

Are We Growing Fast Enough?

A snapshot of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich, Germany

May 3, 2011

RIPE 62, Amsterdam

Overview

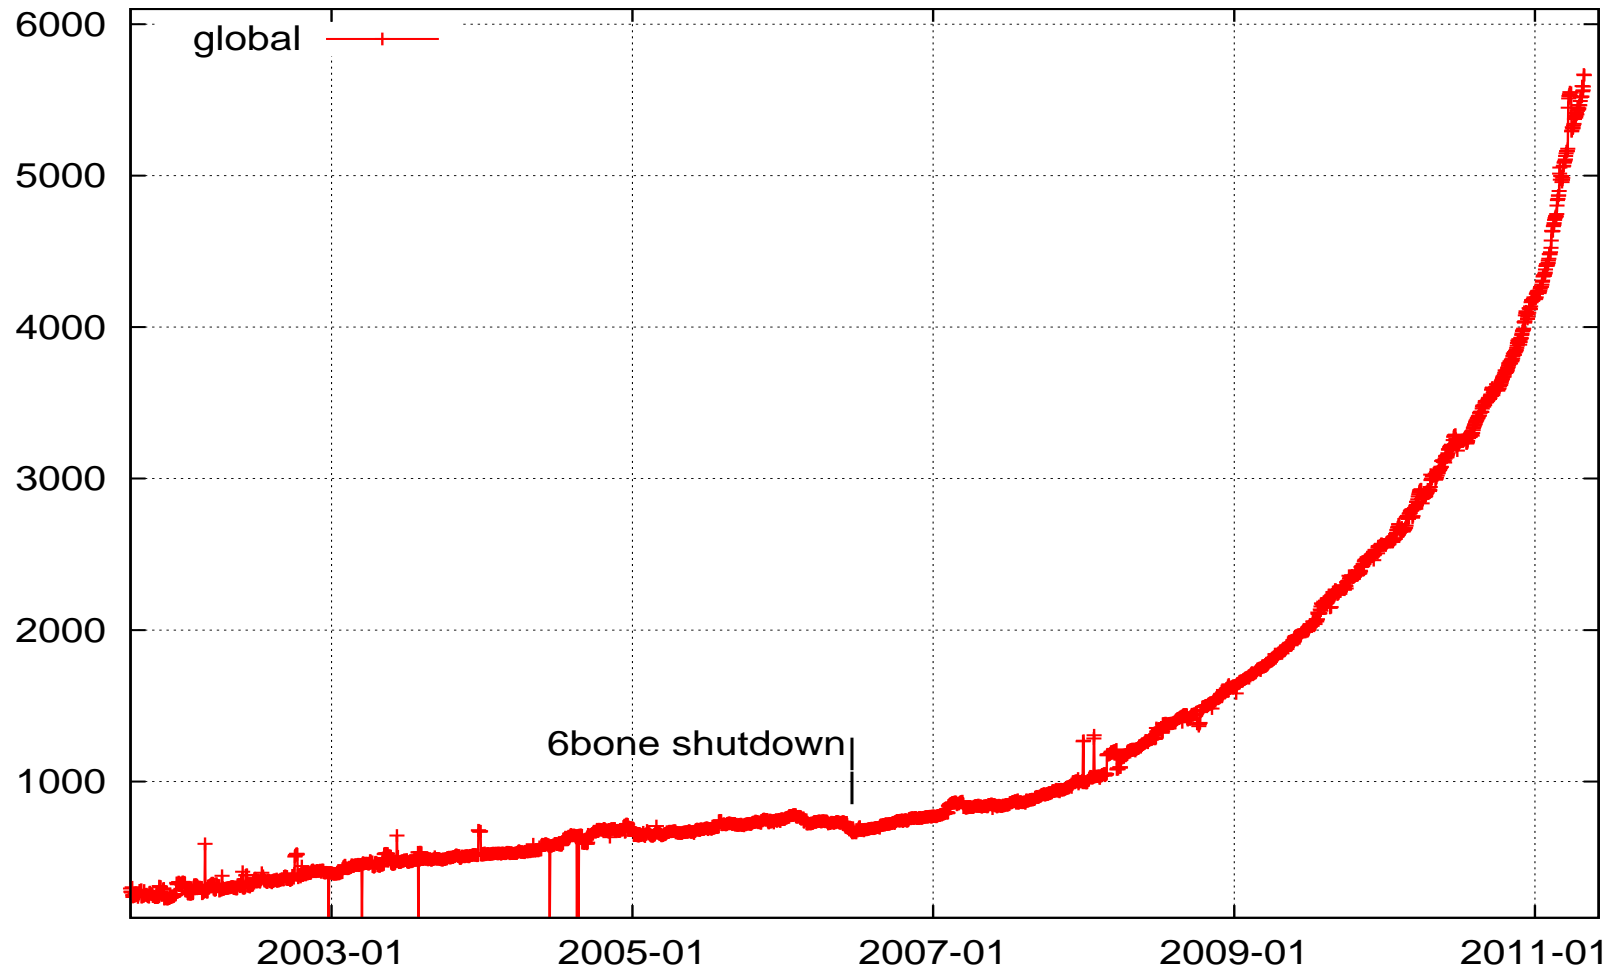
- pictures & trends
- numbers...
- things that should not be there...
- route6 current practices
- references

Slides online at:


