

Inter-Domain Routing Trends

Geoff Huston
gih@apnic.net

APNIC
March 2007

Agenda

1. Some data about the network and BGP over 2006
2. Some observations about the distribution patterns of BGP updates
3. Pointers to some possible areas of further study

IPv4 Stats for 2006

Prefixes	173,800 – 203,800	+17%
Roots	85,800 – 100,800	+17%
Specifics	88,000 – 103,000	+17%
Addresses	87.6 – 98.4 (/8)	+12%
ASNs	21,200 – 24,000	+13%

AS growth – 13%

BGP growth – 17%

Average advertisement size is getting smaller (8,450 – 8,100)

Average address origination per AS is getting smaller (69,600 – 69,150)

Average AS Path length steady (3.4)

AS transit interconnection degree rising (2.56 – 2.60)

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

More interconnections, more specific advertisements

BGP Stats for 2006

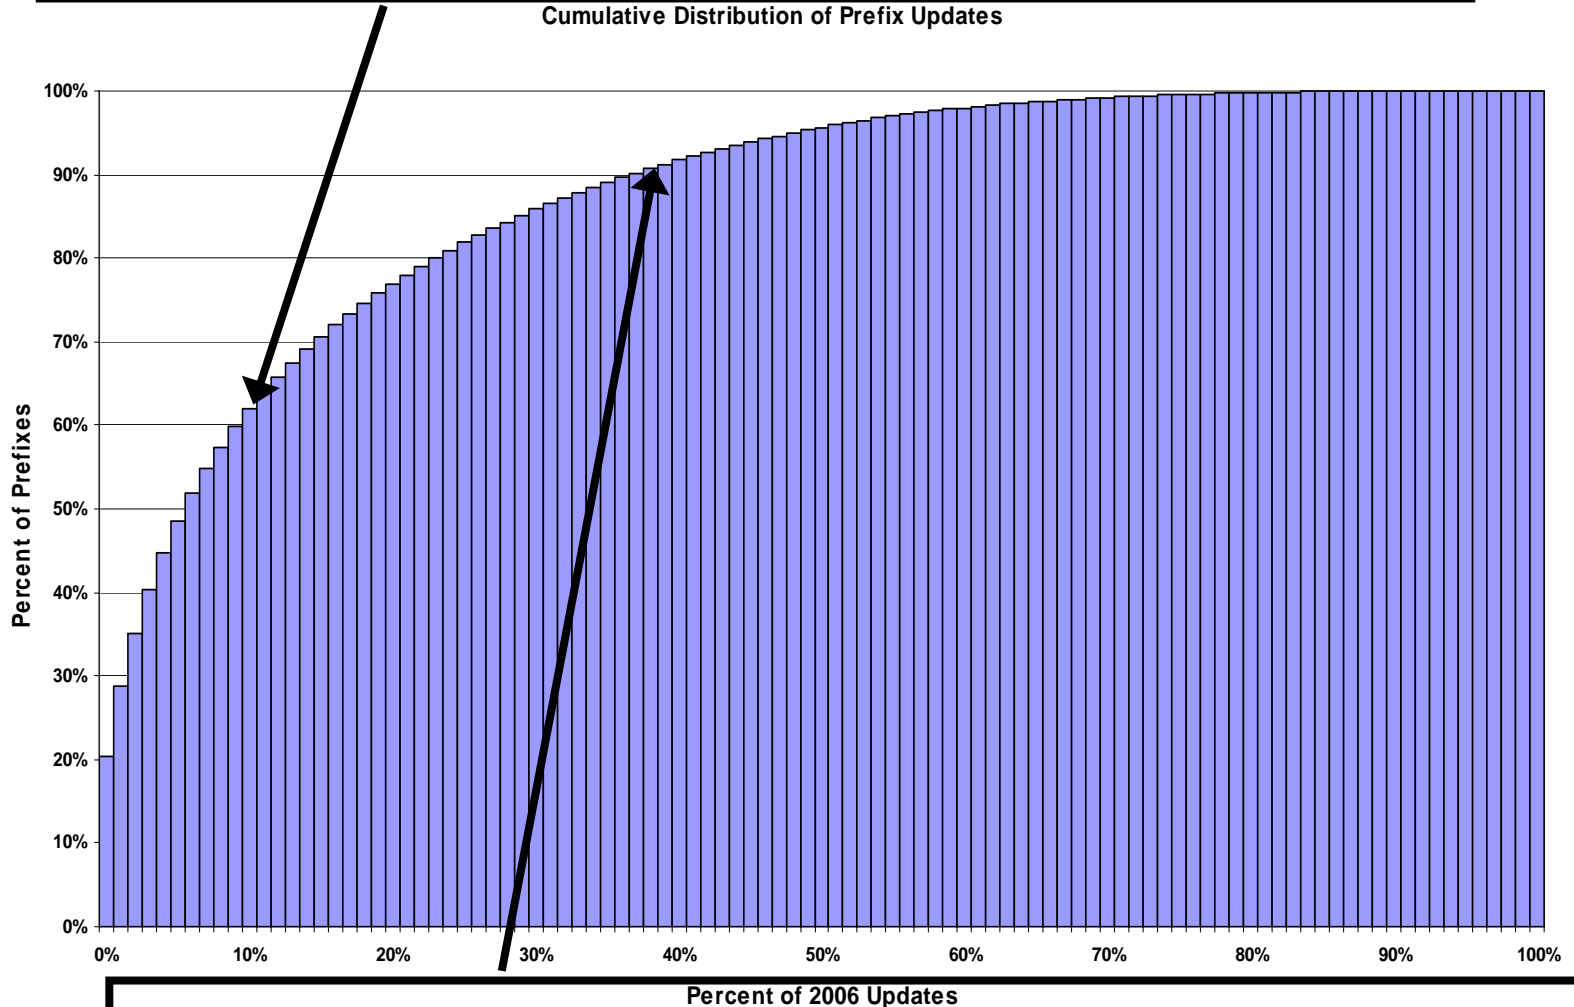
- Number of unique prefixes announced: 354,589
- Prefix Updates: 89,582,323 (average = 2.84 per second)
- Prefix Withdrawals: 30,531,219 (average = 0.96 per second)
- Updated prefixes (year end): 203,635
- Withdrawn prefixes: 150,954
- Average Prefixes per BGP Update: 1.95 (down from 2.1 at the start of 2006)

How “good” is this data?

- Its just one (ordinary) router’s view of a rather complex routing world, not an aggregated view of a larger routing environment. There is some ‘locality’ component in the data.
- Its not located the within the world’s richest connectivity (it may be understating the routing load)
- The data is very noisy (e.g. 150,000 short term (leaked?) prefixes)
- The data is heavily skewed by a ‘heavy tail’ distribution (small number of prefixes and ASs appear to be the subject of a large number of updates)
- So any effort at generating trend data is biased by the small number of these “intense updaters” (making predictive models even more uncertain than normal)

CDF of Updates by Prefix

10% of prefixes are the subject of 60% of updates for 2006



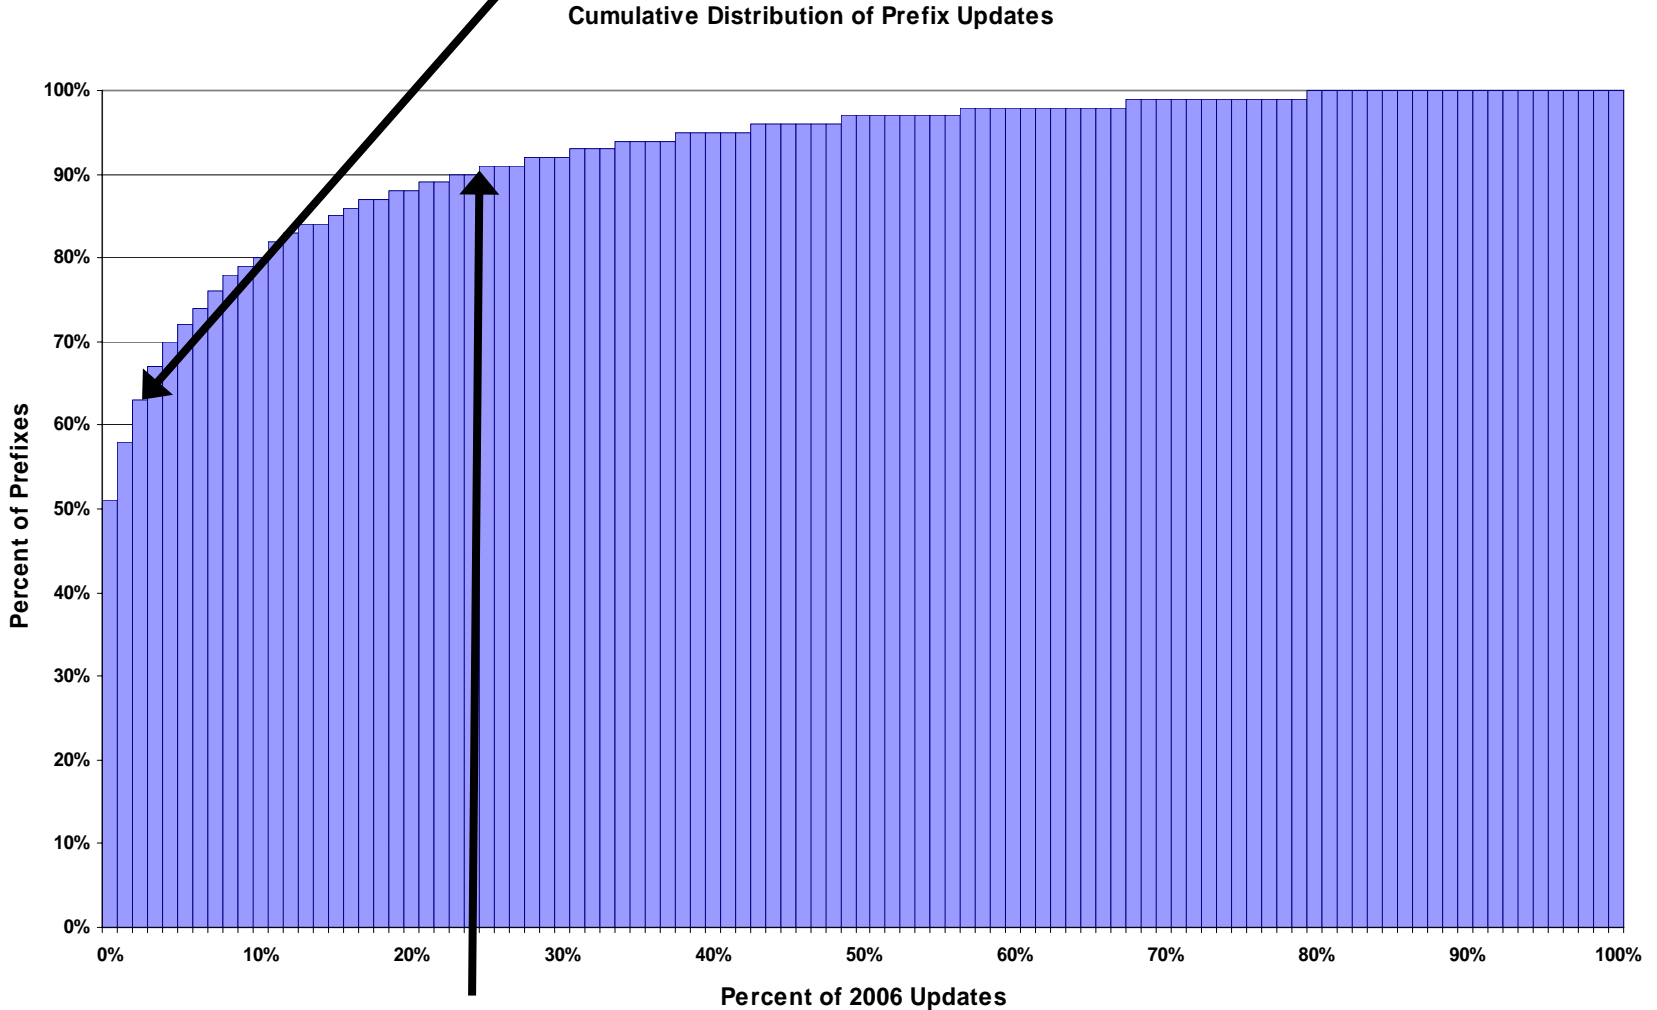
60% of announced prefixes accounted for 10% of the updates

