

Have We Reached 1000 Prefixes Yet?

A snapshot of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich, Germany

May 6, 2008

LACNIC IX, Salvador, Brazil

presented by Roque Gagliano

And the answer is...




Overview

- pictures & trends
- the end of the 6bone
- numbers...
- things that should not be there...
- route6 current practices
- references

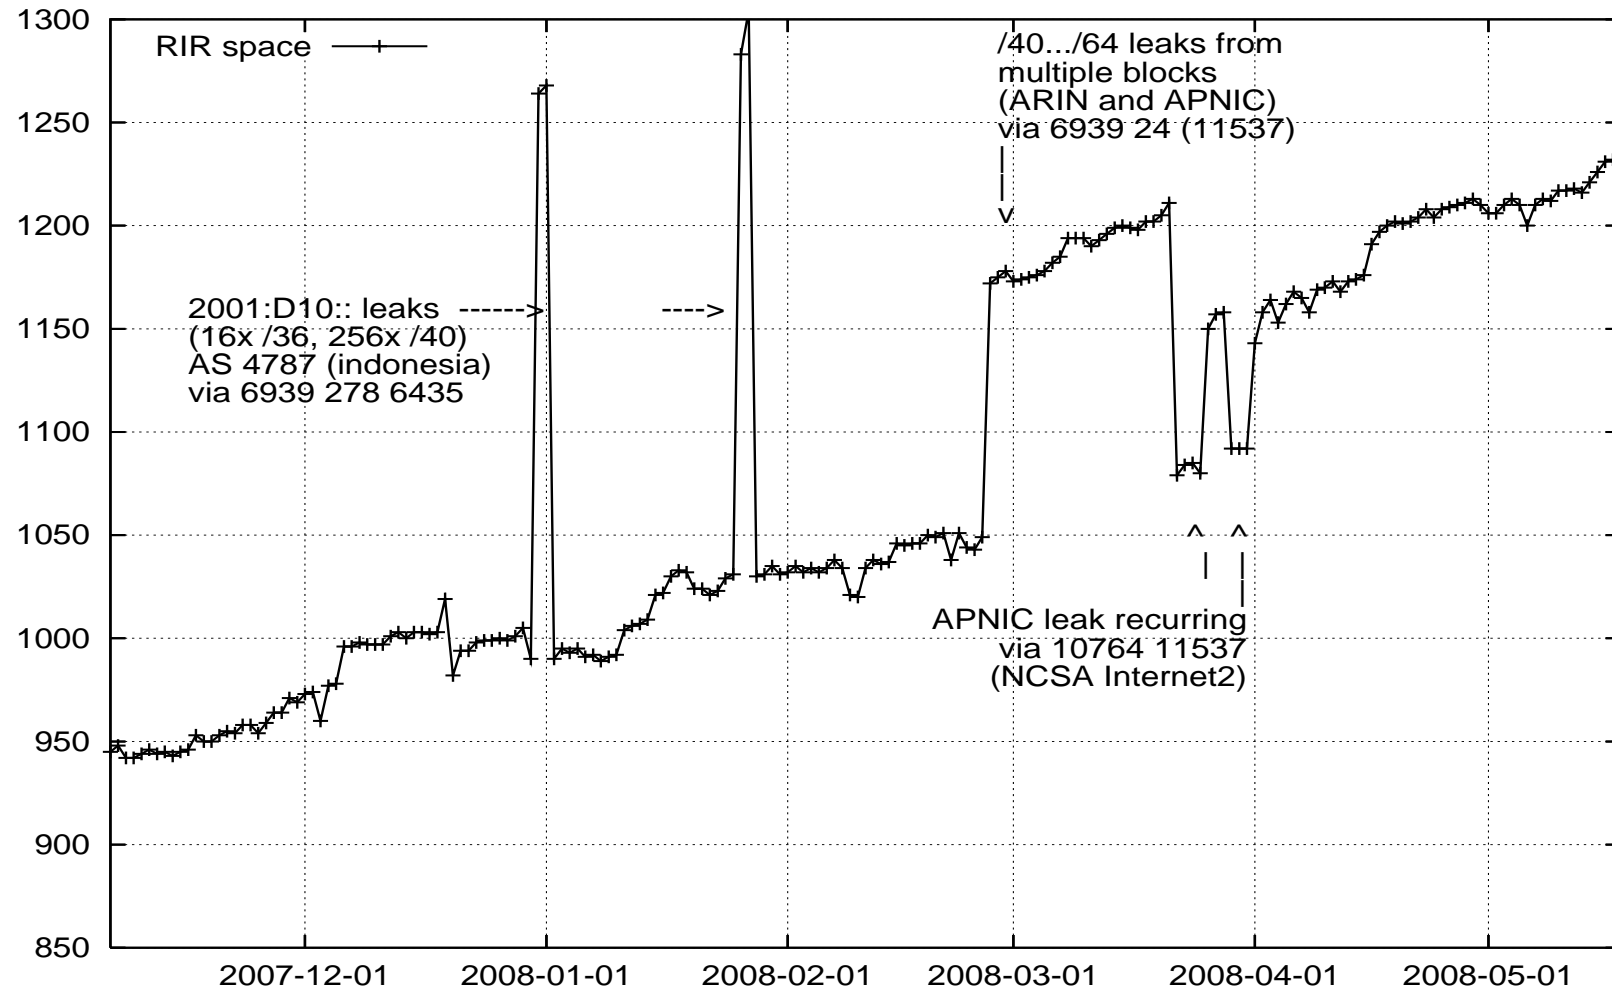
Slides online at:

