


# Comparing IPv4 and IPv6 from the perspective of BGP dynamic activity

Geoff Huston


APNIC

February 2012

# The IPv4 Table: 2004 - now




# The IPv6 Table: 2004 - now



# AS131072 – BGP Updates / day

upd - log scale



V4 - ~100K updates/day

V6 - 1K rising to 10K  
updates/day

# That's unexpected!

- Most models of routing instability view instability as a probabilistic function relating to prefixes and ASs.
  - The corollary is that the greater the number of prefixes and ASs the greater the level of routing activity to maintain a coherent network topology
- So why is IPv4's update rate constant at ~100K updates across a period that has seen the IPv4 table grow from 125K to 400K entries?
- And why is IPv6's update rate growing at a similar rate to the size of the IPv6 table?

# Prefixes vs Updates

- The number of updates generated by a distance vector protocol increases with the distance between the “root cause” and the listening position
- So instead of looking at the number of updates, lets look instead at the number of prefixes that are associated with a changed routing state each day

# AS131072 – Unstable Prefixes /day

pxfs - log scale



V4 - ~20K prefixes

V6 - 100 rising to 1K prefixes

There are two curious aspects of this data:

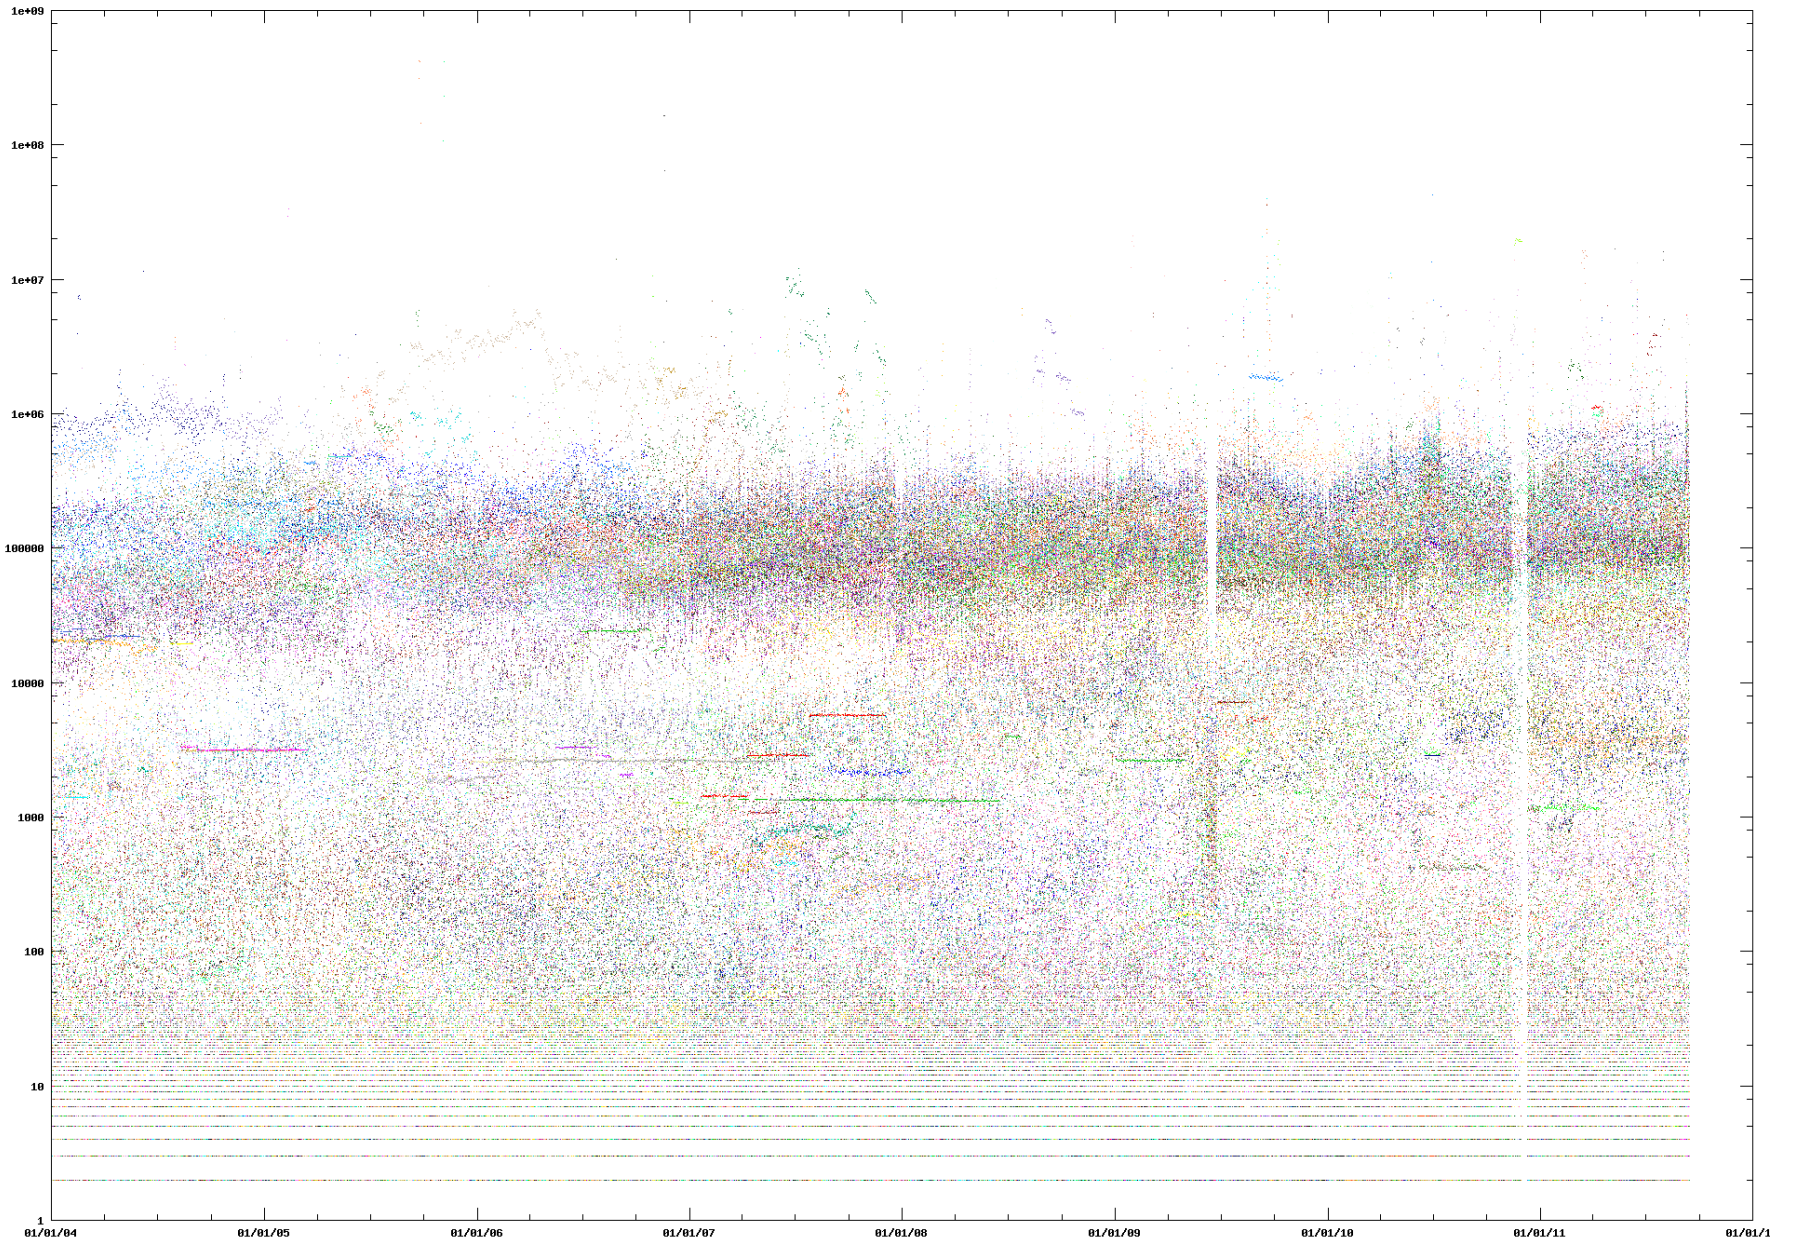
- 1 - Is routing IPv4 really “scale free”?
- 2 - Why is IPv6 different?



