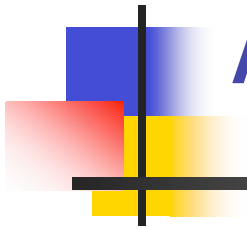

Internet Routing Table Analysis Update

Philip Smith

AfNOG 2009

10th-21st May

Cairo

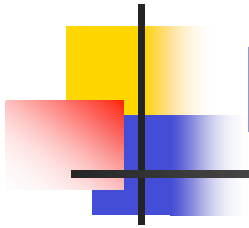

Motivation

- 1998: No one was publishing any Internet routing table analysis
  - Only CIDR-Report reporting on top 20 contributors to routing table, and top 20 bad aggregators
- With support of APNIC, my weekly reporting report started 23rd February 1999:
  - Routing table size
  - CIDR-Report style reporting on a per-RIR basis
  - …and many other interesting features

Routing Report 19 May 2009

BGP routing table entries examined:	288227
Prefixes after maximum aggregation:	136604
Deaggregation factor:	2.11
Unique aggregates announced to Internet:	142283
Total ASes present in the Internet Routing Table:	31269
Prefixes per ASN:	9.22
Origin-only ASes present in the Internet Routing Table:	27184
Origin ASes announcing only one prefix:	13262
Transit ASes present in the Internet Routing Table:	4085
Transit-only ASes present in the Internet Routing Table:	93
Average AS path length visible in the Internet Routing Table:	3.6
Max AS path length visible:	33
Max AS path prepend of ASN (43683)	31
Prefixes from unregistered ASNs in the Routing Table:	460
Unregistered ASNs in the Routing Table:	149
Number of 32-bit ASNs allocated by the RIRs:	146
Prefixes from 32-bit ASNs in the Routing Table:	35
Special use prefixes present in the Routing Table:	0
Prefixes being announced from unallocated address space:	214
Number of addresses announced to Internet:	2046363504
Equivalent to 121 /8s, 249 /16s and 7 /24s	
Percentage of available address space announced:	55.2
Percentage of allocated address space announced:	63.9
Percentage of available address space allocated:	86.4
Percentage of address space in use by end-sites:	77.1
Total number of prefixes smaller than registry allocations:	142300

AfriNIC Region

Prefixes being announced by AfriNIC Region ASes:	5017
Total AfriNIC prefixes after maximum aggregation:	1461
AfriNIC Deaggregation factor:	3.43
Prefixes being announced from the AfriNIC address blocks:	4616
Unique aggregates announced from the AfriNIC address blocks:	1408
AfriNIC Region origin ASes present in the Internet Routing Table:	299
AfriNIC Prefixes per ASN:	15.44
AfriNIC Region origin ASes announcing only one prefix:	92
AfriNIC Region transit ASes present in the Internet Routing Table:	60
Average AfriNIC Region AS path length visible:	3.9
Max AfriNIC Region AS path length visible:	15
Number of AfriNIC addresses announced to Internet:	11367680
Equivalent to 0 /8s, 173 /16s and 117 /24s	
Percentage of available AfriNIC address space announced:	33.9

AfriNIC AS Blocks 36864-37887 & ERX transfers
AfriNIC Address Blocks 41/8, 197/8,

Global per AS prefix count summary

ASN	No of nets	/20 equiv	Max Agg	Description
6389	4297	3647	324	bellsouth.net, inc.
209	2558	4149	607	Qwest
4323	1848	1035	373	Time Warner Telecom
1785	1760	717	139	PaeTec Communications, Inc.
4766	1701	6930	405	Korea Telecom (KIX)
20115	1629	1444	720	Charter Communications
17488	1584	128	104	Hathway IP Over Cable Interne
7018	1492	5903	1026	AT&T WorldNet Services
8151	1447	2862	232	UniNet S.A. de C.V.
6478	1404	310	400	AT&T Worldnet Services
2386	1266	682	920	AT&T Data Communications Serv
4755	1248	464	126	TATA Communications formerly
3356	1224	10981	462	Level 3 Communications, LLC
8452	1203	188	7	TEDATA
11492	1092	174	330	Cable One
9583	1075	86	540	Sify Limited
22773	1063	2604	66	Cox Communications, Inc.
18566	1062	296	10	Covad Communications
19262	996	3998	236	Verizon Global Networks
7011	963	248	553	Citizens Utilities

Number of prefixes announced by prefix length

/1:0	/2:0	/3:0	/4:0	/5:0	/6:0
/7:0	/8:19	/9:10	/10:21	/11:58	/12:166
/13:339	/14:598	/15:1149	/16:10459	/17:4745	/18:8143
/19:17144	/20:20423	/21:20318	/22:25909	/23:25599	/24:150522
/25:840	/26:1003	/27:548	/28:159	/29:37	/30:10
/31:0	/32:8				

May 2009 ↑

May 2008 ↓

Number of prefixes announced by prefix length

/1:0	/2:0	/3:0	/4:0	/5:0	/6:0
/7:0	/8:18	/9:9	/10:16	/11:42	/12:140
/13:281	/14:512	/15:1007	/16:9927	/17:4368	/18:7291
/19:15424	/20:17880	/21:17298	/22:21738	/23:22823	/24:134960
/25:783	/26:926	/27:489	/28:81	/29:9	/30:1
/31:0	/32:9				

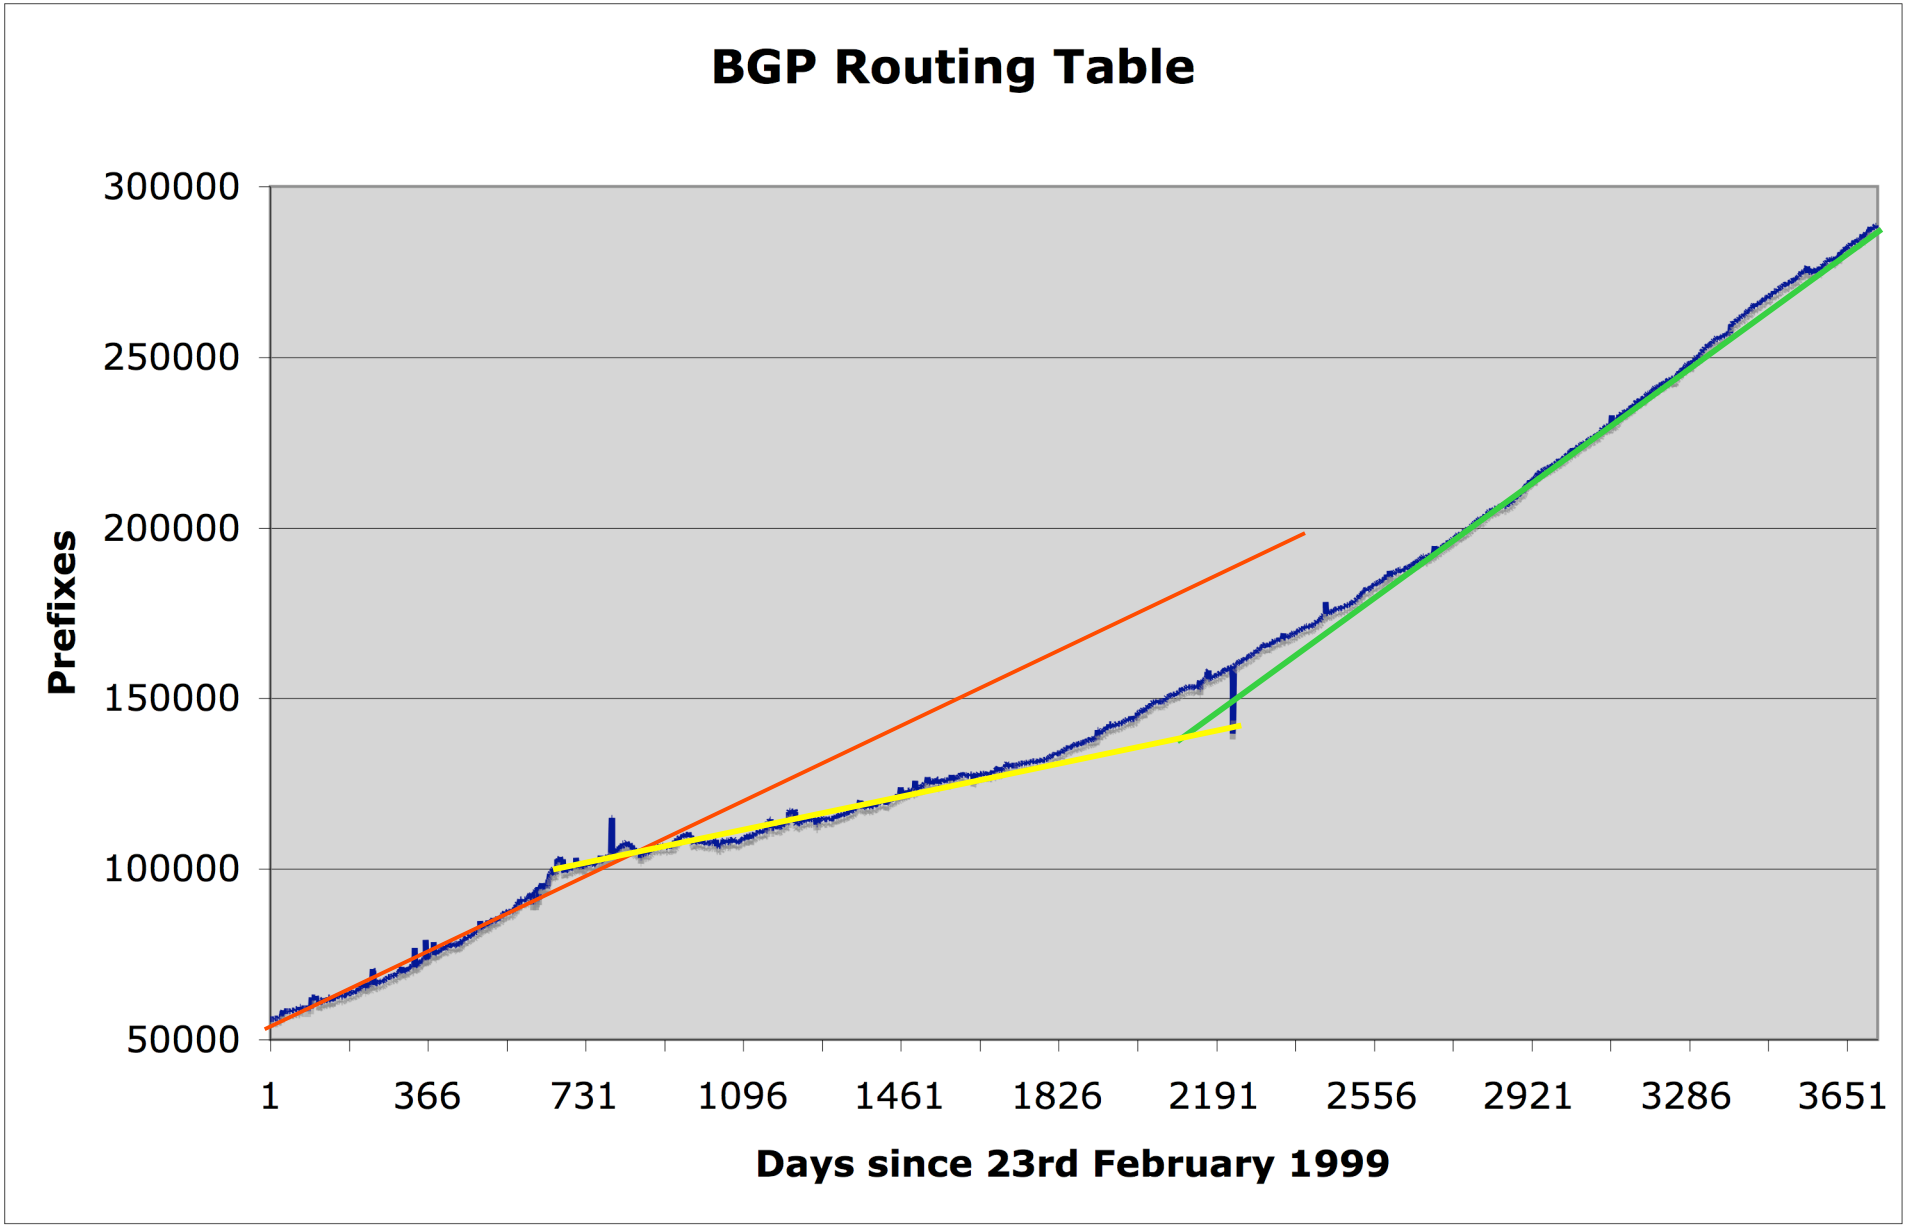
BGP Routing Table
300000
250000
200000
150000
100000
50000
Prefixes
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999

