Internet Status

Philip Smith

<pfs@cisco.com>

Corporate Consulting Engineering

10-12 June 2008

Topics

- IPv4 run-out
- State of the Backbone
- Internet Operations Groups
- Anycast
- 4-byte ASNs
- Security

IPv4 run-out…

…and the history of IPv6

Early Internet History

- Late 1980s
  - Exponential growth of the Internet
- Late 1990: CLNS proposed as IP replacement
- 1991-1992
  - Running out of “class-B” network numbers
  - Exponential growth of the “default-free” routing table
  - Imminent exhaustion of 32-bit address space
- Two efforts – short-term vs. long-term
  - More at “The Long and Windy ROAD”
    http://rms46.vlsm.org/1/42.html

Early Internet History

- CIDR and Supernetting proposed in 1992-3
  - Deployment started in 1994
- IETF “ipng” solicitation – RFC1550, Dec 1993
- Direction and technical criteria for ipng choice – RFC1719 and RFC1726, Dec 1994
- Proliferation of proposals:
  - TUBA – RFC1347, June 1992
  - PIP – RFC1621, RFC1622, May 1994
  - CATNIP – RFC1707, October 1994
  - SIP – RFC1710, October 1994
  - NIMROD – RFC1753, December 1994
  - ENCAPS – RFC1955, June 1996

Early Internet History → 1996

- Other activities included:
  - Development of NAT, PPP, DHCP,…
  - Some IPv4 address reclamation
  - The RIR system was introduced
- → Brakes were put on IPv4 address consumption
- IPv4 32 bit address = 4 billion hosts
  - HD Ratio (RFC3194) realistically limits IPv4 to 250 million hosts
- IPv6 adopted as the replacement of IPv4
  - Protocol development work begun

Recent Internet History
The “boom” years → 2001

- IPv6 Development in full swing
  - Rapid IPv4 consumption
  - IPv6 specifications sorted out
  - (Many) Transition mechanisms developed
- 6bone
  - Experimental IPv6 backbone sitting on top of Internet
  - Participants from over 100 countries
- Early adopters
  - Japan, Germany, France, UK,…

Recent Internet History
The “bust” years: 2001 → 2004

- The DotCom “crash”

  i.e. Internet became mainstream

- IPv4:

  Consumption slowed

  Address space pressure “reduced”

- Indifference

  Early adopters surging onwards

  Sceptics more sceptical

  Yet more transition mechanisms developed

2004 → Today

- Resurgence in demand for IPv4 address space

  17.7% now remaining in IANA free pool

  Exhaustion predictions range from wild to conservative

  …but late 2010 seems realistic at current rates

  …but what about the market for address space?

- Market for IPv4 addresses:

  Creates barrier to entry

  Condemns the less affluent to use of NATs

- IPv6 offers vast address space

  The only compelling reason for IPv6

Current Situation

- General perception is that “IPv6 has not yet taken hold”
  - IPv4 Address run-out is not quite “headline news” yet
    - More discussions and run-out plans proposed
  - Private sector requires a business case to “migrate”
    - No easy Return on Investment (RoI) computation
- But reality is very different from perception!
  - Something needs to be done to sustain the Internet growth
  - IPv6 or NAT or both or something else?

Do we really need a larger address space?

- Internet population

  ~630 million users end of 2002 – 10% of world pop.

  ~1320 million users end of 2007 – 20% of world pop.

  Future? (World pop. ~9B in 2050)

- US uses 81 /8s – this is 3.9 IPv4 addresses per person

  Repeat this the world over…

  6 billion population could require 23.4 billion IPv4 addresses

  (6 times larger than the IPv4 address pool)

- Emerging Internet economies need address space:

  China uses more than 94 million IPv4 addresses today (5.5 /8s)

Do we really need a larger address space?

- RFC 1918 is not sufficient for large environments
  - Cable Operators (e.g. Comcast – NANOG37 presentation)
  - Mobile providers (fixed/mobile convergence)
  - Large enterprises
- The Policy Development process of the RIRs turned down a request to increase private address space
  - RIR membership guideline is to use global addresses instead
  - This leads to an accelerated depletion of the global address space
- 240/4 being proposed as new private address space

IPv6 OS and Application Support

- All software vendors officially support IPv6 in their latest Operating System releases
  - Apple Mac OS X; HP (HP-UX, Tru64 & OpenVMS); IBM zSeries & AIX; Microsoft Windows XP, Vista, .NET, CE; Sun Solaris,…
  - *BSD, Linux,…
- Application Support
  - Applications must be IPv4 and IPv6 agnostic
  - User should not have to “pick a protocol”
  - Successful deployment is driven by Applications

ISP Deployment Activities

- Several Market segments
  - IX, Carriers, Regional ISP, Wireless
- ISP have to get an IPv6 prefix from their Regional Registry
  - www.ripe.net/ripencc/mem-services/registration/ipv6/ipv6allocs.html
- Large carriers planning driven by customer demand:
  - Some running trial networks (e.g. Sprint)
  - Others running commercial services (e.g. NTT, FT,…)
- Regional ISP focus on their specific markets
- Much discussion by operators about transition
  - www.civil-tongue.net/6and4/
  - www.nanog.org/mtg-0710/presentations/Bush-v6-op-reality.pdf

Why not use Network Address Translation?

- Private address space and Network address translation (NAT) could be used instead of IPv6
- But NAT has many serious issues:
  - Breaks the end-to-end model of IP
  - Layered NAT devices
  - Mandates that the network keeps the state of the connections
  - How to scale NAT performance for large networks?
  - Makes fast rerouting & multihoming difficult
  - Service provision inhibited
  - RFC2993 – architectural implications of NAT

Where we are heading

- There is a need for a larger address space
  - IPv6 offers this – will eventually replace NAT
  - But NAT will be around for a while too
  - Market for IPv4 addresses looming also
- Many challenges ahead
  - Internet infrastructure will become dual-stack
  - Using IPv4 & IPv6 at the same time

DUAL STACK NETWORKS

Transition versus Replacement

- Internet is transitioning to IPv6
  - NOT switching off IPv4 one day, switching on IPv6 the next
  - BUT running both protocols at the same time
    - For several years to come
- Dual-stack networks
  - Routers handle IPv4 and IPv6 at the same time
  - As they handled IPv4, IPX, DECnet, AppleTalk, X.25 etc in the past
  - Co-existence, in the same way that IPv4 co-existed with the other protocols in the early days of the Internet

Co-existence

- Clients

  Choose IPv6 transport if available and destination has IPv6 address

  Otherwise continue using IPv4

- Routers

  IPv6 and IPv4

- Servers

  Must support both IPv4 and IPv6

- And this is the challenge facing the industry over the next 5 to 10 years

The challenge

- Dual stack hosts
  - Can see both IPv4 and IPv6 destinations
- IPv6-only hosts
  - How do they access IPv4-only hosts?
  - NAT-PT (translates between IPv4 and IPv6)
- ISP with little or no IPv4 address space
  - NAT IPv4 into upstream provider's public space?
    - Use IPv6 to access IPv6-only destinations
    - What type of NAT device is this?
  - Or **buy** "right to use" IPv4 address space from other ISPs?

Summary

- Depletion of IPv4 IANA pool is reality
- Internet industry needs to consider future strategy to ensure continued growth of the Internet beyond IPv4 depletion

  i.e. what is the ISP business plan beyond 2010?

State of the Backbone

Internet Routing System

- Currently 257000 prefixes
  - Late 1990s all over again
  - Adding over 2500 prefixes per month
  - Router with 512Mbytes RAM mandatory now
- 17.7% of IPv4 address space left
  - 24.6% was left 12 months ago (down 6.9%)
  - 30.9% was left 12 months before that (down 6.3%)
- 16000 16-bit ASNs left
  - Approx 4000 assigned per year