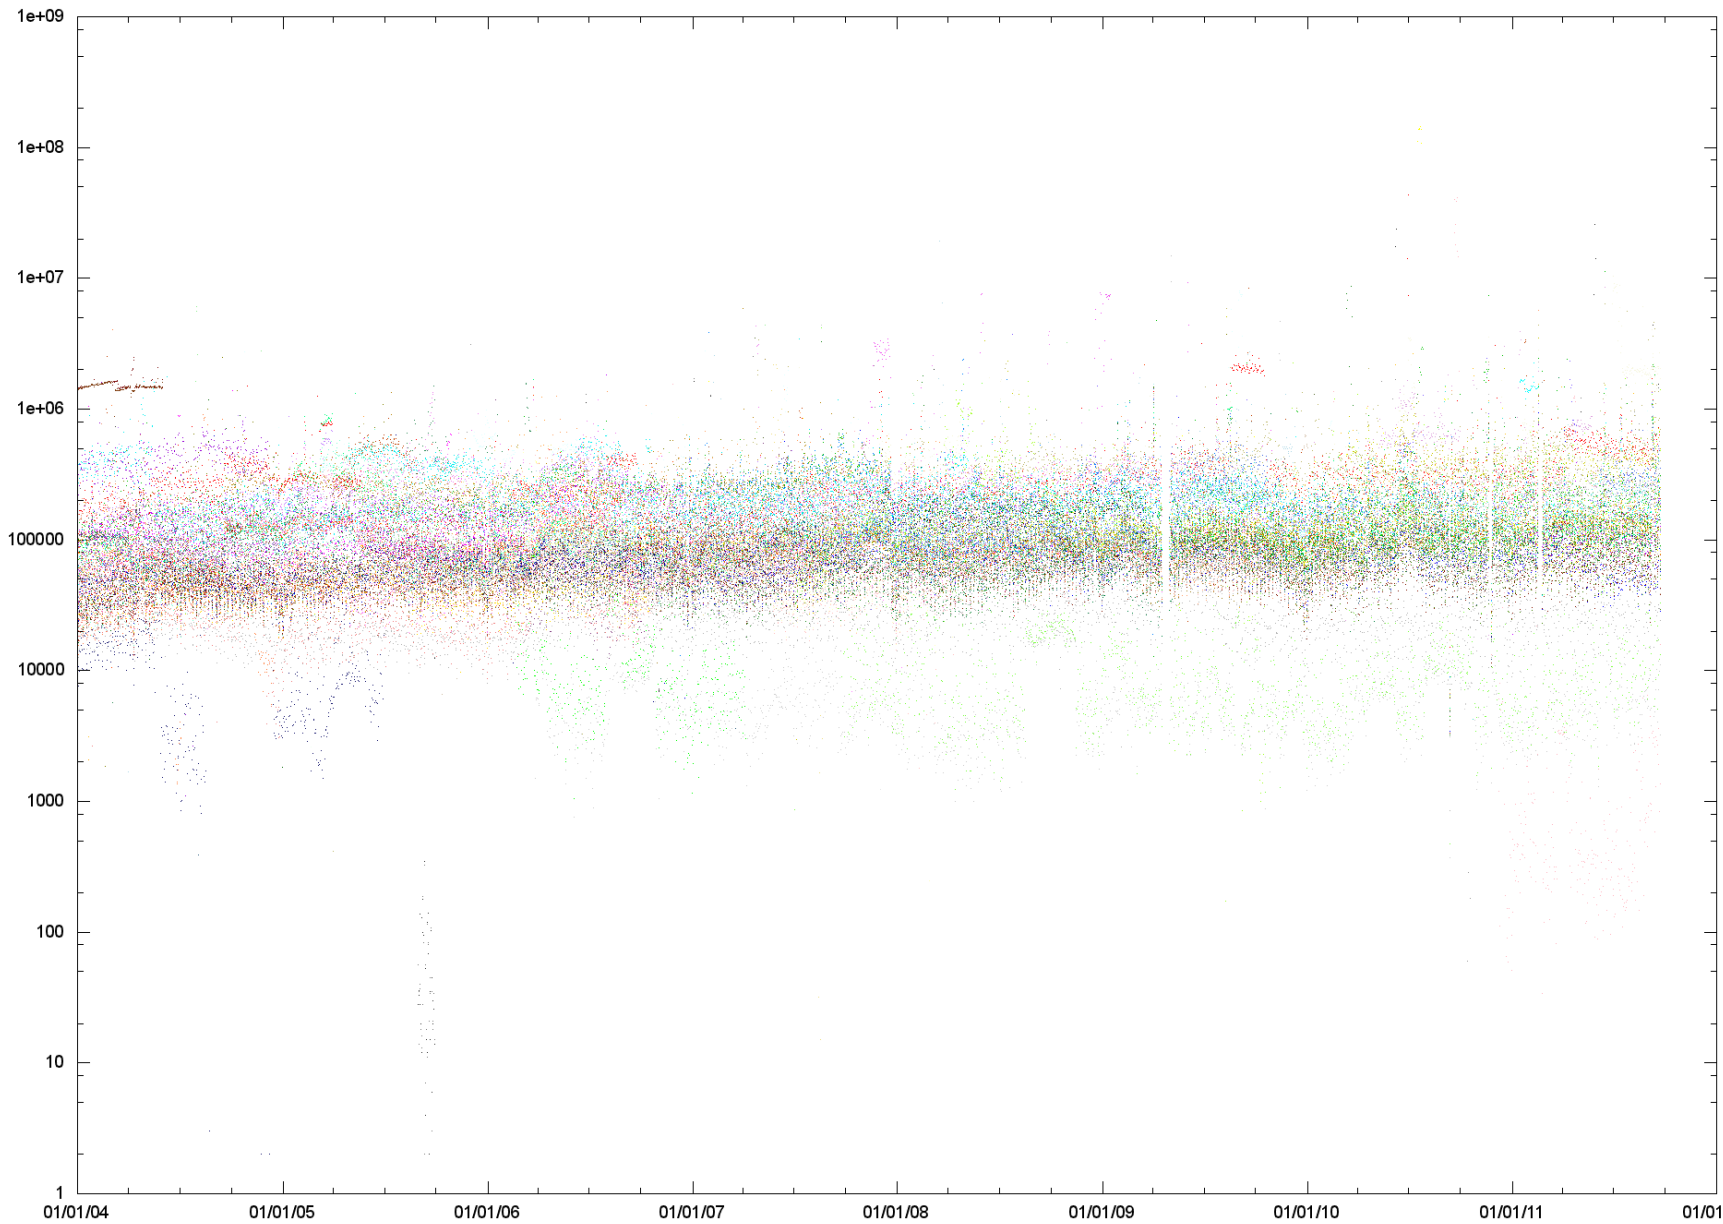
# 1 - Is IPv4 BGP flat-lining?