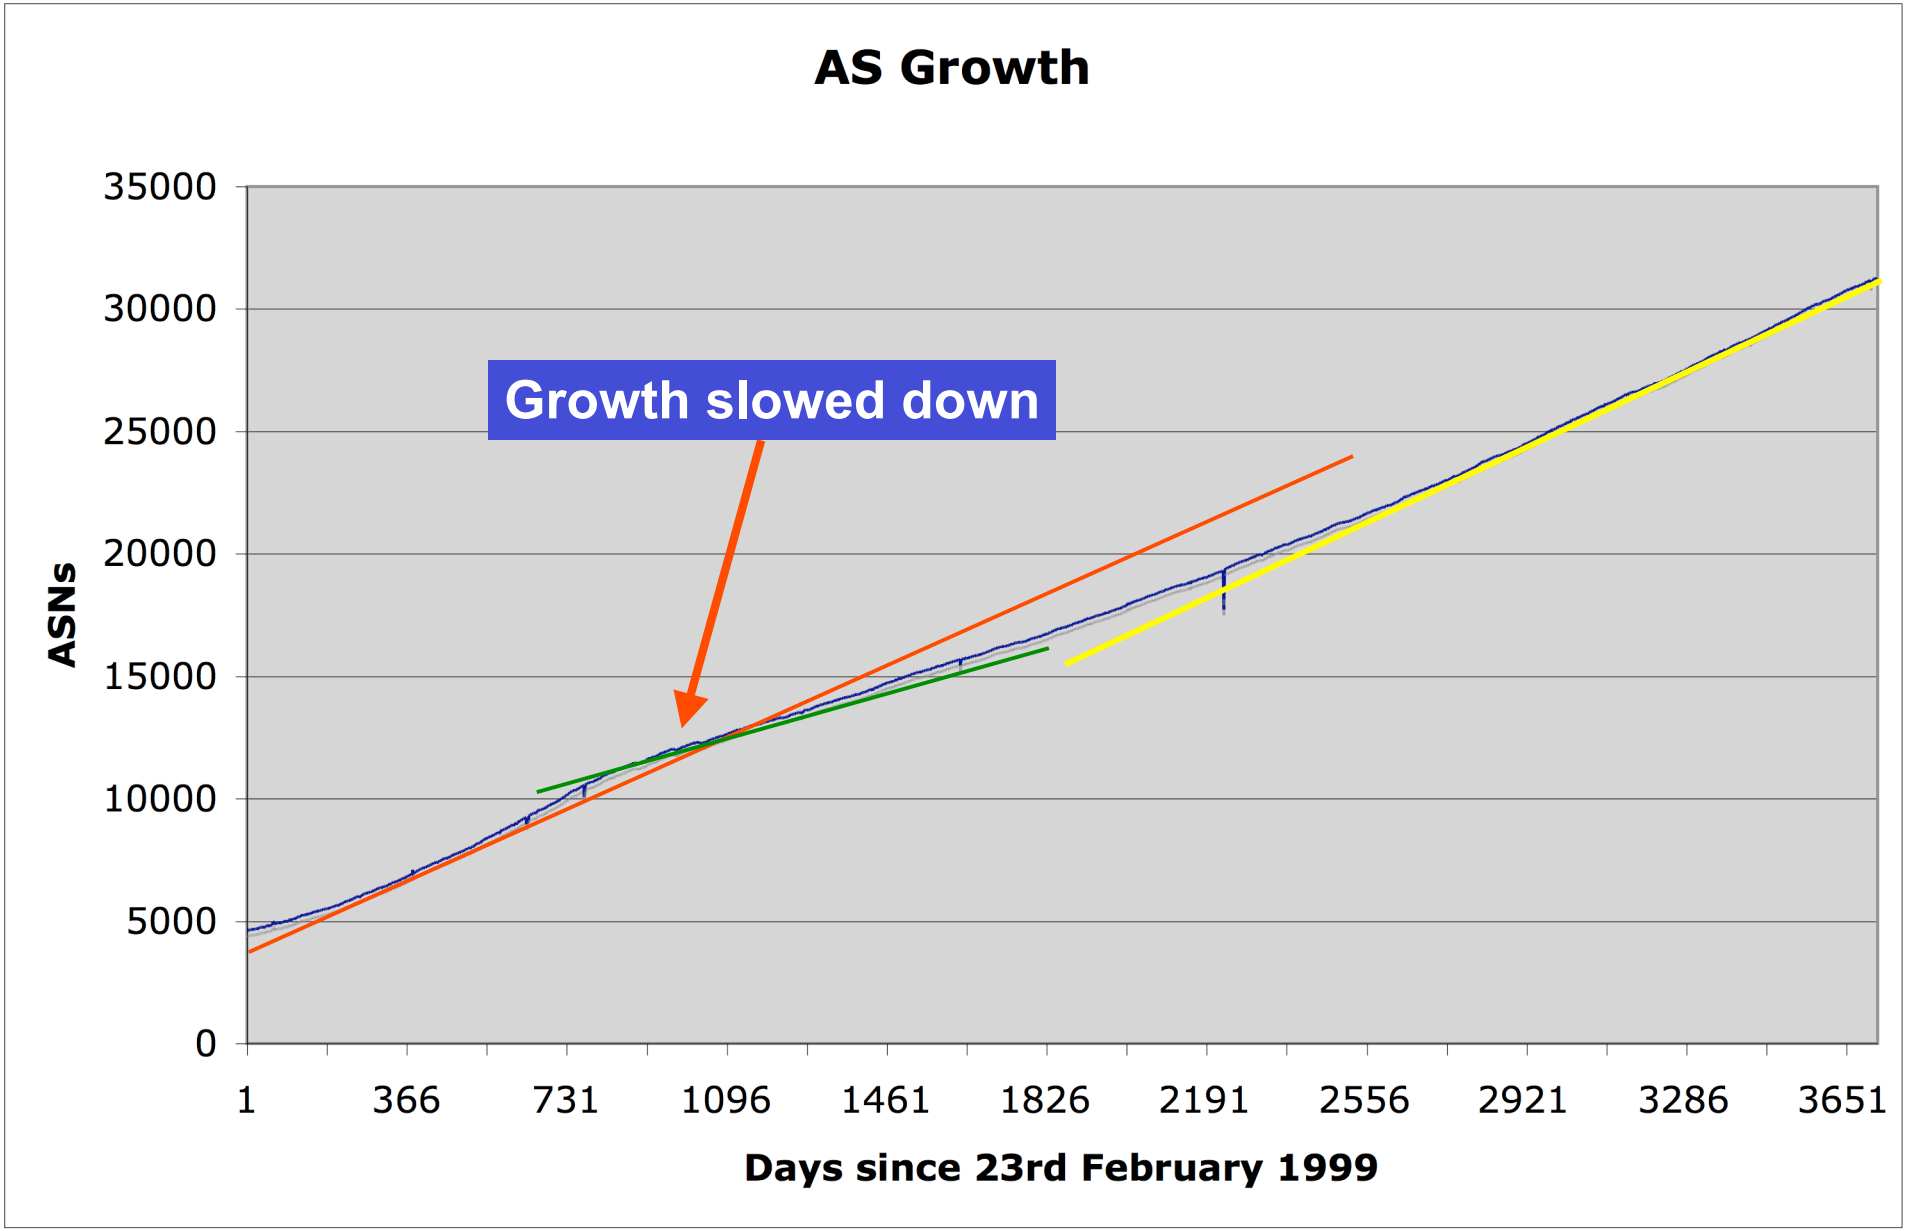
AS Growth
Growth slowed down
ASNs
35000
30000
25000
20000
15000
10000
5000
0
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999