<http://www.space.net/~gert/RIPE/LACNIC-XI-v6-table/>


Graphics: Total Prefixes - 6.5 years




Graphics: zoom into last 6 months




Graphics: trends? (RIR prefixes, 24 months)




Graphics: prefixes by RIR region



Graphics: prefixes by country (RIPE)



Graphics: prefixes by country (APNIC)



Numbers - AS numbers

- as of 2008-05-04: 933 unique AS numbers visible (2007-10: 798)
 - 646 origin-only ASes (no transit paths seen) (549)
 - 253 ASes originate & give transit (215)
 - 34 transit-only ASes (e.g. 57, 1239, 5549, 6667, ...) (34)
- different number of prefixes announced
 - 0 ASes originate 6Bone (3ffe) prefixes (*hooray!*)
 - 775 ASes originate 1 RIR prefix (672)
 - 75 ASes originate 2 RIR prefixes (3 due to /32+/35)
 - 28 ASes originate 3 RIR prefixes
 - 22 ASes with “more than that”, max. is 24 & 34 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:F8:8000::/35	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	Path
*> 2001:3a0:8000::/35	5539 2914 4697 2.3 2.7 ?
*> 2001:df0:2::/48	5539 1273 2497 2.3 i
*> 2001:4810:2000::/35	5539 1273 29748 33437 6.3 i
*> 2403:2000::/32	5539 3257 2497 2500 18146 2.9 i

Total number of prefixes 4

- this needs Quagga or IOS XR to see – others will see “2-byte tunnel AS” 23456:

```
*>i2001:3A0:8000::/35
      2001:7F8::B62:0:1      2914 4697 23456 23456 ?
* i      2001:608:0:FFF:1::6  2914 4697 23456 23456 ?
*      ::FFFF:203.14.5.7      1221 4777 38610 30071 2914 4697 23456 23456 ?
*      2001:470:1FFF:2::      6939 2914 4697 23456 23456 i
*      2001:608:0:3::3        1853 1853 3356 2914 4697 23456 23456 ?
*      2001:420:0:7FF0::1     109 5511 2914 4697 23456 23456 ?
```


- see RFC 4893 for the details (and bug your vendor)

Numbers - Prefixes

As of 2008-05-04: 1214 prefixes in total (2007-10-22: 937)

/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	8	4	4	0	0	0	0
/24..27	9	4	4	1	0	0	0
/28..31	9	2	5	0	2	0	0
/32	780	401	167	156	37	17	2
/33..34	19	6	5	8	0	0	0
/35	38	8	20	10	0	0	0
/36	6	3	0	2	1	0	0
/40	23	7	4	8	2	2	0
/41..42	3	2	0	1	0	0	0
/44..47	12	3	2	7	0	0	0
/48	302	50	83	153	9	4	3
/49..63	2	0	1	1	0	0	0
/64..128	0	0	0	0	0	0	0

