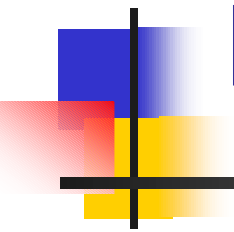

BGP'01

An Examination of the Internet's BGP Table Behaviour in 2001

Geoff Huston

Telstra

2001 - The Prediction

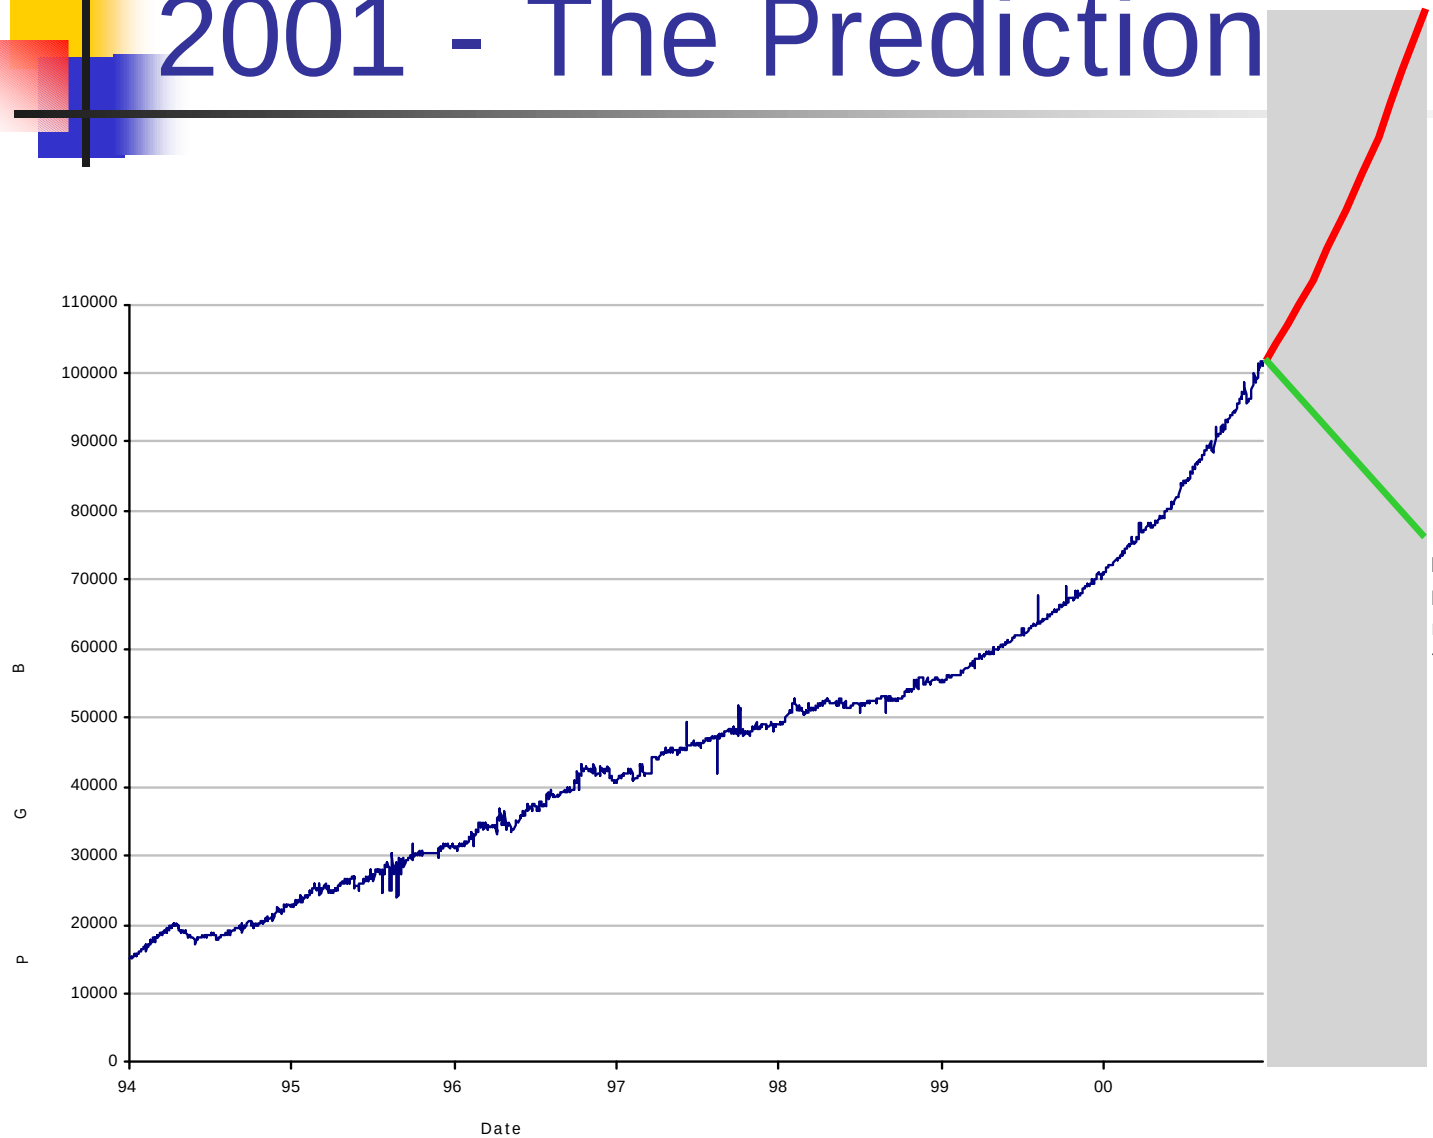

Worst Case
Continued Exponential Growth
150,000 entries by January 2002

Best Case
Elimination of all extraneous routing entries
75,000 entries by January 2002