84.205.76.0/24 – a known control point

- Beacon on a 1 hour cycle
 - 12 UP and 12 DOWN events per 24 hours
 - 4,380 beacon events in 2006 (4380 announces and 4380 withdrawals at origin)
 - Recorded 55,634 BGP update events and 4,423 withdrawals for 2006
 - Every withdrawal at origin caused an average of 11 update messages to reach the observation point throughout the year
 - BGP appears to be a very efficient event amplifier

CDF of Updates by Origin AS

3% of ASs were the origin AS for 60% of all updates



10% of Updates were associated with 75% of Origin ASs

Example: AS17974: 1.3M updates in 2006

- 17974 TELKOMNET-AS2-AP PT TELEKOMUNIKASI INDONESIA
Adjacency: Upstream: 1 Downstream: 0
- Upstream Adjacent AS list
 - AS7713 TELKOMNET-AS-AP PT TELEKOMUNIKASI INDONESIA
Upstream: 5
 - AS9237 HUTCHCA-AS Corporate Access (HK) Ltd.
 - AS11919 LORAL-SKYNET-AR - Loral Skynet Network Services, Inc.
 - AS24077 TMHK-TRANSIT-AS-HK-AP TMHK Global Transit
 - AS7473 SINGTEL-AS-AP Singapore Telecom
 - AS7632 MEGHANTARA-AS-AP PT. Meghantara

So what's going on?

- It would appear that the BGP update rate is being strongly biased by a small number of origins with two forms of behaviour:
 - Supernova
 - Multi-Homing & Traffic Engineering - bursting update rates sustained over weeks / months with a strong component of first hop change and persistent announce and withdrawal of more specifics
 - Background Radiation
 - Unstable configuration states – a configuration which cannot stabilise and for a period of hours or days where the convergence to withdrawal causes continual updates

Where is this heading?

- Can we make BGP “scale” better or are we forced to look at a new routing structure?
- Making BGP “scale”:
 - Is there a more effective mechanism for damping unstable routes and paths and /or damping convergence to withdrawal?
 - Can we encourage widespread use of standard mechanisms that limit the propagation of BGP advertisements?
 - Should we consider alternate ways of BGP coping with withdrawal?
 - Does the “origin withdrawal” attribute added to BGP protocol specification make sense?
 - Should we consider “alternate reachability” selective advertisements that address withdrawal / update patterns in BGP convergence by changing the BGP protocol behaviour?

Changing BGP

- It's now a large system with massive deployment inertia
- Any ‘change’ will require piecemeal deployment capability with benefits realized by those who deploy
 - Which implies that use of backward compatible incremental change with piecemeal deployment is the only feasible approach here
 - The 32-bit ASN transition is a useful case study in changing BGP
 - Capability negotiation for peer setup
 - Transitive opaque attributes to signal additional capabilities (such as origin withdrawal)
 - Local changes to BGP processing

Some themes for further study

- How well do we understand BGP today?
 - More observation points
 - Investigation of known BGP pathologies
 - Control points as observation benchmarks
- How well do we understand the BGP of tomorrow?
 - What metrics provide reasonable indicators?
 - How stable is the time series data?
 - What is the confidence interval of 3 – 5 year predictors?
- How well do we understand the impacts of incremental change to BGP?
 - Modelling connectivity and behaviours
 - Simulation and direct experimentation

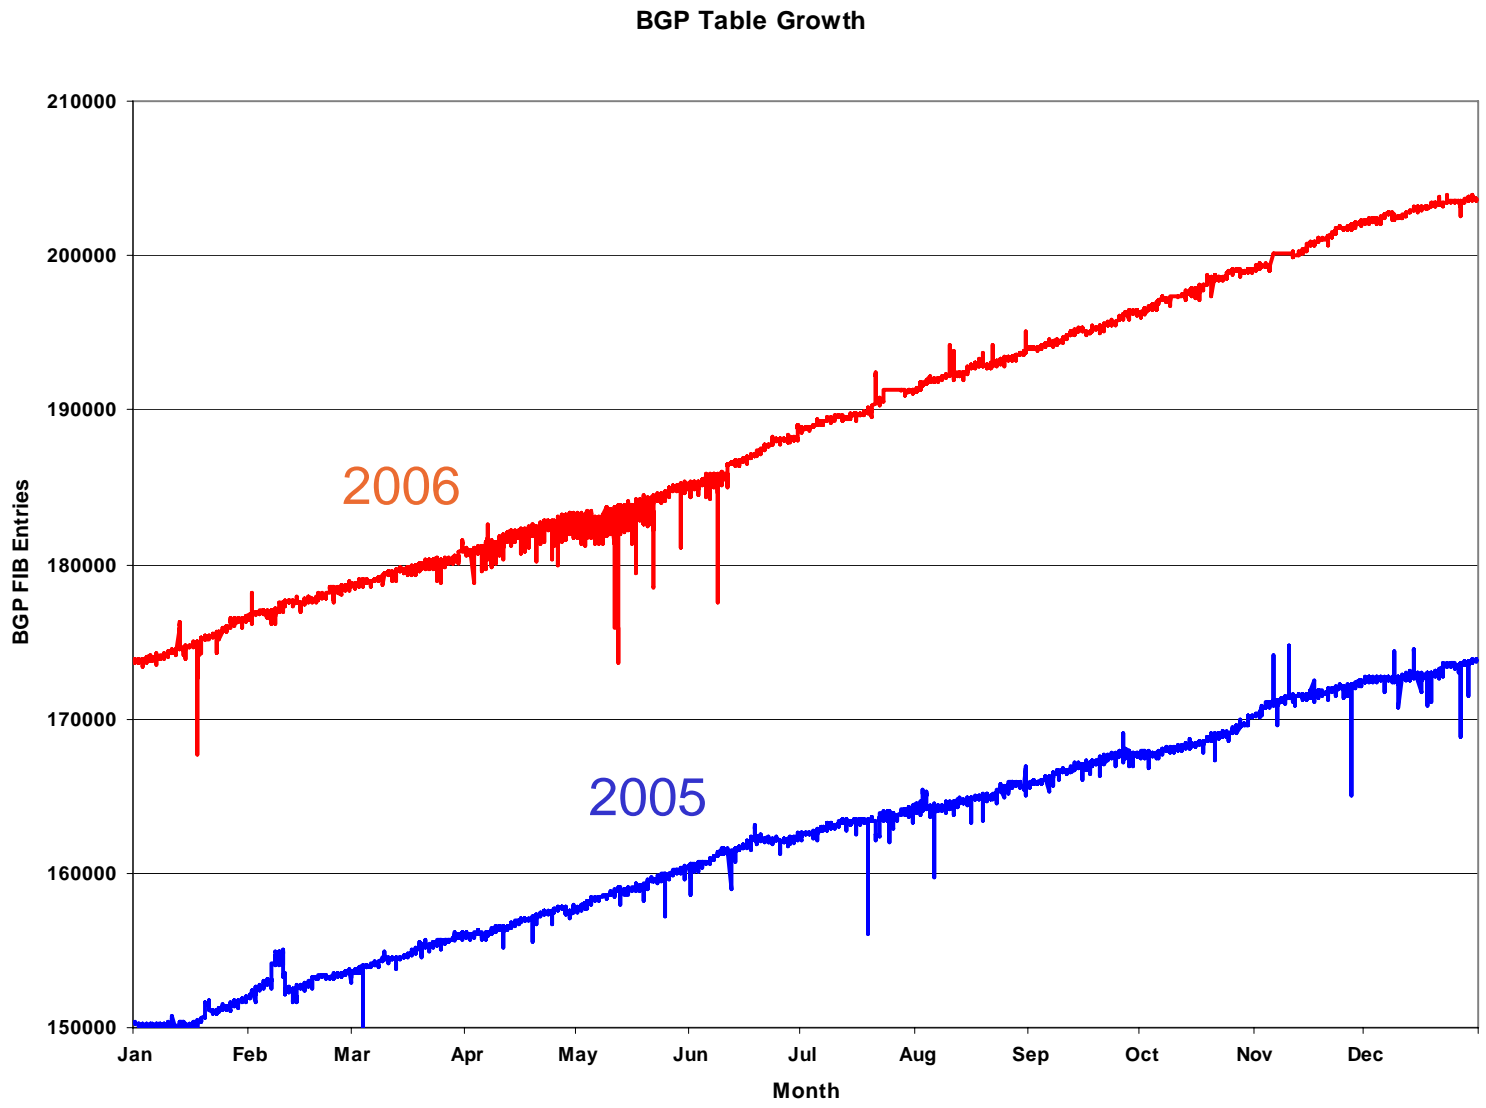
Thank You

Additional Material

The following are some graphs of aspects of BGP activity over 2006, with comparisons to comparable 2005 data

