

BGP: 2008

Geoff Huston
APNIC

Some BGP-related questions:

The Network Operator:

- How big a router should you buy today if you want it to cope with BGP in 3 years time? What FIB size? What processing capability?
- What about 5 years?

The Protocol Engineer:

- Is BGP scaling or failing?
- Do we need to develop a new IDR protocol?
- How much time do we have?

The Network Architecture Researcher:

- Are the Internet's concepts of names and addresses adequate?
- Are alternate models of id / loc split more friendly to routing?
- Is routing scaling an intractable problem within the confines of the current architecture?

An approach to generating some answers

- Use real data about today's network!
 - Understand the actual data about IDR and the operation of BGP
 - Use a long base line of consistent observations of BGP behaviour
 - Analyze the data carefully

BGP measurements

- There are a number of ways to “measure” BGP:
 - Assemble a large set of BGP peering sessions and record everything
 - RIPE NCC’s RIS service
 - Route Views
 - Perform carefully controlled injections of route information and observe the propagation of information
 - Beacons
 - AS Set manipulation
 - Bogon Detection and Triangulation
 - Take a single BGP perspective and perform continuous recording of a number of BGP metrics


AS131072 (or AS2.0) BGP measurement

- Successor to the AS1221 observation point
- Data collection since 1 July 2007 (since 2000 for AS1221)
- Passive data measurement technique (no advertisements or probes)
- Quagga platform, connected to AS4608 and AS4777 via eBGP
- IPv4 and IPv6 simultaneous
- Archive of all BGP updates and daily RIB dumps
- Data and reports are continuously updated and published:
<http://bgp.potaroo.net>


BGP: 2008

- Analysis of BGP data collected at AS131072 from 1 January 2008 to 31 December 2008