<http://www.space.net/~gert/RIPE/R62-v6-table.pdf>

Prefixes in BGP Table

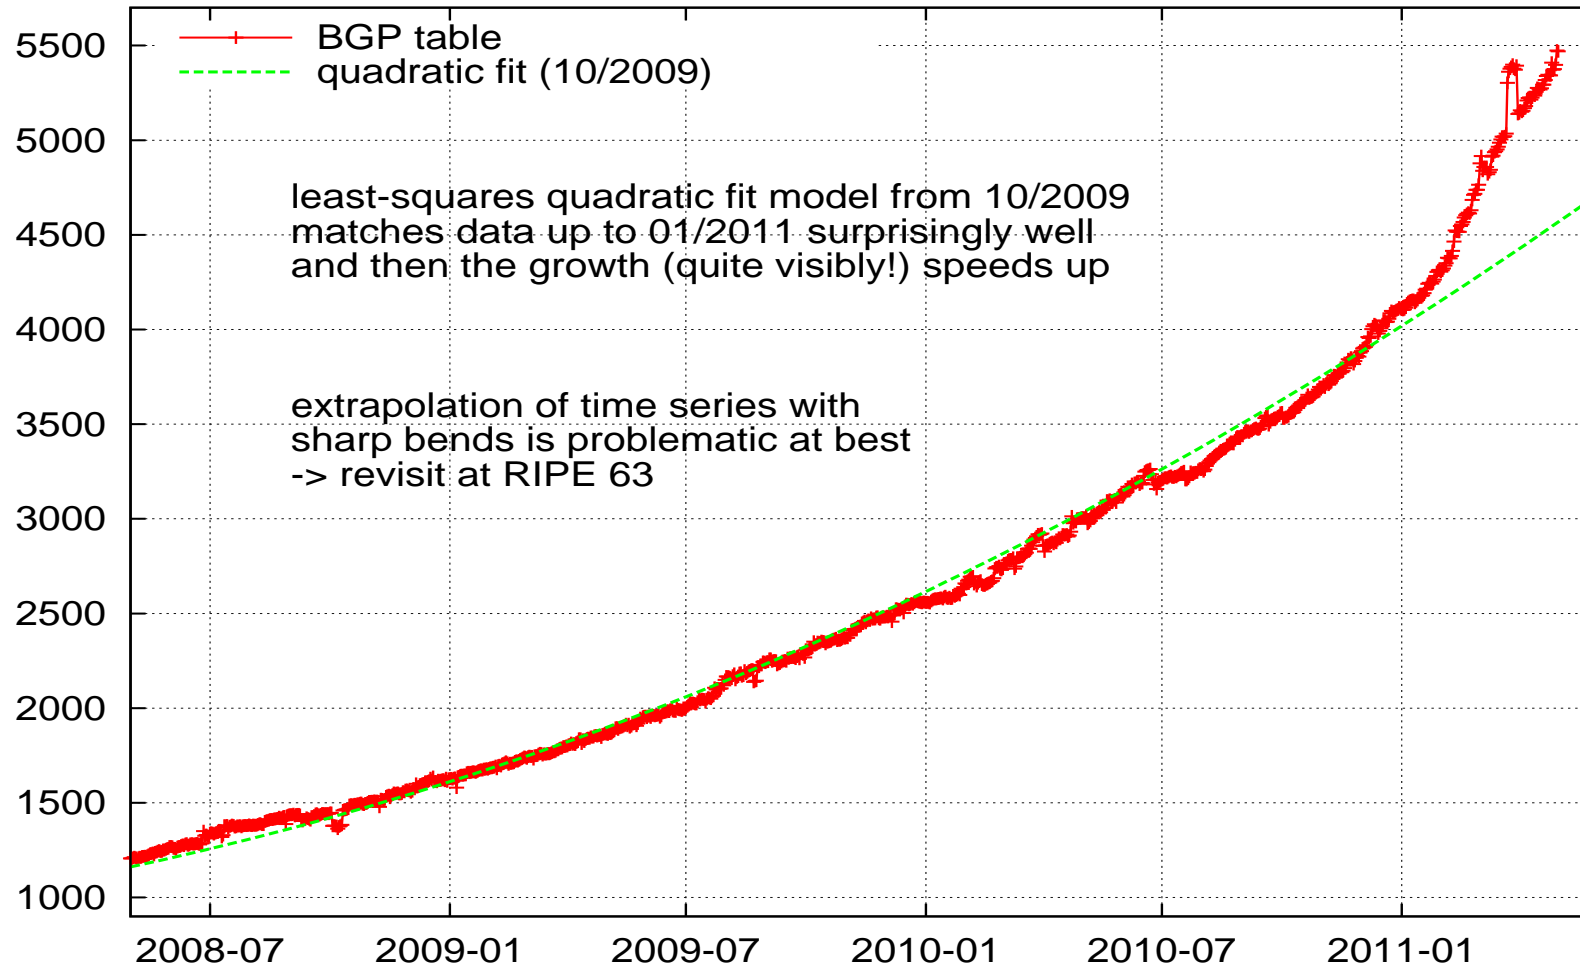
Graphics: Total Prefixes - 9.5 years




Graphics: zoom into last 6 months




Graphics: trends? (36 months)



Graphics: prefixes by RIR region




Graphics: prefixes by country (RIPE)



