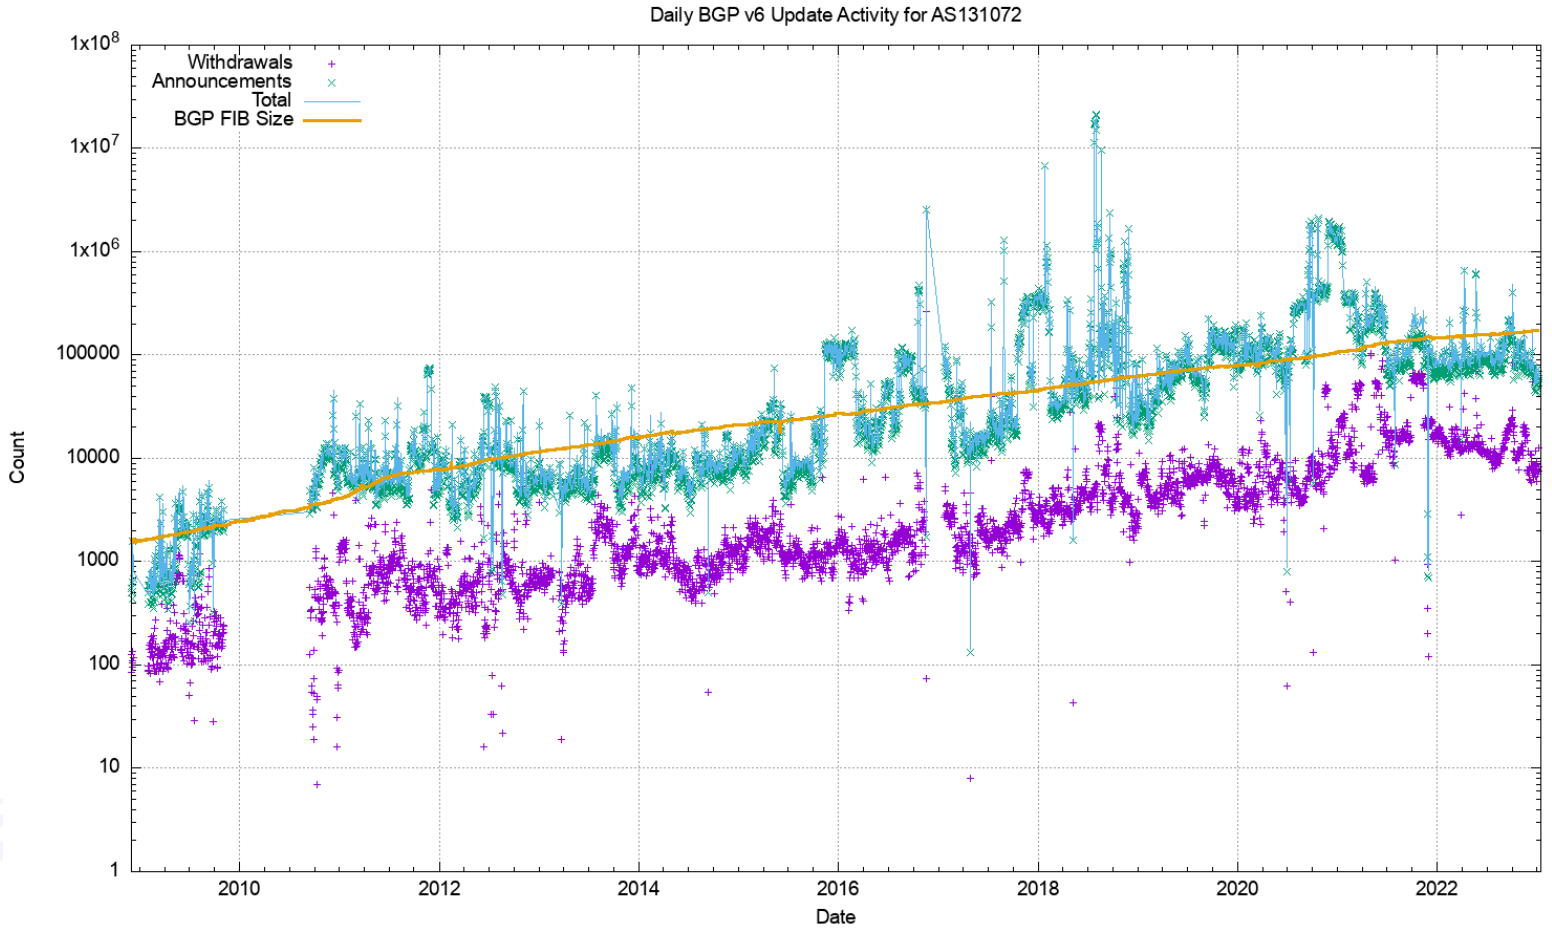

V6 BGP Updates

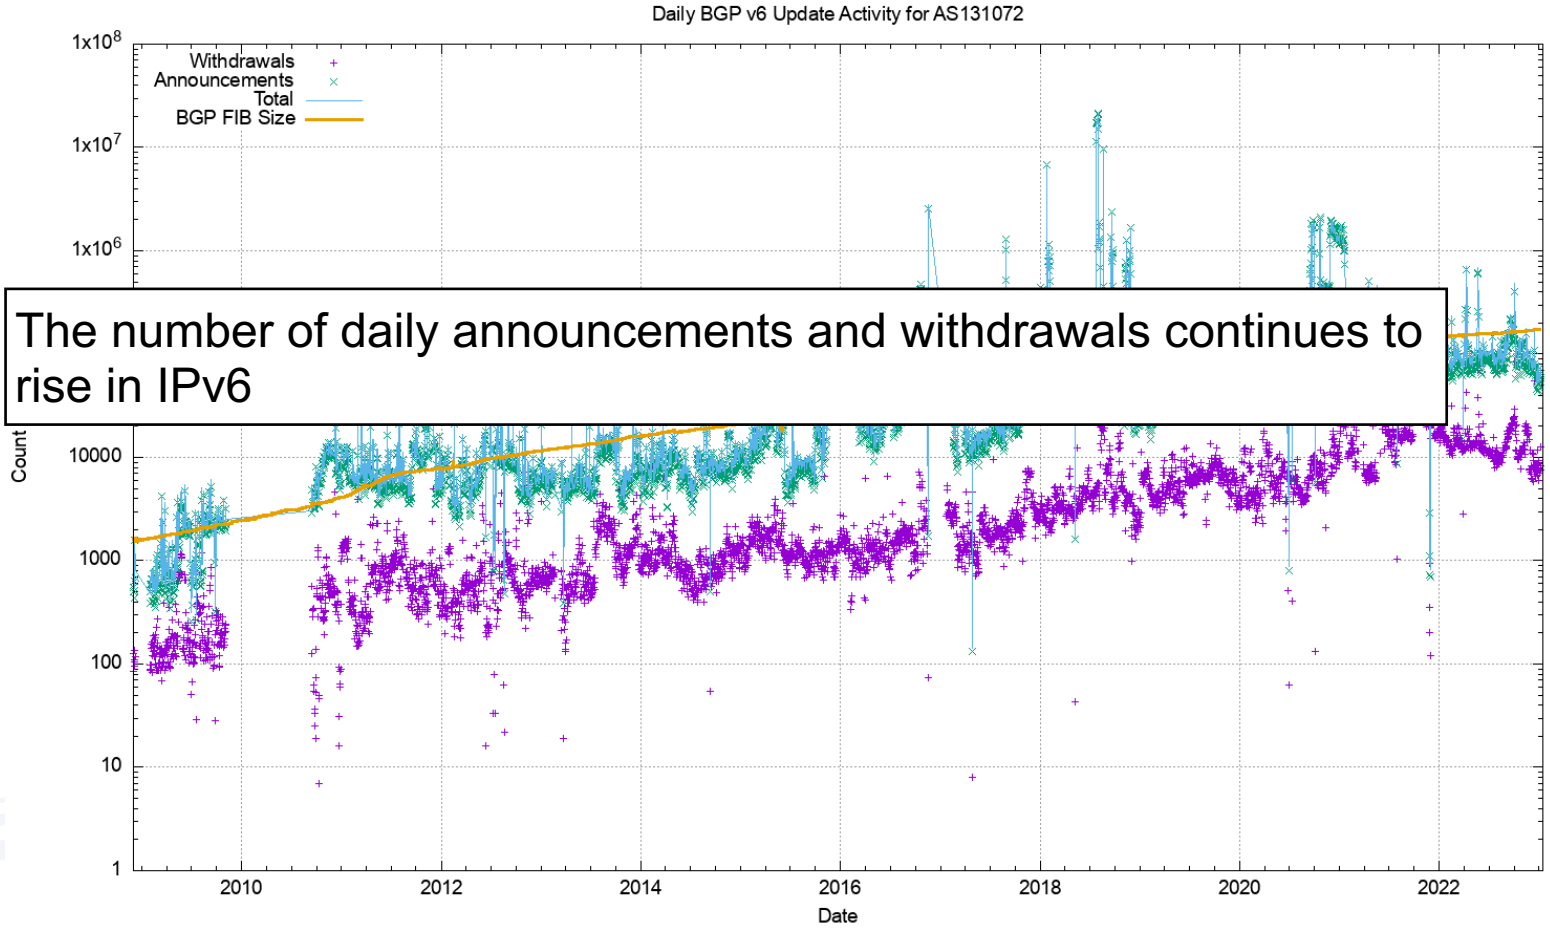

