Have We Reached 1000 Prefixes Yet?

A snapshot of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich, Germany

Oct 16th, 2007

NANOG 41, Albuquerque, New Mexico

presented by CJ Aronson

Overview

- pictures & trends
- the end of the 6bone
- numbers...
- references

Slides online at:
http://www.space.net/~gert/RIPE/NANOG41-v6-table/

Graphics: Total Prefixes - 6 years

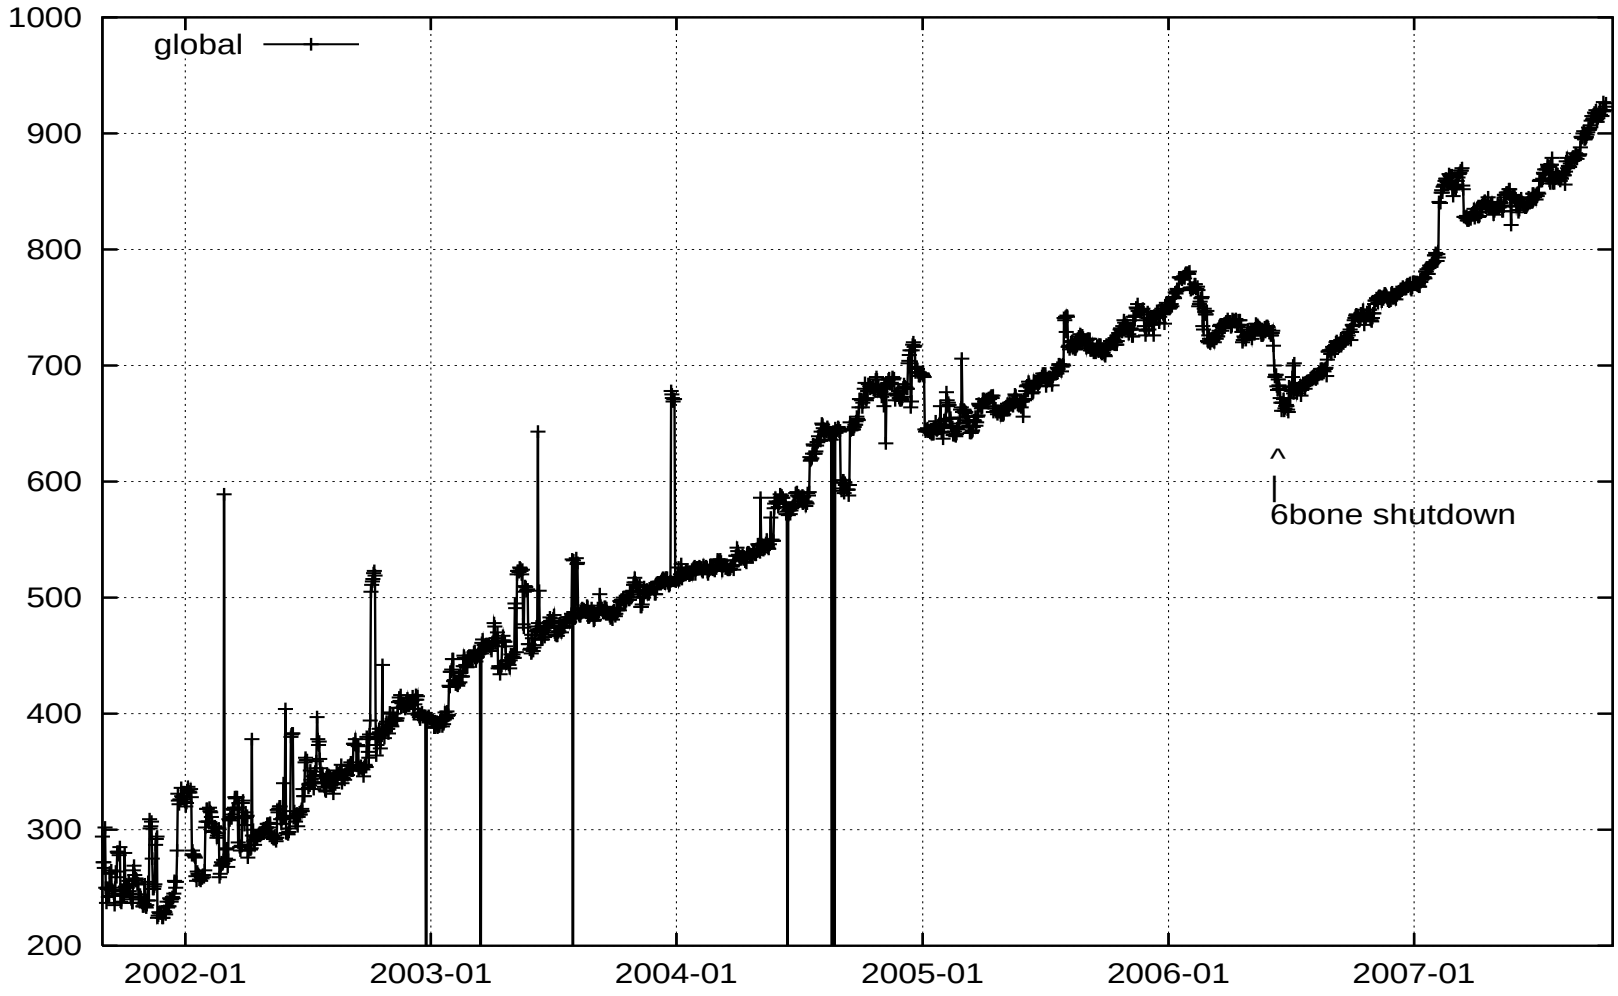

Graphics: RIR vs. 6Bone Prefixes - 6 years

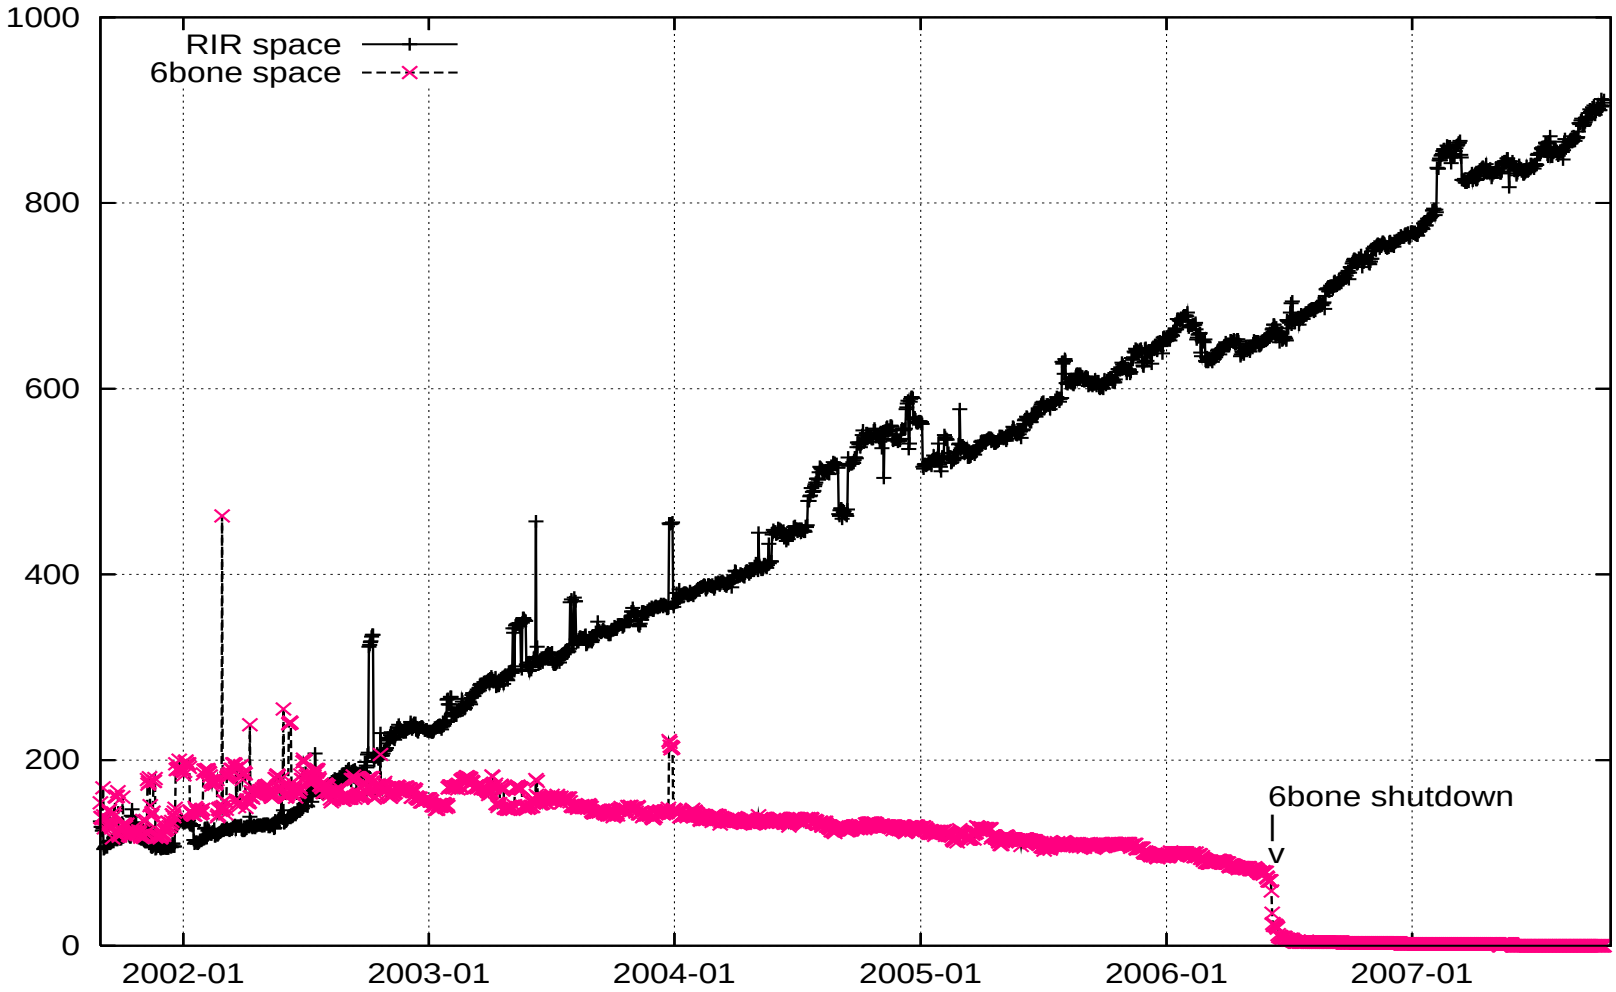

Graphics: zoom into last 6 months

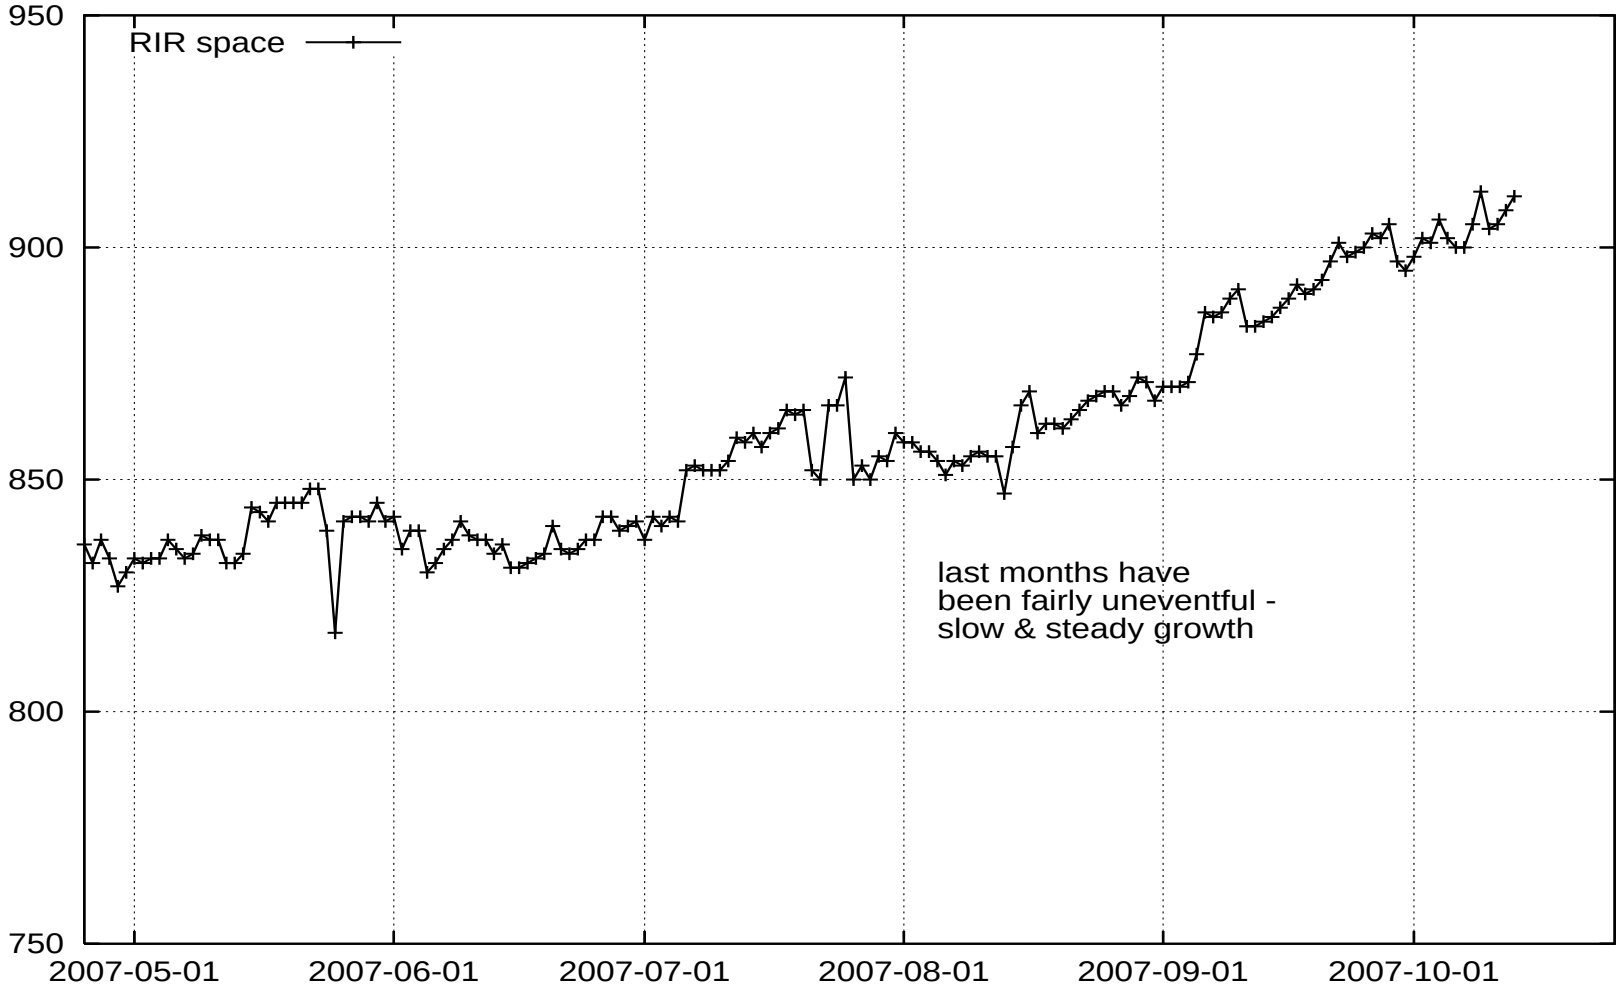

Graphics: The End Of The 6bone

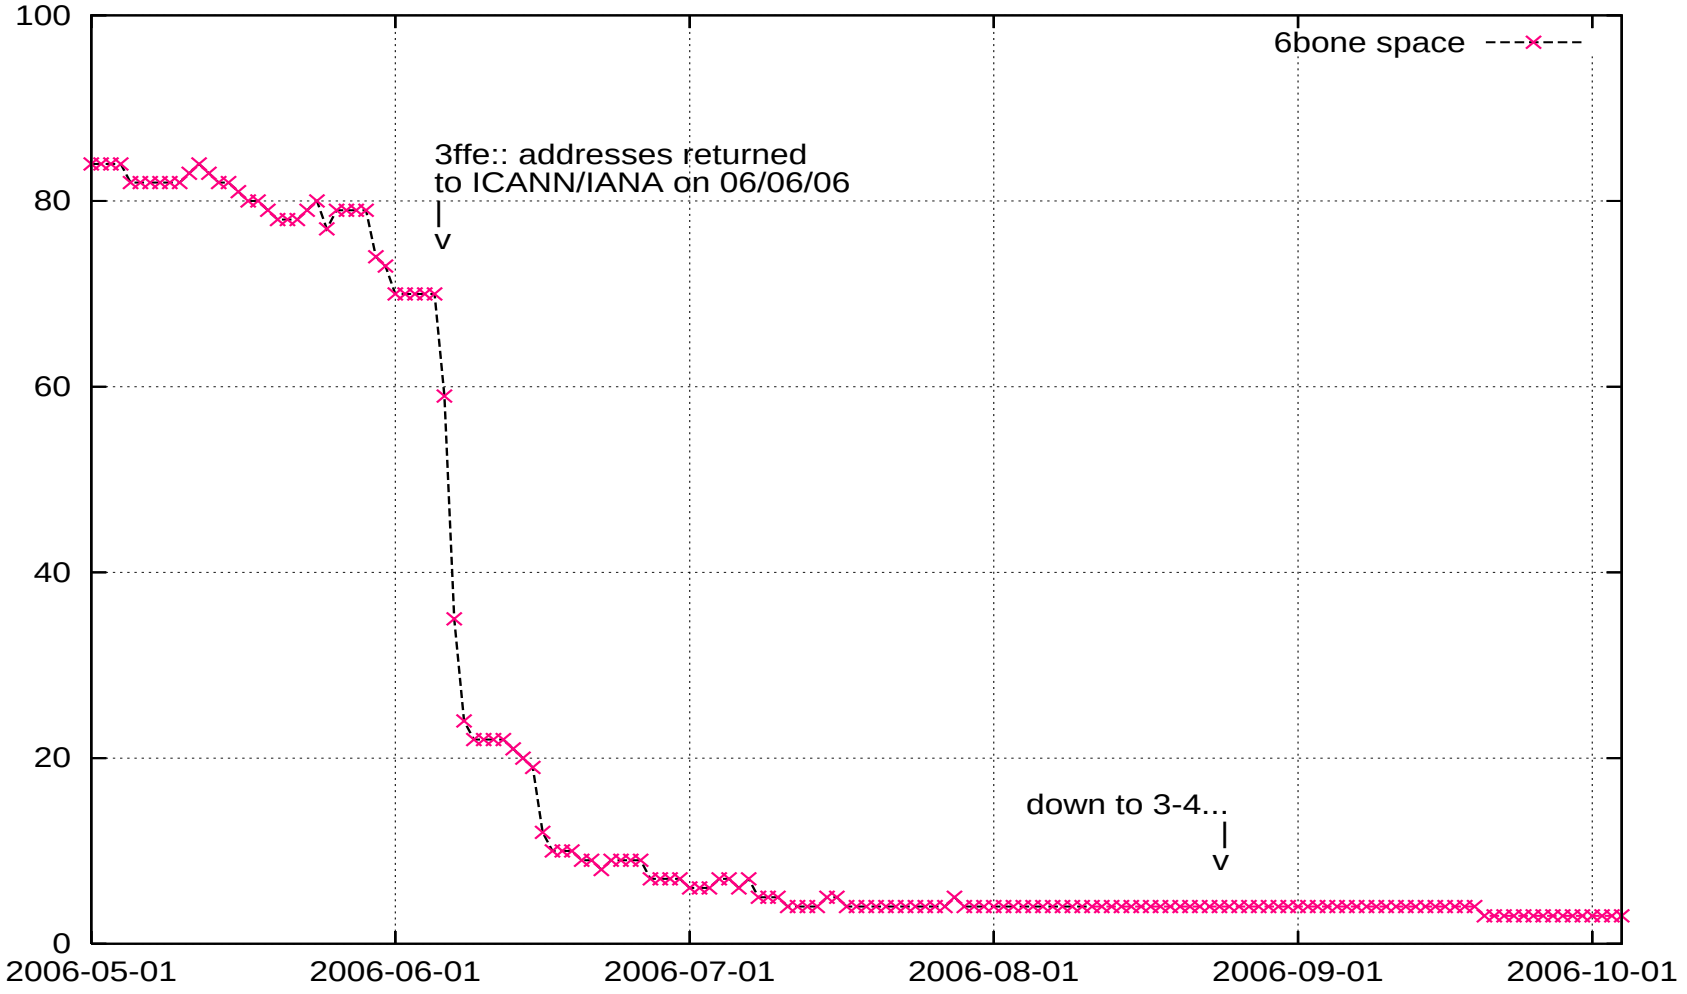

The End Of The 6bone

- on 06/06/06, the 3FFE addresses allocated to the 6Bone test network have been returned to ICANN/IANA (rfc3701)
- this means: there are no official address holders from 3FFE anymore, anybody still announcing space is an address hijacker
- at AS 5539, there are no 3FFE prefixes visible anymore :-)
- GRH (grh.sixxs.net) still sees one single path:

```
*  3FFE::/24          2A01:B8::E              8978 5609 4555 i
```

- 8978 = vatican.it, 5609 = cselt.it – out of clue error?
- still relevant: please stop using 3FFE transfer networks
- please *STOP* giving transit to 3FFE announcements!

Graphics: trends? (RIR prefixes, 24 months)