IPv4 BGP Prefix Count




IPv4 BGP Prefix Count




IPv4 Routed Address Span




IPv4 Routed Address Span



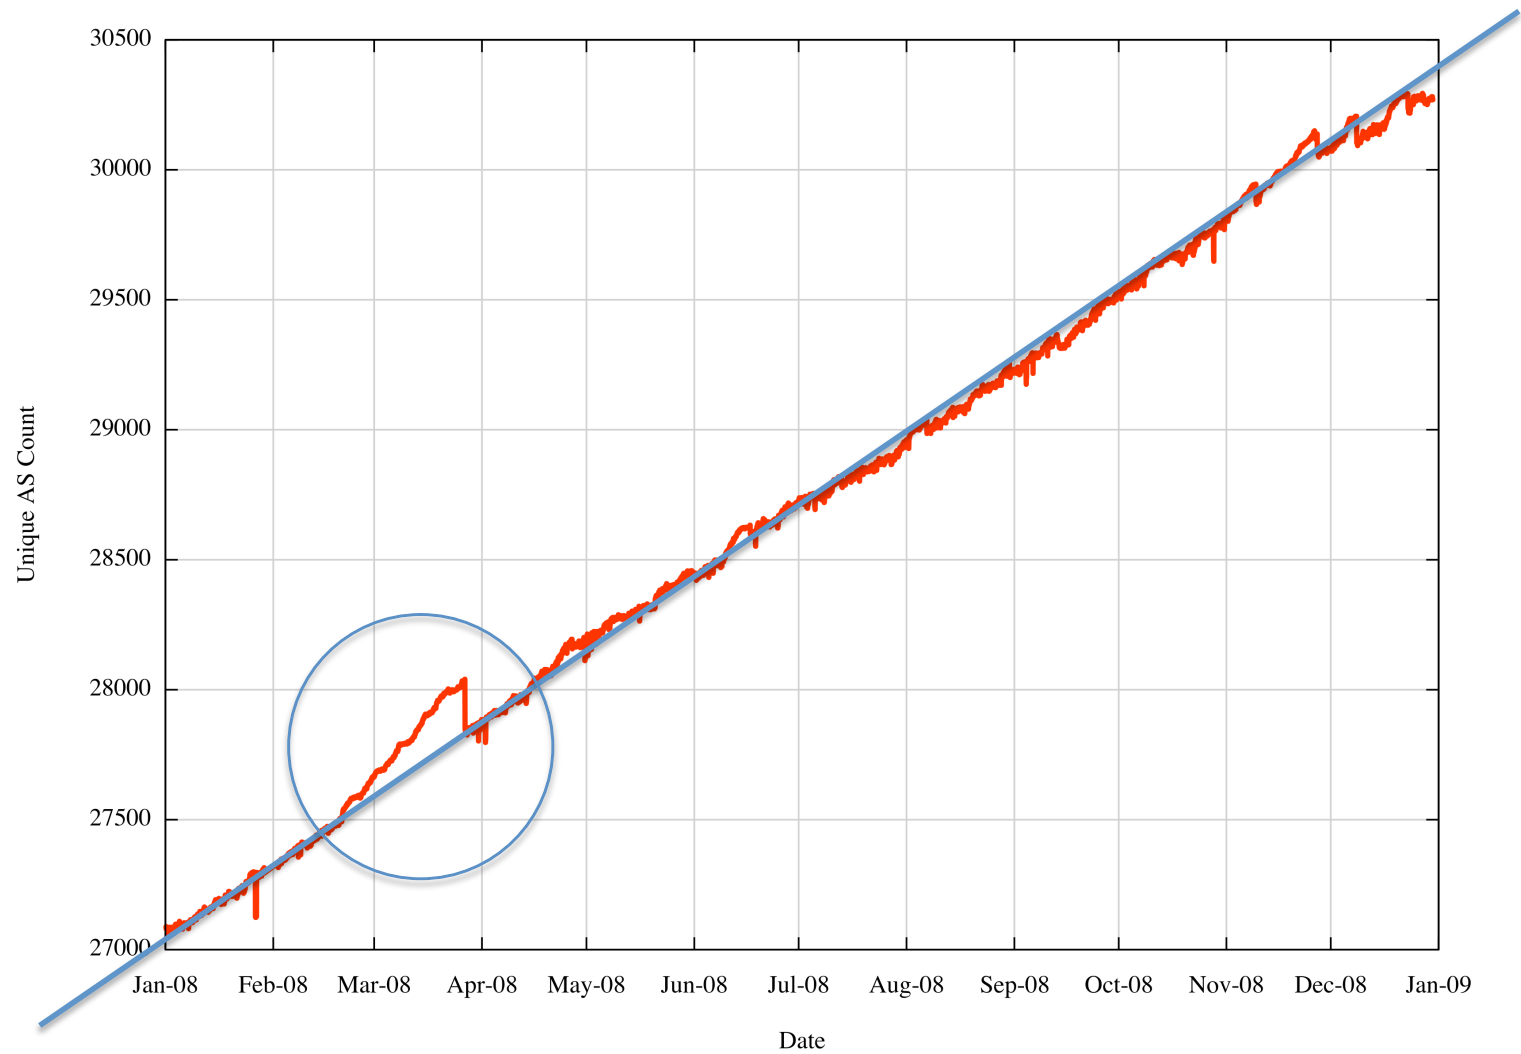
IPv4 Routed Address Span



IPv4 Routed AS Count



IPv4 Routed AS Count




IPv4 Vital Statistics for 2008

	Jan-08	Dec-08	
Prefix Count	245,000	286,000	17%
Roots	118,000	133,000	13%
More Specifics	127,000	152,000	20%
Address Span	106.39	118.44	11%
AS Count	27,000	30,300	11%
Transit	3,600	4,100	14%
Stub	23,400	26,200	11%

Some Observations


- Growth in IPv4 deployment slowed considerably as of the end of April 2008
 - Is this a possible consequence of the financial crash of 2008?
- Fragmentation of the IPv4 routing space continues to grow at a faster pace than underlying growth of the network itself

