

Impressions

An overview of the global IPv6 routing table

Gert Döring, SpaceNet AG, Munich

May 5th, 2004

RIPE 48, Amsterdam

Overview

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

Slides online at: <http://www.space.net/~gert/RIPE/R48-v6-table/>

Numbers - AS numbers

- as of 2004/05/03: 421 unique AS numbers visible (409)
 - 263 origin-only ASes (no transit paths seen) (232)
 - 144 ASes originate & give transit (156)
 - 14 transit-only ASes (e.g. 1717, 2153, 2603, 3301, ...) (21)
- mixture of RIR (2001::) and 6Bone (3FFE::) space announced
 - 258 ASes originate 1 RIR prefix (239)
 - 50 ASes originate 1 6Bone prefix (54)
 - 43 ASes originate 1 6Bone + 1 RIR prefix (42)
 - 27 ASes originate 2 RIR prefixes (9 due to /32+/35)
 - 29 ASes with “more than that”, maximum is 6 prefixes
- 12 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?

2001:360::/32 1221 i

2001:360:4001::/48 1221 ?

2001:360:4002::/48 1221 ?

- 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 2004/05/03: 546 prefixes in total (2004/01/24: 524)

/n	global	RIR space	6bone	6to4	(2004/01/24)
/16	1	0	0	1	(1 0 0 1)
/24	39	0	39	0	(42 0 42 0)
/27	1	1	0	0	(1 1 0 0)
/28	40	0	40	0	(41 0 41 0)
/32	350	319	31	0	(324 293 31 0)
/33	2	2	0	0	(3 3 0 0)
/34	0	0	0	0	(0 0 0 0)
/35	42	42	0	0	(45 43 2 0)
/36-/39	1	1	0	0	(0 0 0 0)
/40	5	4	1	0	(3 3 0 0)
/41-/45	5	5	0	0	(3 3 0 0)
/48	59	37	22	0	(57 35 22 0)
/52-/60	0	0	0	0	(0 0 0 0)
/64	1	0	1	0	(4 1 3 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	2001:608:0:3::9	3320 1752 i
*		2001:948:0:F00F::1	2603 1741 i
*		3FFE:C00:8023:19::1	109 i
*		3FFE:8150:0:1::17	9044 559 i
*		2001:608:0:3::7	1930 i
*		3FFE:1108:40A:FFFF::1:2	3274 1741 i
*		2001:7F8:2:8001::2	1752 i
*		2001:7F8::3349:0:1	13129 9033 i
*		2001:7F8::CB9:0:1	3257 12859 ?
*		2001:450:1:2001::AA	3549 20965 559 i

- this is **fine**, anycast relay approach, see RFC3068
- this is just a snap-shot from AS5539, not all relays visible
- 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 - 32 months