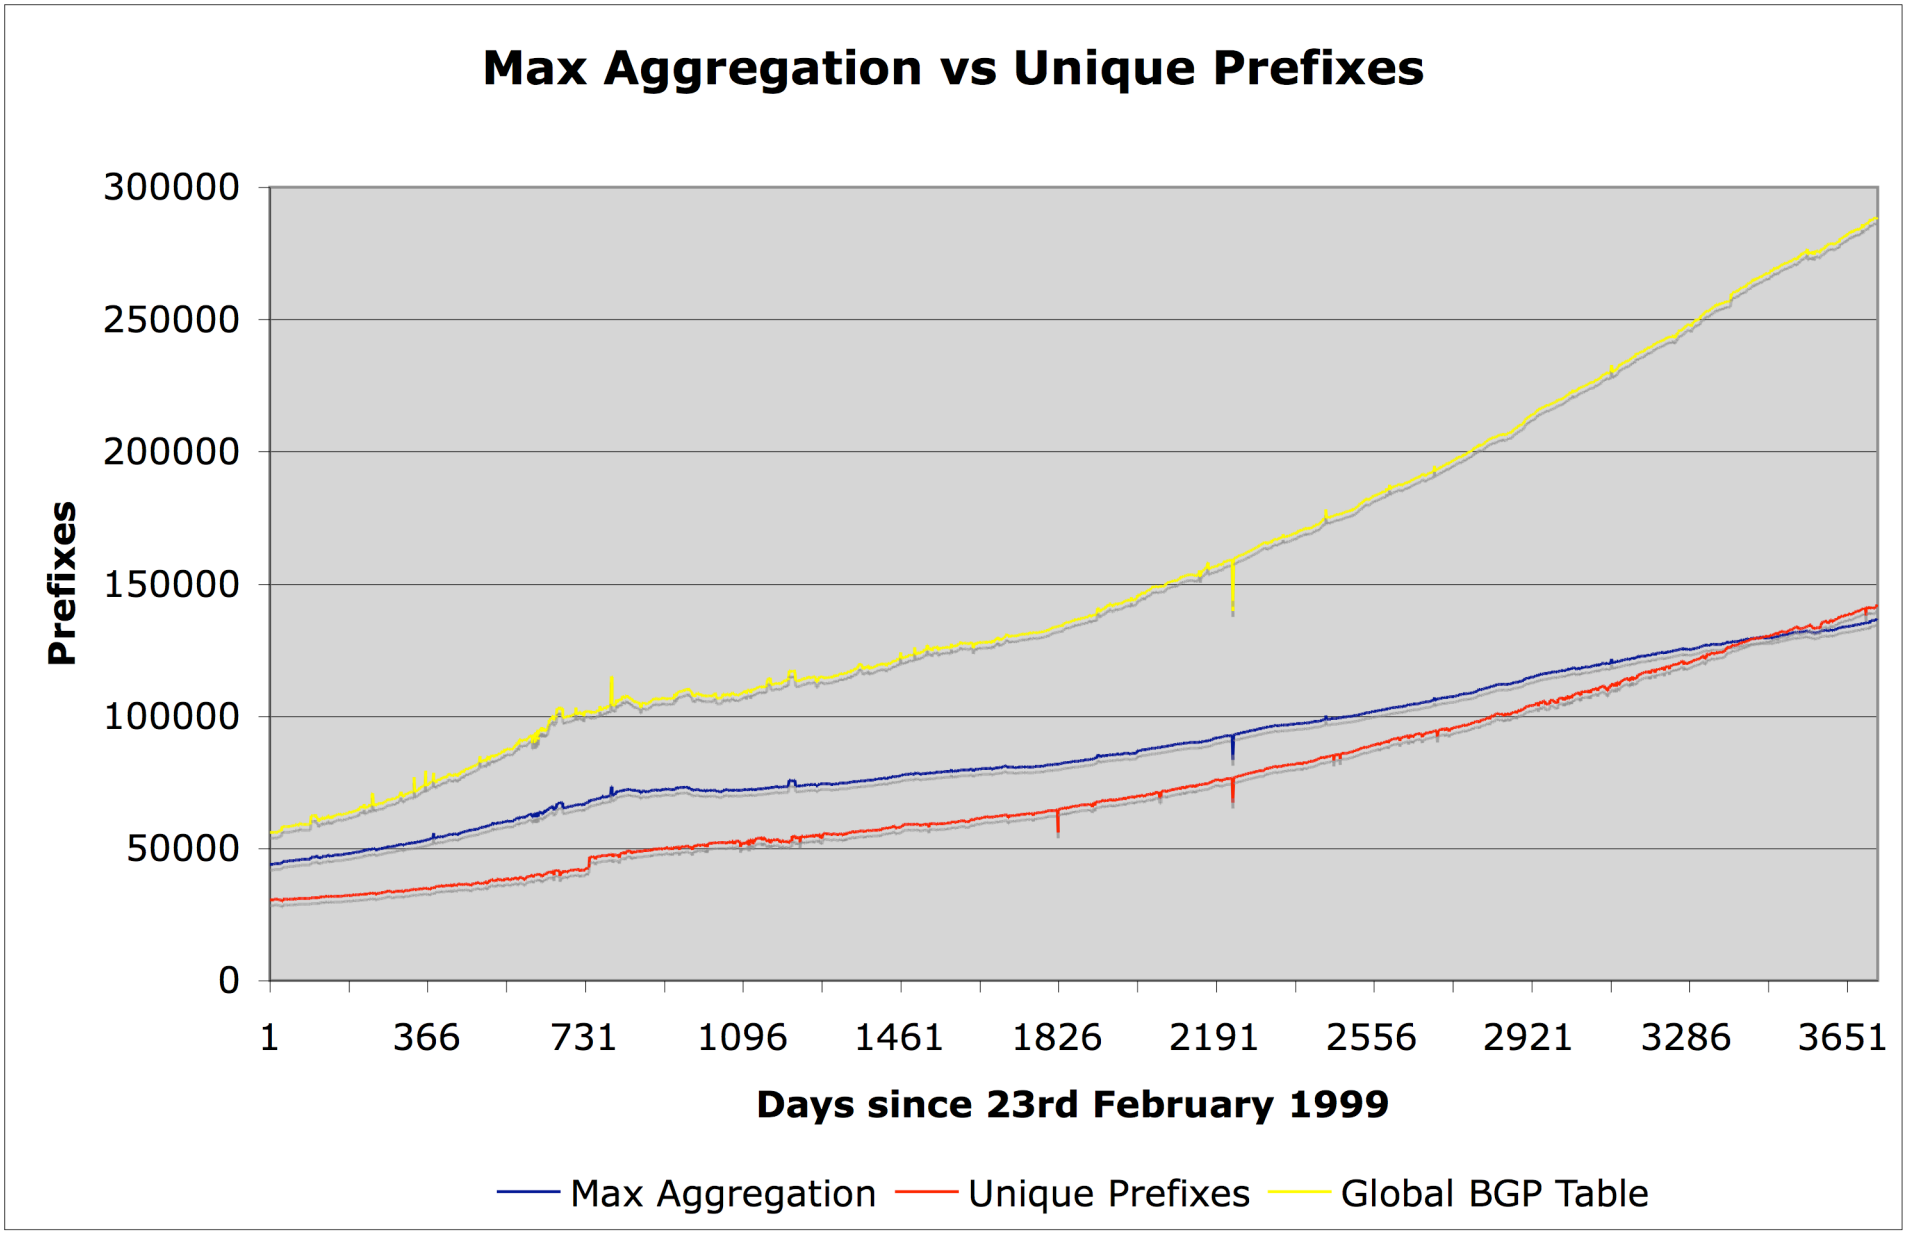
Max Aggregation vs Unique Prefixes
300000
250000
200000
150000
100000
50000
0
Prefixes
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999
Max Aggregation
Unique Prefixes
Global BGP Table

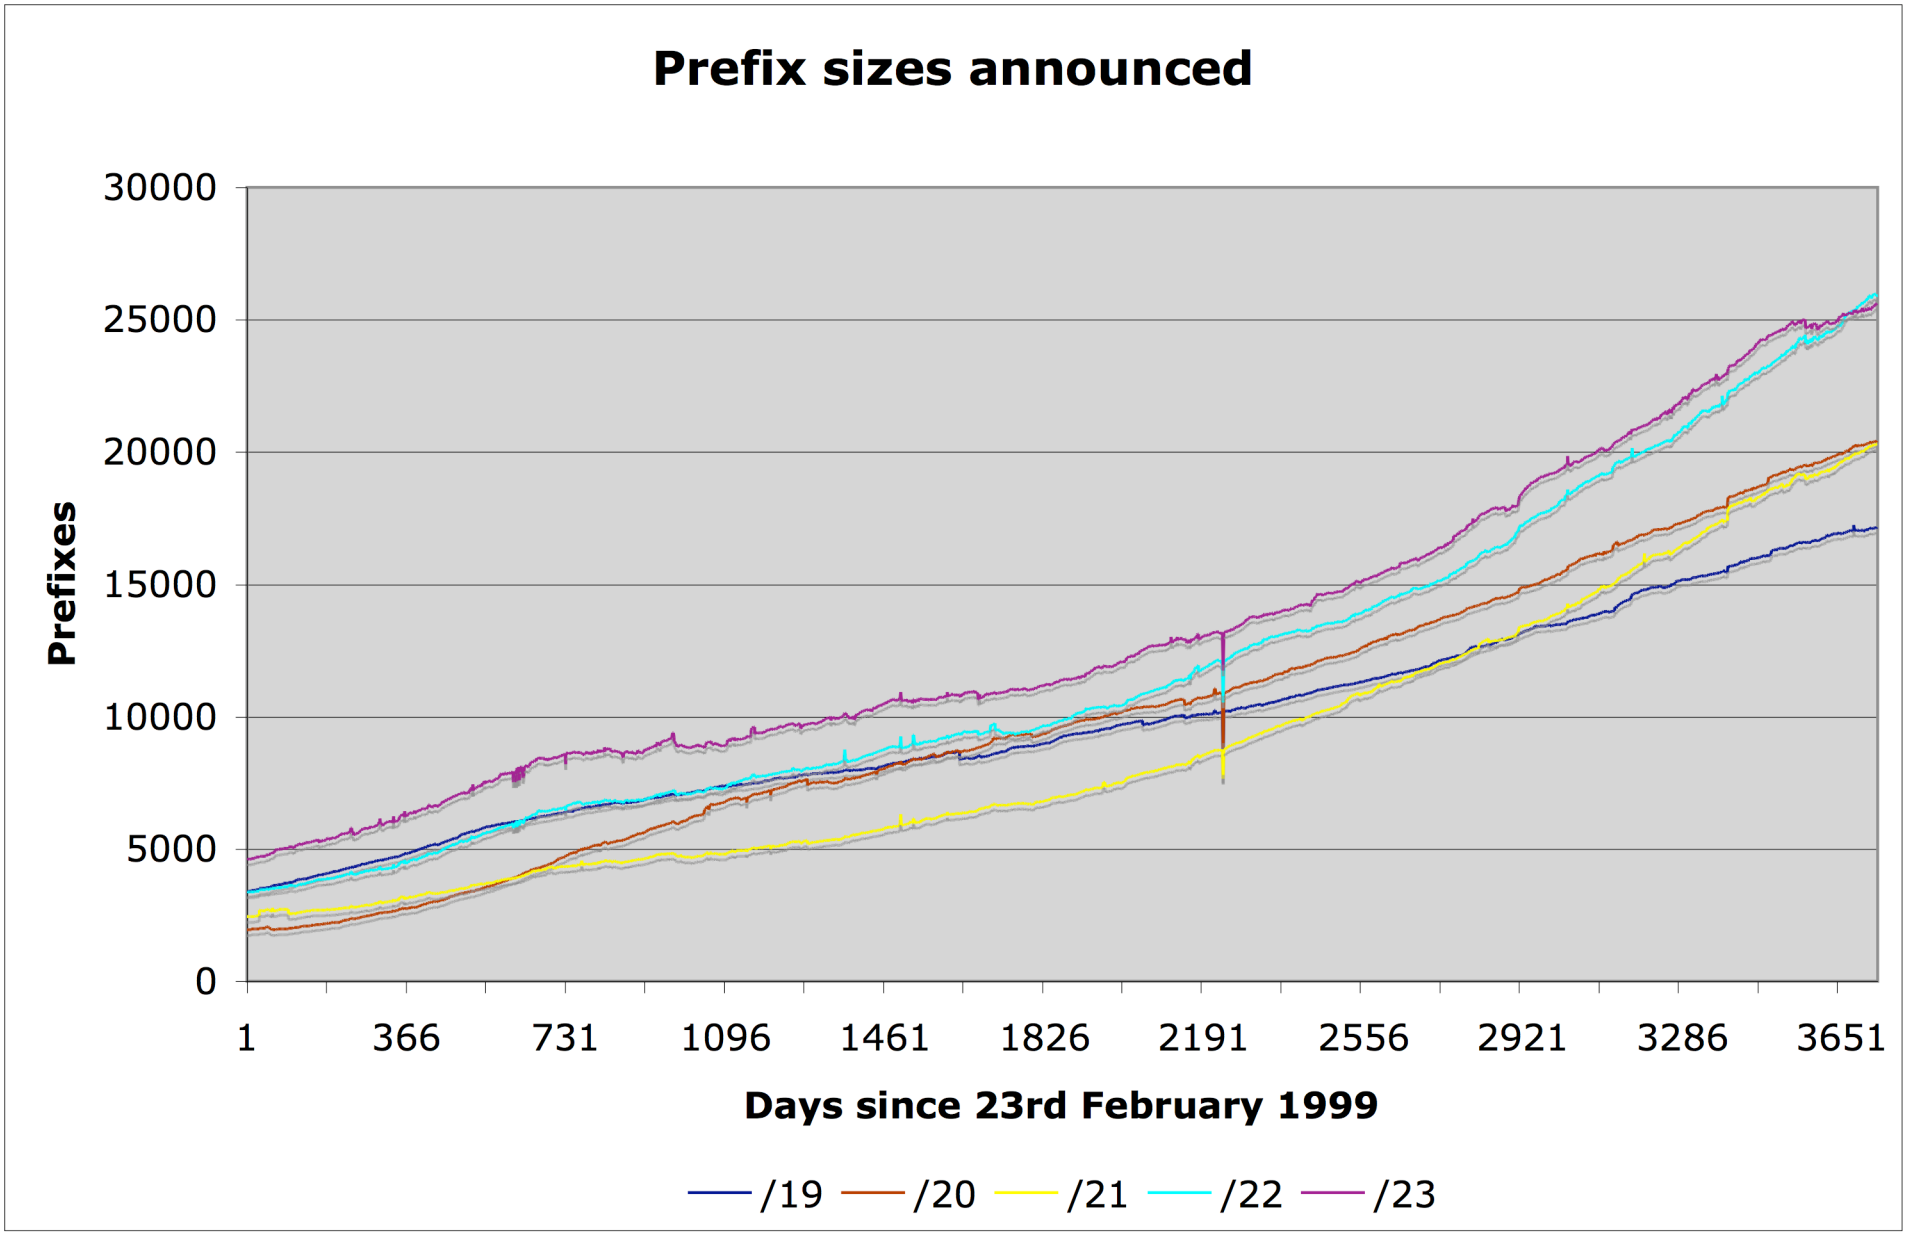
Prefix sizes announced
Prefixes
30000
25000
20000
15000
10000
5000
0
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999
/19
/20
/21
/22
/23

