Impressions

An overview of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich

January 27th, 2004

RIPE 47, Amsterdam

Overview

- numbers
- pictures & trends
- things that should not be there. . .
- conclusions & recommendations
- references

Slides online at: http://www.space.net/~gert/RIPE/R47-v6-table/

Numbers - AS numbers

- as of 2004/01/24: 409 unique AS numbers visible
  - 232 origin-only ASes (no transit paths seen)
  - 156 ASes originate & give transit
  - 21 transit-only ASes (e.g. 1237, 1717, 2603, 3301, . . .)
- mixture of RIR (2001::) and 6Bone (3FFE::) space announced
  - 239 ASes originate 1 RIR prefix
  - 54 ASes originate 1 6Bone prefix
  - 42 ASes originate 1 6bone + 1 RIR prefix
  - 22 ASes originate 2 RIR prefixes (7 due to /32+/35)
  - 31 ASes with “more than that”, maximum is 4 prefixes
- 13 ASes still announce their prefix as /32 and /35
- note: all paths observed from AS5539

ASes - why are people announcing 2 prefixes?

- 6bone to RIR migration: 1 6bone, 1 RIR prefix, *temporary*

```
2001:420::/35          109 i
3FFE:C00::/24          109 i
```

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

```
2001:760::/32          3549 20965 137 i
2001:760::/35          680 6880 137 i
```

- experiments and/or leaks?

```
2001:700::/32          2603 224 i
2001:700:E000::/35     680 6680 2603 224 i
2001:700:FFFF::/48     1853 1853 1853 1853 1853 6680 2603 224 i
```

- multi-uplink-/multi-homing-experiments?

```
2001:610:140::/48      3549 1200 i
2001:7B8:200::/48      3549 1200 i
2001:7F8:1::/48        3549 1200 i
```

- mergers and acquisitions, different business units, ...

```
2001:1428::/32         3303 i
2001:918::/32          3303 i
```

Numbers - Prefixes

As of 2004/01/24: 524 prefixes in total (2003/09/03: 485)

/n	global	RIR space	6bone	6to4	(2003/09/13)
/16	1	0	0	1	(1 0 0 1)
/24	42	0	**42**	0	(45 0 45 0)
/27	1	1	0	0	(0 0 0 0)
/28	41	0	**41**	0	(42 0 42 0)
/32	324	**293**	**31**	0	(266 238 28 0)
/33	3	3	0	0	(4 3 0 1)
/34	0	0	0	0	(1 0 0 1)
/35	45	**43**	2	0	(45 45 0 0)
/36-/39	0	0	0	0	(4 3 0 1)
/40	3	3	0	0	(4 4 0 0)
/42-/45	3	3	0	0	(3 3 0 0)
/48	57	35	22	0	(66 36 26 4)
/52-/60	0	0	0	0	(0 0 0 0)
/64	4	1	3	0	(4 2 2 0)
/65-/128	0	0	0	0	(0 0 0 0)

6to4 - 2002::/16

- 6to4 prefix 2002::/16 anycast prefix - *multiple* origin ASes

```
   Network    Next Hop                Path
*  2002::/16  3FFE:8150:0:1::17       9044 559 i
*             2001:608:0:3::7         1930 i
*             2001:948:0:F00F::1      2603 1741 i
*             3FFE:C00:8023:19::1     109 i
*             2001:608:0:3::3         1853 1853 1853 1853 1853 6680 786 ?
*             2001:608:0:3::9         3320 1275 8379
*             3FFE:1108:40A:FFFF::1:2 3274 1741 i
*             2001:7F8::3349:0:1      13129 9033 i
*             2001:7F8::CB9:0:1       3257 786 ?
*             2001:7F8::1A0B:0:1      6667 790 3327 i
```

- this is **fine**, anycast relay approach, see RFC3068
- relay in AS 1752 has disappeared. 768, 9033 and 3327 are new
- some research on non-publically visible 6to4 relays by David Malone (dwmalone@maths.tcd.ie): approximately 42 relays found. Good start, but more relays would be useful.
- **all** more-specific pfxs have been removed - thanks, Jeroen!