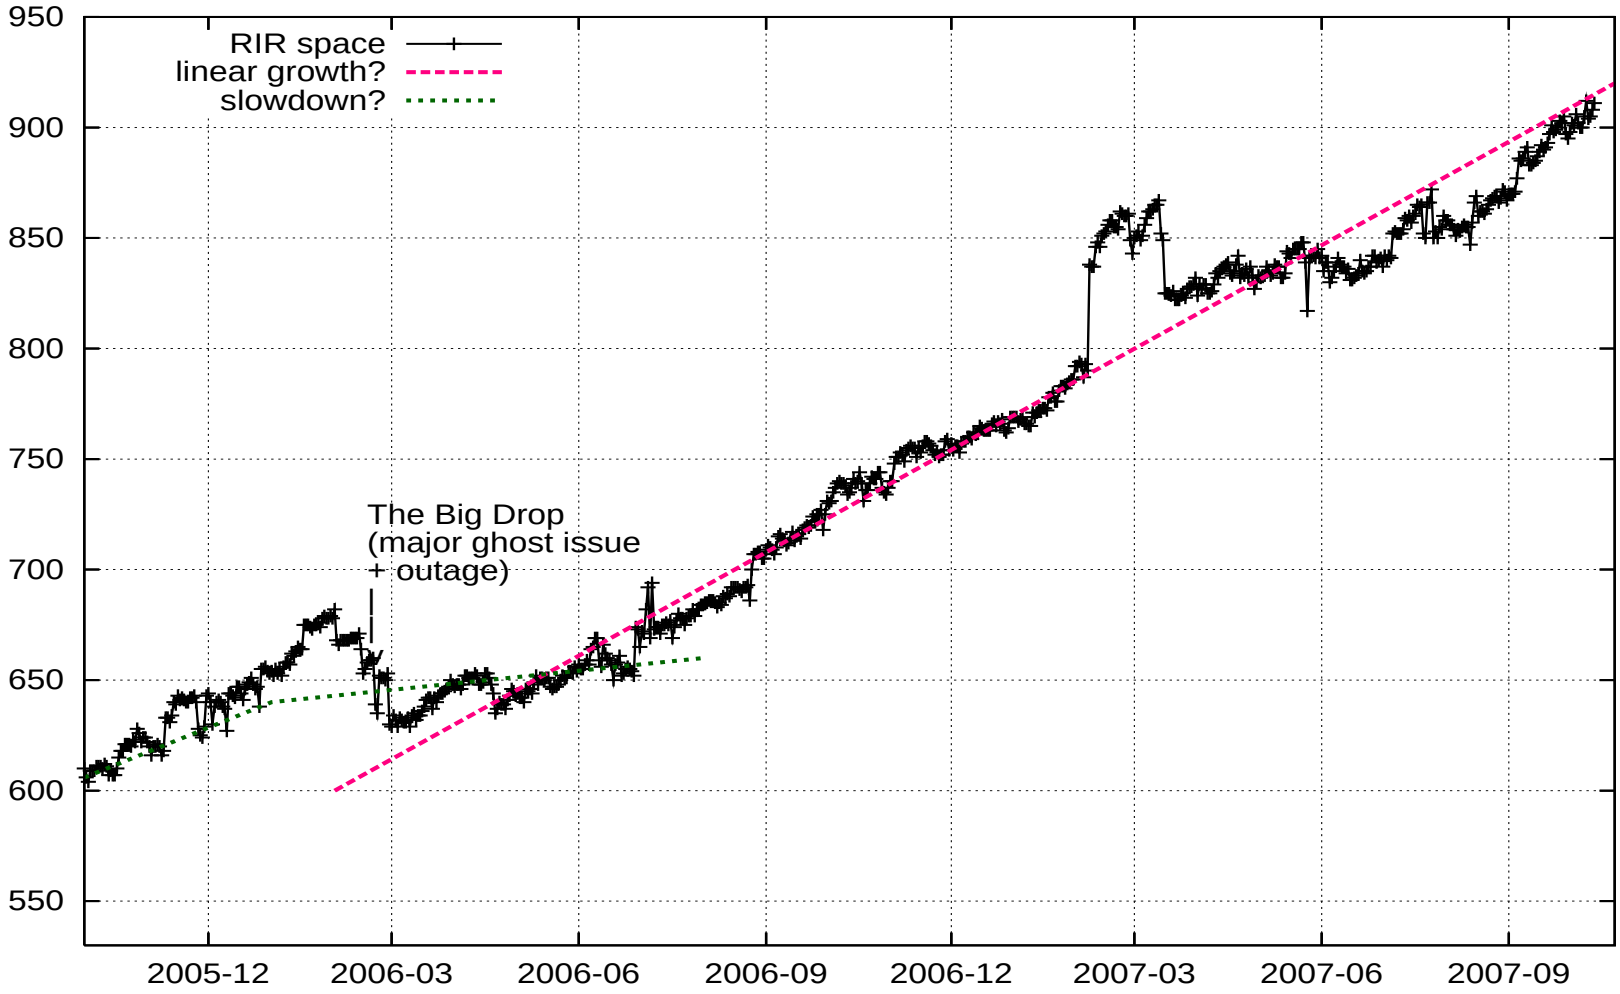

Graphics: prefixes by RIR region

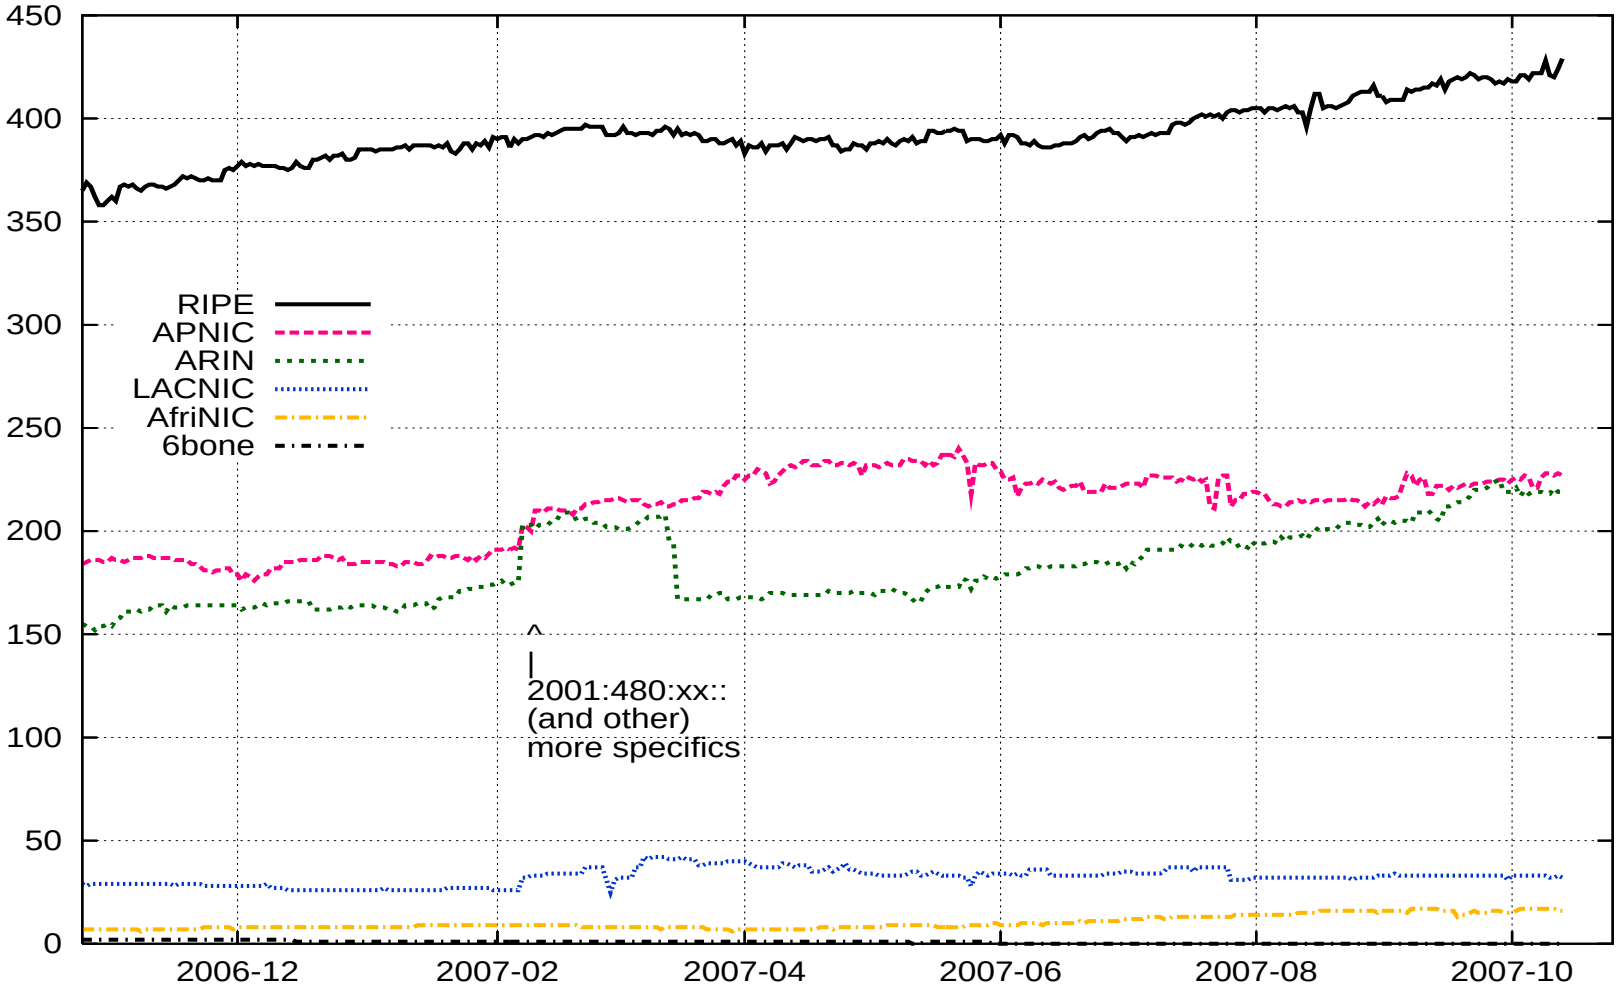