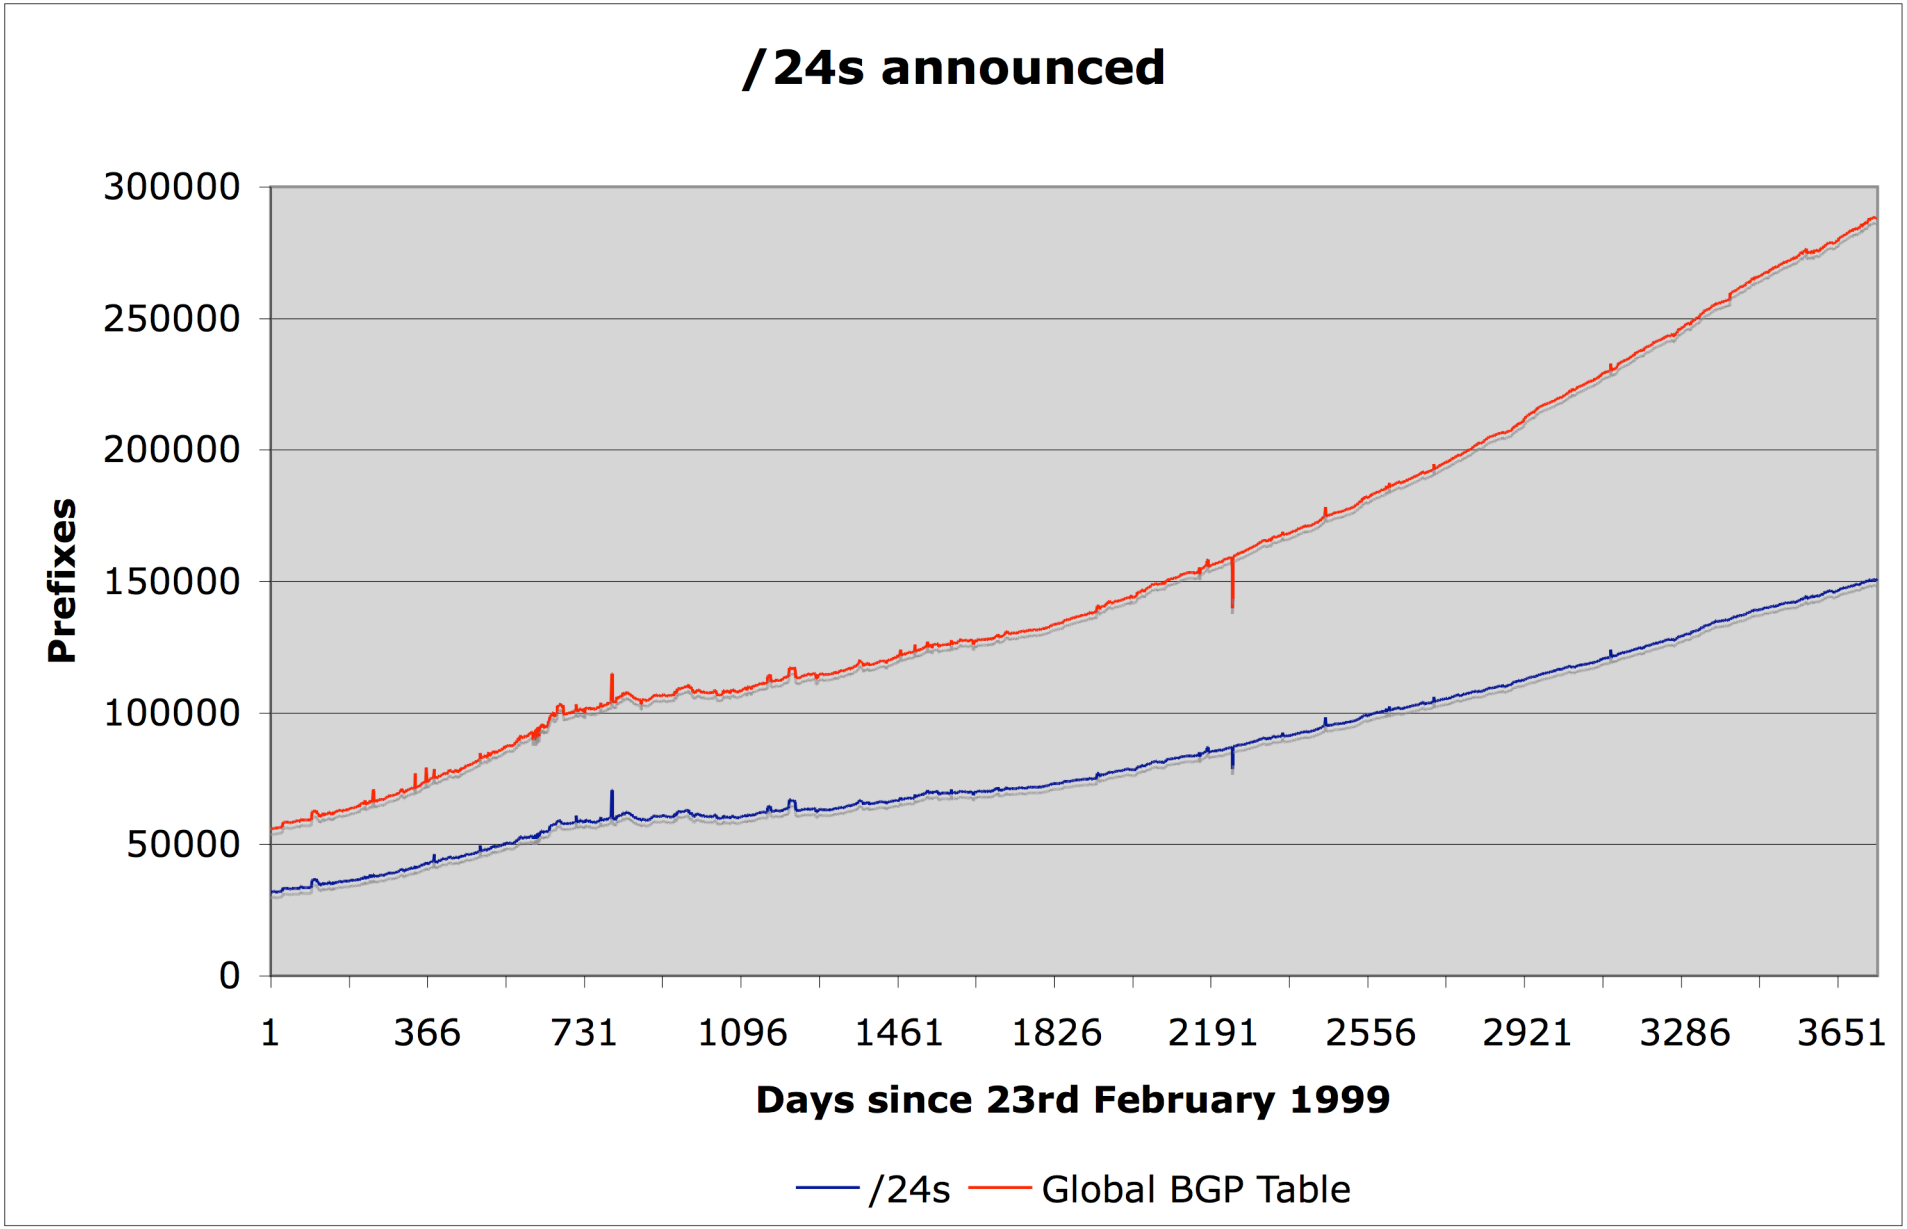
/24s announced
300000
250000
200000
150000
100000
50000
0
Prefixes
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999
/24s
Global BGP Table