2001 - What Happened

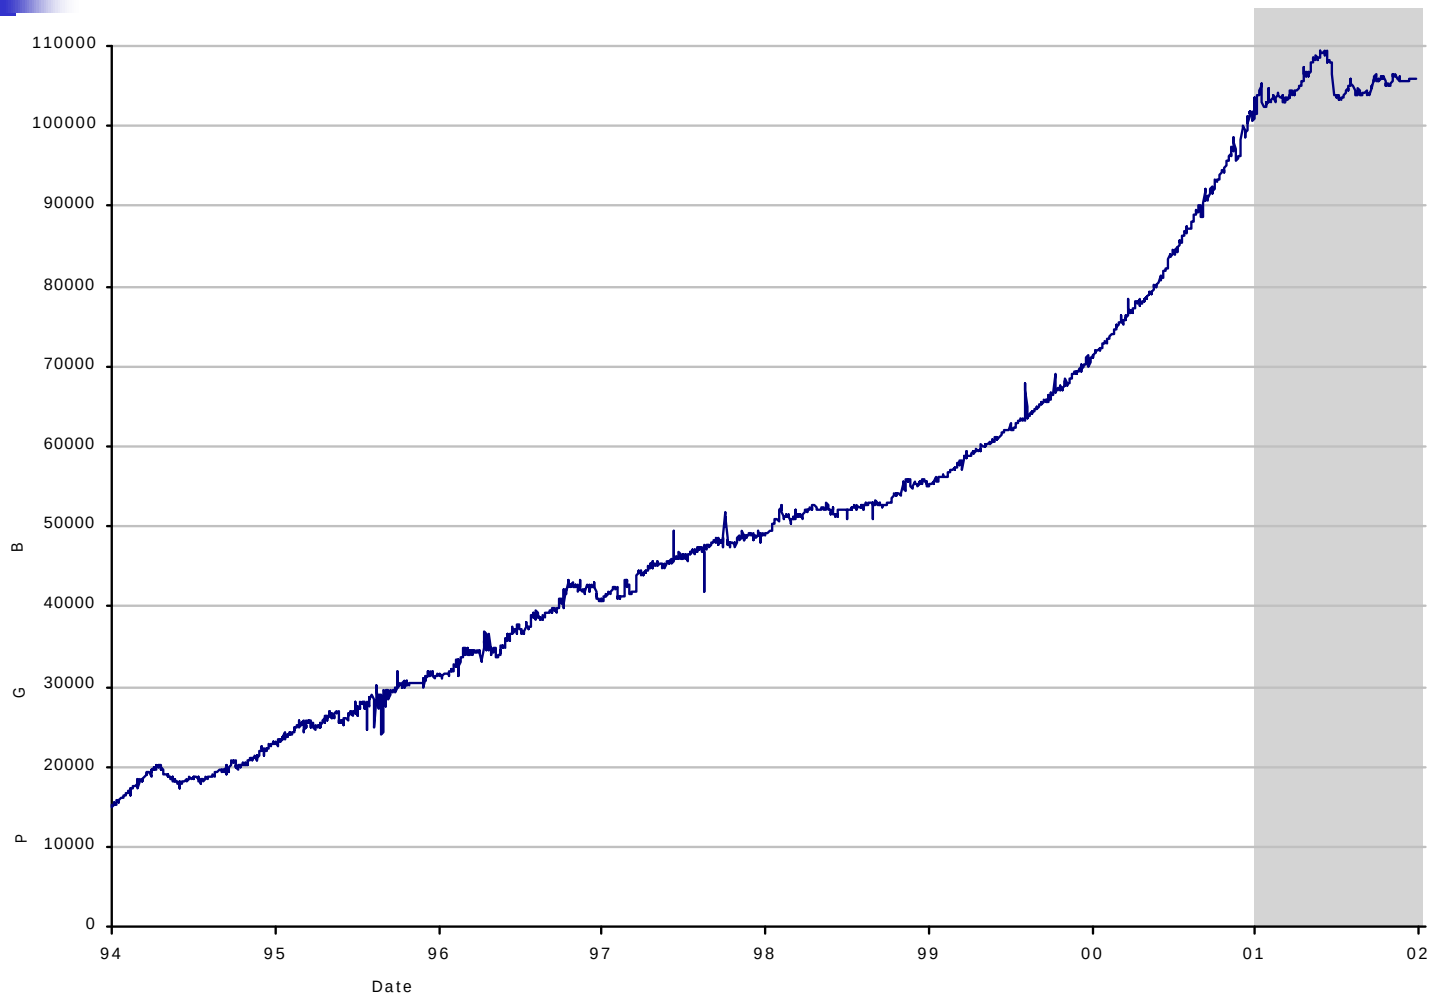

2001 - Route Views' View

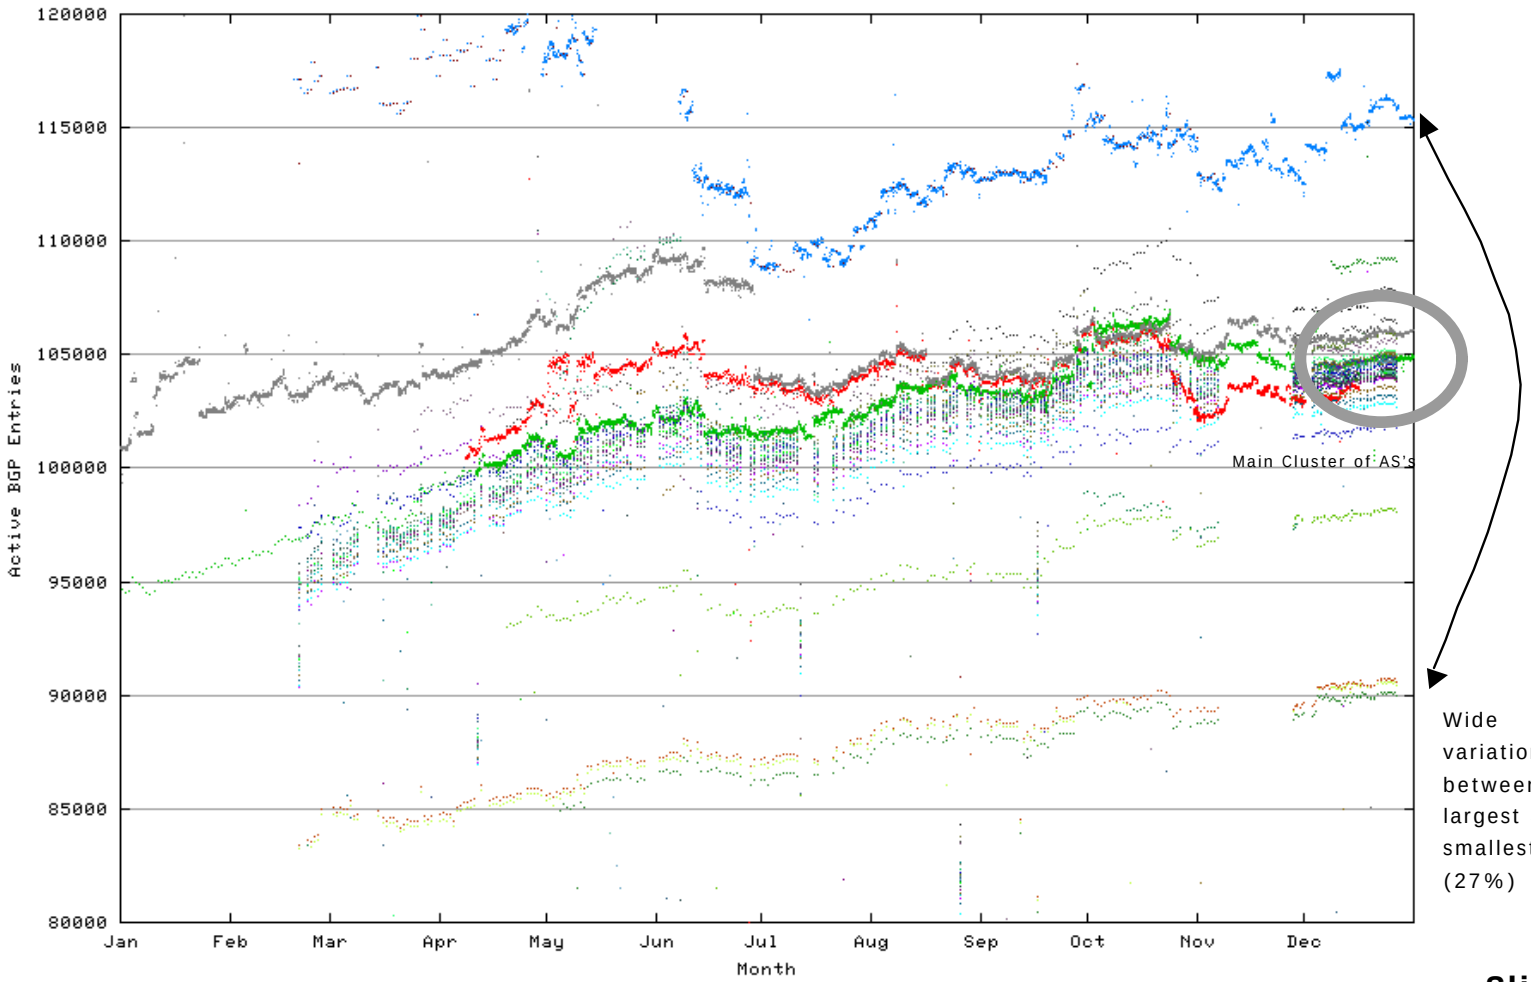

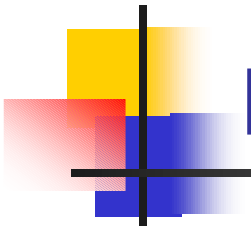

BGP in 2001

- Growth in Internet table size contained at roughly 105,000 entries through the year

Is this a stable state?

For how long?

Will exponential growth resume?

If so, at what rate?