Graphics: Total Prefixes - 28 months

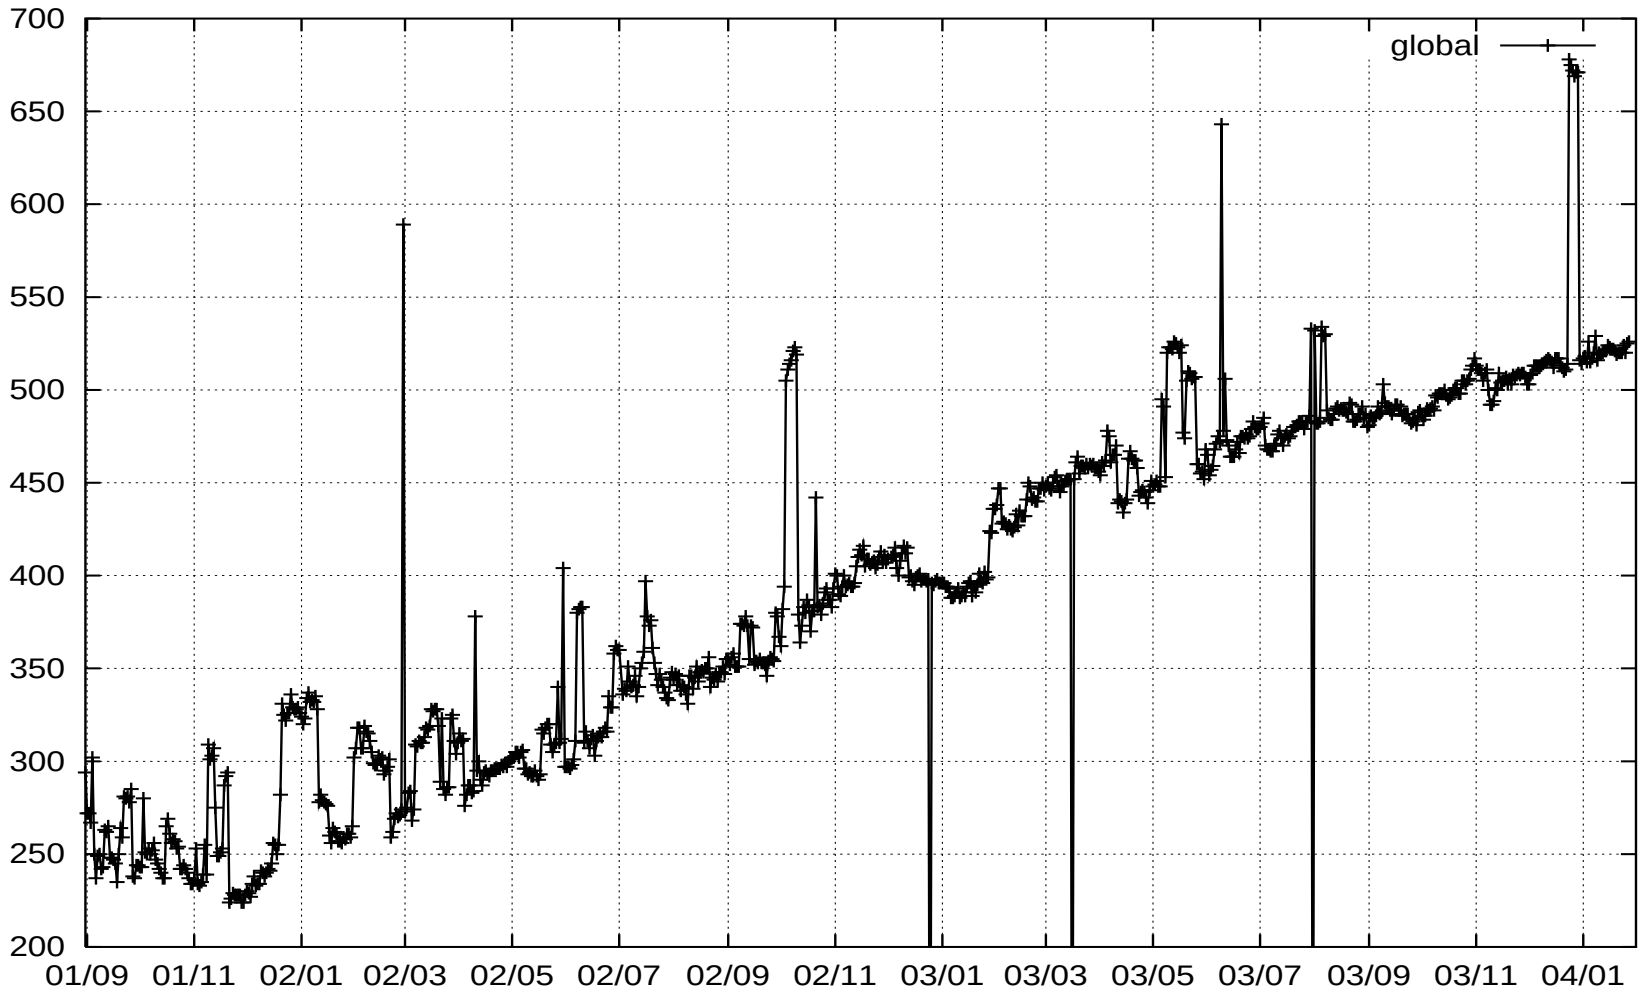

Graphics: RIR vs. 6Bone Prefixes - 28 months

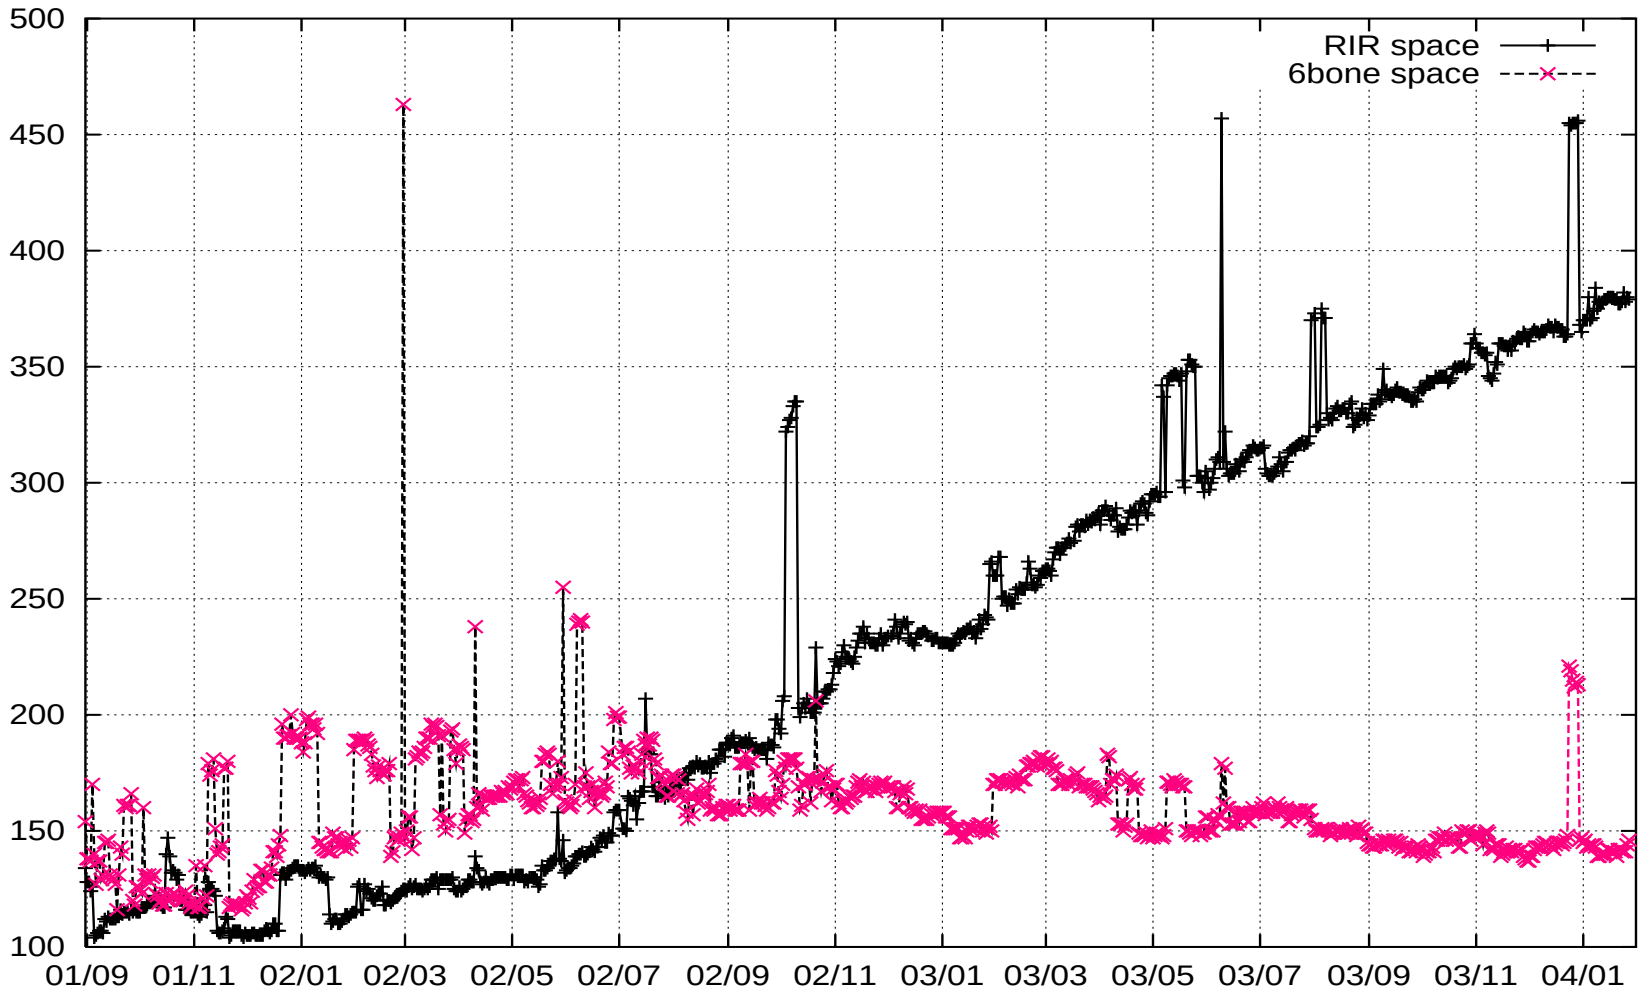

Graphics: RIR vs. 6Bone Prefixes - 4 months

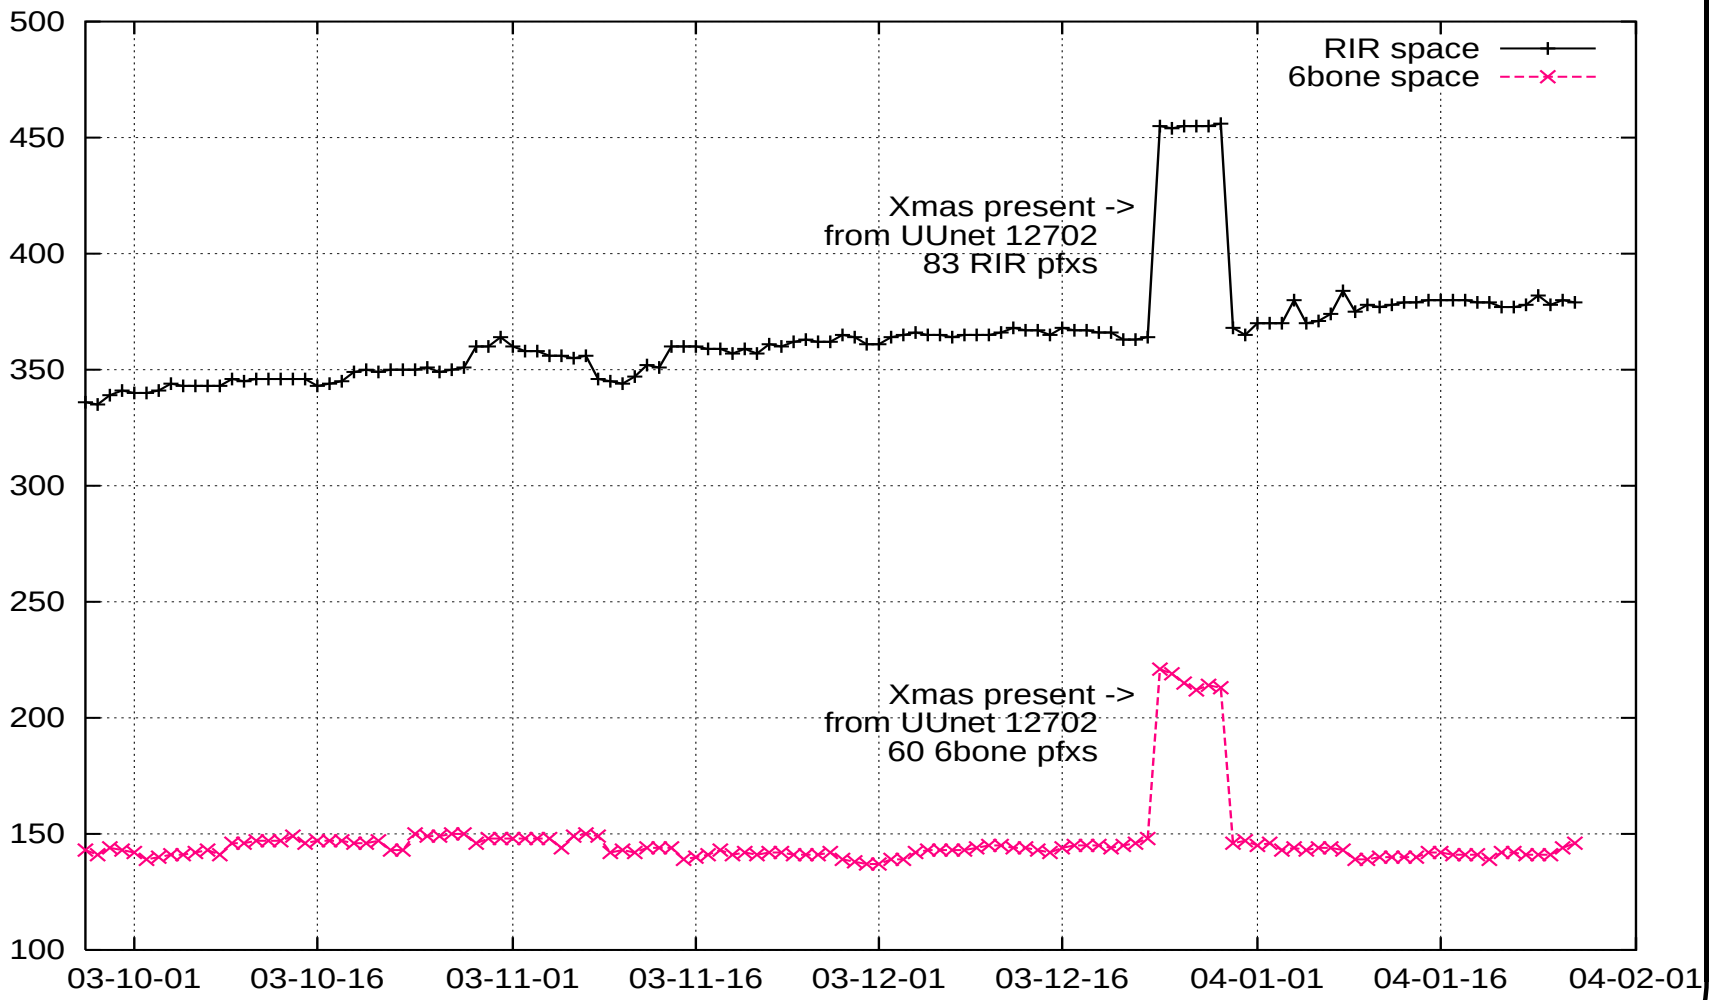

Graphics: RIR /35s vs. /32s

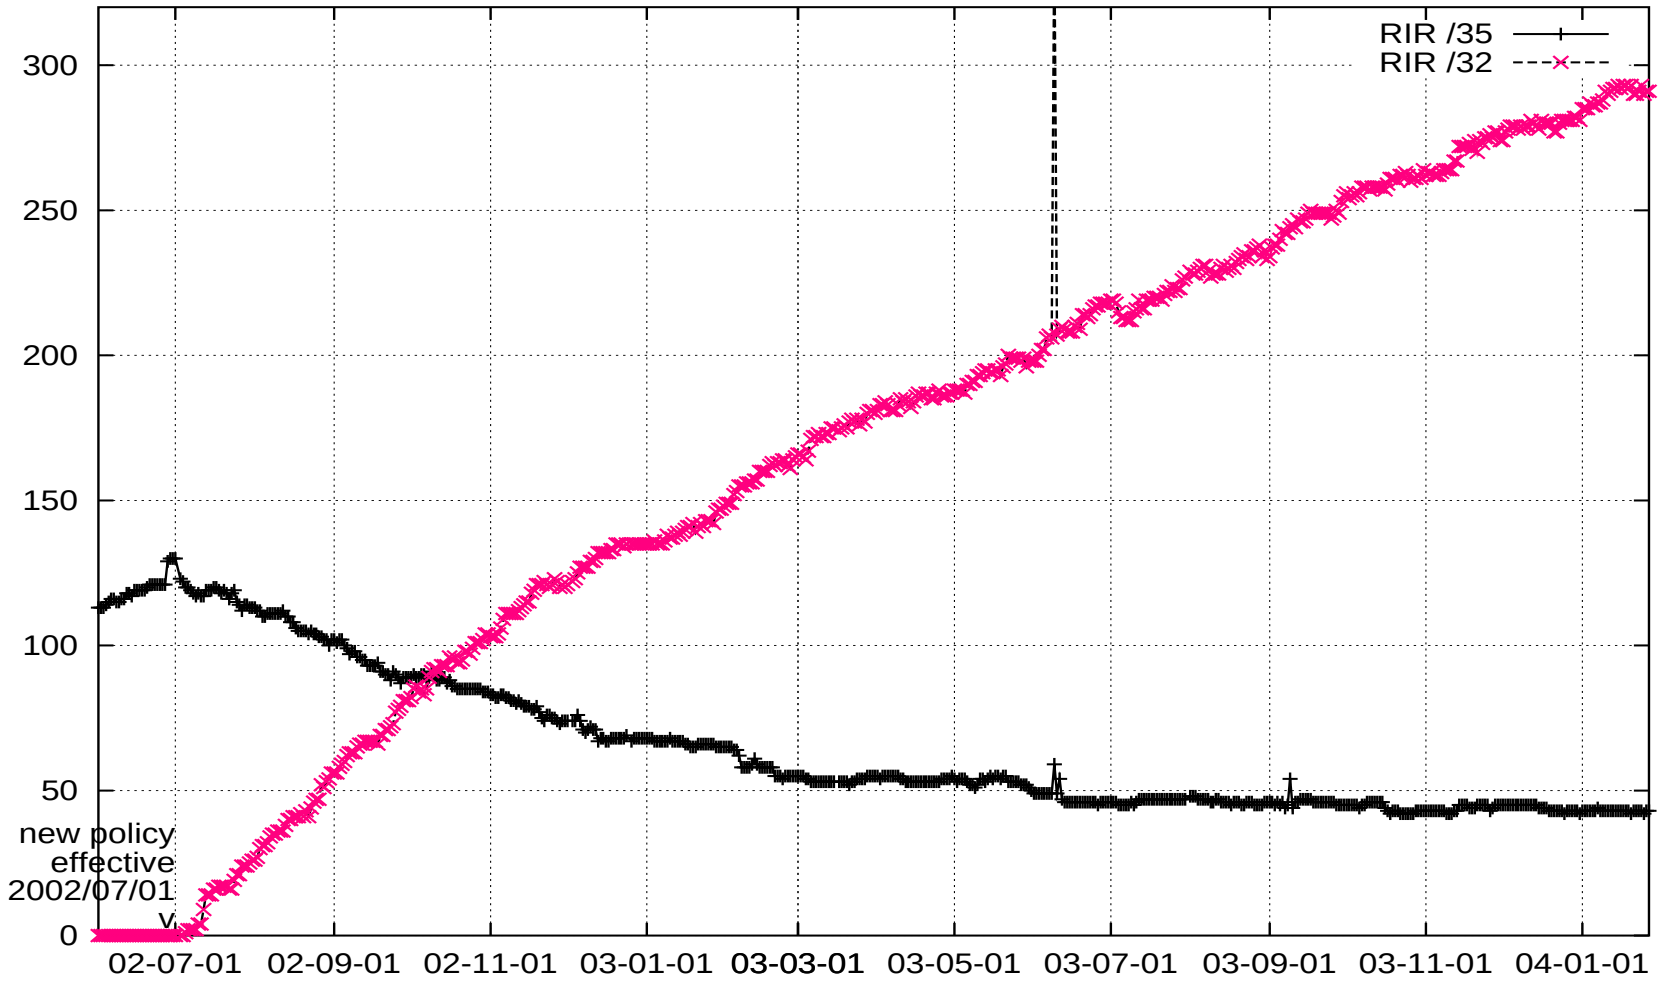

Numbers: RIRs, Allocations, ...

- 519 LIR blocks out of 2001::/16 allocated by RIRs:
  ARIN 95(76), APNIC 132(112), RIPE 286(238), LACNIC 6(3)
  as of 2004/01/27 (2003/09/03: 429)
- plus some IXP and other microallocations
- new IPv6 assignment to I.root (2001:07FE::/32), not visible yet
- B, F, H, I, K, M Root-Servers have IPv6 addresses (some are visible on www.root-servers.org, some are already in BGP)
- 319(271) allocations visible
- allocations take up 334(283) routes: 43(45) /35s, 291(238) /32s
  15(16) allocations visible as /32 *and* /35
- new IPv6 block for RIPE since 2004/04/12: **2001:1A00::/23**
- first RIR /27 allocation ever on 2003/11/24 (CH-SUNRISE)