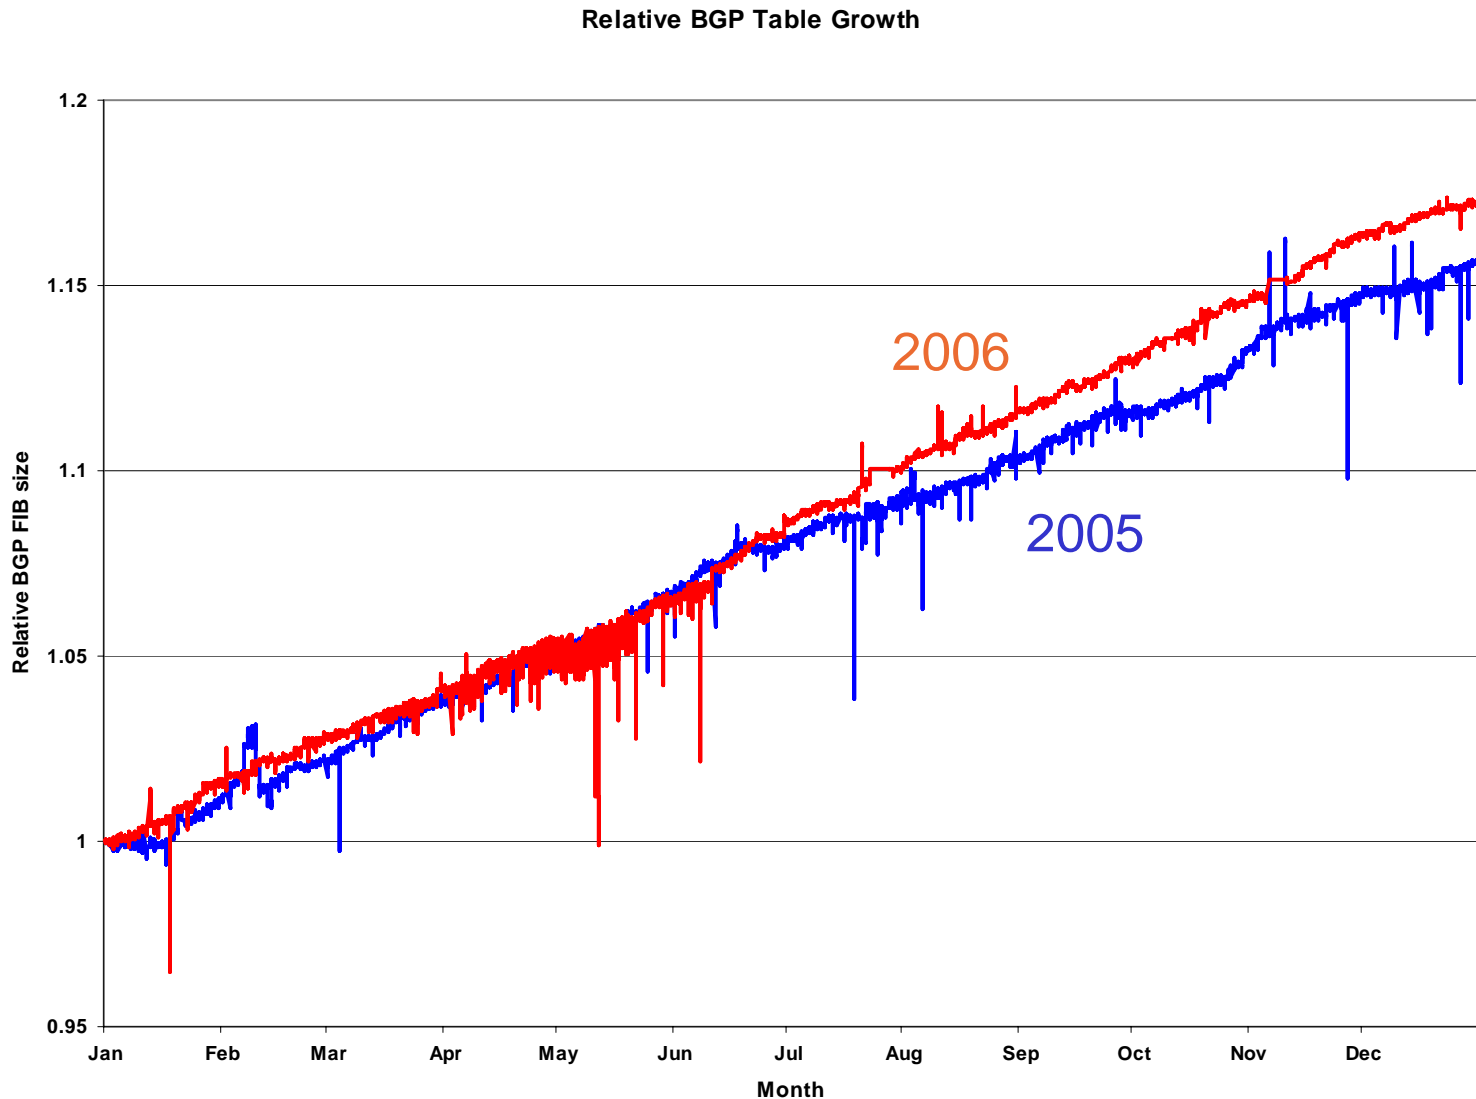
IPv4 in pictures

Total Advertised BGP Prefixes



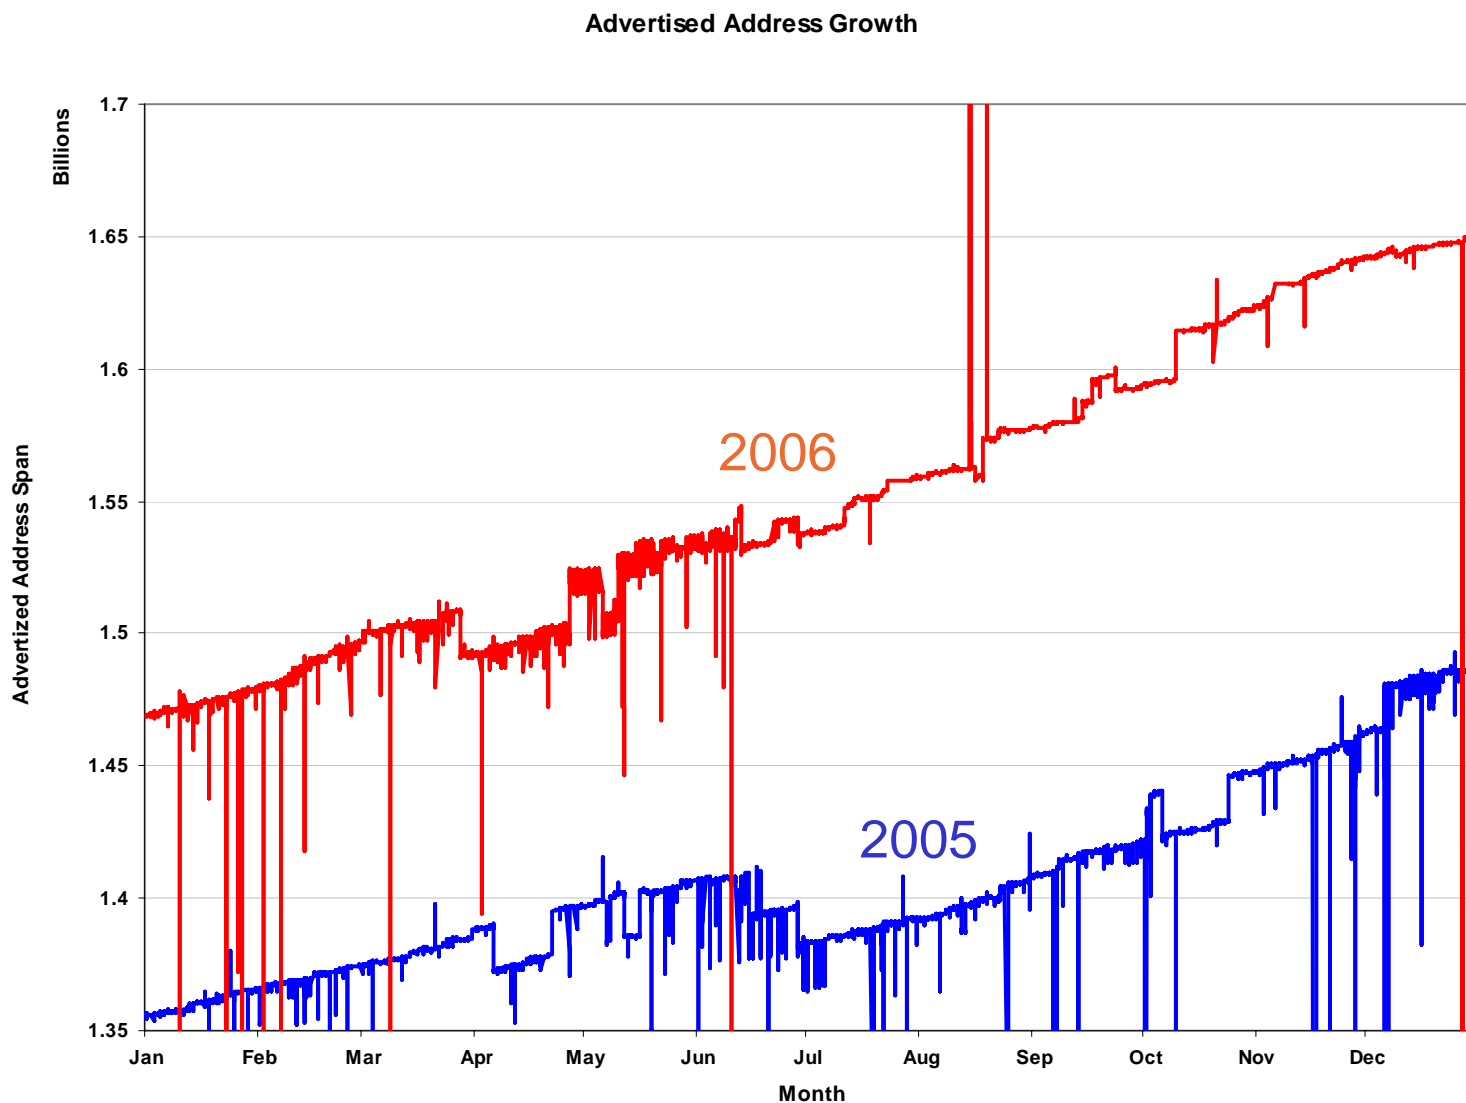
IPv4 in pictures

Relative Growth: 2005 to 2006



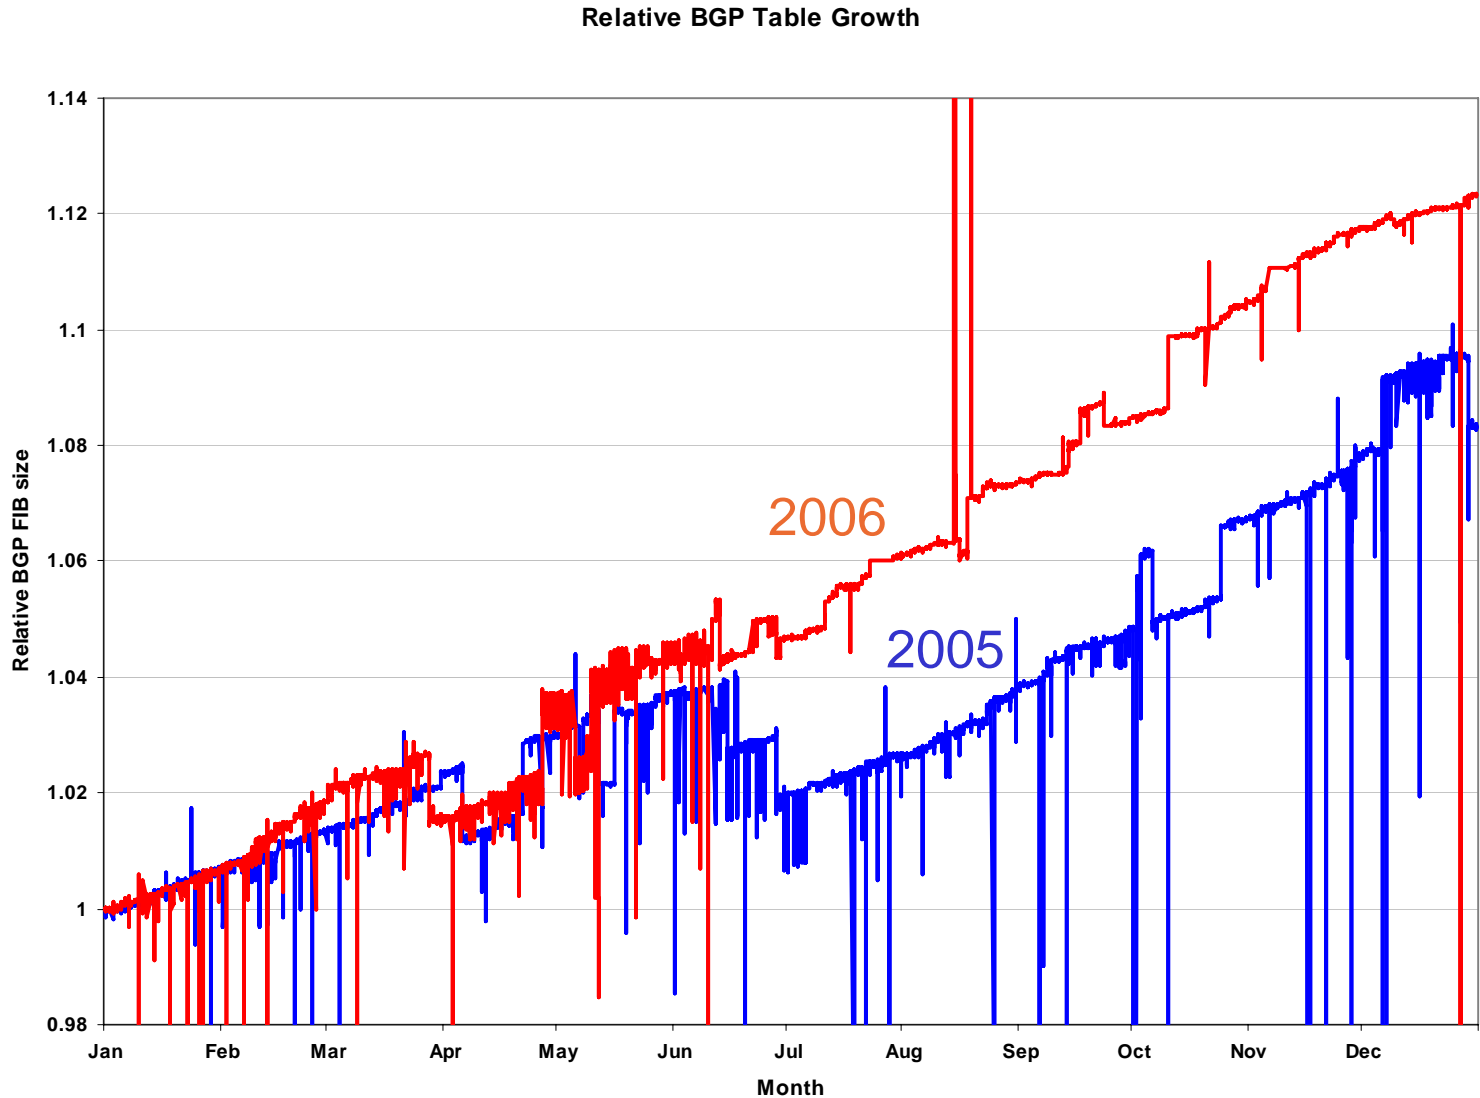
IPv4 in pictures

Total Advertised IPv4 Address Span



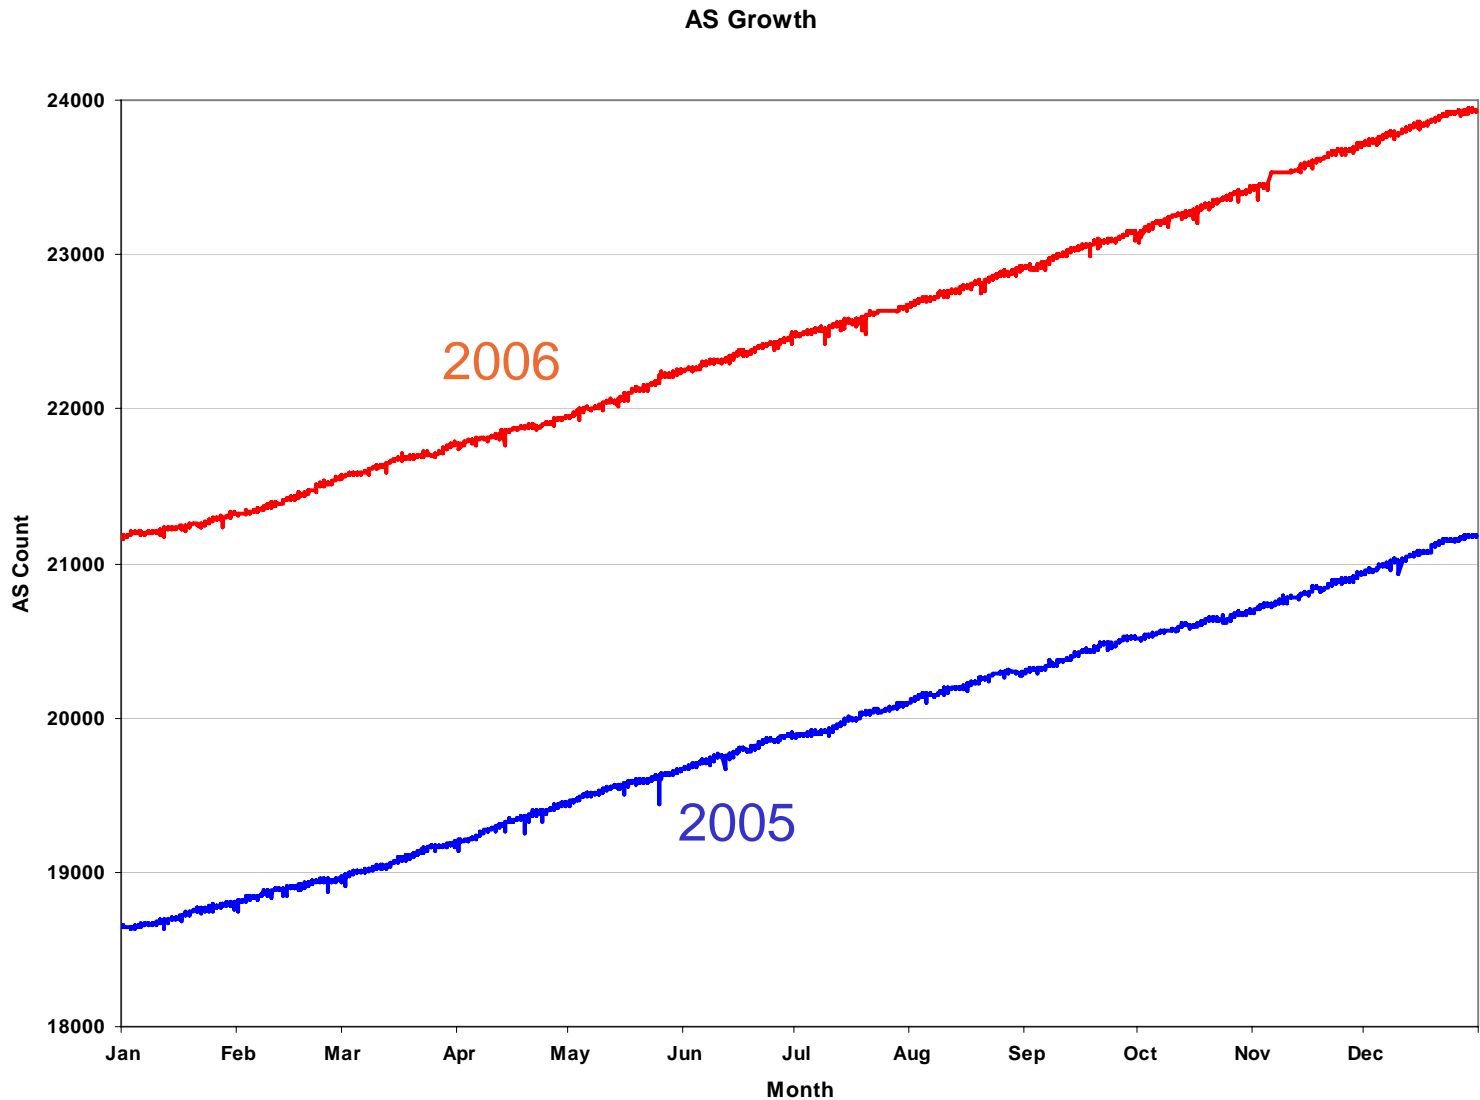