Graphics - Prefixes



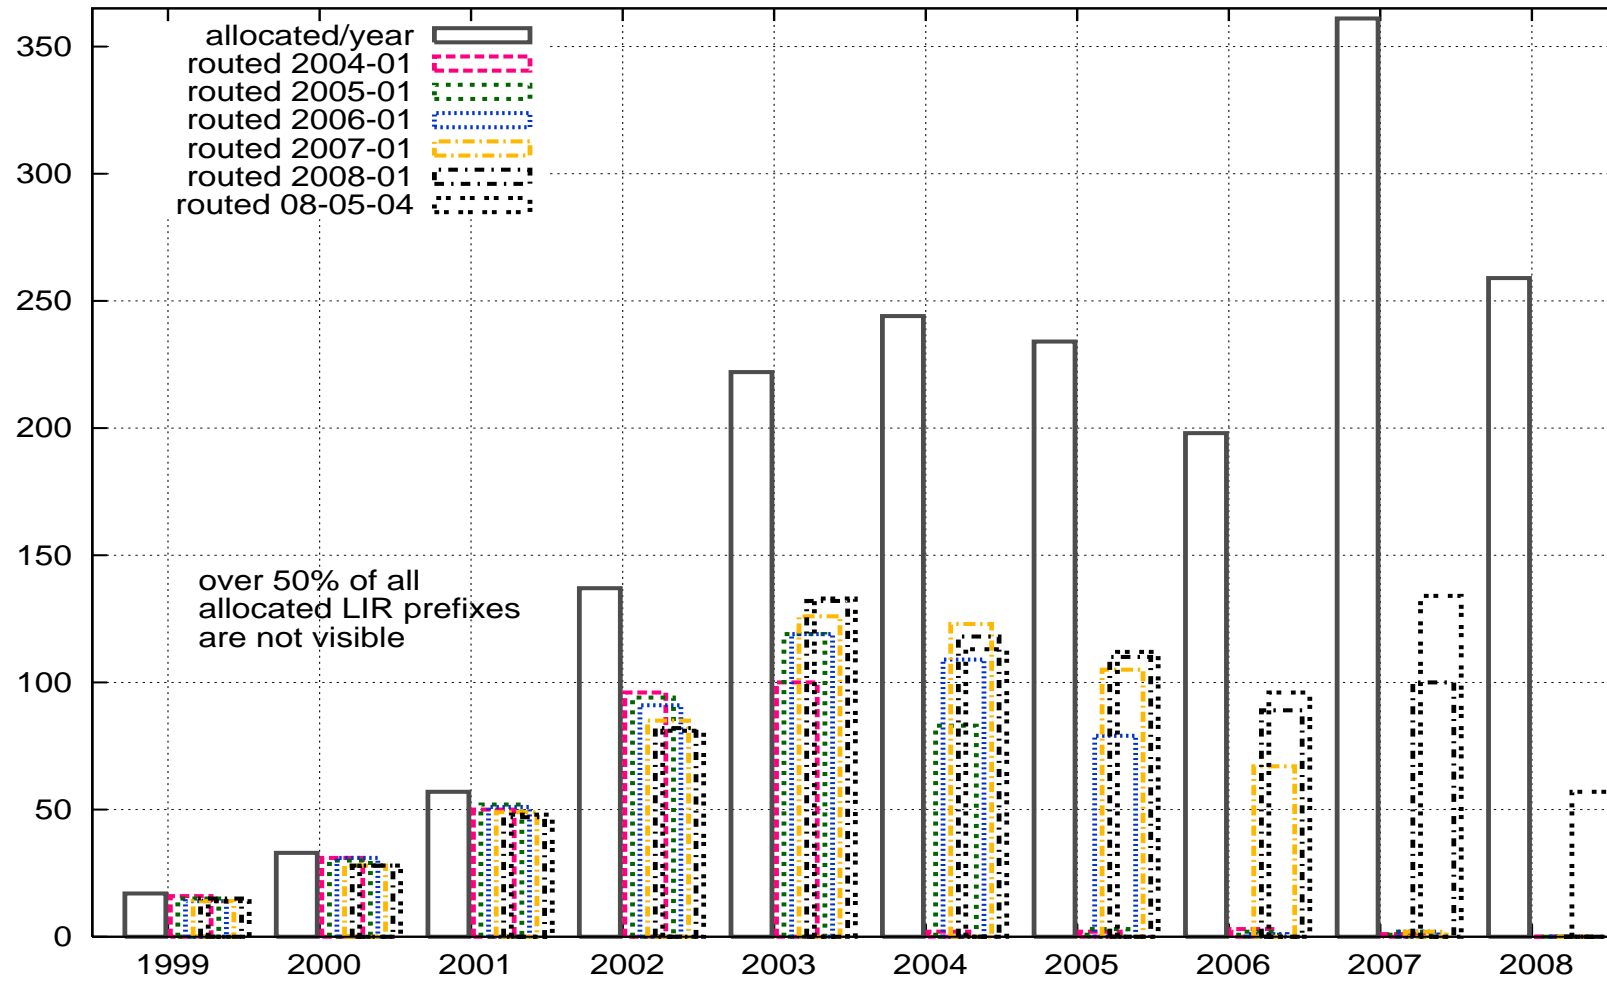
Numbers: RIRs, Allocations, ...