Graphics: prefixes by RIR region

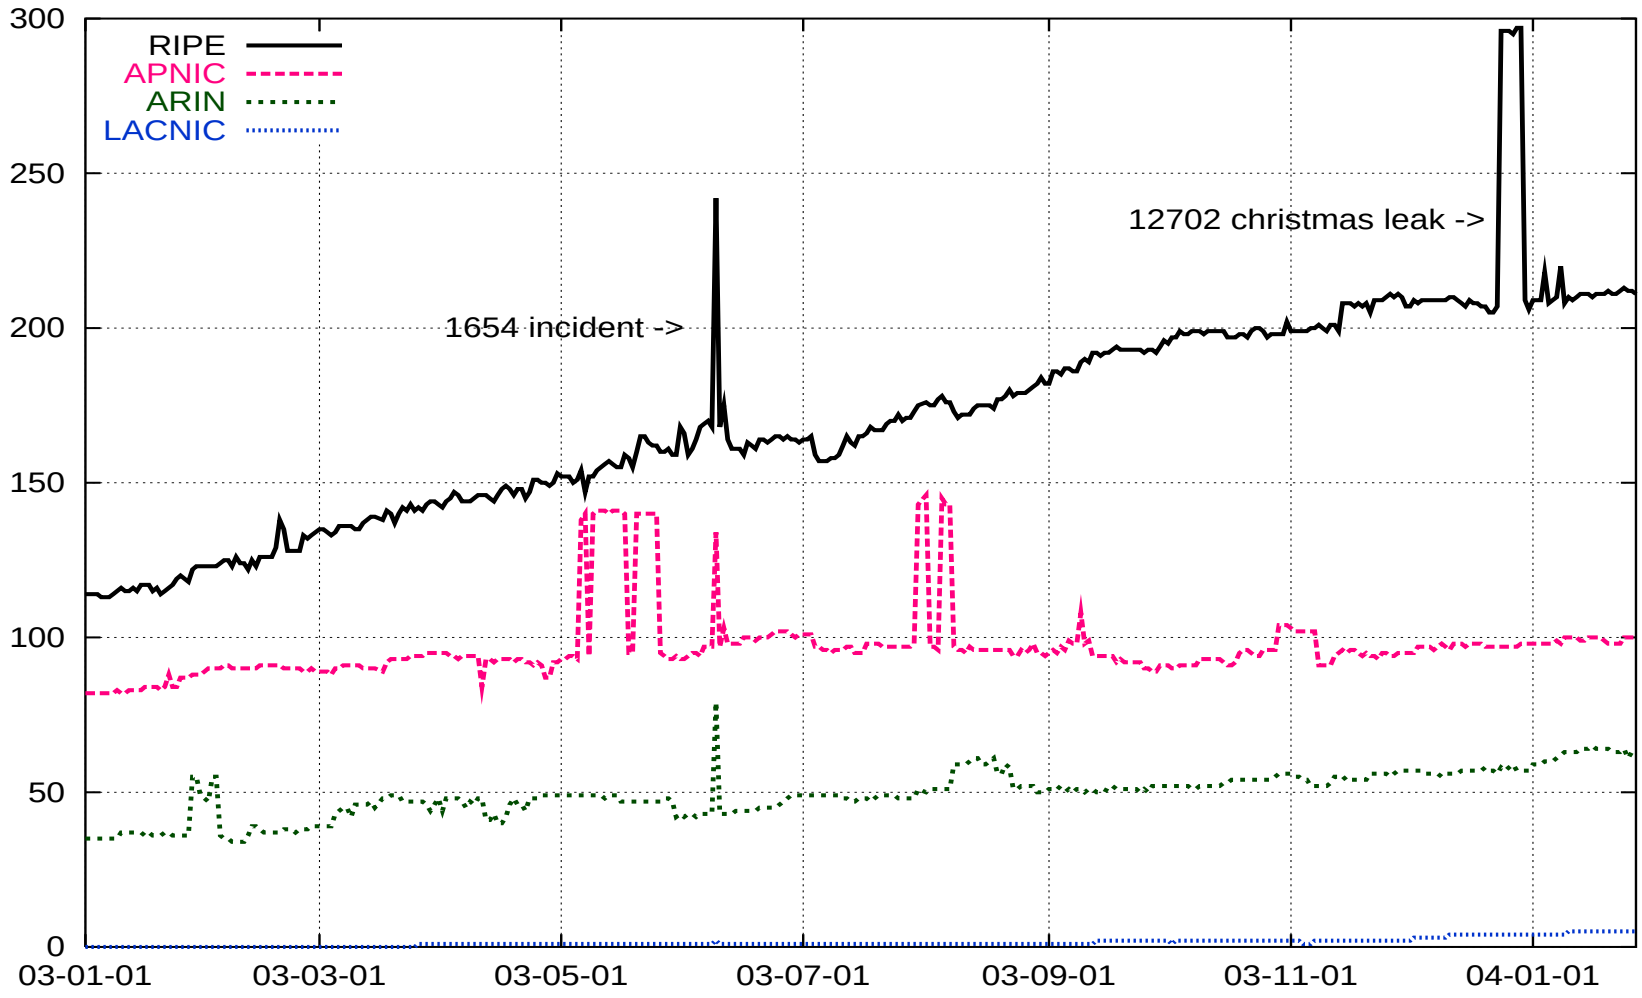

Graphics: prefixes by country (RIPE)