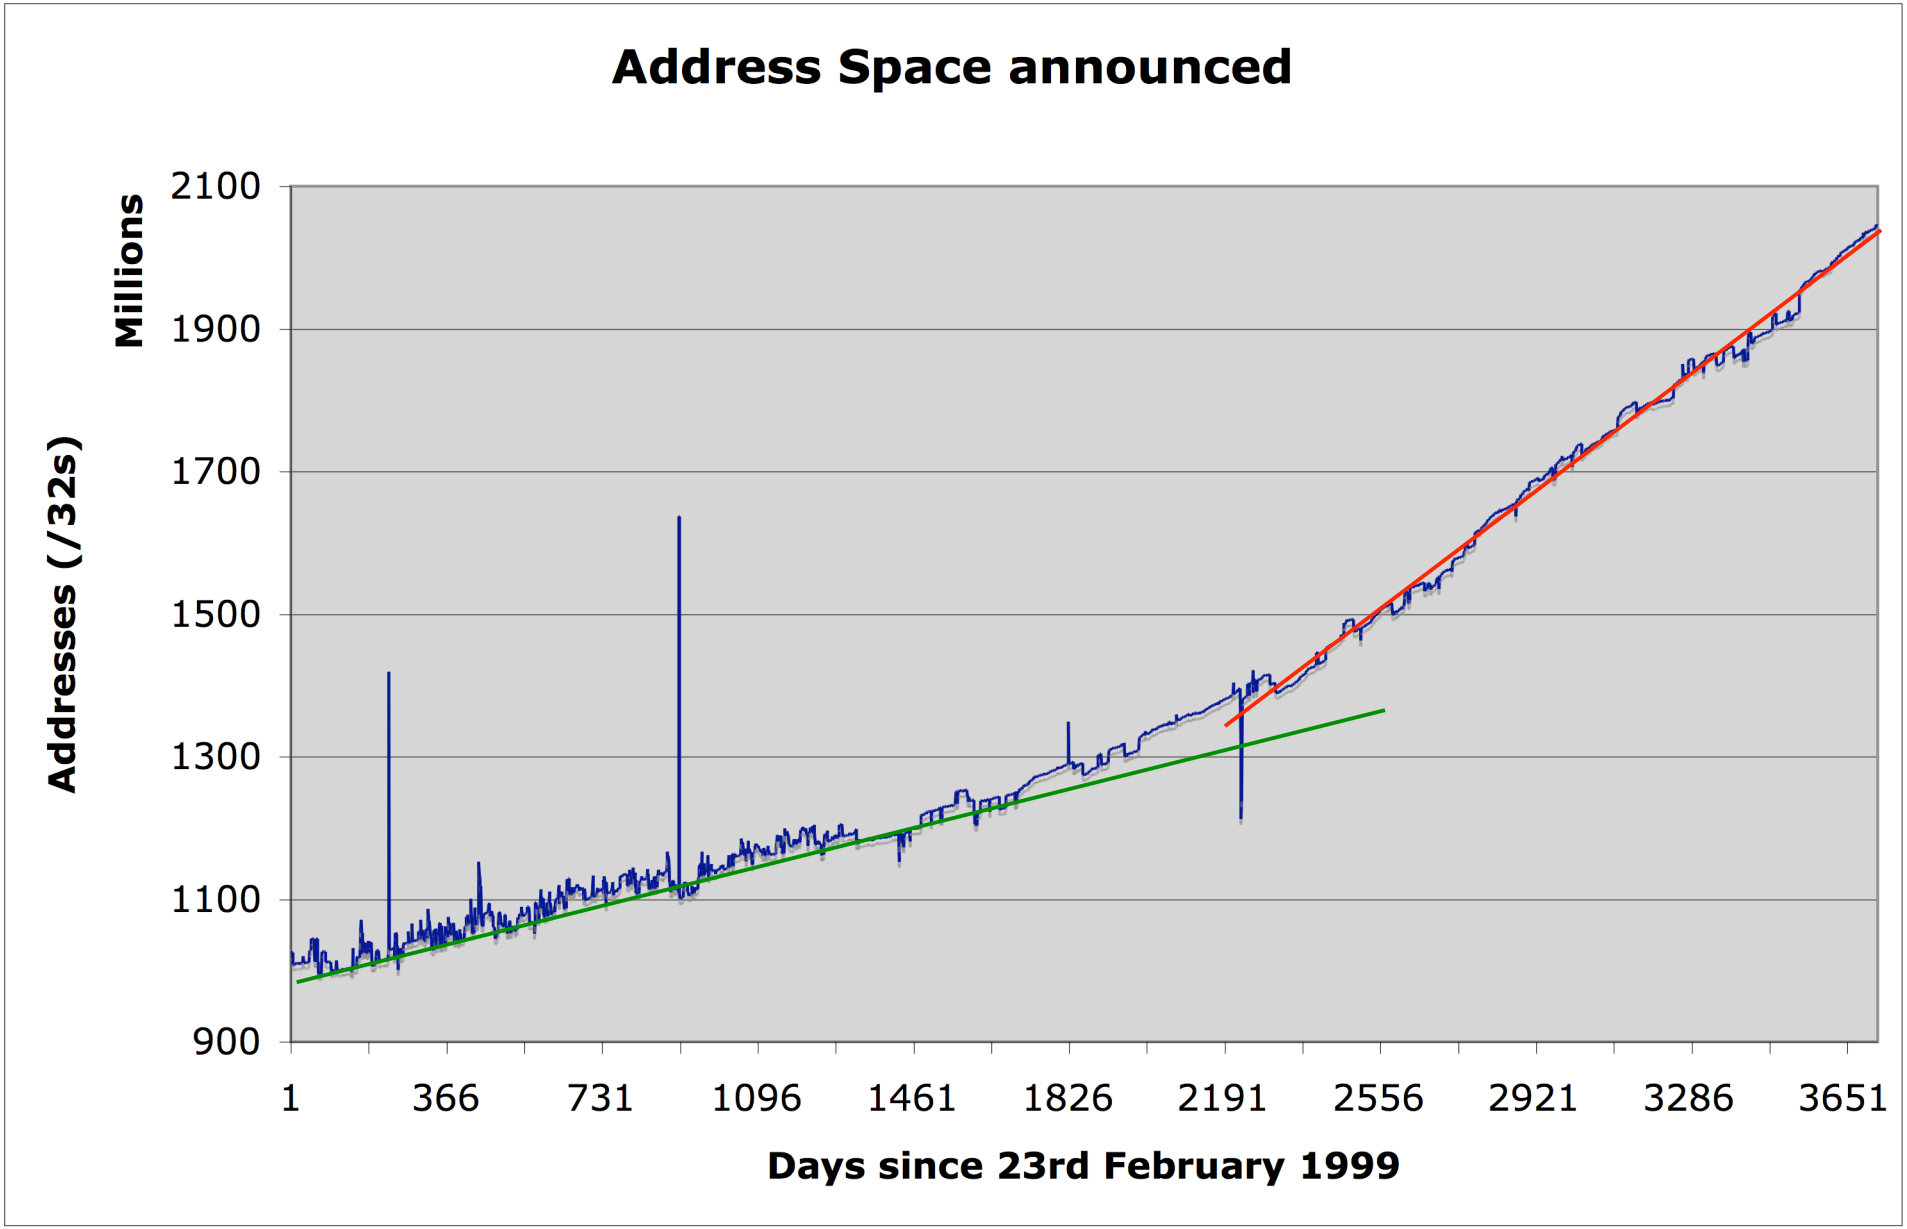
Address Space announced
Addresses (/32s)
Millions
2100
1900
1700
1500
1300
1100
900
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999

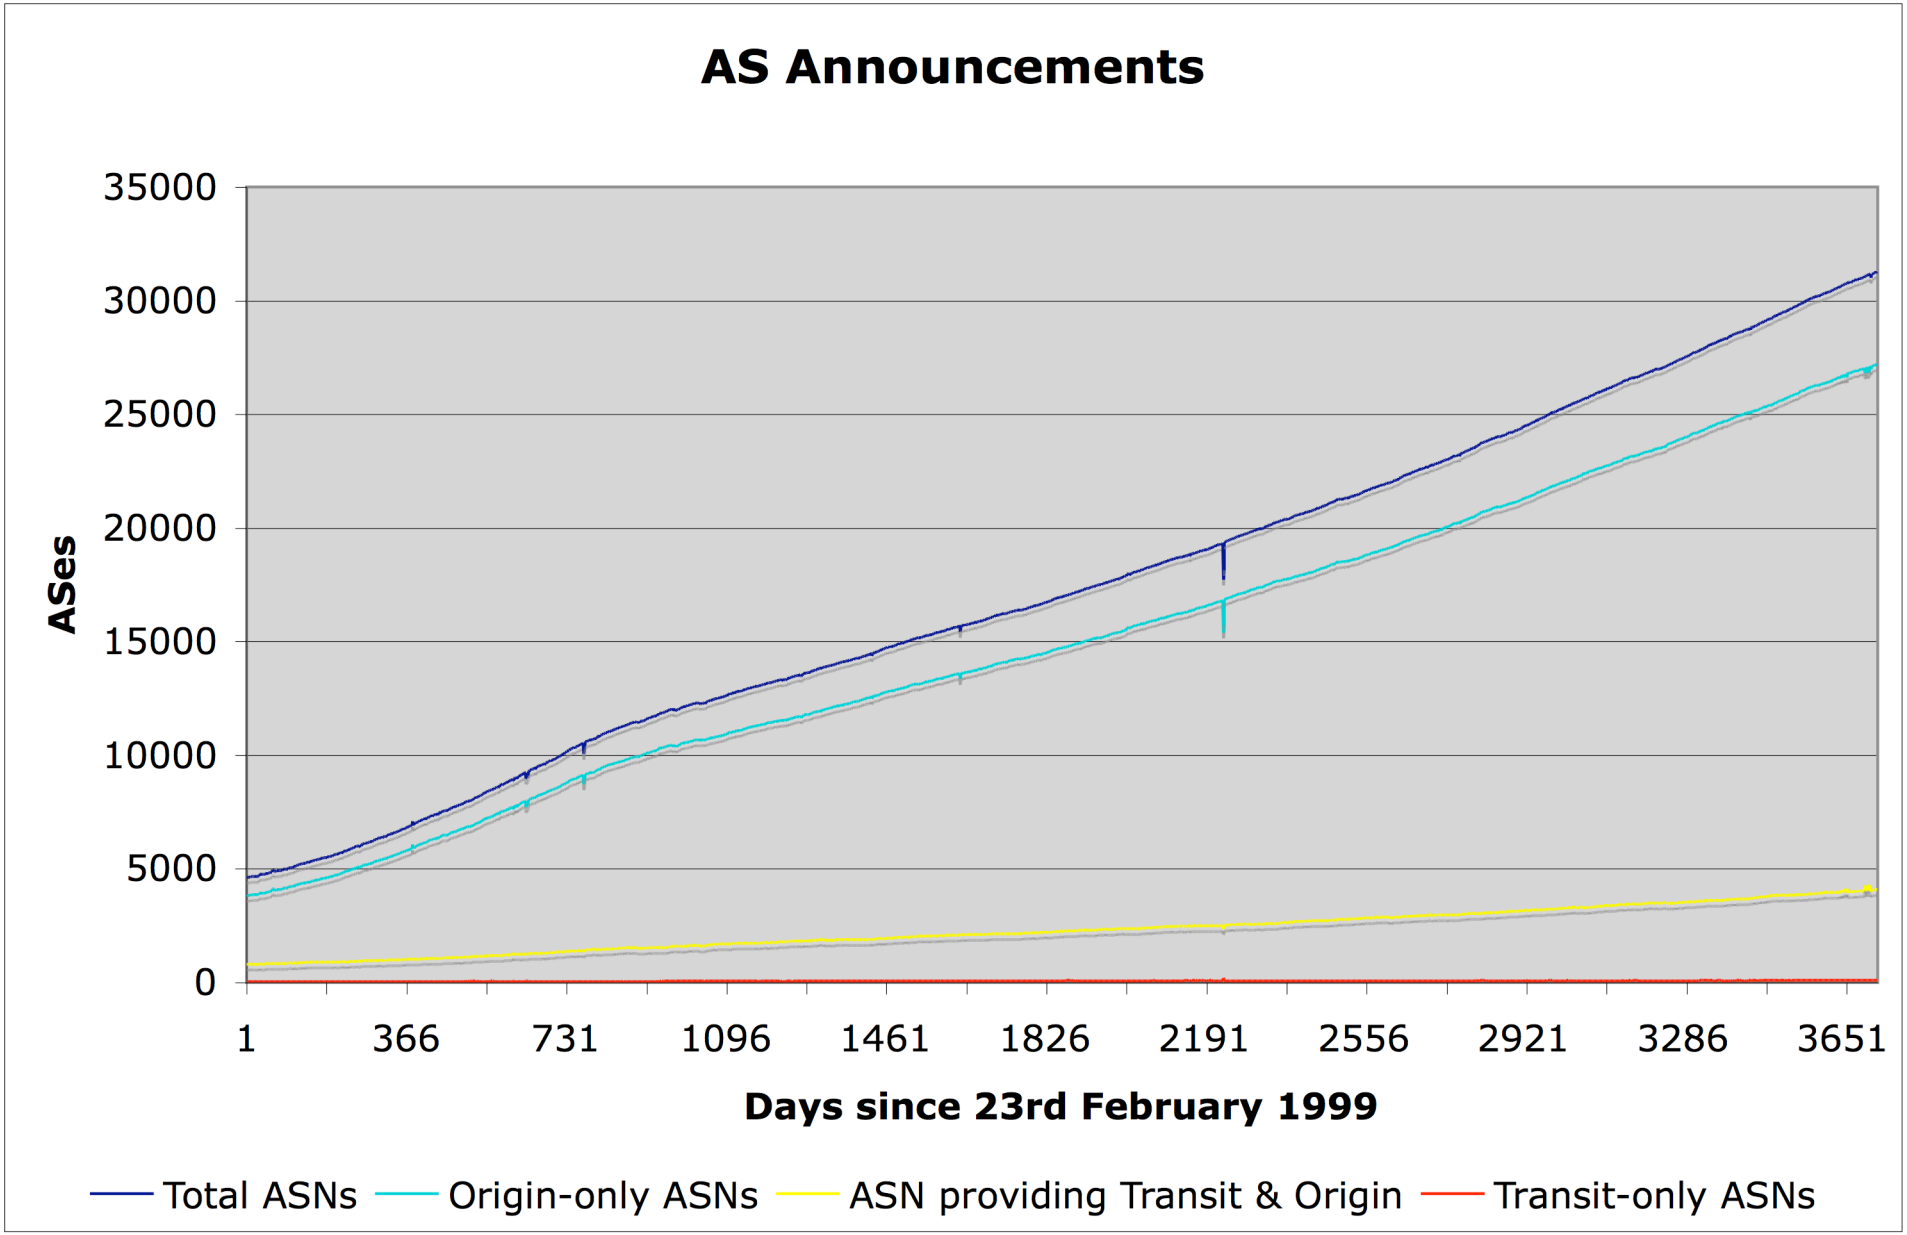
AS Announcements
35000
30000
25000
20000
15000
10000
5000
0
ASes
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999
Total ASNs
Origin-only ASNs
ASN providing Transit & Origin
Transit-only ASNs

