

Are We Growing Fast Enough?

A snapshot of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich, Germany

Oct 5, 2009

RIPE 59, Lisbon

(not presented at all, just available online)

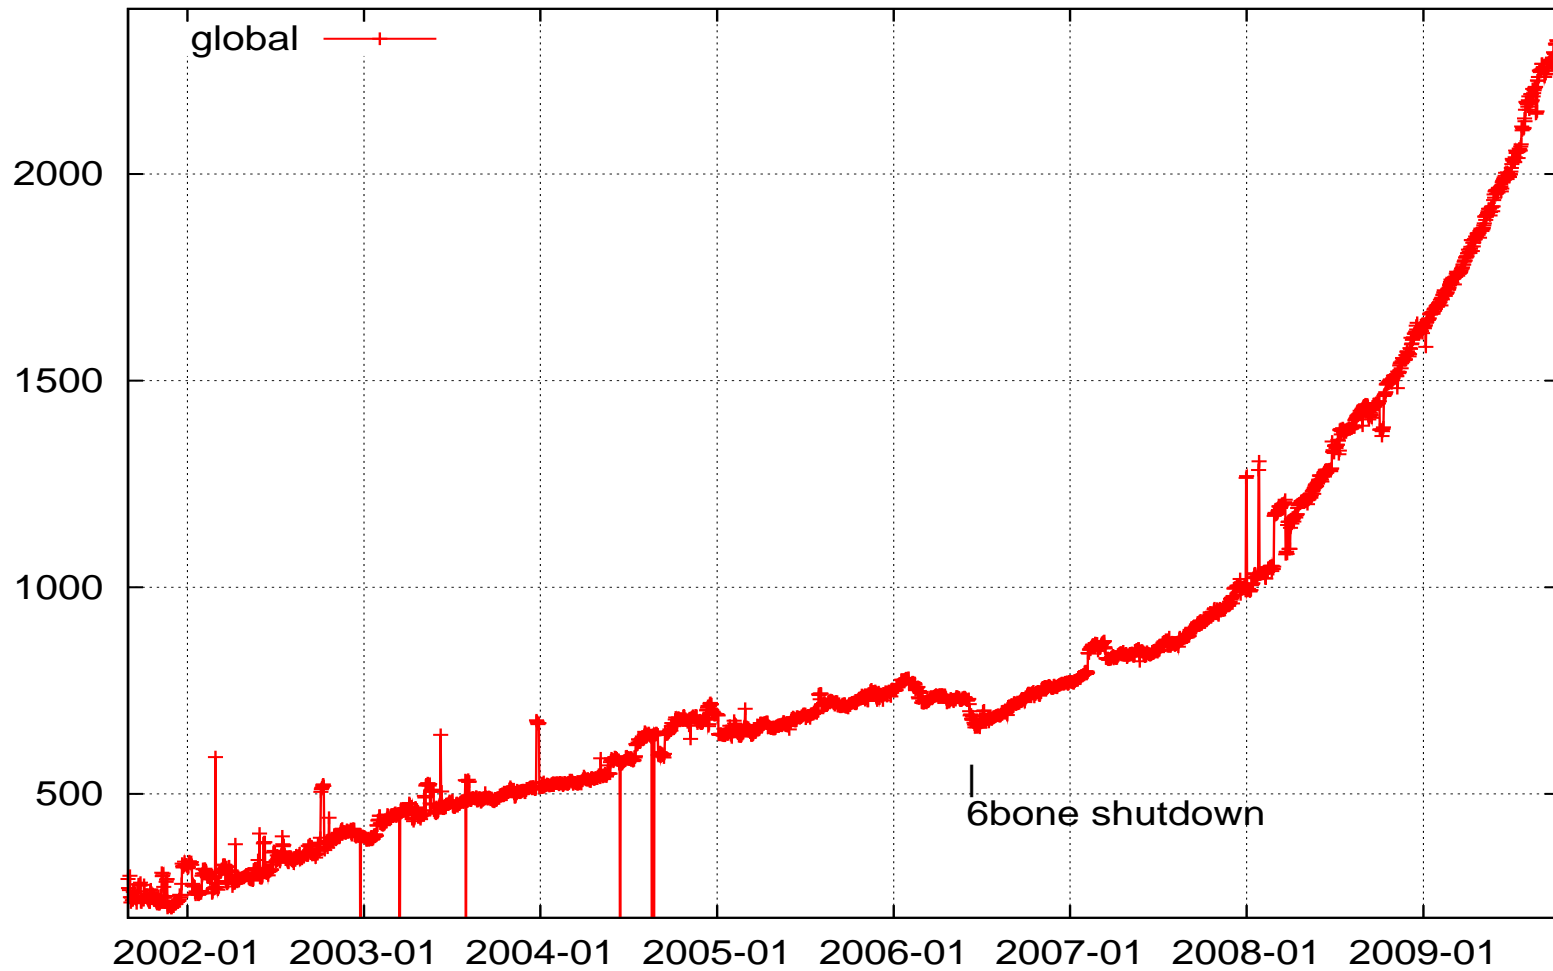
Overview

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


Slides online at: <http://www.space.net/~gert/RIPE/R59-v6-table/>

Prefixes in BGP Table


Graphics: Total Prefixes - 7 years




Graphics: zoom into last 6 months




Graphics: trends? (33 months)