Graphics: prefixes by country (RIPE)

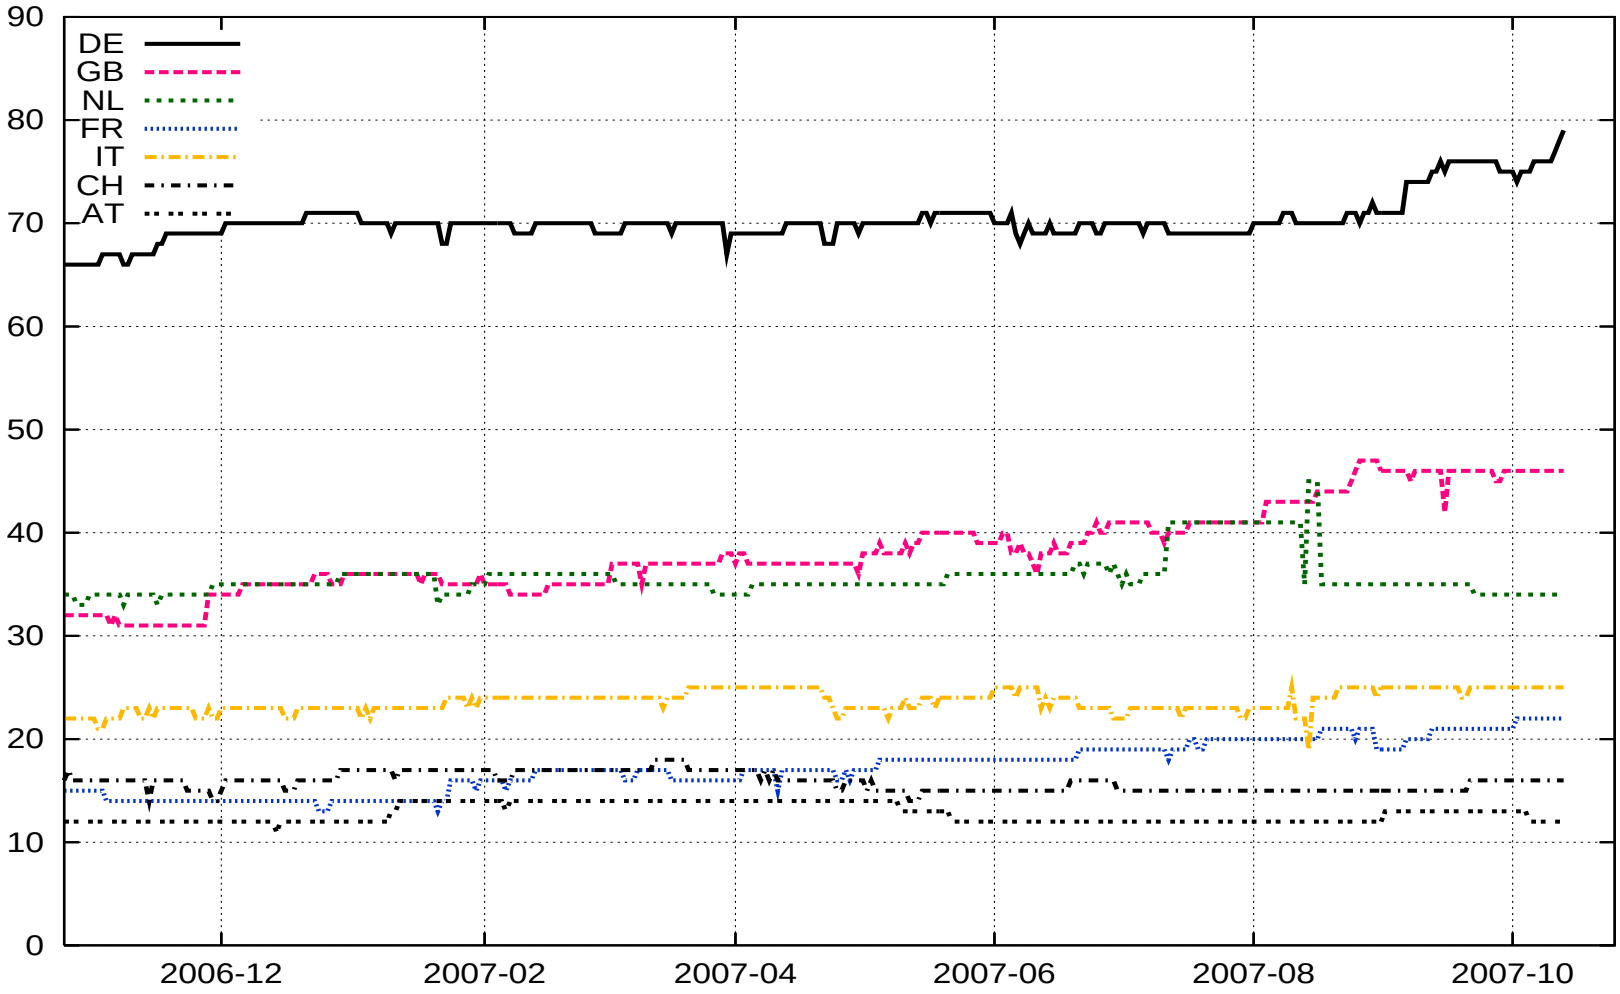

Graphics: prefixes by country (ARIN)

