


2006 – (Almost another) BGP Year in Review

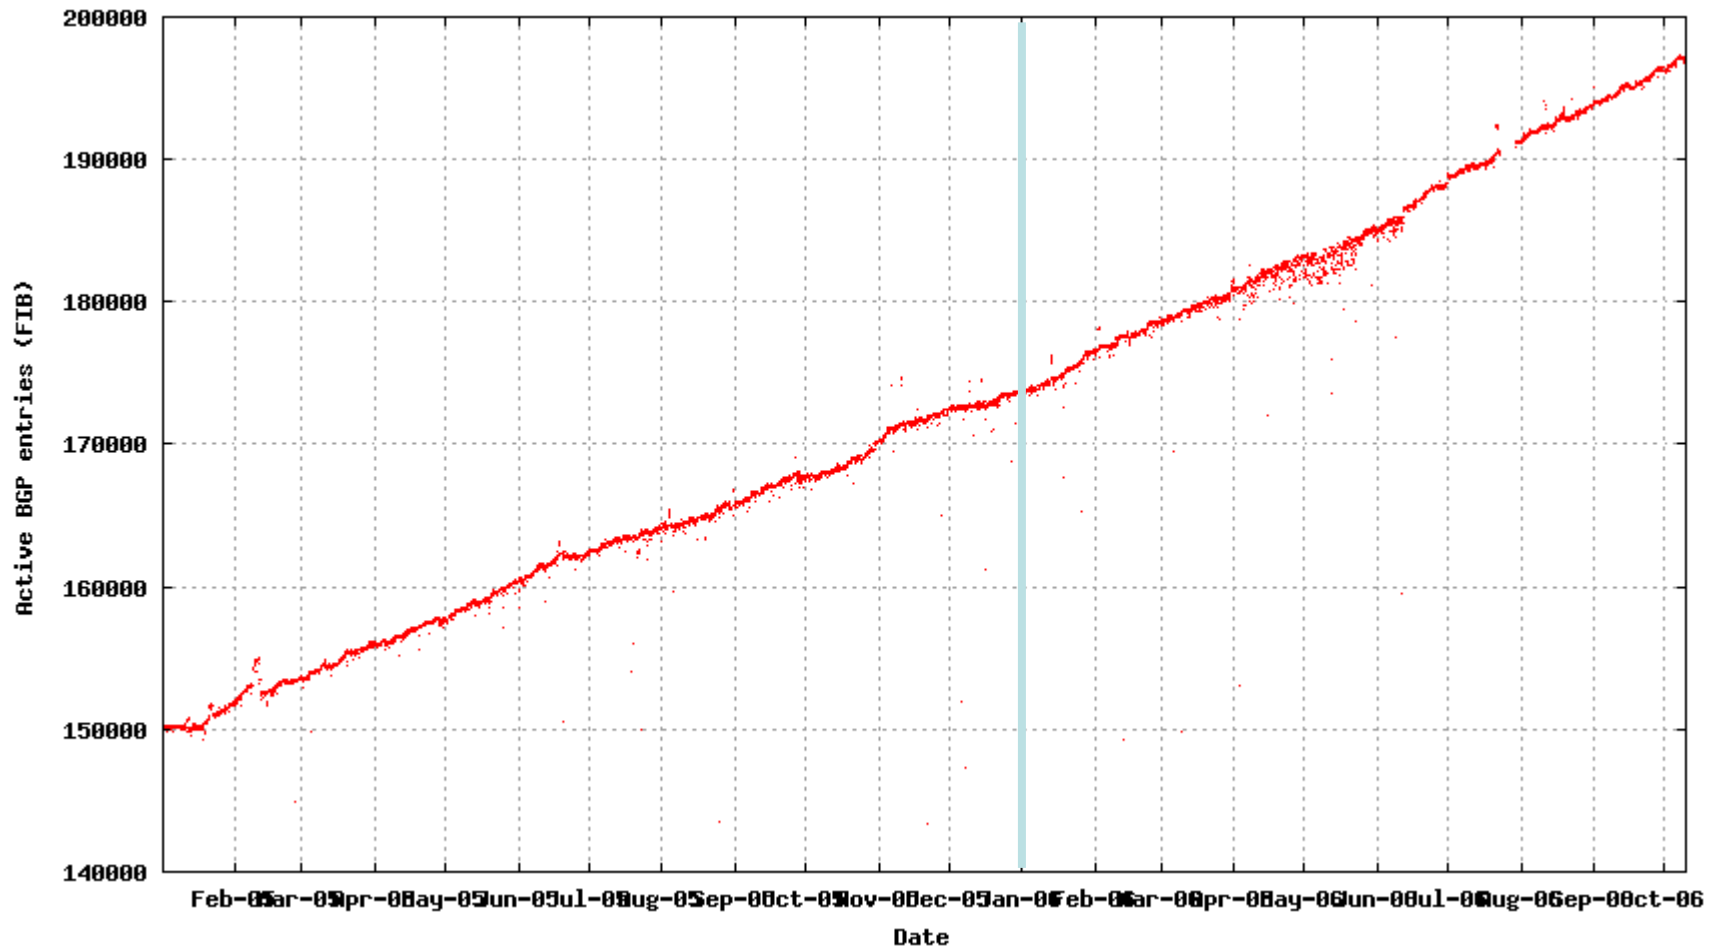
A BRIEF update to the 2005 report



18 October 2006
IAB Routing Workshop
Geoff Huston
APNIC

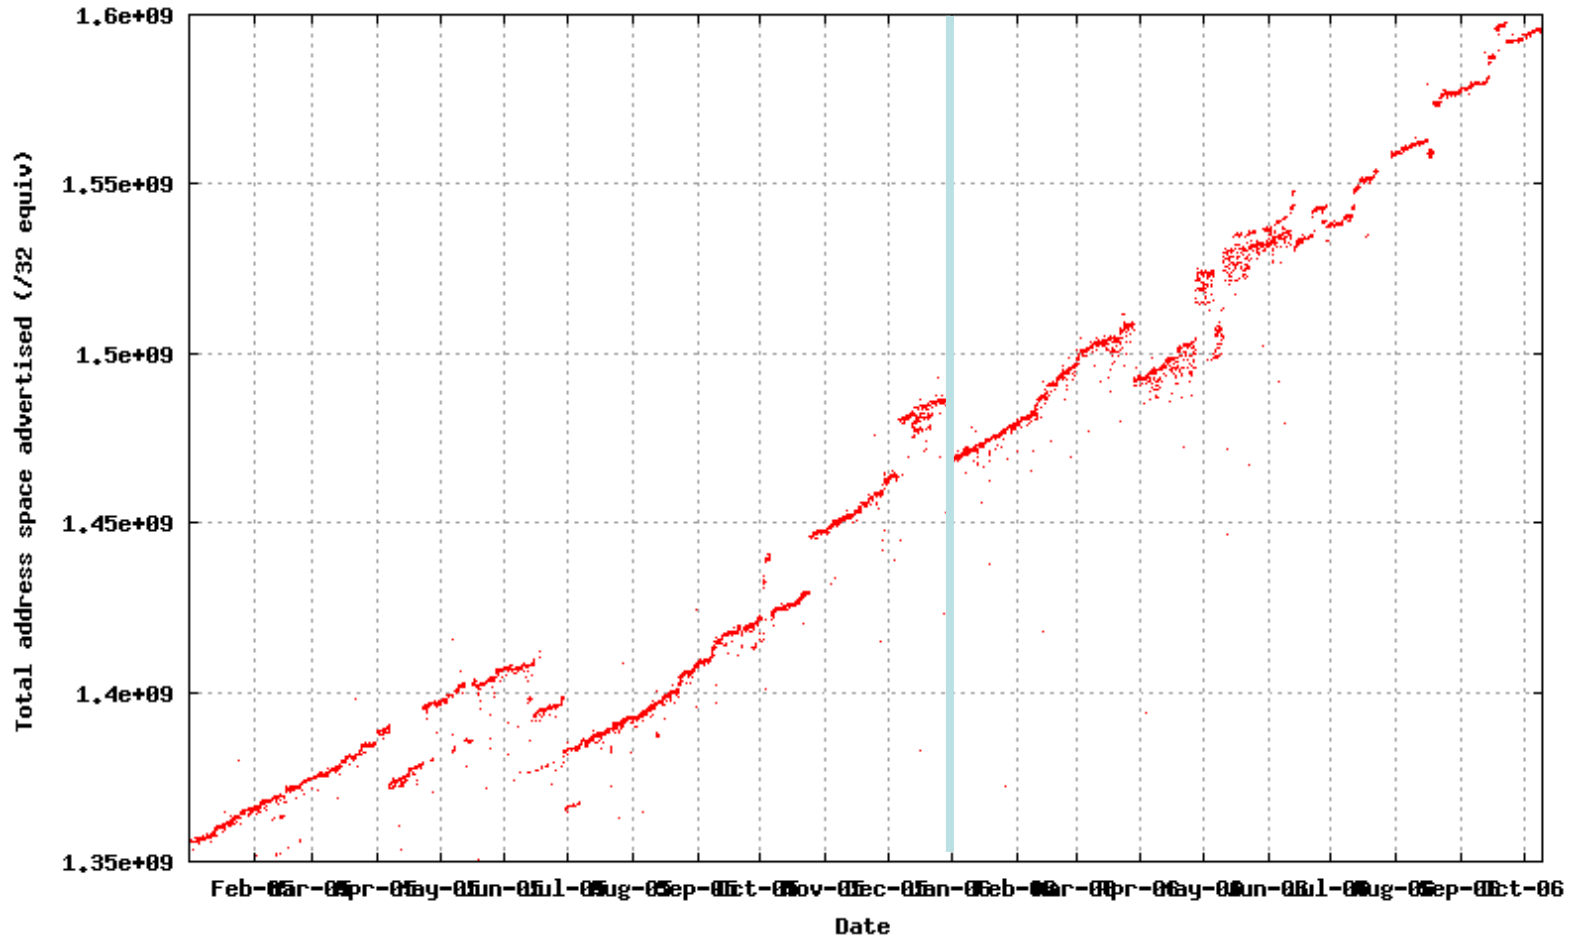
IPv4 in 2005/6

Total Advertised BGP Prefixes




IPv4 in 2005/6

Total Advertised Address Span



IPv4 in 2005/6

Total Advertised AS Numbers



IPv4 – Vital Statistics for 2005

Prefixes	148,500 – 175,400	+18%	26,900
Roots	72,600 – 85,500	+18%	12,900
Specifics	77,200 – 88,900	+18%	14,000
Addresses	80.6 – 88.9 (/8)	+10%	8.3 /8s
ASNs	18,600 – 21,300	+14%	2,600

Average advertisement size is getting smaller

Average address origination per AS is getting smaller

Average AS Path length steady at 3.5

AS transit out-degree degree up

The IPv4 network continues to get denser, with finer levels of advertisement granularity.

More interconnections, more specific advertisements

IPv4 – Vital Statistics for 2006 (est)

Prefixes	175,400 – 201,320	+15%	25,920
Roots	85,500 – 99,060	+16%	13,560
Specifics	88,900 – 102,260	+14%	12,360
Addresses	88.9 – 96.4 (/8)	+ 8%	7.5 /8s
ASNs	21,300 – 23,626	+11%	2,326

Average advertisement size is getting smaller (/18.8 to /19.0)

Average address origination per AS is getting smaller (/16.1 to /16.0)

Average AS Path length steady at 3.3 AS hops

AS transit connectivity out-degree steady at 4.4 AS peers *

The IPv4 network continues to get denser, with finer levels of advertisement granularity.


More interconnections, more specific advertisements

* Peer interconnections are invisible to this measurement system


So what?

Its not the numbers per se – it's the relentless consistent inflation of the routing system that these numbers represent that motivates consideration of where we are heading

