BGP in 2021

Geoff Huston
APNIC Labs

The Highlights

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

28 Years of Routing the Internet

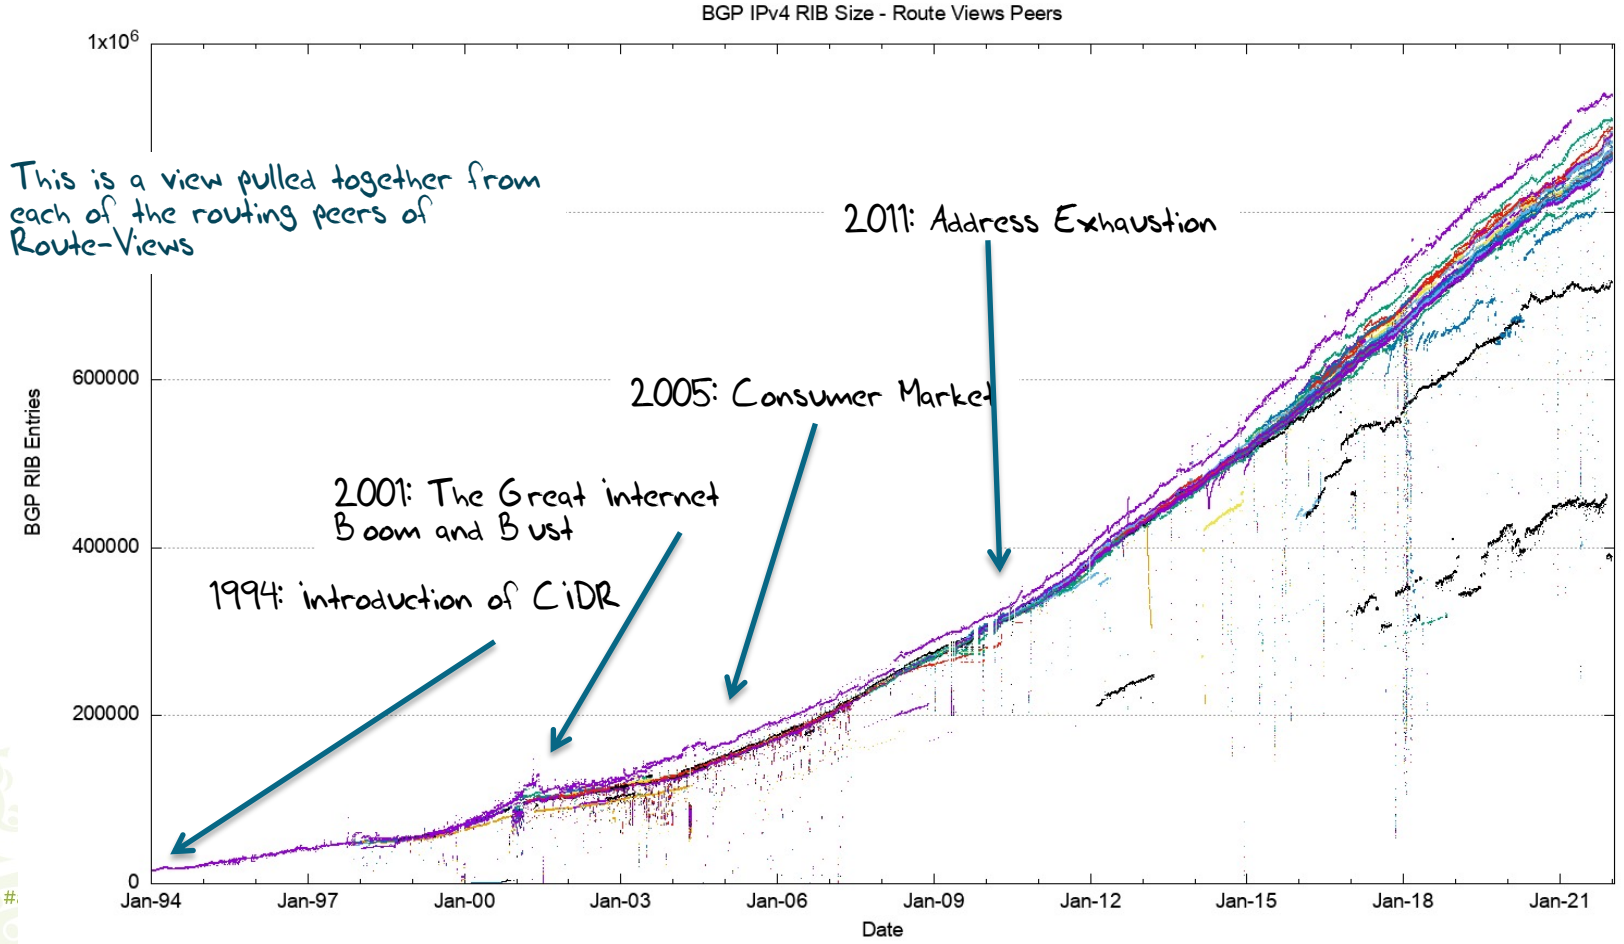

2021 in detail

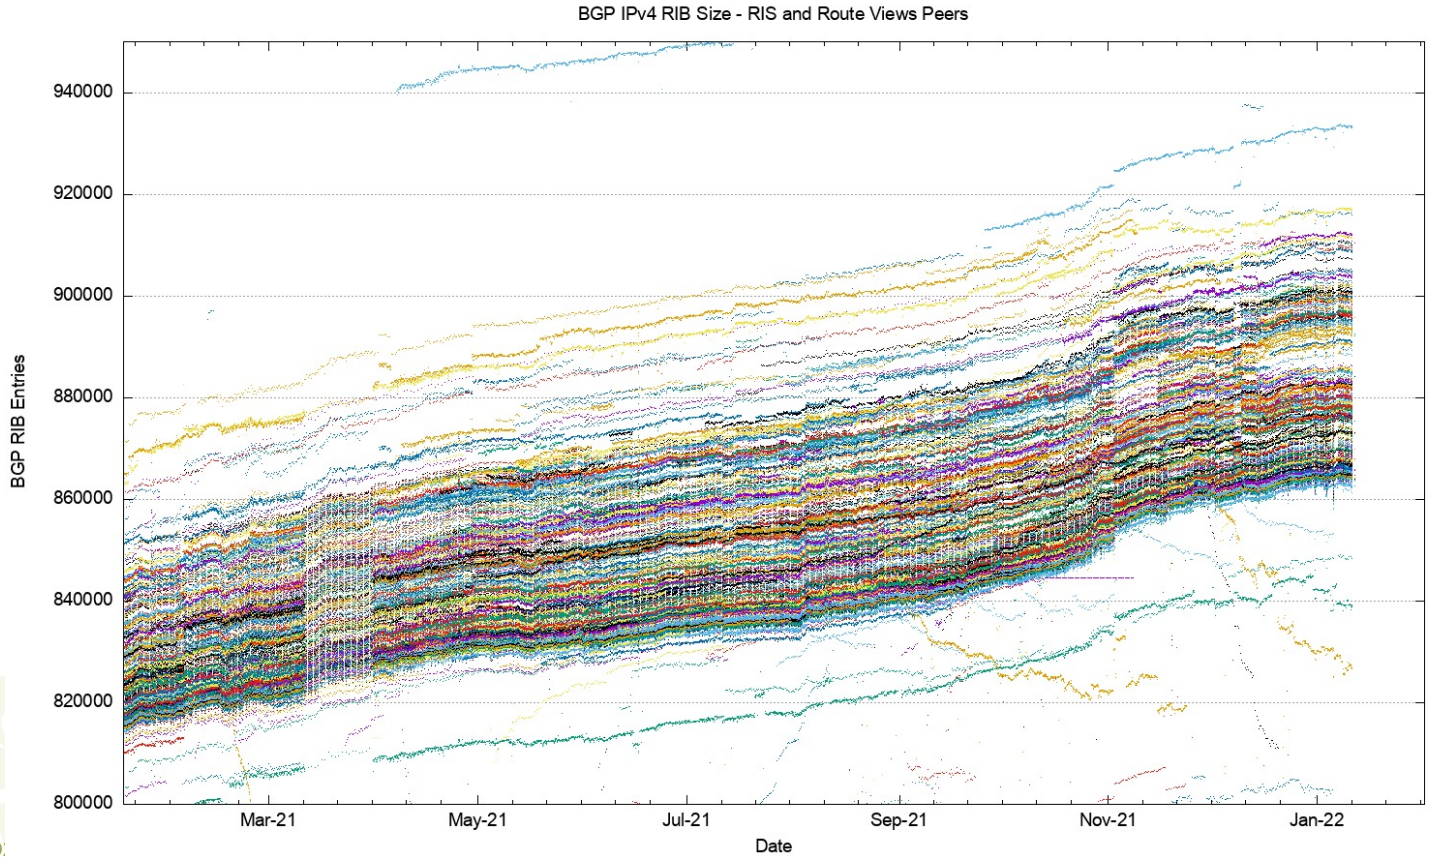

2021 in detail

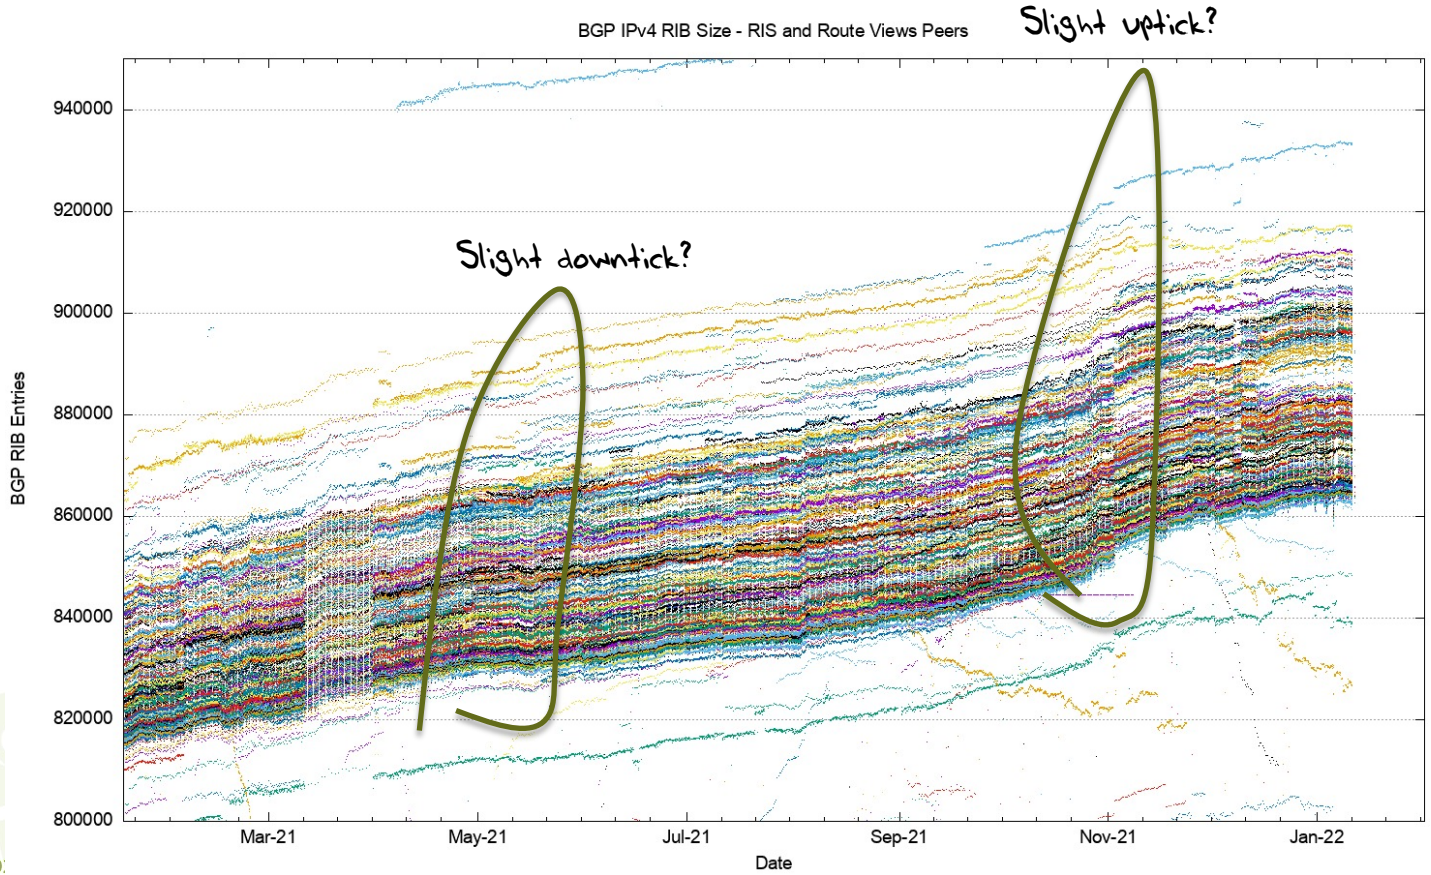

2021: Assigned vs Recovered

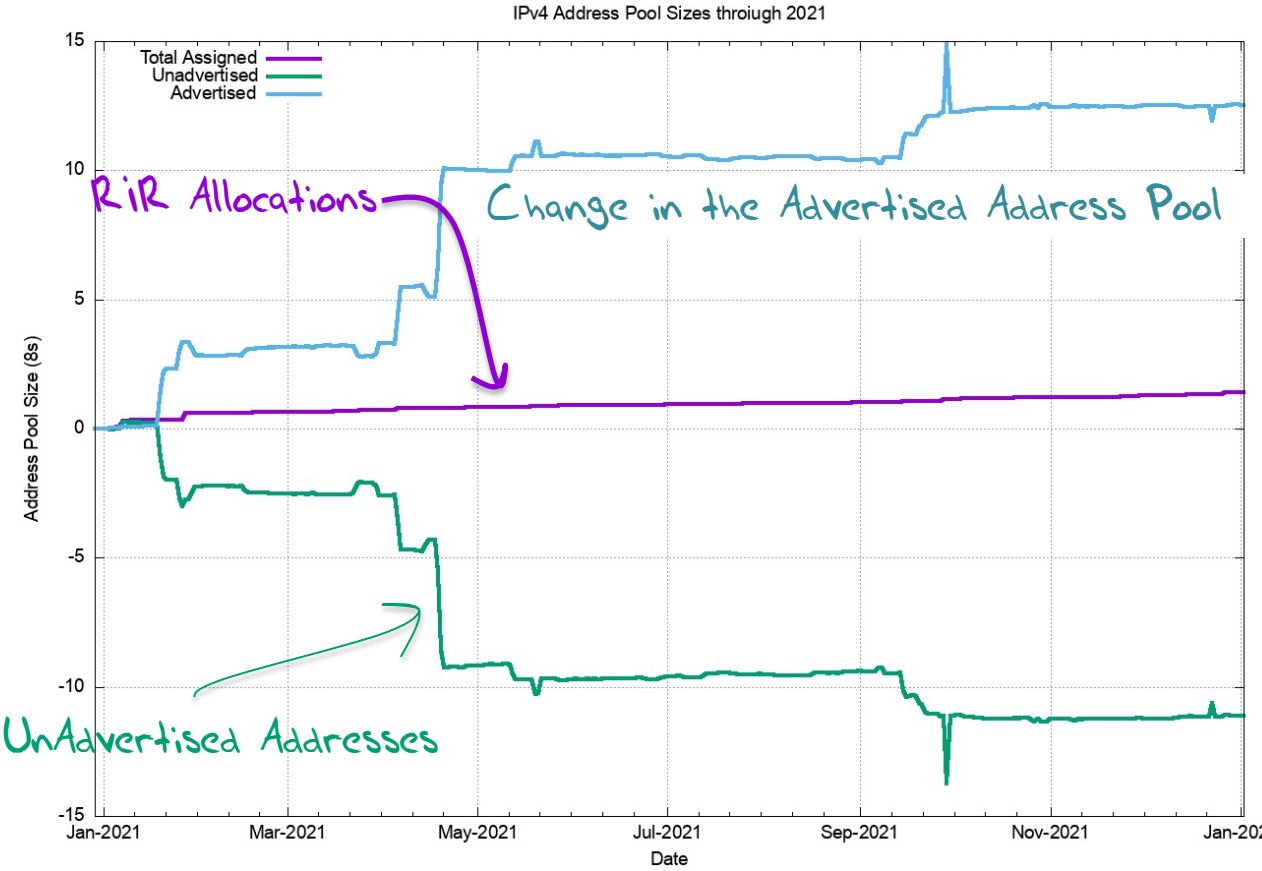

218M advertised

24M allocated

186M recovered

What happened in 2021 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 906,000 by the end of 2021
- The pace of growth of the routing table was slightly lower than the rolling 5-year average, with **40,000 new entries in 2021** (was 52,000 in 2020)
- The AS position was slightly lower with **2,400 new AS's advertised in 2021** (was 3,400 in 2020)
- Transit relationships have not changed materially over 2021 for most networks
- **The overall growth trends are slowing down in 2021**

