



Tracking the Internet's BGP Table

Geoff Huston


Telstra

December 2000




Methodology

- The BGP table monitor uses a router at the boundary of AS1221 which has a default-free eBGP routing table
 1. Capture the output from “show ip bgp” every hour
 2. Perform analysis of the data
(and then discard the raw dump!)
 3. Update reports at www.telstra.net/ops/bgp




Phases of Growth




Growth Characteristics

- Short term route fluctuation is an absolute value (not a % of total routes) of 1,000 – 2,000 routes




Routed Address Space



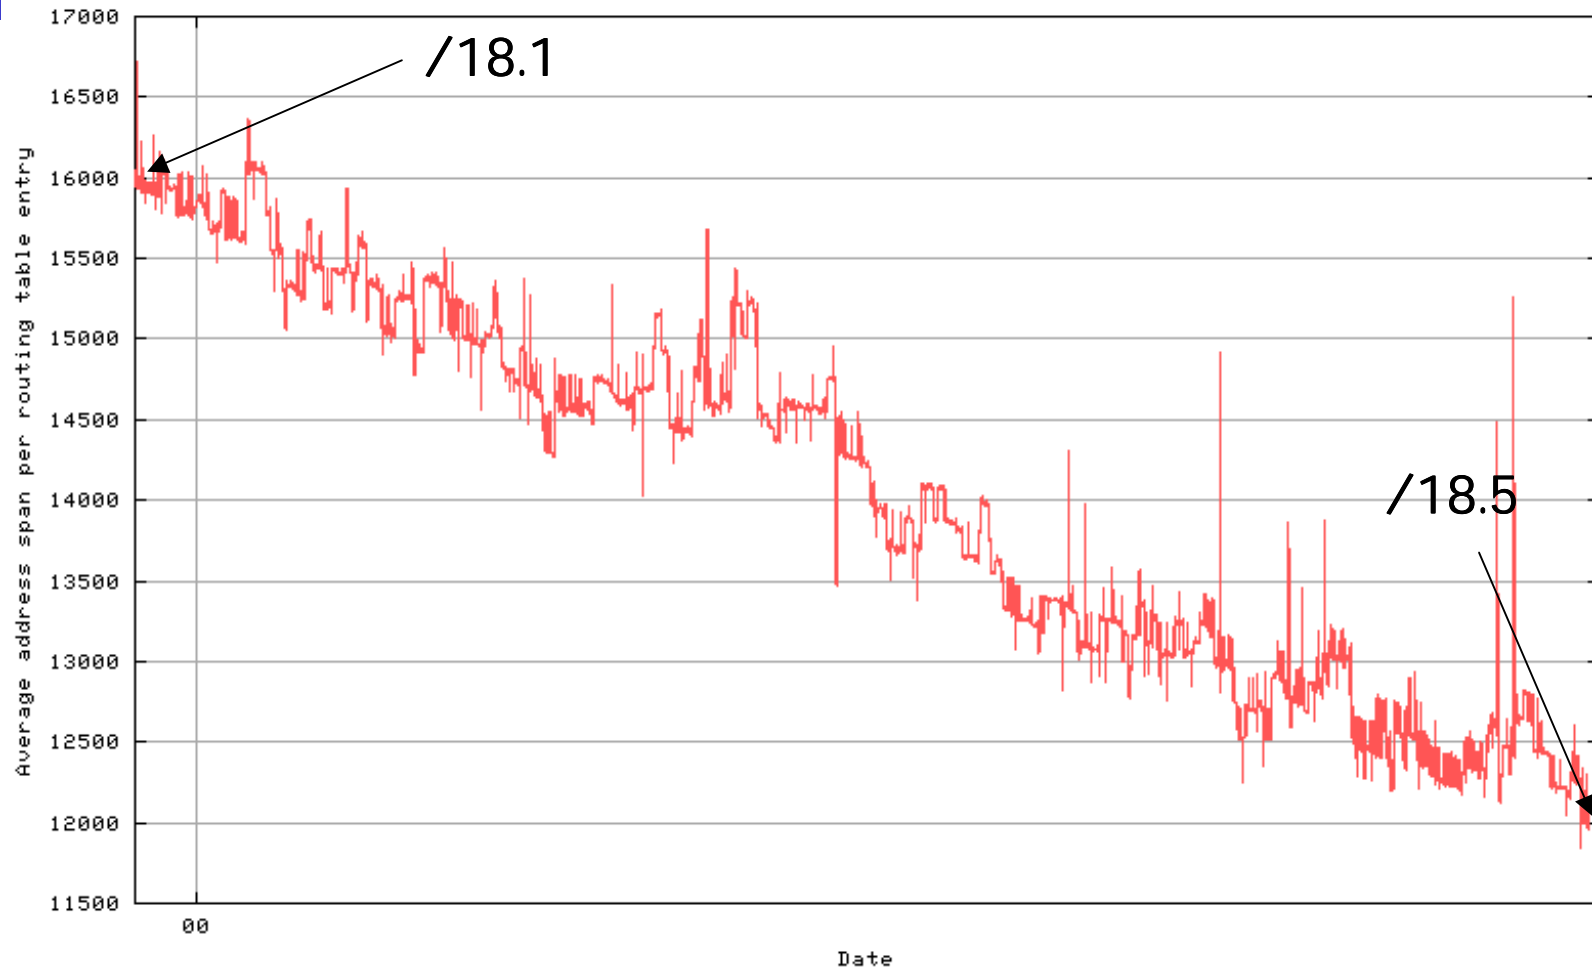
Fluctuation is due to announcement / withdrawals of /8 prefixes
11 months of data does not provide clear longer growth characteristic