Or is it just some strange anomaly in AS131072?

– So I looked at the RIS and Route-View data sets



Daily Update rate for RIS peers – 2004 to 2011



Daily Update rate for Route Views peers – 2004 to 2011

# Its not me!

- Its not completely flat
- But it's everywhere
  - The long term growth rate of the dynamic activity of eBGP in IPv4 is far lower than the growth rate of the default-free routing table.

# BGP Updates

- There are two components to BGP update activity:
  1. Convergence updates as BGP searches for a new stable “solution”
  2. The update relating to the “primary” event
- In an ever expanding network both BGP update components should be rising
  - But the total number of updates is not rising in IPv4

# Convergence

- BGP is a distance vector protocol
- This implies that BGP may send a number of updates in a tight “cluster” before converging to the “best” path
- This is clearly evident in withdrawals and convergence to (longer) secondary paths

# For Example

Withdrawal at source at 08:00:00 03-Apr of 84.205.77.0/24 at MSK-IX, as observed at AS 2.0

Announced AS Path: <4777 2497 9002 12654>

Received update sequence:

08:02:22 03-Apr + <4777 2516 3549 3327 12976 20483 31323 12654>


08:02:51 03-Apr + <4777 2497 3549 3327 12976 20483 39792 8359 12654>

08:03:52 03-Apr + <4777 2516 3549 3327 12976 20483 39792 6939 16150 8359 12654>

08:04:28 03-Apr + <4777 2516 1239 3549 3327 12976 20483 39792 6939 16150 8359 12654>


08:04:52 03-Apr - <4777 2516 1239 3549 3327 12976 20483 39792 6939 16150 8359 12654>

1 withdrawal at source generated a convergence sequence of 5 events, spanning 150 seconds



Average Convergence Time for RIS peers – 2004 to 2011






Average Convergence Time for Route Views peers – 2004 to 2011

# IPv4 Convergence Behaviour

- IPv4 BGP convergence times and average convergence update counts have been constant for the past 7 years
- This implies that a critical aspect of the network's topology has also been held constant over the same period

# AS Path is Constant

- The AS Path Length has been constant for many years, implying that the convergence effort has also remained constant




Per-peer average  
AS Path Length as  
Measured by Route-Views  
Peers, 1998 - 2011