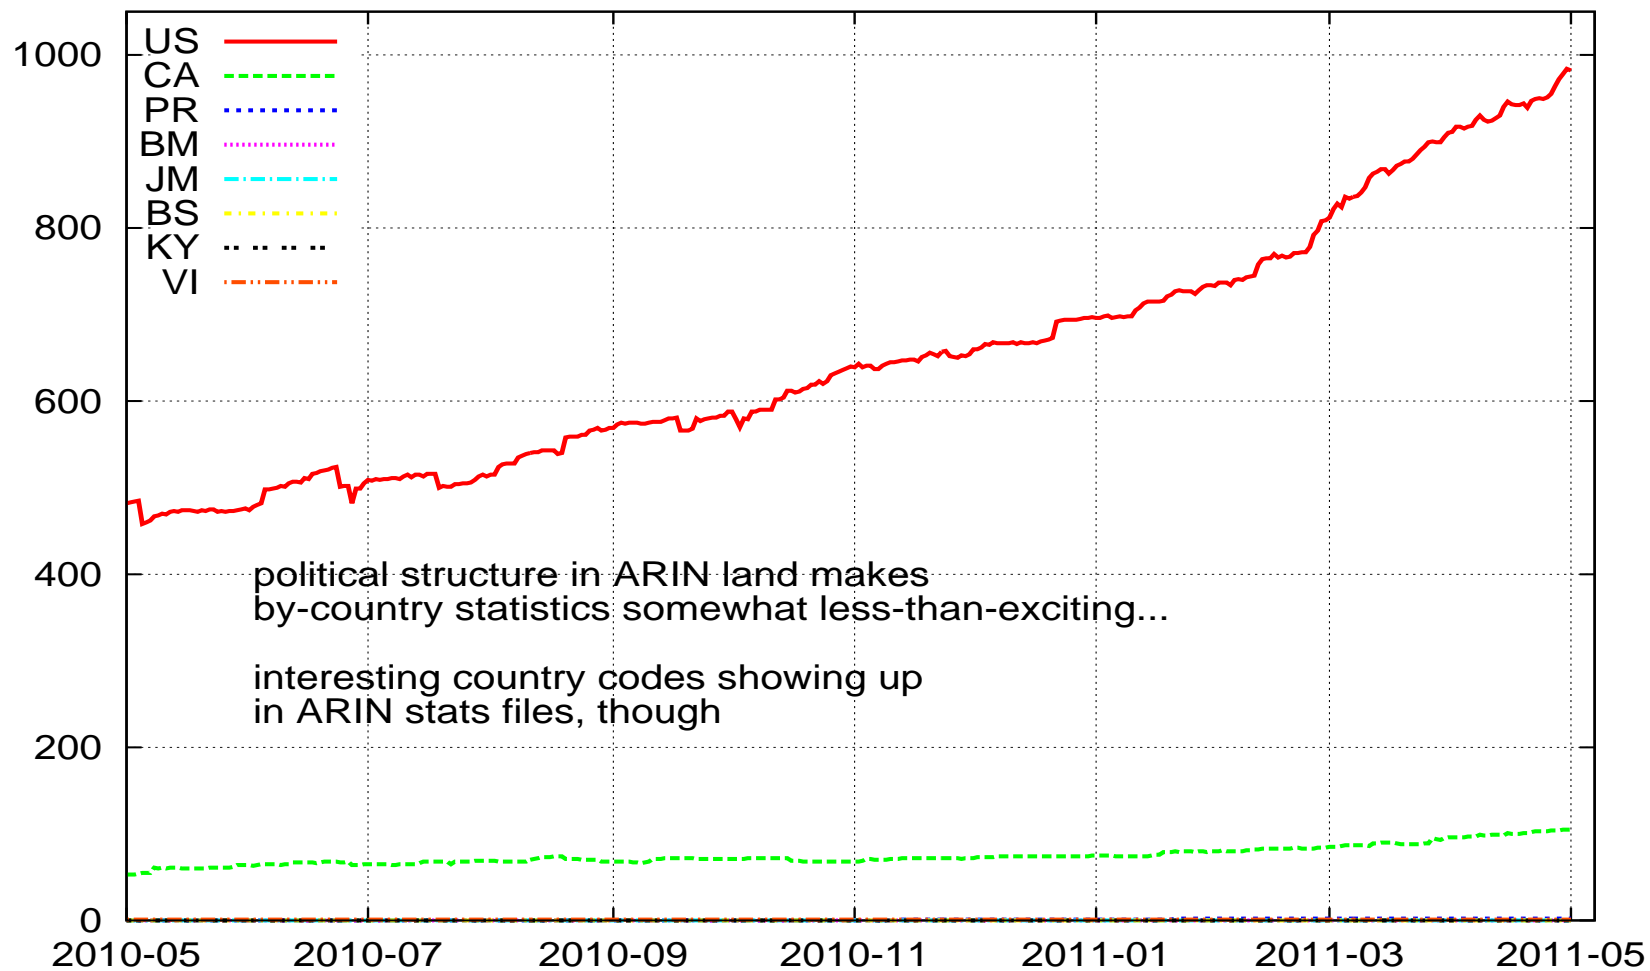
note: graph shows top 8 countries by number of RIR allocations

Graphics: prefixes by country (APNIC)




note: graph shows top 8 countries by number of RIR allocations

Graphics: prefixes by country (ARIN)




note: graph shows top 8 countries by number of RIR allocations

Graphics: prefixes by country (LACNIC)



note: graph shows top 8 countries by number of RIR allocations

Graphics: prefixes by country (AfrinIC)