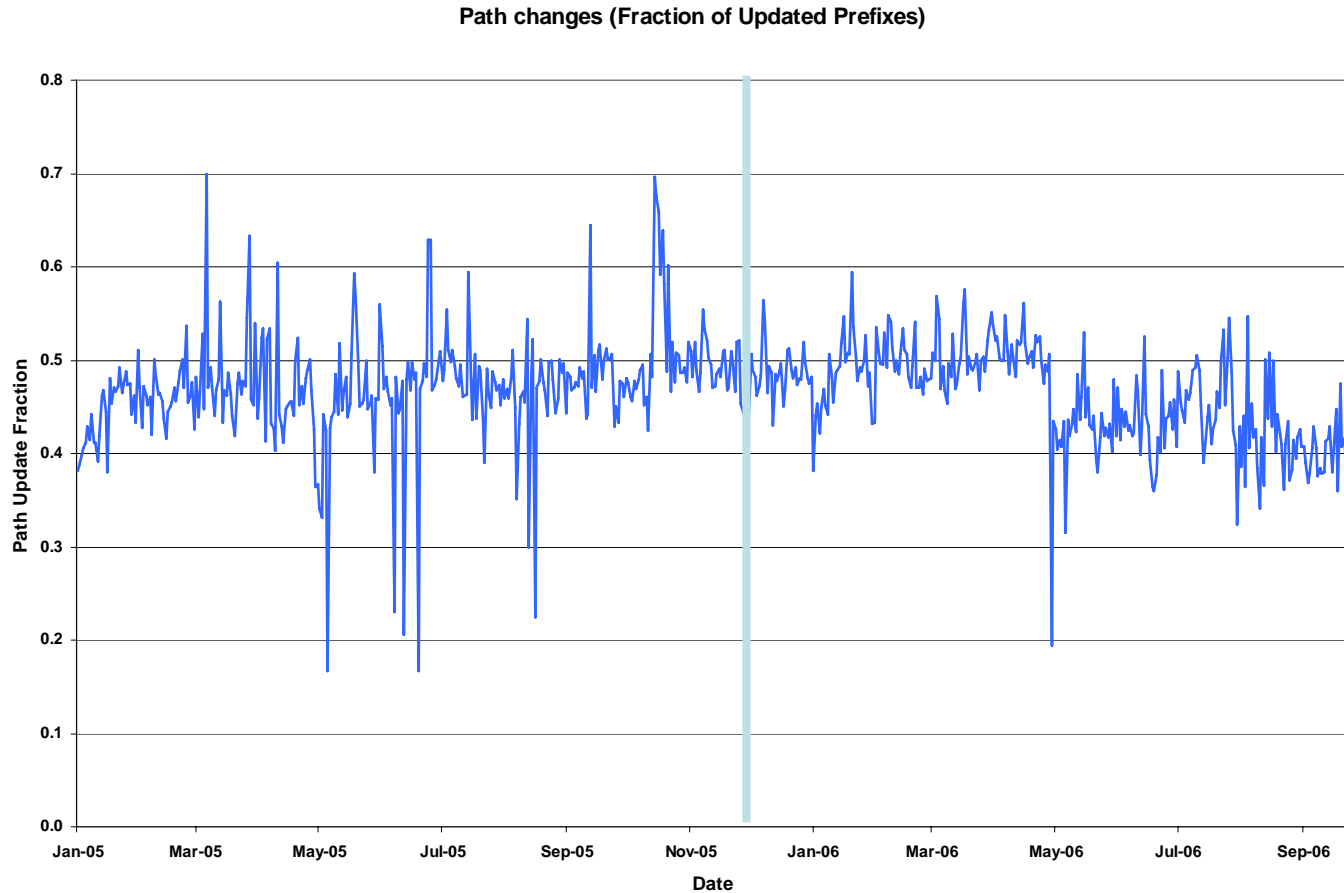
Daily Prefix Update Rate



Daily Prefix Withdrawal Rate





Daily Path Change Ratio



Update distribution is Heavy Tail

- 14 days ending 16 October 2006



1% of prefixes are associated with 20% of all prefix updates
1% of all origin ASes are associated with 40% of all updates

A couple of Intensive Updaters

- TTNET AS9121 (interview)
 - Mid-level aggregator (Turkish Telecom)
 - 5 upstream transits
 - 1 noisy transit SDH link with 600Mbps traffic
 - High convergence update load on link failure
- TELEKOMUNIKASI INDONESIA AS1794
 - Transit via AS7713
 - 4 upstream transits (3 are multi-homed mid-level)
 - Traffic engineering on de-aggregated /24 prefixes?
- Both are extensively multi-homed lower-tier transit providers!

So what's going on?

- It appears that the BGP update rate is being strongly biased by a small number of origins with two forms of behaviour:
 - **Traffic Engineering** - consistent update rates sustained over weeks / months with a strong component of first hop change and persistent announce and withdrawal of more specifics
 - **Unstable configuration states** – a configuration which cannot stabilise and for a period of hours or days the update rate is extremely intense

What's going astray?

- The combination of deaggregation, traffic engineering, unstable transit circuits and pruned Best Path BGP Distance Vector advertisements is causing massive update load in areas of high density BGP interconnection

Issue Identification

- Currently, BGP load generation is not an “everywhere, all at once” problem – most of the network (70% of prefix advertisements) is highly stable most of the time – as it should be
- For the others: ...

Issue Identification

- Reachability is combined with Best Path Selection
 - Traffic Engineering is performed through routing
- Some configurations are unstable (peer TE feedback loops?)
- There is no routing economy (no ‘natural’ pushback)
- There is inadequate base knowledge of routing at the transit edges
- There is no clear view of the impact of local actions from ‘the other side’
- There is no ‘natural’ dampening to mitigate routing inflation