The Highlights

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

The Route-Views View of IPv6

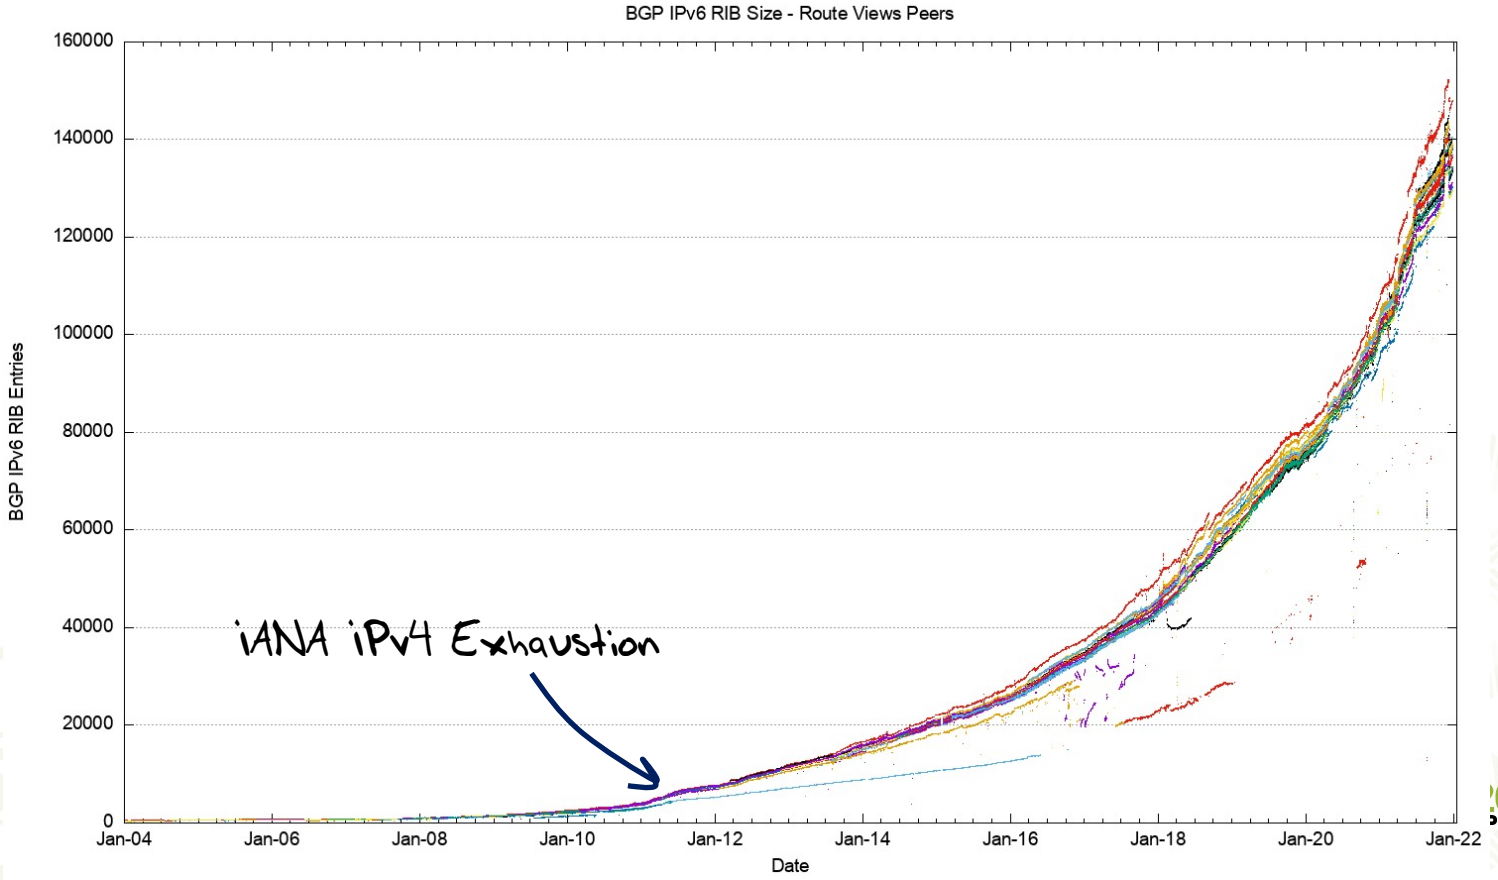

2021 in Detail

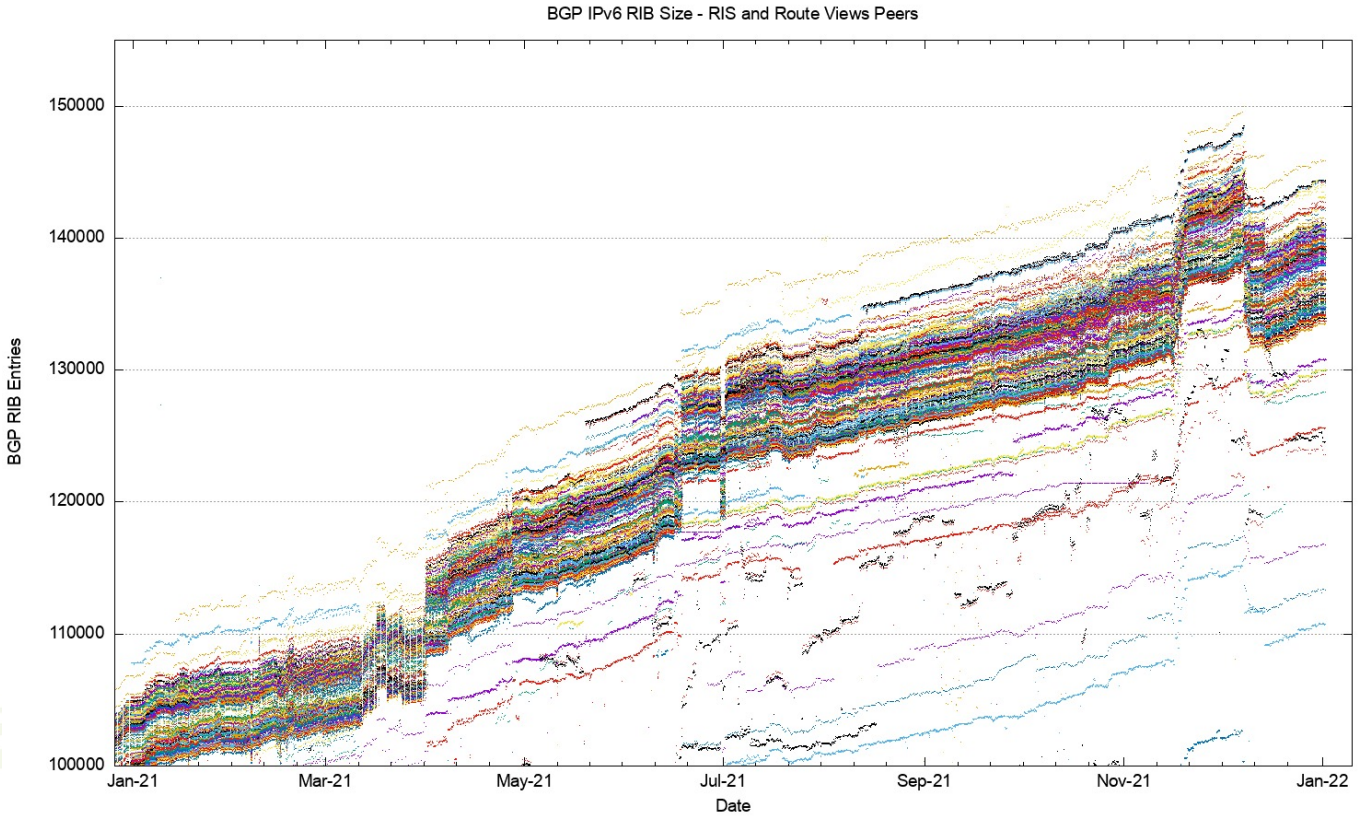

2021 in Detail

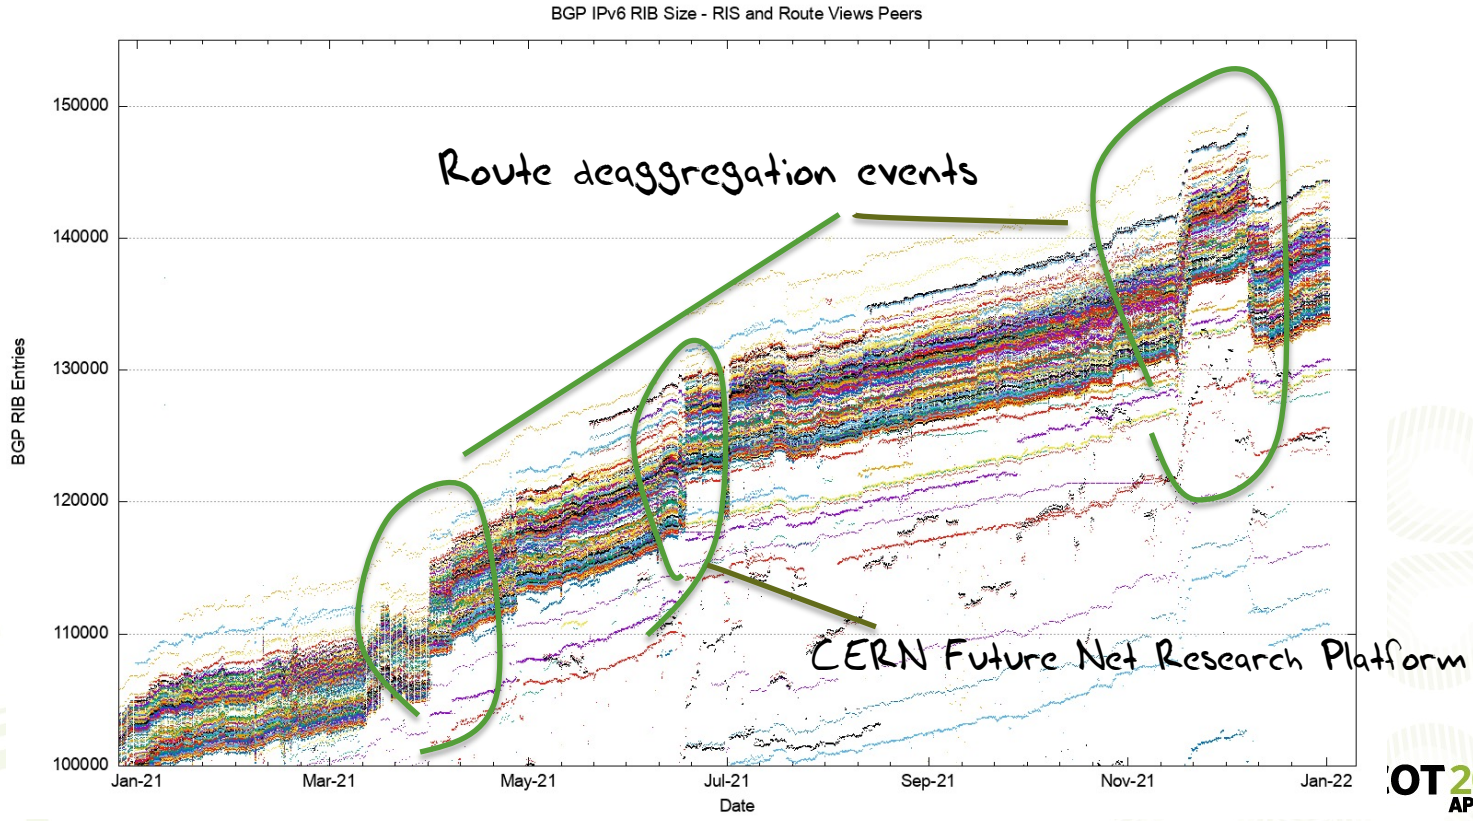

V6 in 2021

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **41,000 route entries p.a.**
  - With a couple of deaggregation leaks along the way!

- 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 *2018*	*699,000*	
2019	*760,000*	
2020	*814,000*	
2021	*866,000*	
2022	*906,000*	*908,000*
2023		*956,000*
2024		*1,004,000*
2025		*1,052,000*
2026		*1,100,000*
2027		*1,148,000*

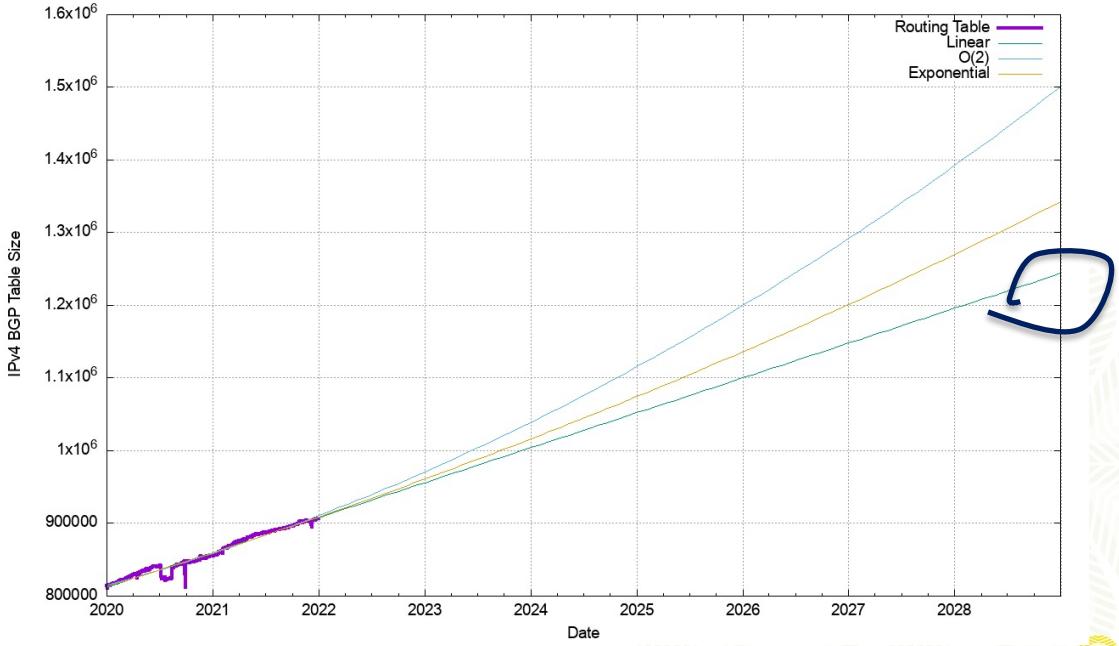

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2018*	*45,000*	
2019	*62,000*	
2020	*79,000*	
2021	*104,000*	
2022	*147,000*	*152,000*
2023	*185,000*	*214,000*
2024	*233,000*	*302,000*
2025	*261,000*	*426,000*
2026	*299,000*	*601,000*
2027	*366,000*	*847,000*

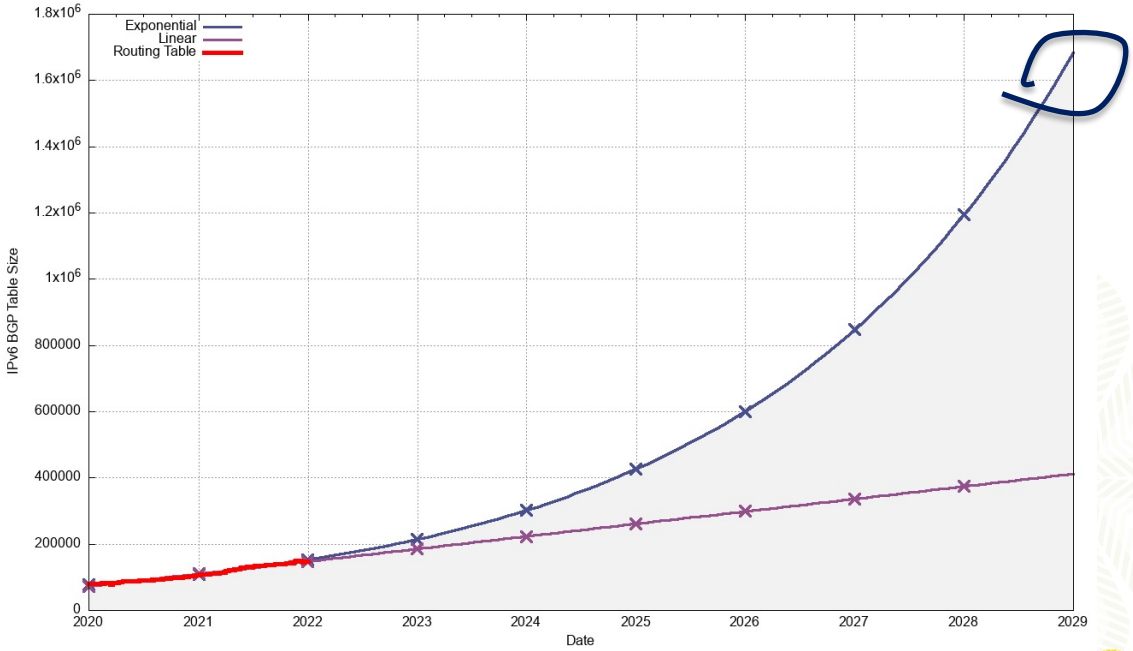

Note that the IPv6 tables are 128bits wide – i.e. 4x the size of the IPv4 tables!

BGP Table Growth

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

These two tables will require the same storage/lookup size in around 2 years time, 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