2001 – Main Cluster Behaviour

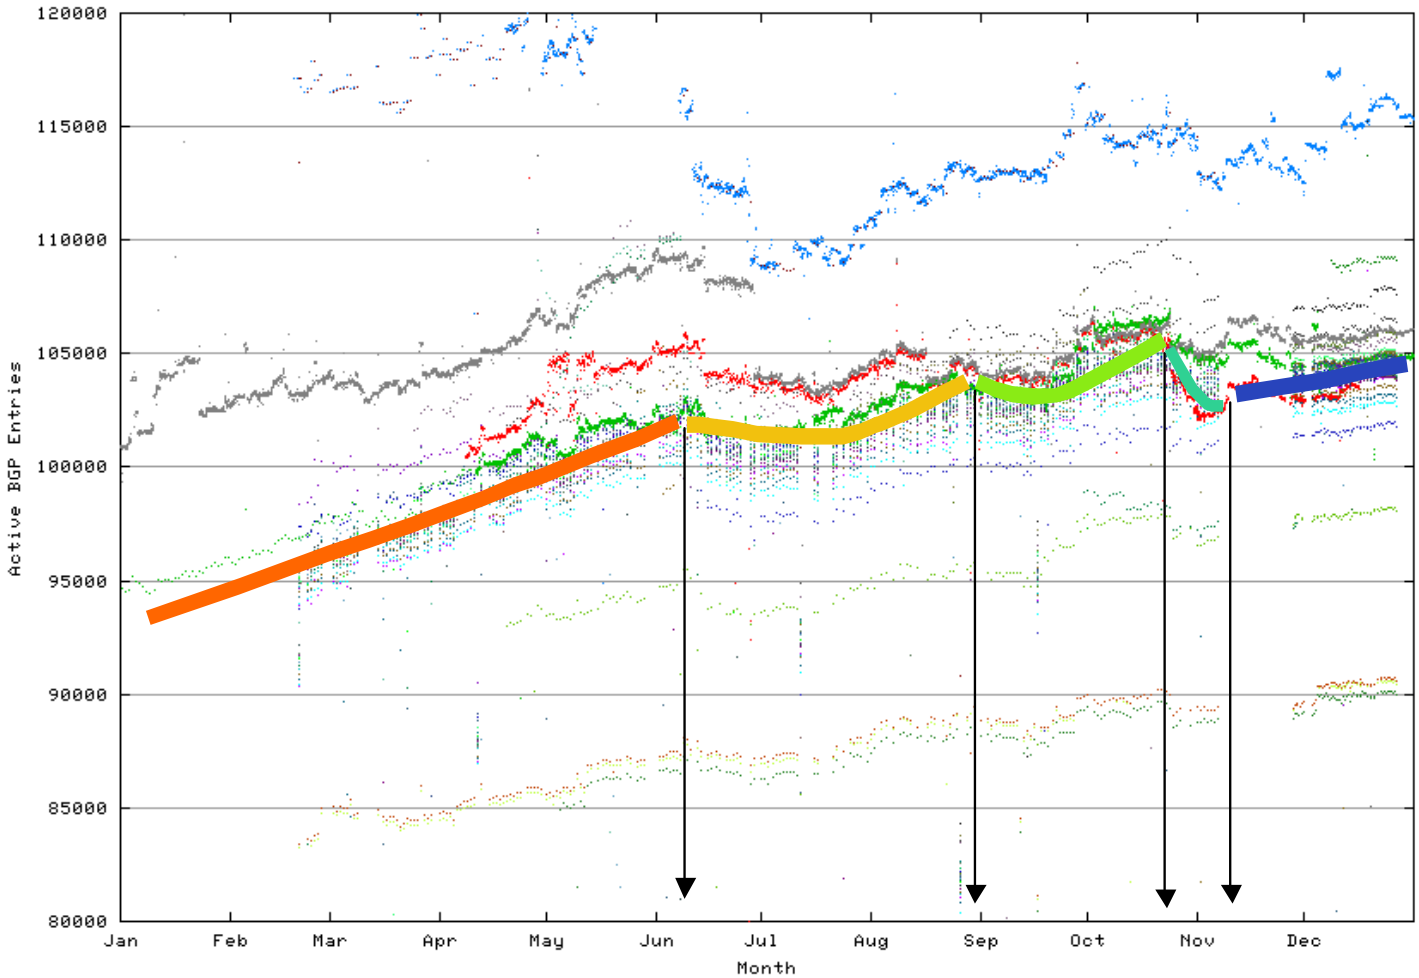

Has the Internet Stopped Growing in 2001?

- A number of other metrics do not show the same pattern as the number of BGP table entries:
  - Total routed address space
  - Number of AS's
  - Number of "root" prefixes in the BGP table