note: graph shows top 8 countries by number of RIR allocations

AS Numbers

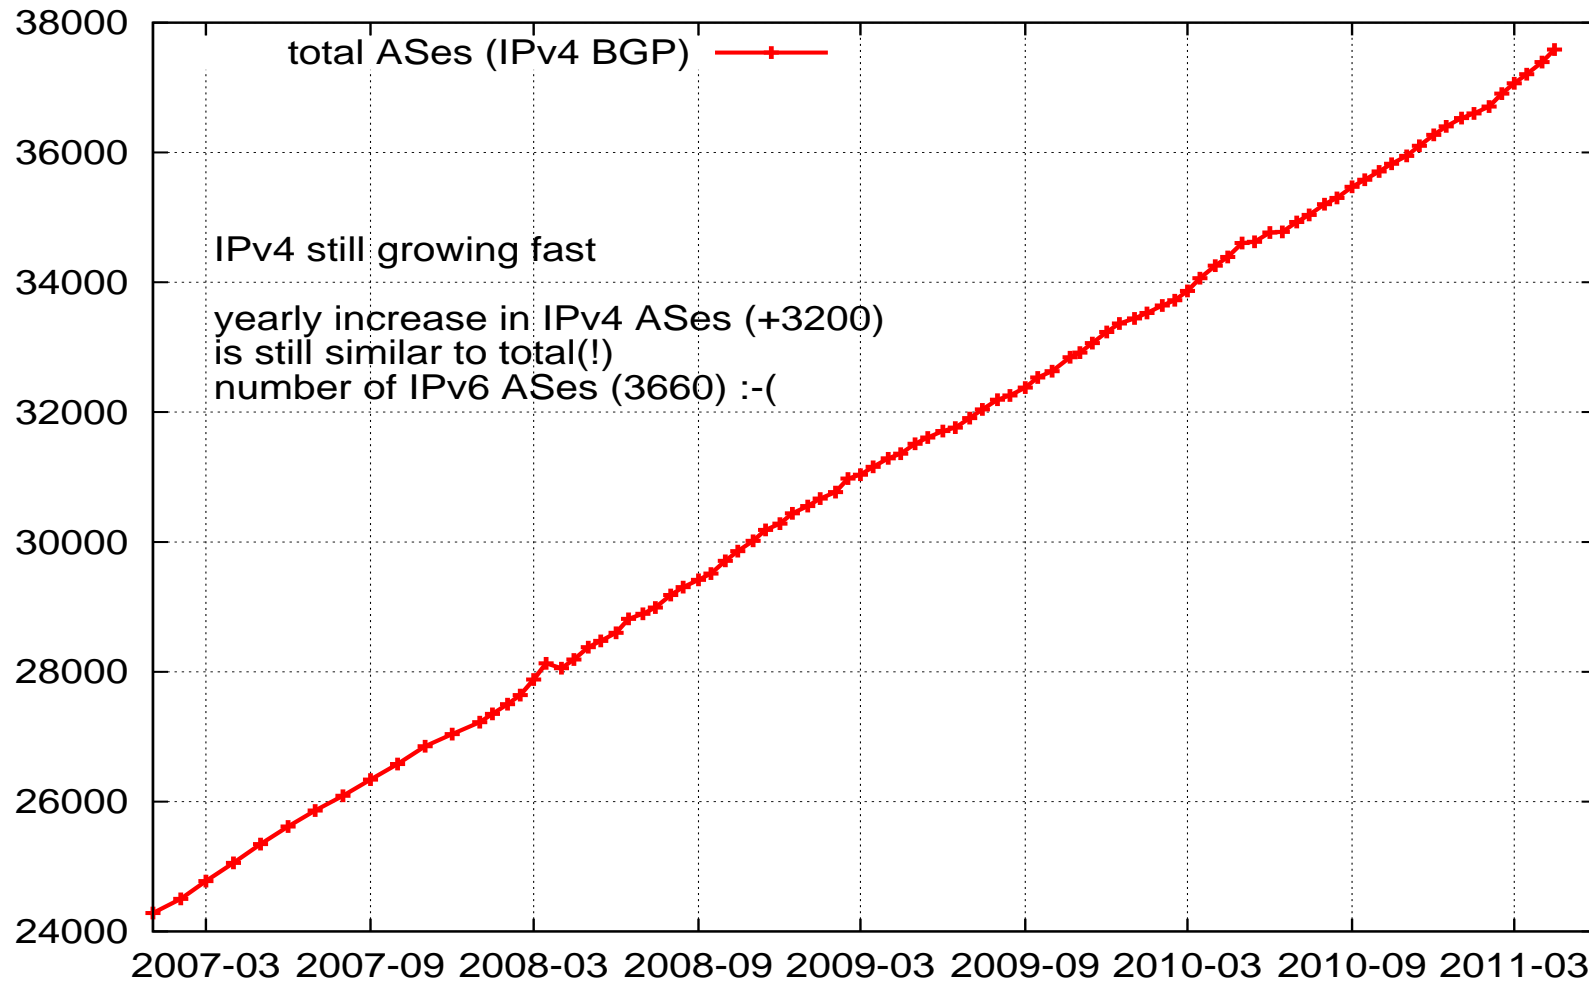
Numbers - AS numbers

- as of 2011-04-30: 3661 unique AS#s visible (2009-10-04: 1674)
 - 2777 origin-only ASes (no transit paths seen) (1207)
 - 831 ASes originate & give transit (419)
 - 53 transit-only ASes (e.g. 1998, 2698, 3754, 12702, ...) (48)
- different number of prefixes announced
 - 3009 ASes originate 1 prefix (2457)
 - 324 ASes originate 2 prefixes
 - 121 ASes originate 3 prefixes
 - 46 ASes originate 4 prefixes
 - 108 ASes with “more than that”, max. is 55 & 87 prefixes
- note: all paths observed from AS5539
- note: numbers skewed by ~67 x 4-byte-ASes (old collector sw)