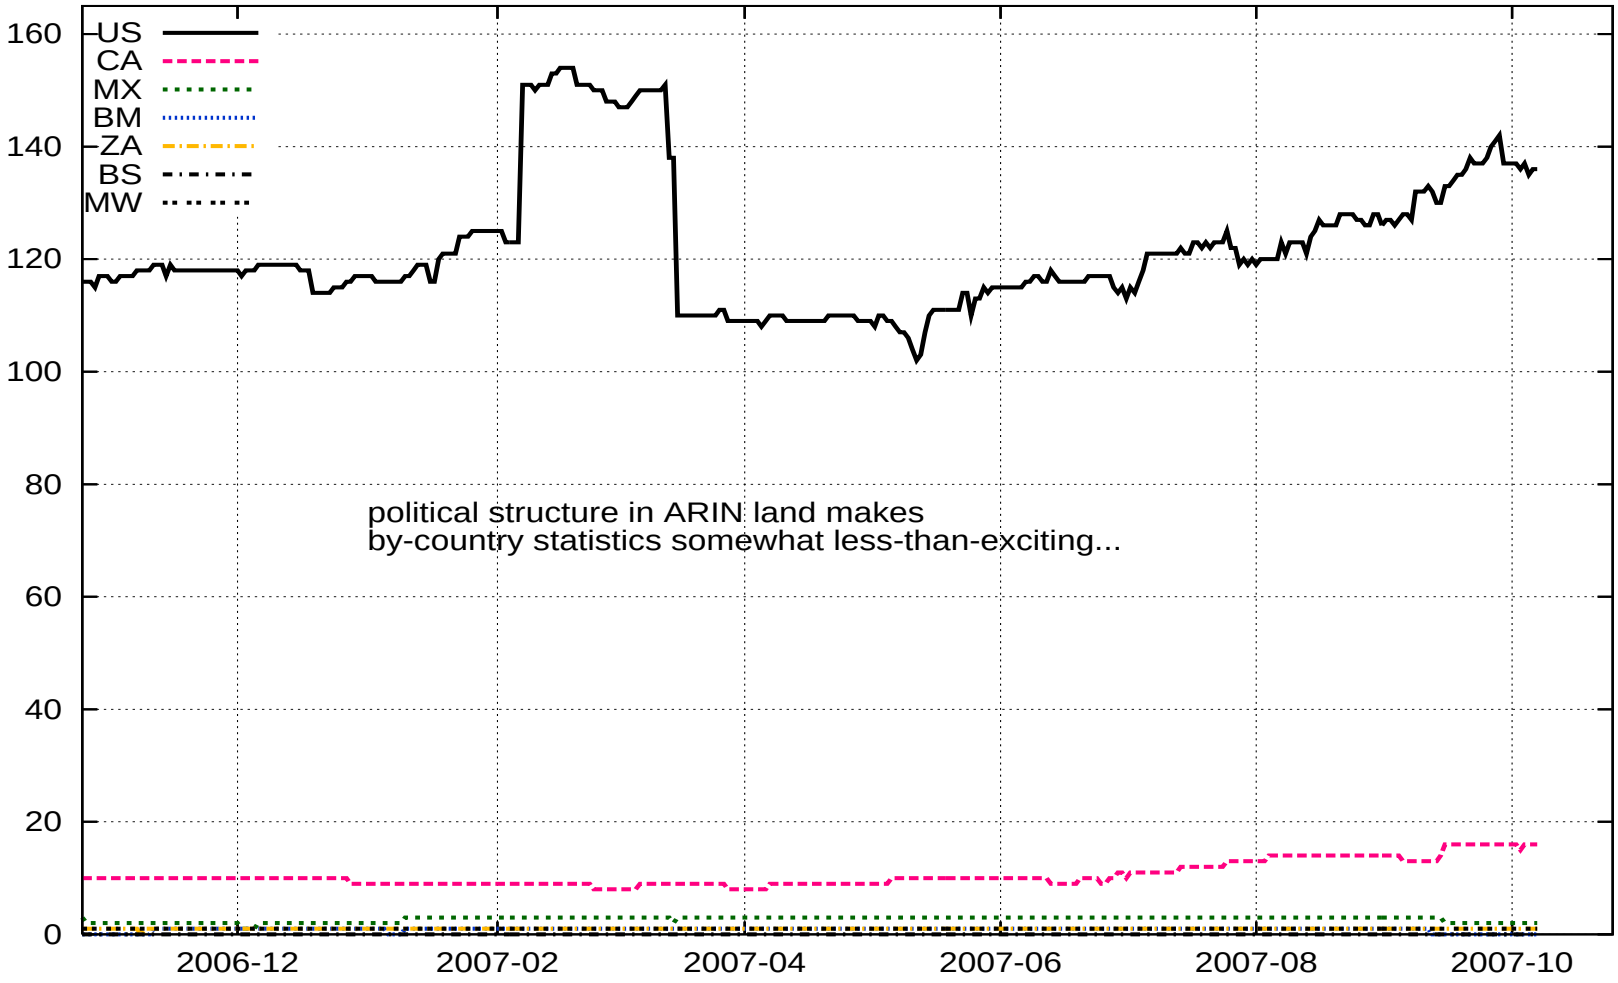

Numbers - AS numbers

- as of 2007-10-04: 798 unique AS numbers visible (2007-05: 720)
  - 552 origin-only ASes (no transit paths seen) (488)
  - 211 ASes originate & give transit (200)
  - 35 transit-only ASes (e.g. 57, 2153, 5549, 6667, . . .) (32)
- different number of prefixes announced
  - 0 ASes originate 6Bone (3ffe) prefixes (*hooray!*)
  - 673 ASes originate 1 RIR prefix (604)
  - 54 ASes originate 2 RIR prefixes (3 due to /32+/35)
  - 22 ASes originate 3 RIR prefixes
  - 14 ASes with “more than that”, maximum is 6 prefixes
- 3 ASes still announce their prefix as /32 and /35
- note: all paths observed from AS5539

ASes - why are people announcing 2+ prefixes

- /35 to /32 migration: 2 RIR prefixes, *temporary (?)*

```
2001:420::/35          109 i
2001:420::/32          109 i
```

- ISP/LIR address space plus IXP prefixes

```
2001:5000::/21          1273 i    (C&W LIR space)
2001:7F8:2B::/48        1273 i    (IXP: INXS HAM)
2001:7F8:2C::/48        1273 i    (IXP: INXS MUC)
```

- mergers and acquisitions, business units, customer pfxs, ...

```
2001:218::/32           2914 i      NTT JP
2001:418::/32           2914 i      NTT America
2001:49F0::/32          2914 i      FDCServers
2001:728::/32           2914 i      Verio Europe
2610:150::/32           2914 i      Sharktech Internet
2610:F8::/32            2914 i      Command Information Inc.
```

- networks with multiple sites and multiple PI prefixes

```
2001:502:100E::/48     2914 12008 i  UltraDNS
2001:502:2EDA::/48     2914 12008 i  UltraDNS
2001:502:4612::/48     2914 12008 i  UltraDNS
2001:502:AD09::/48     2914 12008 i  UltraDNS
2001:502:D399::/48     2914 12008 i  UltraDNS
2001:502:F3FF::/48     2914 12008 i  UltraDNS
```

ASes - 32 bit ASNs showing up

- sidetrack: some 32 bit AS numbers already active

```
   Network          Next Hop              Path
*> 2001:7fb:fd00::/48
                    ::ffff:194.97.146.46  5539 1273 12859 12654 3.7 i
*> 2001:df0:2::/48  ::ffff:194.97.146.46  5539 3257 2497 2.3 i
*> 2001:4810:2000::/35
                    ::ffff:194.97.146.46  5539 1273 29748 33437 6.3 i
Total number of prefixes 3
```

- this needs Quagga, OpenBGPd or IOS XR to see – others will see "2-byte tunnel AS" 23456:

```
*  2001:7FB:FD00::/48
                 2001:420:0:7FF0::1     109 30071 6939 12859 12654 23456 i
*                2001:470:1FFF:2::      6939 12859 12654 23456 i
* i              2001:7F8::CB9:0:1      3257 12859 12654 23456 i
*                ::FFFF:203.14.5.7      1221 30071 6939 12859 12654 23456 i
* i              2001:5001:100:16::1    1273 12859 12654 23456 i
```