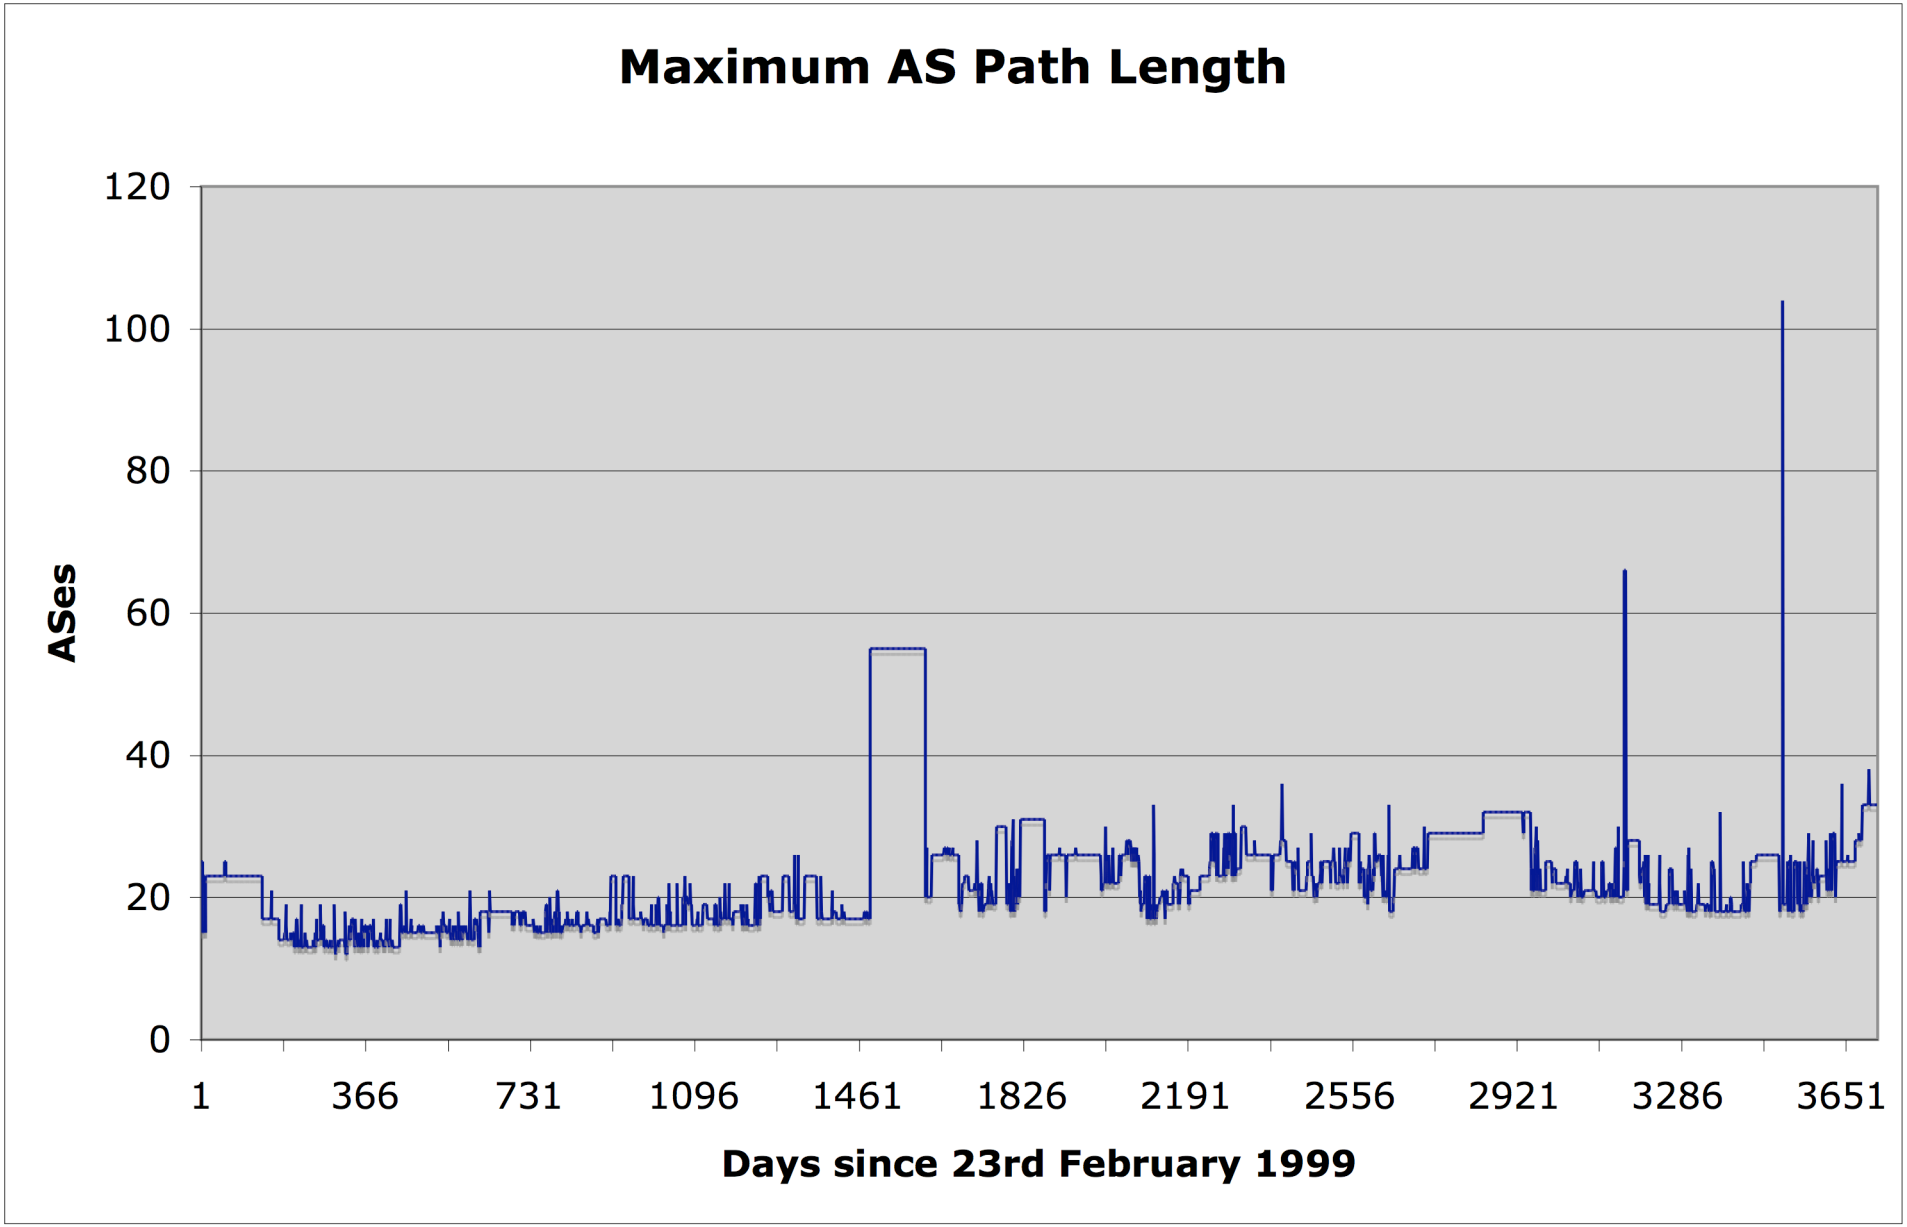
Maximum AS Path Length
ASes
120
100
80
60
40
20
0
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999