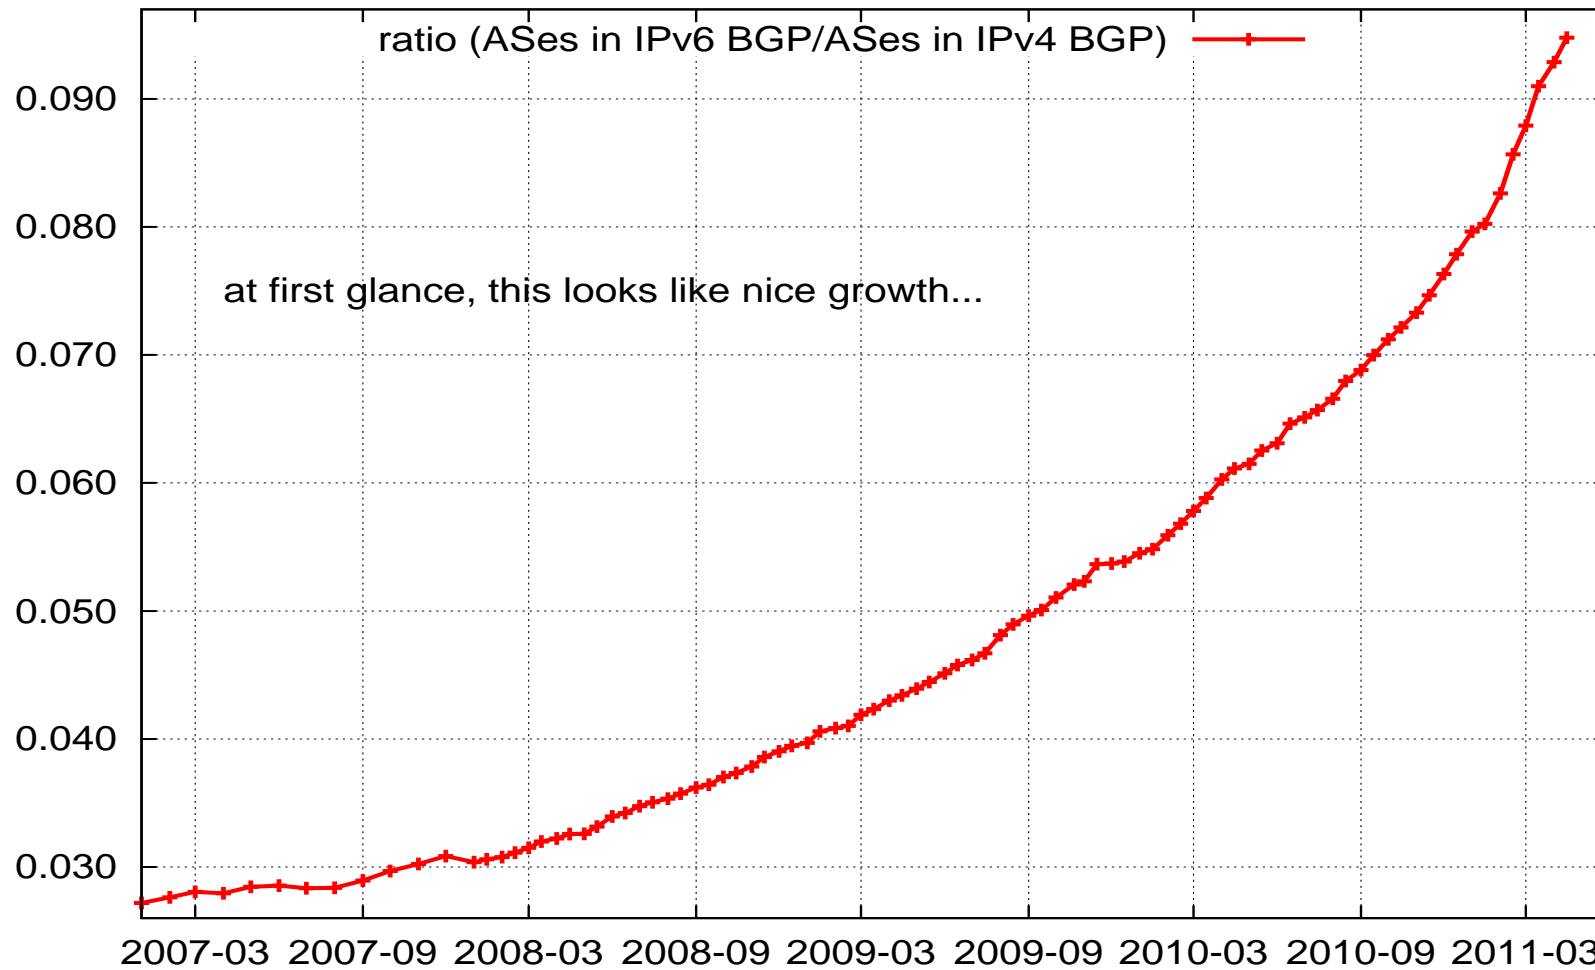
Graphics: AS Numbers (v6 BGP)




Graphics: AS Numbers (v4 BGP)



Graphics: AS Number Ratio (v6 BGP/v4 BGP)



Graphics: AS Number Ratio (v6 BGP/v4 BGP)




Numbers: ASes and AS ratio

- so we see nice growth in ASes participating in IPv6 BGP
- ...but IPv4 BGP growth is still very fast
- ratio of “ASes with IPv6” to “ASes with IPv4” *is* growing
- ratio is not growing fast enough!
- at the current growth rate, it will take 6+ years for all ASes to have IPv6
- **brace yourself for the impact!**

(...insert small picture of train wreck here)

Graphics: Prefixes per AS (v4+v6)



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
2001:500:13::/48	2914 i	ARIN PI
2001:728::/32	2914 i	Verio Europe

- networks with multiple sites and multiple PI prefixes

2001:500:16::/48	3257 6453 12041 i	Afilias
2001:500:17::/48	12041 i	Afilias
2001:500:18::/48	12041 i	Afilias
2001:500:19::/48	2914 12041 i	Afilias
2001:500:1A::/48	3257 6453 12041 i	Afilias
2001:500:1B::/48	12041 i	Afilias