Internet Size: Routed Address Space

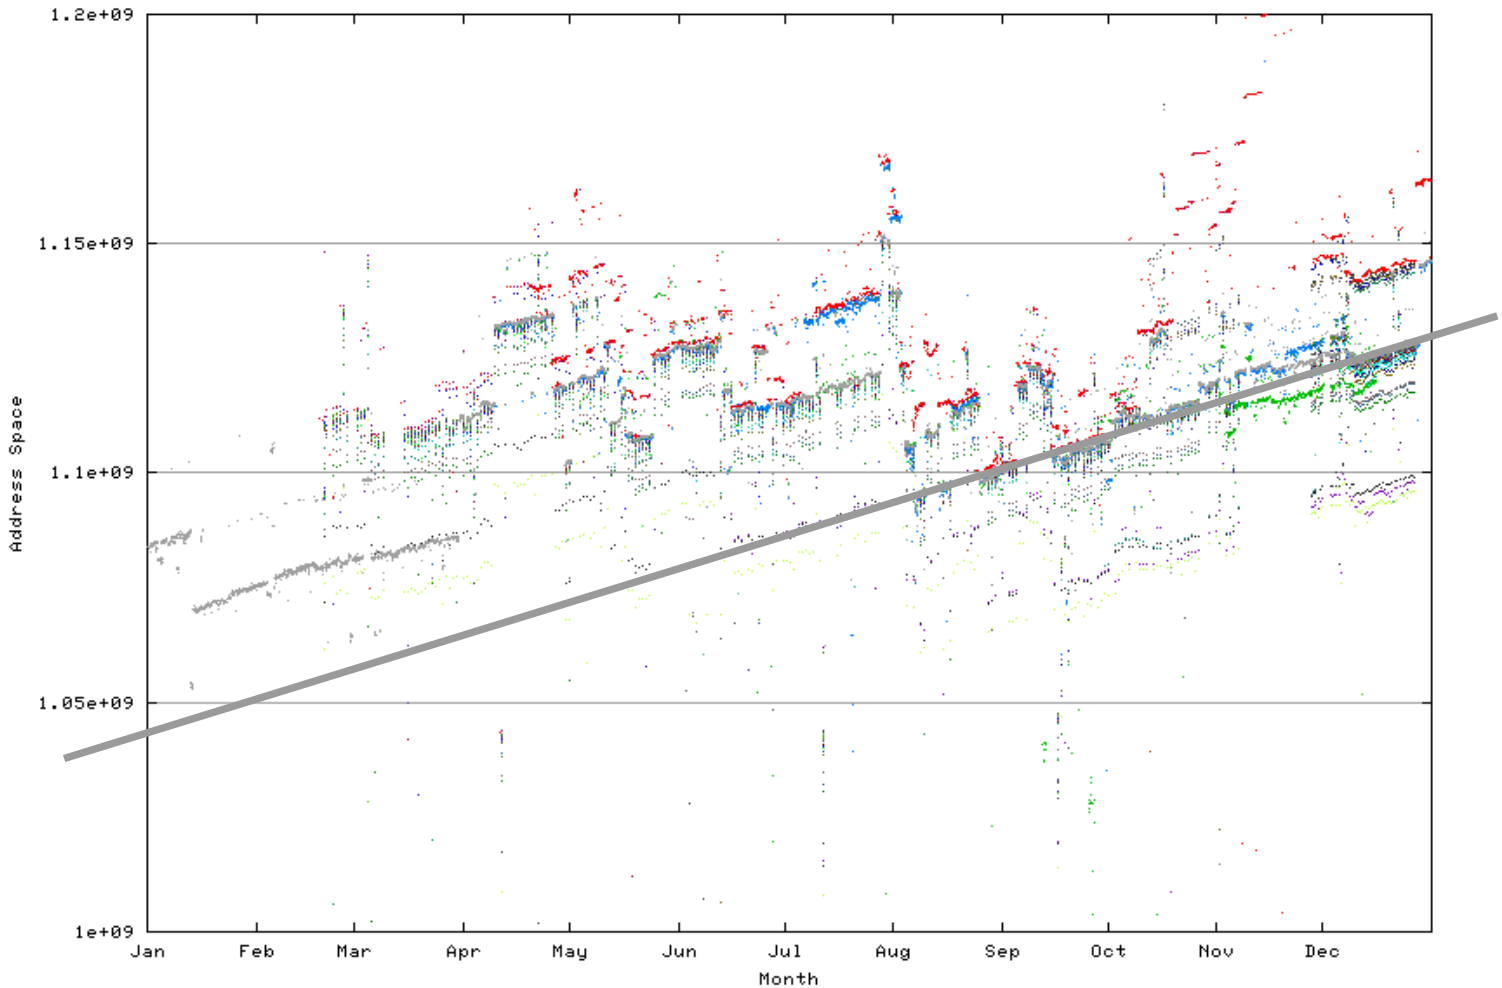

Steady growth in routed address space at an annual rate of 8%

Number of AS's

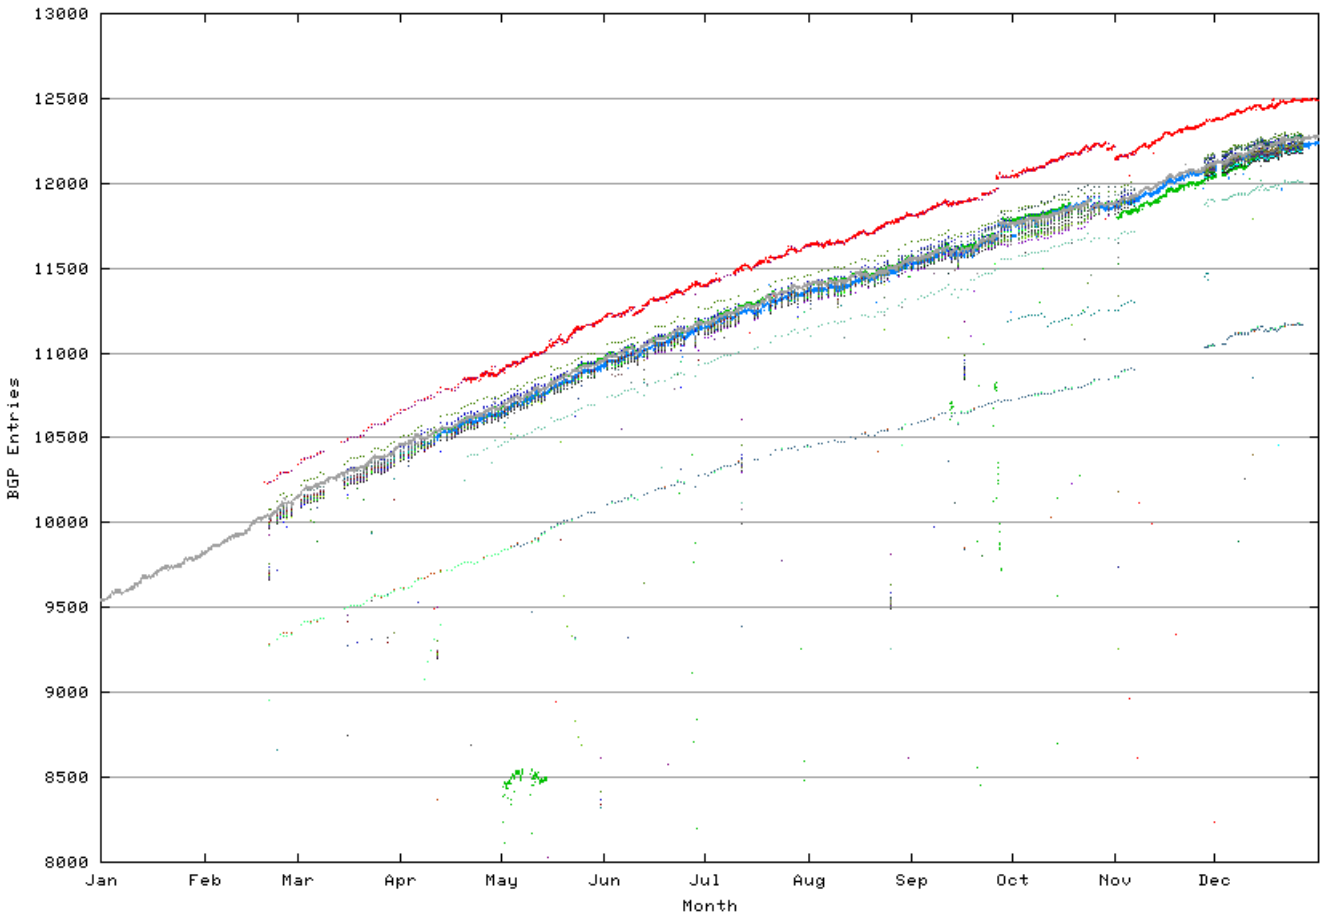

- AS's grew by 25% over the year
- Note span of visible AS's (11,200 – 12,500)
  - Not every AS is visible to all other AS's

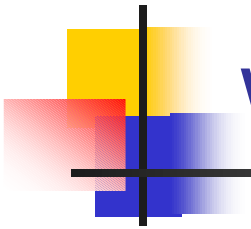

What Happened…

- The Internet continued to grow in 2001
- The routing space appeared to be better managed in 2001
  - Less routing “noise”
  - Better adherence to hierarchical aggregation in the routed address space

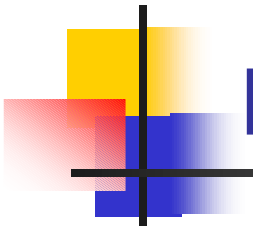

Per-Prefix views

- Some 60% of the routing table are /24 or smaller
- “Better” management of the routing space would see the relative numbers of small-sized prefixes declining
- And we have observed this in 2001…..

Relative percentage of /24 prefixes in the Routing Table

- /24 prefixes have declined by 3 – 4 % over 2001

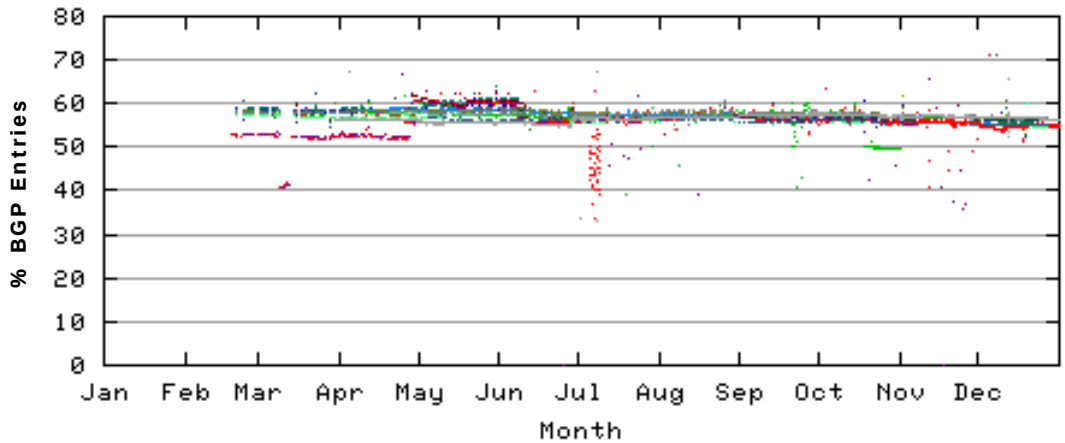

/24 Prefixes

- Largely steady at 60,000 entries for the year

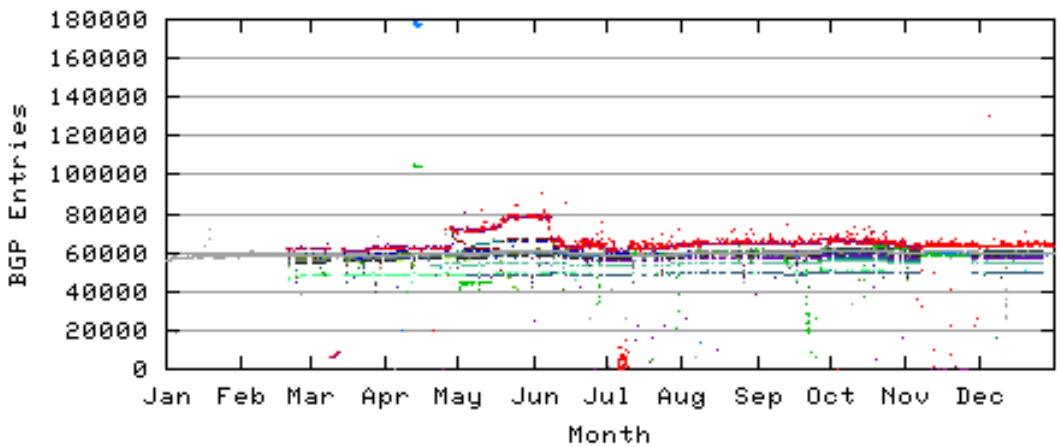

/20 Prefixes

- Grew from 4200 entries to 6100 entries (45% growth)
- Even growth throughout the year