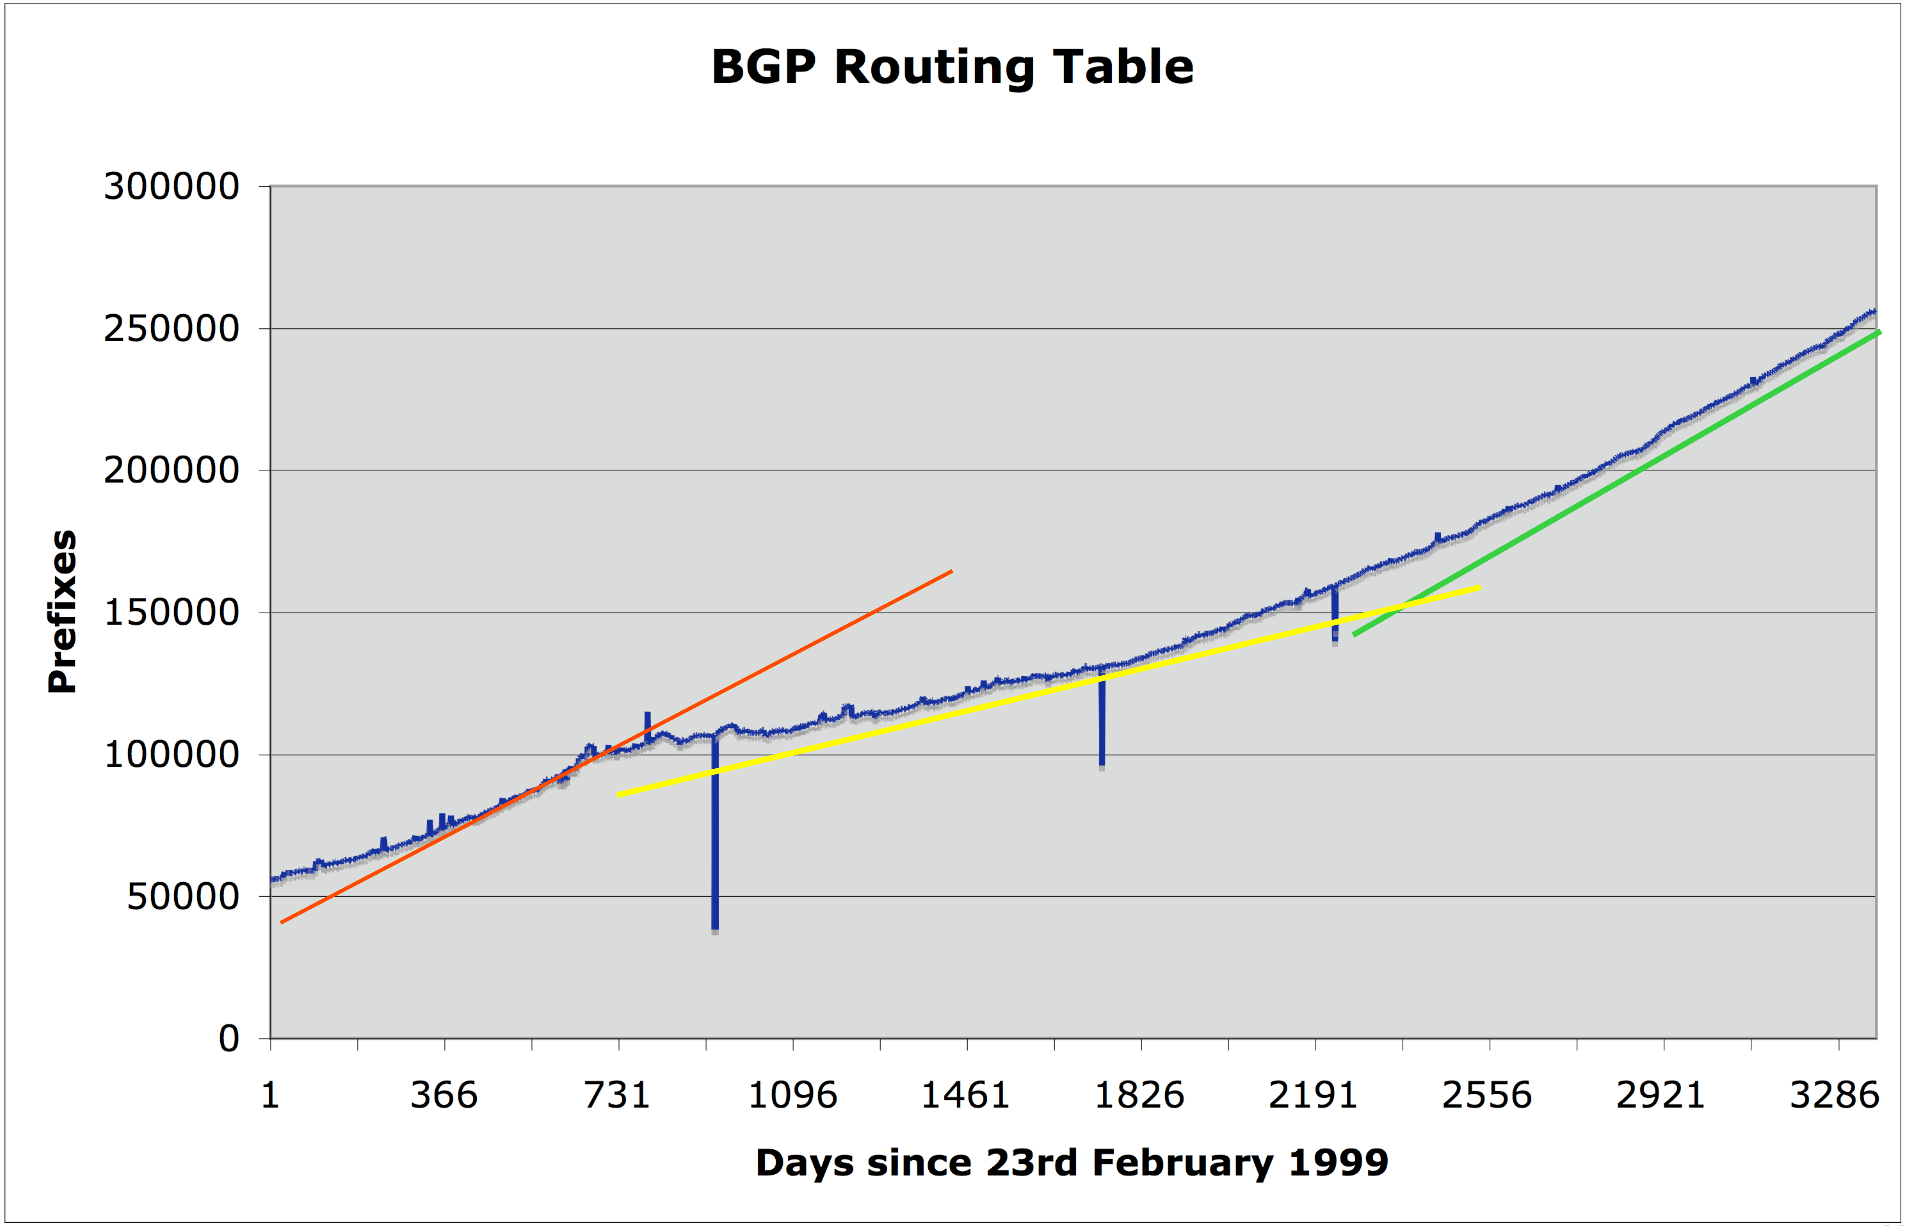
BGP Routing Table
300000
250000
200000
150000
100000
50000
0
Prefixes
1
366
731
1096
1461
1826
2191
2556
2921
3286
Days since 23rd February 1999

Internet Routing System

- Who is contributing most to the size of the table?

  Not only the biggest ISPs…!

ASN	No of nets	/20 equiv	MaxAgg	Description
6389	2208	3103	182	bellsouth.net, inc.
4755	1642	387	89	Videsh Sanchar Nigam Ltd. Aut
4323	1471	1045	377	Time Warner Telecom
2386	1439	658	872	AT&T Data Communications Serv
7018	1372	5815	981	AT&T WorldNet Services
8151	1238	2464	227	UniNet S.A. de C.V.
11492	1228	148	12	Cable One
9498	1174	550	61	BHARTI BT INTERNET LTD.
9583	1112	111	412	Sify Limited
17488	1078	71	94	Hathway IP Over Cable Interne