Routed Address Space (/8 Corrected)

Annual compound growth rate is 7% p.a.
Most address consumption today appears to be occurring behind NATs




Average size of a routing table entry




The BGP routing table is growing at a faster rate than the rate of growth of announced address space

Number of AS's in the table




Exponential growth is evident in a longer term view of the AS deployment rate

AS Number Trend Models




Best fit model is an exponential model using 12 months of data

AS Number Use - Extrapolation



Number of distinct AS Paths






Observations for 99/00

- Low growth in the number of routed addresses
0.6% growth / month (7% / year)
- High growth in number of route advertisements
3% growth / month (42% / year)
- High growth in number of AS's
3.5% growth / month (51% / year)

Multi-homing on the rise?


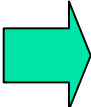
- Track rate of CIDR "holes" – currently 40% of all route advertisements are routing 'holes'



This graph tracks the number of address prefix advertisements which are part of an advertised larger address prefix







Prefix Growth – Aug 00 to Oct 00

/16	6553	->	6670	absolute growth =	117,	relative =	1.79%
/17	889	->	936	absolute growth =	47,	relative =	5.29%
/18	1763	->	1884	absolute growth =	121,	relative =	6.86%
/19	5704	->	5984	absolute growth =	280,	relative =	4.91%
	/20		3423 -> 3854	absolute growth =	431,	relative =	12.59%
/21	3621	->	3856	absolute growth =	235,	relative =	6.49%
/22	5415	->	5870	absolute growth =	455,	relative =	8.40%
/23	7298	->	7788	absolute growth =	490,	relative =	6.71%
/24	49169	->	52449	absolute growth =	3280,	relative =	6.67%
	/25		208 -> 436	absolute growth =	228,	relative =	109.62%
/26	334	->	606	absolute growth =	272,	relative =	81.44%
/27	469	->	667	absolute growth =	198,	relative =	42.22%
/28	357	->	452	absolute growth =	95,	relative =	26.61%
/29	579	->	764	absolute growth =	185,	relative =	31.95%
/30	746	->	1026	absolute growth =	280,	relative =	37.53%

The largest significant relative growth in recent times is /20, tracking the allocation policy change in the RIRs

While the absolute number is low, the largest relative growth is in /25 prefixes, and /25 to /30 represent the greatest area of prefix growth in relative terms








Conjectures....

- BGP table size will continue to rise exponentially
- Multi-homing at the edge of the Internet is on the increase
- The interconnectivity mesh is getting denser
 - The number of AS paths is increasing faster than the number of AS's
 - Average AS path length remains constant
- AS number deployment growth will exhaust 64K AS number space in August 2005 if current growth trends continue



More conjecturing....

- Inter-AS Traffic Engineering is being undertaken through routing discrete prefixes along different paths -- globally (the routing mallet!)
- RIR allocation policy (/19, /20) is driving one area of per-prefix length growth in the aggregated prefix area of the table
- BUT - NAT is a **very** common deployment tool
- NAT, multihoming and TE is driving even larger growth in the /24 prefix area



And while we are having such a good time conjecturing...

- Over 12 months average prefix length in the table has shifted from /18.1 to /18.5
- More noise (/25 and greater) in the table, but the absolute level of noise is low (so far)
- Most routing table flux is in the /24 to /32 prefix space – as this space gets relatively larger so will total routing table flux levels
 - “Flux” here is used to describe the cumulative result of the withdrawals and announcements



This is fun – lets have even more conjectures...

- CIDR worked effectively for four years, but its effective leverage has now **finished**
- Provider-based service aggregation hierarchies as a model of Internet deployment structure is more theoretic than real these days

i.e. provider based route aggregation is leaking like a sieve!