Graphics - Prefixes

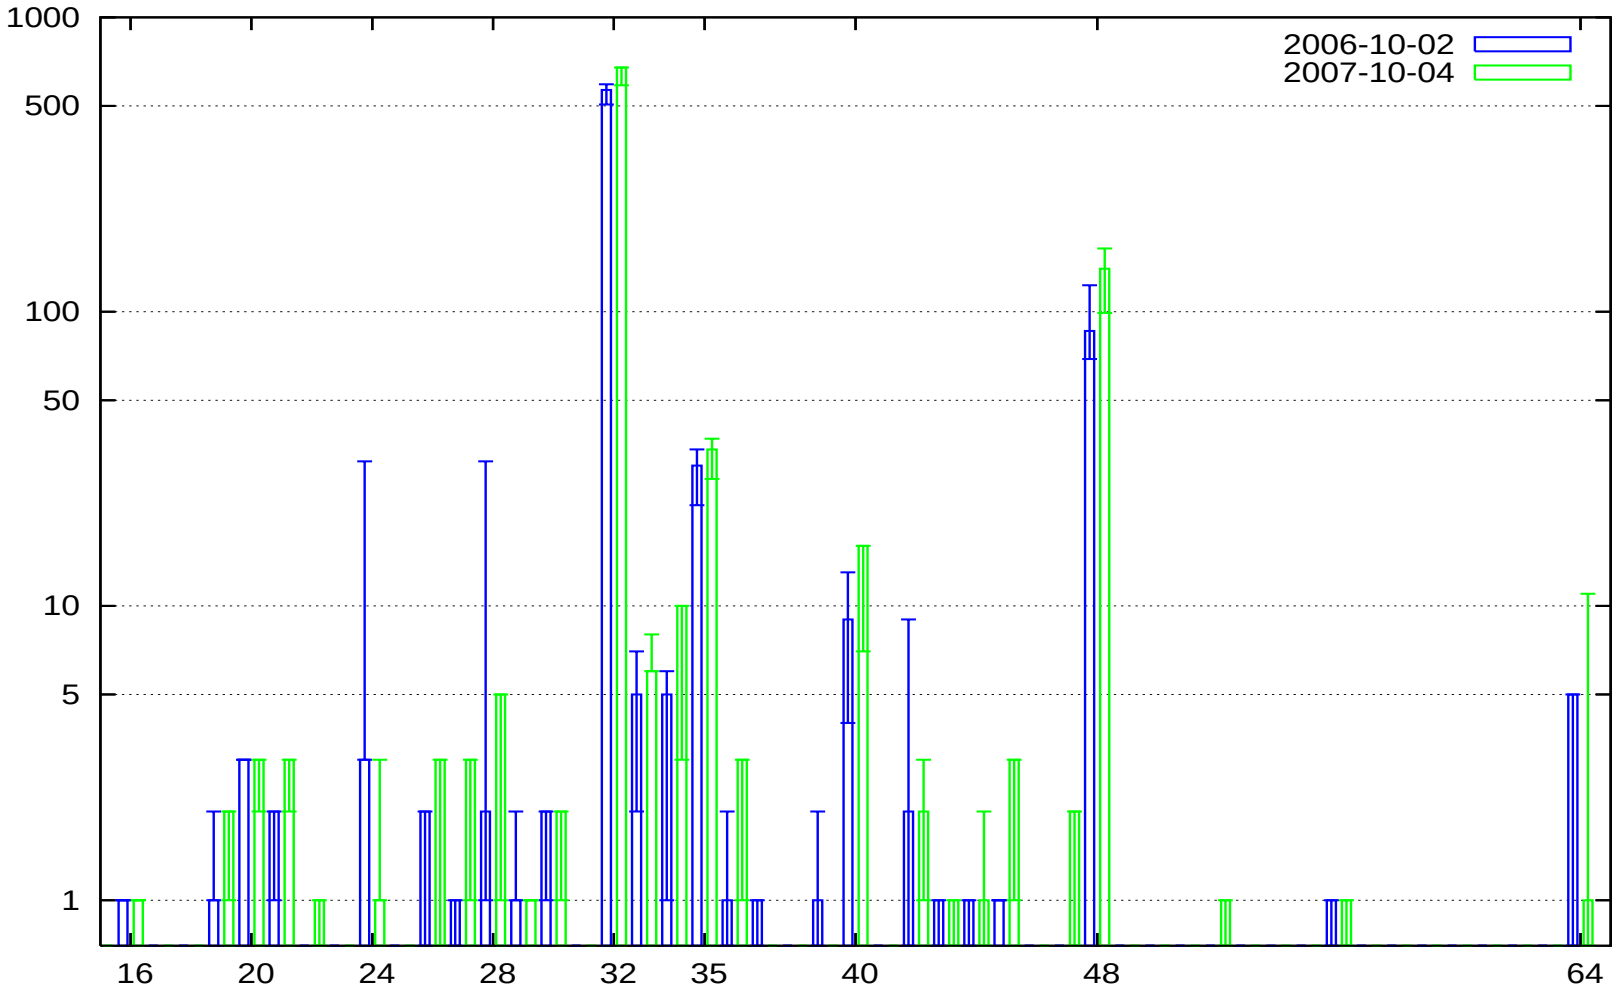

Numbers - Prefixes

As of 2007/10/04: 921 prefixes in total (2007-05-06: 838)

/n	global	RIPE	APNIC	ARIN	Lacn.	Afri.	oth
/16	1	0	0	0	0	0	1
/19	2	2	0	0	0	0	0
/20..23	7	3	4	0	0	0	0
/24..27	7	2	4	1	0	0	0
/28..31	8	1	4	0	3	0	0
/32	675	352	162	116	30	14	1
/33..34	16	5	4	7	0	0	0
/35	34	8	20	6	0	0	0
/36	3	2	0	1	0	0	0
/40	16	7	4	4	0	1	0
/42	2	2	0	0	0	0	0
/43	1	0	0	1	0	0	0
/44..47	6	2	0	4	0	0	0
/48	142	34	25	79	0	2	0
/49..63	2	0	1	1	0	0	0
/64..128	1	1	0	0	0	0	0

Numbers: RIRs, Allocations, ...

- On 2007-10-13, 1433 LIR blocks (2000::/4) allocated by RIRs:

RIR	alloc.	members	perc.	on 2007-05-06
ARIN	312	~ 2901	10.8%	247 (+26%)
APNIC	315	~ 2561	12.3%	287 (+10%)
RIPE	682	~ 5217	13.1%	617 (+11%)
LACNIC	86	~ 718	12.0%	74 (+16%)
AfriNIC	38	~ ??	??	28 (+36%)

- note: not counting /48 microallocs and /35⇒/32 extentions
- actual *percentage* with IPv6 similar among regions
- 689 (R54: 547) allocations visible in routing table (*only 48%!*)