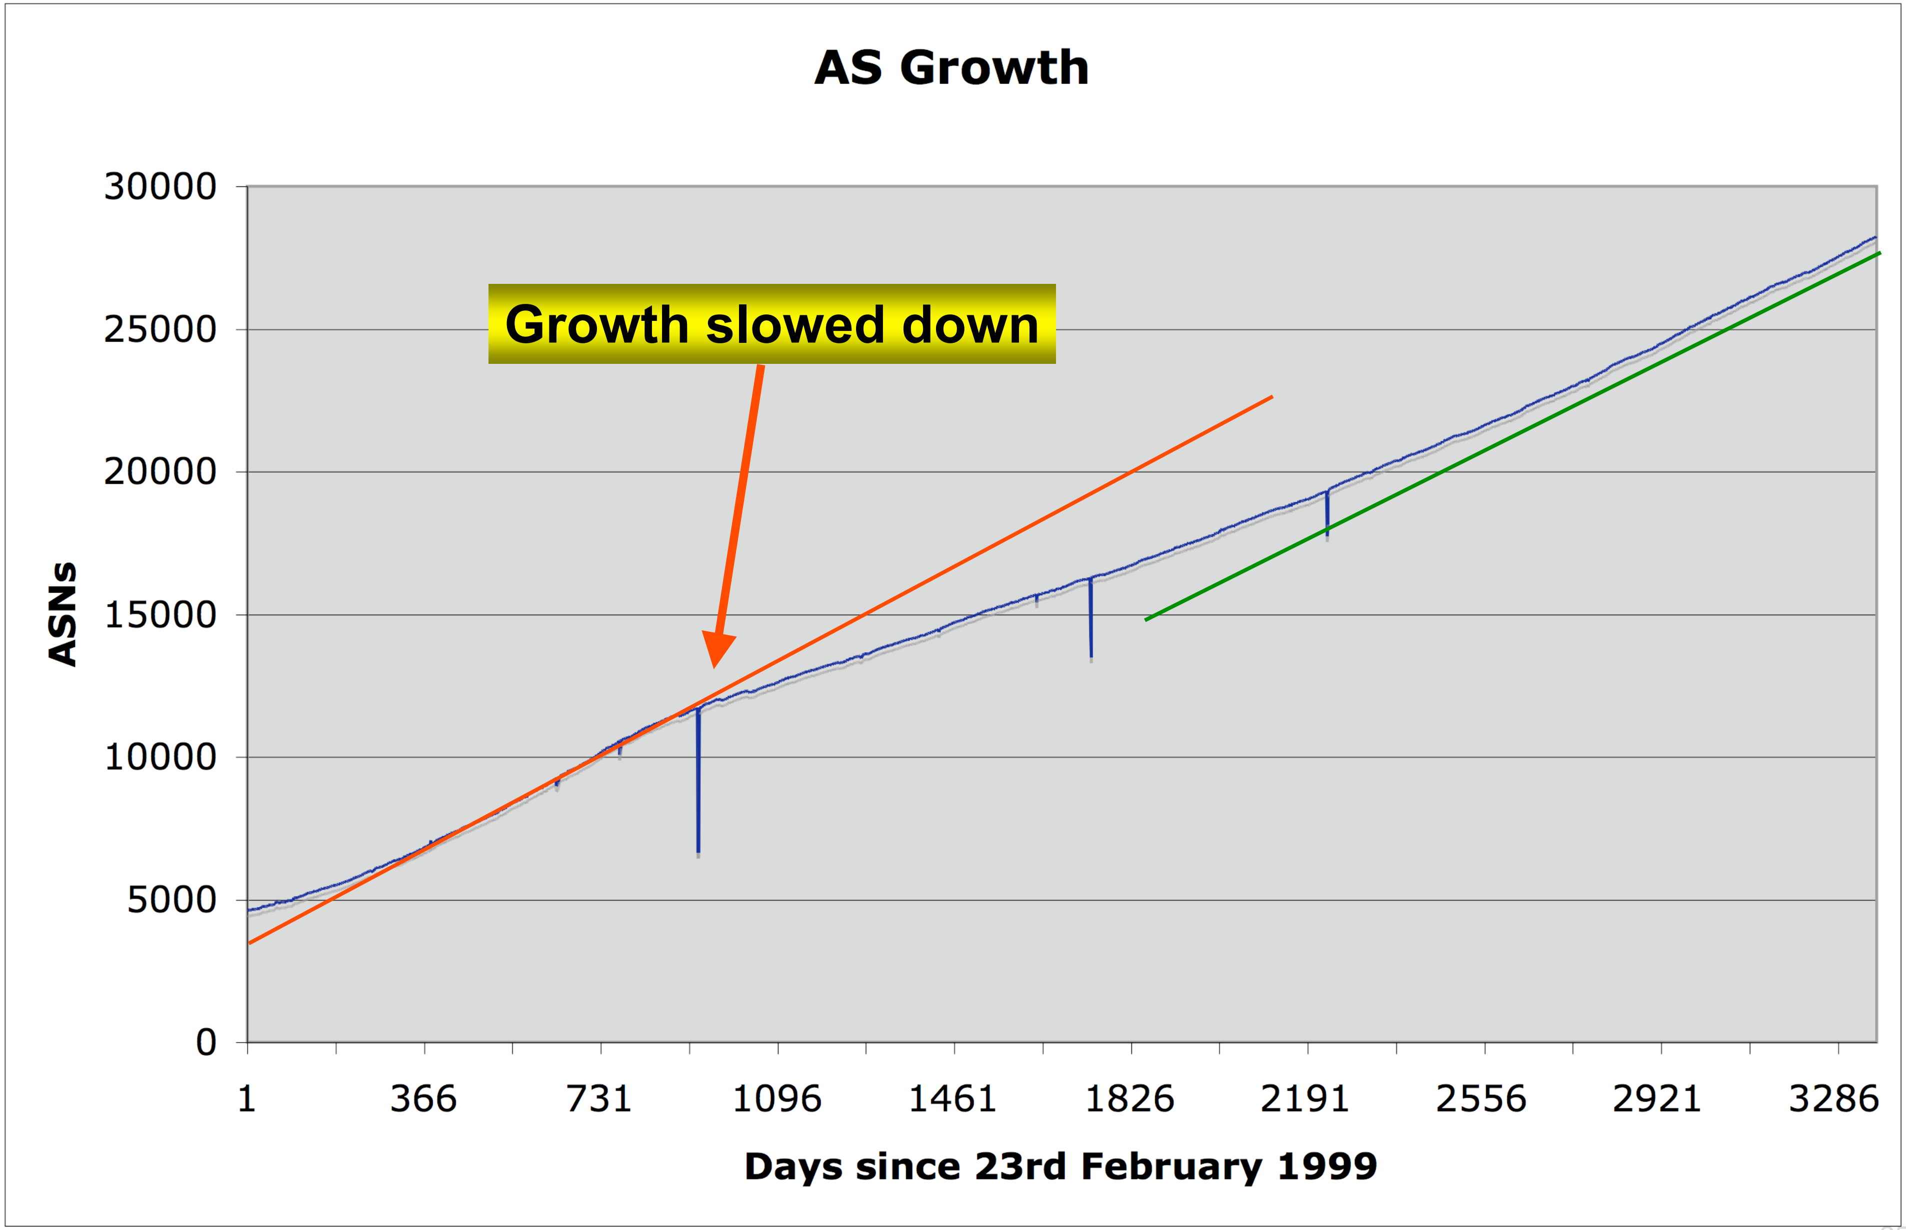
AS Growth
Growth slowed down
30000
25000
20000
15000
10000
5000
0
ASNs
1
366
731
1096
1461
1826
2191
2556
2921
3286
Days since 23rd February 1999

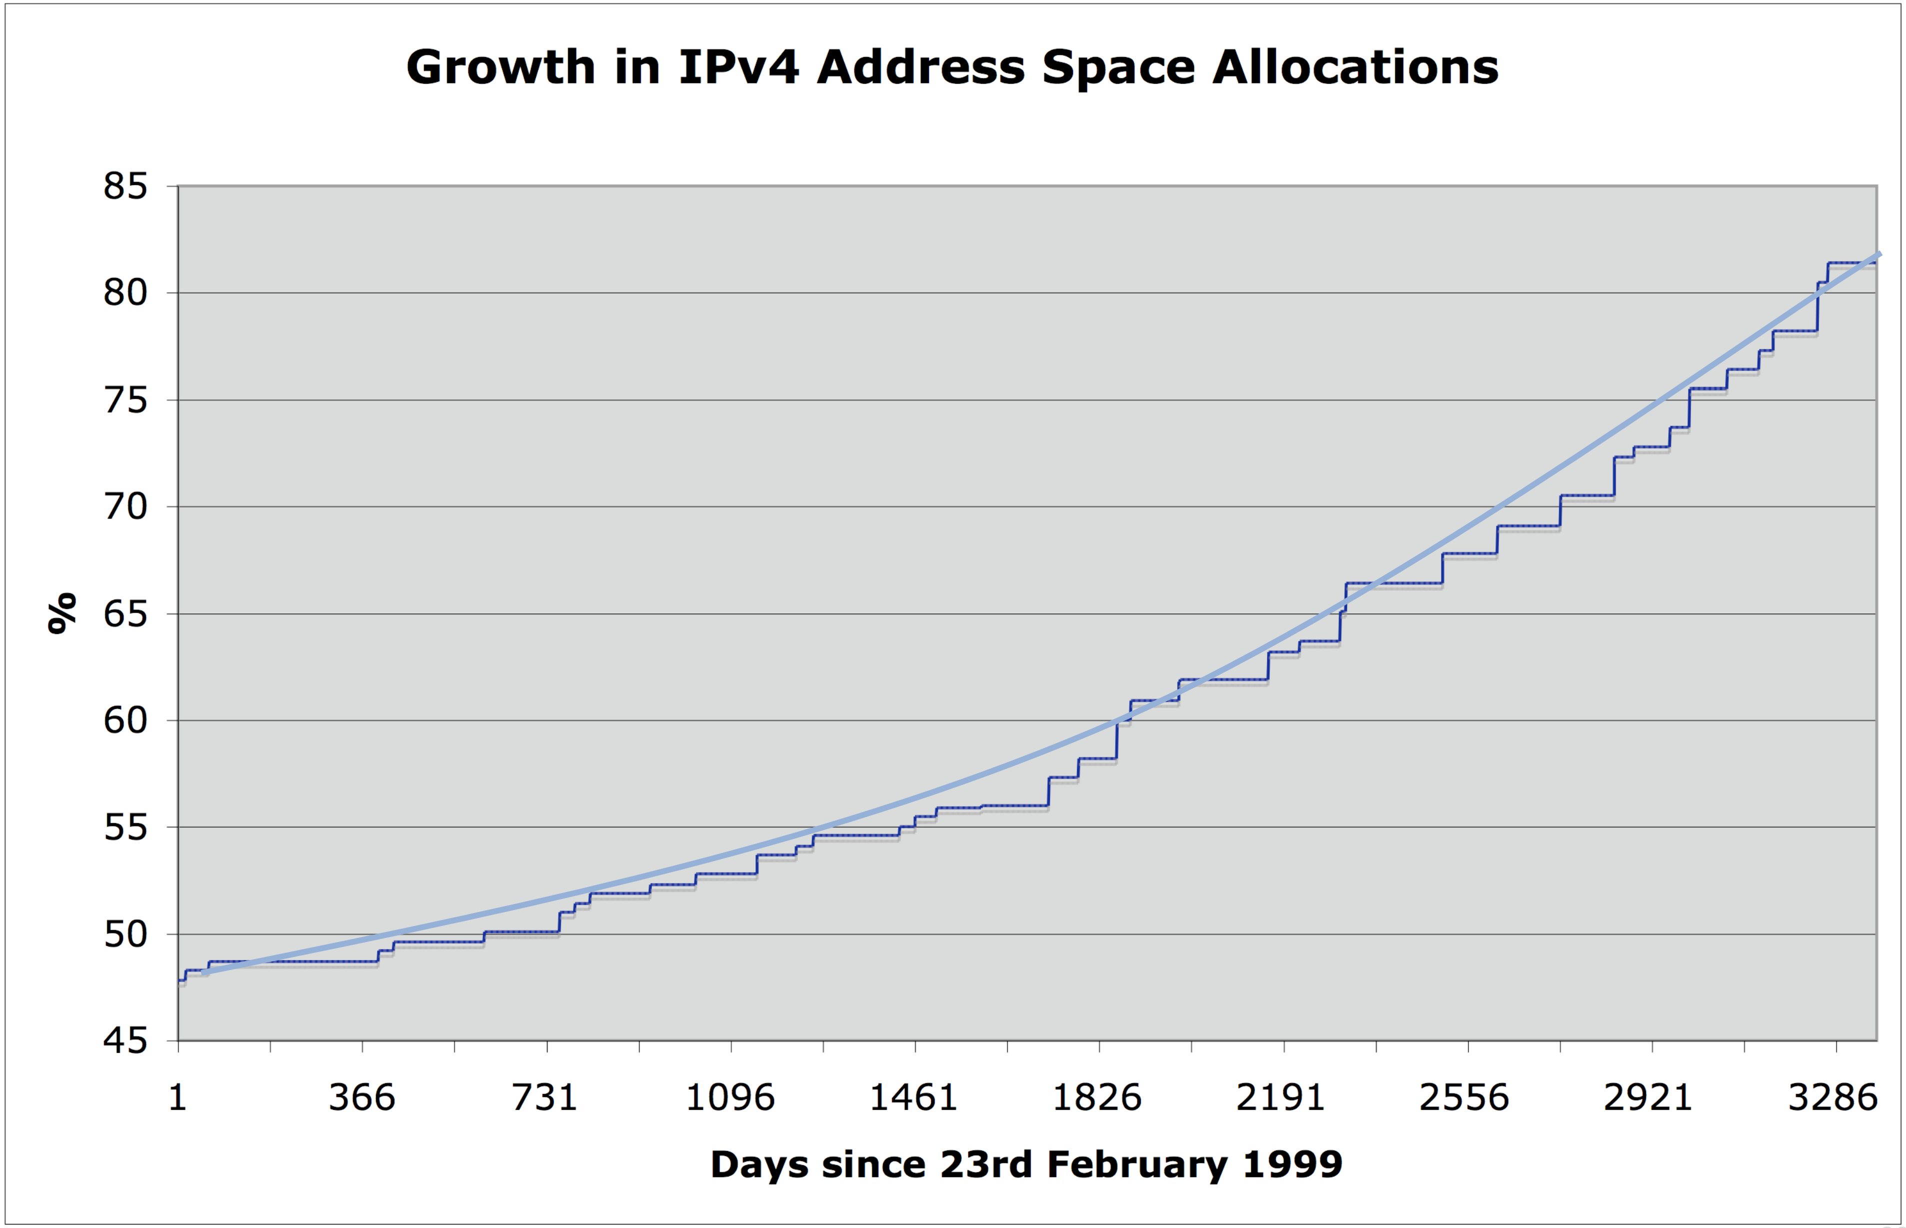
Growth in IPv4 Address Space Allocations
85
80
75
70
65
60
55
50
45
%
1
366
731
1096
1461
1826
2191
2556
2921
3286
Days since 23rd February 1999