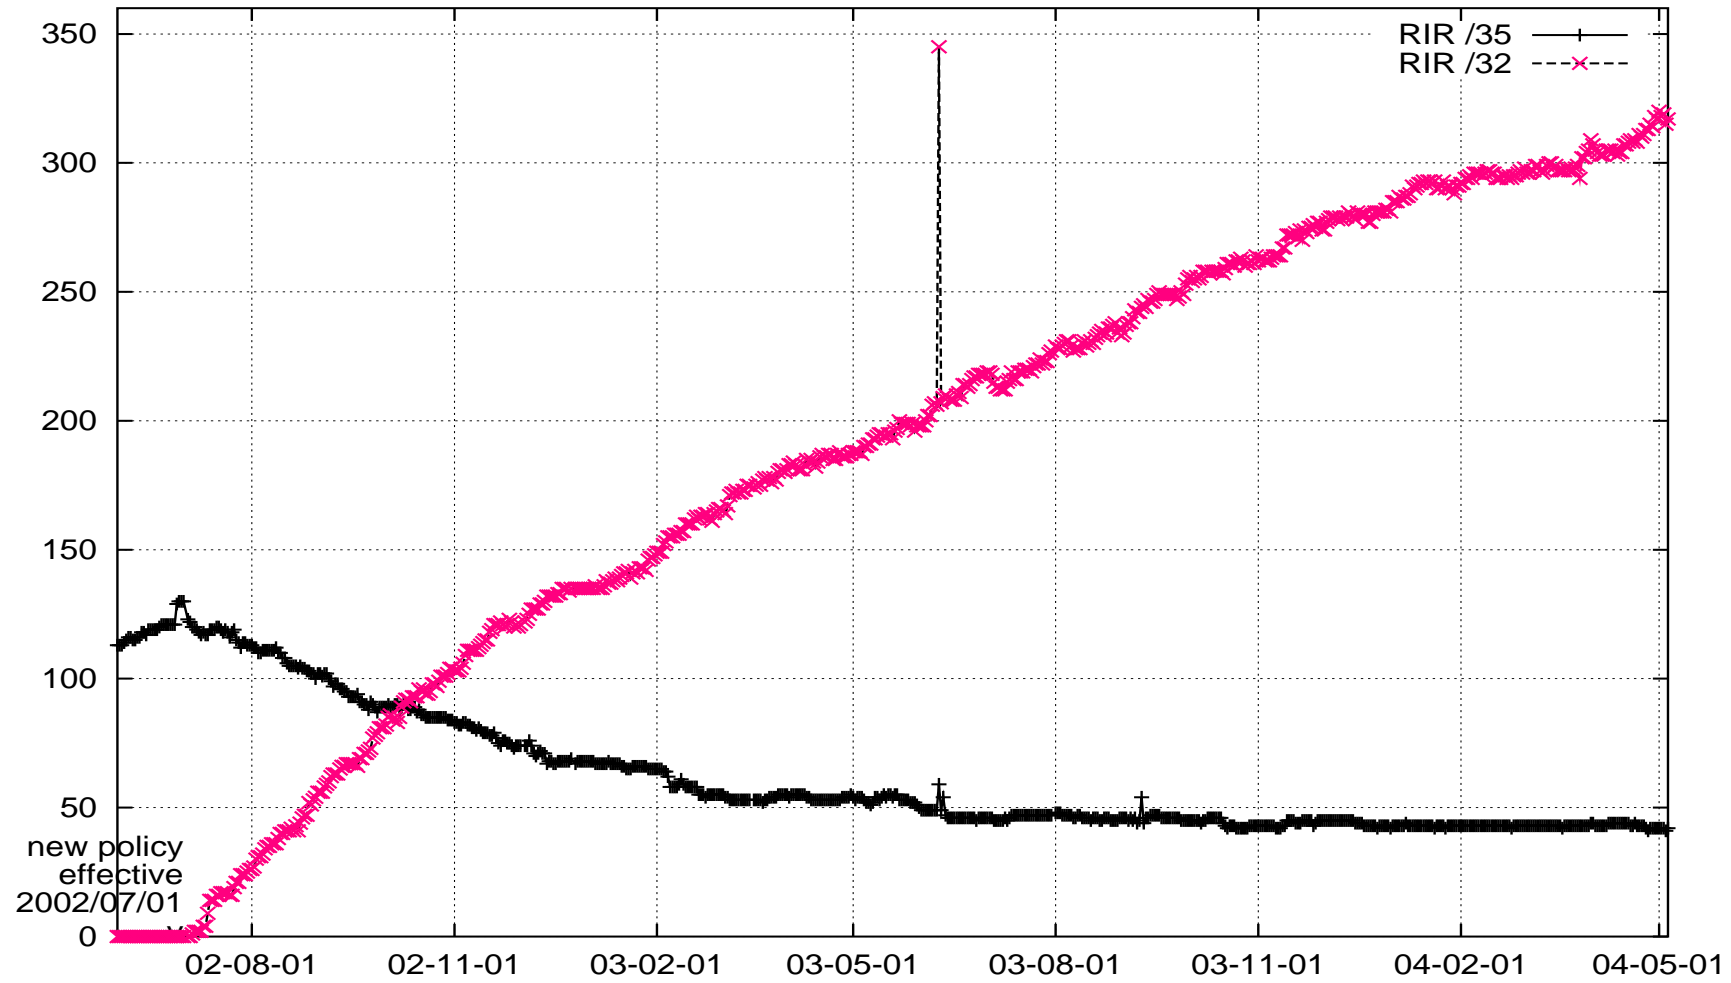
Graphics: RIR vs. 6Bone Prefixes - 32 months