IPv4 prefix distribution




IPv6 BGP in 2008

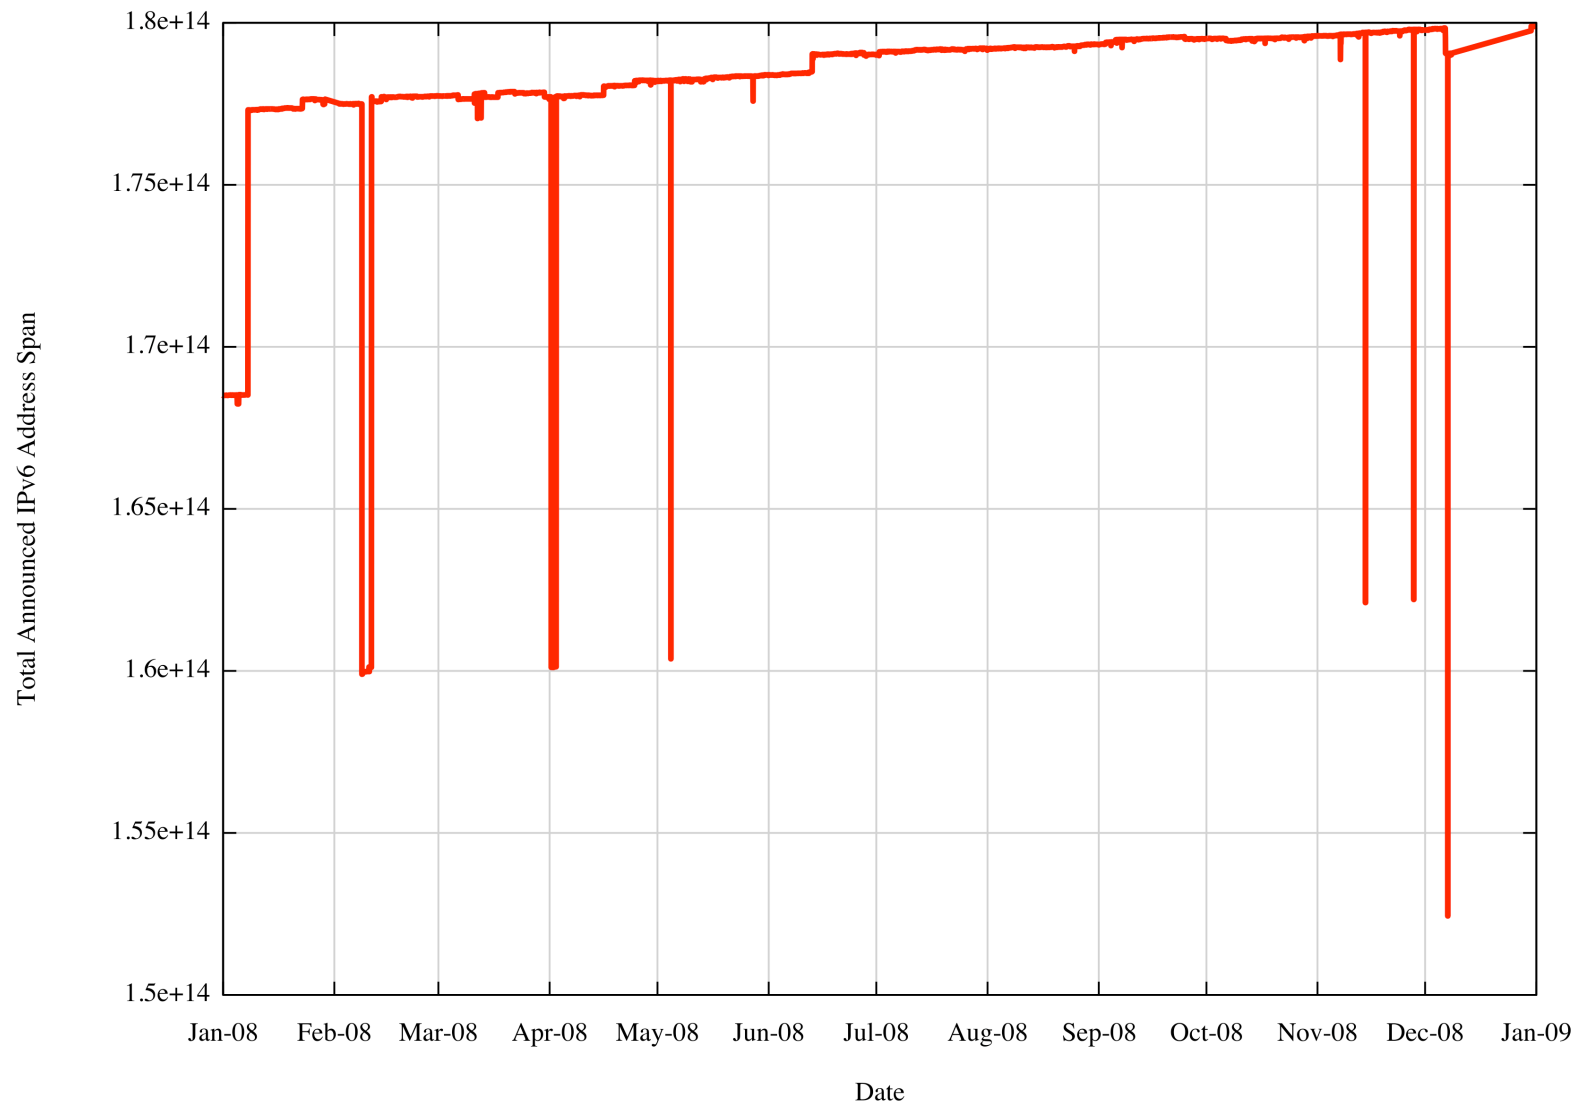
IPv6 BGP Prefix Count




IPv6 BGP Prefix Count




IPv6 Routed Address Span




IPv6 Routed Address Span




IPv6 Routed Address Span



IPv6 Routed AS Count



IPv6 Routed AS Count



IPv6 Vital Statistics for 2008


	Jan-08	Dec-08	
Prefix Count	1,050	1,600	52%
Roots	840	1,300	55%
More Specifics	210	300	43%
Address Span	/16.67	/16.65	1%
AS Count	860	1,230	43%
Transit	240	310	29%
Stub	620	920	48%

Trends and Projections

BGP Table Size Projection


- 4 year projection of the IPv4 routing table size

IPv4 Table Size - 12 months




daily average and 60 day sliding window applied to smooth the data

IPv4 Table Size - 60 months




Daily Growth Rates




First order
differential of
the smoothed
data

Daily Growth Rates




IPv4 Table Size

Quadratic Growth Model



IPv4 Table Size Quadratic Growth Model - Projection



BGP Table Size Predictions

now	285,000 entries
12 months	335,000 entries
24 months	388,000 entries
36 months*	447,000 entries
48 months*	512,000 entries


* These numbers are very dubious due to IPv4 address exhaustion pressures. It is probable that the number will be larger than this

Back in 2006 ...

- This modelling work of the IDR table size was performed at the end of 2005 to generate a 3 and 5 year projection