Internet Routing System

- Deaggregation is a serious concern
  - Not just Developing Regions though
  - ChinaNet – no longer in top ten following hard work cleaning up
  - Covad – deliberate deaggregation for "security reasons"
  - India is in a bad way

ASN	No of nets	Net Savings	Description
4755	1642	1553	Videsh Sanchar Nigam Ltd. Aut
11492	1228	1216	Cable One
9498	1174	1113	BHARTI BT INTERNET LTD.
4323	1471	1094	Time Warner Telecom
18566	1044	1034	Covad Communications
8151	1238	1011	UniNet S.A. de C.V.
17488	1078	984	Hathway IP Over Cable Interne
1785	1028	923	AppliedTheory Corporation
22773	946	884	Cox Communications, Inc.
6478	951	799	AT&T Worldnet Services

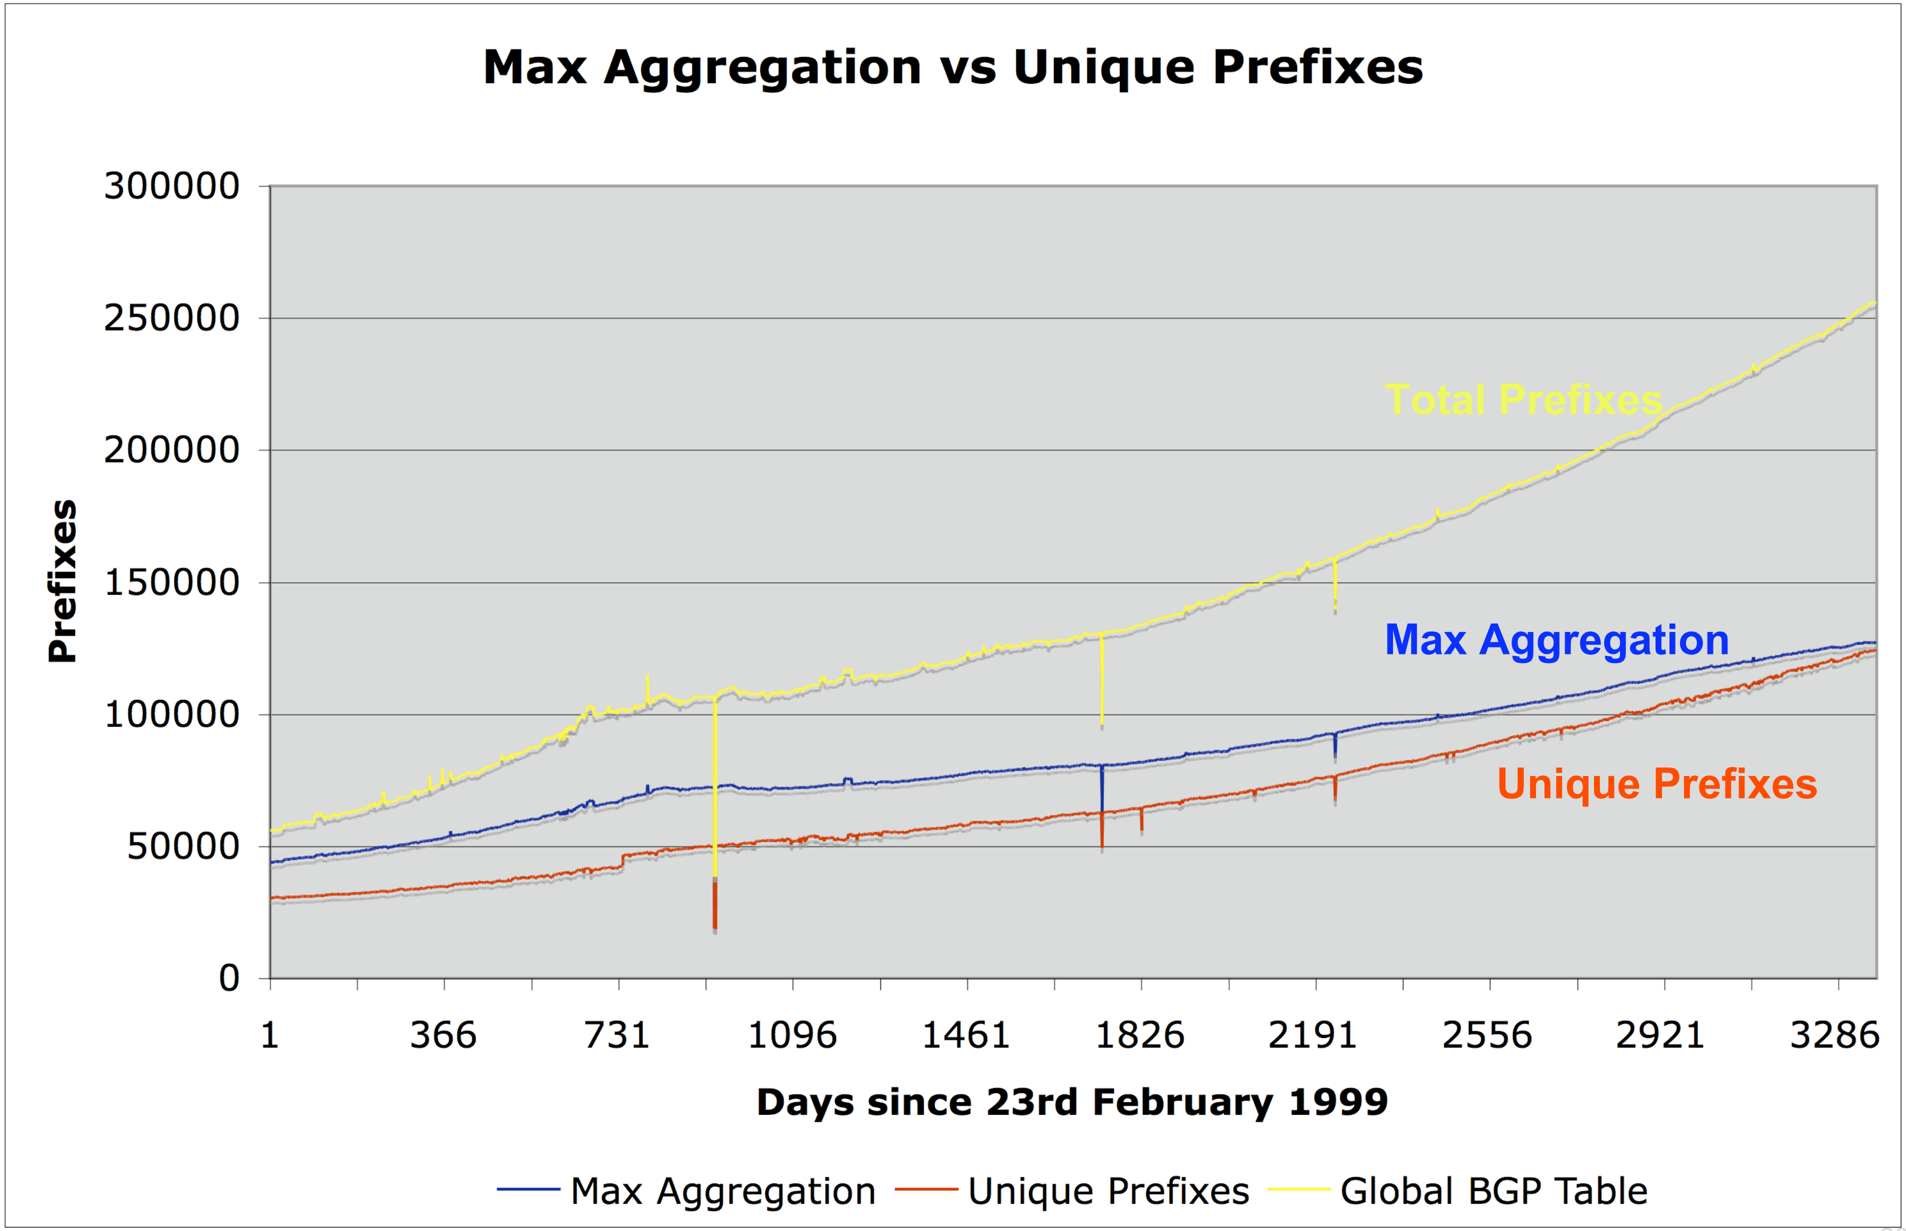

Max Aggregation vs Unique Prefixes
300000
250000
200000
150000
100000
50000
0
Prefixes
Total Prefixes
Max Aggregation
Unique Prefixes
1
366
731
1096
1461
1826
2191
2556
2921
3286
Days since 23rd February 1999
Max Aggregation
Unique Prefixes
Global BGP Table