Graphics: RIR vs. 6Bone Prefixes - 4 months



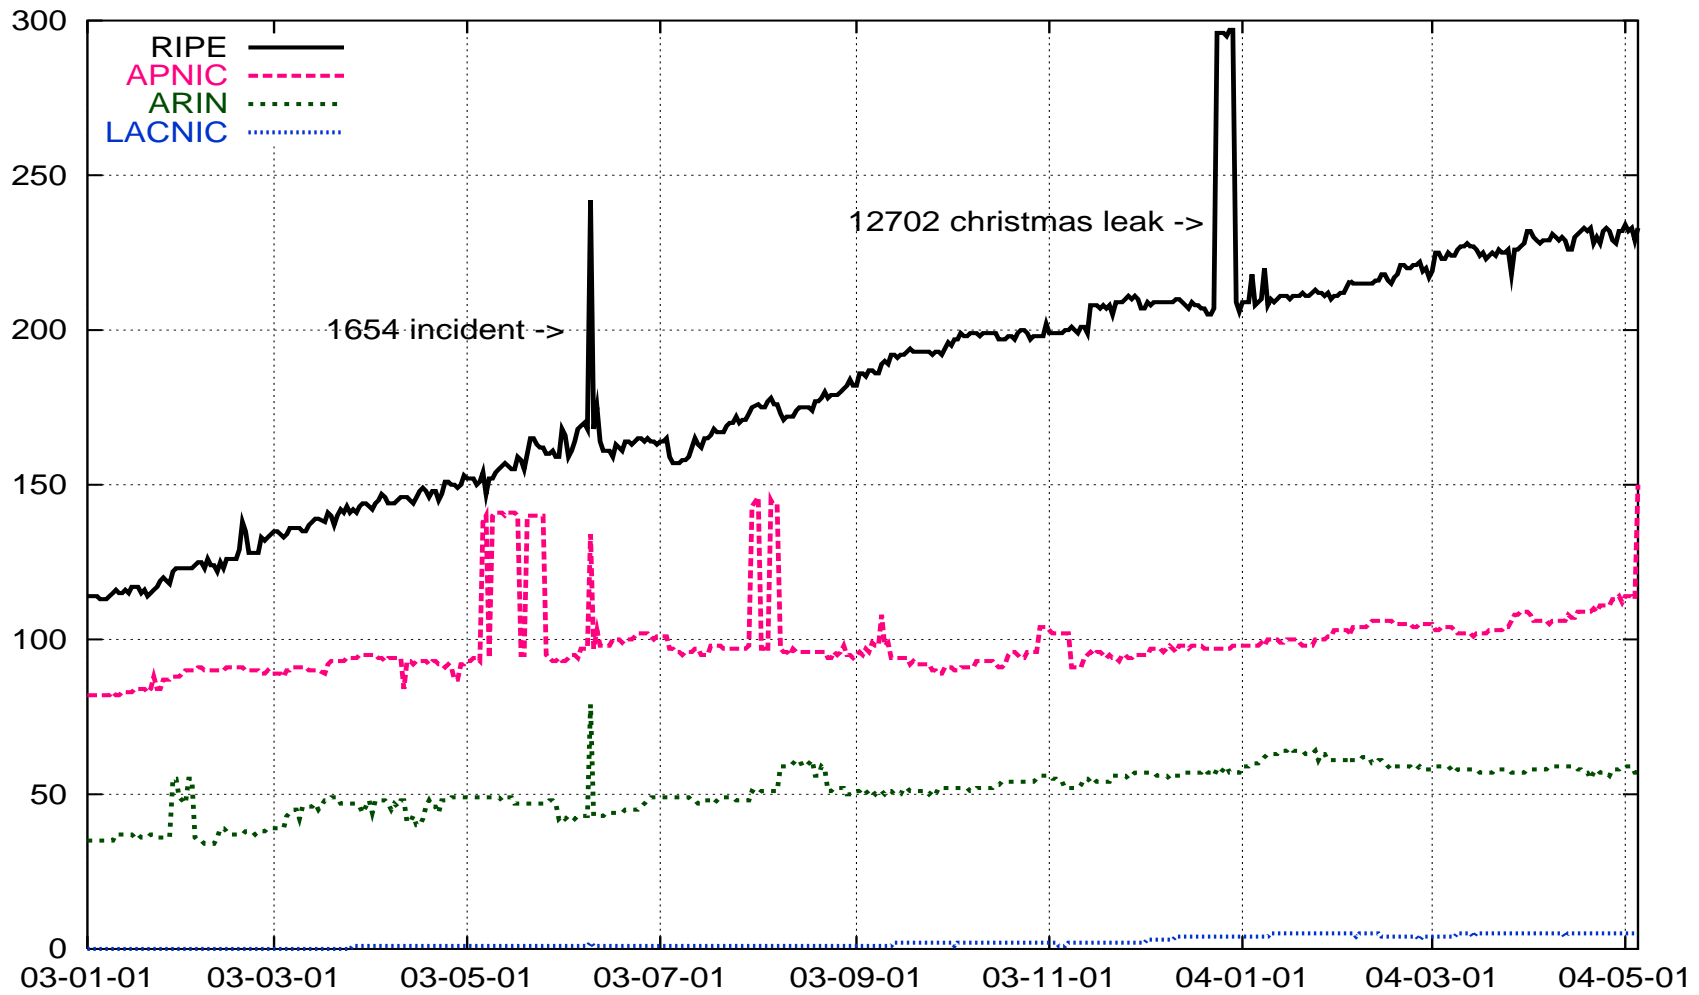
Graphics: RIR /35s vs. /32s




Numbers: RIRs, Allocations, ...