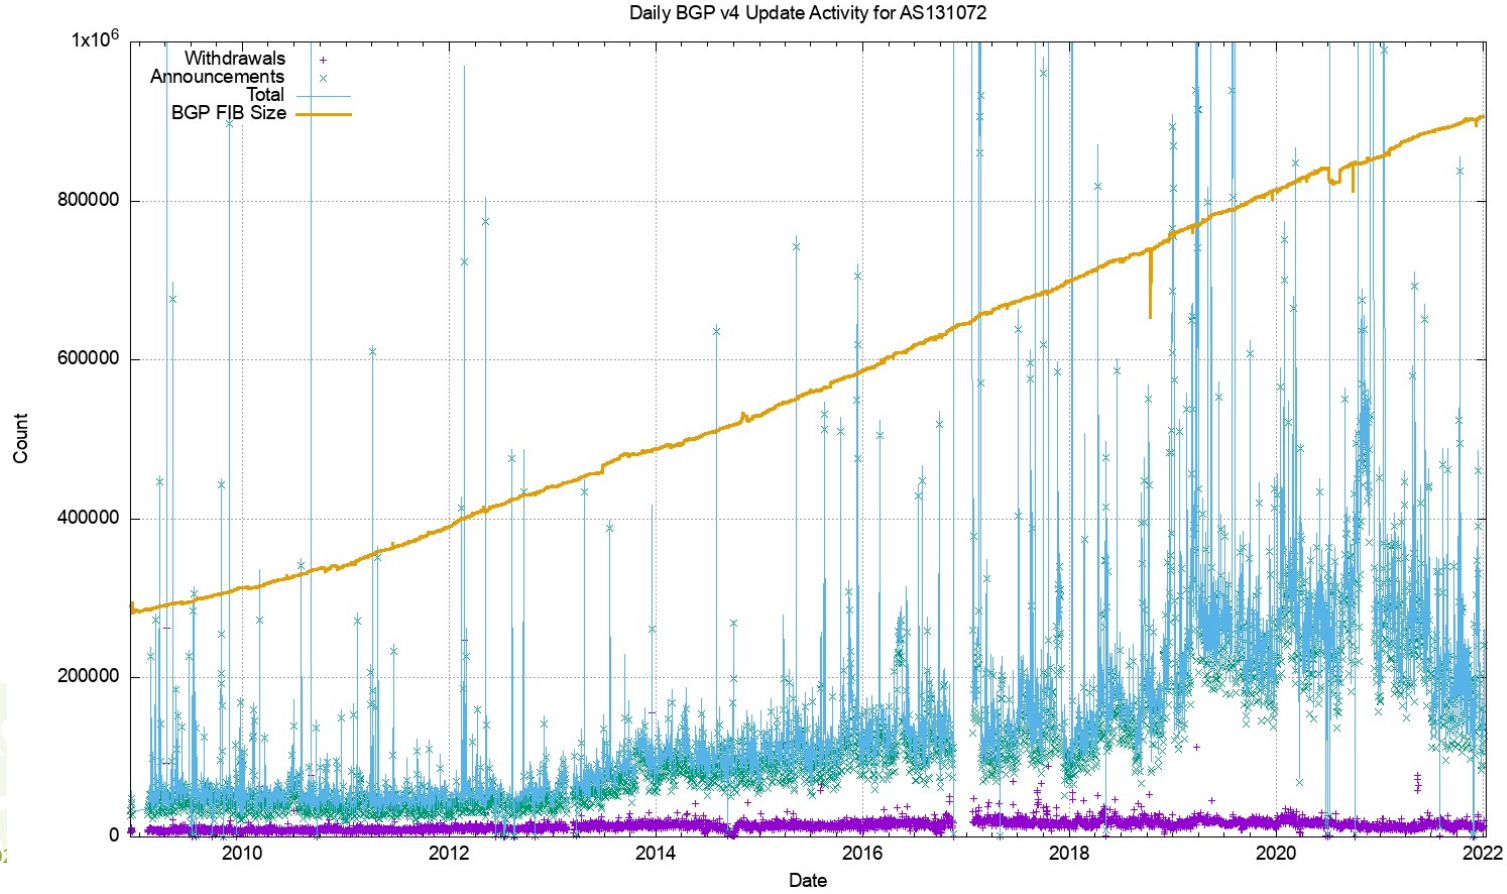

IPv4 BGP Convergence Performance

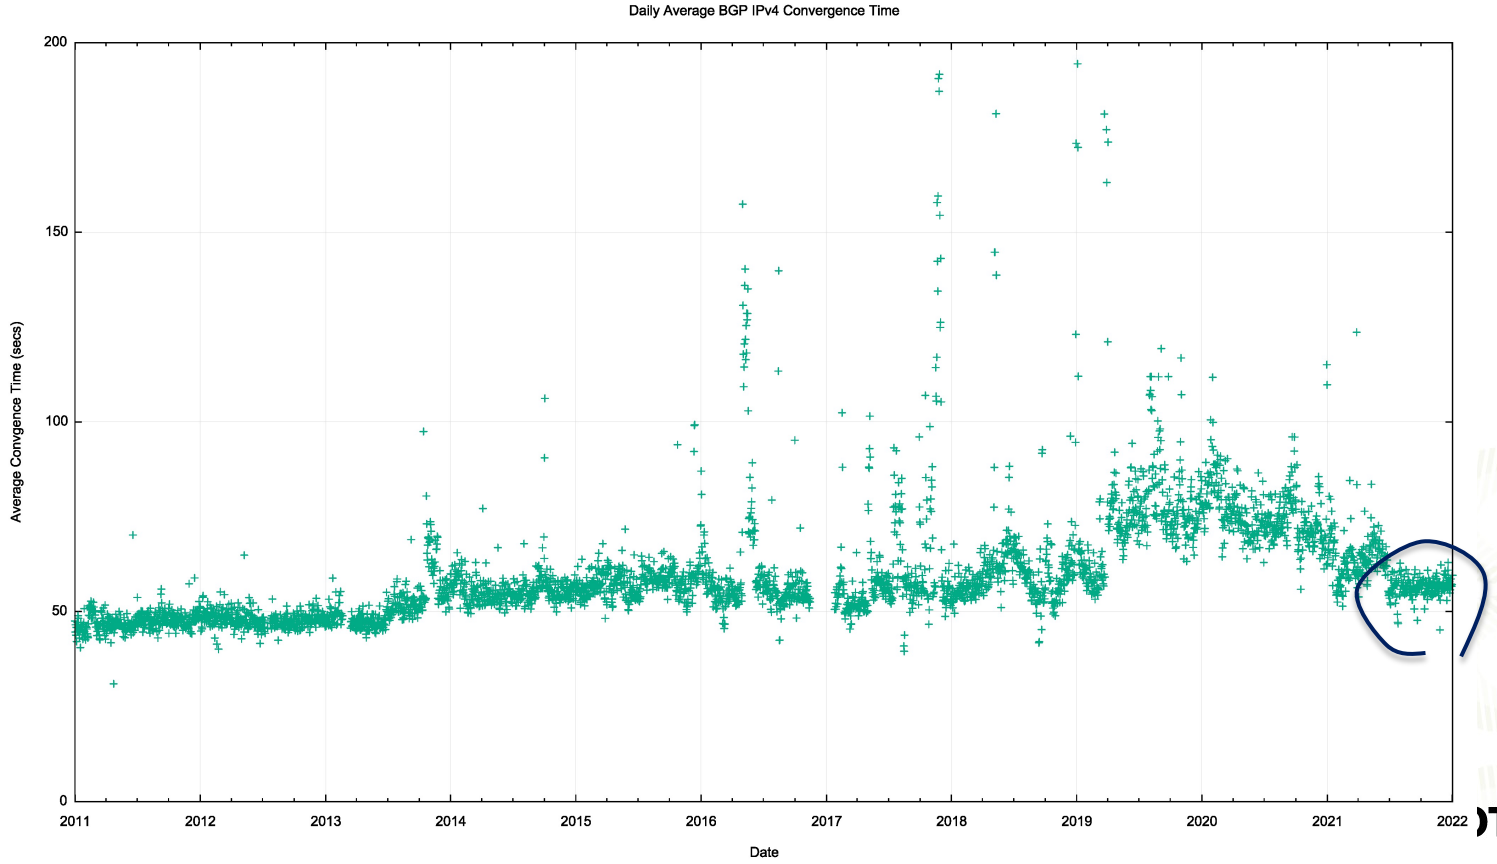

Updates in IPv4 BGP

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
- 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

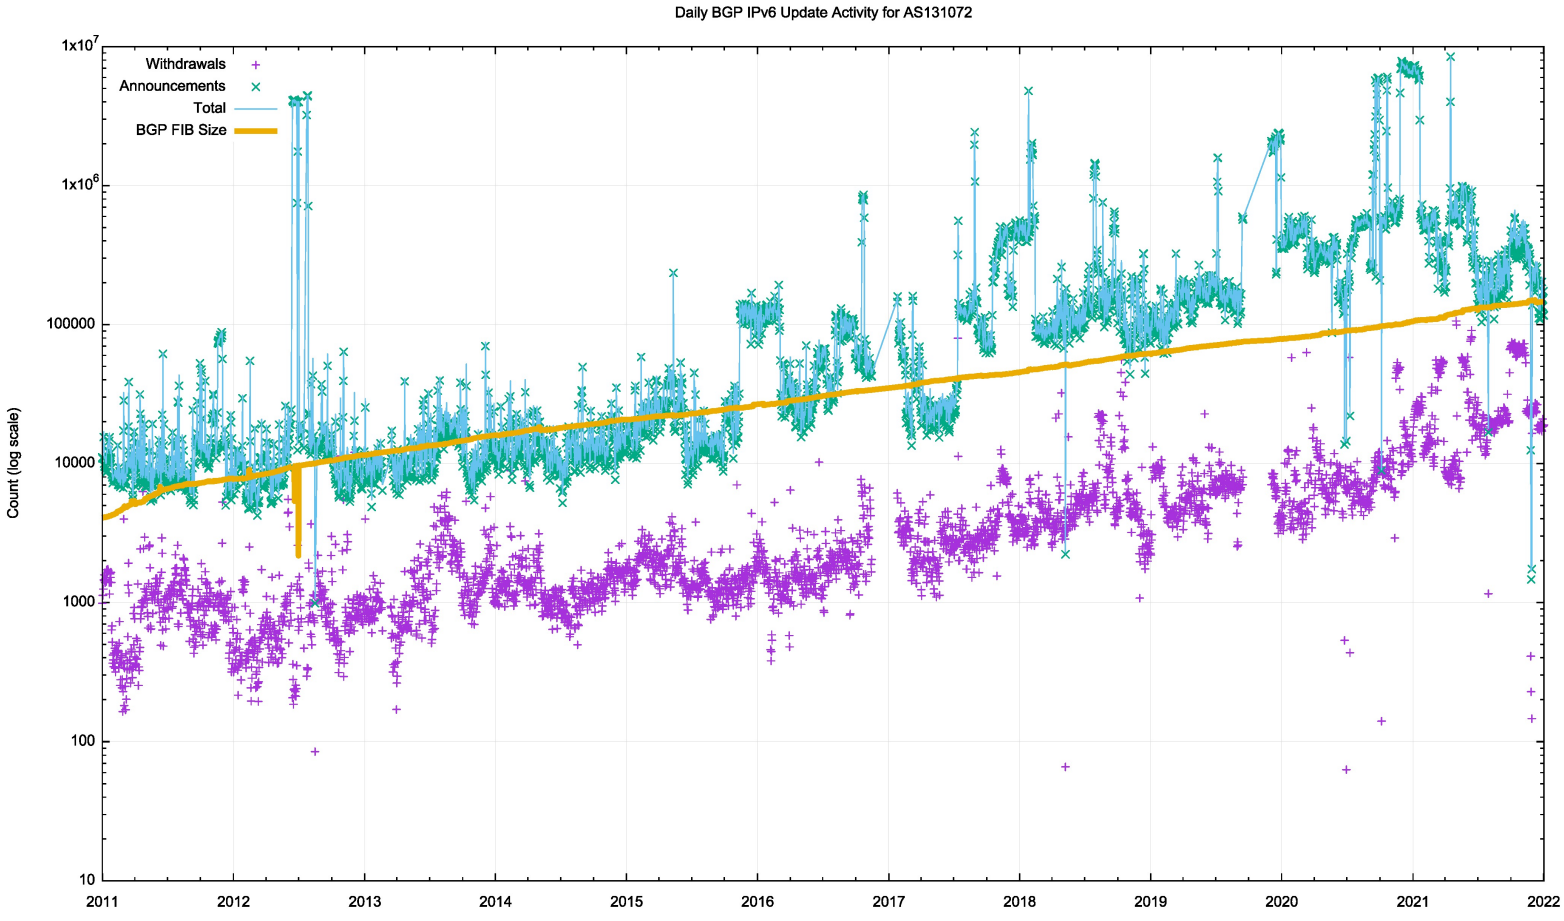

V6 Unstable Prefixes

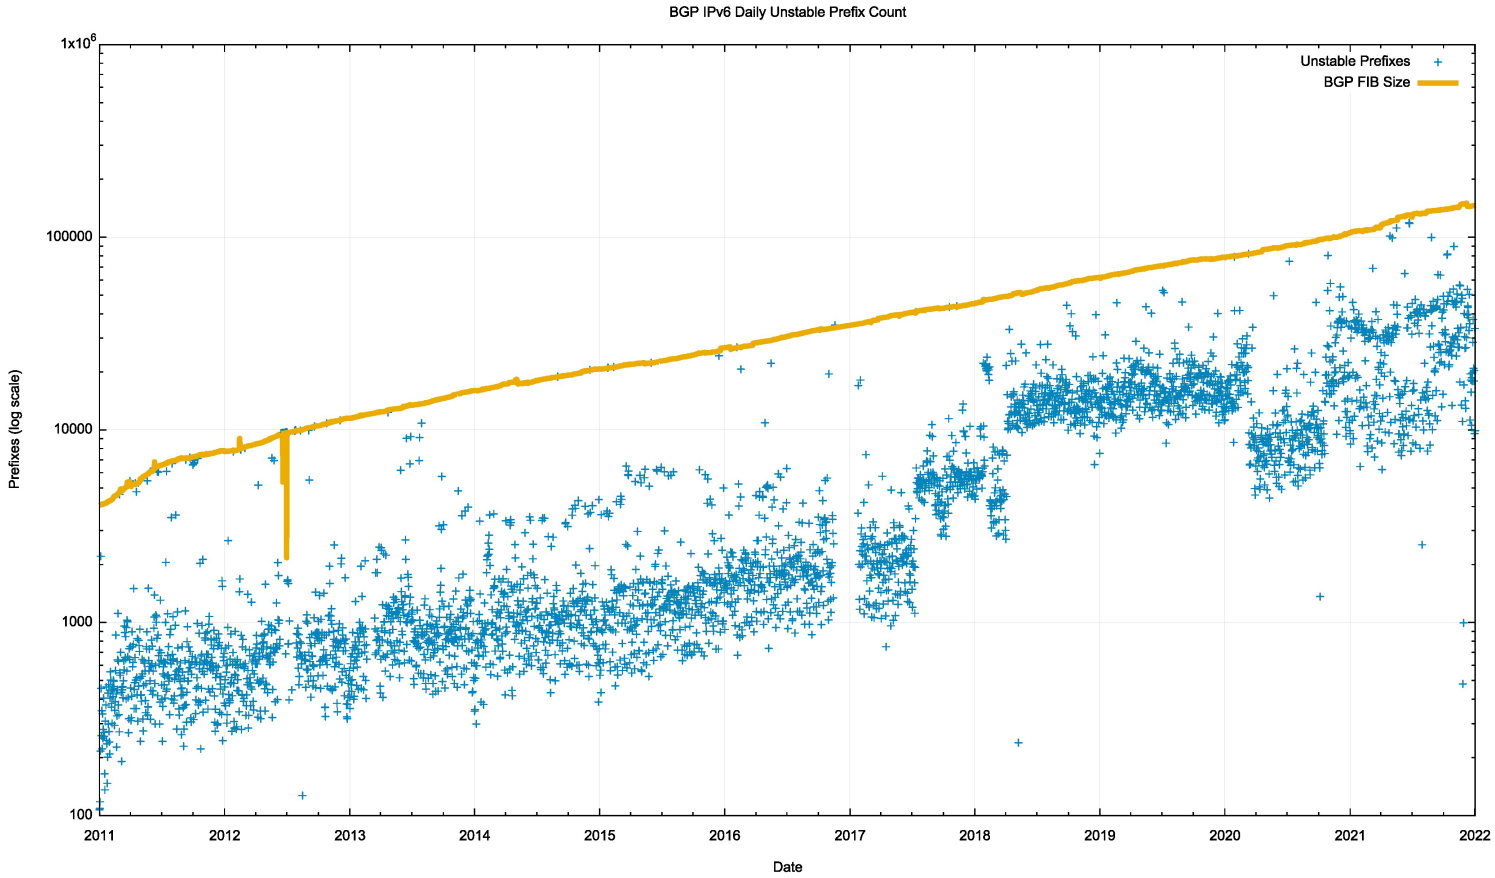

V6 Convergence Performance

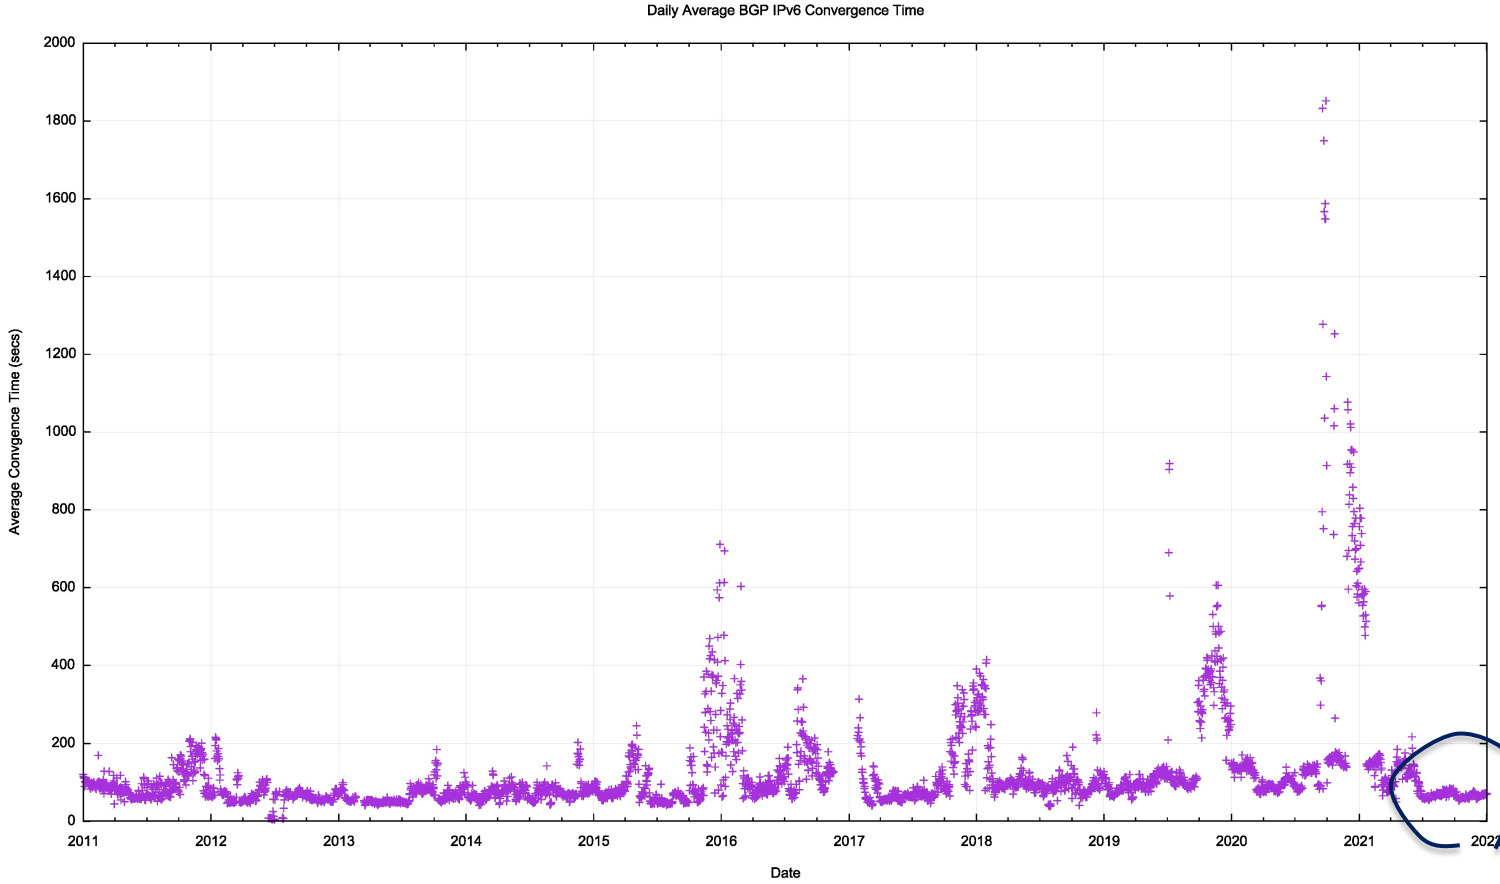

Updates in IPv6 BGP

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. The V6 network was more stable than ever in the last half of 2021
- Just 2 AS's generated 70% of the BGP update load in the last 2 weeks of 2021. Instability is still concentrated in a small number of pathological cases.

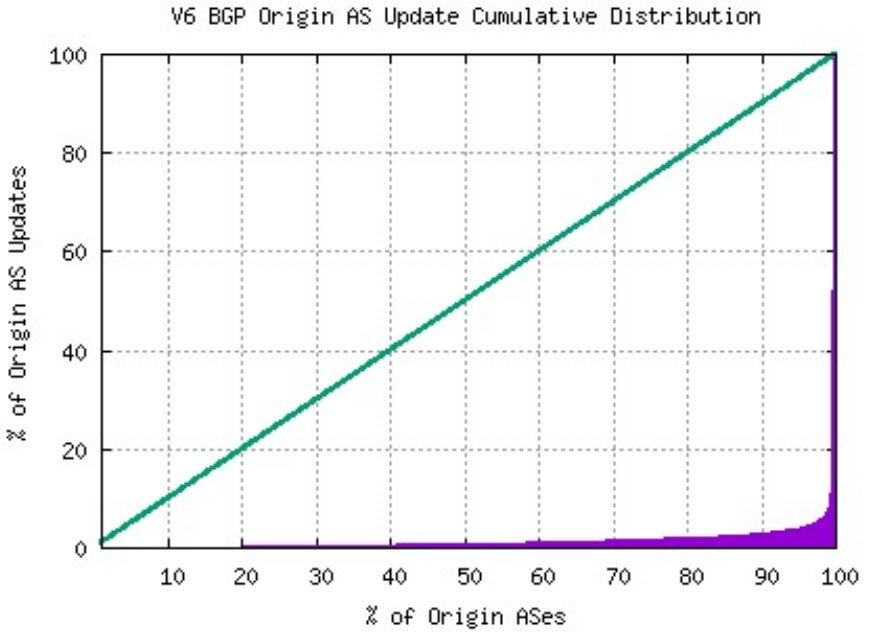

The Highlights

- 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”

- 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 233K 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 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 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!