Route Aggregation Recommendations

- RIPE Document — RIPE-399

  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
  - Training, tutorials, Project 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
- (Oh, and all 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”

June 2008

Total Prefixes

- Global BGP Table
  257k prefixes
- Europe & Middle East
  56k prefixes
- North America
  118k prefixes
- Asia & Pacific
  59k prefixes
- Africa
  4k prefixes
- Latin America & Caribbean
  20k prefixes

Deaggregation Factor

- Global Average
  2.01
- Europe & Middle East
  1.64
- North America
  1.81
- Asia & Pacific
  2.67
- Africa
  3.25
- Latin America & Caribbean
  3.95

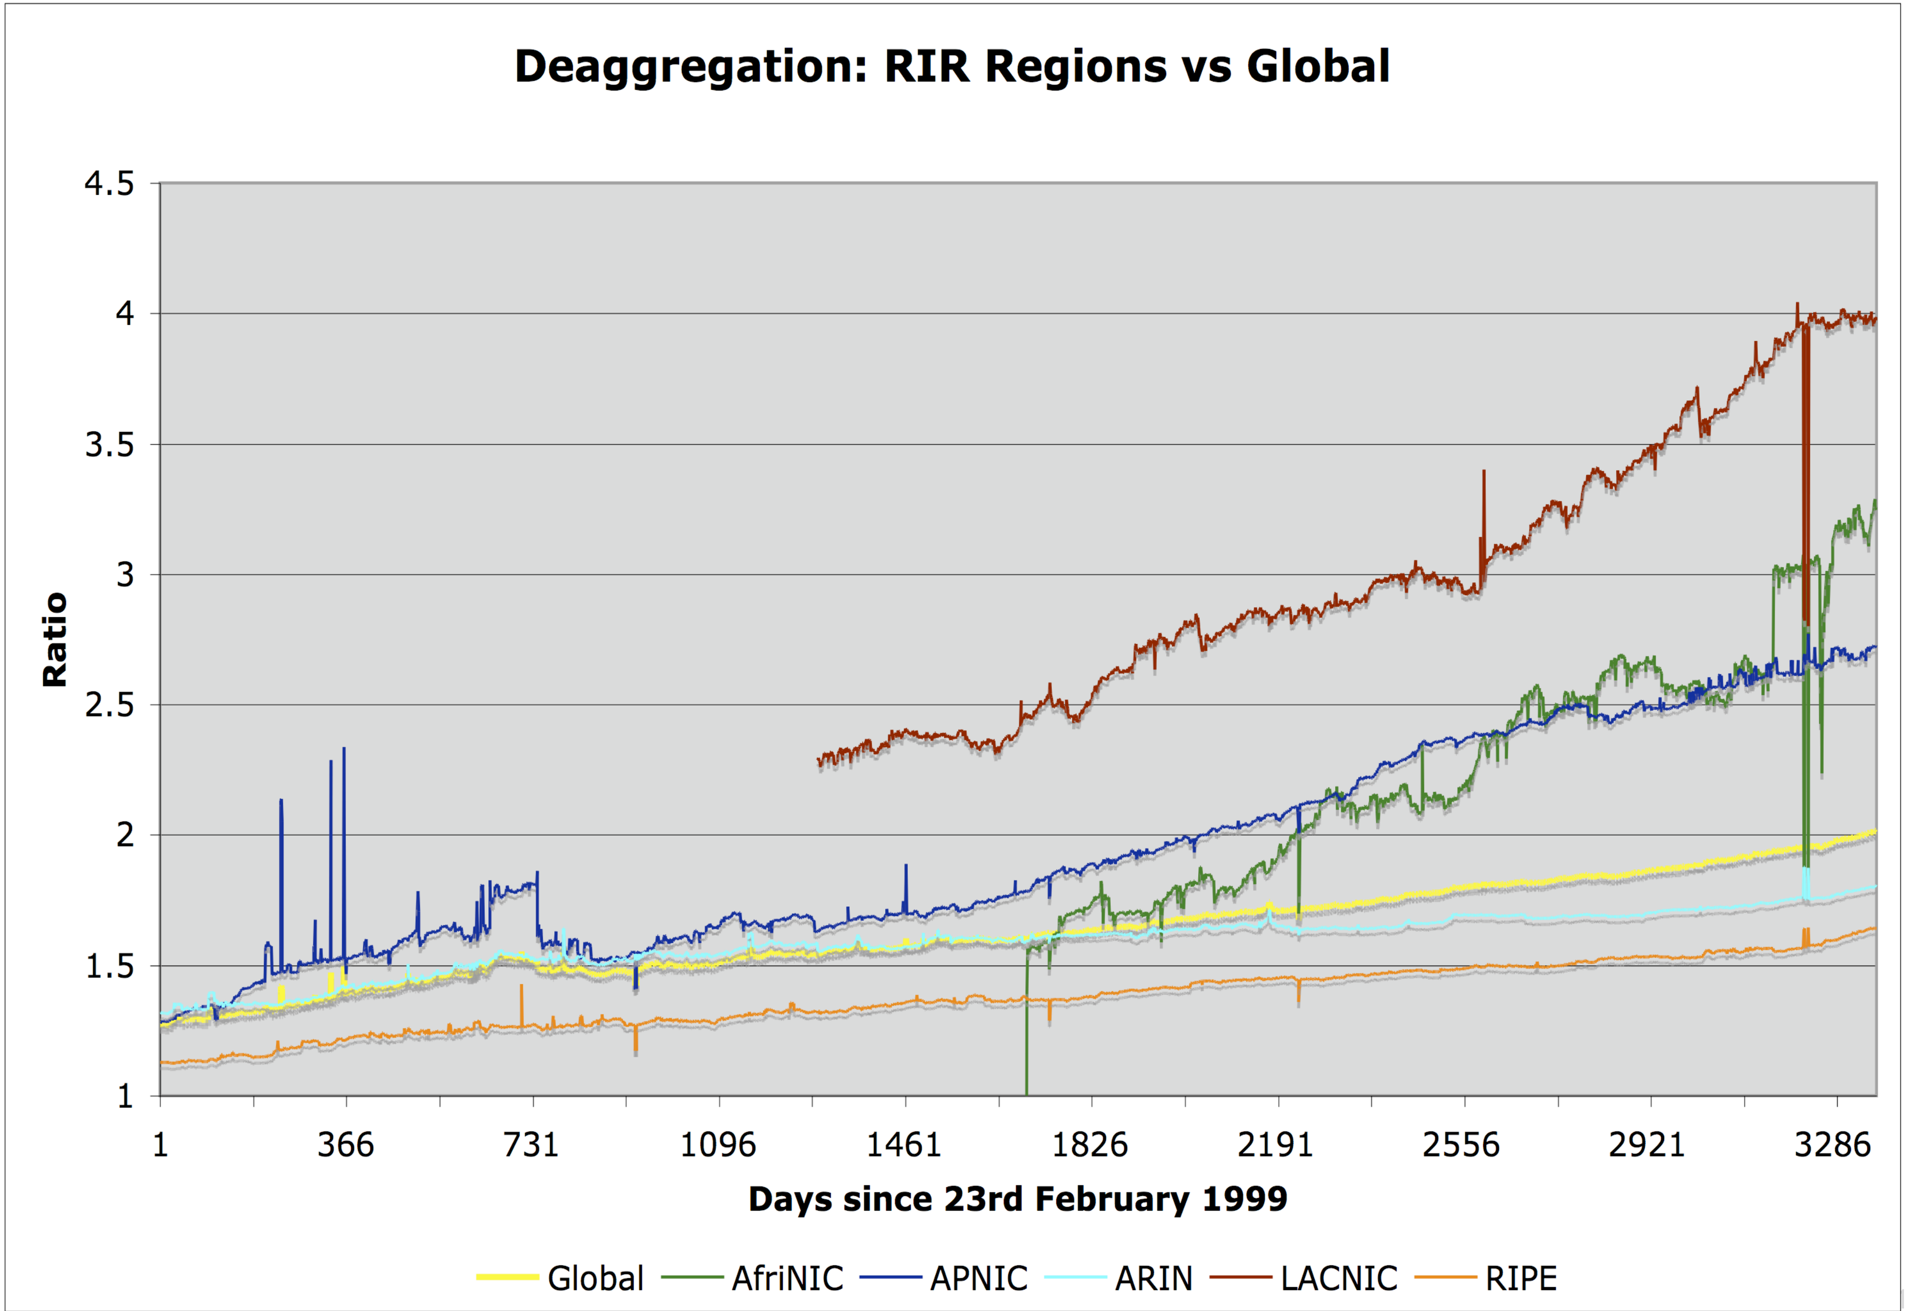
Deaggregation: RIR Regions vs Global
Ratio
4.5
4
3.5
3
2.5
2
1.5
1
1
366
731
1096
1461
1826
2191
2556
2921
3286
Days since 23rd February 1999
Global
AfriNIC
APNIC
ARIN
LACNIC
RIPE

Observations

- Range of operational “practices” between RIR regions
  - “Newer” Internet is growing rapidly
    - As is the deaggregation there
- 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

IPv6 Routing Table?