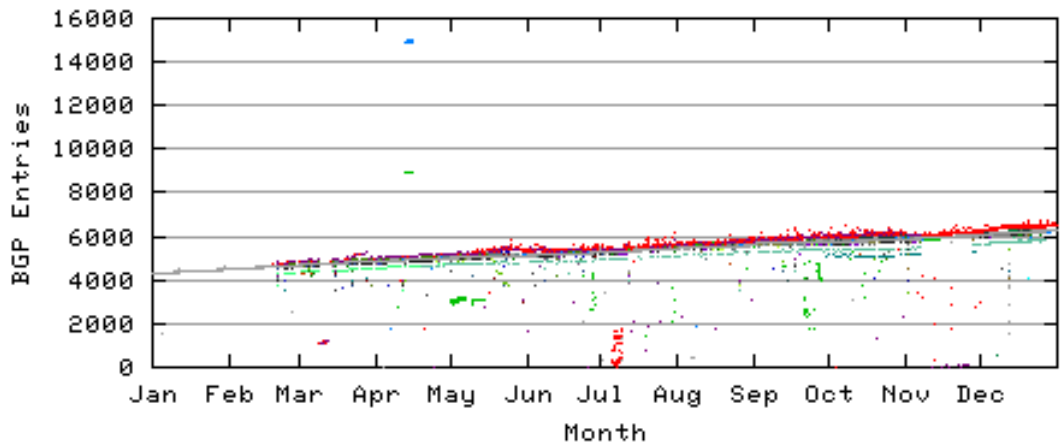

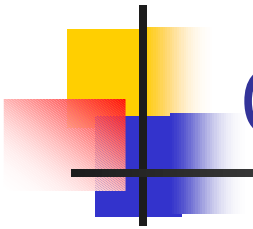

Changes in the Routing Table

- No major table growth from small prefixes (/24 and smaller)
- Table growth occurred using RIR allocation prefix sizes (/18 through /20)
- Growth in /18 - /20 prefix numbers even through the year

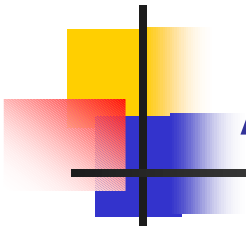

A “Root” Table Entry

- Is not part of an enclosing aggregate
- May contain any number of more specific entries
  - irrespective of AS Path of the specific
- Is the minimal spanning set of entries using a strict view of address / routing hierarchies
- Provides a view of the “best case” of the hierarchical model