The number of daily announcements and withdrawals continues to rise in IPv6

V6 Unstable Prefixes

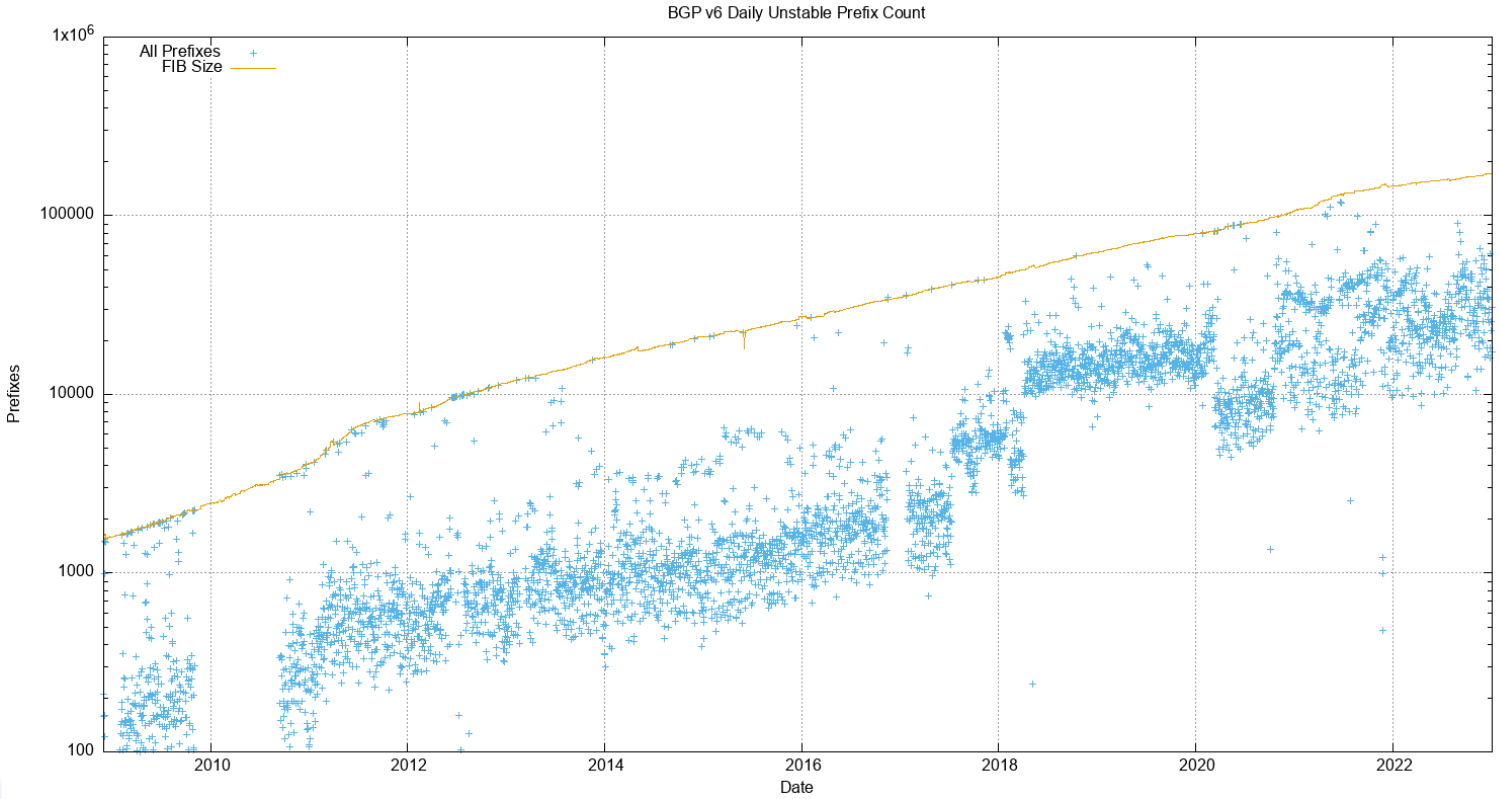

V6 Unstable Prefixes

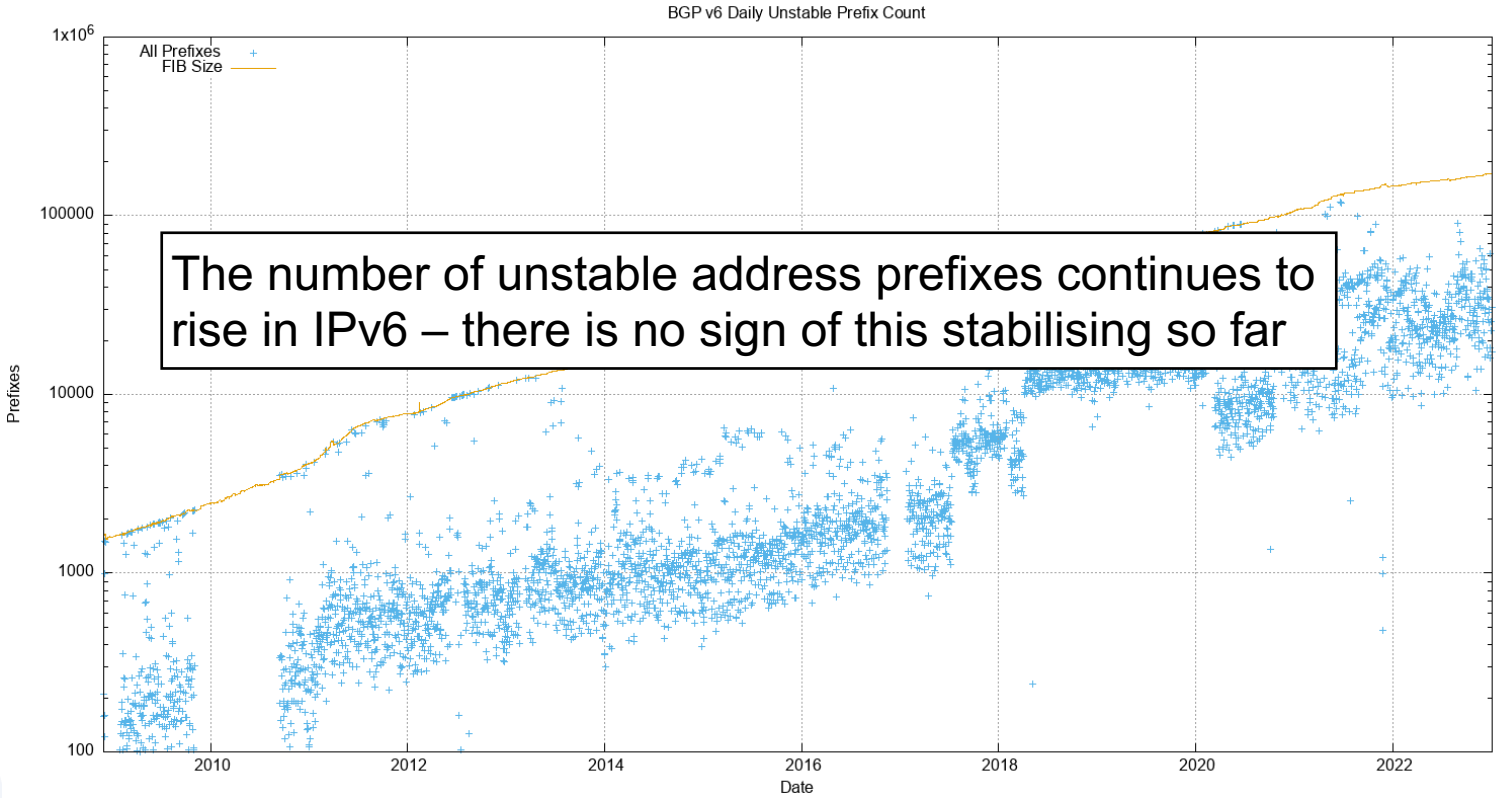

The number of unstable address prefixes continues to rise in IPv6 – there is no sign of this stabilising so far