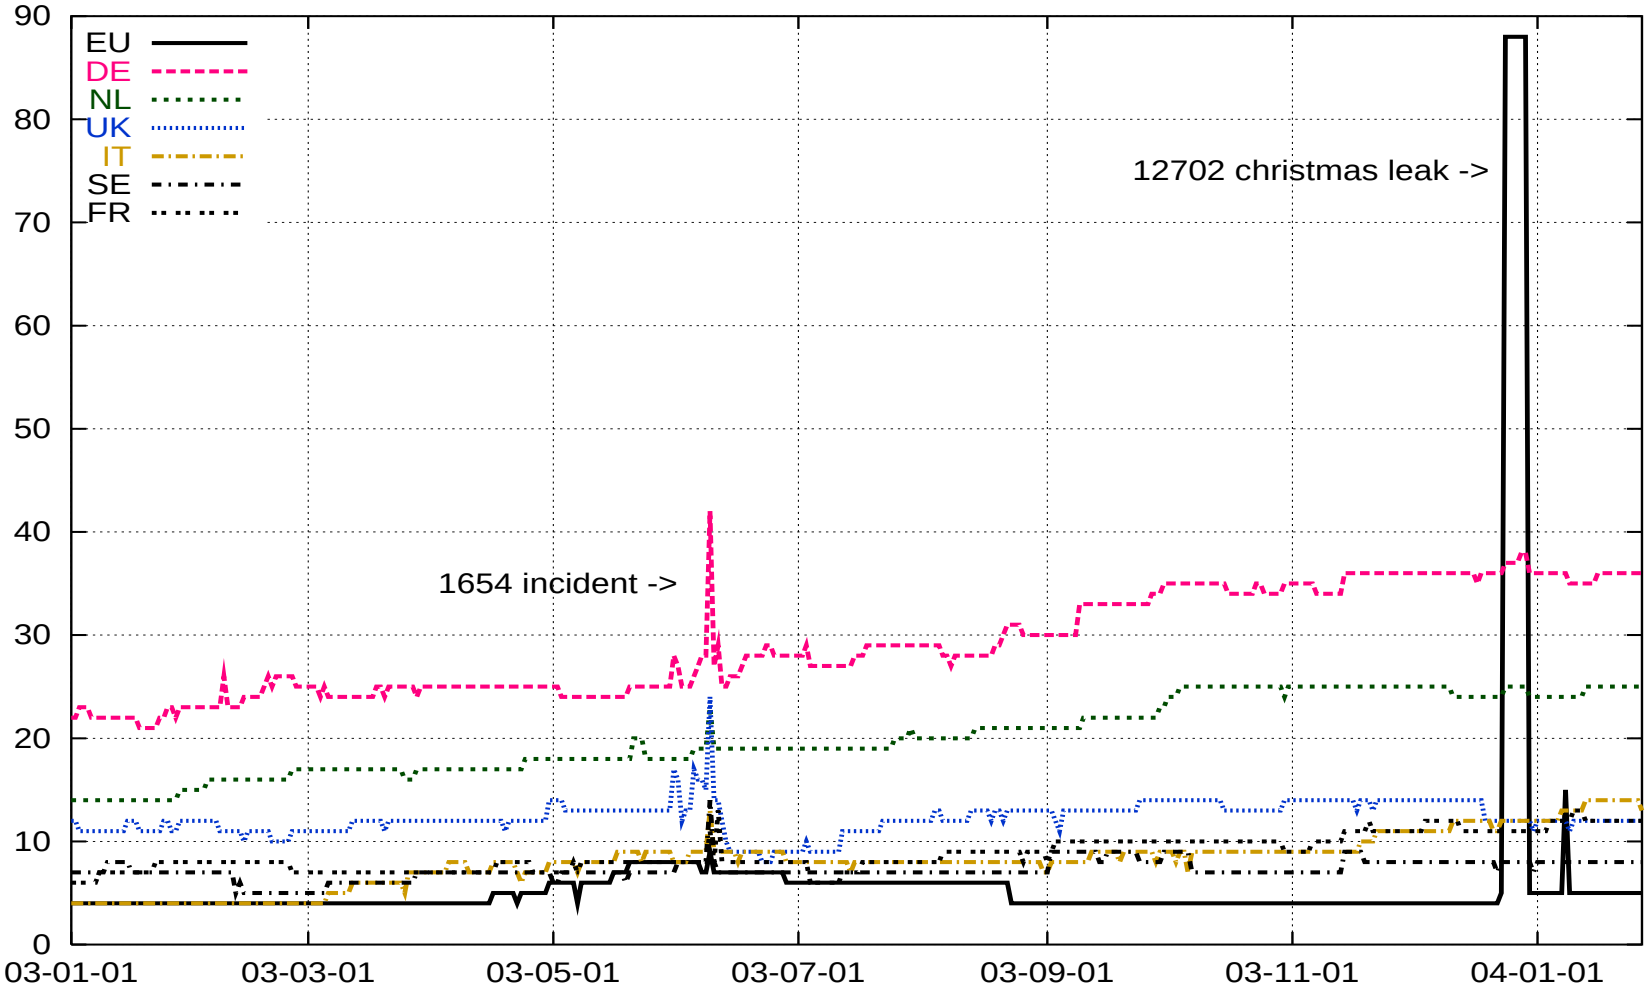

Graphics: Allocated vs. Routed

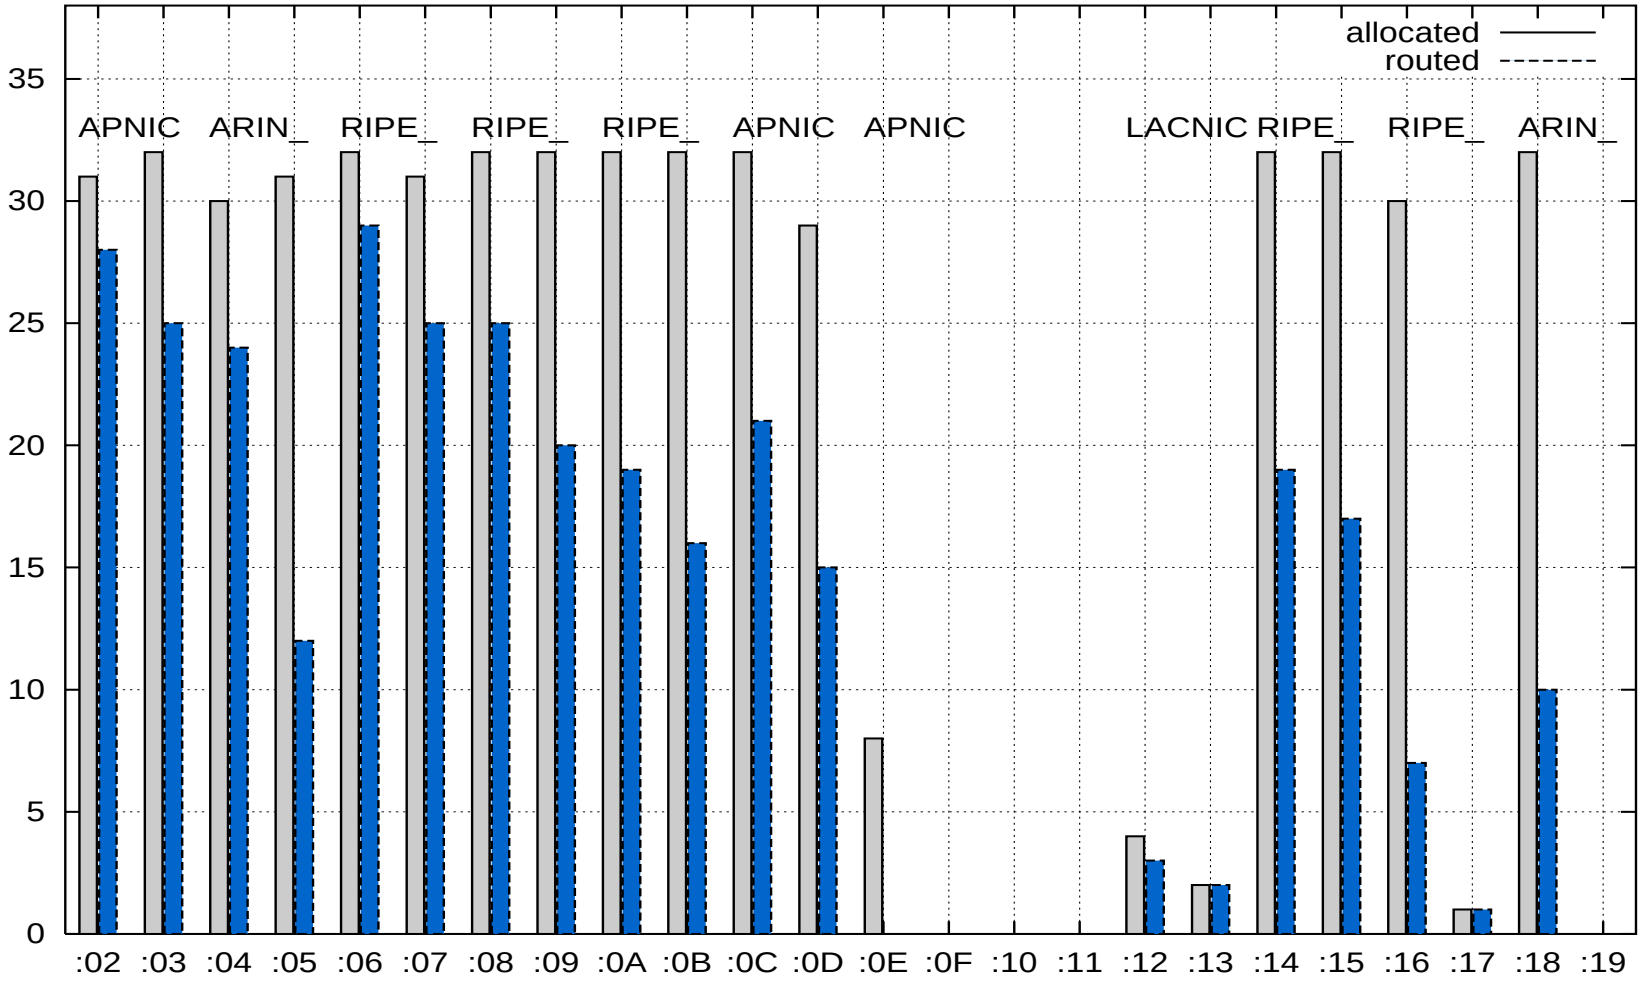

Interesting Observations (1) - Ghost Busting

```
   Network          Path
*> 2001:200::/32    3549 2500 i
*                   1752 4725 2500 i
*                   1930 20965 11537 2500 i

*  3FFE:2200::/24
*        3549 2915 7660 11537 3425 293 5609
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
*        1752 12853 1275 5609
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
*        1930 1752 12853 1275 5609
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
*        3274 16023 12702 15694 9112 12968 1752 12853 1275 5609
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
*>       9044
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
*        4555 5609
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
*        1853 1853 1853 1853 1853 6680 1103 11537 3425 293 5609
              5623 10318 10566 15589 513 6939 4725 4697 3320 6830 6175 2497 2500 i
```

Ghosts = BGP withdrawal bug, caused by old and buggy software.