The Complete Pack

27 Years of Routing the Internet

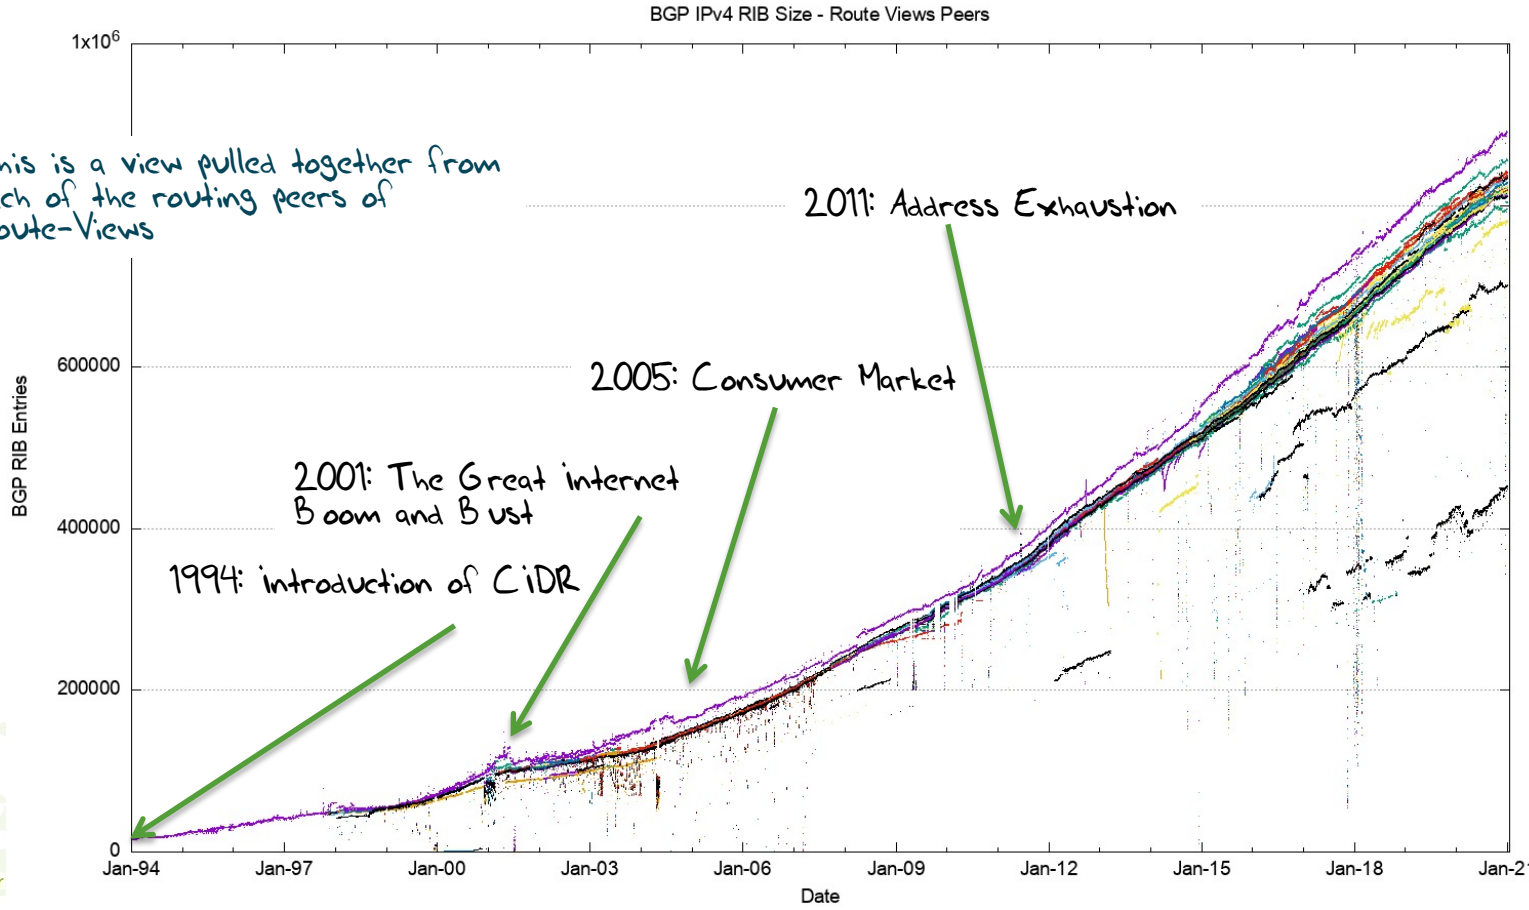

2017-2022 in detail

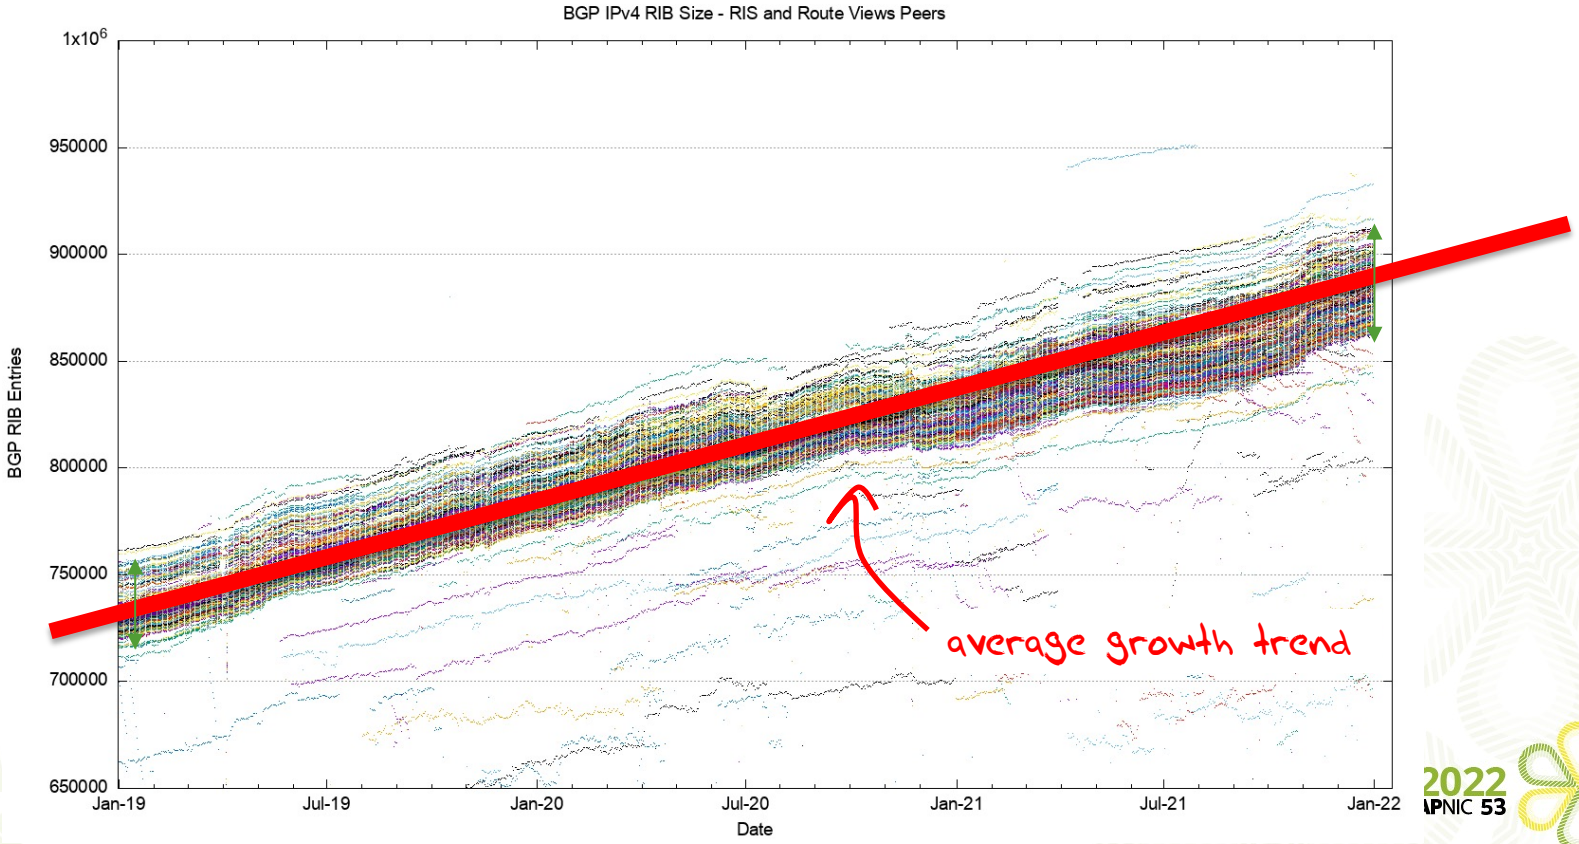

2021 in detail

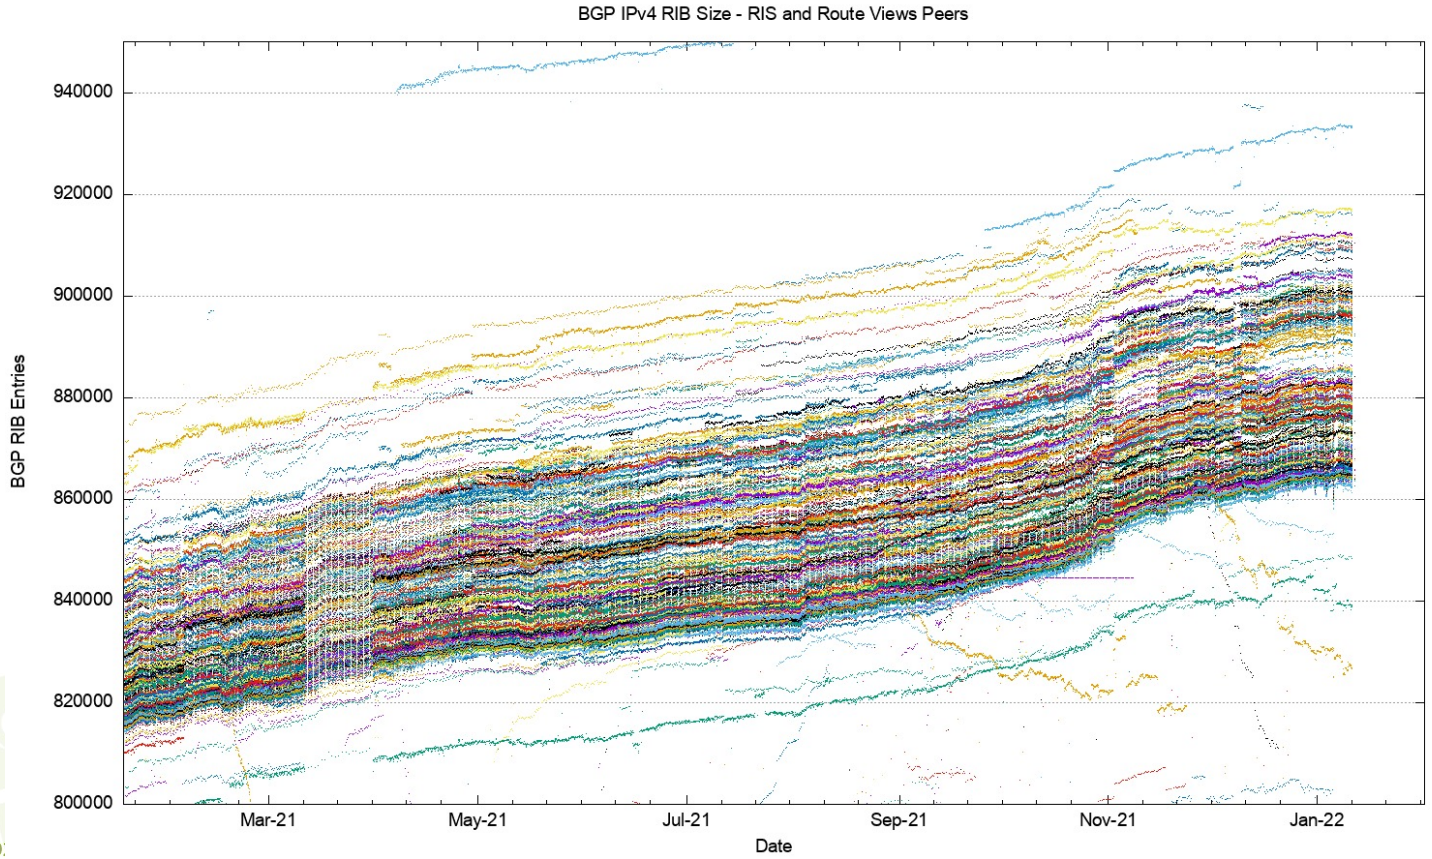

Routing Indicators for IPv4

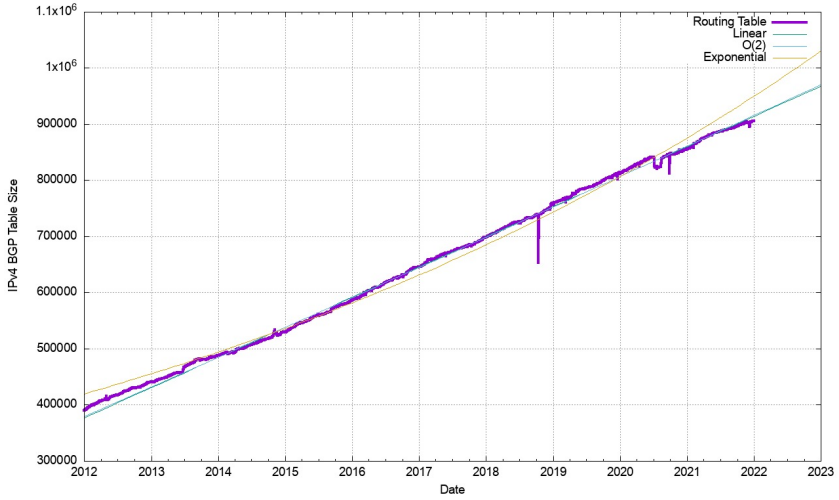

Routing prefixes – growing by some 40,000 prefixes per year

AS Numbers– growing by some 2,400 prefixes per year

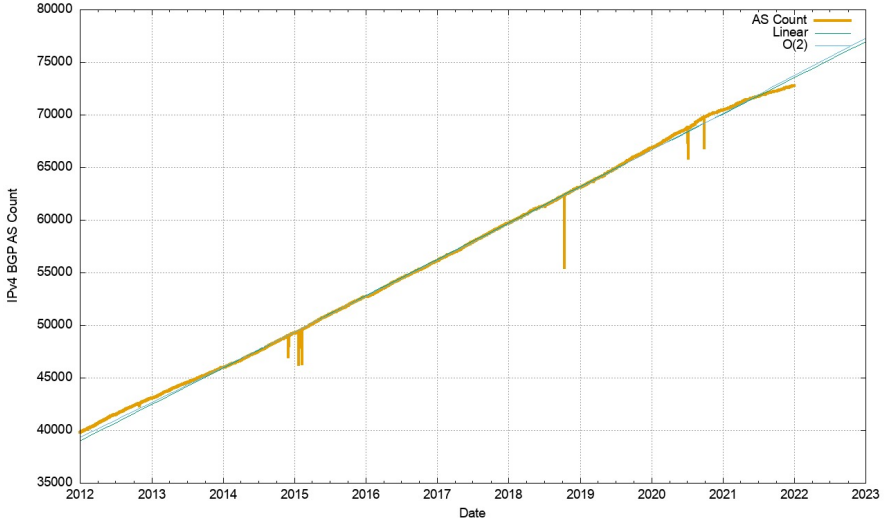

Routing Indicators for IPv4

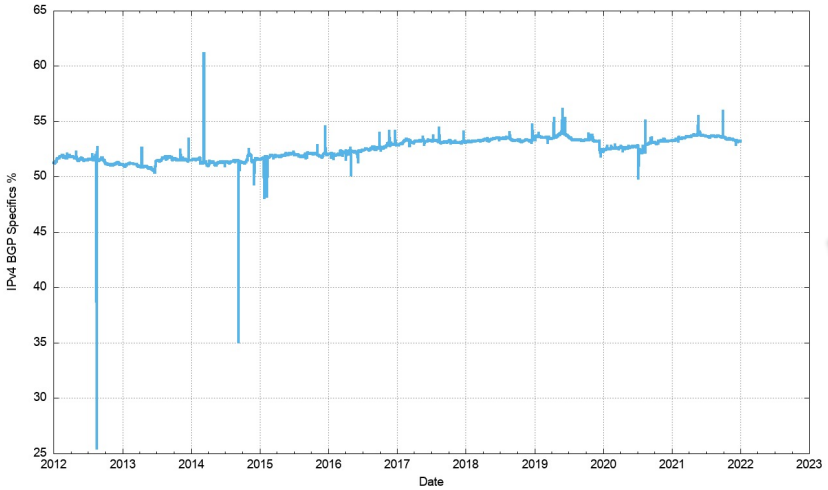

More Specifics are still taking up slightly more than one half of the routing table

But the average size of an IPv4 routing advertisement continues to shrink

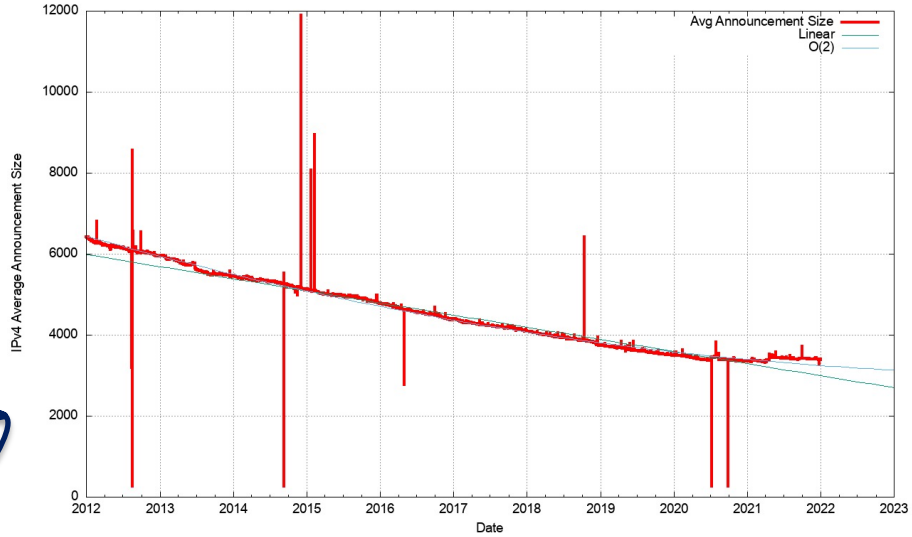

Routing Indicators for IPv4

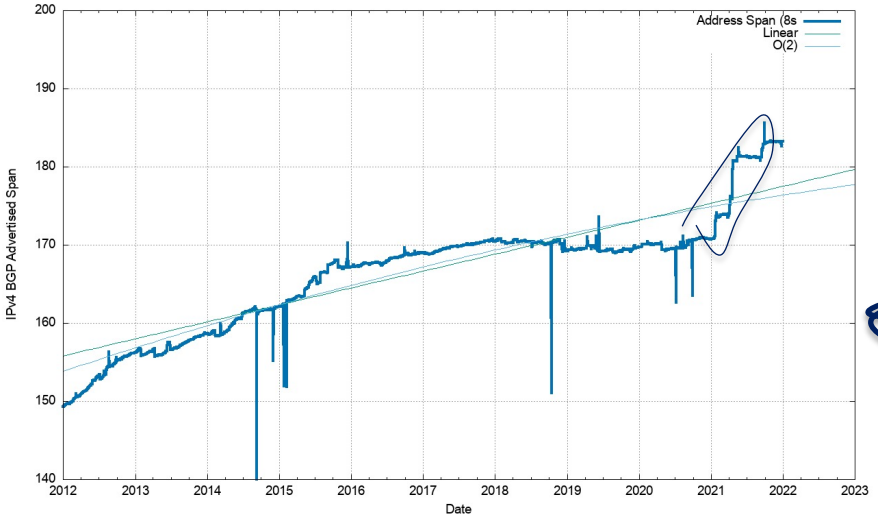

The US DoD advertised its collection of legacy /8 prefixes in 2021

The "shape" of inter-AS interconnection appears to be trending to a more compact topology