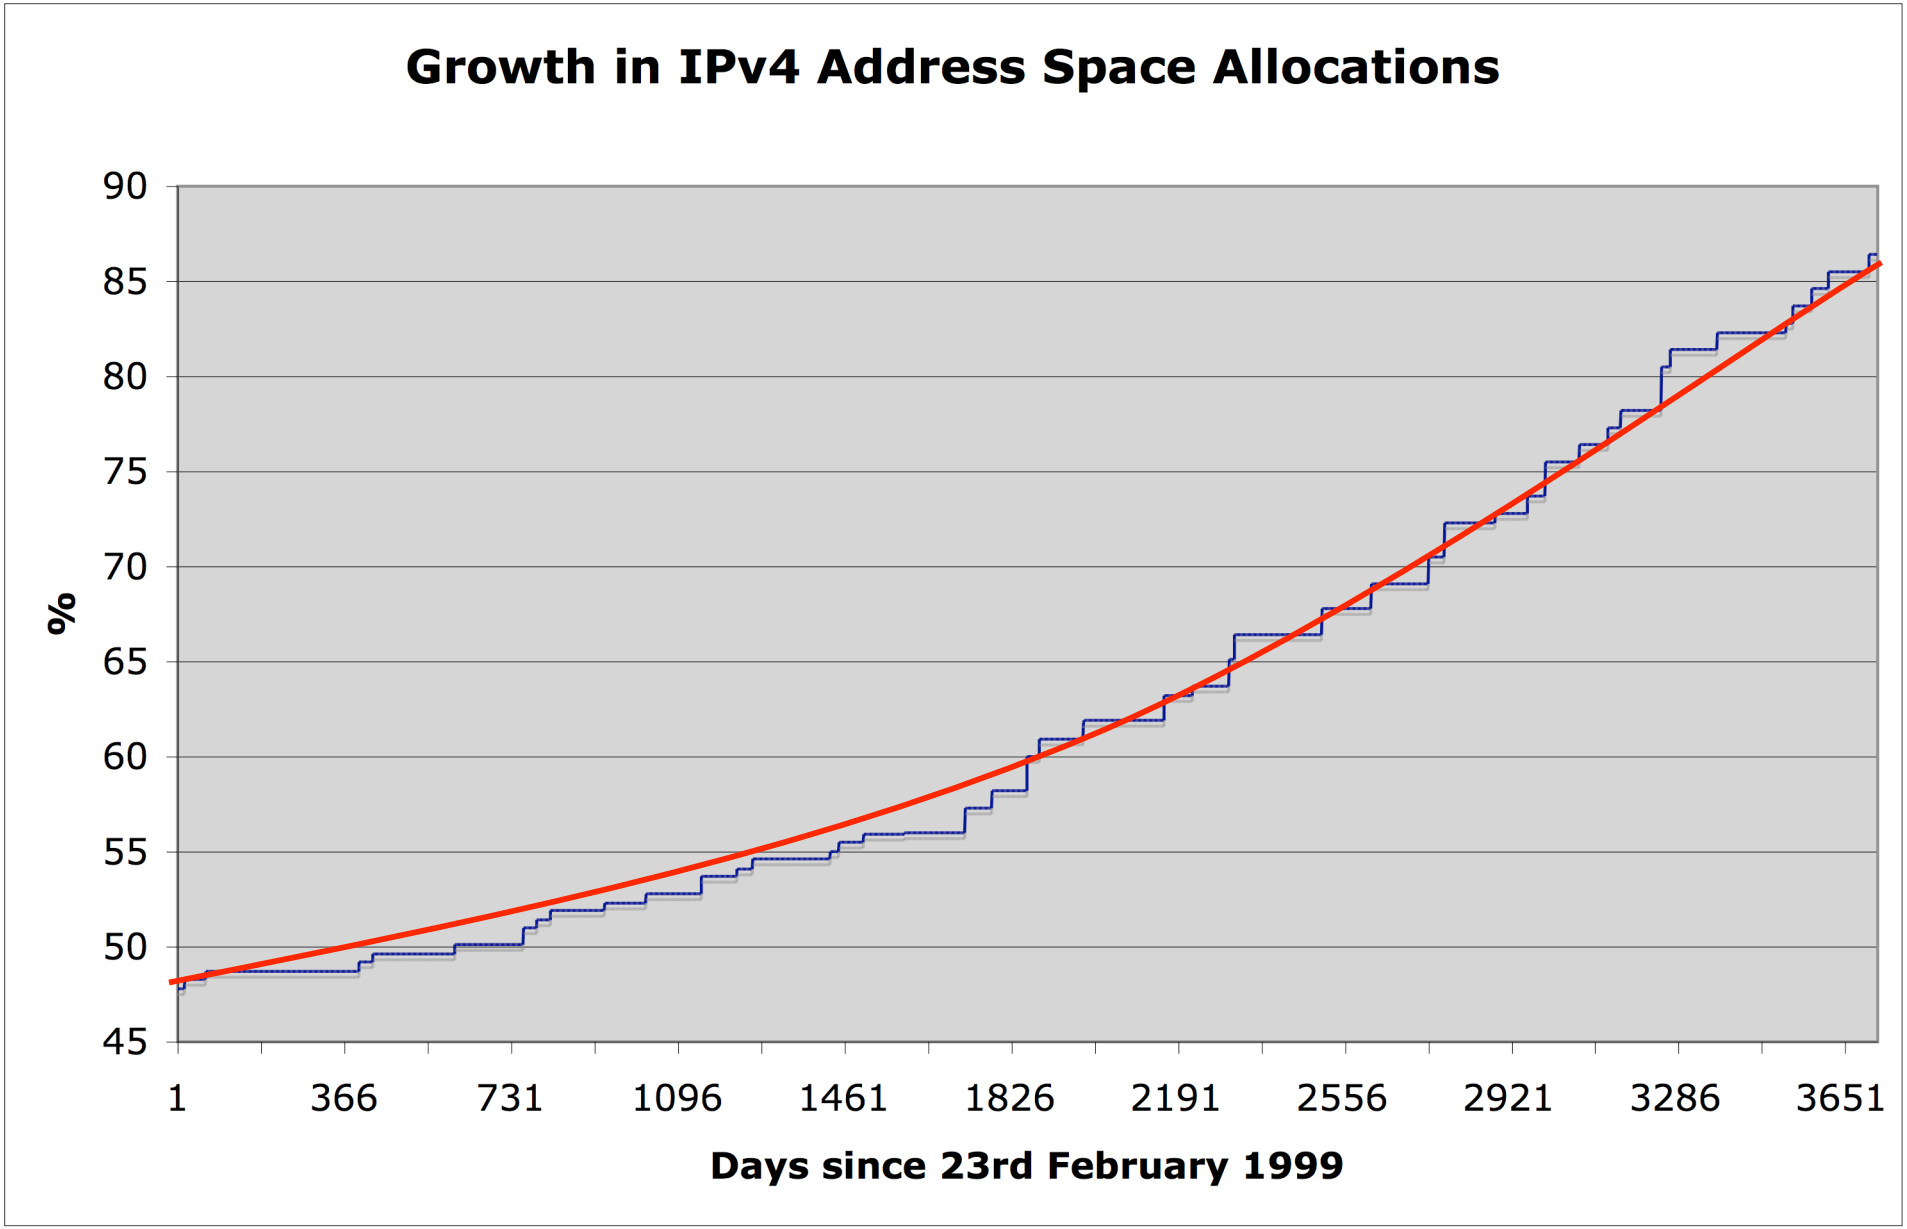
Growth in IPv4 Address Space Allocations
90
85
80
75
70
65
60
55
50
45
%
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999

Route Aggregation Recommendations

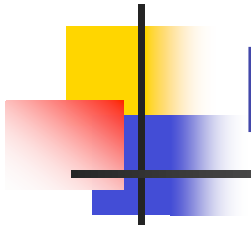

- RIPE Document — RIPE-399
  - http://www.ripe.net/ripe/docs/ripe-399.html
- Discusses:
  - History of aggregation
  - Causes of de-aggregation
  - Impacts on global routing system
  - Available Solutions
  - Recommendations for ISPs

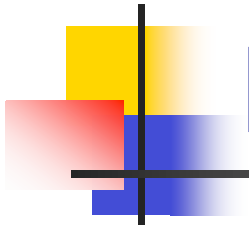

History:

- Classful to classless migration
  - Clean-up efforts in 192/8
- CIDR Report
  - Started by Tony Bates to encourage adoption of CIDR & aggregation
  - Mostly ignored through late 90s
  - Now part of extensive BGP table analysis by Geoff Huston
- Introduction of Regional Internet Registry system and PA address space

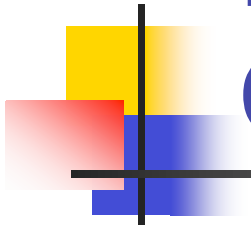

Deaggregation: Claimed causes (1):

- Routing System Security
  - "Announcing /24s means that no one else can DOS the network"
- Reduction of DOS attacks & miscreant activities
  - "Announcing only address space in use as rest attracts 'noise'"
- Commercial Reasons
  - "Mind your own business"

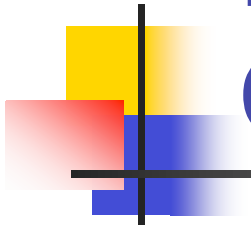

Deaggregation: Claimed causes (2):

- Leakage of iBGP outside of local AS
  - eBGP is NOT iBGP – how many ISPs know this?
- Traffic Engineering for Multihoming
  - Spraying out /24s hoping it will work
  - Rather than do any **real engineering**
- Legacy Assignments
  - "All those pre-RIR assignments are to blame"
  - In reality it is both RIR and legacy assignments

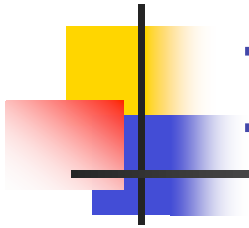

Impacts (1):

- Router memory
  - Shortens router life time as vendors underestimate memory growth requirements
  - Depreciation life-cycle shortened
  - Increased costs for ISP and customers
- Router processing power
  - Processors are underpowered as vendors underestimate CPU requirement
  - Depreciation life-cycle shortened
  - Increased costs for ISP and customers

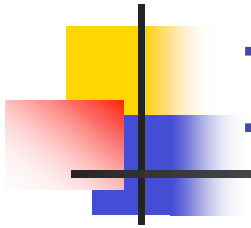

Impacts (2):

- Routing System convergence
  - Larger routing table → slowed convergence
  - Can be improved by faster control plane processors — see earlier
- Network Performance & Stability
  - Slowed convergence → slowed recovery from failure
  - Slowed recovery → longer downtime
  - Longer downtime → unhappy customers

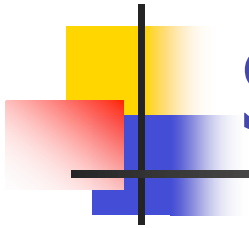

Solutions (1):

- CIDR Report
  - Global aggregation efforts
  - Running since 1994
- Routing Table Report
  - Per RIR region aggregation efforts
  - Running since 1999
- Filtering recommendations
  - BCP38, training, tutorials, Team Cymru,…
- “CIDR Police”

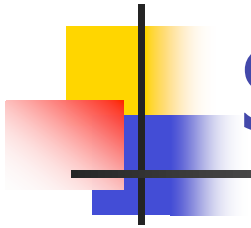

Solutions (2):

- BGP Features:
  - NO_EXPORT Community
  - NOPEER Community
    - RFC3765 — but no one has implemented it
  - AS_PATHLIMIT attribute
    - Still working through IETF IDR Working Group
  - Provider Specific Communities
    - Some ISPs use them; most do not

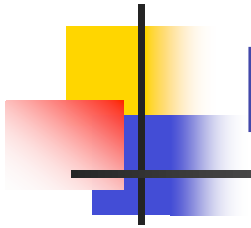

RIPE-399 Recommendations:

- Announcement of initial allocation as a single entity
- Subsequent allocations aggregated if they are contiguous and bit-wise aligned
- Prudent subdivision of aggregates for Multihoming
- Use BGP enhancements already discussed
- (All of this applies to IPv6 too)