Long paths stay *mostly unchanged* in the table for weeks.

Don't confuse with BGP count-to-infinity (= paths change quickly).

Observations (2) - Christmas leaks

```
   Network                 Next Hop                Path
*  2001:600::/32           2001:360:1:2::1         1221 109 109 6939 3257 12702 i
*>i                        2001:7F8::319E:0:1      12702 i
*                          2001:608:0:3::D         5430 12702 i
*                          2001:470:1FFF:2::       6939 3257 12702 i

*> 2001:600:4:10::/64     2001:470:1FFF:2::       6939 6939 5424 12702 ?
*> 2001:600:4:11::/64     2001:470:1FFF:2::       6939 6939 5424 12702 ?
*> 2001:600:4:12::/64     2001:470:1FFF:2::       6939 6939 5424 12702 i
...
*  2001:600:100F::/48     2001:608:0:3::D         5430 3549 12702 ?
*                          3FFE:C00:8023:19::1     109 12702 ?
*>i                        2001:450:1:2001::AA     3549 12702 ?
*                          2001:470:1FFF:2::       6939 6939 5424 12702 ?
...
*> 3FFE:2500:0:401::/64   2001:470:1FFF:2::       6939 6939 5424 12702 ?
*> 3FFE:2500:302::/64     2001:470:1FFF:2::       6939 6939 5424 12702 ?
*> 3FFE:2500:324::/64     2001:470:1FFF:2::       6939 6939 5424 12702 ?
```

Due to a router crash / configuration issue, 12702 leaked 143 pfxs on Dec 24th. Upon notification, this was quickly fixed. Thanks!

Overall, people are aggregating (and filtering!) pretty well.

Observations (3) - 2500 path truncation issue

```
   Network   Next Hop                  Path
*  2001:308::/32
             2001:608:0:3::9           3320 1752 5511 4716 7668 i
*            2001:7F8:2:8001::2        1752 5511 4716 7668 i
*            2001:608:0:3::7           1930 20965 11537 7660 4716 7668 i
*            3FFE:1108:40A:FFFF::1:2   3274 790 790 6830 8472 1752 5511 2500 i  <<<
*            3FFE:C00:8023:19::1       109 6939 4716 7668 i

*  2001:310::/32
             2001:608:0:3::9           3320 1752 5511 2500 4694 i
*            2001:608:0:3::7           1930 20965 11537 7660 2500 4694 i
*            3FFE:1108:40A:FFFF::1:2   3274 790 790 6830 8472 1752 5511 2500 i  <<<

*  2001:C40::/32
             2001:608:0:3::9           3320 1752 5511 4691 18098 i
*>i          2001:650:F807::2118:1     8472 1752 5511 2500 i                    <<<
*            ::FFFF:203.14.5.73        1221 4777 4691 18098 i
```

Something in the outbound path from WIDE (2500) seems to be truncating AS paths. 2500's BGP tables are correct, but e.g. opentransit (5511) sees prefixes 'originated' from 2500. Hitachi bug?

People are actively looking into this.

Interesting Observations (4) - Invalid AS numbers

```
   Network     Next Hop                Path
*  2001:568::/32
*>i            2001:650:F807::DE9:1     3561 145 6435 6342
                                          64600 4787 64702 29155 3263 852 e
*              3FFE:8150:0:1::17        9044 5424 1275 15589 33 6435 6342
                                          64600 4787 64702 29155 3263 852 e
*              2001:608:0:3::7          1930 20965 3549 2500 2497 6175 6435 6342
                                          64600 4787 64702 29155 3263 852 e

*  2001:D10::/32
*              2001:478:FFFF::1         4555 109 6342 64600 4787 i
*              2001:7F8:2:8001::2       1752 5408 10566 3758 4787 i
*>             3FFE:C00:8023:19::1      109 6342 64600 4787 i

*  3FFE:3100::/24
*              3FFE:C00:8023:19::1      109 6342 64600 4787 1930 i
*>             2001:608:0:3::7          1930 i
```

again: prolonged leaks of private ASns to the global table :-(

private ASns only seen as transit ASes. Why?

News (?)

- 6bone (3FFE:...) going away, end date: 2006/06/06
- no martian (bogon) networks seen since 2002/10/21
- private/unallocated AS numbers reappearing ?!
- ghost routes are reappearing ?!
- are “early” IPv6 networks already deteriorating?
- more people actually look at traceroutes and fix things
- overall structure really improving, towards production quality (to be defined as: IPv6 path is no worse than the IPv4 path)
- US region catching up on allocations, but still lacking far behind on actually advertised routes

Where to go from here?

- more work needed on filtering recommendations
- more work on "routing BCP" recommendations (→ routing wg)
- still **much** cleanup work to do ("bad" tunnels, filters, unsolicited transit relations)
- bug your upstream providers to offer native IPv6 upstream
- have an eye on traceroute(6)s to find out which ways packets are travelling, and resolve stupid paths if possible
- consider de-peering non-useful peers (bad tunnels)
- *talk* to your peers and help them fix their stuff

IPv6 routing recommendations

- MIPP project recommendations:
  - no peerings over 'bad' tunnels (high RTTs / 3rd parties)
  - apply incoming prefix filters to peers
  - filter private ASn and overly long paths
- do not give unrestricted IPv6 transit to peers unless asked to
- do not take IPv6 transit from too many upstreams
- avoid taking your single upstream over intercontinental tunnel

References

- Ghost Route Hunter: http://www.sixxs.net/tools/grh/
- Merit 6bone routing report: http://www.merit.edu/mail.archives/html/6bone-routing-report/
- List of IPv6 blocks allocated by the RIRs: http://www.ripe.net/rs/ipv6/ipv6allocs.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/R47-v6-table/

Questions?

gert@space.net