Graphics: prefixes by RIR region



Graphics: prefixes by country (RIPE)



Graphics: prefixes by country (APNIC)




AS Numbers


Numbers - AS numbers

- as of 2009-10-04: 1674 unique AS#s visible (2009-05-02: 1382)
 - 1207 origin-only ASes (no transit paths seen) (970)
 - 419 ASes originate & give transit (368)
 - 48 transit-only ASes (e.g. 760, 1125, 1659, 6667, ...) (44)
- different number of prefixes announced
 - 1403 ASes originate 1 prefix (1157)
 - 131 ASes originate 2 prefixes (2 due to /32+/35)
 - 31 ASes originate 3 prefixes
 - 25 ASes originate 4 prefixes
 - 35 ASes with “more than that”, max. is 32 & 34 prefixes
- 2 ASes still announce their prefix as /32 and /35
- note: all paths observed from AS5539


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 2009-10-02: 2232 prefixes in total (2009-05-02: 1840)

/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	9	5	4	0	0	0	0
/24..27	10	5	4	1	0	0	0
/28..31	12	2	6	1	3	0	0
/32	1393	779	246	284	60	21	3
/33..34	28	13	9	6	0	0	0
/35	34	10	19	5	0	0	0
/36	25	4	15	6	0	0	0
/37..39	4	0	1	3	0	0	0
/40..41	57	8	13	30	1	2	3
/42..47	68	13	37	17	0	1	0
/48	557	99	152	263	16	27	0
/49..63	9	0	8	1	0	0	0
/64..128	23	3	13	6	1	0	0

Graphics - Prefixes / Size



Allocations


Numbers: RIRs, Allocations, ...

- On 2009-10-02, 2952 LIR blocks (PA space) allocated by RIRs:


RIR	alloc.	members	perc.	on 2008-10-12
ARIN	674	~ 3465	19.5%	487 (+38%)
APNIC	552	~ 3233	17.1%	430 (+28%)
RIPE	1546	~ 6428	24.1%	1052 (+50%)
LACNIC	121	~ 888	13.6%	101 (+20%)
AfriNIC	59	~ 350(?)	16.9%	52 (+13%)

- percentage with IPv6 *used* to be similar (~ 15%),
- much higher relative growth in RIPE and ARIN region(?!)
 - 1381 (R57: 949) allocations visible in routing table (*only 47%!*)

Graphics: Allocations over Time



Graphics: Allocated vs. Routed (LIRs)



Allocated vs. Routed - by region & class

RIR	type	alloc.	visible	perc.	subnets	
ARIN	LIR	677	285	42%	146	
	IXP	25	0	0%	0	
	Critical Inf.	88	34	39%	56	
	PI	251	60	24%	42	(*)
APNIC	LIR	554	247	45%	226	
	IXP	20	1	5%	2	
	PI	75	25	33%	26	(*)
RIPE	LIR	1548	785	51%	96	
	IXP	81	17	21%	0	
	Anycast DNS	17	9	53%	0	
	PI	28	15	54%	0	
LACNIC	LIR	122	47	39%	30	(NIR)
	Crit.Inf.+PI	12	3	25%	0	
AfriNIC	LIR	59	21	36%	27	
	PI	13	3	23%	0	

Numbers: some notable allocations

- US DoD got a LOT of space... (2008/05/06)

2608::/22

2608:4000::/22

2608:8000::/22

...

260f:c000::/22

260f:d000::/22

260f:f000::/22

- effectively this is “2608::/13”, but in disguise
- not really new, but has not been mentioned in this talk yet

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 8 years later, there *still* are weird things...

Eastern Asian Research

```

Network          Path
*> 2001:2B8::/32  3257 2497 4725 6939 17832 i
                  109 5511 10764 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
...
*> 2001:3C8::/32   109 5511 10764 11537 22388 7660 24475 4621 i
* 2001:3C8:100D::/48 109 5511 10764 11537 4621 4621 4621 4621 4621 ?
* 2001:3C8:1202::/48 109 5511 10764 11537 4621 4621 4621 4621 4621 ?
* 2001:3C8:1303::/48 109 5511 10764 11537 4621 4621 4621 4621 4621 ?
* 2001:3C8:9007::/48 109 5511 10764 11537 4621 4621 4621 4621 4621 ?
...