# BGP Updates

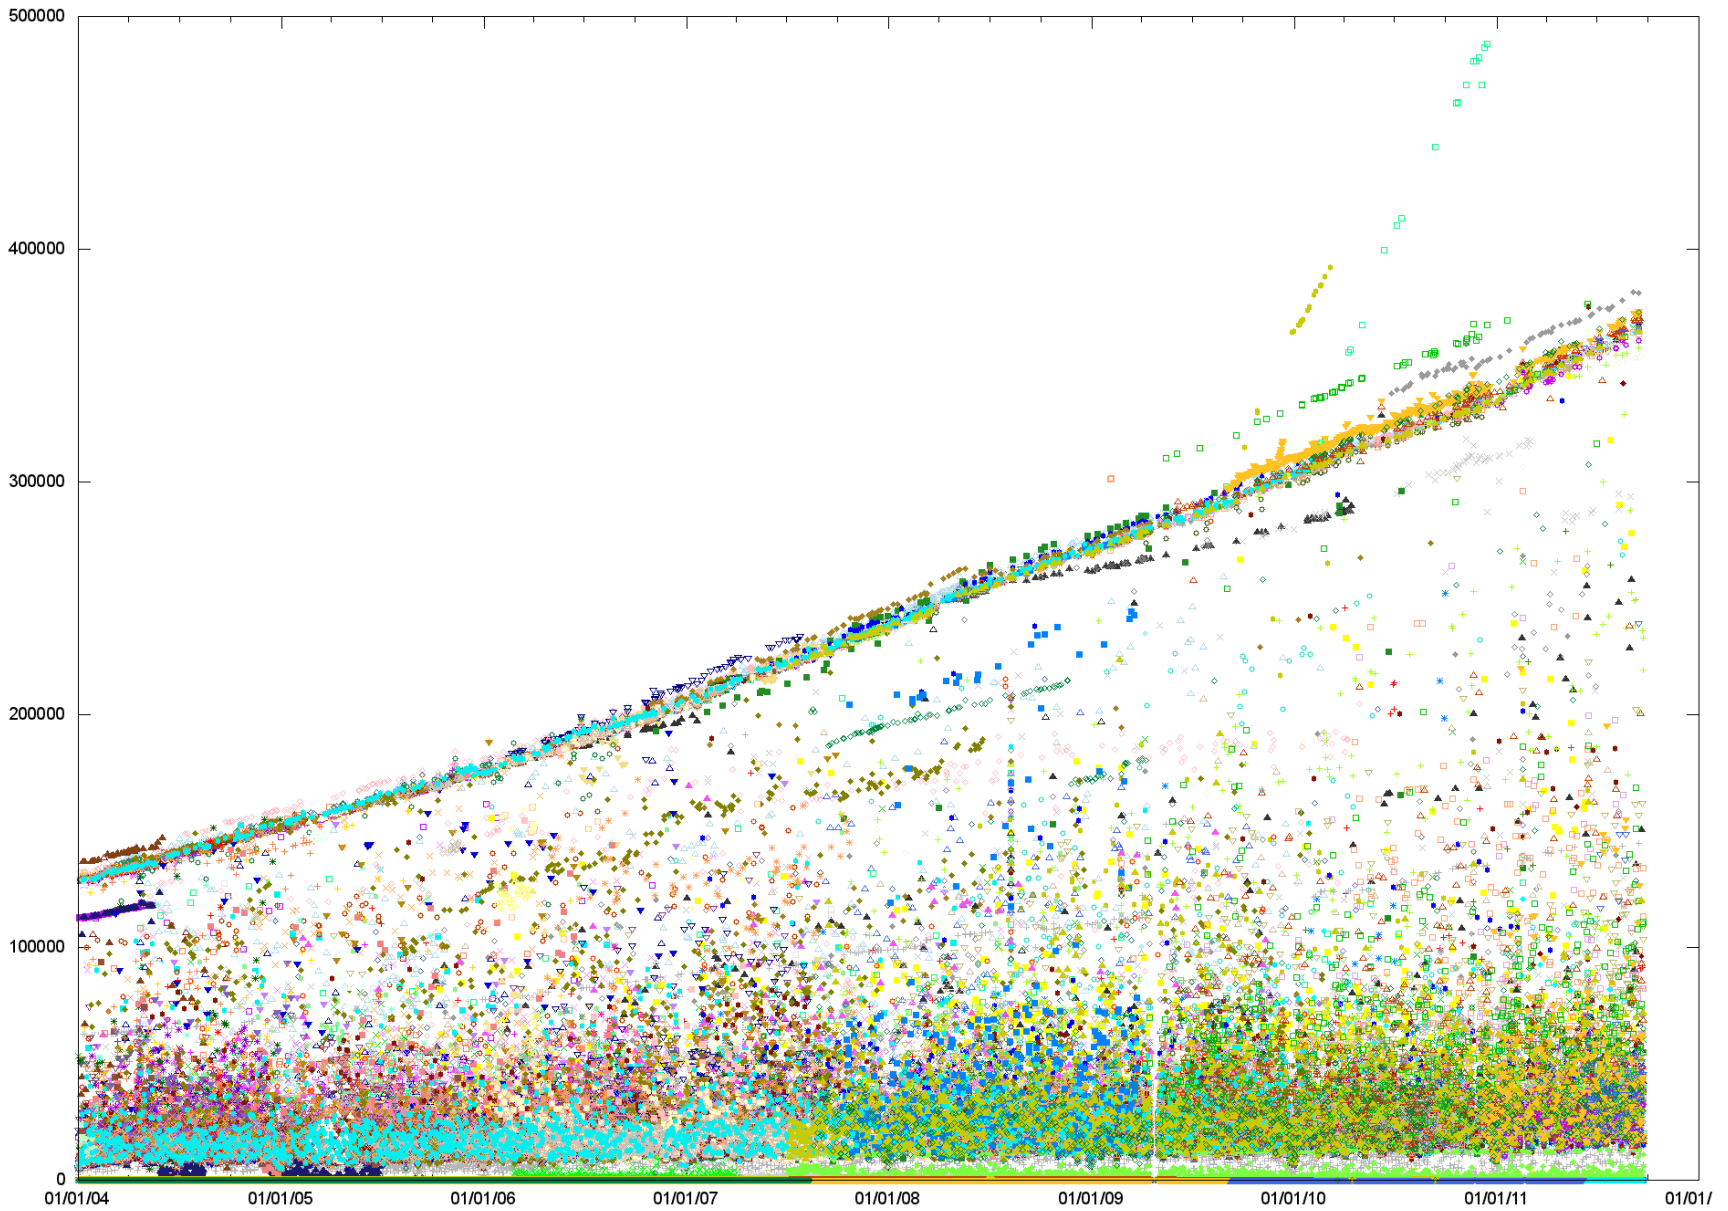
- There are two components to BGP update activity:
  1. Convergence updates as BGP searches for a new stable “solution”
    - AS Path lengths have been steady as the Internet grows by increasing the density of interconnection, not by increasing average AS Path length
  2. The update relating to the “primary” event

# Unstable Prefixes

- Are we seeing the same prefixes exhibiting instability multiple times per day, or different prefixes?
- What's the profile of instability from the perspective of individual prefixes?