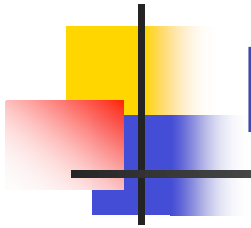

Looking at Deaggregation

- CIDR Report
  - www.cidr-report.org
  - Encourages aggregation following CIDRisation of Internet
  - Today: extensive suite of reports and tools covering state of BGP table
- Routing Report
  - BGP table status on per RIR basis
  - Original CIDR Report and a whole lot more

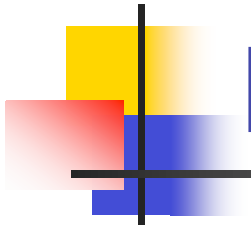

Deaggregation Factor

- Routing Report
  - One summary takes BGP table and aggregates prefixes by origin AS
    - Called "Max Aggregation" in report
  - Global and per RIR basis
    - http://thyme.apnic.net/current/
- New Deaggregation Factor:
  - Measure of Routing Table size/Aggregated Size
  - Global value has been increasing slowly and steadily since "records began"

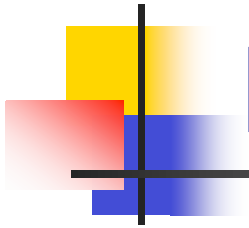

May 2009

Total Prefixes

- Global BGP Table
  - 288k prefixes
- Europe & Middle East
  - 62k prefixes
- North America
  - 125k prefixes
- Asia & Pacific
  - 68k prefixes
- Africa
  - 5k prefixes
- Latin America & Caribbean
  - 24k prefixes

Deaggregation Factor

- Global Average
  - 2.11
- Europe & Middle East
  - 1.69
- North America
  - 1.89
- Asia & Pacific
  - 2.80
- Africa
  - 3.43
- Latin America & Caribbean
  - 4.03

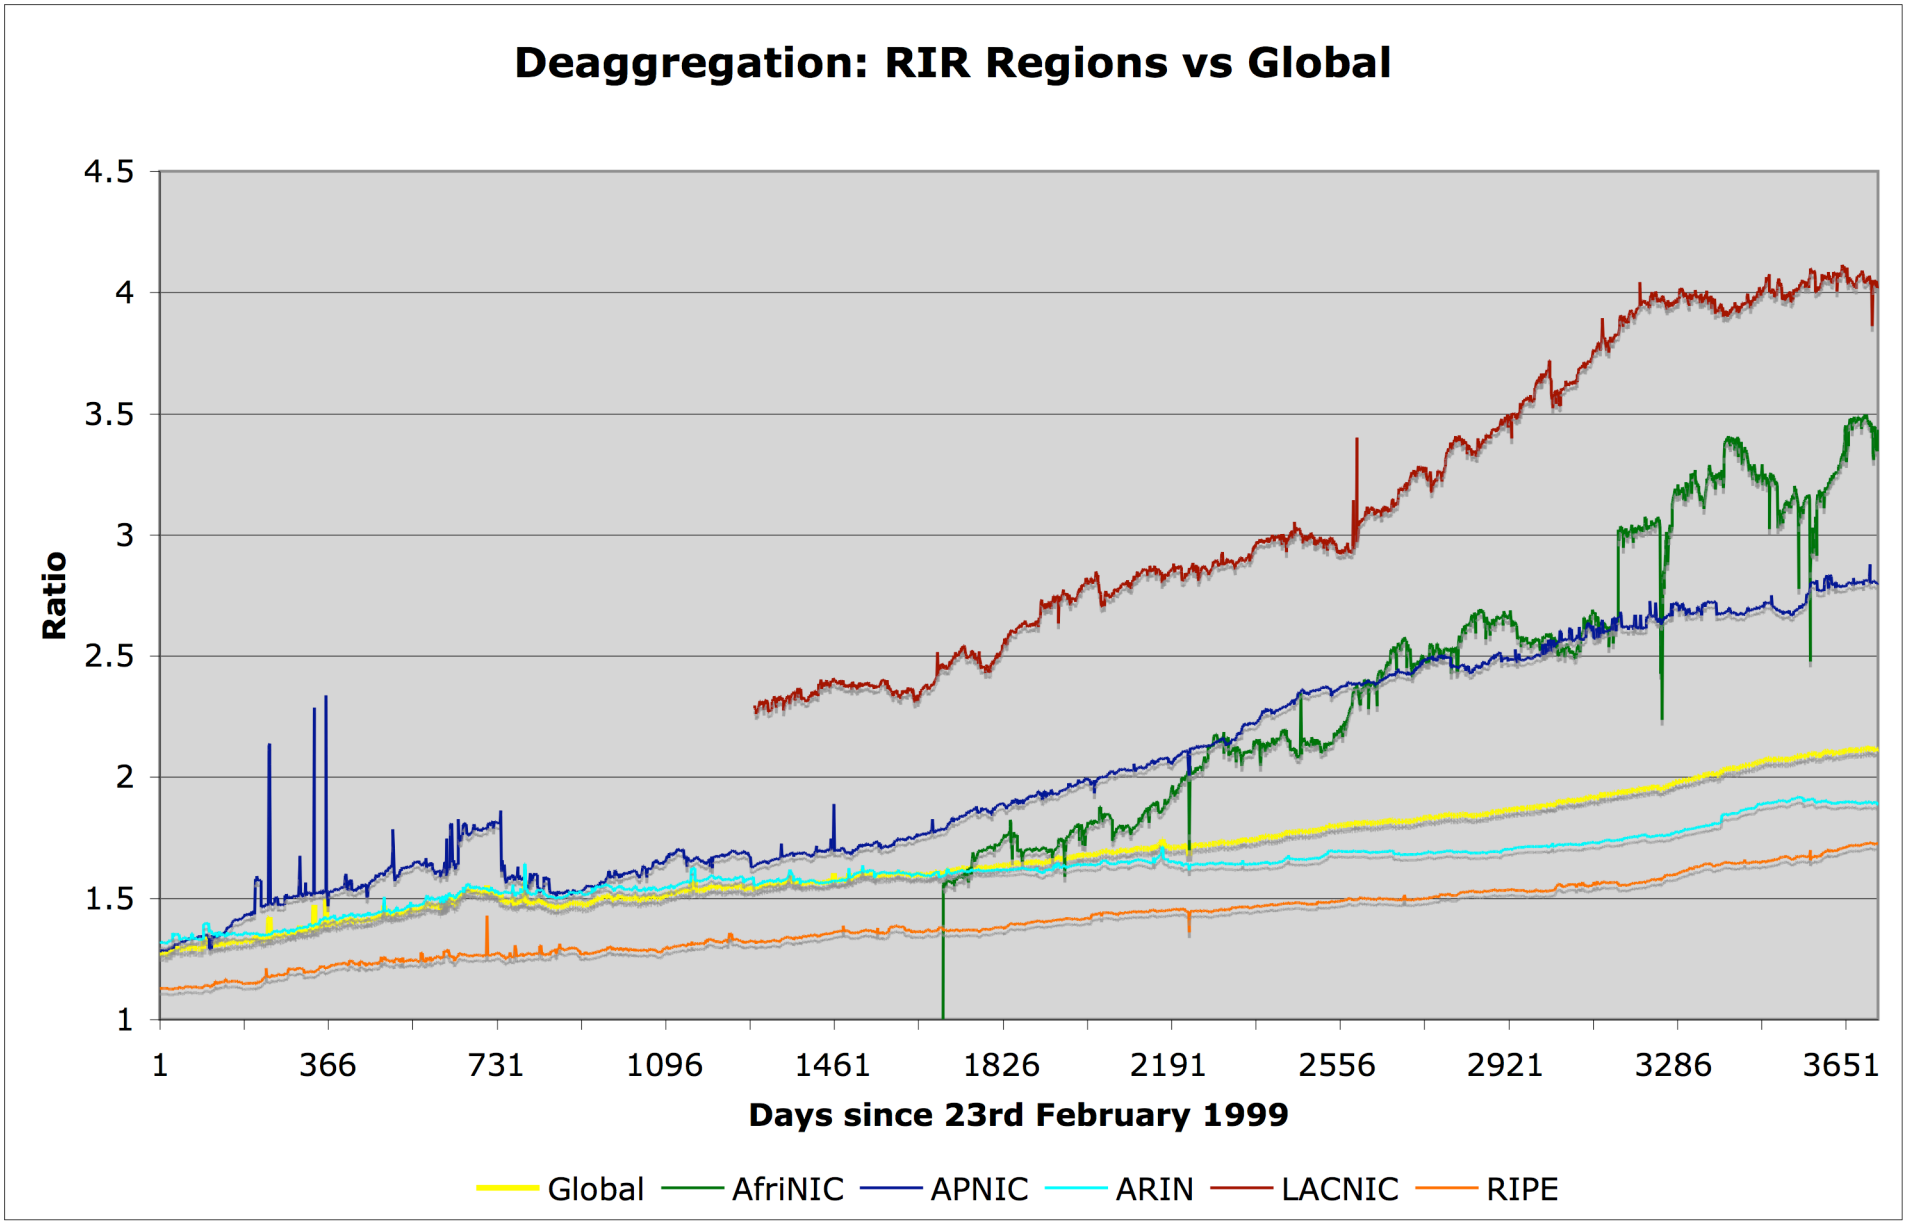
Deaggregation: RIR Regions vs Global
Ratio
4.5
4
3.5
3
2.5
2
1.5
1
1
366
731
1096
1461
1826
2191
2556
2921
3286
3651
Days since 23rd February 1999
Global
AfriNIC
APNIC
ARIN
LACNIC
RIPE

Africa Aggregation Savings Summary

ASN	No of Nets	Savings	Description
8452	1203	1196	TEDATA
24863	862	824	LINKdotNET AS number
20858	309	304	This AS will be used to conne
6713	160	148	Itissalat Al-MAGHRIB
33783	152	144	EEPAD TISP TELECOM & INTERNET
29571	138	130	Ci Telecom Autonomous system
5536	123	114	Internet Egypt Network
33776	116	109	Starcomms Nigeria Limited
2018	244	101	Tertiary Education Network
24835	107	98	RAYA Telecom - Egypt
20928	72	70	Noor Advanced Technologies AS
15475	71	66	Nile Online
15706	61	57	Sudatel Internet Exchange Aut
5713	115	49	Telkom SA Ltd
29975	62	47	Vodacom
37054	70	44	Data Telecom Service
8094	46	43	PUKNET
3741	278	40	The Internet Solution
12455	43	40	Jambonet Autonomous system
33774	77	37	AS Number for Telecom Algeria

http://thyme.apnic.net/current/data-CIDRnet-AFRINIC

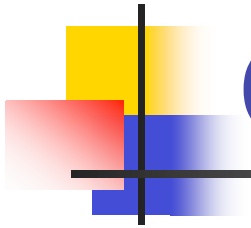

Observations

- Range of operational "practices" between RIR regions
  - Deaggregation by newer ISPs & developing regions is growing rapidly
  - Is harming the **entire** Internet
- RIPE-399 is only a recommendation
  - Hopefully all the RIRs will include pointers with each address allocation
  - Hopefully more ISPs will pay attention to it
  - Training is there — most ISPs choose to ignore it

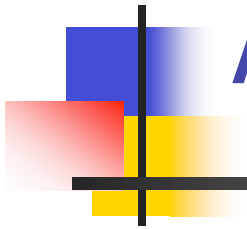

Internet Routing Table Analysis Update

Questions?