Number of BGP “Roots” in 2001

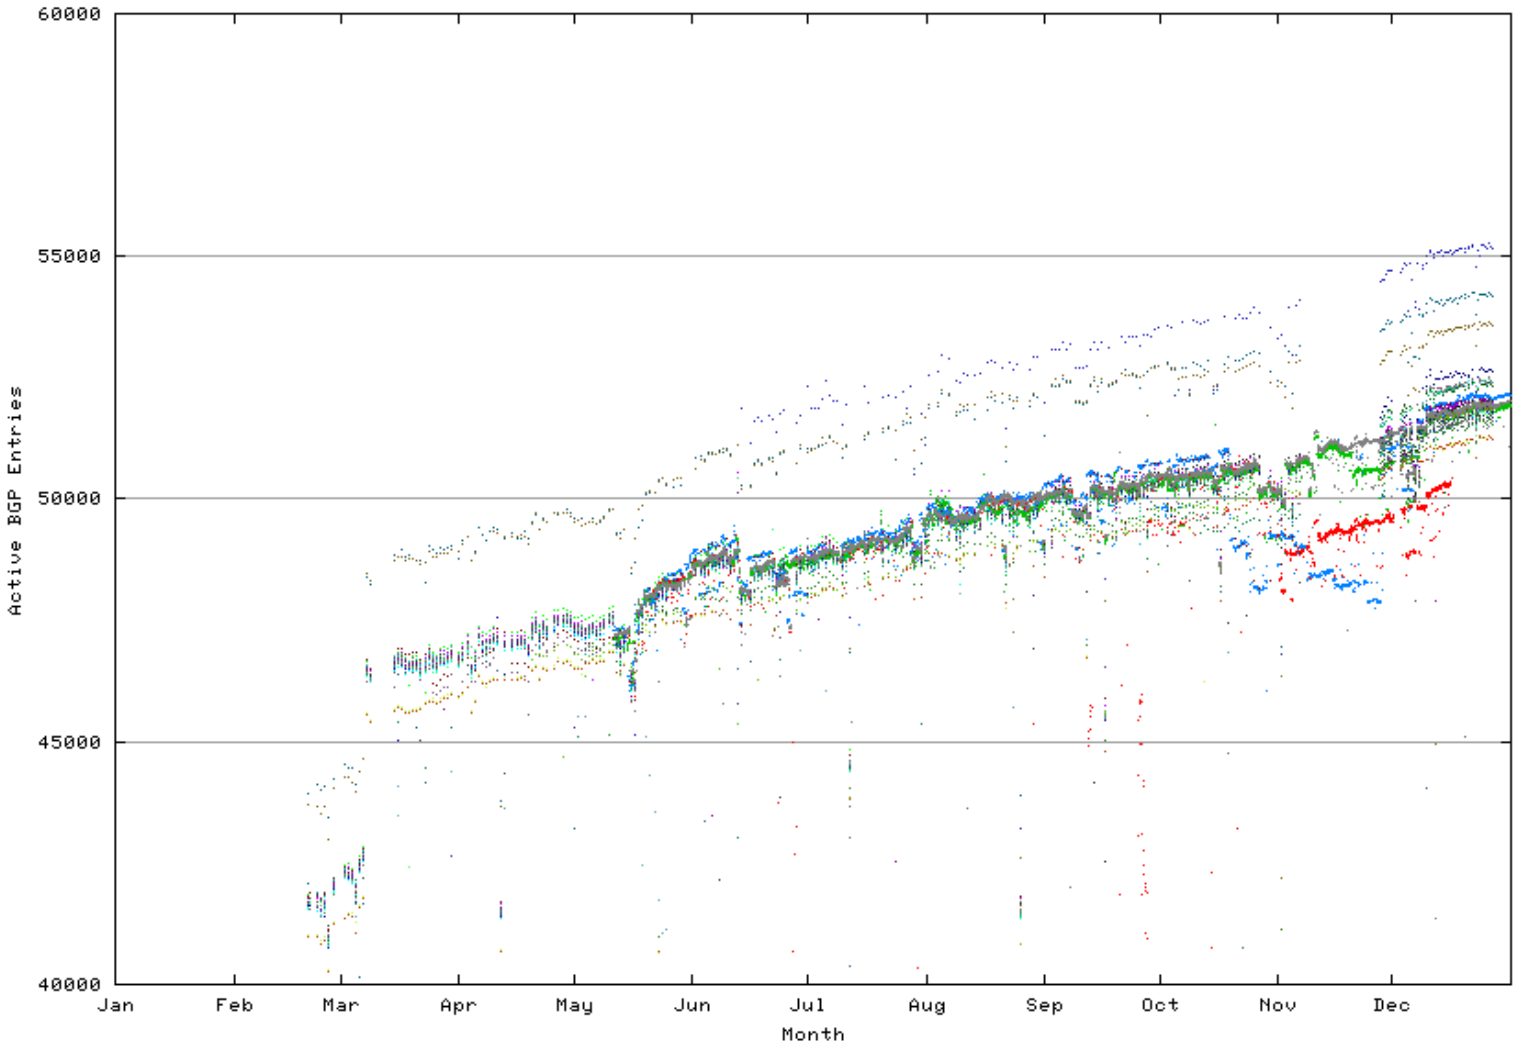

More Specifics (non-Roots) as a percentage of the table size

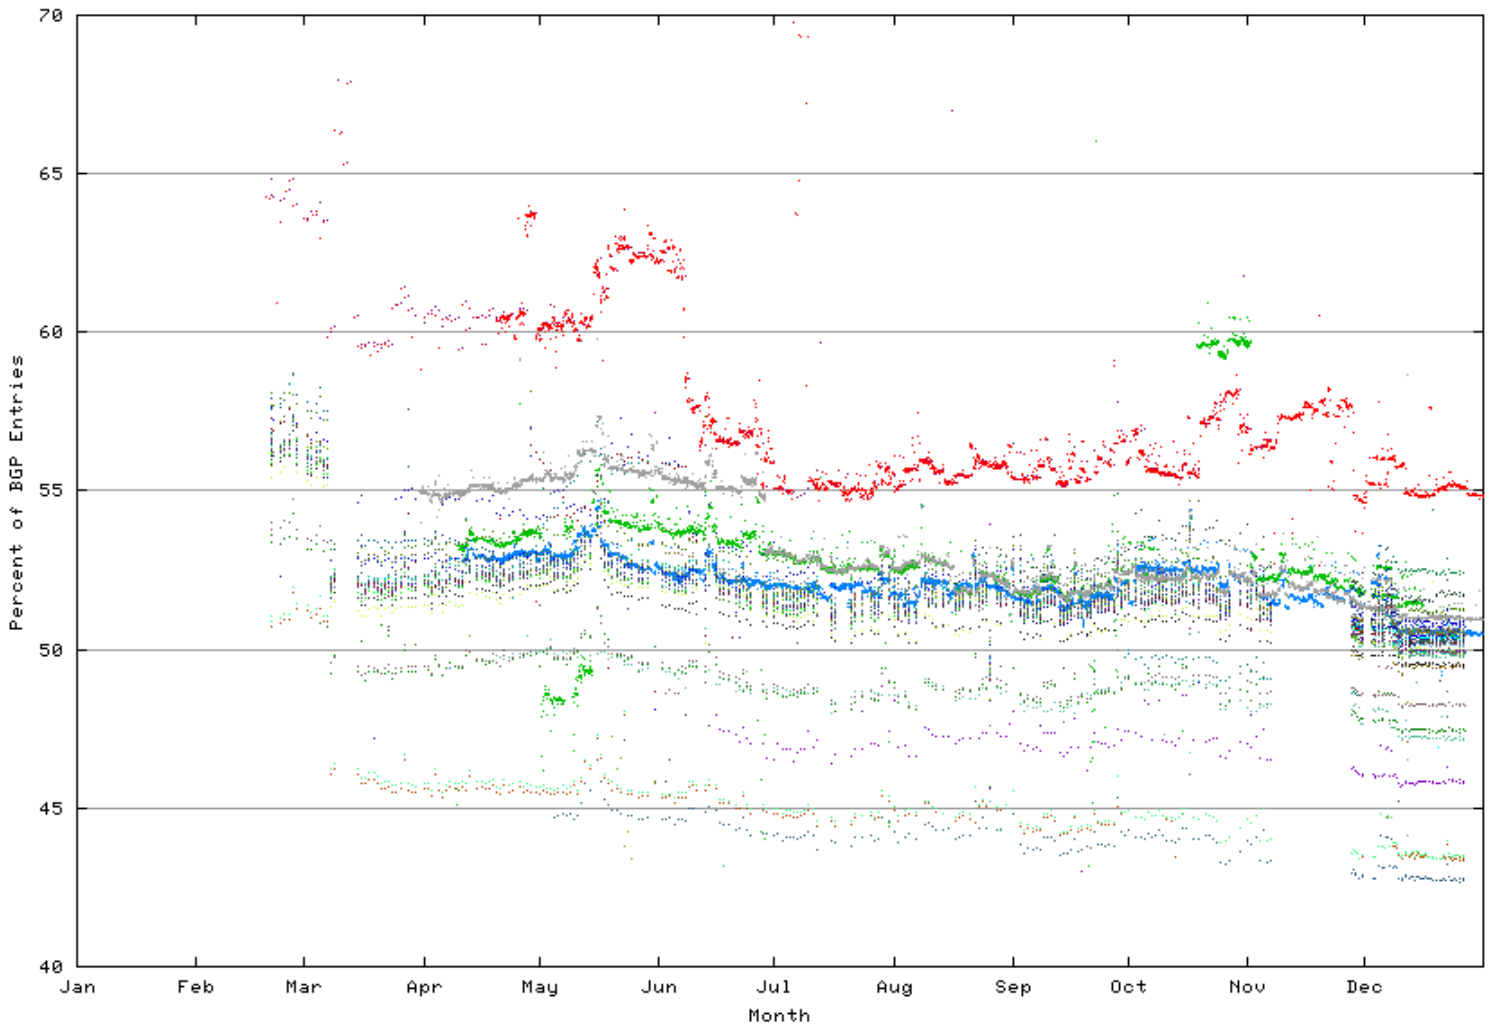

Whats Happening

- More specific entries in the routing table are declining in relative terms
- Possibly due to:
  - increasing amount of prefix-length route filtering
  - Increasing peer pressure to conform to RIR-allocated prefixes
  - Better understanding in the operator community of how to manage the routing space

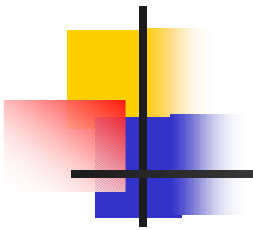

Interconnectivity Density

- Compare number of AS's to average AS path length
- A uniform density model would predict an increasing AS Path length ("Radius") with increasing AS's
- Increasing density predicts a constant or declining average AS Path Length

Average AS Path Length

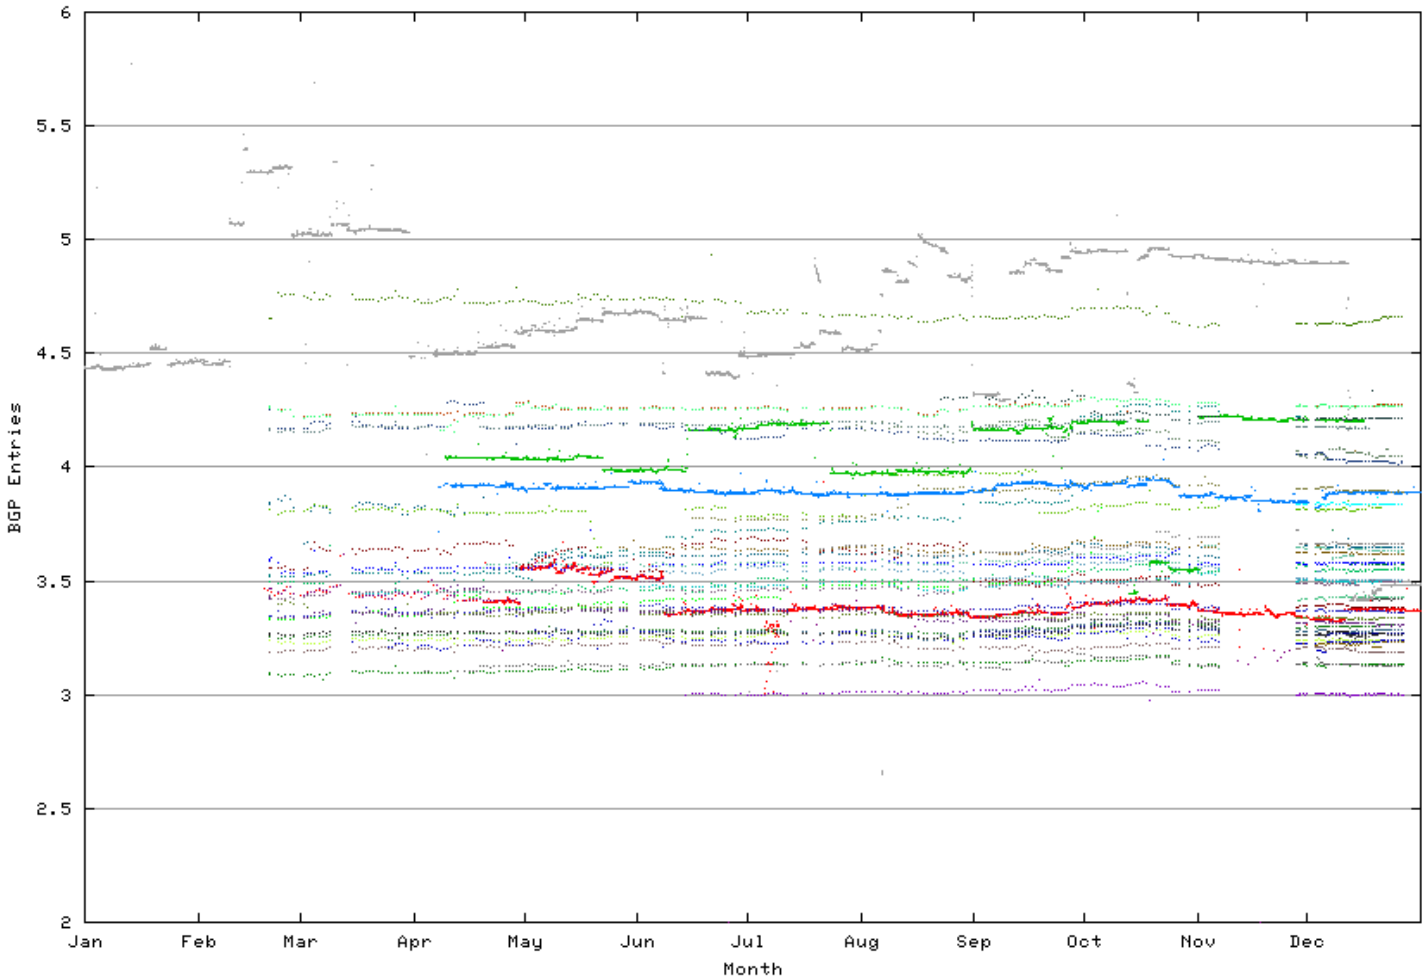

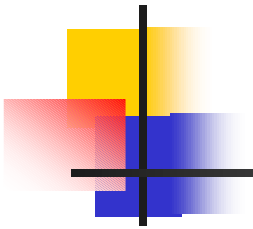

Interconnectivity Density

- Average number of per-AS interconnections was steady across 2001
  - Although the route views data is noisy due to the issues of
    - Dependence of the data on the number of BGP peer sessions
    - External exported view masks some level of local peer interconnection
    - Heavy tail distribution within the data