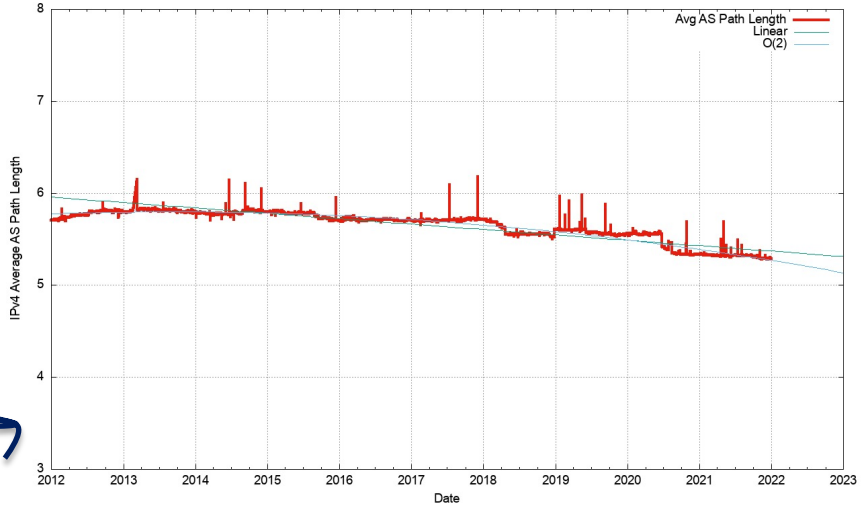

AS Adjacencies (as seen by AS131072)

59,044 out of 72,930 ASNs have 1 or 2 AS Adjacencies (82%)

2,425 ASNs have 10 or more adjacencies

9 ASNs have >1,000 adjacencies

6,059	AS6939	HURRICANE - Hurricane Electric, Inc., US
4,977	AS3356	LEVEL3 - Level 3 Communications, Inc., US
3,743	AS174	COGENT-174 - Cogent Communications, US
2,015	AS6461	ZAYO Bandwidth, US
1,776	AS3257	GTT-Backbone, DE
1,742	AS7018	ATT-INTERNET4 - AT&T Services, Inc., US
1,372	AS1299	TWELVE99 Arelion (fka Telia Carrier), SE
1,243	AS2914	NTT America, US
1,023	AS12389	ROSTELECOM, RU

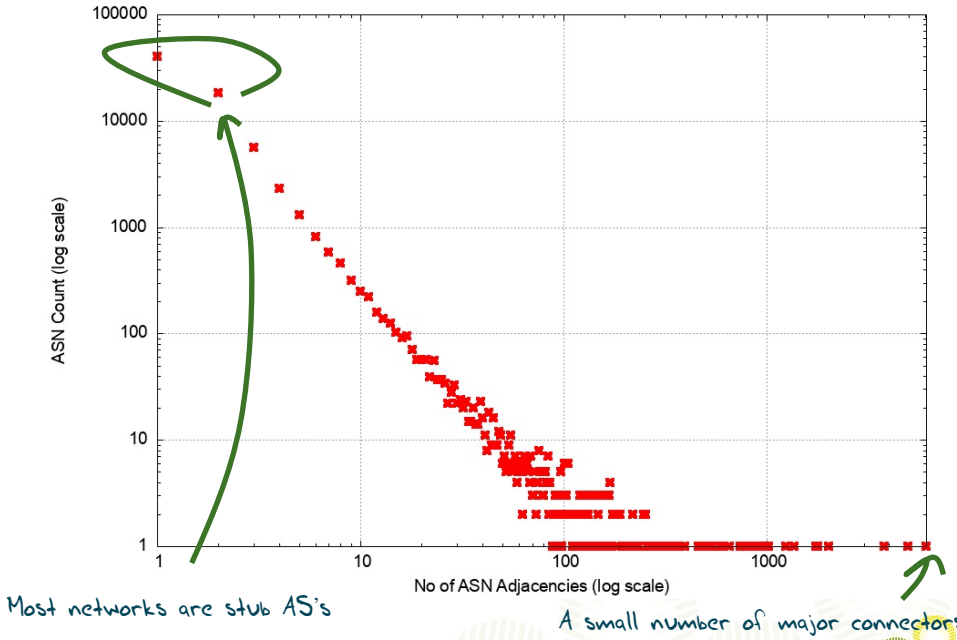

What happened in 2021 in V4?

- 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 906,000 by the end of 2021
- The pace of growth of the routing table was slightly lower than the rolling 5-year average, with **40,000 new entries in 2021** (was 52,000 in 2020)
- The AS position was slightly lower with 2,400 new AS's advertised in 2021 (was 3,400 in 2020)
- Transit relationships have not changed materially over 2021 for most networks
- The overall growth trends are slowing down in 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

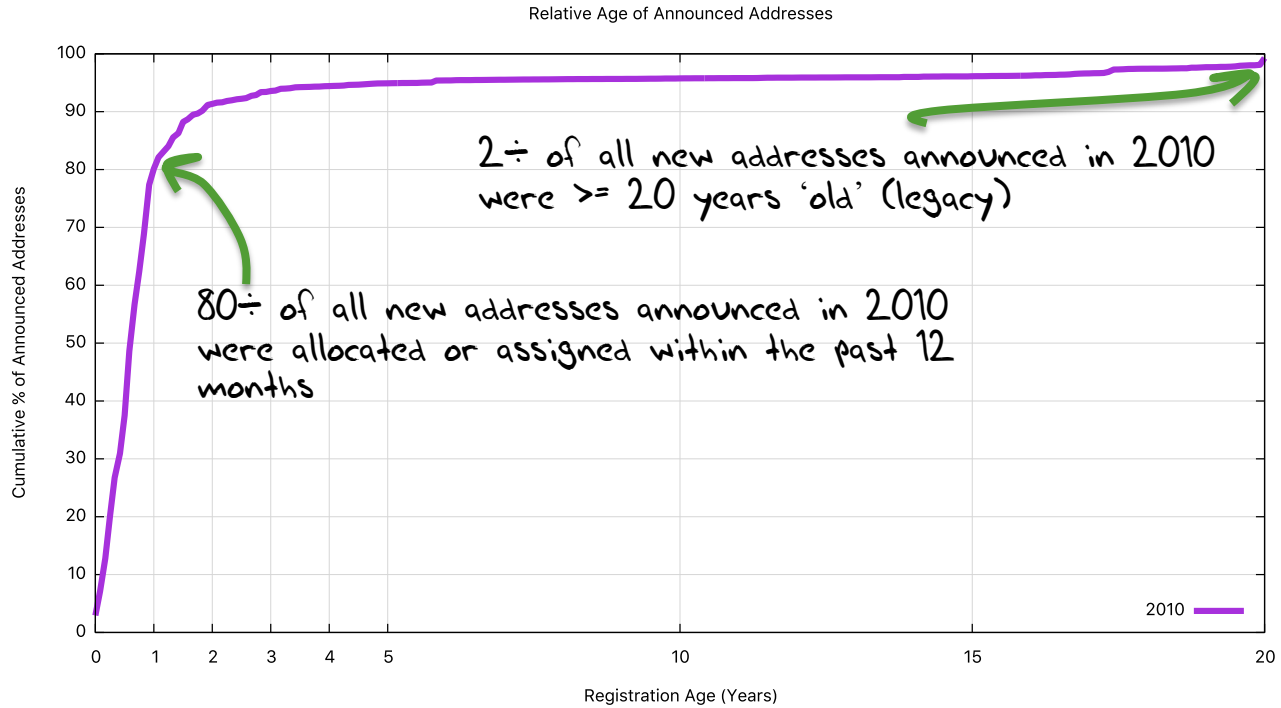

Advertised Address "Age"