- DFZ is around 1241 prefixes (Cisco IPv6 gateway)
- OCCAID backbone across US, edge into Japan and EU
  - Has normalised most remaining “weird” transit paths
- ISPs waiting for customers, customers waiting for ISPs
- IPv6 BGP Table report by Geoff Huston:
  - http://bgp.potaroo.net/v6/as2.0/index.html
  - End of 6bone very noticeable early 2006
  - Rapid growth in table from mid 2007

IPv6 Routing Table Growth

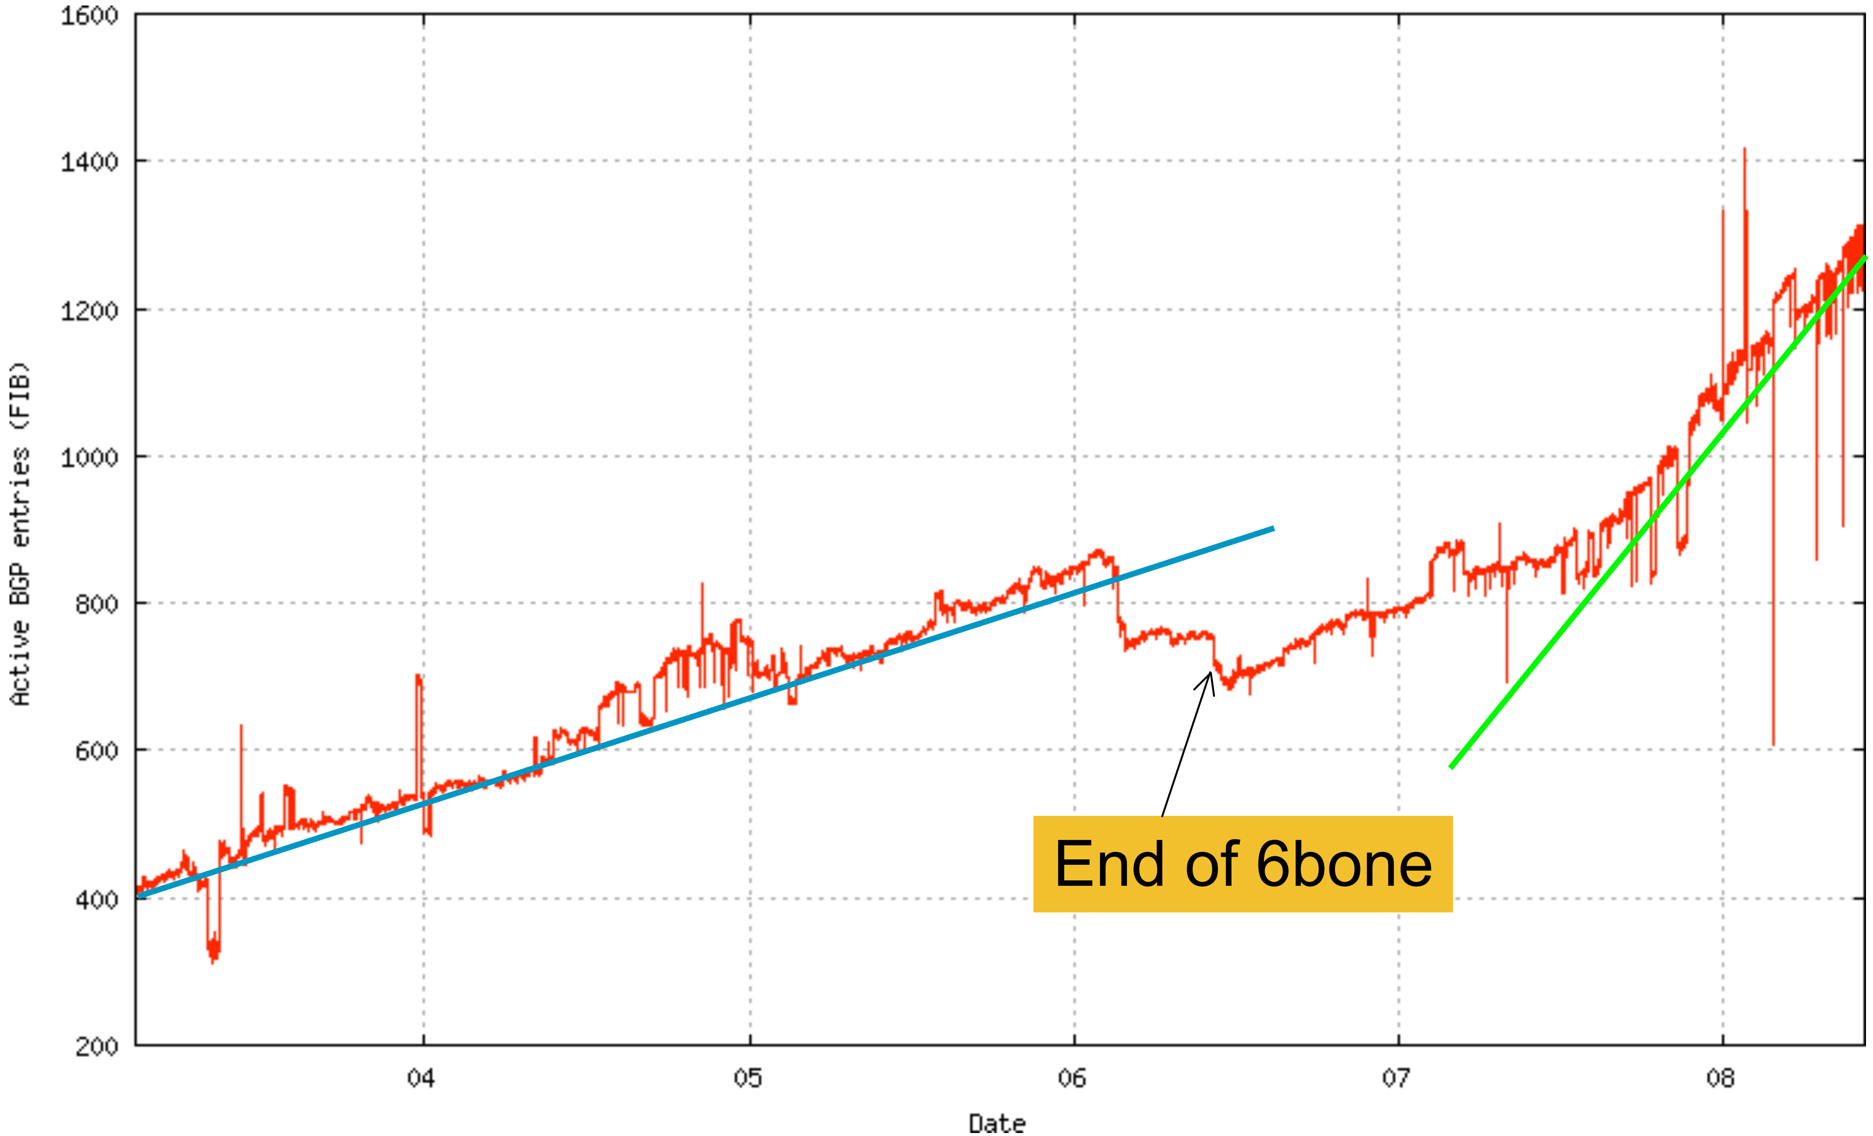

Summary

- Size of routing table is growing rapidly again
- Deaggregation is growing rapidly
  - Especially in newer parts of Internet
- IPv6 routing table growth is now showing signs of picking up momentum

Internet Operations Groups

Internet Operations Groups

- Where network engineers and operators meet their colleagues
  - Peering & Business relationships
  - Industry relationship
  - Technology discussions
  - Operational best practices
  - Compare experiences (supplier, operational,…)
  - Purchasing decisions influenced
  - Routing software feature requests worked out
  - Jobs fair
  - **Keeping the Internet Working**

Regional Internet Operations Groups

- NANOG – North America
- APRICOT – Asia & Pacific Region
- SANOG – South Asia
- MENOG – Middle East
- PacNOG – Pacific Islands
- RIPE – Europe
- AfNOG – Africa
- LACNOG – Latin America

Country Network Operations Groups

- NZNOG – New Zealand
- JANOG – Japan
- CNNOG – China
- PhNOG – Philippines
- AusNOG – Australia
- SWINOG – Switzerland
- UKNOF – United Kingdom
- …

New NOGs

- NOG creation is a recent phenomenon
  - Local Language
  - Local Culture – Internet is not just American culture
  - Local Needs
- SANOG and NZNOG are common models
  - Low overheads, low fees
  - Too much temptation to introduce bureaucracy in newer NOGs
- Potential newcomers:
  - Central Asia
  - Caribbean
  - Many countries…

More information

- Other NOGs:

  www.bugest.net/nogs.html

- Meeting planning:

  "Network Education and Training Calendar of Events" maintained by the NSRC

  ws.edu.isoc.org/calendar

  Intention is to avoid conflicts and overlaps

Anycast Root Nameservers & DNS

Anycast DNS

- Anycast:
  - Multiple instances of the identical service visible in multiple parts of the Internet
  - Individual devices share the same global IP address
  - Routing system chooses service closest to the end-user
- DNS Anycast Advantages
  - Insulates DNS against DoS attacks
  - Improves DNS lookup performance
  - Located at IXPs meaning low latency to end users
  - Lower cost than providing one ultimate resilient node one place

Anycast DNS

- DNS Root Nameservers

  Many of the operators now anycast the DNS service (at least F-root, I-root, J-root, K-root and M-root are visible in many parts of the world)

    http://root-servers.org/

  Root Nameservers are also reachable over IPv6 now

    The glue added earlier this year by IANA

    www.iana.org/reports/root-aaaa-announcement.html

- GTLD and CCTLD Nameservers

  Many cctld and gtld operators now anycast their DNS services (e.g. Verisign, PCH)

Other new things

More changes afoot…

Autonomous Systems

- An AS is a collection of networks with same routing policy, ownership, trust, admin control
- An ASN is a 16-bit integer identifying the AS