V6 Convergence Performance

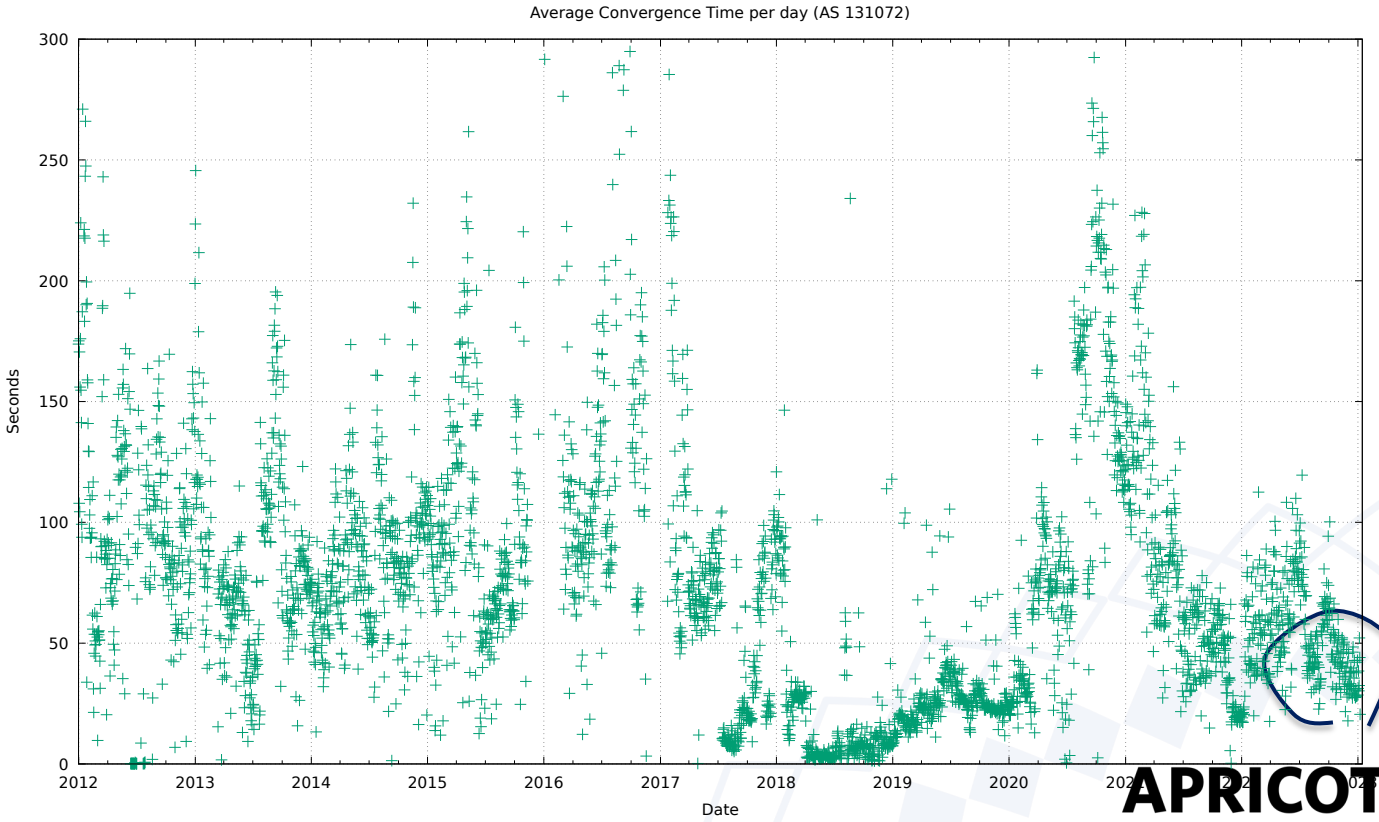

V6 Convergence Performance

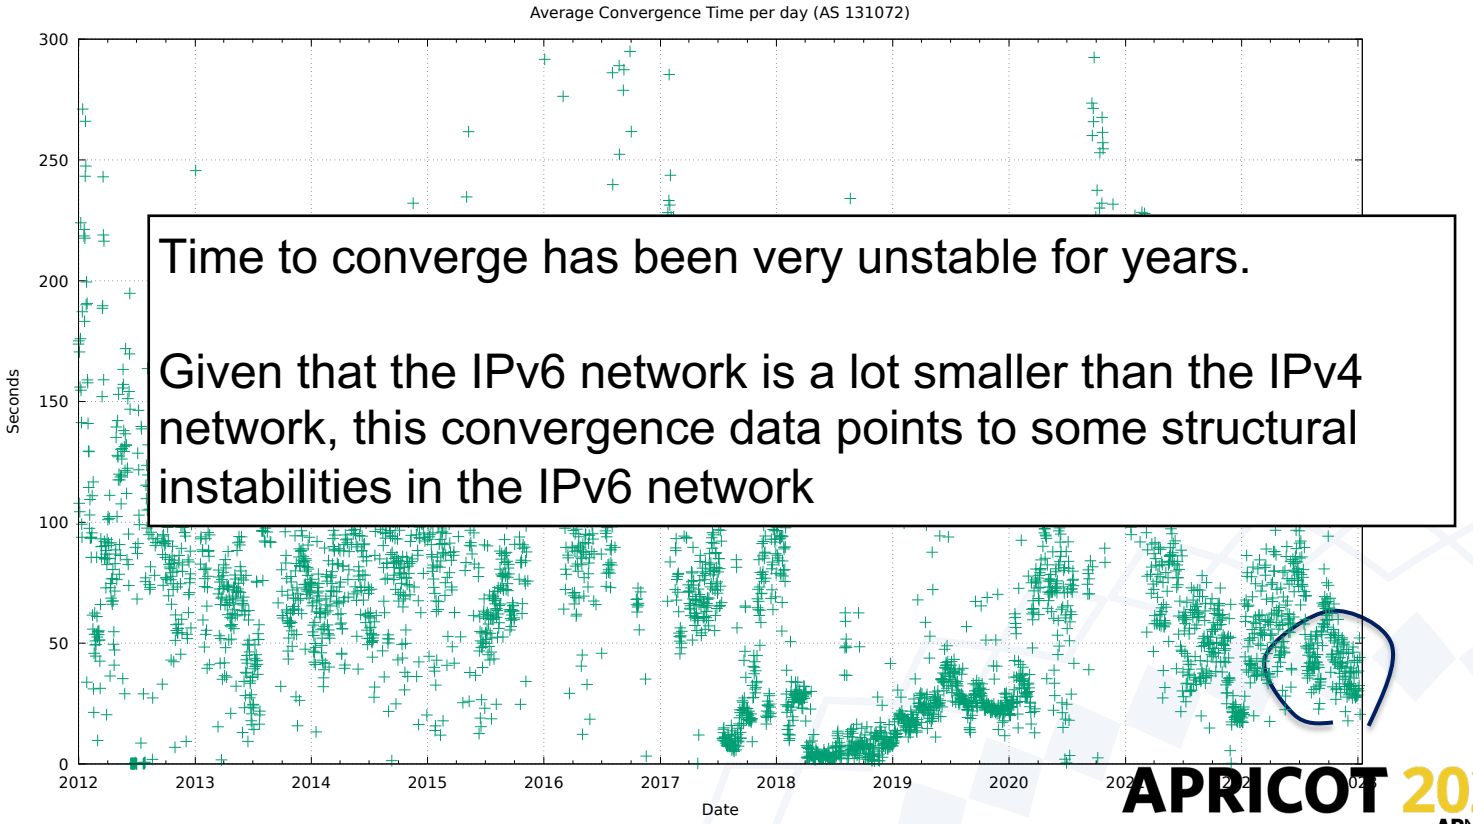

Time to converge has been very unstable for years.

Given that the IPv6 network is a lot smaller than the IPv4 network, this convergence data points to some structural instabilities in the IPv6 network

Updates in IPv6 BGP

It's improving …

- Compared to IPv4, the IPv6 network has exhibited 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 2 AS's generated 70% of the BGP update load in the last 2 weeks of 2022. IPv6 routing instability is still concentrated in a small number of pathological cases.

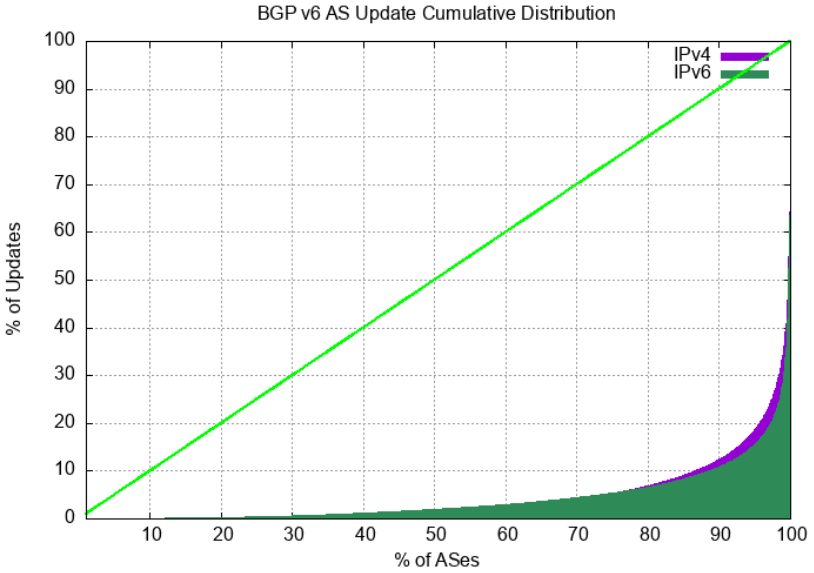

The Highlights

- 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 inter-AS 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”

Routing Futures

- 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 lesser-used routes to a slower / cheaper lookup path. This approach may also become common in very high-capacity line cards

Some Practical Suggestions

Know your network's limits

- Understand your routing hardware's line card FIB capacity in the default-free parts of your network

Some Practical Suggestions

Know your network's limits

Review your routers' settings

- Review your IPv4 / IPv6 portioning in the FIB tables - a dual-stack eBGP router will need 1M 32-bit IPv4 slots and 320K 128-bit IPv6 slots for a full eBGP routing table in line cards in 2 years time if they are using a full eBGP FIB load (plus internal routes of course). That's roughly the same memory footprint for IPv4 and IPv6!

Some Practical Suggestions

Know your network's limits

Review your routers' settings

Default routes can be helpful

- Judicious use of **default** routes in your internal network may allow you drop this high speed line card memory requirement significantly

Some Practical Suggestions

Know your network's limits

Review your routers' settings

Default routes can be helpful

Time for hot caching in line card FIBs?

- Using a hot cache for line card FIB cache would reduce the high-speed TCAM memory requirement significantly without visible performance cost

Some Practical Suggestions

Know your network's limits

Review your routers' settings

Default routes can be helpful

Time for hot caching in line card FIBs?

That's it!