IPv4 in pictures

Relative Growth: 2005 to 2006



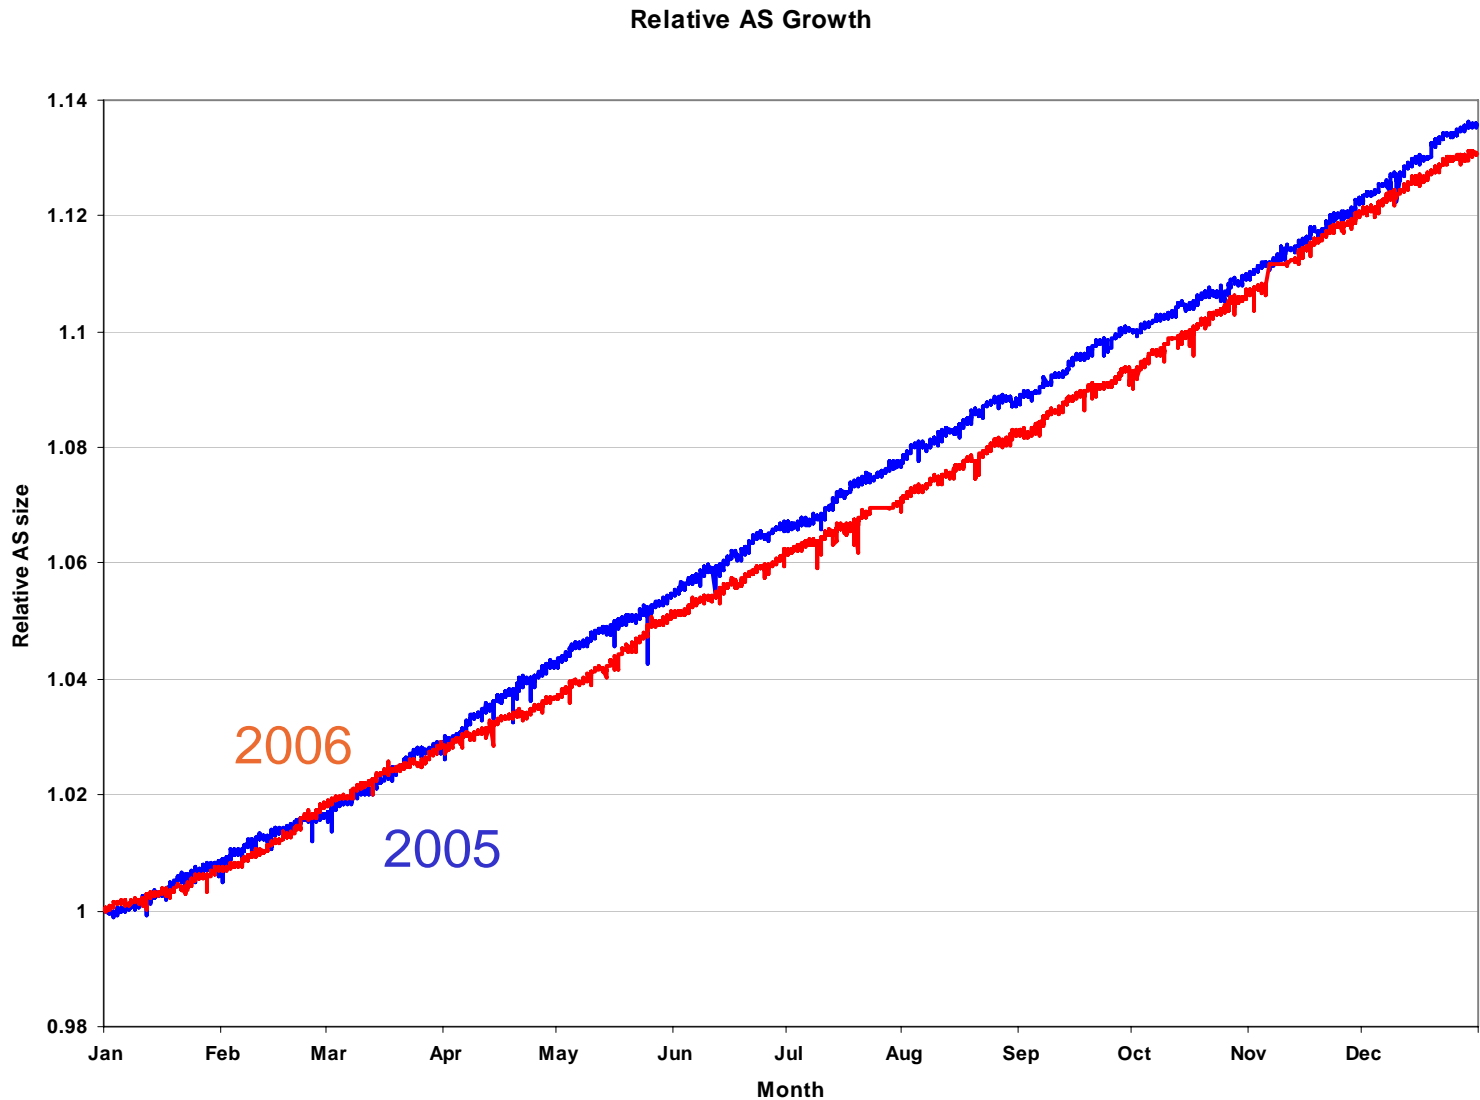
IPv4 in pictures

Total Advertised AS Numbers

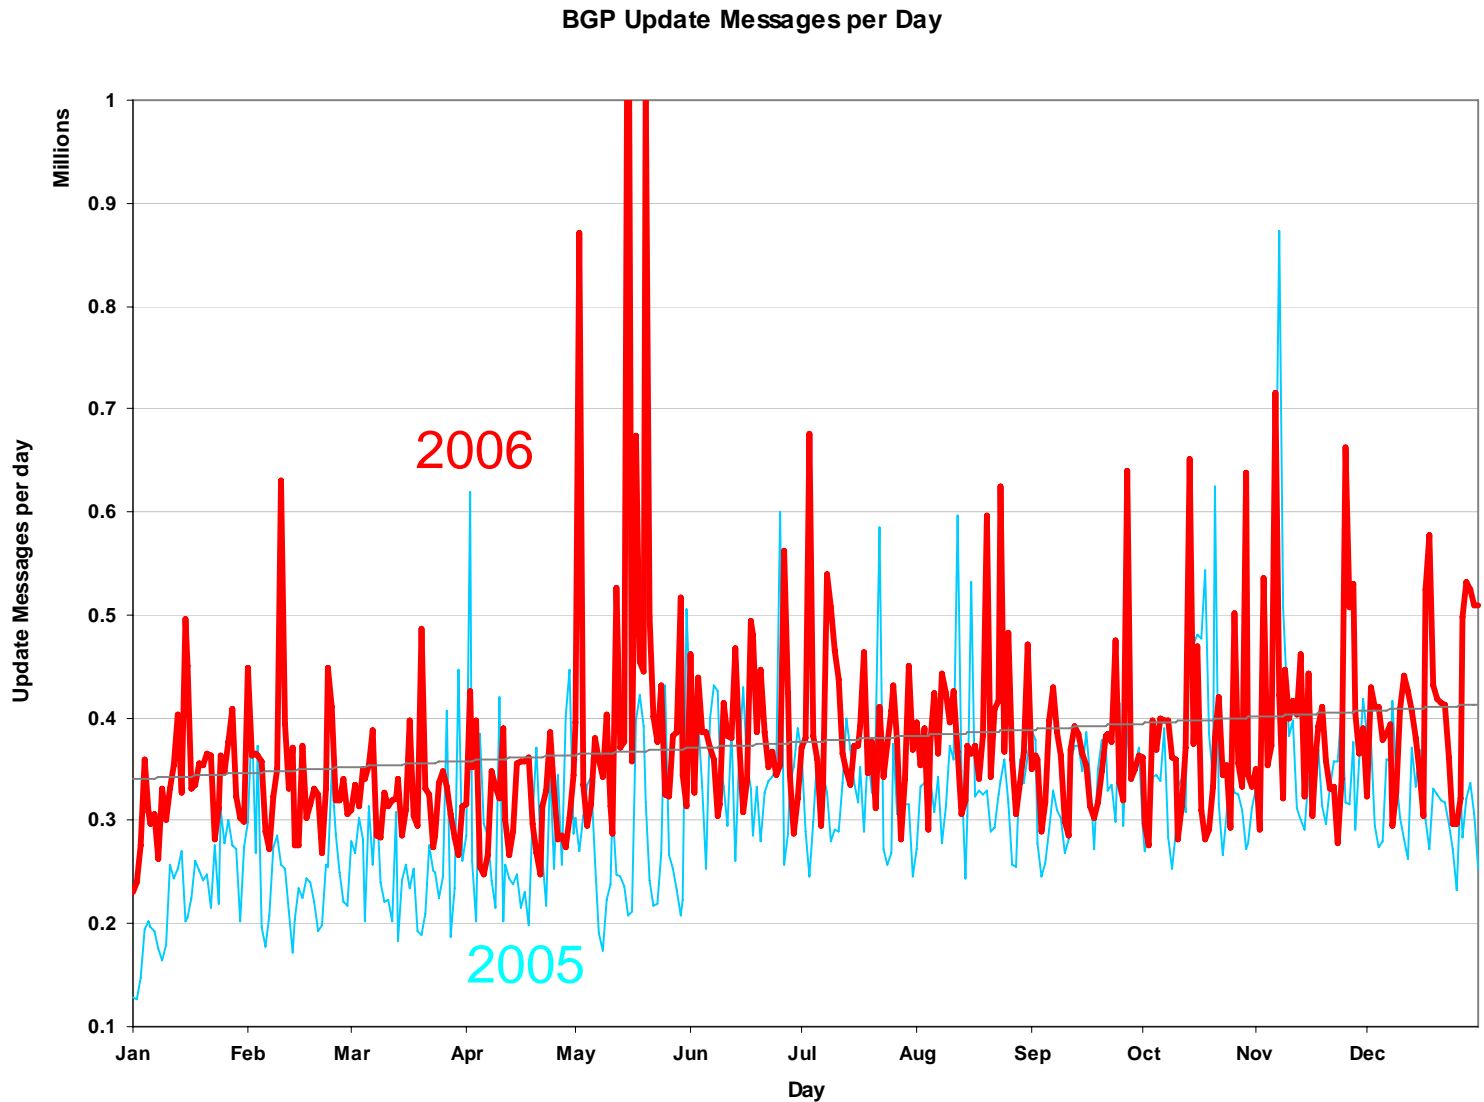


IPv4 in pictures

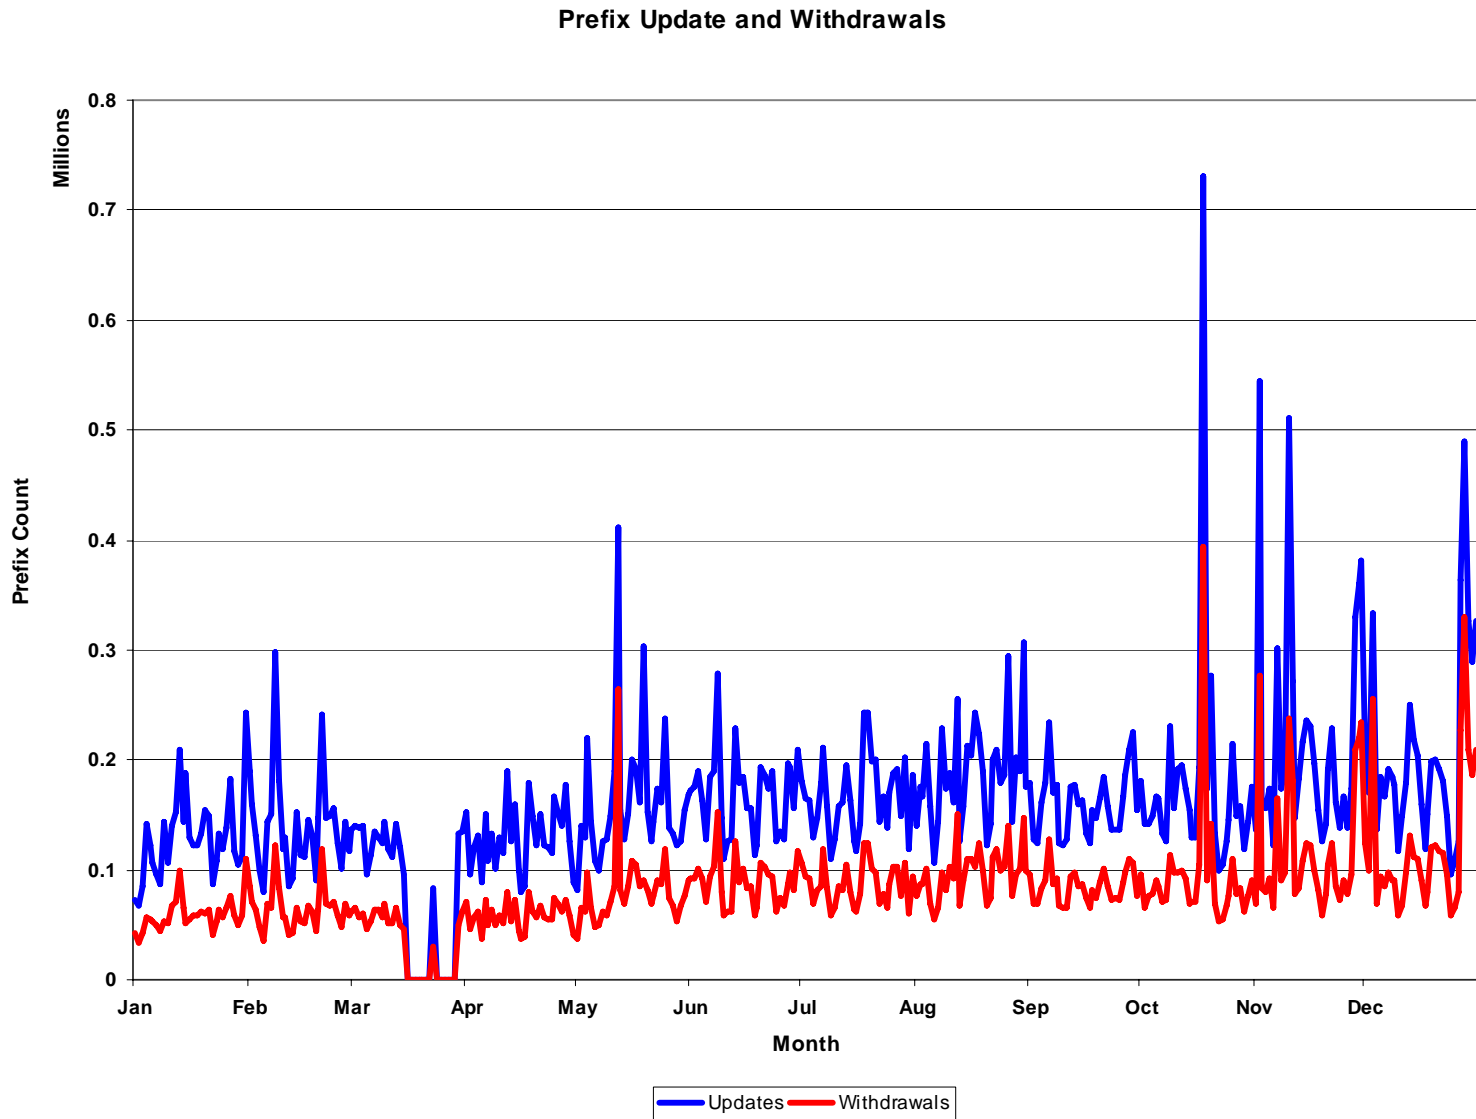
Relative Growth: 2005 to 2006



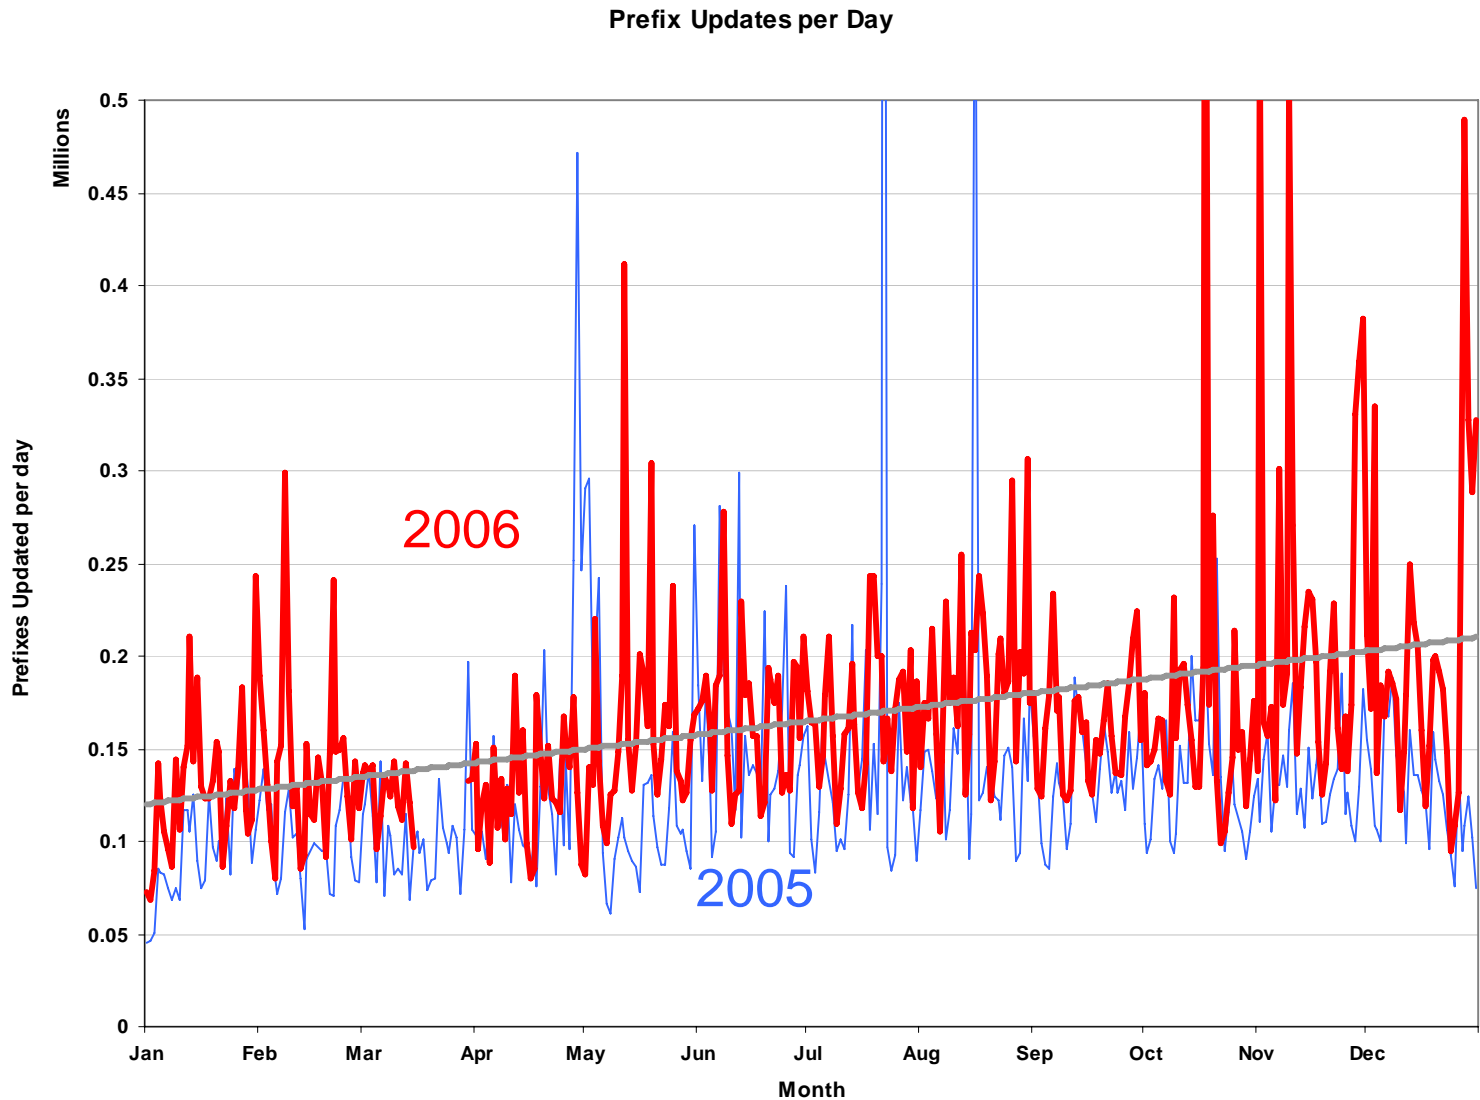