- On 2008-05-04, 1790 LIR blocks (2000:: $/4$) allocated by RIRs:

RIR	alloc.	members	perc.	on 2007-05-06
ARIN	408	~ 2916	14.0%	312 (+31%)
APNIC	366	~ 2742	13.3%	315 (+17%)
RIPE	878	~ 5525	15.8%	617 (+42%)
LACNIC	95	~ 832	11.4%	86 (+10%)
AfriNIC	43	~ 423	10.2%	38 (+13%)

- note: not counting $/48$ microallocs and $/35 \Rightarrow /32$ extentions
- actual *percentage* with IPv6 similar among regions
- 799 (R55: 689) allocations visible in routing table (*only 45%!*)

Graphics: Allocated vs. Routed



Allocated vs. Routed - by region & class

RIR	type	alloc.	visible	perc.	subnets
ARIN	LIR	384	157	41%	73
	IXP	21	0	0%	0
	Critical Inf.	51	24	47%	32
	Internal Inf.	2	0	0%	0
	PI	145	24	17%	37 (*)
APNIC	LIR	362	178	49%	107
	IXP	21	2	10%	0
	PI	16	6	38%	2
RIPE	LIR	878	411	47%	65
	IXP	68	13	19%	0
	Anycast DNS	6	3	50%	0
LACNIC	LIR	94	36	38%	14
	Crit.Inf.+PI	5	1	20%	0
AfriNIC	LIR	42	17	40%	2
	PI	7	4	57%	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

- 9 new IPv6 PI networks from APNIC (4 in BGP)
 - 2001:DE8:2::/48 VoIP eXchange Point Indonesia
 - 2001:DE8:4::/46 Equinix Asia Pacific
 - 2001:DF0:5::/48 NTT West IPv6 Labs (!?)
 - 2001:DF0:6::/48 VOCUS Pty Ltd, Sydney
 - 2001:DF0:7::/48 Microsoft Japan
 - 2001:DF0:8::/48 JPNIC registry services
 - 2001:DF0:9::/48 BarNetwork ISP, Sydney
 - 2001:DF0:a::/48 APNIC Training IPv6 Address
 - 2001:DF0:20::/43 HP IPv6 networks for Asia Pacific region
- 4 IPv6 PI networks from LacNIC (1 in BGP)
 - 2001:13c7:6000::/48 Camara Colombiana de Inform. y Telecom.
 - 2001:13c7:6001::/48 Camara Argent. de Base de Datos y .. Linea
 - 2001:13c7:7000::/48 NIC Mexico
 - 2001:13c7:7001::/48 LACNIC

Numbers: notable allocations (2)