  1-64511 are assigned by the Regional Internet Registries

  64512-65534 are private ASNs and should never be used on the Internet

  0 and 65535 are reserved

- Current allocations up to 48127 have been made to the RIRs

  Only another 16000 left

ASN status

- The pool of 16-bit ASNs will soon be exhausted

  Analysis at www.potaroo.net/tools/asns/

  Estimates are that the 16-bit ASN pool will be exhausted late 2010

- Work started in 2001 to extend the ASN pool to 32-bits

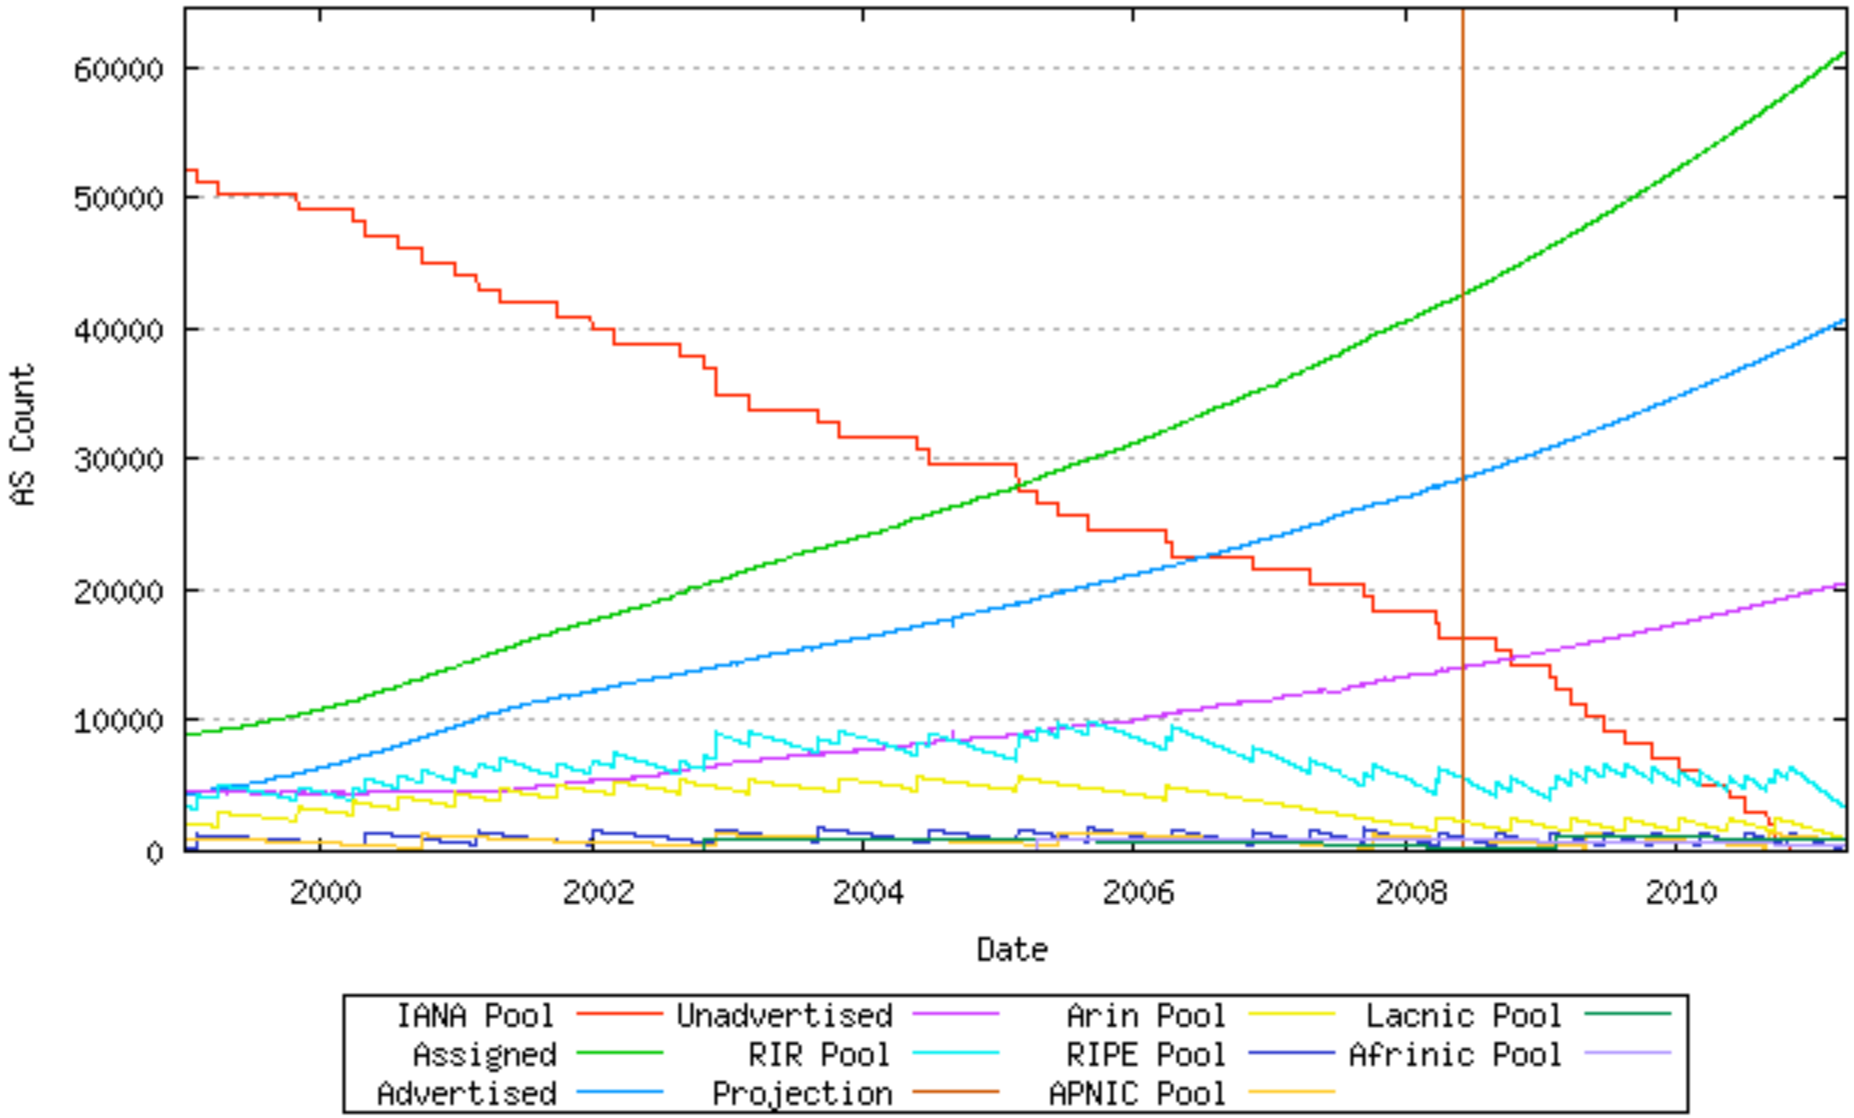

Source: http://www.potaroo.net/tools/asns/fig28.png

32-bit ASNs

- Standards documents (drafts)

  Description of 32-bit ASNs

  www.rfc-editor.org/rfc/rfc4893.txt

  Proposal for the representation of 32-bit ASNs

  www.ietf.org/internet-drafts/draft-michaelson-4byte-as-representation-05.txt

  New extended community

  www.ietf.org/internet-drafts/draft-rekhter-as4octet-ext-community-01.txt

- AS 23456 is reserved as interface between 16-bit and 32-bit ASN world

Getting a 32-bit ASN

- Sample RIR policy
  - www.apnic.net/docs/policy/asn-policy.html
- From 1st January 2007
  - 32-bit ASNs available on request
- From 1st January 2009
  - 32-bit ASNs assigned by default
  - 16-bit ASNs only available on request
- From 1st January 2010
  - No distinction – ASNs assigned from 32-bit pool

Representation

- 32-bit ASNs extend the pool:
  - 0-65535 extended to 0-4294967295
- Still discussion on representation of 65536-4294967295 range
- Some favour X.Y:
  - For 65536-4294967295, or
  - For 0-4294967295
  - But how will regular expressions work?
- Some favour traditional format
  - But gets bulky to handle when numbers get v big

IANA Assignments

- 0.0 - 0.65535 16-bit ASN block
- 2.0 - 2.1023 APNIC
- 3.0 - 3.1023 RIPE NCC
- 4.0 - 4.1023 LACNIC
- 5.0 - 5.1023 AfriNIC
- 6.0 - 6.1023 ARIN
- Remainder are reserved or held by IANA

Changes (1)

- 32-bit ASNs are backwardly compatible with 16-bit ASNs
- There is no flag day
- You do NOT need to:
  - Throw out your old routers
  - Replace your 16-bit ASN with a 32-bit ASN

Changes (2)

- You do need to be aware that:
  - Your customers will come with 32-bit ASNs
  - ASN 23456 is not a bogon!
  - You will need a router supporting 32-bit ASNs to use a 32-bit ASN
- If you have a proper BGP implementation, 32-bit ASNs will be transported silently across your network

How does it work (1)?

- Local router only supports 16-bit ASN
- Remote router uses 32-bit ASN
- BGP peering initiated:

    Remote asks local if 32-bit supported (BGP capability negotiation)

    When local says "no", remote then presents AS23456

    Local needs to be configured to peer with remote using AS23456

How does it work (2)?