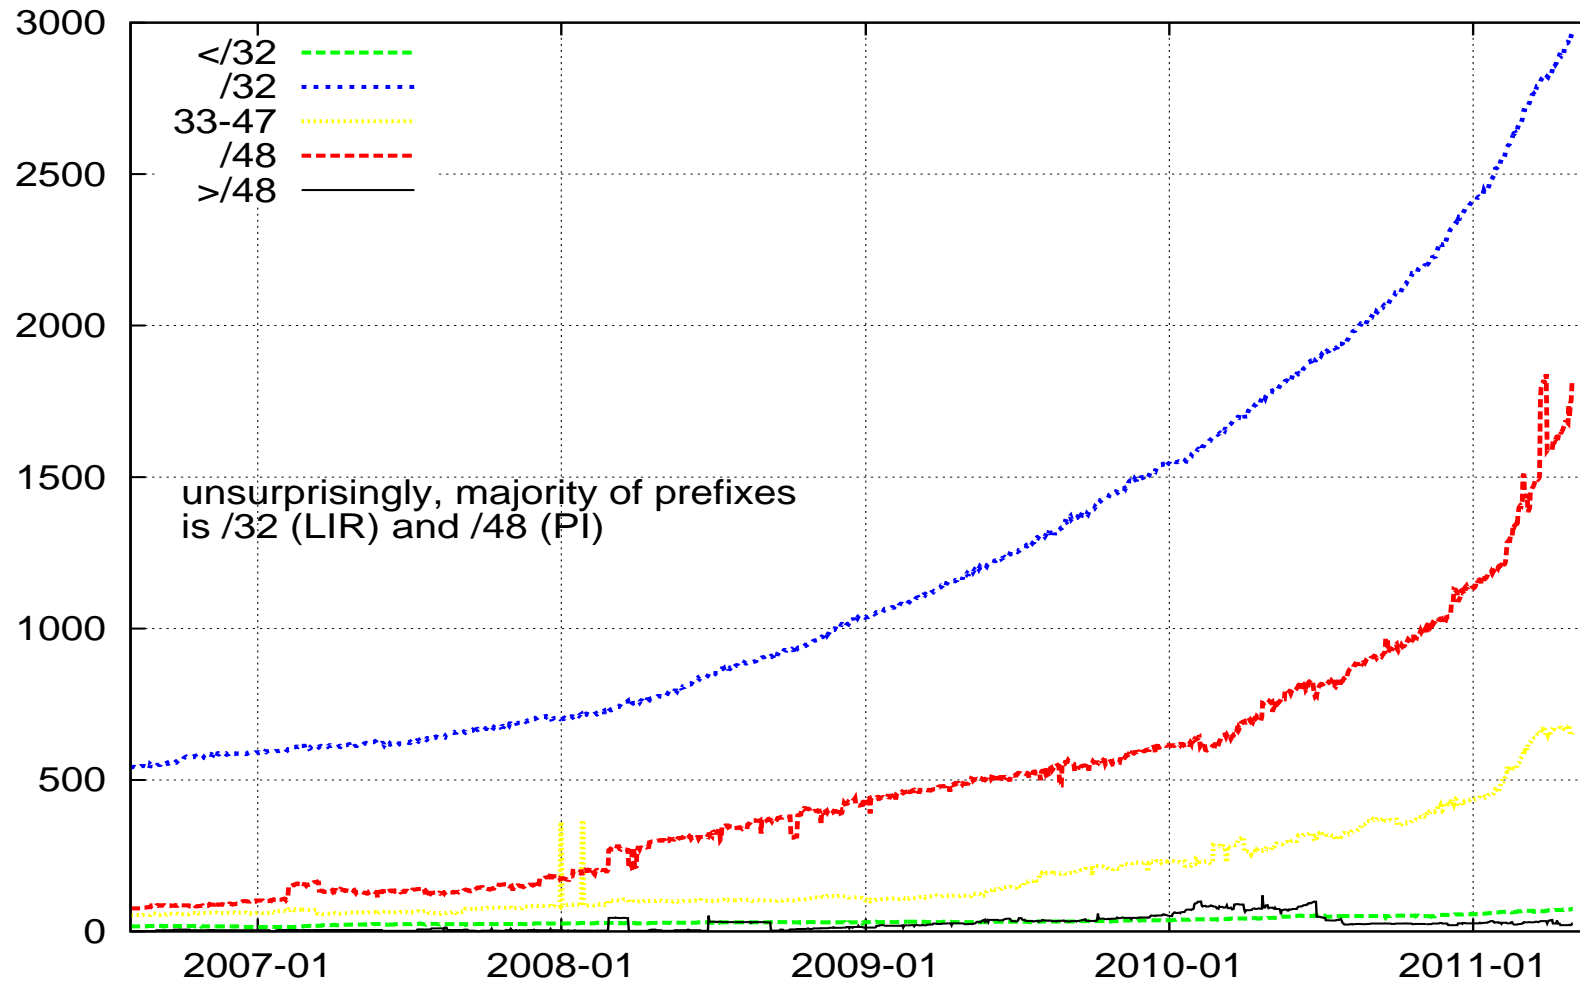
Looking at Prefixes again

Numbers - Prefixes

As of 2011-01-23: 5530 prefixes in total


/n	global	RIPE	APNIC	ARIN	Lacn.	Afri.	oth
/12,/16	2	0	1	0	0	0	1
/19	2	2	0	0	0	0	0
/20..23	14	7	7	0	0	0	0
/24..27	22	12	5	3	0	2	0
/28..31	34	9	9	9	5	0	2
/32	2963	1666	470	646	144	34	3
/33..34	88	45	17	18	8	0	0
/35	75	19	32	16	8	0	0
/36	119	15	55	28	21	0	0
/37..39	28	7	4	17	0	0	0
/40..41	182	36	36	100	8	2	0
/42..47	161	43	33	82	2	1	0
/48	1814	491	485	765	42	31	0
/49..63	7	0	1	6	0	0	0
/64..128	19	16	0	0	0	0	3

Graphics - Prefixes / Size




Allocations

Graphics: Allocations over Time



Graphics: Allocated vs. Routed (LIRs)



Allocated vs. Routed - by region & class

RIR	type	alloc.	visible	perc.	subnets	
ARIN	LIR	1365	596	44%	631	
	IXP	31	1	3%	0	
	Critical Inf.	127	53	42%	118	
	PI	638	178	28%	113	(*)
APNIC	LIR	1300	470	36%	585	
	IXP	20	3	15%	0	
	PI	298	65	22%	31	(*)
RIPE	LIR	3071	1681	55%	353	
	IXP	97	27	28%	2	
	Anycast DNS	29	20	69%	0	
	PI	394	241	61%	18	
LACNIC	LIR	146	45	31%	147	(NIR)
	Crit.Inf.+PI	74	77	38%	16	
AfriNIC	LIR	131	36	28%	29	
	PI	38	5	13%	0	

Notable Allocations

- numerous “large” allocations to Telcos in RIPE land last year:
 - 2a02:1200::/27 CH Swisscom IP-Plus (2010/01)
 - 2a02:1400::/26 SE Telenor Sverige AB (2010/03)
 - 2a02:1800::/24 BE Telenet N.V. (2010/02)
 - 2a02:2e00::/27 ES Jazz Telecom S.A. (2010/11)
 - 2a02:2f00::/28 RO RCS & RDS SA (2010/11)
 - 2a02:3000::/23 DE Telefonica O2 Germany (2010/11)
 - 2a02:8100::/27 DE Kabel Deutschland Breitband (2011/02)
 - 2a02:8200::/27 DE EWE Tel GmbH (2011/02)
 - 2a02:8400::/25 FR LDCOM (2010/12)
 - 2a02:8800::/24 GB Virgin Media Limited (2011/02)
 - 2a02:9000::/23 ES Telefonica de Espana SAU (2011/03)
- similar allocations seen in other RIR regions
- this is a *very good* sign - hints at serious deployment plans

Weirdos

- this talk started out of a remark, “hey, there are some weird things in the BGP table, people might want to know about them”
- well, some 9 years later, there *still* are weird things...