- 2 new IPv6 PI networks from AfriNIC (0 in BGP)
 - 2001:43F8:50::/48 BNR Consulting, South Africa
 - 2001:43F8:60::/48 Kenya Internet Exchange point
- 51 new “direct” assignments from ARIN, 4+(32) in BGP
 - 2620:0:a00::/43 HP-V6-Internet (32x /48)
 - 2620:0:1000::/40 Google IT IPv6 (?!)
- new DNS anycast assignments from RIPE (2001:678::/29)
 - 2001:0678:5::/48 .TM, TM Domain Registry Ltd (2007/10/10)
 - 2001:0678:6::/48 .LT, Kaunas University (2007/12/11) (in BGP)
- (a few) more “very large” allocations seen:
 - 2408::/22 NTT East JP
 - 2a00:800::/25 Tele2 / SWIPnet SE
 - 2a01:c00::/26 HanseNet DE (DSL provider)
 - 2a01:e00::/26 Proxad FR (DSL provider)
- ⇒ **check your BGP filters!!**

Interesting Observations - New Year's Spikes

```

Network          Path
* 2001:D10::/32  1221 30071 2914 10204 4787 i
*                6939 2914 10204 4787 i
*                109 6175 4787 i
*>i             3257 6175 4787 i

* 2001:D10::/36  6939 278 6435 4787 i
* 2001:D10::/40  6939 278 6435 4787 i
* 2001:D10:100::/40 6939 278 6435 4787 i
* 2001:D10:200::/40 6939 278 6435 4787 i
...
* 2001:D10:E000::/36 6939 278 6435 4787 i
* 2001:D10:E000::/40 6939 278 6435 4787 i
* 2001:D10:E100::/40 6939 278 6435 4787 i
...
* 2001:D10:F000::/36 6939 278 6435 4787 i
...
* 2001:D10:FE00::/40 6939 278 6435 4787 i
* 2001:D10:FF00::/40 6939 278 6435 4787 i

```

- for 2x 2 days, AS 4787 (Indonesia) leaked 16 /36s and 256 /40s
- path via LavaNet Honolulu, Red Academia Mexico, HE.NET
- ⇒ please *check* what you give transit for!

Interesting Observations - Routing Out Of Space

Network	Path
2001:2B8::/32	109 4555 4555 6939 24 11537 17579 1237 17832 i 3257 2497 2500 7660 22388 11537 17579 1237 17832 i 1221 38610 38610 18203 18084 2500 7660 22388 11537 17579 1237 17832 i
2001:2B8:90::/48	6939 24 11537 17579 1237 i
2001:2B8:94::/48	6939 24 11537 17579 1237 i
2001:2B8:9A::/48	6939 24 11537 17579 1237 i
2001:2B8:9C::/48	6939 24 11537 17579 1237 i
...	
2001:4D0:8100::/40	6939 24 i
2001:4D0:8102:20::/64	6939 24 i
2001:4D0:8102:32::/64	6939 24 i
2001:4D0:8102:44::/64	6939 24 i
2001:4D0:8102:52::/64	6939 24 i
...	
2001:504:B:10::/64	6939 24 i
2001:504:B:11::/64	6939 24 i
2001:504:B:30::/64	6939 24 i
...	

- NASA (AS 48) leaked a bunch of more-specifics to HE.NET
- worse, Internet2 (AS 11537) leaks *Korean* more specifics
- leak at NASA fixed after notification by Bernhard Schmidt

Interesting Observations - Research in Progress...?

```
Network          Path
* > 2001:2B8::/32  3257 2497 2500 7660 22388 11537 17579 1237 17832 i
* 2001:2B8:90::/48 109 5511 10764 11537 17579 1237 i
* 2001:2B8:94::/48 109 5511 10764 11537 17579 1237 i
* 2001:2B8:9A::/48 109 5511 10764 11537 17579 1237 i
* 2001:2B8:9C::/48 109 5511 10764 11537 17579 1237 i
...
* > 2001:3C8::/32  3257 2497 2500 7660 24287 24490 24475 4621 i
* 2001:3C8:9007::/48 109 5511 10764 11537 4621 4621 ?
* 2001:3C8:9009::/48 109 5511 10764 11537 4621 4621 ?
* 2001:3C8:9103::/48 109 5511 10764 11537 4621 4621 ?
* 2001:3C8:C103::/48 109 5511 10764 11537 4621 4621 ?
...
```

- Internet2 (AS 11537) now leaks via NCSA and OpenTransit
- Korean RENS 1237 and 17579 only have I2 as transit :(
- Thai REN 4621 has other transit, but still leaks via I2 :(
- long-standing problem: research networks with poor connectivity, due to political issues or plain disinterest

Interesting Observations - PI funnies...