Graphics: Allocated vs. Routed

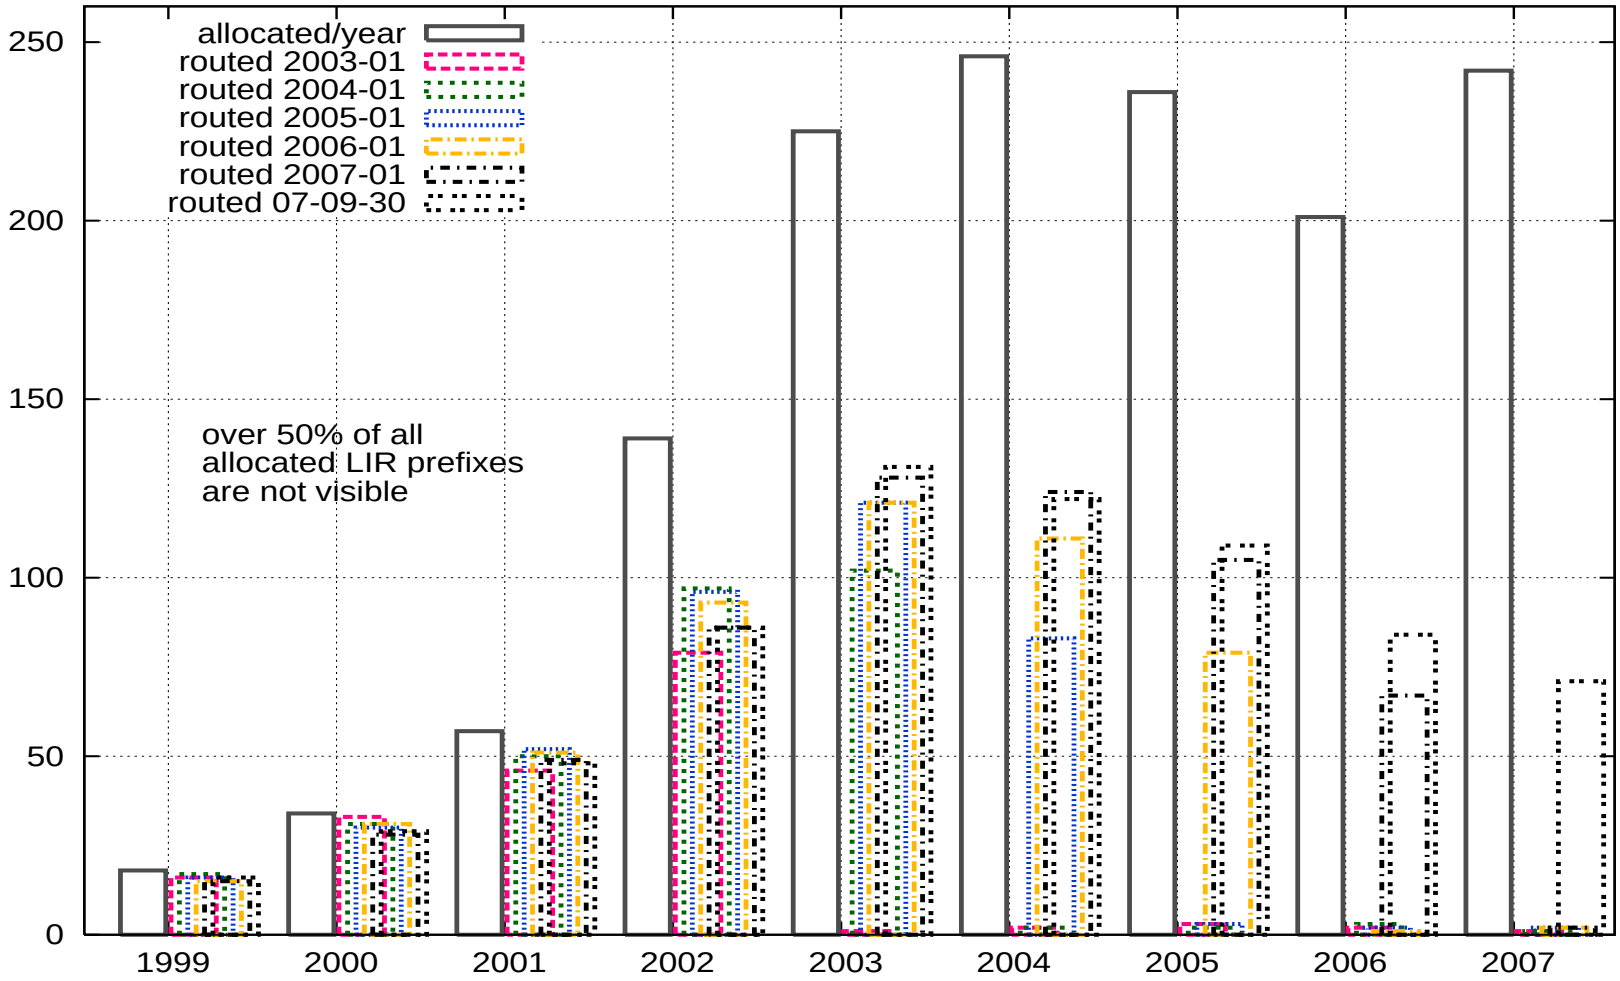

Allocated vs. Routed - by region & class

RIR	type	alloc.	visible	perc.	subnets
ARIN	LIR	292	117	40%	58
	IXP	20	0	0%	0
	Critical Inf.	42	20	48%	8
	Internal Inf.	2	0	0%	0
	PI	92	16	17%	1
APNIC	LIR	312	173	55%	52
	IXP	21	1	5%	0
	PI	7	2	29%	0
RIPE	LIR	674	358	53%	48
	IXP	64	12	19%	0
	Anycast DNS	5	3	60%	0
LACNIC	LIR	85	30	35%	3
	Critical Inf.	4	0	0 %	0
AfriNIC	LIR	33	14	42%	1
	PI	5	2	40%	0

Allocated vs. Routed - reasons?

- "early adopters" already losing interest in IPv6?
- "prepare for the future" allocations?
- "for internal use" allocations? (some, yes)
- distribution of non-announced prefixes does not show any specific characteristic, like "academia" vs. "commercial networks" etc.
- some delay between prefix allocation and announcement is to be expected (expect some more statistics in this space...)
  - but this cannot explain effects seen on 2003+2004 allocations – about 40% don't show up after over 3 years...

Numbers: notable allocations - PI news

- 6 IPv6 PI networks from APNIC (2 in BGP)
  - 2001:DE8::/48 Triple T Global Net, TH
  - 2001:DD8::/48 University of the South Pacific, Fiji
  - 2001:DF0::/47 University of Auckland, NZ
  - 2001:DF0:2::/48 Quadnet, NTT Laboratory Network, JP
  - 2001:DF0:3::/48 Crown Research Institute, NZ
  - 2001:DF0:4::/48 University of Waikato, NZ
- 5 IPv6 PI networks from AfriNIC (2 in BGP)
  - 2001:43F8::/48 Tanzania Internet Exchange, TZ
  - 2001:43F8:10::/48 KENIC
  - 2001:43F8:20::/48 Ubuntunet Alliance, ZA
  - 2001:43F8:30::/48 UniForum, ZA
  - 2001:43F8:40::/48 descr: Testing Reachability for PI /48s
- 90 “direct” assignments (PI) from ARIN so far, 15 in BGP
- $\Rightarrow$ **check your BGP filters!!**

Numbers: notable allocations (2)

- (a few) more "very large" allocations seen:
  - 2a00:2000::/22 British Telecom, UK (2007-08-29)
  - 2401:6000::/20 Australian Govt Dpt. of Defense (2007-08-10)
- Allocations ICANN ⇒ RIRs since RIPE 52

Prefix	RIR	Date	Comment
2400:0000::/12	APNIC	2006-10-03	
2600:0000::/12	ARIN	2006-10-03	
2800:0000::/12	LACNIC	2006-10-03	
2A00:0000::/12	RIPE NCC	2006-10-03	
2C00:0000::/12	AfriNIC	2006-10-03	

- http://www.iana.org/assignments/ipv6-unicast-address-assignments

Interesting Observations (1) - DTAG hijack?

```
   Network            Next Hop                  Path
*  2003::/19         2001:420:0:7FF0::1         109 5511 3320 i
*                    2001:470:1FFF:2::          6939 2497 3257 3320 i
*                    2001:7F8:2:8001::2         1752 3320 i
*>i                  2001:7F8:2C:1000:0:A500:3320:1
                                                3320 i
*                    2001:608:0:3::9            3320 i

*> 2003:8FE:0:A012::/64
                     2001:470:1FFF:2::          6939 5623 7018 2386 ?
```

- This really doesn't look "right" - no inet6num, no route6 object, AS path through completely unrelated ASes.
- 5623, 7018, 2386 = AT&T, 6939 = Hurricane Electric
- only paths via 6939 are visible at GRH looking glass
- $\Rightarrow$ please *check* what you give transit for!

Interesting Observations (2) - Ghost Busting

```
     Network          Path
2007-05-09:
  * i2001:18B0::/32  3257 6939 1280 3557 3741 i
  *                  1752 2914 3557 3557 3557 3741 i
  *                  1221 30071 3557 3741 i
  *                  6939 1280 3557 3741 i
  *                  109 30071 3557 3741 i

2007-05-10 .. 2007-05-22:
  *  2001:18B0::/32  1752 3356 3257 6939 2516 7660 2500 2497 1273
                                               3303 2914 3557 3557 3741 i
  *>i                3303 2914 3557 3557 3741 i
  *                  6939                2516 7660 2500 2497 1273
                                               3303 2914 3557 3557 3741 i
  *                  109 6453 8002       2516 7660 2500 2497 1273
                                               3303 2914 3557 3557 3741 i
```

- Ghosts = BGP withdrawal bug, caused by *buggy software*.
  Long paths can stay *mostly unchanged* in the table for weeks.
- don't confuse with BGP count-to-infinity (= paths *change*).
- Cisco has been able to reproduce & fix bug: CSCsc59089

Interesting Observations (3) - Accidential "Hijack"

```
   Network             Path
*>i2001:200::/32       2914 2500 i                          <<< normal path
*                      1221 30071 3557 2500 i

*  2001:6E0::/32       1221 30071 6175 2497 2500 i
*>i                    3257 8954 i

*  2001:740::/32       1221 30071 6175 17715 6435 278 6939 2516 7660 2500 i
*                      6939 2516 7660 2500 2914 8472 i
*>i                    1273 8472 i

*  2001:808:E000::/35 1221 30071 6175 17715 6435 278 6939 2516 7660 2500 i
*  2001:AA8::/32       1221 30071 6175 17715 6435 278 6939 2516 7660 2500 i
*  2001:1450::/32      1221 30071 6175 2497 2500 i
*  2001:1498::/32      1221 30071 6175 2497 2500 i
...
*  2001:1820::/32      1221 30071 6175 2497 2500 i
*  2001:1B70::/32      1221 30071 6175 17715 6435 278 6939 2516 7660 2500 i
*  2001:4130::/32      1221 30071 6175 4555 6939 2516 7660 2500 i
*  2001:4B20::/32      1221 30071 6175 17715 6435 278 6939 2516 7660 2500 i
```

- this has happened before, but the cause is unknown
- theory: combination of ghosting and AS path truncation