Update Message Rate



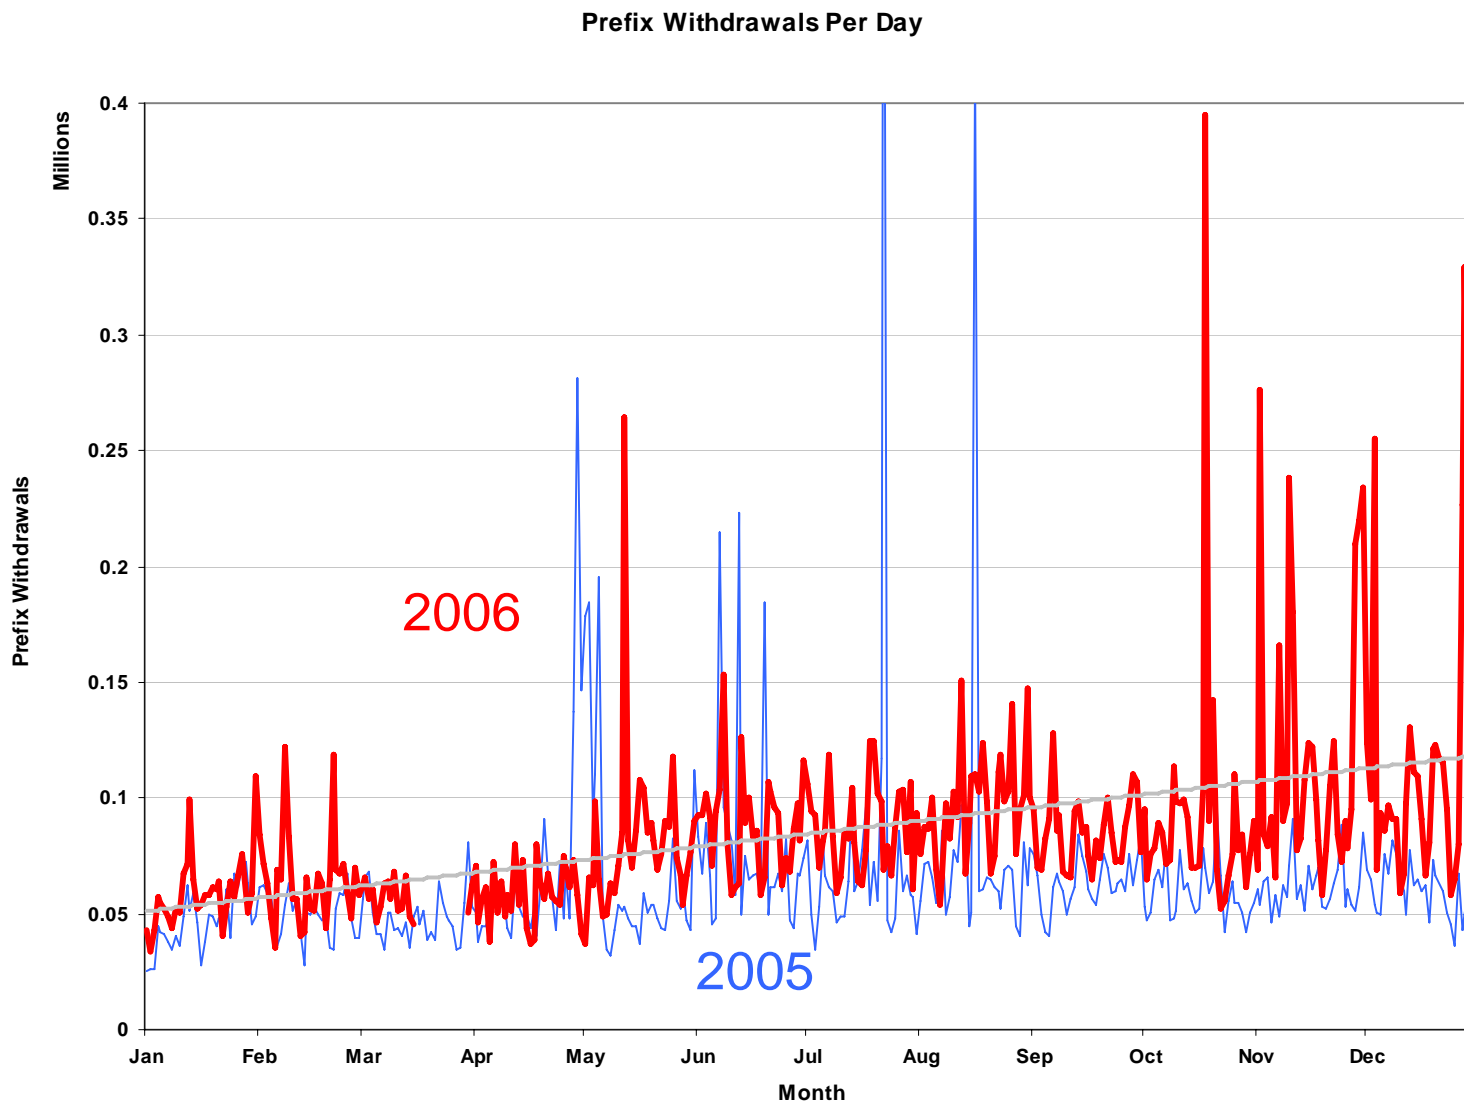
Prefix Update and Withdrawal Rates



Prefix Update Rates



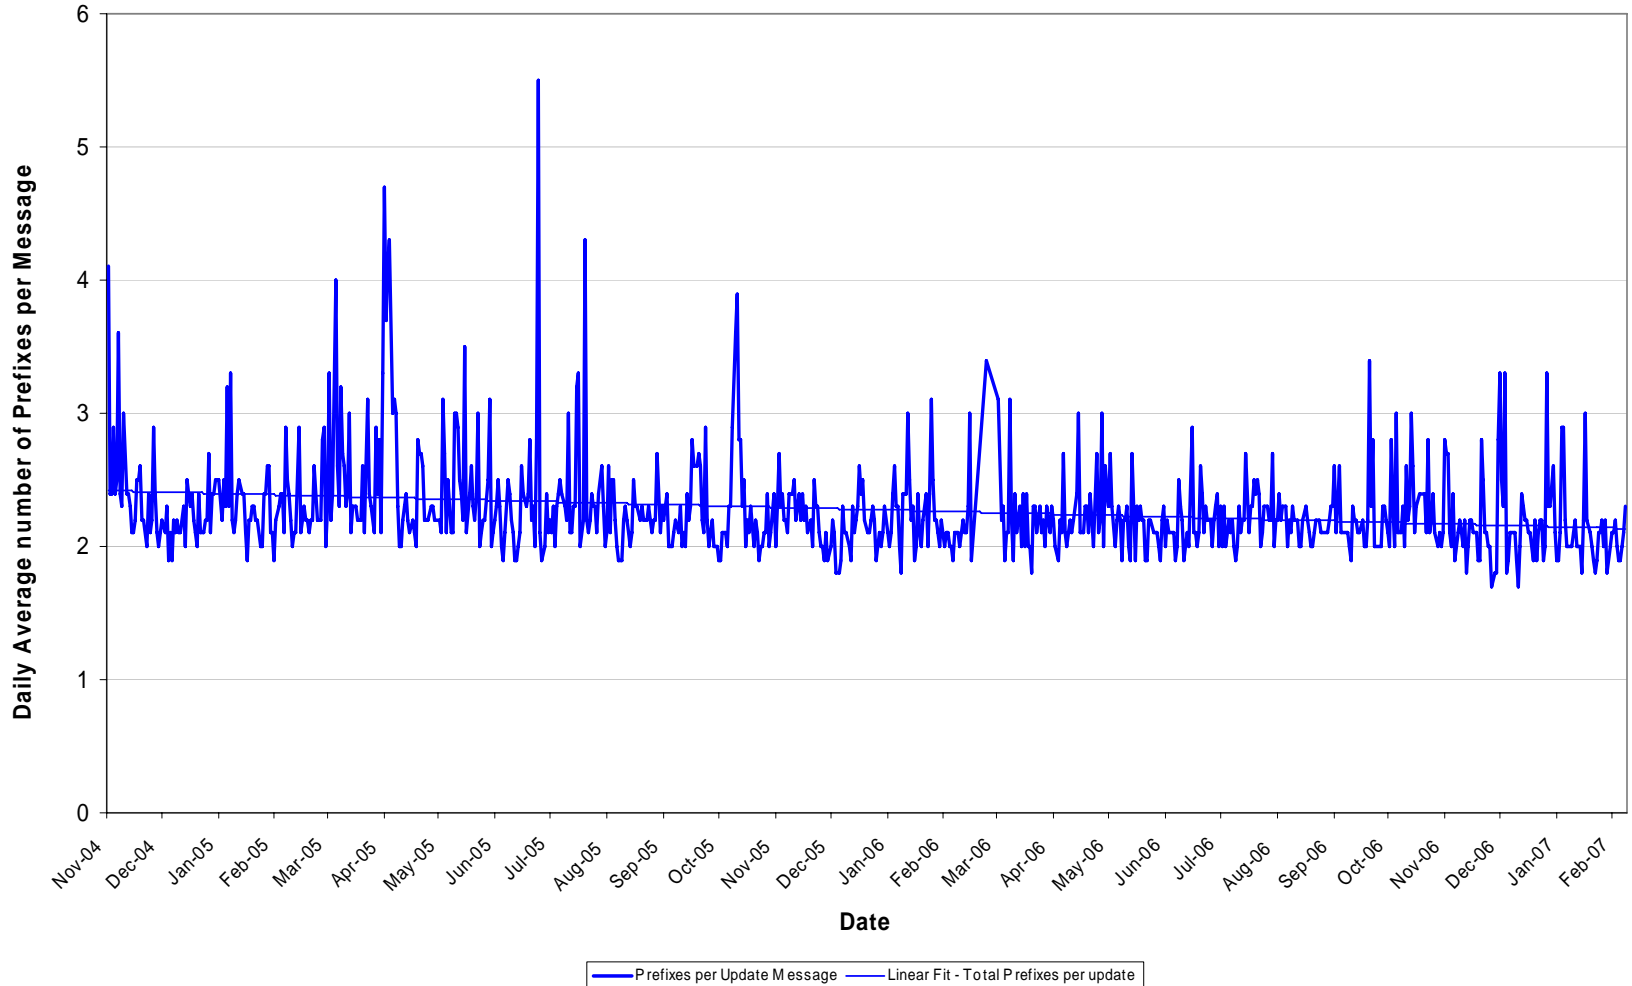
Withdrawal Rates



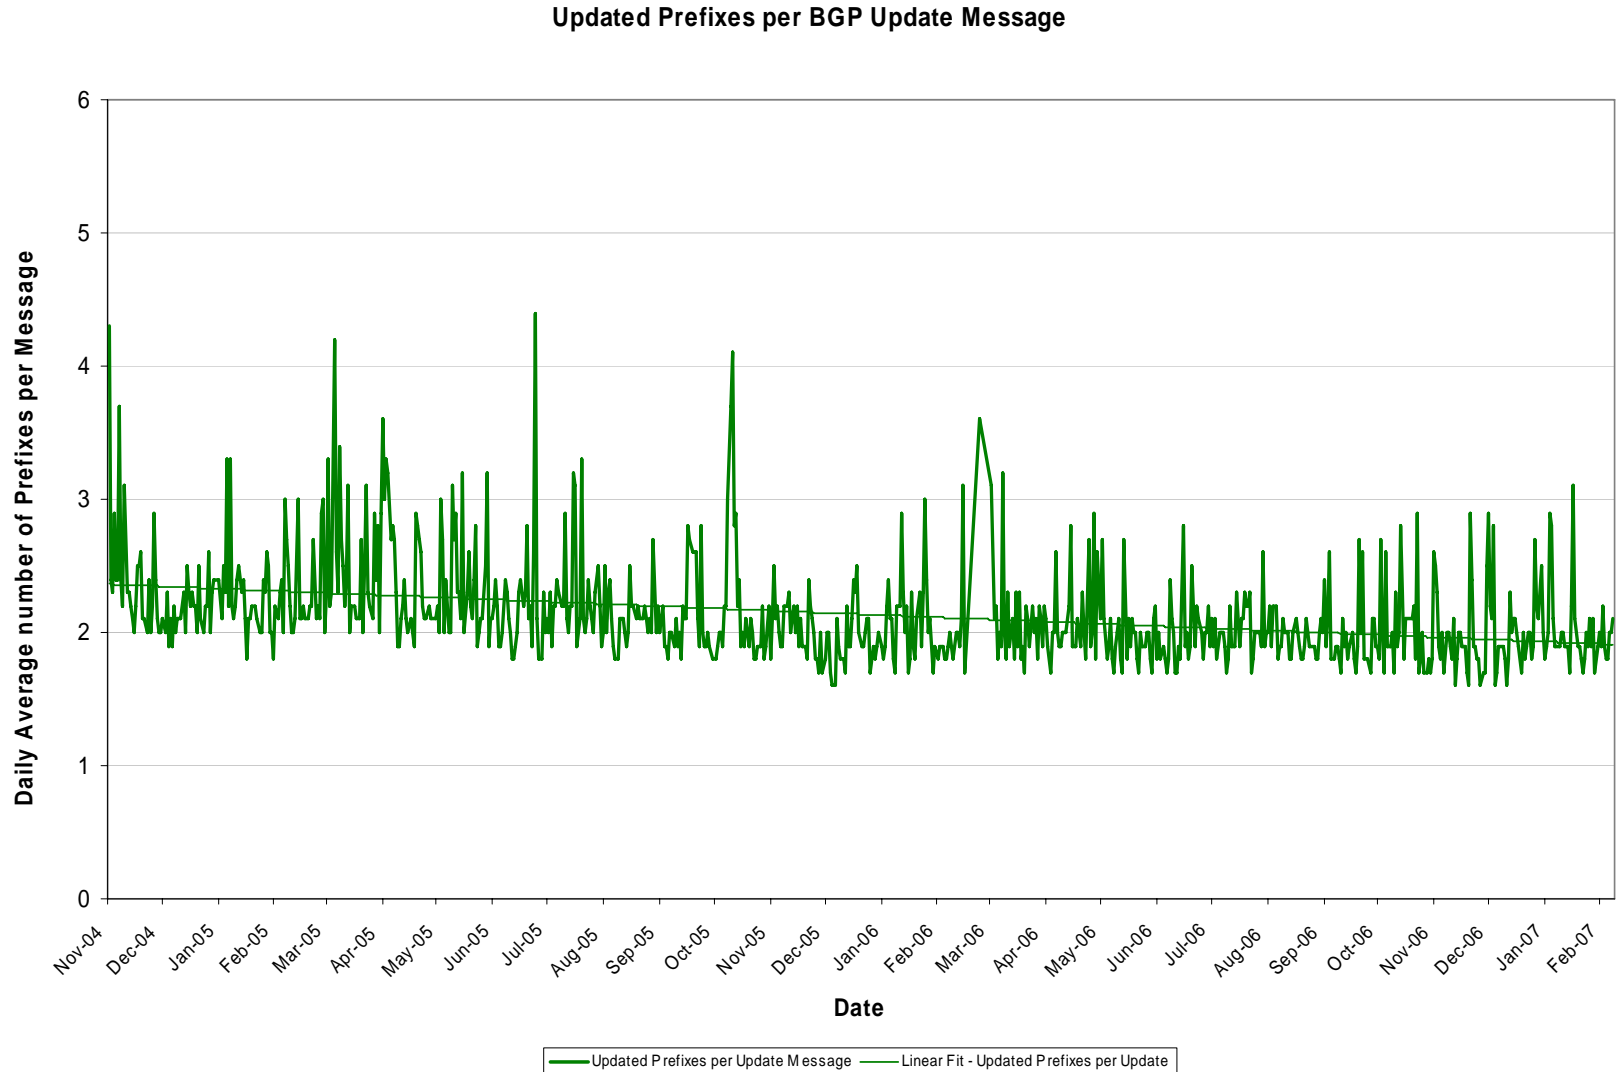
Average Prefixes per BGP Update



Prefixes per BGP Update Message