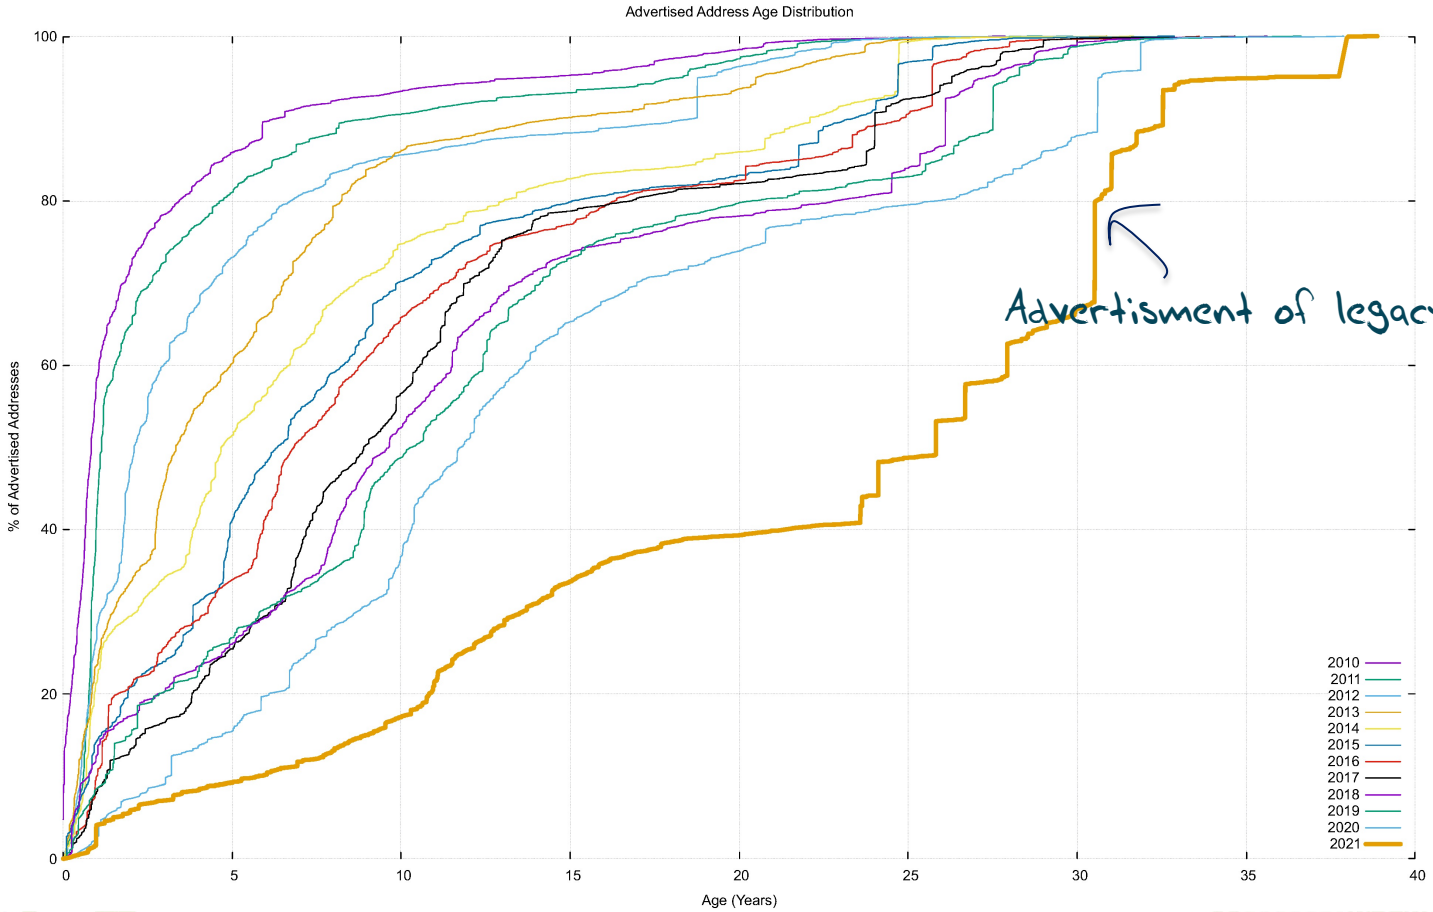

Advertisment of legacy addresses

IPv4 Advertised vs Unadvertised

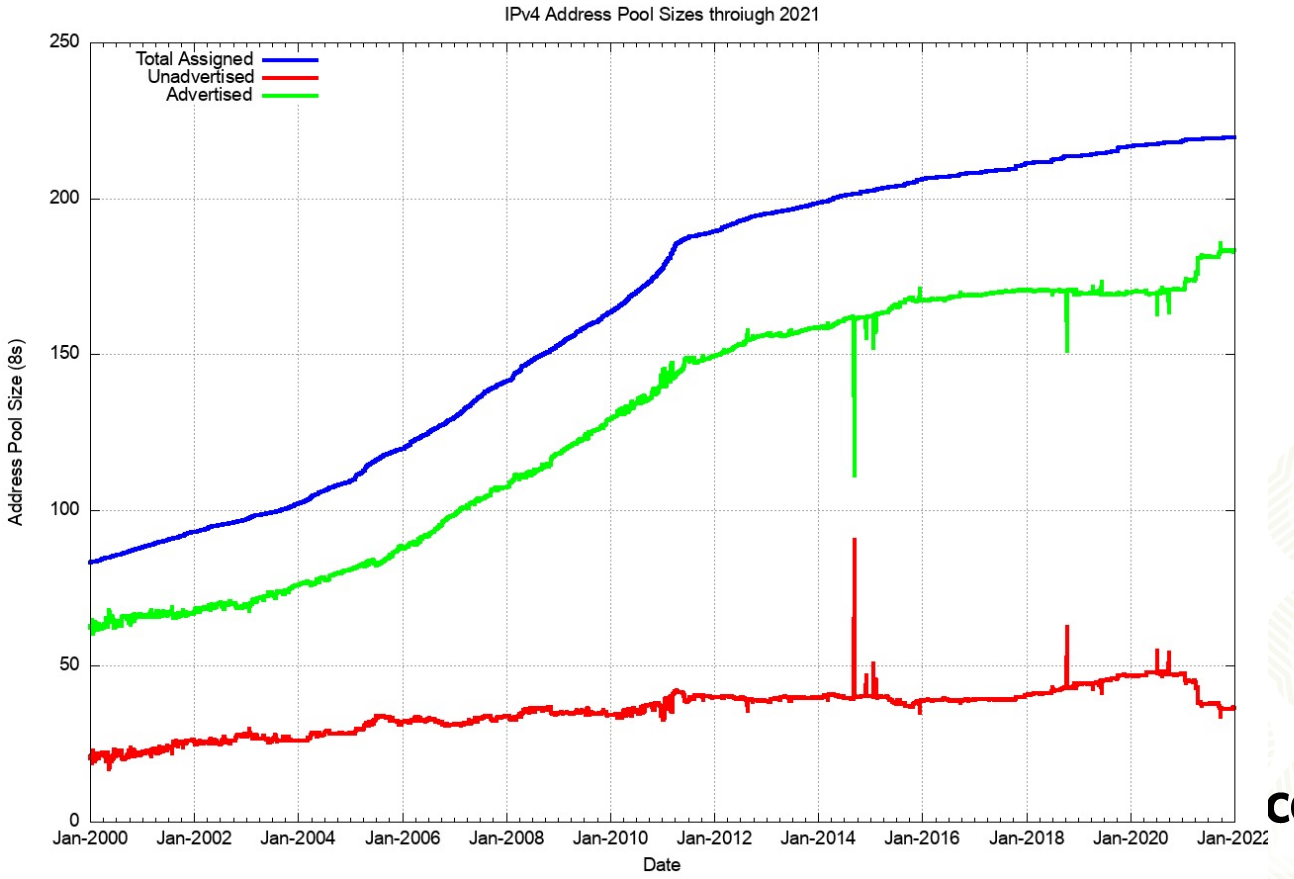

2010 - 2021: Unadvertised Addresses

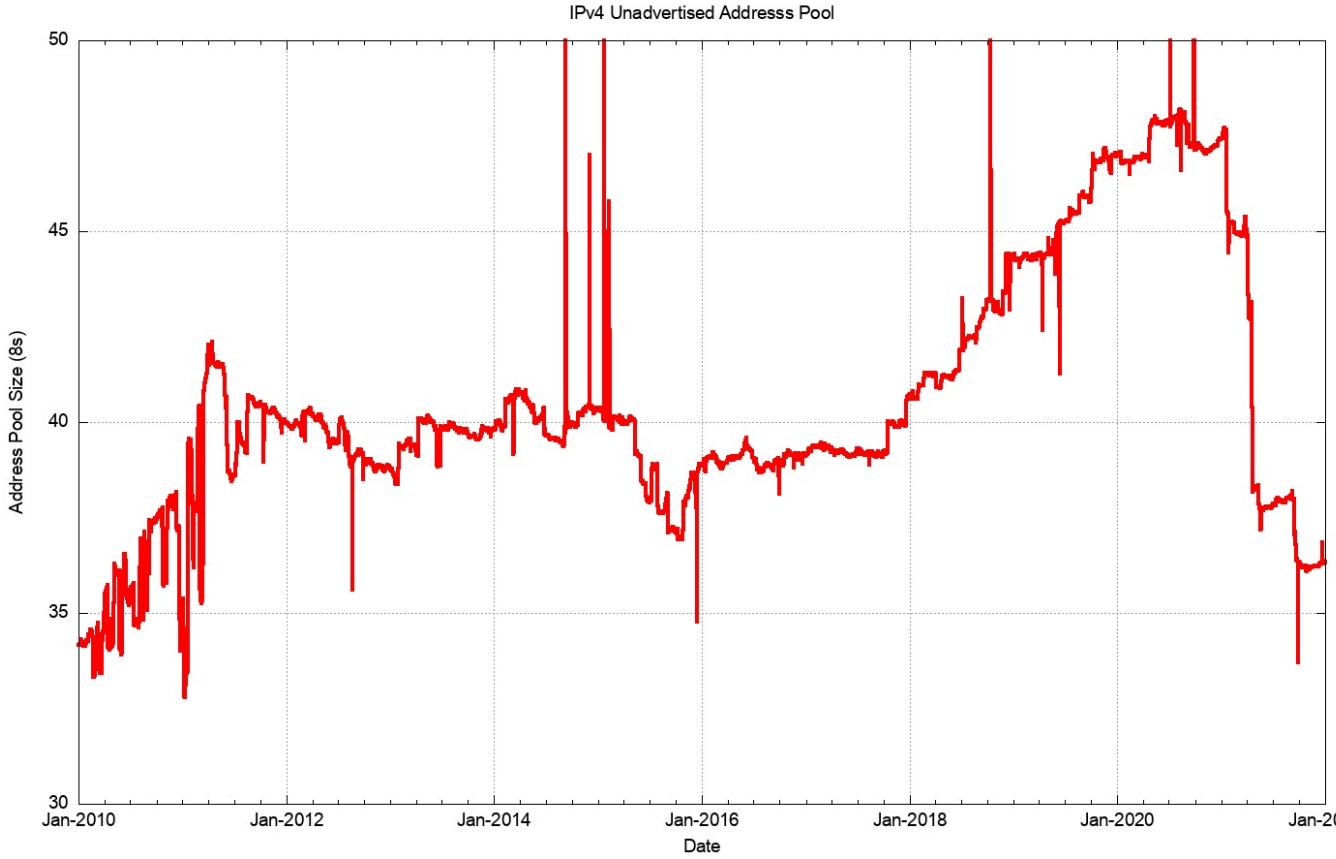

2021: Assigned vs Recovered

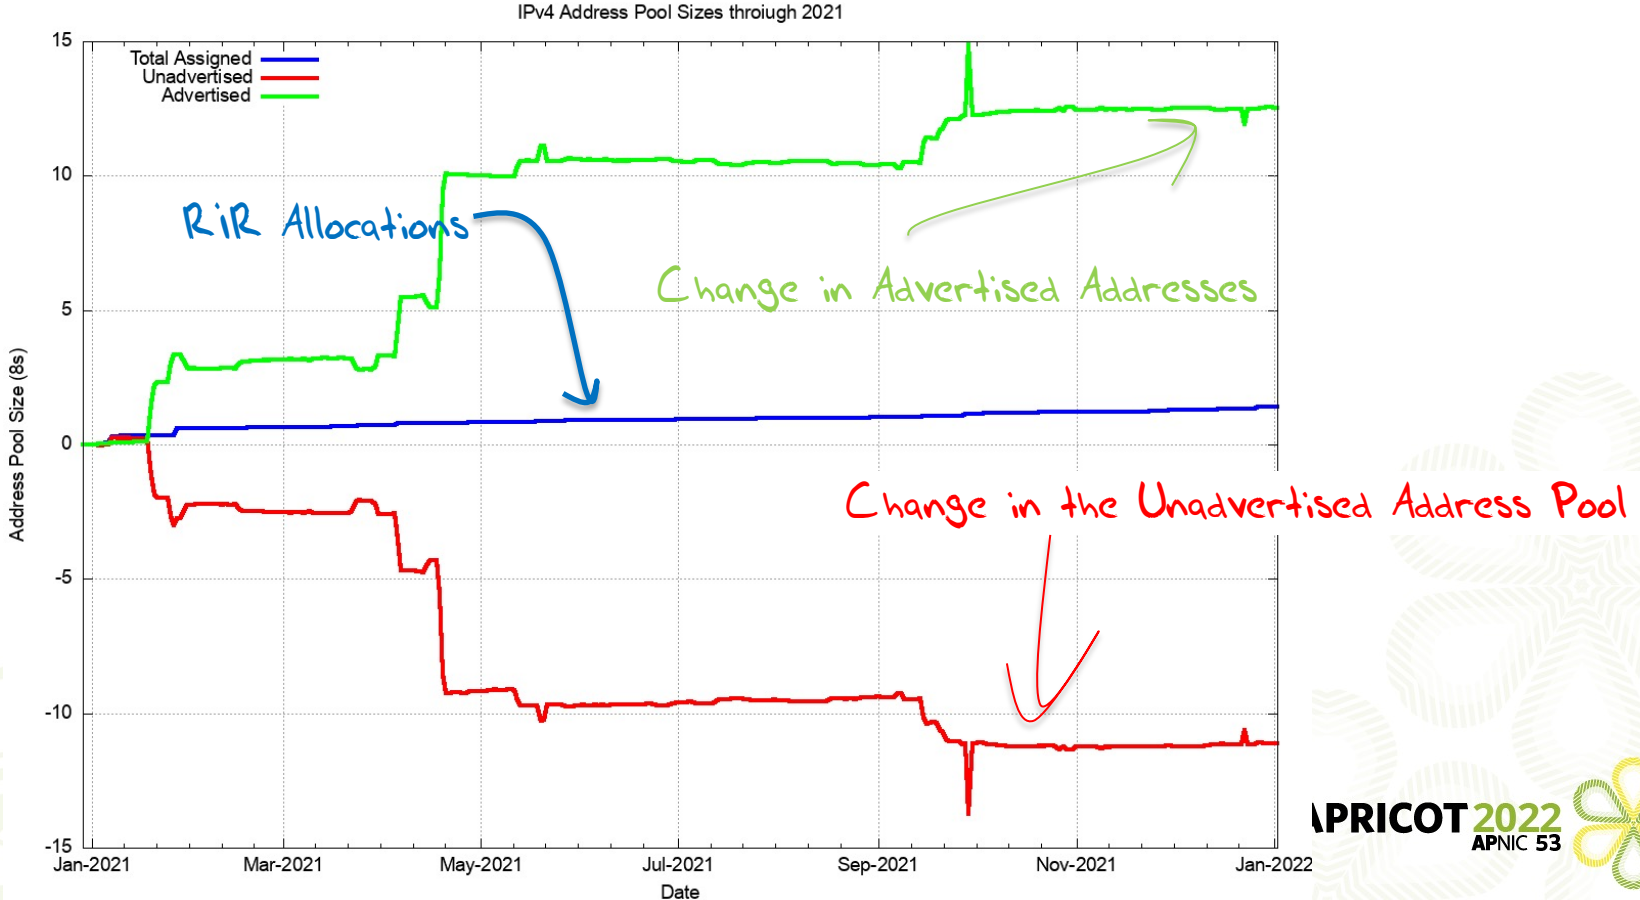

V4 in 2021

- 199.6M addresses were **added** to the routing table across 2020
- 1.1M addresses were **assigned** by RIRs in 2019
- And a net of 198.5M addresses were **drawn** from the pool of unadvertised addresses

The major shift in 2021 was the advertisement of the previously dormant legacy /8s that were assigned to the US DoD in the early days of the ARPANET/Internet

The Route-Views View of IPv6

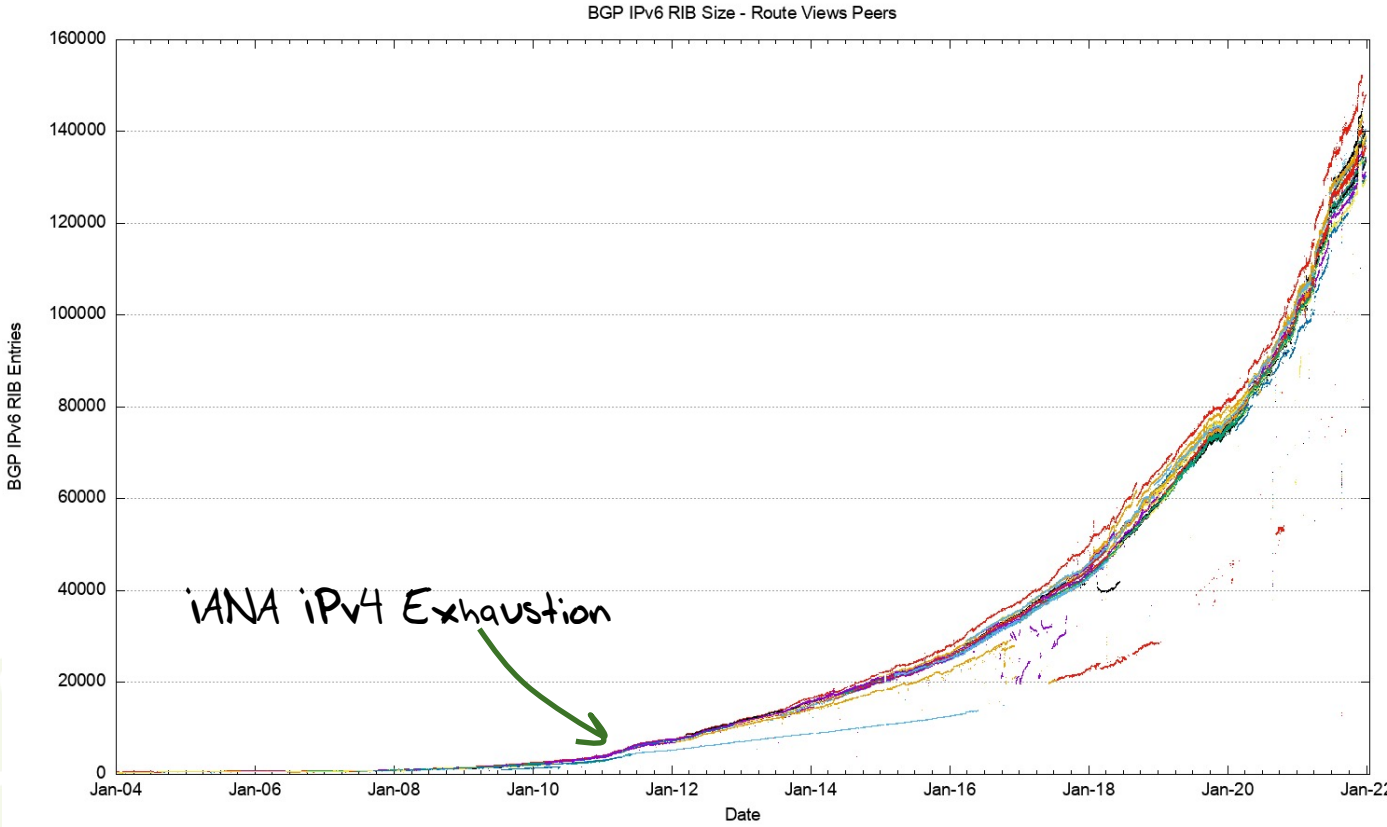

2020-2022 in Detail

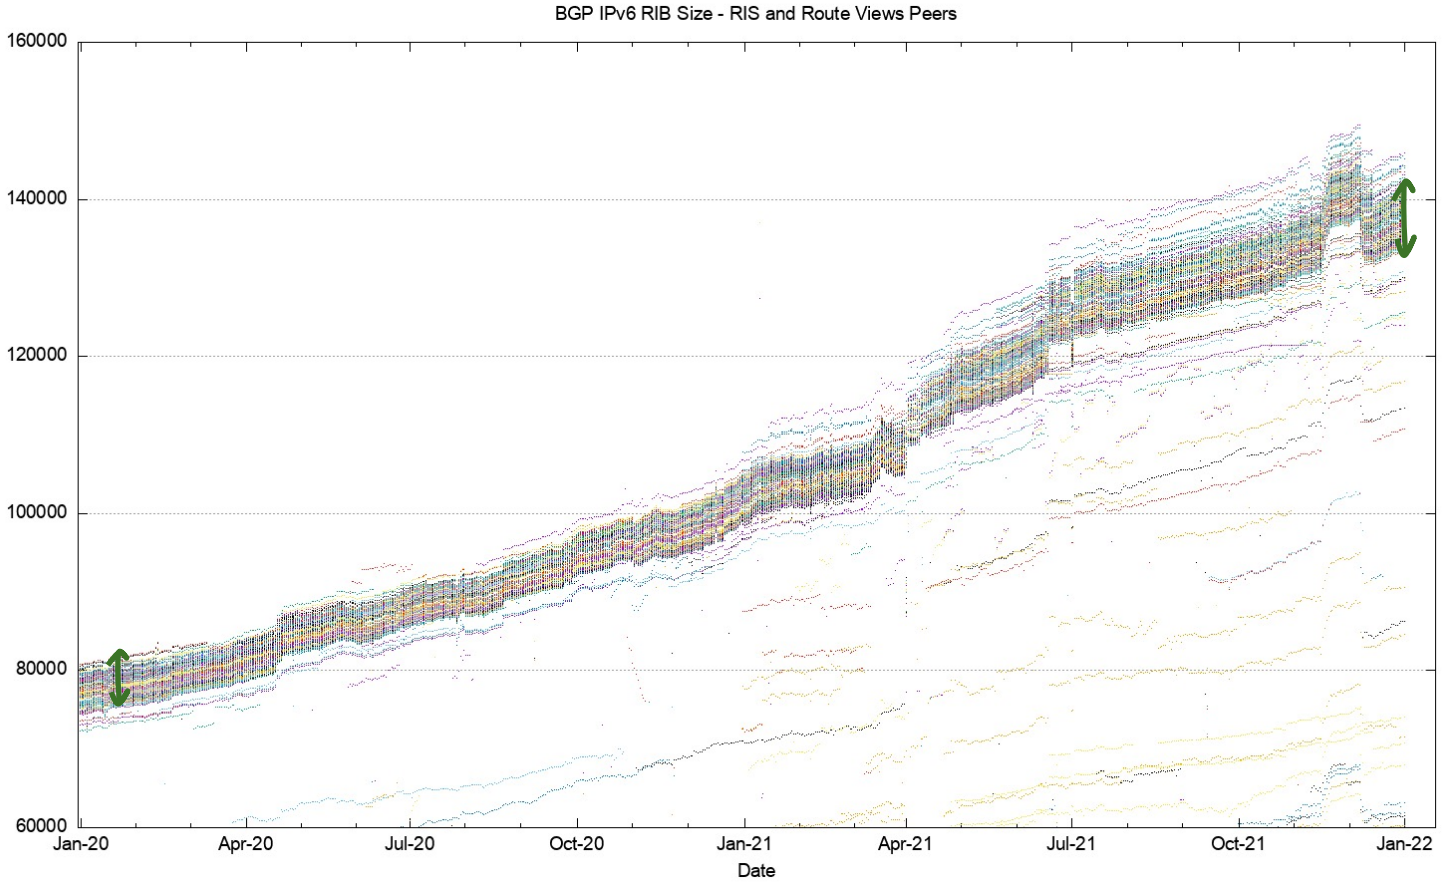

Routing Indicators for IPv6

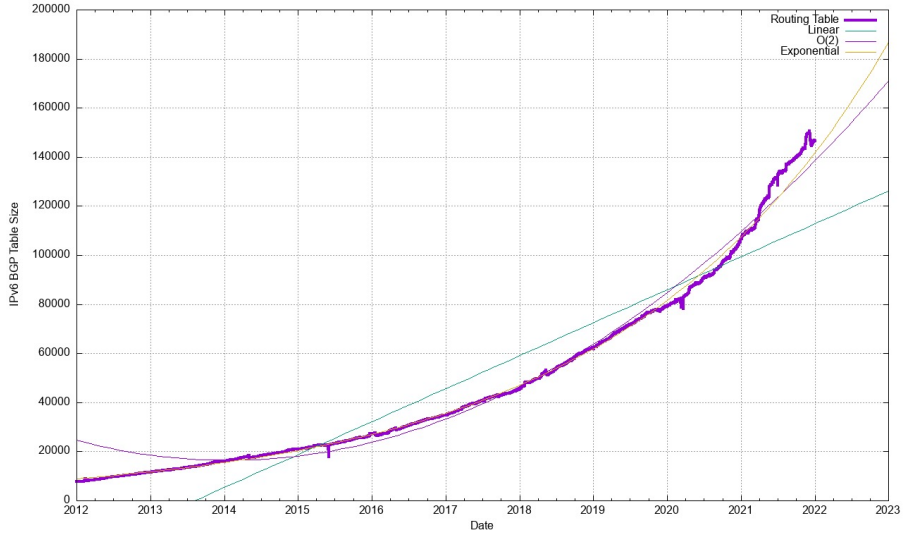

Routing prefixes – growing by some 46,000 prefixes per year

AS Numbers – growing by some 2,400 ASNs per year – with 6,740 added in 2021

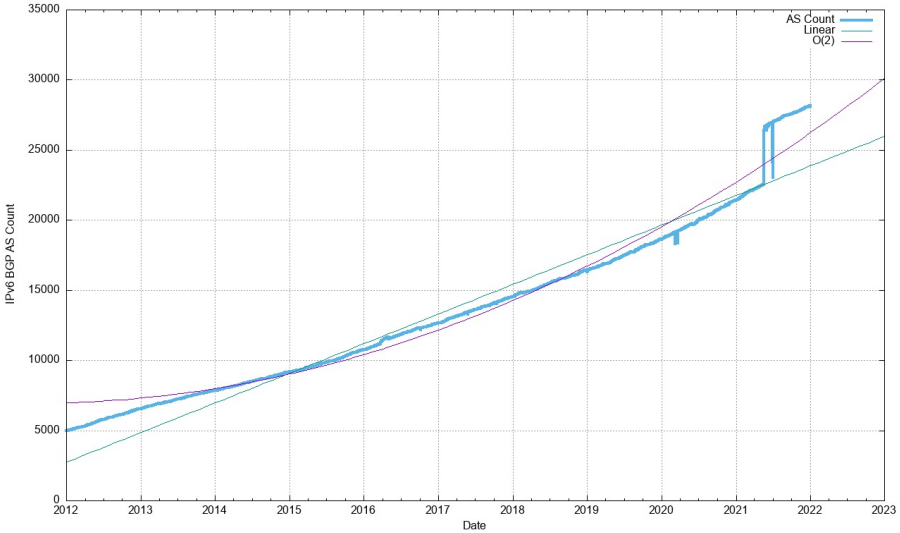

Routing Indicators for IPv6

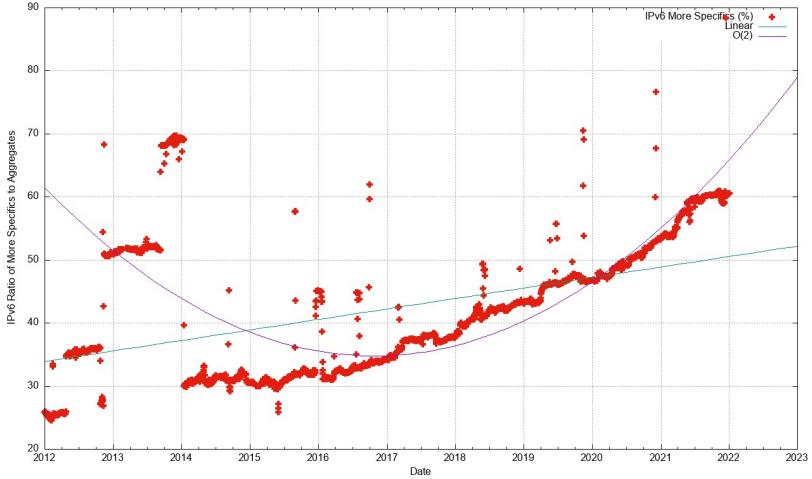

More Specifics now take up 60÷ of the routing table

The average size of a routing advertisement is getting smaller

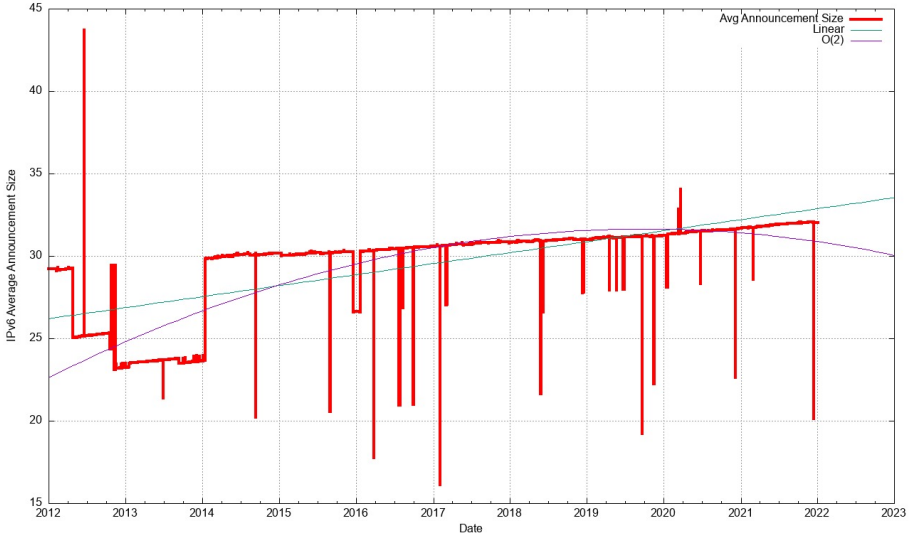

Routing Indicators for IPv6

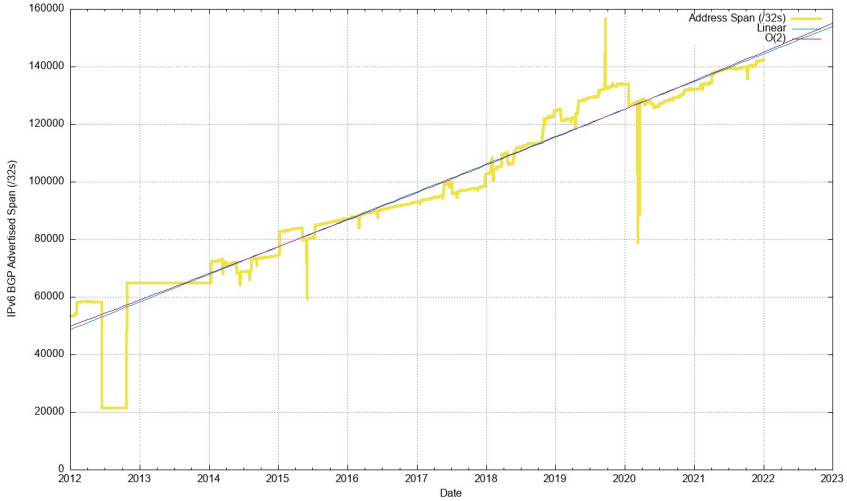

Advertised Address span is growing at an exponential rate

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

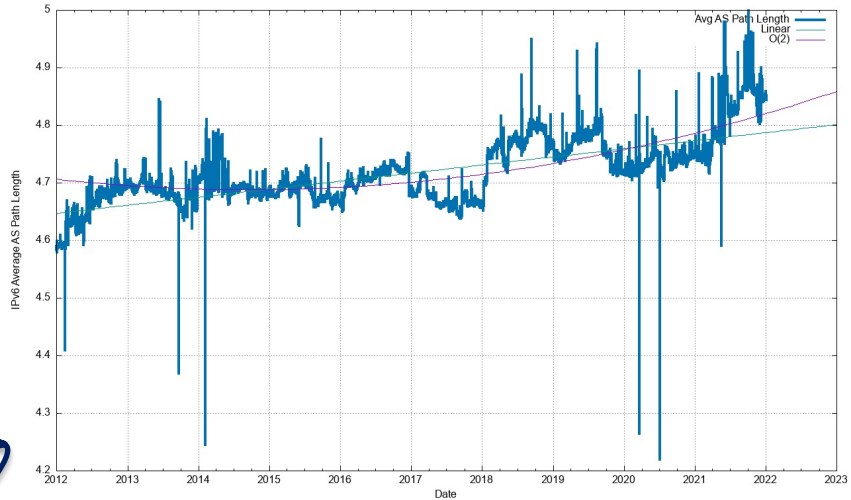

AS Adjacencies (vantage point: AS131072)