- BGP peering initiated (cont):

  BGP session established using AS23456

  32-bit ASN included in a new BGP attribute called AS4_PATH

    (as opposed to AS_PATH for 16-bit ASNs)

- Result:

  16-bit ASN world sees 16-bit ASNs and 23456 standing in for 32-bit ASNs

  32-bit ASN world sees 16 and 32-bit ASNs

Example:

- Internet with 32-bit and 16-bit ASNs
- AS-PATH length maintained

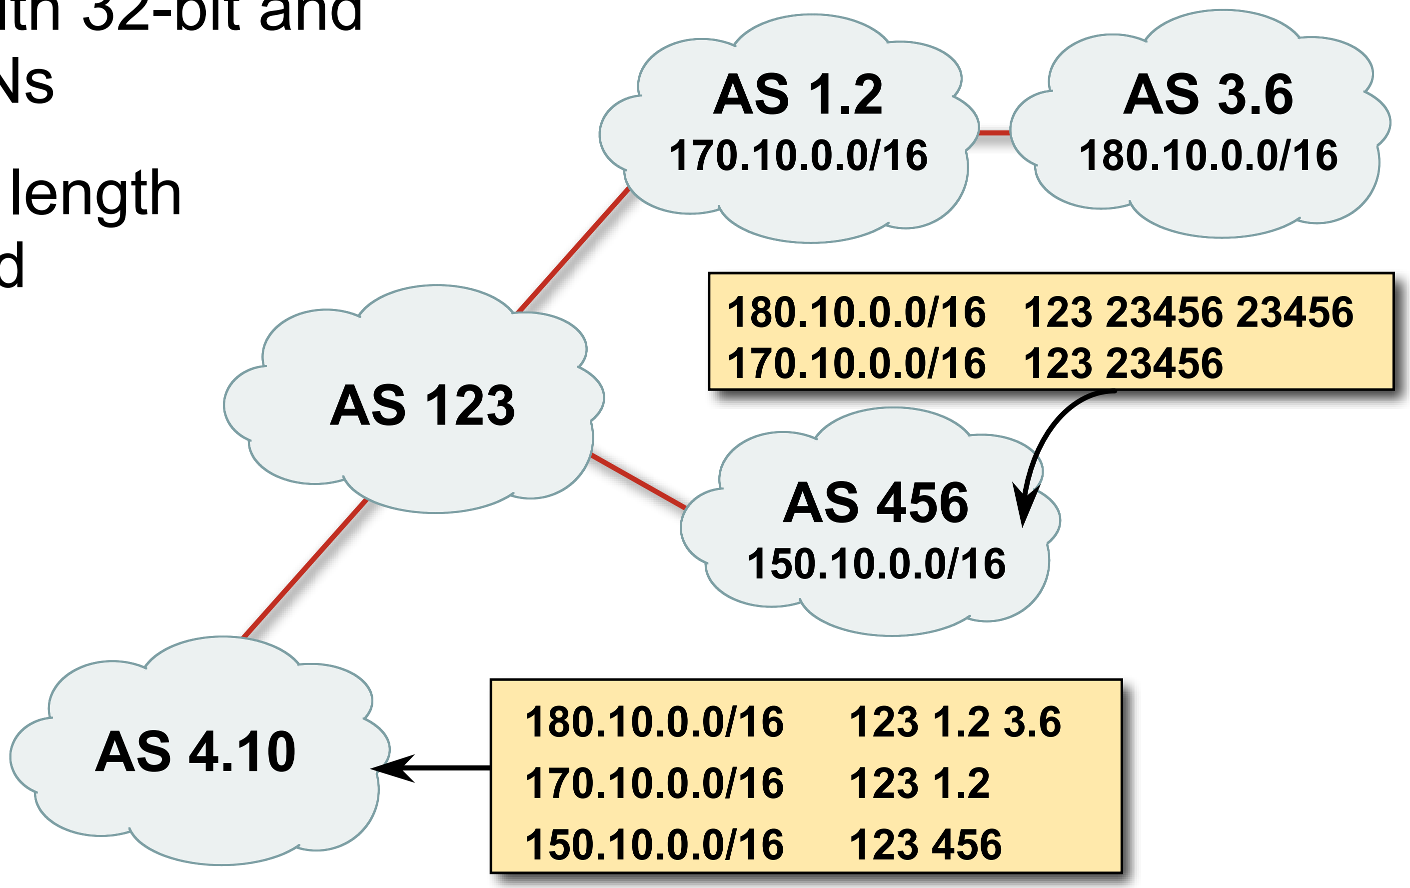

What has changed?

- Two new BGP attributes:
  - AS4_PATH
    - Carries 32-bit ASN path info
  - AS4_AGGREGATOR
    - Carries 32-bit ASN aggregator info
  - Well-behaved BGP implementations will simply pass these along if they don’t understand them
- AS23456 (AS_TRANS)

What do they look like?

- IPv4 prefix originated by AS 1.202

  In 32-bit ASN world:

```
# bgpctl show rib 203.10.62.0/24
flags: * = Valid, > = Selected, I = via IBGP, A = Announced
origin: i = IGP, e = EGP, ? = Incomplete

flags destination     gateway    lpref med aspath origin
*>    203.10.62.0/24  147.28.0.1   100   0 0.3130 0.1239 0.4637 0.1221 1.202 i
```

  In 16-bit ASN world:

```
router# sh ip bgp 203.10.62.0

   Network          Next Hop            Metric LocPrf Weight Path
*> 203.10.62.0      202.249.2.169                      0 2497 4637 1221 23456 i
```

What do they look like?

- IPv6 prefix originated by AS 2.9

```
RP/0/0/CPU0:as4byte#sh bgp ipv6 uni 2403:2000::/32
BGP routing table entry for 2403:2000::/32
Versions:
  Process            bRIB/RIB  SendTblVer
  Speaker                  93          93
Paths: (1 available, best #1)
  Not advertised to any peer
  Path #1: Received by speaker 0
  109 6175 2497 2500 18146 2.9
    2001:420:0:8001::1 from 2001:420:0:8001::1 (204.69.200.22)
      Origin IGP, localpref 100, valid, external, best
RP/0/0/CPU0:as4byte#
```

Implementation status (June 08)

- Cisco
  - IOS-XR 3.4 onwards
  - IOS coming in 12.2SRE (end 08) & 12.5T (mid 09)
- Quagga (patches for 0.99.6)
- OpenBGPd (patches for 3.9 & 4.0)
- JunOSe 4.1.0 and JunOS 9.1
- Redback

4-byte ASNs

- IPv4 BGP table:

```
   Network            Next Hop          Path
*> 66.117.63.0/24      203.119.0.116     2.0 4777 2497 2828 29748 33437 6.3 i
*> 192.26.93.0/24      203.119.0.116     2.0 4777 2497 2914 4697 2.3 i
*> 193.5.68.0/23       203.119.0.116     2.0 4777 2516 8928 8758 3.13 i
*> 193.31.7.0/24       203.119.0.116     2.0 4777 2516 13237 5539 3.3 i
*> 195.47.195.0/24     203.119.0.116     2.0 12654 3257 8495 3.16 i
*> 196.1.15.0/24       196.216.2.49      5.1 i
*> 202.255.47.0/24     203.119.0.116     2.0 4777 2516 7667 2.6 i
```

- IPv6 BGP table:

```
   Network               Next Hop            Path
*> 2001:3a0:8000::/35  2001:420:0:8001::1     6453 2914 4697 2.3 2.7 ?
*> 2001:df0:2::/48     2001:420:0:8001::1     6453 2914 4697 2.3 i
*> 2001:4810:2000::/35 2001:420:0:8001::1     30071 33437 6.3 i
*> 2403:2000::/32      2001:420:0:8001::1     6175 2497 2500 18146 2.9 i
```

ASN conclusion

- The “forgotten” depletion
- The Internet will not break
- Your network will not break
- If you have an ASN today:

    You don’t need to change anything

    32-bit ASNs appear as AS 23456

- If you have no ASN today:

    Your routers will need 32-bit ASN support after 1st January 2009

The Security Community

Internet Security

- Security is more than firewalls and virus scanners
- Security includes:
  - Bogon filters
  - Handling DoS attacks
  - NSP-SEC Community

Bogon Filters

- Bogon is a prefix which should not appear in the Internet Routing System

    e.g. 10.0.0.0/8, 169.254.0.0/16

- Filtering

    Static filters - not recommended

    Dynamic feed, e.g. from Team Cymru (**www.cymru.com**)

    www.team-cymru.org/Services/Bogons/routeserver.html

    Supplies BGP feed of the paths/prefixes to be dropped

Protecting Infrastructure

- Best Practice configuration
  - Protecting router control plane
    - Control plane separation from data path
  - Protecting BGP sessions
  - **www.team-cymru.org/ReadingRoom/Documents/**
- Out of band access
  - Remote access without depending on the IP backbone

Handling DoS attacks

- Remote triggered blackhole filtering
    - ISP routers are set to drop traffic to specified destinations
    - These "specified destinations" are injected and passed around the infrastructure by BGP routing protocol
    - They have a "next-hop" address configured to point to the null or drop interface on each core router
    - Very efficient means of sinking DOS attacks before they hit the customer
    - Injection of prefix is done by NOC operated router, injecting /32 (or whatever) of destination being attacked
- Configuration concept v similar to Team Cymru Bogon Routeserver feed

Service Provider Security

- 1990s saw rapid growth of Internet
  - Getting established and financial profit came before quality and professional service
- Early 2000s saw bigger threats to Internet infrastructure
  - DoS against routers and high profile servers/services
  - Packet amplification attacks
- Responses
  - Formation of the ISP Security Community (NSP-SEC)
  - Development of more techniques and robust network design to thwart abuse of Internet infrastructure

NSP-Sec Community

- Security specialists from ISPs work together to enhance the security of the Internet
  - Closed trust based community
  - Membership invitation only by existing members
- Membership
  - Ask your upstream ISPs to recommend you for membership
  - All ISPs should participate
  - Key security personnel of vendors participate (e.g. Cisco PSIRT)
- Regional and country NSP-Sec organisations now forming

Conclusions

Summary

- Internet Infrastructure more fragile than ever
  - New entrants have lower skills
  - Long term SPs now treat as commodity
  - Vendors focus on edge solutions less than on core infrastructure
- NOG growth outside of North America
  - Means that more operators in more parts of the world are talking to each other – the real distributed Internet