Unstable Prefixes for RIS peers – 2004 to 2011



Unstable prefixes for Route Views peers – 2004 to 2011


# Unstable Prefixes

- Over the past 4 years the number of unstable prefixes lies between 20,000 – 50,000 prefixes per day
- How “stable” is this set of unstable prefixes?
  - Are they the same prefixes?
  - Are they equally noisy?
  - What are the characteristics of this “noise”?



# Prefix Instability Duration

Prefix Instability Distribution



# Prefix Instability

- Prefix Instability is generally short lived
  - 90% of all prefixes are unstable for 2 days or less
  - 6 prefixes are persistently unstable – these are beacon prefixes.
- The distribution of the duration of prefix instability at a coarse level (per day) appears to be a power law distribution (see Zipfs' Law)

# The Flat World of IPv4

- The average number of convergence events appears to be basically flat for the past couple of years
  - The growth rate appears to be far lower than the growth rate of the routing table itself
- The number of unstable prefixes per day is also relatively long term constant
  - but the individual prefixes themselves are unstable for 1 – 2 days on average


# Why is BGP in IPv4 so Flat?

- The convergence amplification factor is governed by the bounded diameter of the Internet
- But why hasn't the number of unstable prefixes grown in line with the growth in the table size? What is limiting this behaviour of the routing system? Why 20-50K unstable prefixes per day? Why not 100K? Or 5K? What is bounding this observed behaviour?

## 2 – Why is IPv6 Different?

- Now lets return to the comparison of IPv4 and IPv6...


# AS131072 – Average time to converge (secs)



Since mid 2009  
AS131072 has been  
seeing dramatically  
higher average  
convergence times in  
IPv6

# AS131072 – average number of updates to converge

cv g2u - log scale



Since mid 2009  
AS131072 has been  
seeing the average  
number of updates per  
instability event rise to  
2x – 10x the IPv4 rates

# AS131072 Observations


The IPv4 network appears to exhibit scale-free properties:  
while the number of advertised entries has triple from 125K to 400K, the number of updates, the number of unstable routes and the time to converge are all stable over the period

The IPv6 network appears to have scaling properties:  
approx 10% of announced prefixes are unstable each day, and the time to converge is getting longer


Lets look at the RIPE NCC's RIS data set, and the Route Views archive data set to see what other AS's have observed over this period