23,243 out of 28,328 ASNs have 1 or 2 AS Adjacencies (82%)
914 ASNs have 10 or more adjacencies
5 ASNs have >1,000 adjacencies

5,734	AS6939	HURRICANE - Hurricane Electric, Inc., US
3,970	AS38255	FITI, China Education and Research Network CN
1,472	AS3356	LEVEL3 - Level 3 Communications, Inc., US
1,325	AS1299	TWELVE99 Arelion, (fka Telia Carrier), SE
1,232	AS174	COGENT-174 - Cogent Communications, US

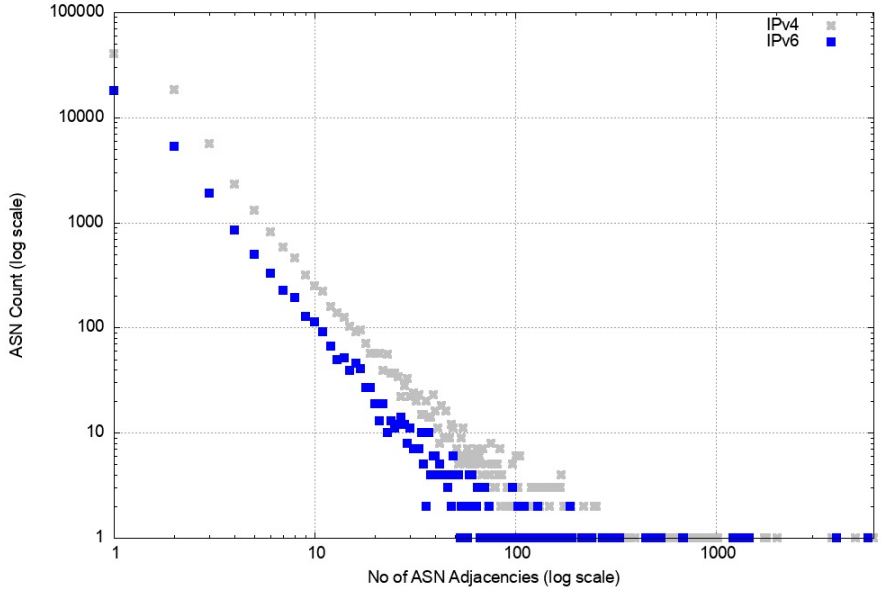

V6 in 2021

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently growing at some **41,000 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

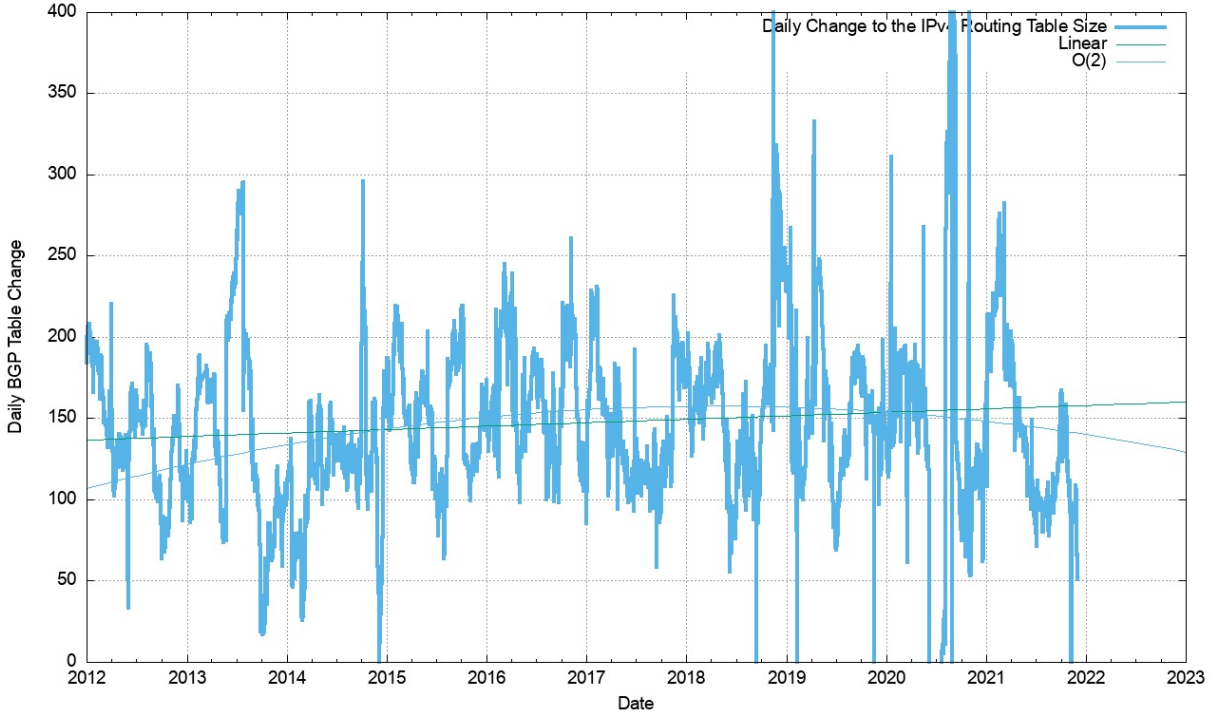

Growth in the V4 network appears to be constant at a long-term average of 150 additional routes per day, or some 54,000 additional routes per year. Recent growth rates are lower

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	908,000
2023		956,000
2024		1,004,000
2025		1,052,000
2026		1,100,000
2027		1,148,000

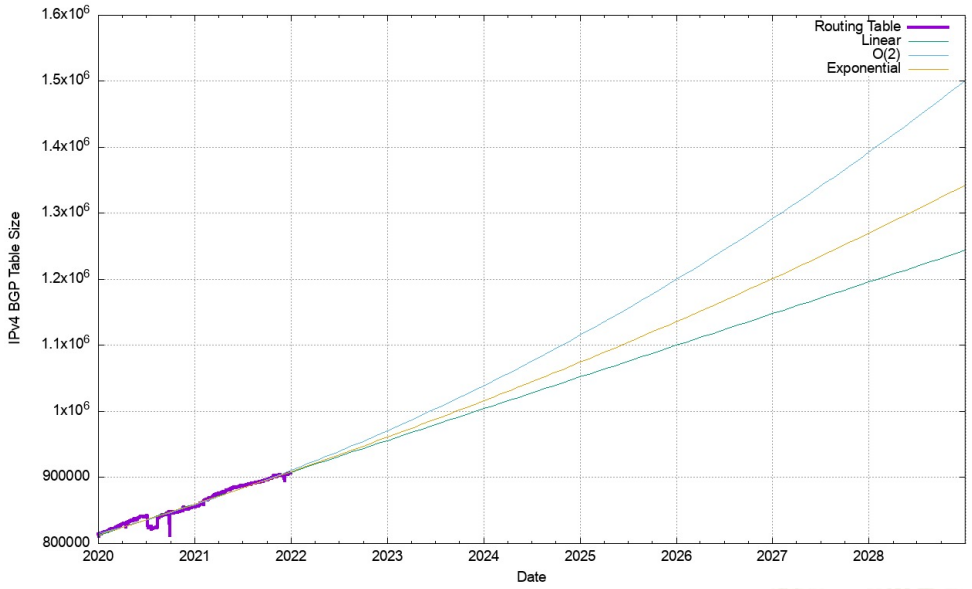

V6 - Daily Growth Rates

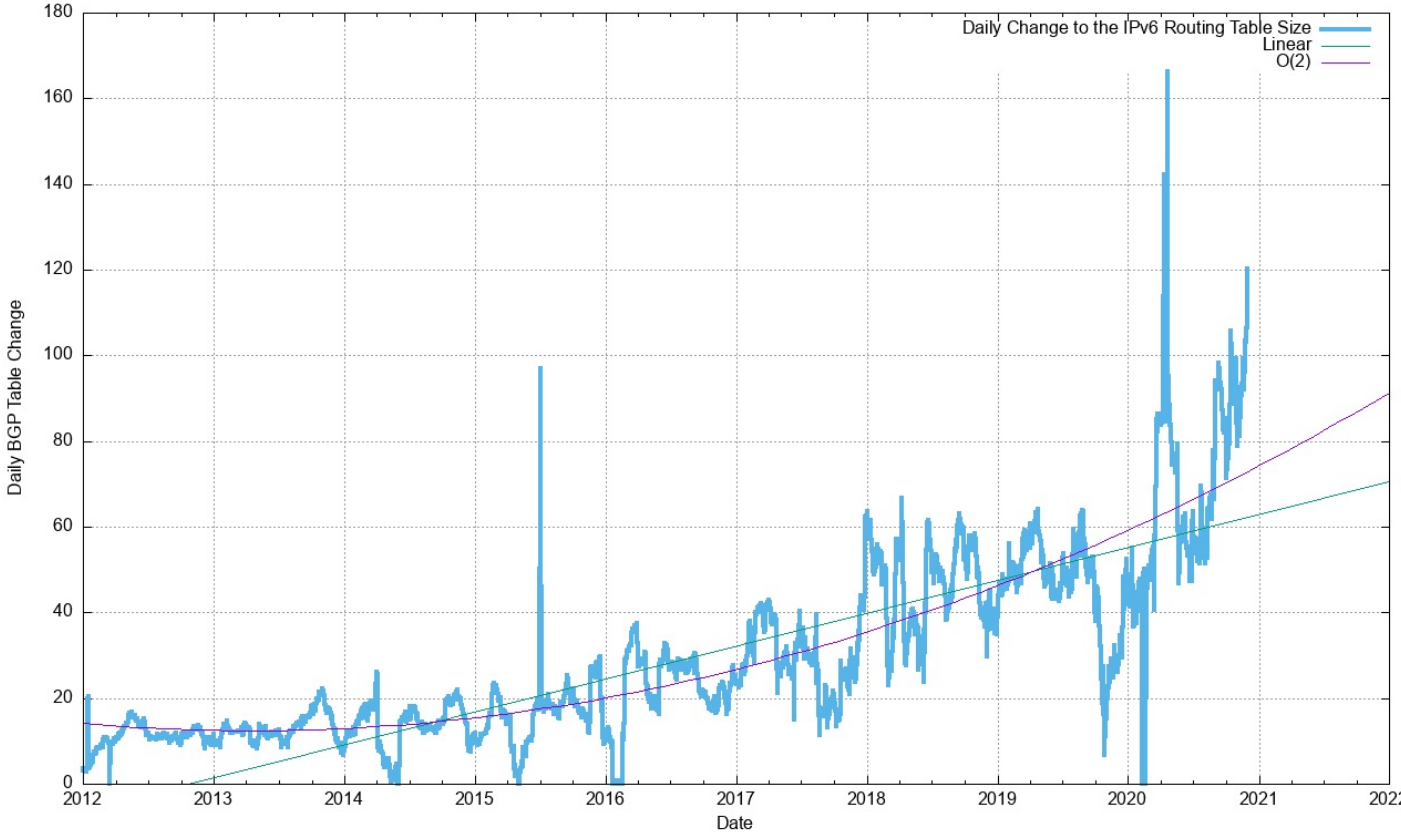

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2018*	*45,000*	
2019	*62,000*	
2020	*79,000*	
2021	*104,000*	
2022	*147,000*	*152,000*
2023	*185,000*	*214,000*
2024	*233,000*	*302,000*
2025	*261,000*	*426,000*
2026	*299,000*	*601,000*
2027	*366,000*	*847,000*

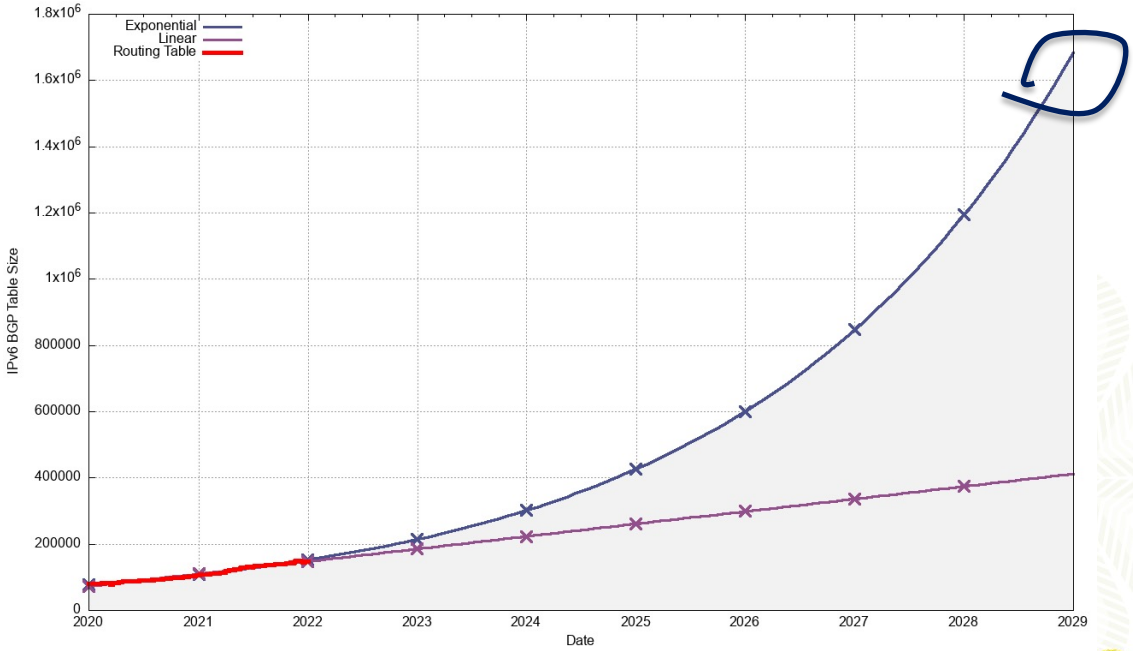

Note that the IPv6 tables are 128bits wide – i.e. 4x the size of the IPv4 tables!

BGP Table Growth

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

These two tables will require the same storage/lookup size in around 2 years time, 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

BGP Updates

- What about the rate of updates in BGP?