2006 prediction

RIB SIZE - Predictive Model



BGP Table Size Predictions

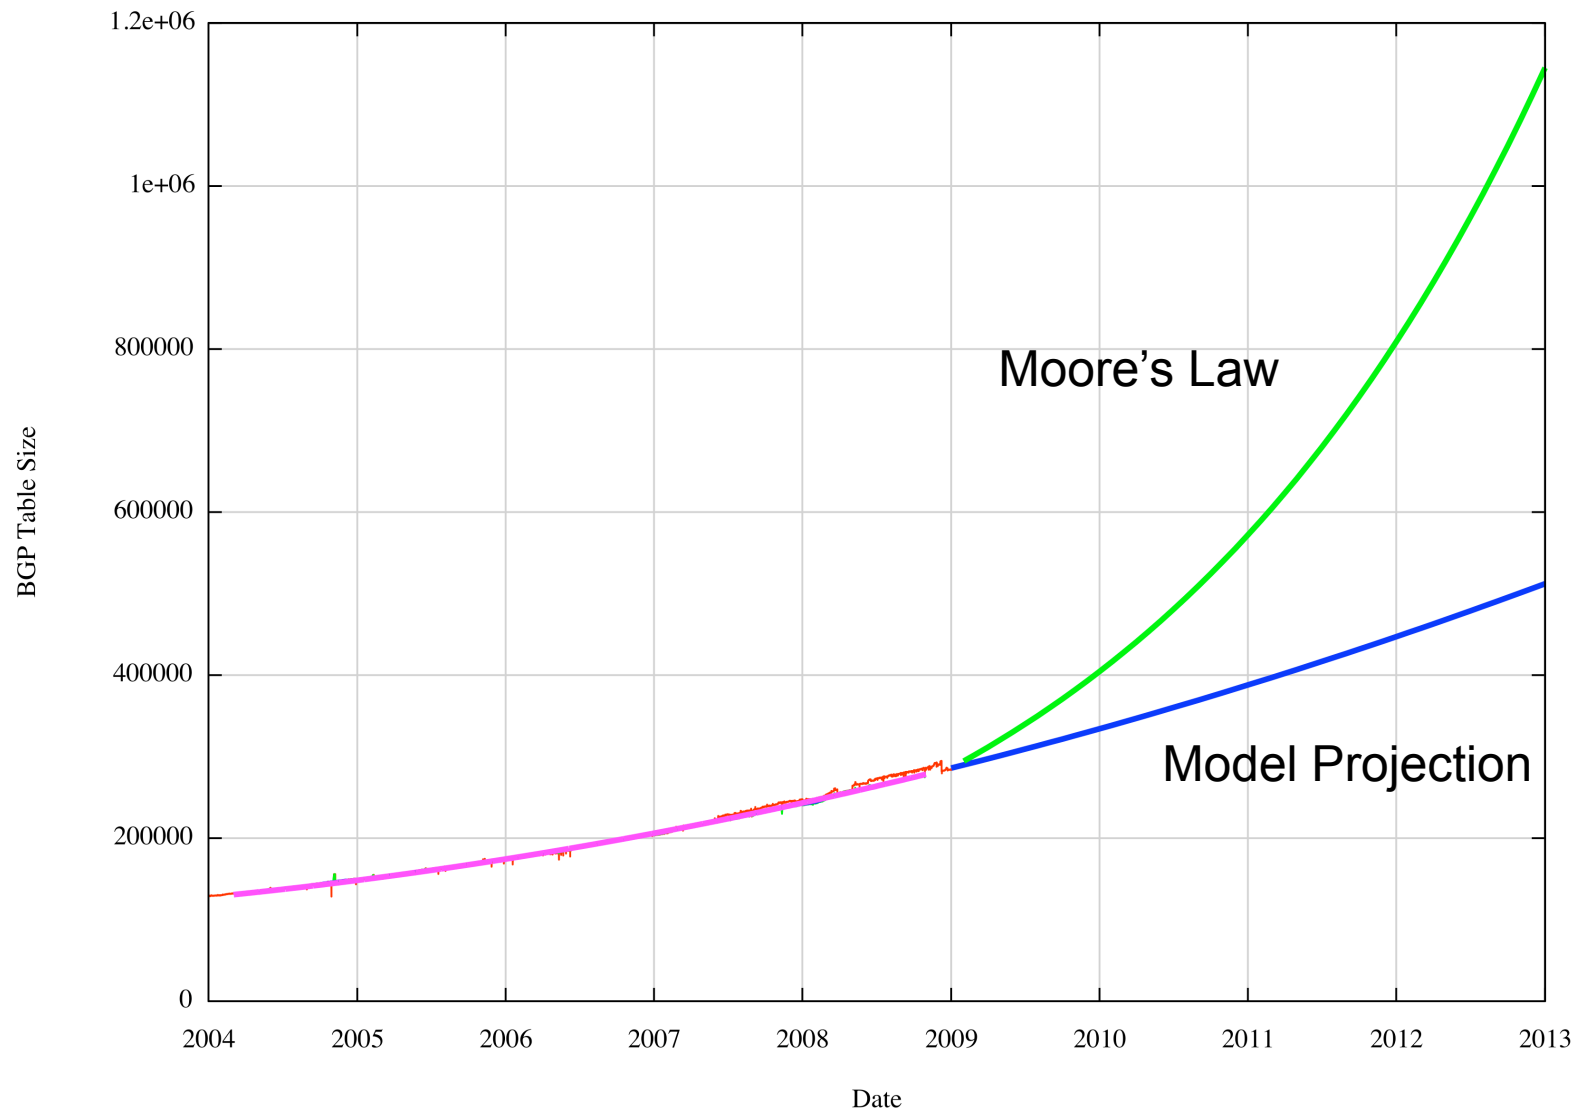
now	285,000 entries (2006: 275,000)
12 months	335,000 entries
24 months	388,000 entries (2006: 370,000)
36 months*	447,000 entries
48 months*	512,000 entries

