

Are We Growing Fast Enough?

A snapshot of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich, Germany

Jan 31, 2011

NANOG 51, Miami

presented by CJ Aronson

Overview


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

Slides online at:


<http://www.space.net/~gert/RIPE/NANOG51-v6-table.pdf>

Prefixes in BGP Table


Graphics: Total Prefixes - 9.5 years




Graphics: zoom into last 6 months




Graphics: trends? (38 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 (AfrNIC)




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

AS Numbers


Numbers - AS numbers

- as of 2011-01-23: 2981 unique AS#s visible (2009-10-04: 1674)
 - 2214 origin-only ASes (no transit paths seen) (1207)
 - 706 ASes originate & give transit (419)
 - 61 transit-only ASes (e.g. 276, 567, 1998, 6667, ...) (48)
- different number of prefixes announced
 - 2457 ASes originate 1 prefix (1403)
 - 265 ASes originate 2 prefixes
 - 77 ASes originate 3 prefixes
 - 42 ASes originate 4 prefixes
 - 79 ASes with “more than that”, max. is 39 & 40 prefixes
- note: all paths observed from AS5539
- note: numbers skewed by ~36 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: 4210 prefixes in total (2009-10-02: 2232)


/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	14	7	7	0	0	0	0
/24..27	14	9	3	2	0	0	0
/28..31	30	9	8	7	5	0	1
/32	2482	1397	408	528	112	31	6
/33..34	77	36	20	13	8	0	0
/35	54	14	32	8	0	0	0
/36	78	24	36	16	2	0	0
/37..39	13	4	1	6	0	0	0
/40..41	118	19	10	85	3	2	0
/42..47	101	32	21	56	0	0	0
/48	1188	344	282	517	36	8	1
/49..63	11	3	3	5	0	0	0
/64..128	17	10	2	5	0	0	0

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	1137	487	43%	430	
	IXP	29	0	0%	0	
	Critical Inf.	122	46	38%	63	
	PI	466	129	28%	492	(*)
APNIC	LIR	1142	414	36%	334	
	IXP	20	1	5%	3	
	PI	248	53	31%	28	(*)
RIPE	LIR	2579	1410	55%	264	
	IXP	94	25	27%	0	
	Anycast DNS	28	19	68%	1	
	PI	273	159	58%	13	
LACNIC	LIR	122	39	32%	89	(NIR)
	Crit.Inf.+PI	74	23	32%	17	
AfriNIC	LIR	104	31	30%	5	
	PI	28	5	18%	0	

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
* D000::/8	109 30071 6762 28716 i
* i	3257 6762 28716 i
* i	1273 6762 28716 i

- in Table 2009/11/14-2010/01/15
- D000:* is completely outside IETF assigned unicast space
- upstream provider 6762 happily provides transit

More Fat Fingers

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

- in Table 2010/12/10-2011/01/23
- 2008:* is not allocated by IANA to any RIR yet (*bogon!*)
- upstream provider 6939 happily provides transit

Really Fat Fingers

Network	Path
*>i1A:F000::/24	3257 174 376 376 376 36786 i
* i	3257 174 376 376 376 36786 i
* i1000::/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

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 2011-01-25:
 - 1253 BGP routes from RIPE region
 - 2447 route6: objects in RIPE DB
- correlation?
 - multiple origin route6's (26 2002::/16, 15x 2001::/32, ...)
 - \Rightarrow 2376 route6 objects for *unique* prefixes
 - 108 route6 objects for prefixes from *other* RIRs...
- so...

route6 correlation (2)

- ... and this is what I found:

RIPE prefix, route6 ok	1569	:-)
RIPE prefix, route6 missing	271	!!!
RIPE prefix, route6 origin mismatch	68	
RIPE prefix, BGP inconsistant AS	2	
route6 objects without BGP route	631	???
other region, route6 ok	81	
other region, route6 missing	2192	
other region, route6 origin mismatch	27	
other region, BGP inconsistant AS	16	

- \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!

2001:600:1C0::/48	BGP: 12702 12234
2001:67C:28::/48	BGP: 13237 8954 3943
2001:67C:34::/48	BGP: 6939 39700
2001:688::/32	BGP: 1273 5511
2001:6B8::/32	BGP: 3257 5609
2001:728::/32	BGP: 2914
2001:7F0:1:34::/64	BGP: 13237
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
2001:4CE8::/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!)

```
2001:8D0:2000::/35    BGP: 3320 1299 4589 20707
                      route6: origin: AS16210
2001:930::/32        BGP: 6939 1299 20965 8517 8517 8517
                      route6: origin: AS8386
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
2A02:1668::/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!