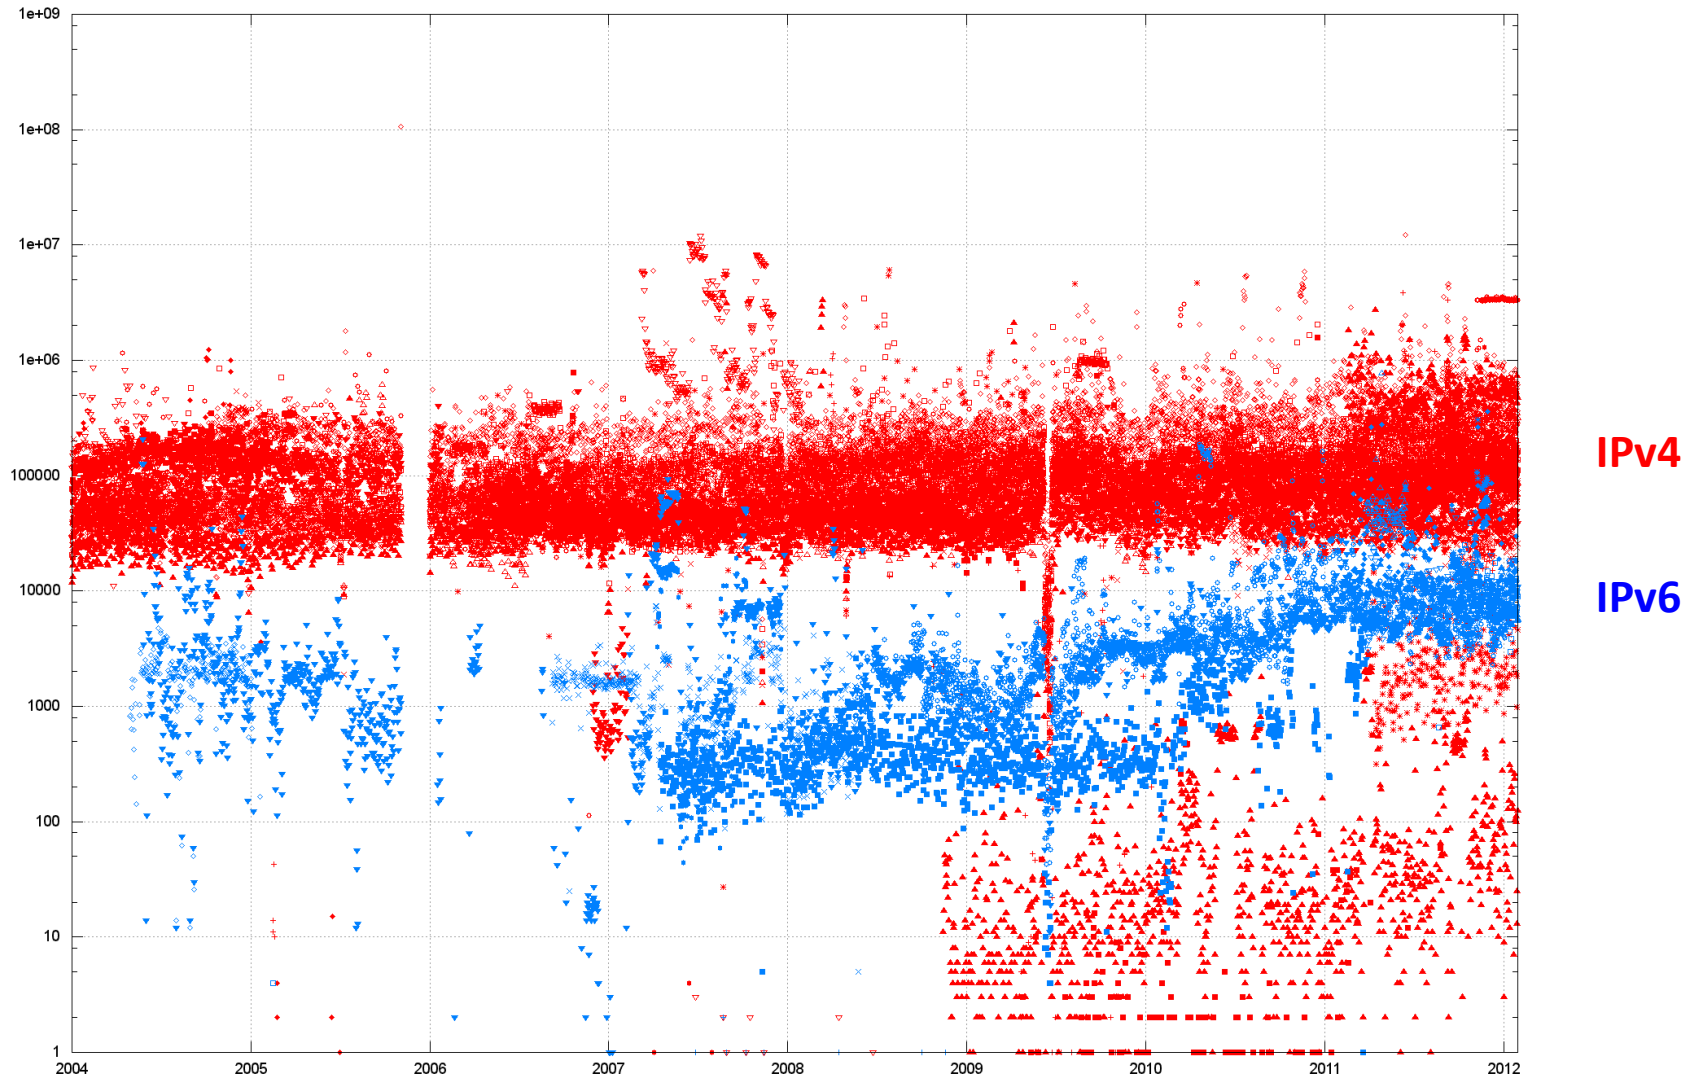
# BGP IPv4 Updates / Day – RIS




# BGP IPv6 Updates / Day – RIS



# BGP Updates / Day – RIS




# BGP Updates / Day – Route Views




# It's still not just me

- A possible explanation was that AS131072 was seeing amplified IPv6 routing instability due to instability in its IPv6 routing
- But we see the same behaviour across the larger set of IPv6 BGP peers
- What about convergence behaviours in Ipv6?