* These numbers are very dubious due to IPv4 address exhaustion pressures. It is probable that the number will be larger than this.

Moore's Law

- As long as growth rates stay within the general parameters of Moore's Law the unit cost of the routing function should not escalate
 - use a function of doubling every two years as the application of Moore's Law
 - assuming that Moore's law continue to hold


Projections against Moore's Law




BGP Update Behaviour

- What trends are visible in the number of BGP updates?


BGP Update Behaviour




BGP Updates - Announce and Withdrawals




Daily Updates




Best Fit to Updates




BGP Updates - Extended Data Set




BGP Updates - Extended Data Set




BGP Update Projection




Daily Withdrawals



Best Fit to WDLs




BGP Withdrawal Projection



Why is this?

- Why are the levels of growth in BGP updates not proportional to the size of the routing table?
 - growth rates of BGP updates appear to be far smaller than the growth rate of the routing space itself
- What is going on?


Average AS Path Length has remained constant



Why is this?

- The convergence instability factor for a distance vector protocol like BGP is related to the AS path length, and average AS Path length has remained steady in the Internet for some years
- Taking MRAI factors into account, the number of received Path Exploration Updates in advance of a withdrawal is related to the propagation time of the withdrawal message. This is approximately related to the average AS path length
- Today's Internet is more densely interconnected, but is not any more “stringier”
- This implies that the number of protocol path exploration transitions leading to a prefix withdrawal is relatively stable over time

BGP is being swamped by TE noise



Why is this?

- The major component of dynamic BGP load is not an artifact of the larger routing space, but a case of BGP path manipulation at origin for TE purposes from a small subset of origin AS's
 - what is being measured in updates is not implicitly related to network component stability, but more likely to be related to path manipulation related to TE

A prudent view of capacity planning for 4 - 5 years into BGP's future

- A FIB size of 750,000 entries of IPv4, and a total of 1M entries of IPv4 + IPv6 FIB space
- A BGP update load of 2x the current rates
- No major changes in the BGP use profile, no major changes in the distribution of BGP updates

Some Closing Opinions

- The BGP sky is **not** falling

The 2008 BGP data appears to indicate that the prospects of the imminent death of BGP through routing table inflation appear to be vastly exaggerated

- The inflation rate of the routing table remains well under Moore's law
- The rate of increase of processed updates in the gather data appears to be minimal
- The rate of increase of withdrawals is slightly larger, but is still slight
- The network is, on the whole, very stable and BGP is not under immediate stress in terms of scaling pressures

- But this does not model post V4 address exhaustion scenarios
 - But there is no evidence to be alarmed about excessive fragmentation of the routing space at this stage

A Word of Caution

- This is a simple exercise in statistical curve fits, not a demand level simulation of the players routing environment.
- This exercise does not factor in any IPv4 address exhaustion considerations and scenarios around address movement that may alter the picture of fragmentation of the routing space.
- However the AS growth projections are a strong indicator of underlying industry dynamics in terms of discrete routing entities, and these projections show a modest growth component

This means that while the projections are very weak in the period of 2011 and beyond, there are grounds to take a conservative view of BGP growth in this phase of the Internet's evolution

Thank You

Questions?