Average number of AS Neighbours

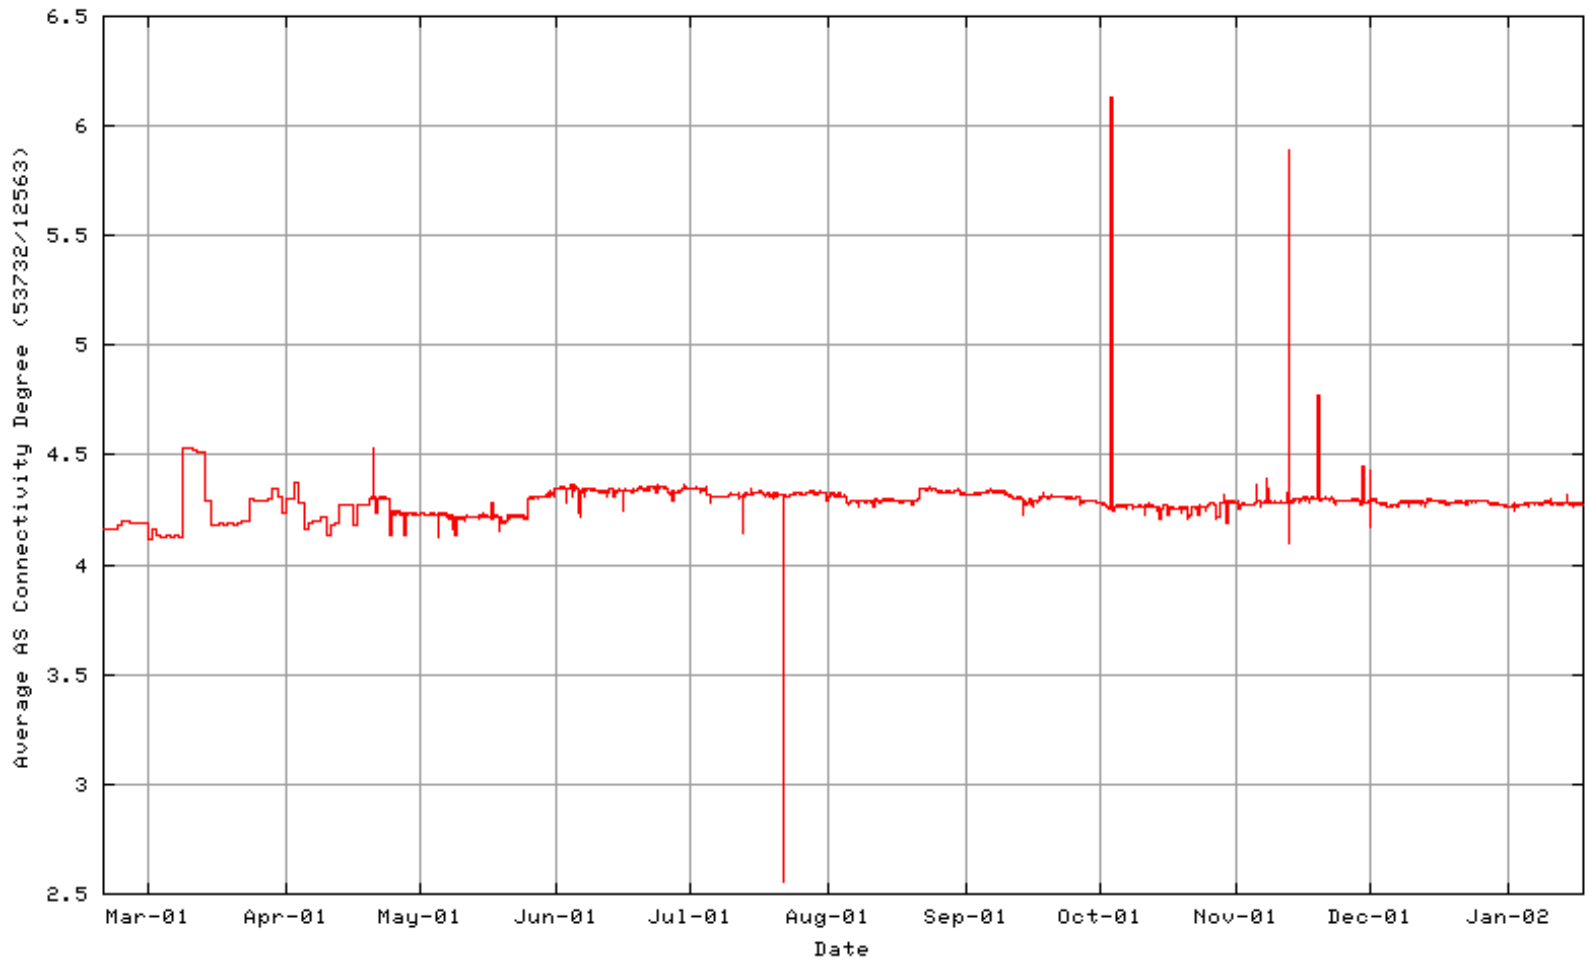

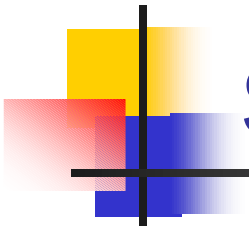

Stability of the BGP Table

- Measure rate of announcements + withdrawals + path updates
- Compare relative update rate per prefix length to the relative number of prefixes of that length
  - >1 implies higher than average update rate (less stable)
  - <1 implies lower than average update rate (more stable)

Stability Rates - /24 and /19

/19 Update rate

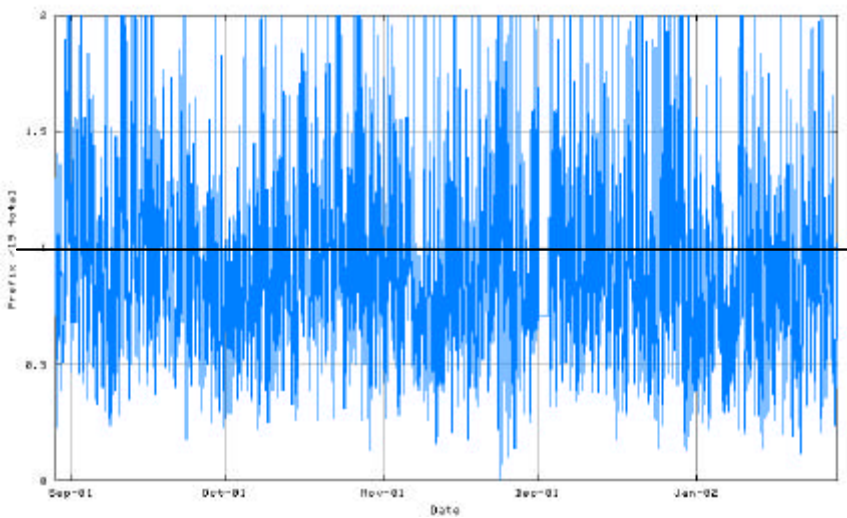

/24 Update rate

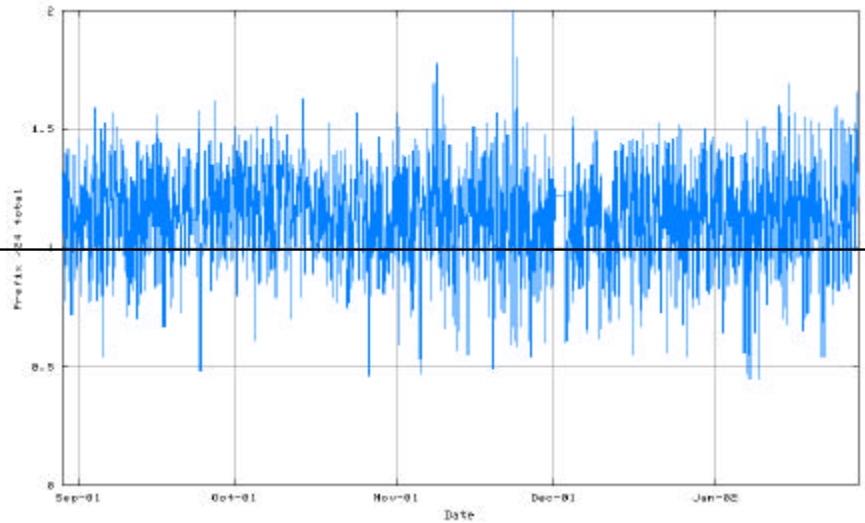

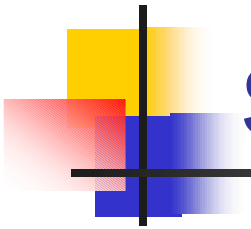

Stability Rates

- Smaller prefixes tend to contribute greater relative update load levels than larger prefixes
- Decreasing relative number of small prefixes is improving BGP stability levels (slightly)