```

- Korean REN 1237 and 17579 has proper transit these days
- but leaks more-specifics via Internet2 (11537), NCSA and OpenTransit (unchanged, since over a year)
- Thai REN 4621 seems to have no “real” IPv6 transit at all
- long-standing problem: research networks with poor connectivity, due to political issues or plain disinterest

Egyptian Research

```
      Network          Path
* 2001:4300::/32      6939 6175 i
*>                      3257 6175 i
*> 2001:4300:2001::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
*> 2001:4300:2002::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
*> 2001:4300:2003::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
*> 2001:4300:2004::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
...
*> 2001:4300:2009::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
*> 2001:4300:2010::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
...
*> 2001:4300:2019::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
*> 2001:4300:2020::/48 13237 6453 6939 3491 4635 17579 11537 33789 24863 i
```

- 10764 = KREONET2, 11537 = I2
33789 = Egyptian MCIT Internet2
24863 = LINKdotNET AS number
- can you recognize a pattern?
- (this *might* be legitimate, but there is no IRR data to check)

Some Chinese Research

```
Network          Path
* > 2001:CC0::/32 1273 23911 18344 37944 i
                    6939 23911 18344 37944 i
* > 2001:CC0:2000::/36 1273 23911 18344 37944 i
* > 2001:CC0:3000::/36 1273 23911 18344 37944 i
* > 2001:CC0:4000::/36 1273 23911 18344 37944 i
* > 2001:CC0:5000::/36 1273 23911 18344 37944 i
...
* > 2001:CC0:F000::/36 1273 23911 18344 37944 i

* 2001:C68::/32      6939 23911 4134 i
* >                  1273 23911 4134 i
* > 2001:C68:100::/40 1273 23911 4134 i
* > 2001:C68:200::/40 1273 23911 4134 i
* > 2001:C68:500::/40 1273 23911 4134 i
...
```

- 37944 = “China Science and Technology Network”
- 4134 = “China Telecom”
- 23911 = “China Next Generation Internet” (CNGI-CERNET)

The Ghosts Are Back


- wrong prefix briefly announced by a typo...

Network	Path
*>i2024:5800::/32	3549 3561 7473 9911 i
*>i2024:5800:100::/40	3549 3561 7473 9911 i
*>i2024:5800:200::/40	3549 3561 7473 9911 i
*>i2024:5800:400::/40	3549 3561 7473 9911 i

- ... lingers as Ghost Route in the table forever.
 - announced only on 2009/05/20
 - still visible at 2009/10/04 (!)
 - BGP withdraw message not propagated to peers
- most likely buggy Cisco IOS at 3561 or 3549
- (CSCsu59917 and CSCsu03167, in 12.2(33)SXH3 + 12.4T)

Route Registry - route6

Graphics: route6 objects vs. routes seen



route6 correlation (RIPE region)

- on 2009-10-04:
 - 952 BGP routes from RIPE region
 - 1205 route6: objects in RIPE DB
- correlation?
 - multiple origin route6's (17x 2002::/16, 13x 2001::/32, ...)
 - ⇒ 1173 route6 objects for *unique* prefixes
 - 82 route6 objects for prefixes from *other* RIRs...
- so...

route6 correlation (2)

- ... and this is what I found:

RIPE prefix, route6 ok	747	:-)
RIPE prefix, route6 missing	185	!!!
RIPE prefix, route6 origin mismatch	15	
RIPE prefix, BGP inconsistant AS	5	
route6 objects without BGP route	329	???
other region, route6 ok	59	
other region, route6 missing	1213	
other region, route6 origin mismatch	23	
other region, BGP inconsistant AS	21	

- \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!

2001:610:118::/48	BGP: 13237 6453 6939 12759 41692 1140
2001:610:240::/42	BGP: 3257 3333
2001:678:D::/48	BGP: 21385 1853 49488
2001:688::/32	BGP: 1273 5511
2001:6B8::/32	BGP: 3257 5609
2001:6E0::/32	BGP: 3257 8954
2001:718::/32	BGP: 3257 2852
2A02:3C8:BAD::/48	BGP: 6939 43892
2A02:3C8:BAD:BEEF::/64	BGP: 6939 43892
2A02:3C8:BA51::/48	BGP: 6939 43892
2A02:3C8:DEAD::/48	BGP: 6939 43892

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”...?

```
2001:4C98::/32  AS35062,  changed: 20050526
2001:4CE8::/32  AS250,    changed: 20060316
2001:4D20::/32  AS39290,  changed: 20070315
2001:4D58::/32  AS25186,  changed: 20051028
2001:4D80::/32  AS5606,   changed: 20051018
2001:4D98::/32  AS12429,  changed: 20051026
```

examples: route6 origin mismatch (BAD!)

```
2001:7F8:24::/48      BGP: 13030 8758 20612 42476
                      route6: origin: AS20612
2001:1668::/32       BGP: 6939 25358 15830
                      route6: origin: AS12600
2001:1AD0:8000::/33  BGP: 3257 24953
                      route6: origin: AS43556
2001:1B20::/32       BGP: 21385
                      route6: origin: AS8665
2A01:7C0::/32        BGP: 553
                      route6: origin: AS12294, origin: AS8359 (2x!)
2A02:18::/32         BGP: 44042 42652
                      route6: origin: AS44569
2A02:D68::/32        BGP: 6939 12778
                      route6: origin: AS35471
```

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

Questions?

gert@space.net

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

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