Fat Fingers

Network	Path
* 2008:5A0::/32	109 6939 27792 i
*	3320 6939 27792 i
*>i	6939 27792 i

- in Table 2010/12/10-2011/05/02
- 2008:* is not allocated by IANA to any RIR yet (*bogon!*)
- cross-check with AS number: should be **2800:5A0::/32**
- upstream provider 6939 happily provides transit
- *fixed* in 30 minutes after talking to HE.NET - thanks!

More Fat Fingers

Network	Path
*> 2403:E200::/32	1221 38610 38610 23899 23962 45543 i
*> 2403:E200:00::/48	1221 38610 38610 23899 23962 45543 i
*> 2403:E200:01::/48	1221 38610 38610 23899 23962 45543 i
*> 2403:E200:02::/48	1221 38610 38610 23899 23962 45543 i
*> 2403:E200:03::/48	1221 38610 38610 23899 23962 45543 i
...	
*> 2403:E200:FD::/48	1221 38610 38610 23899 23962 45543 i
*> 2403:E200:FE::/48	1221 38610 38610 23899 23962 45543 i
*> 2403:E200:FF::/48	1221 38610 38610 23899 23962 45543 i

- in Table 2011/03/23-2011/03/30
- LIR announced /32 plus 256x /48 more specifics
- smells like a “measuring artefact” via VNIX and APNIC – “real” announcement of /32 (+ 2 /48) via 6939 starts at 2011/04/08

Really Fat Fingers / Router Bugs?

Network	Path
*>i 1A:F000::/24	3257 174 376 376 376 36786 i
* i	3257 174 376 376 376 36786 i
* i 1000::/8	1273 174 376 376 376 i
*>i	1273 174 376 376 376 i
* 10C0:1100::/24	109 6175 1239 174 376 376 376 851 i
*>i	1273 174 376 376 376 851 i

- in Table 2010/05/15-2010/07/06
- 1***:* is not allocated by IANA to any RIR yet (*bogon!*)
- upstream provider 174 happily provides transit
- AS36786 also announces 2620:0:1AF0::/48

Do You Need More?

- there's a few more bogon prefixes in the IPv6 BGP table...
- typos at leaf nodes **do** happen
- but open and unfiltered transit setup at upstream ISP should be highly embarrassing - **you** are the experts, don't permit typos to propagate to the world
- prefix filtering for IPv6 BGP customers is straightforward "BGP 101"
- OTOH, the slow start in setting up IPv6 IRRs in most regions wasn't helping things either
- as of today, at least the 3 big RIRs have functional IPv6 IRR-DBs, so *there is no excuse anymore for not filtering incoming customer BGP announcements*

Most Specific Wins

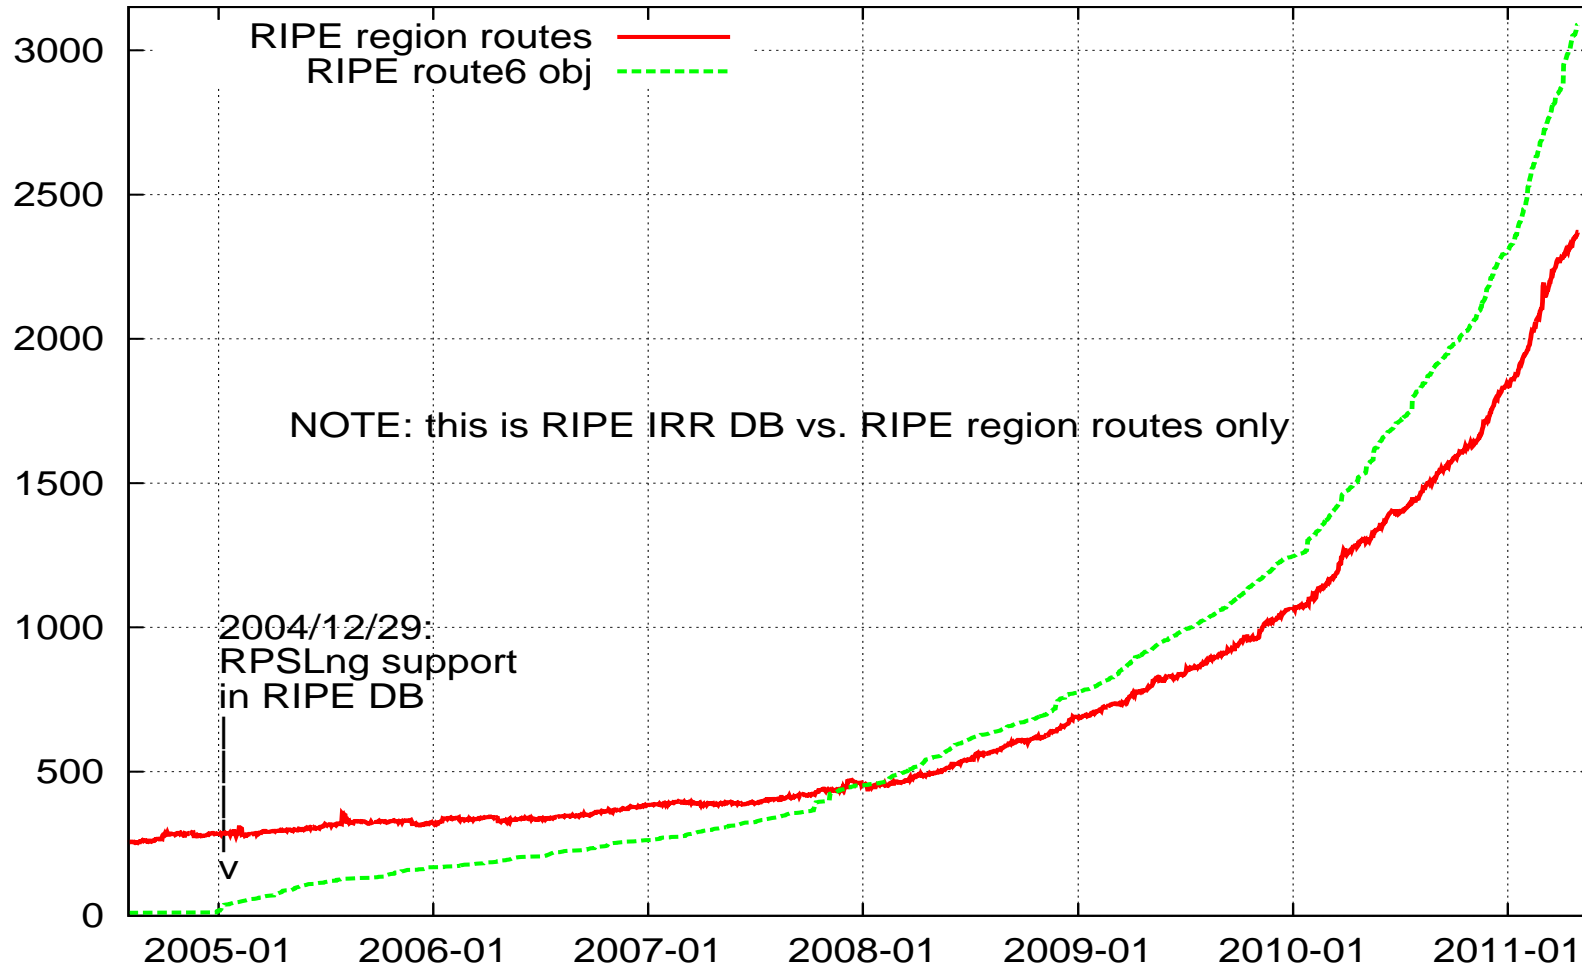
- and the winner is...