- wondering why AS 33 is announcing 34 prefixes...

Network	Path
2001:DF0:20::/44	3257 6175 33 i
2001:DF0:30::/44	3257 6175 33 i
2620:0:A00::/48	3257 6175 33 i
2620:0:A01::/48	3257 6175 33 i
2620:0:A02::/48	3257 6175 33 i
2620:0:A03::/48	3257 6175 33 i
2620:0:A04::/48	3257 6175 33 i
2620:0:A05::/48	3257 6175 33 i
2620:0:A06::/48	3257 6175 33 i
...	
2620:0:A1D::/48	3257 6175 33 i
2620:0:A1E::/48	3257 6175 33 i
2620:0:A1F::/48	3257 6175 33 i

- PI /43 from APNIC, PI /43 from ARIN
- checking BGP paths and traceroutes suggest “same router”
- so one has to wonder why?

Interesting Observations - Zero matters!

- from 2008-03-28 ... 2008-04-05 we saw...

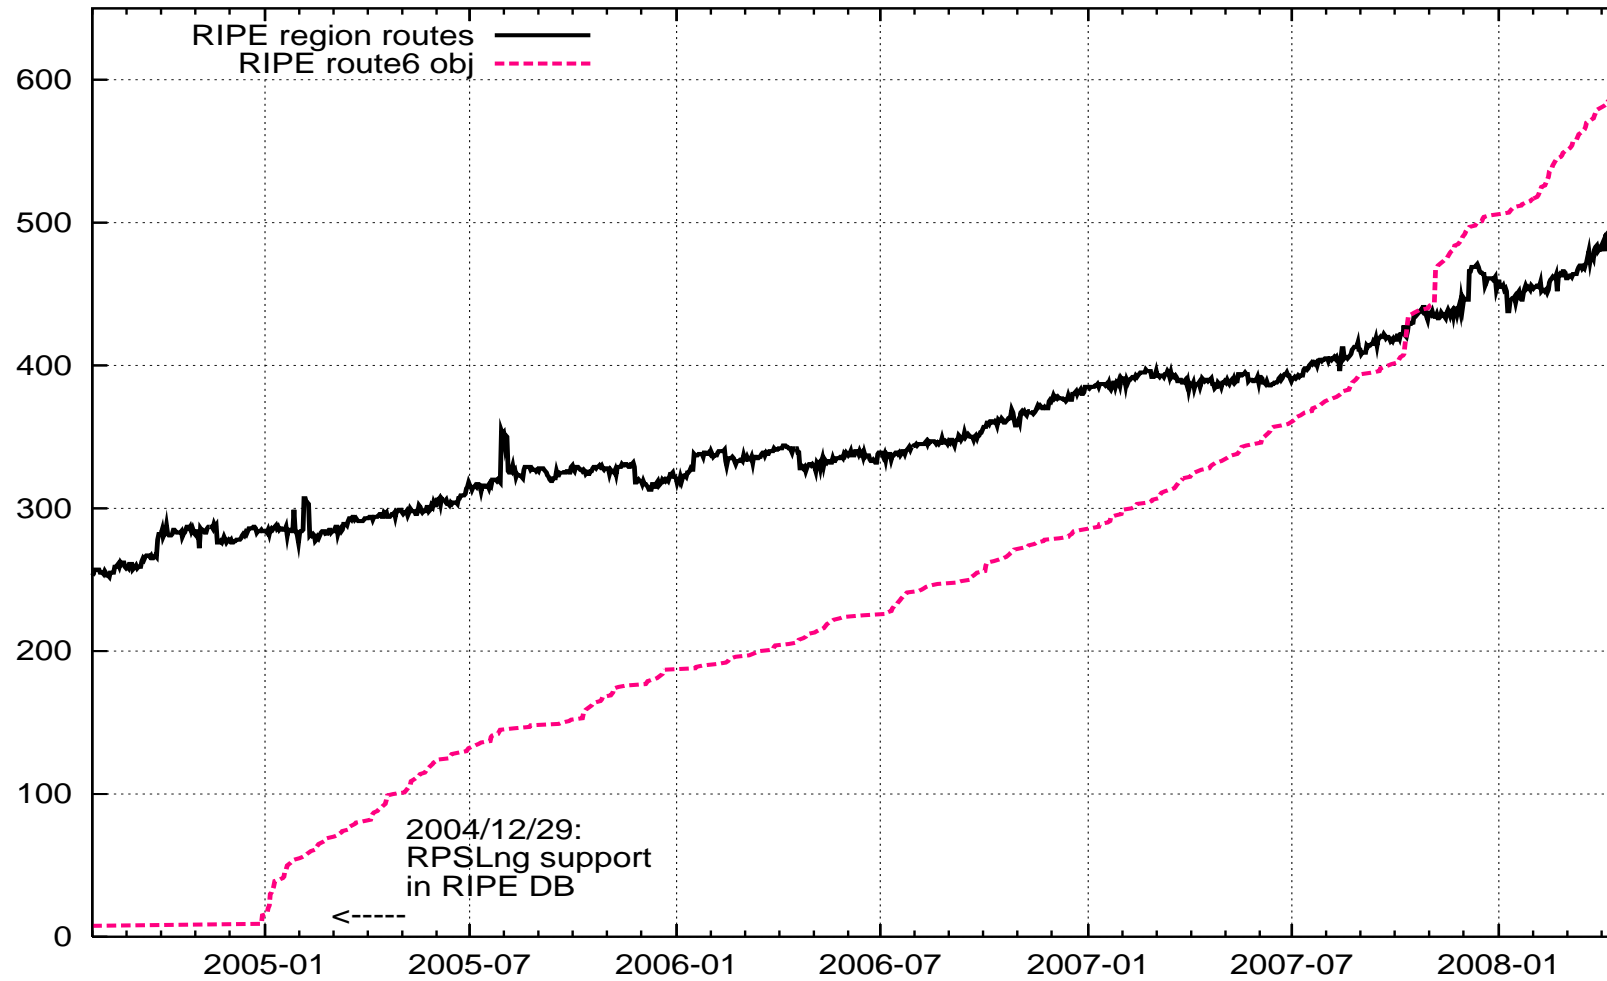
Network	Path
2607:F35::/32	6939 27630 i
	1273 6939 27630 i
	109 30071 6939 27630 i
	1853 3356 3257 6939 27630 i

- ... and now...

Network	Path
2607:F350::/32	6939 27630 i
	1273 6939 27630 i
	109 30071 6939 27630 i
	1853 3356 3257 6939 27630 i

- looks very much like a typo, but should not have been visible
- ⇒ better downstream BGP filters needed!

Graphics: route6 objects vs. routes seen



route6 correlation (RIPE region)

- on 2008-05-04:
 - 492 BGP routes from RIPE region
 - 630 route6: objects in RIPE DB
- correlation?
 - multiple origin route6's (11x 2002::/16, 5x 2001::/32, ...)
 - \Rightarrow 612 route6 objects for *unique* prefixes
 - 43 route6 objects for prefixes from *other* RIRs...
- so...

route6 correlation (2)

- ... and this is what I found:

RIPE prefix, route6 ok	358	:-)
RIPE prefix, route6 missing	125	!!!
RIPE prefix, route6 origin mismatch	8	
RIPE prefix, BGP inconsistant AS	6	
route6 objects without BGP route	201	???
other region, route6 ok	40	
other region, route6 missing	679	
other region, route6 origin mismatch	3	
other region, BGP inconsistant AS	9	

- \Rightarrow 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

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/LACNIC-XI-v6-table/>

Questions?

gert@space.net