Updated Prefixes per BGP Update Message



Most Active Prefixes for 2006

Top 10 Prefixes

	<u>Prefix</u>	<u>Updates</u>	<u>Flaps</u>	<u>Re-Homes</u>
1.	209.140.24.0/24	210,574	151,145	1
2.	61.4.0.0/19	101,901	93,843	35
3.	61.0.0.0/8	89,768	70,863	5,541
4.	81.212.141.0/24	69,688	53,445	12,715
5.	203.199.128.0/19	63,606	51,076	8,592
6.	152.74.0.0/16	61,409	45,532	0
7.	84.205.65.0/24	59,744	44,792	8,454
8.	81.212.149.0/24	59,150	49,159	8,575
9.	193.242.123.0/24	57,717	34,974	16,468
10.	84.205.76.0/24	55,634	41,526	9,110

Active ASNs

Top 10 AS

	<u>AS</u>	<u>Updates</u>	<u>Flaps</u>	<u>Re-Homes</u>
1.	17974	1,340,344	983,667	1,819
2.	9121	783,879	542,965	199,409
3.	855	748,611	489,035	8,484
4.	702	517,723	379,880	96,121
5.	15611	517,243	337,669	2,556
6.	8151	425,852	288,042	33,666
7.	12654	396,924	295,083	49,567
8.	4323	393,687	275,477	130,056
9.	4621	370,478	278,650	21,116
10.	17557	368,689	248,680	121,211