Network	Path
* 2001:1438:10:5::20/124	13237 8881 8881 196870 i
*>i	8881 196870 i
*> 2A02:28:4:2::/121	13237 8359 i
*> 2A02:28:4:2::1:0/127	13237 8359 i
*> 2A02:28:4:2::4:0/125	13237 8359 i

- I think the winner is AS196870 - long AS#, *and* long prefix!
- but seriously: whatever the outcome of the discussions about “how much deaggregation is OK?” - the usefulness of prefixes longer than a /48 in the global table is doubtful
- did I mention “filters on customer BGP sessions” already?

Route Registry - route6

Graphics: route6 objects vs. routes seen



route6 correlation (RIPE region)

- on 2011-04-30:
 - 2371 BGP routes from RIPE region
 - 3088 route6: objects in RIPE DB
- correlation?
 - multiple origin route6's (26 2002::/16, 18x 2001::/32, ...)
 - \Rightarrow 3001 route6 objects for *unique* prefixes
 - 182 route6 objects for prefixes from *other* RIRs...
- so...

route6 correlation (2)

- ... and this is what I found:

RIPE prefix, route6 ok	1933	:-)
RIPE prefix, route6 missing	331	!!!
RIPE prefix, route6 origin mismatch	93	
RIPE prefix, BGP inconsistant AS	4	
route6 objects without BGP route	805	???
other region, route6 ok	145	
other region, route6 missing	2986	
other region, route6 origin mismatch	28	
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.)

examples: route6 missing (bad!)

- just a few random examples - leaks? - but /32s as well!

2001:660:3005::/48	BGP: 6939 2483
2001:67C:28::/48	BGP: 13237 8954 3943
2001:67C:98::/48	BGP: 13237 15600
2001:688::/32	BGP: 6939 5511
2001:6B8::/32	BGP: 3257 5609
2001:728::/32	BGP: 2914
2001:808::/35	BGP: 13237 23456 8364 9112
2A00:1250::/32	BGP: 13030 48885
2A02:2230:100::/40	BGP: 1273 174 25467
2A02:2230:101::/48	BGP: 1273 174 25467
2A02:2230:1300::/40	BGP: 1273 174 25467

examples: route6, but no prefix

- maybe prefixes announced with limited scope only?

```
2001:648:2050:1000::/52  University of the Aegean Chios
2001:648:2050:2000::/52  University of the Aegean Samos
2001:648:2050:3000::/52  University of the Aegean Rhodes
2001:648:2050:4000::/52  University of the Aegean Syros
2001:648:2050:5000::/52  University of the Aegean Lemnos
2001:648:2050:6000::/52  University of the Aegean Athens
```

- these could be preparations for “soon to be announced”...?

```
2001:4C38::/32  AS15435,  changed: 20071003
2001:4D20::/32  AS39290,  changed: 20070315
2001:4D58::/32  AS25186,  changed: 20051028
2001:4D80::/32  AS5606,   changed: 20051018
2001:4D98::/32  AS12429,  changed: 20051026
2A00:2380::/25  AS2856,   changed: 20090916
```

examples: route6 origin mismatch (BAD!)

```
2001:8D0:2000::/35    BGP: 3320 1299 4589 20707
                      route6: origin: AS16210
2001:1AD0:8000::/33  BGP: 3257 24953
                      route6: origin: AS43556
2001:1B20::/32       BGP: 21385
                      route6: origin: AS8665
2A02:18::/32         BGP: 6939 5577 42652
                      route6: origin: AS44569
2A02:660::/32        BGP: 3320 5483 9125
                      route6: origin: AS29433
2A02:1668::/328     BGP: 6939
                      route6: origin: AS51088
2A02:24D8::/32       BGP: 13237 3356 21473
                      route6: origin: AS15945
```

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

route6 usage example

- it's as easy as this...

```
gert@moebius$ bgpq3 -l small-customer-in -6 AS5539
no ipv6 prefix-list small-customer-in
ipv6 prefix-list small-customer-in permit 2001:608::/32
```

```
gert@moebius$ bgpq3 -l large-customer-in -6 AS-SPACENET-V6
no ipv6 prefix-list large-customer-in
ipv6 prefix-list large-customer-in permit 2001:500:66::/48
ipv6 prefix-list large-customer-in permit 2001:608::/32
...
ipv6 prefix-list large-customer-in permit 2001:1488::/32
ipv6 prefix-list large-customer-in permit 2001:14d8::/32
...
```

(then copy-paste to router, or use in provisioning scripts)

- building filters is not hard (provided route6s are registered)
- <http://snar.spb.ru/prog/bgpq3/>

Questions?

gert@space.net

...and to answer the question from the title page:

NO, we are **not** growing fast enough!