IPv4 BGP Updates

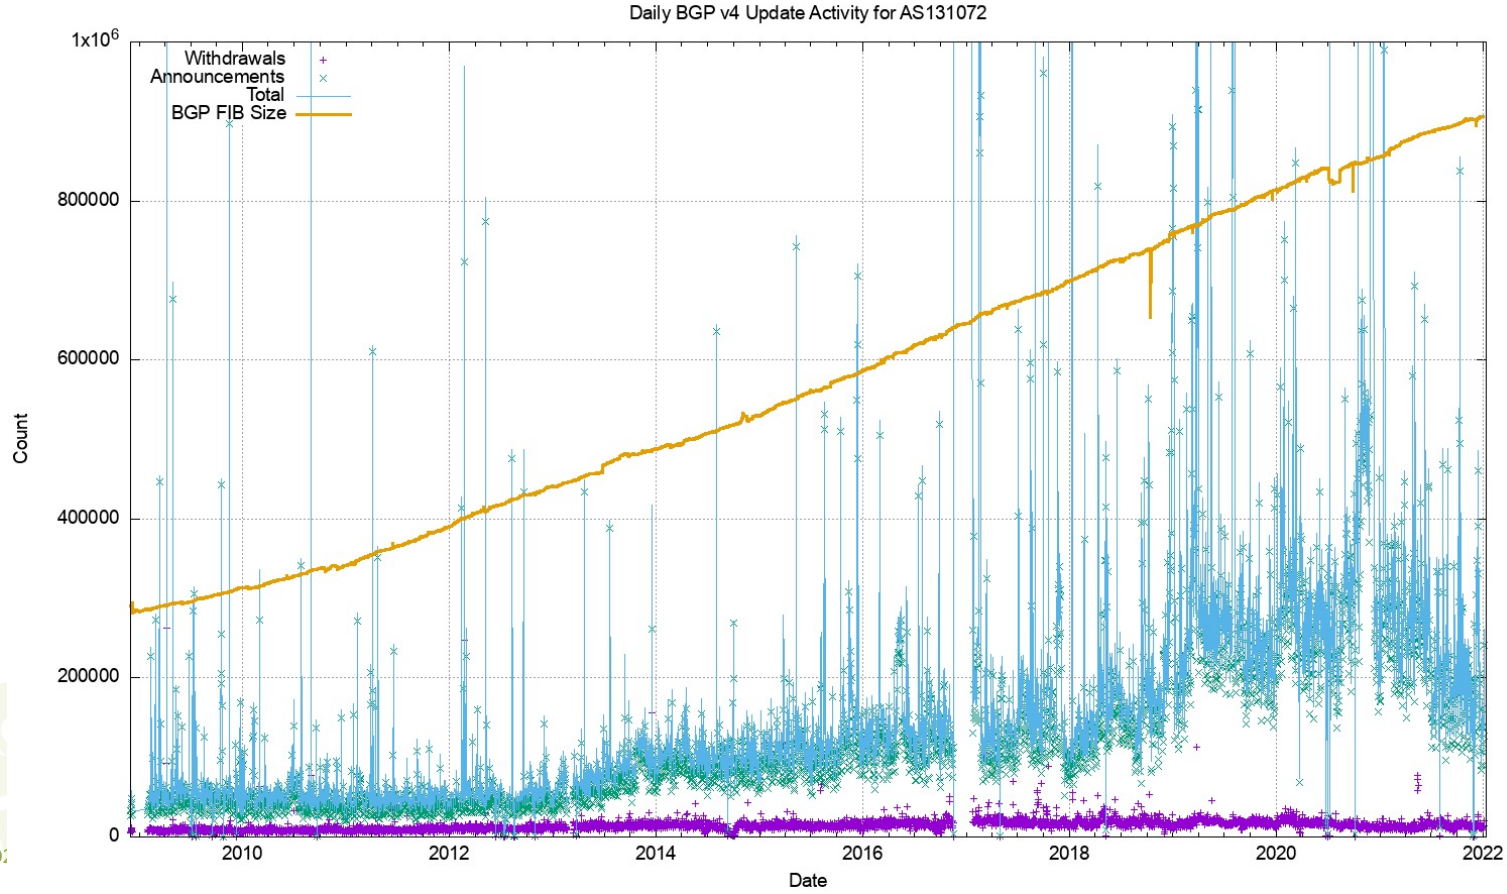

IPv4 BGP Convergence Performance

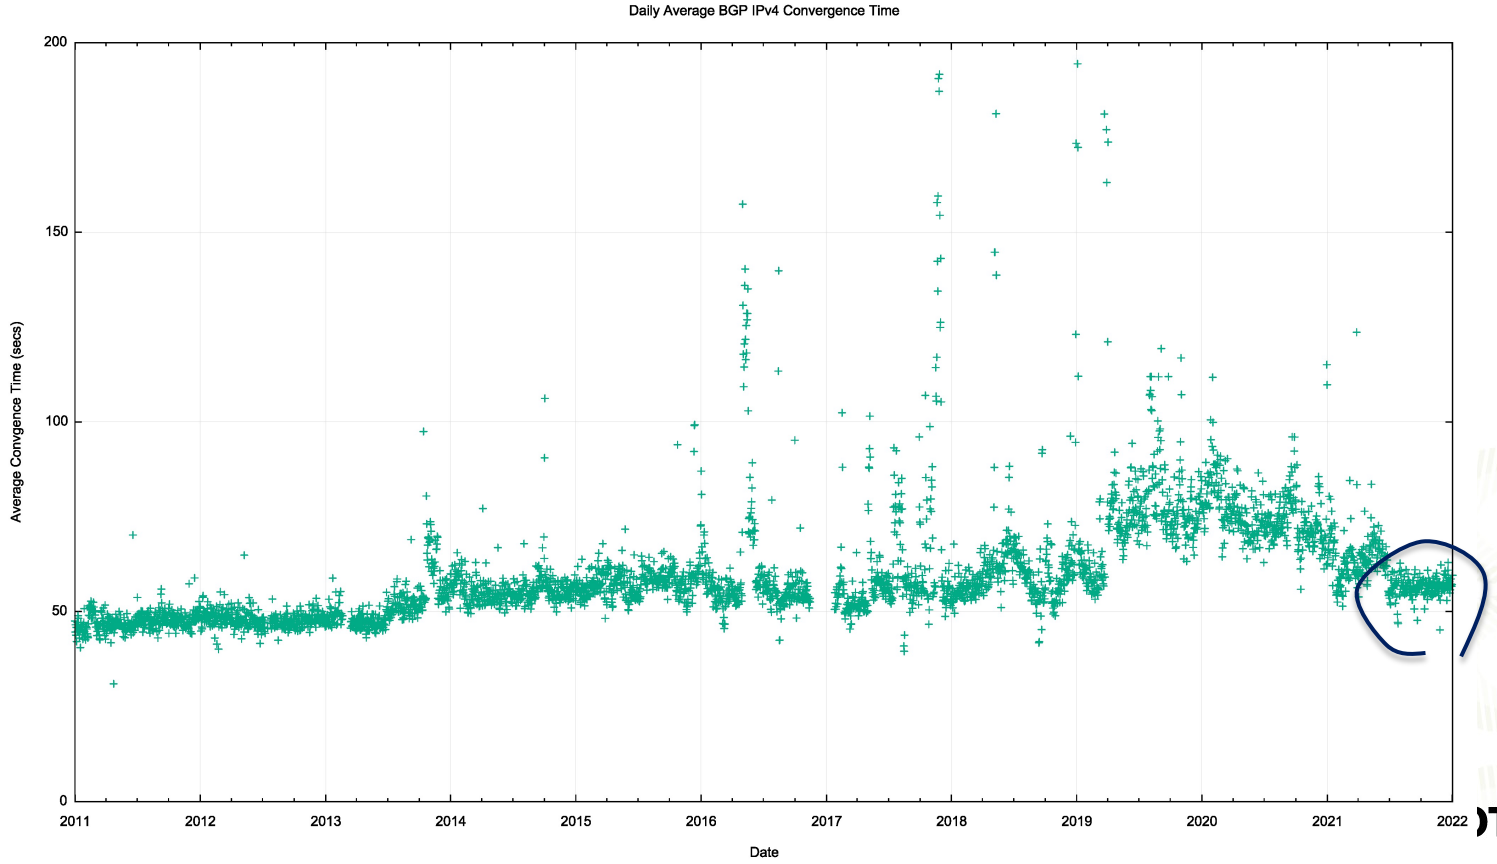

Updates in IPv4 BGP

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
- 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

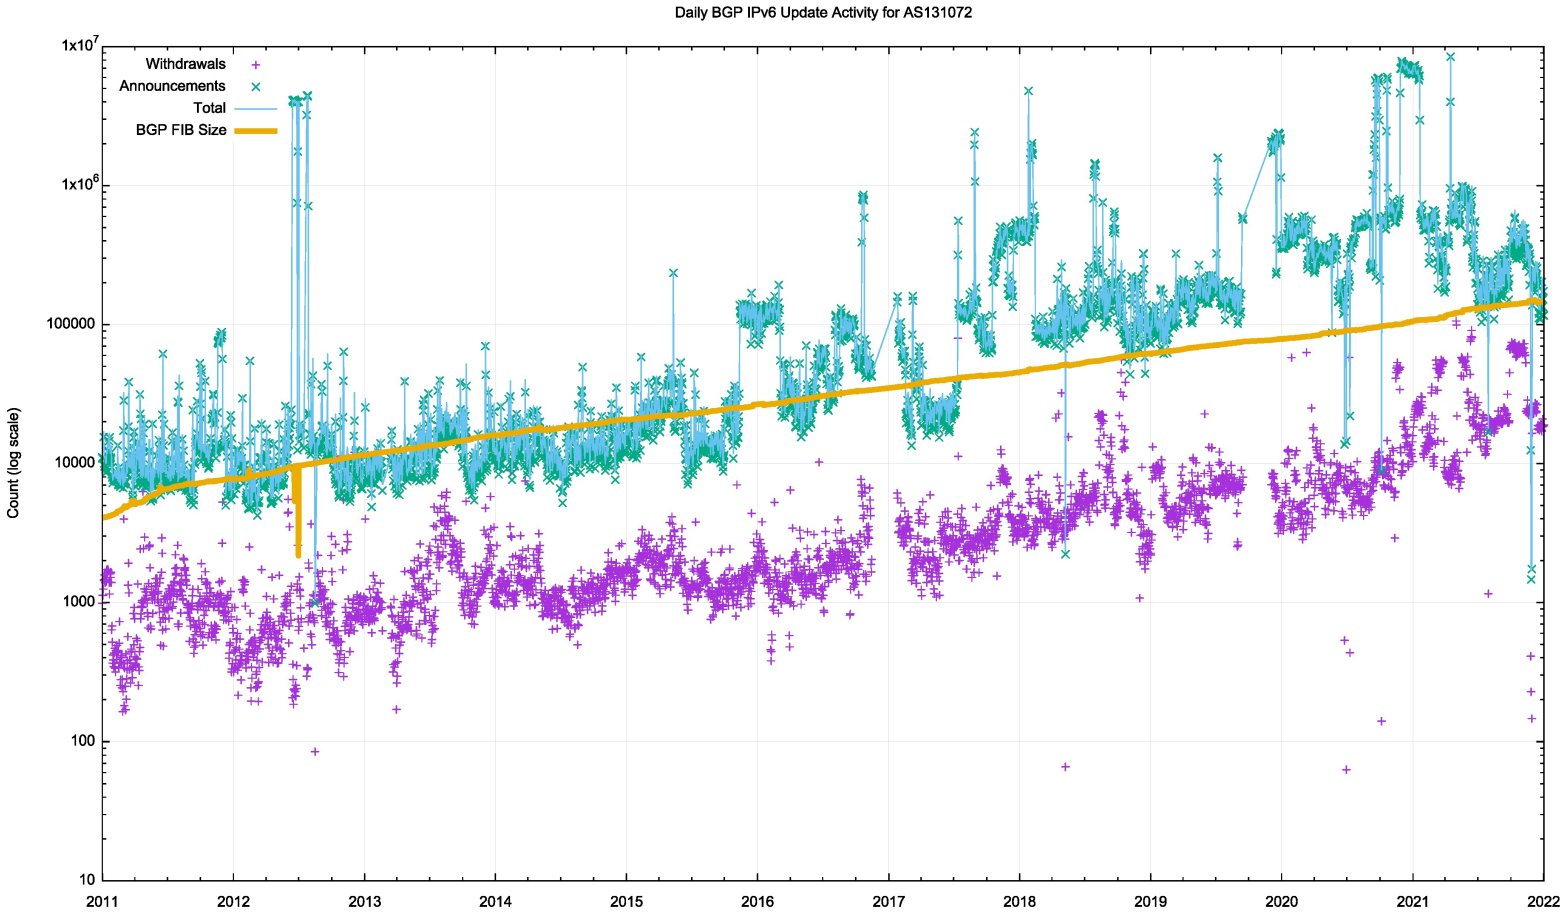

V6 Unstable Prefixes

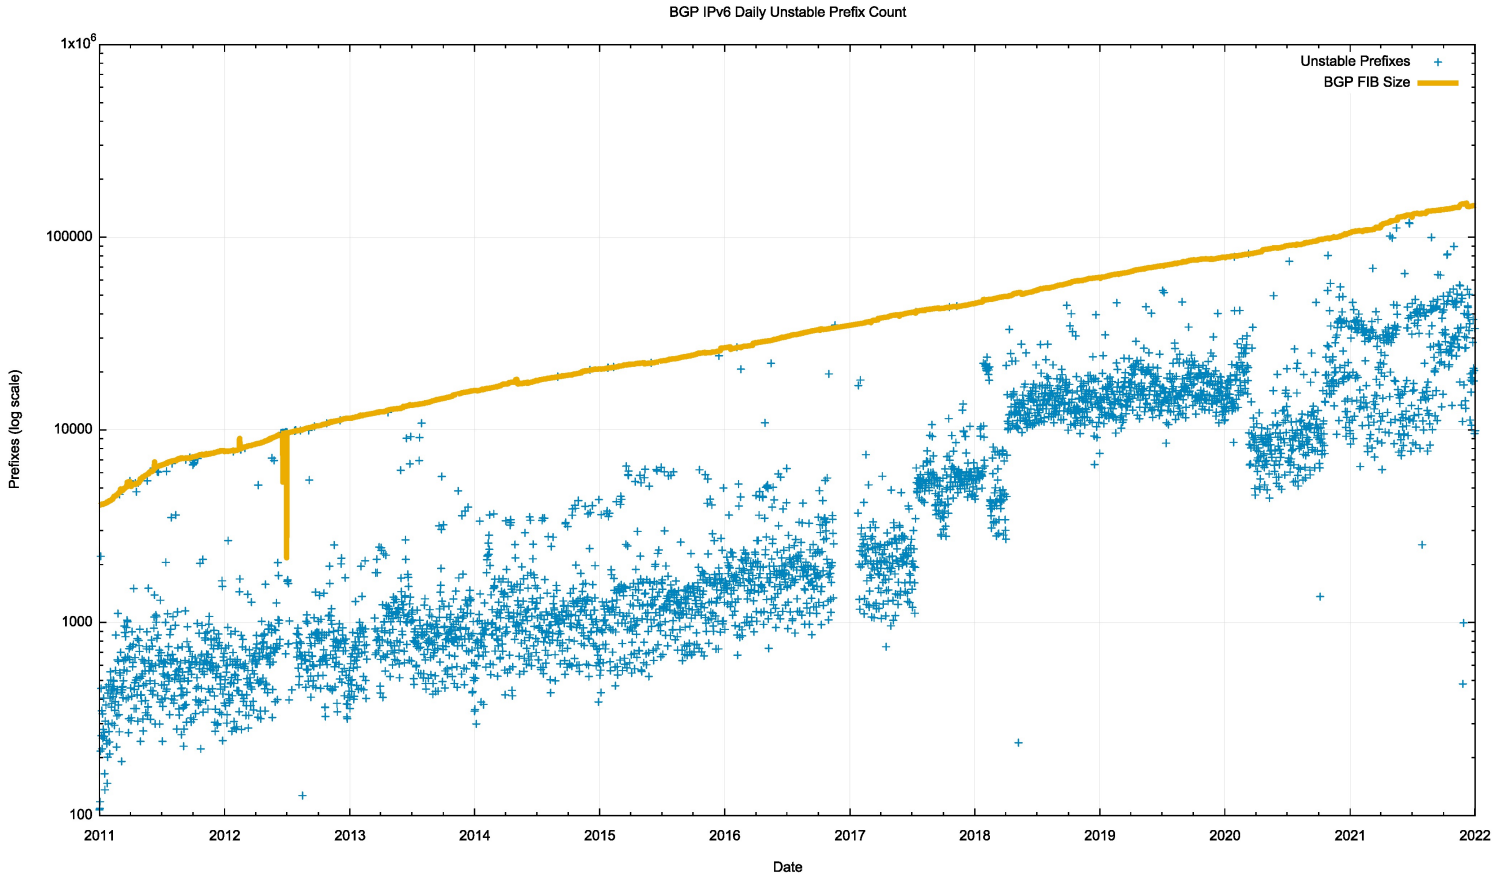

V6 Convergence Performance

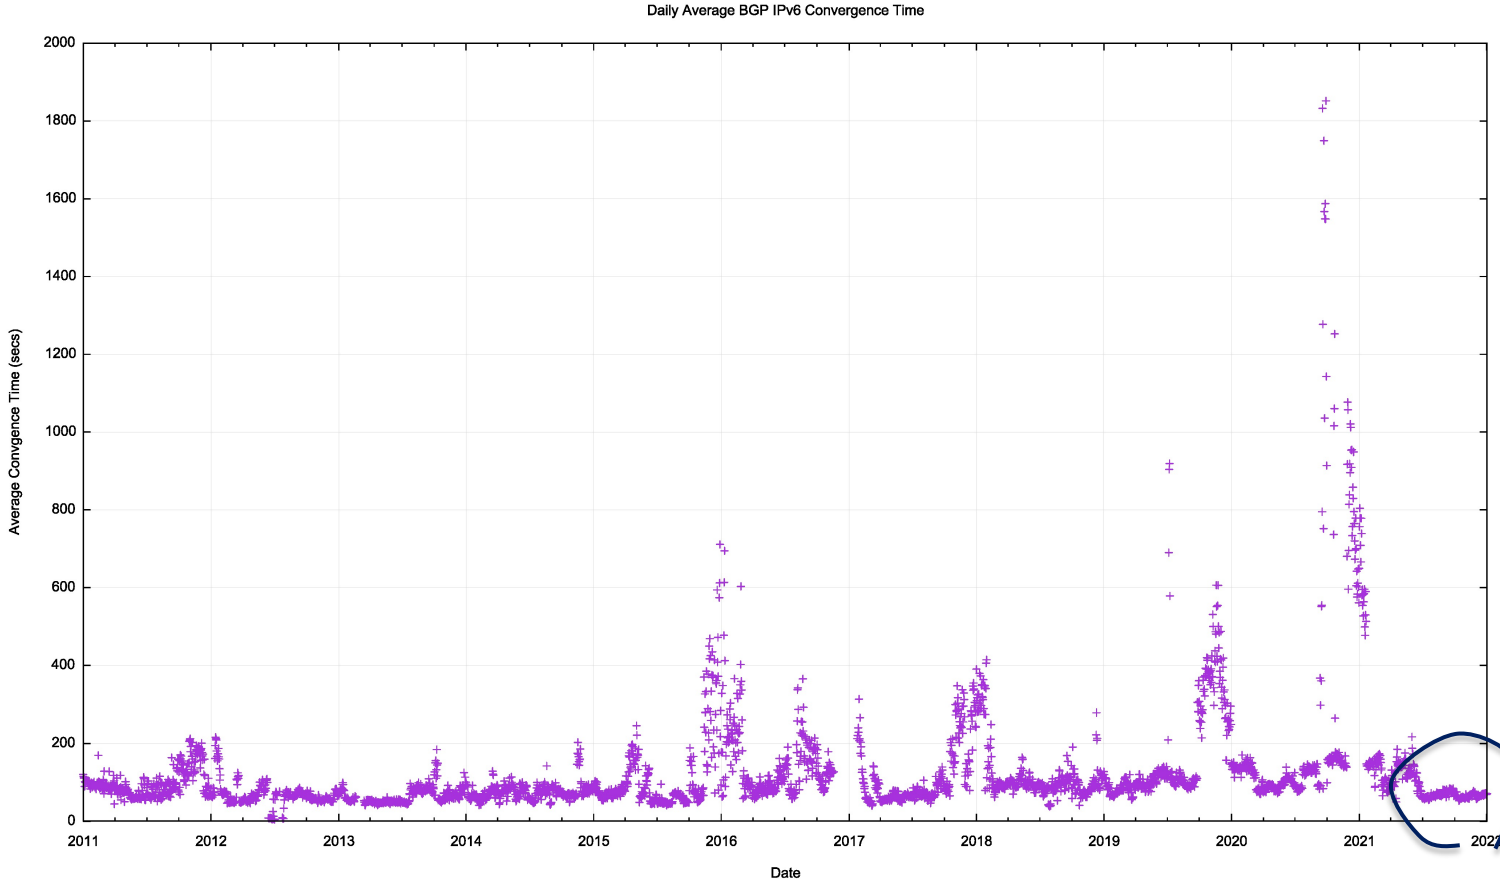

Updates in IPv6 BGP

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. The V6 network was more stable than ever in the last half of 2021
- Just 2 AS's generated 70% of the BGP update load in the last 2 weeks of 2021. Instability is still concentrated in a small number of pathological cases.

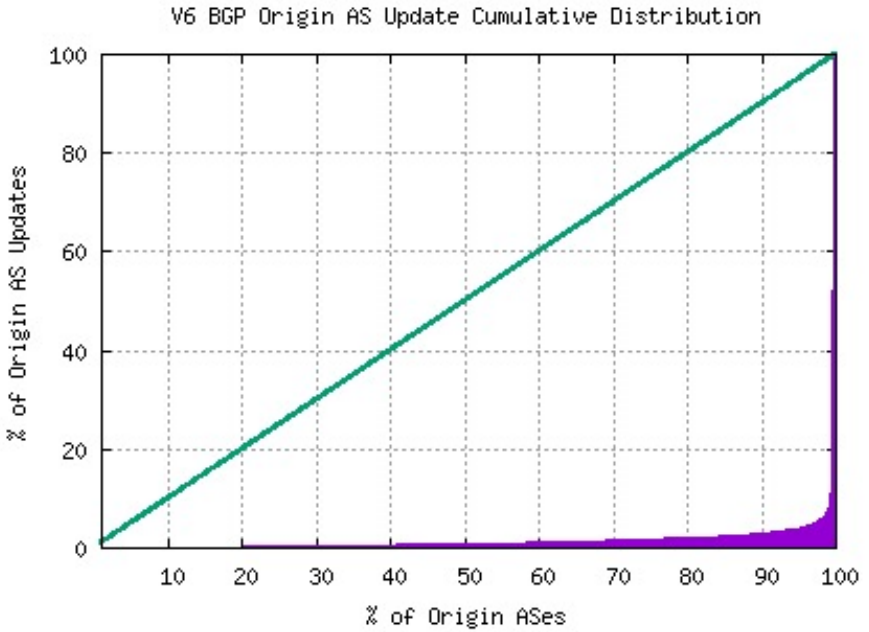

The Highlights

- 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”

- 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 233K 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 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 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?


2022
APRICOT
APNIC 53
21 February – 3 March 2022
#apricot2022