Something for the Routing Police...

```
HOST: svr02.teleport-iabg.de      Loss%   Snt   Last   Avg  Best  Wrst StdDev
  1. backbone2-gige-0-3-15.telepo  0.0%    20    0.7   4.6   0.5  47.3  11.8
  2. mchn-s1-rou-1030.DE.euroring  0.0%    20    2.5   2.6   2.2   2.8   0.2
  3. hmb-s2-rou-1030.DE.eurorings  0.0%    20   17.5  17.4  17.0  17.8   0.2
  4. sl-bb1v6-nyc-t-28.sprintv6.n  0.0%    20  108.4 108.4 107.9 108.9   0.2
  5. sl-bb1v6-rly-t-1003.sprintv6  0.0%    20  184.3 184.5 184.1 185.8   0.4
  6. 2001:ca0:1::1:1               5.0%    20  495.6 495.6 488.4 499.5   2.9
  7. tunnel-chttl-lavanoc.lava.ne  5.0%    20  500.1 496.2 489.8 500.1   2.6
  8. 3ffe:8070:1:13::1            15.0%    20  608.5 570.0 558.8 608.5  13.0
  9. 2001:1228:11b:90a::1         10.0%    20  811.2 819.2 790.9 974.5  39.6
...
 17. 2001:1338:ffff::2            21.1%    19  656.5 700.2 650.3 1032. 108.4
 18. ns1.nic.ve                   15.8%    19  666.0 705.0 651.1 976.8 102.8

   Network           Path
2001:1338::/32  8767 3549 6175 17715 6435 278 18592 27750 27807 20312 i
2001:410::/32   8767 3549 6175 17715 6435 278 18592 27750 6509 i
                286 1273 6830 6830 6830 6830 6939 6939 6939 6939 2516
                 7660 22388 11537 6509 i
```

- some NRENs “do not need” reasonable-quality IPv6 upstream
- global connectivity happens via “tunnel full-table leaks”
- net result is extremely poor user experience

Graphics: route6 objects vs. routes seen

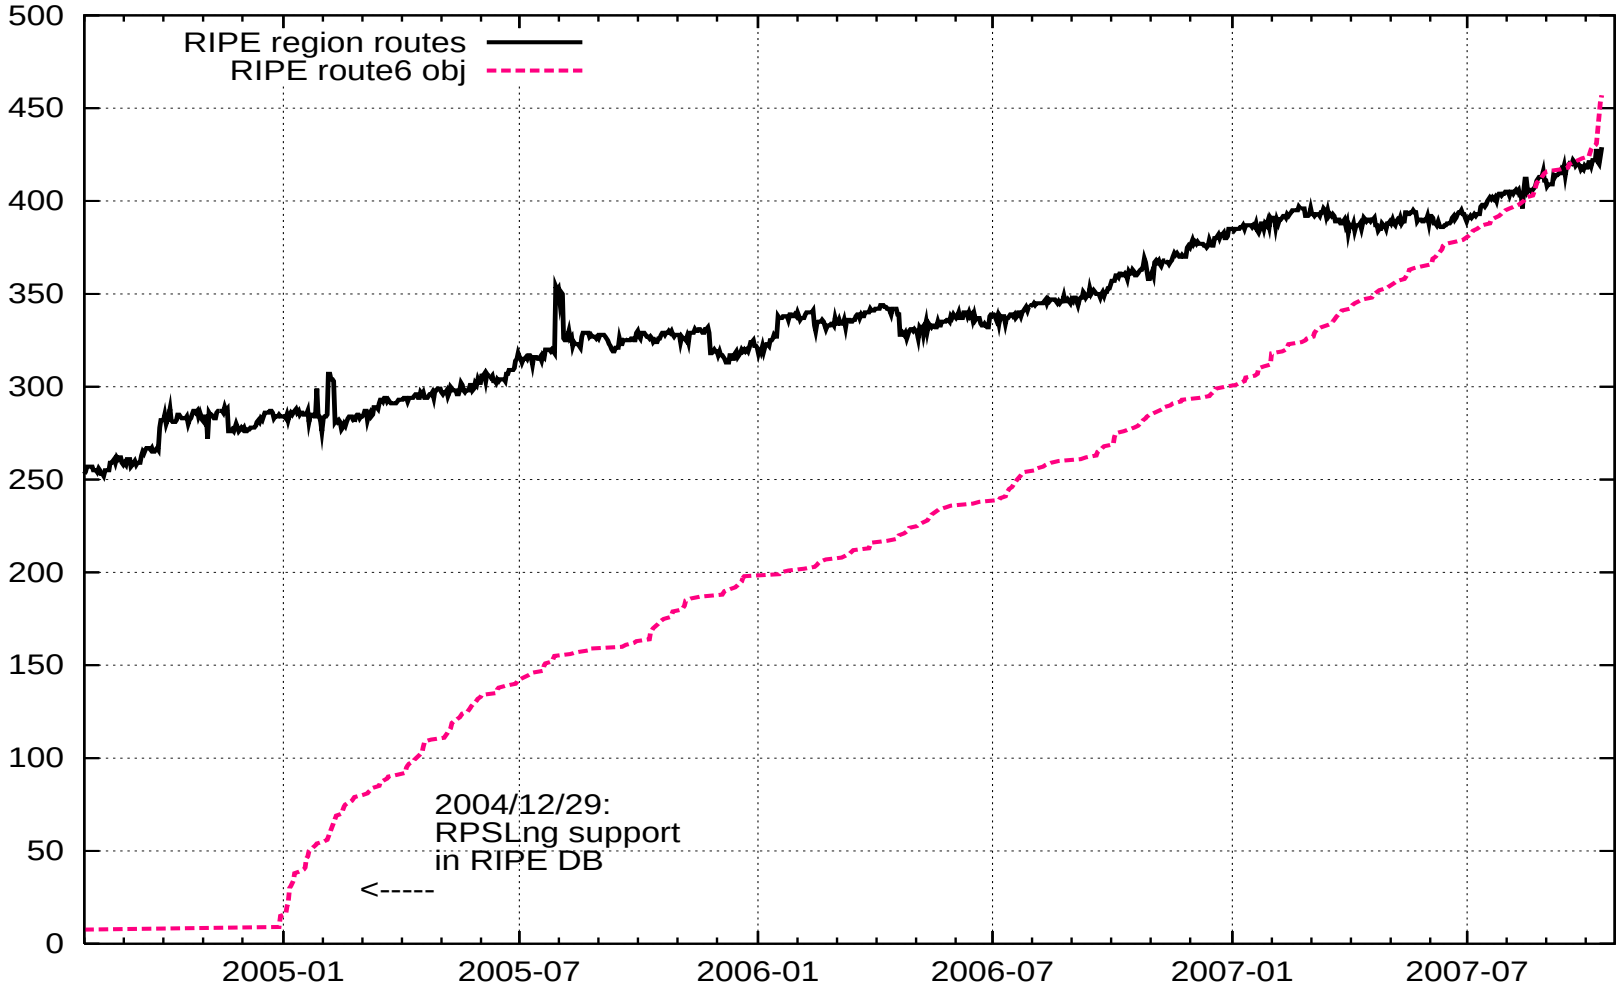

route6 correlation (RIPE region)

- on 2007-10-04:
  - 421 BGP routes from RIPE region
  - 425 route6: objects in RIPE DB
- correlation?
  - multiple origin route6's (9x 2002::/16, 5x 2001::/32, ...)
  - ⇒ 409 route6 objects for *uniqe* prefixes
  - 274 route6: objects exactly matching BGP prefixes
    - (origin AS not checked yet, only prefix match)
  - 135 route6: objects without BGP prefix (?!)
  - 147 BGP prefixes without route6: object :-(
- ⇒ close-up view shows "more work needed"
- in other RIR regions, situation is worse (no IRR DBs yet, etc.)

route6 object example

- it's as easy as this...

```
route6:         2001:608::/32
descr:          DE-SPACE-2001-0608
descr:          SpaceNET AG, Munich
origin:         AS5539
notify:         noc@space.net
mnt-by:         SPACENET-N
changed:        gert@space.net 20041230
source:         RIPE
```

- strongly recommended, helps upstream/peer ASes build decent BGP filters, based on IRR data

new tool: GRH Longest Distance Routing

- http://www.sixxs.net/archive/sixxs/2007-04-01-GRH-LongestDistanceRouting.html
- approximate the *geographical* AS path length for a given prefix
- originally meant as a April Fool's joke, but actually it's quite useful to quickly find *really* bad paths
- and the winner is...

  `2001:200:a000::/35 25441 3257 3549 6939 2516 7660 22388 11537 2500`

  at 40760 km (Ireland, Germany, NL, US, JP, US and Japan), and

  `2001:200:a000::/35 1836 3549 6939 2516 7660 22388 11537 2500`

  at 39500 km (Switzerland, NL, US, JP, US, and Japan)
- kudos goes to Jeroen Massar

References

- Ghost Route Hunter: http://www.sixxs.net/tools/grh/
- List of IPv6 blocks allocated by the RIRs:
  http://www.ripe.net/rs/ipv6/stats/index.html
- MIPP (minimum peering policy) project:
  http://ip6.de.easynet.net/ipv6-minimum-peering.txt
- IPv6 sample prefix filter page
  http://www.space.net/~gert/RIPE/ipv6-filters.html
- Slides are available at:
  http://www.space.net/~gert/RIPE/NANOG41-v6-table/

Questions?

gert@space.net