# Average Convergence Time (RIS)



# Average Convergence Time (R-V)




# Average Convergence Updates (RIS)






# Average Convergence Updates (R-V)



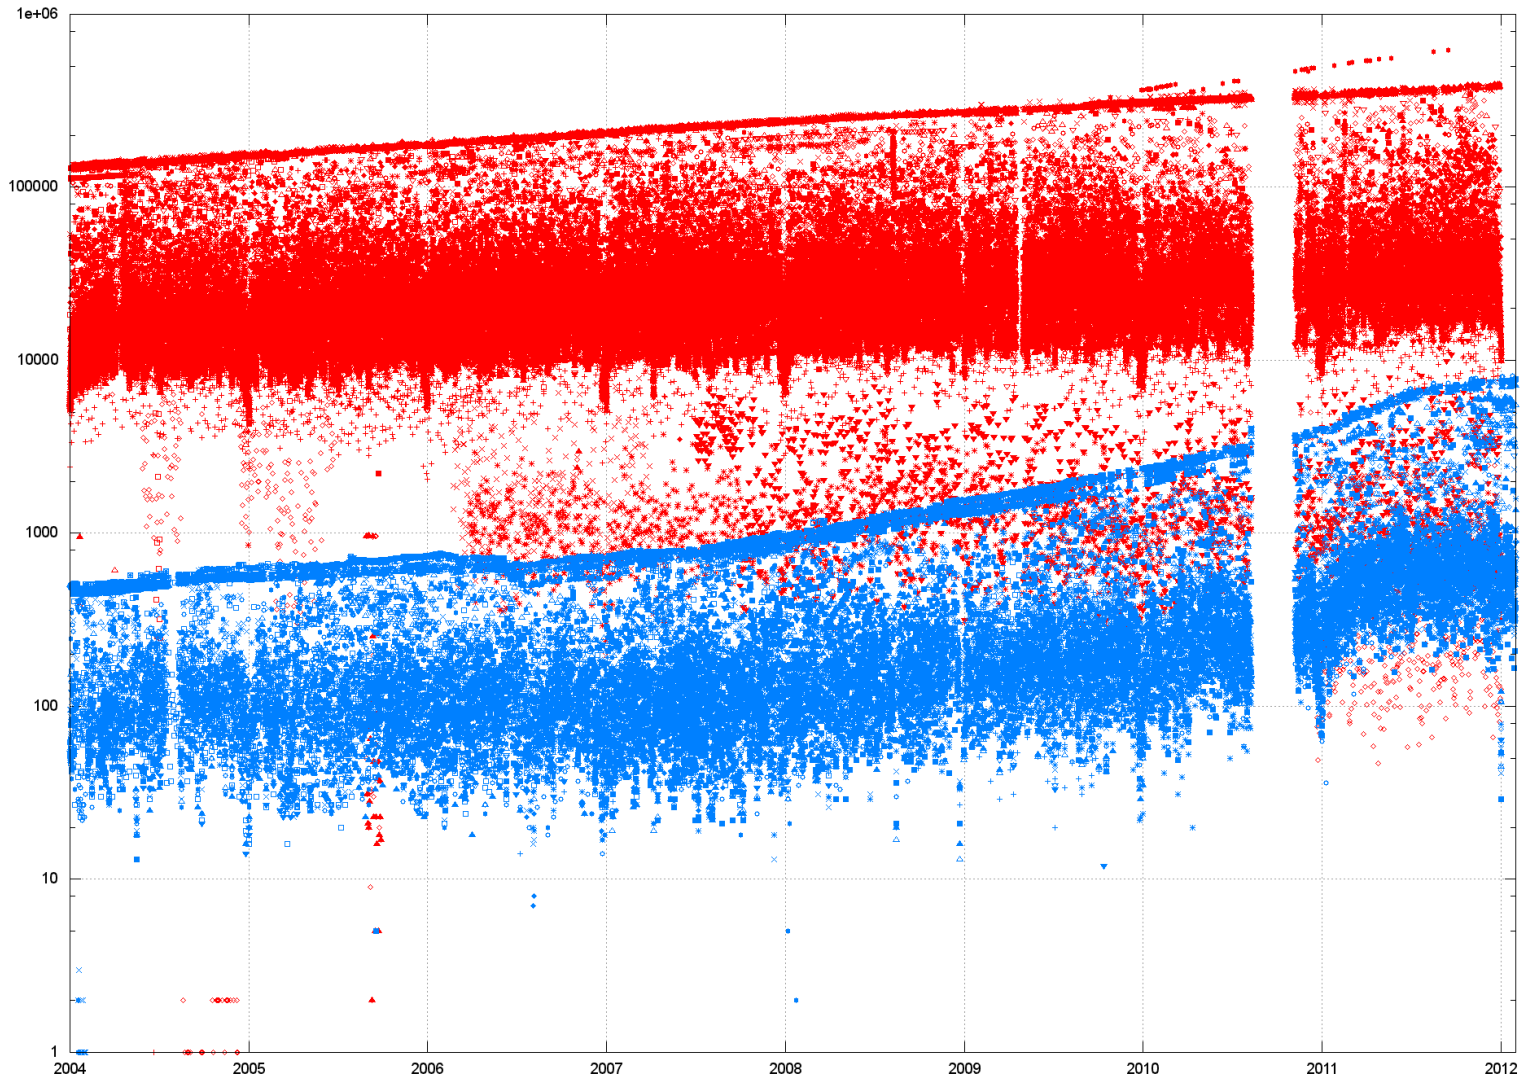
# IPv6 Instability

- It appears that the level of instability in Ipv6 rose significantly for many IPv6 peers in mid 2009, and is only coming back down in recent months
- Did the number of unstable prefixes rise at the same time?

# Unstable Prefixes (RIS)



# Unstable Prefixes (R-V)



# Some observations

The per-AS views of the IPv6 network are more heterogeneous than the IPv4 network. Different AS's see significantly different levels of instability and convergence behaviour in IPv6

The macroscopic properties of instability also differ. IPv6 prefix instability is running at 10% of prefixes. Over the same period network instability in IPv4 has dropped from 10% to 5% of advertised prefixes

# And some questions

- Are instability events in IPv4 and IPv6 linked, or are the dynamic routing behaviours in the two protocols entirely distinct?
- Are there co-incident IPv4 and IPv6 instability events with the same origin AS or AS peering?
- To what extent do IPv6 transit tunnels affect the stability of the IPv6 network? Is the higher variance of per-AS views attributable to the relative use of tunneled transit routes?
- What other factors would drive up relative route instability in IPv6?

Your questions?