- 595 LIR blocks out of 2001::/16 allocated by RIRs:
ARIN 107(95), APNIC 152(132), RIPE 330(286), LACNIC 6(6)
as of 2004/05/03 (2004/01/27: 519)
- plus some IXP and other microallocations
- B, F, H, I, K, M Root-Servers have IPv6 addresses (some are visible on www.root-servers.org, some are already in BGP)
- 344(319) allocations visible
- allocations take up 353(334) routes: 38(43) /35s, 315(291) /32s
14(15) allocations visible as /32 and /35
- 2001:0D70::/30 allocated to NTTWEST-IPv6 (2003/09)
- RIPE NCC has /23 and /20 requests in queue

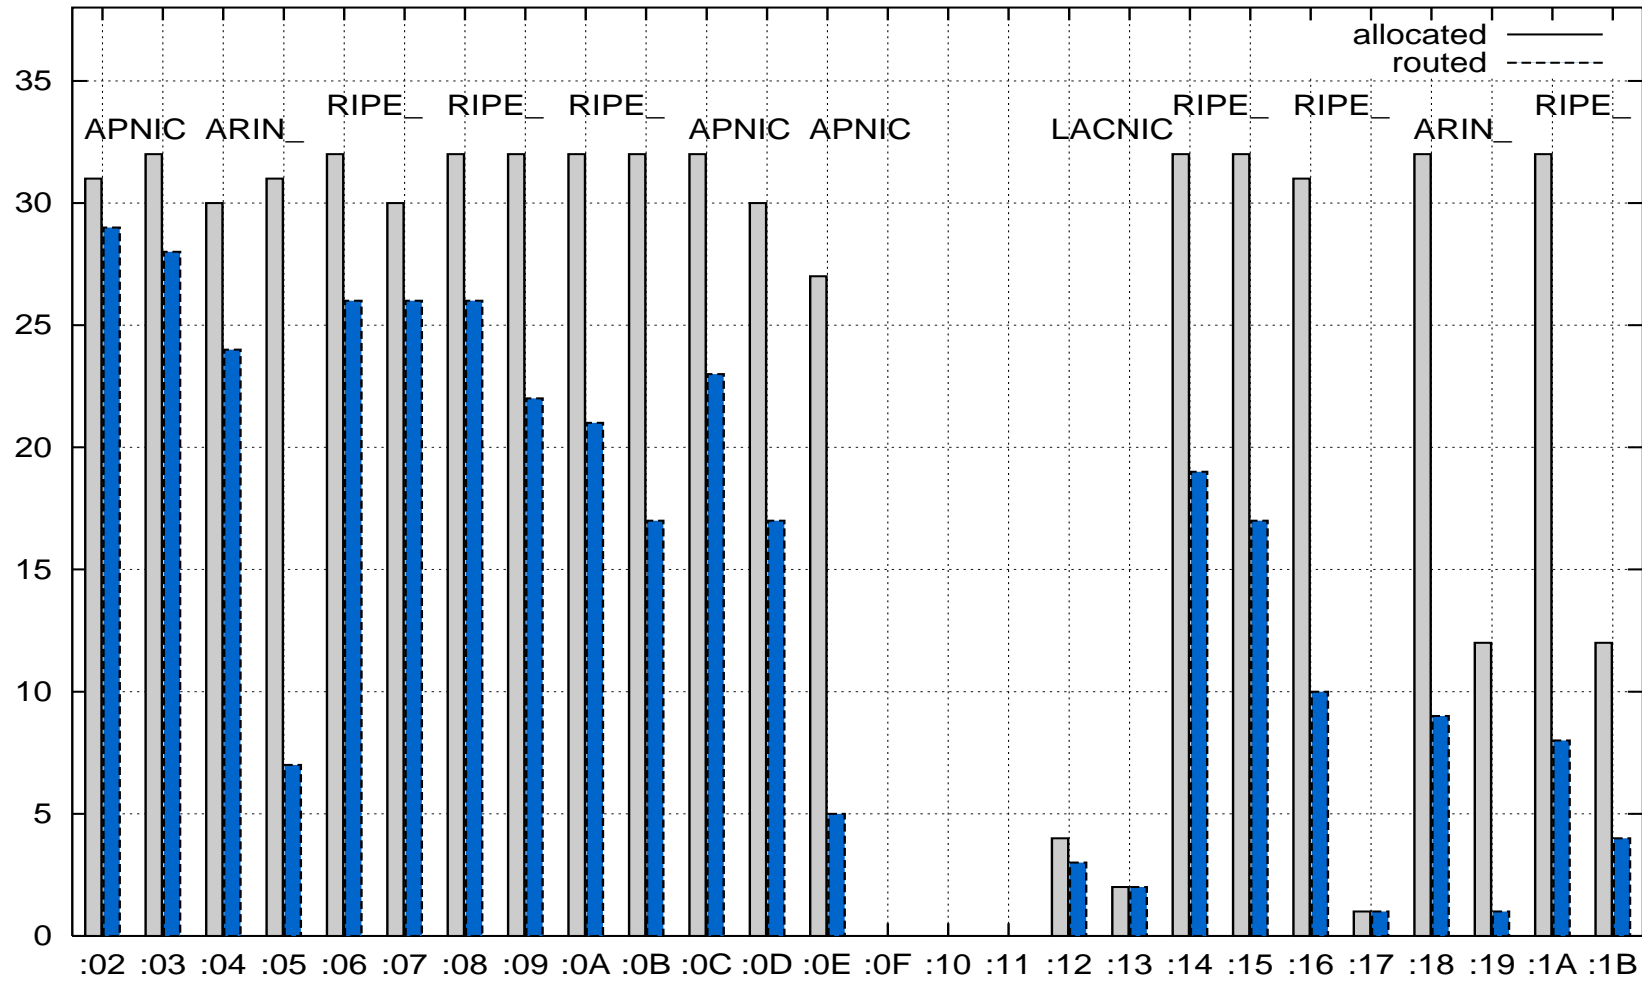
Graphics: prefixes by RIR region



Graphics: prefixes by country (RIPE)



Graphics: Allocated vs. Routed



Interesting Observations (1) - Accidental Hijack

Network	Next Hop	Path
* 2001:6c8::/32	3ffe:8150::1	9044 10566 29657 ?
*> 2001:6c8::/32	> 3ffe:401c:0:3:20c:ceff:fe05:da0e	29657 ?
* 2001:6c8::/32	2001:468:ff:121d::2	11537 3425 293 6175 3292 i
* 2001:6c8::/32	2001:770:8::	1213 11537 3425 293 6175 3292 i
* 2001:6c8::/32	2001:1548:1:10::4	12565 5609 4555 6830 3292 i
* 2001:6c8::/32	2001:15a8:1:1::6	29449 3320 6830 3292 i
* 2001:6c8::/32	2001:15f8:1::1	25384 1752 8472 6830 3292 i