BGP Update Rate

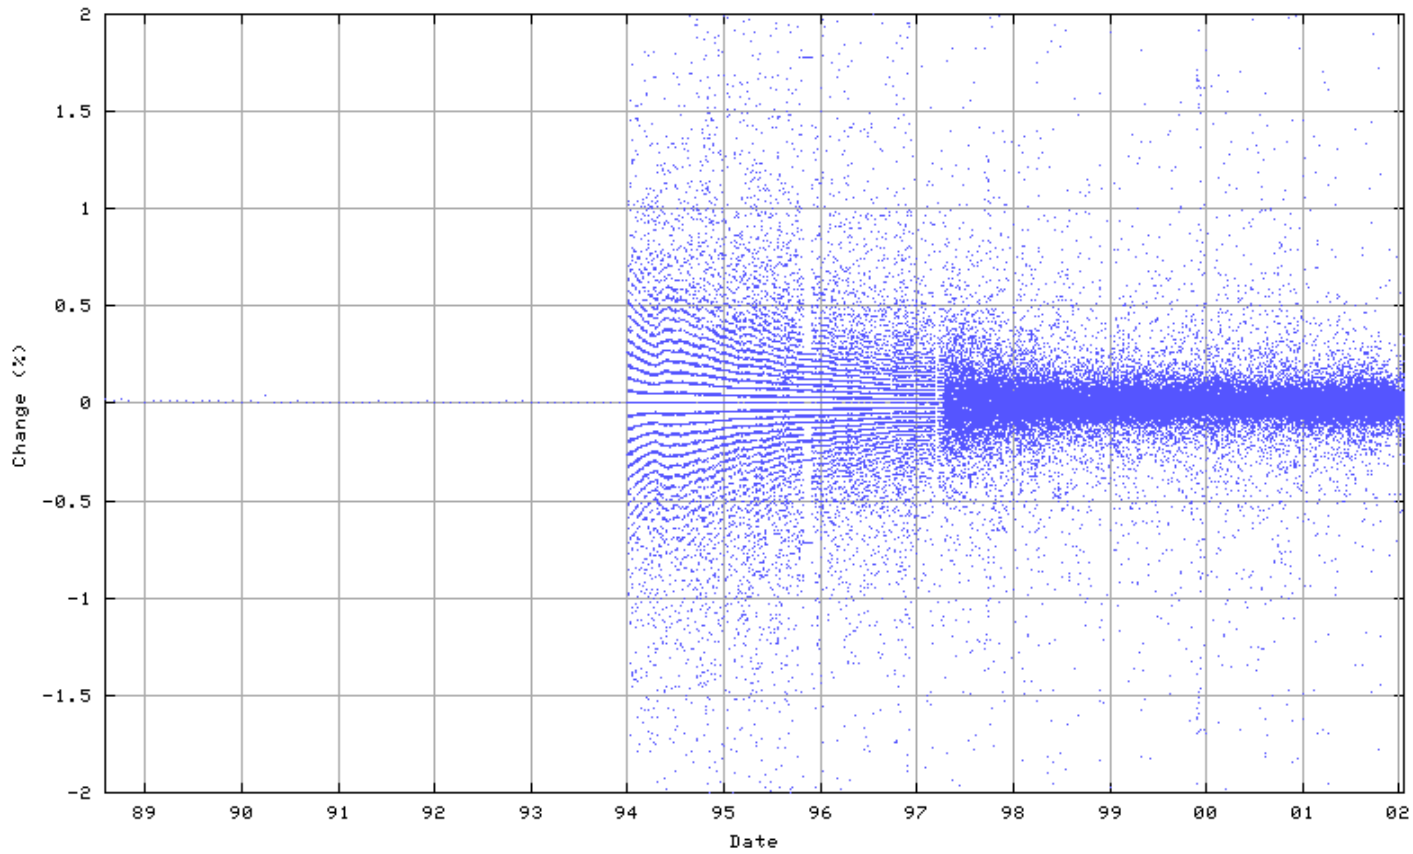

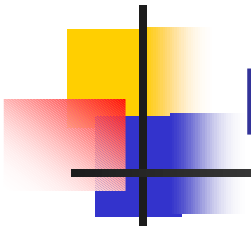

BGP Update Rate

- Proportion of BGP table entries updated each hour is decreasing over time
- The BGP table is becoming more stable
  - Protocol implementation maturity
  - Widespread deployment of flap damping
  - Greater levels of circuit reliability (?)

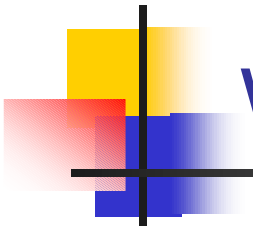

What Happened

- "Base" growth rate of root prefixes was **15%** in 2001
- Growth rate of AS's was **25%** in 2001
- Growth rate of routed address space was **8%** in 2001

- By comparison, annual growth rate of the BGP table for the previous 2 years was **55%**

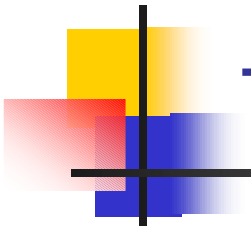

The Good News

- BGP Table growth has been slowed down considerably
- This is largely the result of more care in routing announcements, coupled with more widespread prefix length route filters.

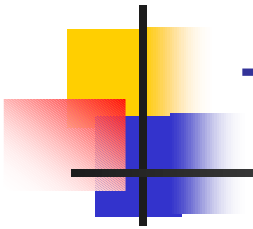

The Not So Good News

- Insufficient data to determine if this is a short term growth correction that will be followed by a resumption of exponential growth
  - Multi-homing, TE, mobility all contribute to a requirement for non-aggregatable atomic entries to be non-locally routed.

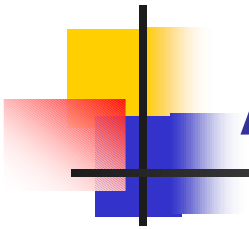

A Useful Agenda (1)

- Stress the value in widespread adoption of operational best practices in BGP
  - Route aggregation
  - Prefix length filtering
  - Advertisements that align with RIR allocation units
  - Flap damping
  - Soft refresh

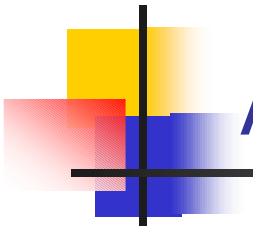

A Useful Agenda (2)

- Understand what metrics of the IDR space are important to track
  - Network Size and Topology
  - The relationship between connectivity policy and topology
  - The relationship between address deployment and connectivity
  - Dynamic properties of the routing system system

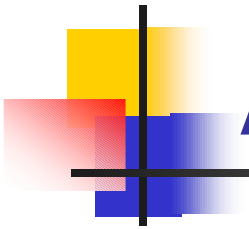

A Useful Agenda (3)

- Define the desireable properties of an inter-domain routing system
- Clearly understand the difference between policy mediated best path computation and the dynamic resource management requirements associated with traffic engineering and QoS
  - and be prepared to admit that doing 1 out of 3 is still better than doing 0 out of 3!

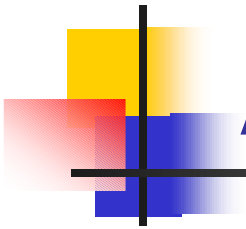

A Useful Agenda (4)

- Examine potential alternative approaches to Inter-Domain Routing systems that may offer superior scaling properties and greater flexibility in scope