* 2001:6c8::/32	3ffe:4017::1	24776 25358 3549 6830 3292 i

- correct origin AS is 3292
- caused by static route and static → BGP redistribution
- fixed in less than 20 hrs - thanks!

Interesting Observations (2) - Martian Leaks

Network	Next Hop	Path
*> ::/0	2001:7F8:1::A500:1103:1	1103 i
...		
*> 1000::/8	2001:610:25:5062::62	1103 i
* 1000::/8	2001:668:0:1:34:49:6900:40	3257 1103 i
* 1000::/8	2001:728:0:1000::f000	33 2914 6680 1103 i
* 1000::/8	2001:608:0:fff::6	399 5539 3257 1103 i
* 1000::/8	2001:780:0:2::6	12337 3257 1103 i
...		
* 1000:2550:100::/40	2001:728:0:1000::f000	33 2914 6680 1103 i
* 1000:2550:200::/40	2001:1418:1:400::1	12779 3549 6939 3257 1103 i
* 1000:2550:300::/40	2001:15a8:1:1::6	29449 6939 3257 1103 i
* 1000:2550:400::/40	2001:608:0:fff::6	399 5539 3257 1103 i

- likely to be caused by some Cisco/6PE weirdness
- noticed and fixed in very short time (thanks!)
- only documented martian leak since 2002/10/21 :-)
- but: displays potential for improving BGP prefix filters...

Observations (3) - weird AS path leaks

Network	Path
2001:a40::/32	12779 6175 145 7580 10566 13944 25396 15444 29678 29678 29158 25489 25484 12731 5539 1930 9186
12931	8472 8903 16091 513 9044 10566 13944 25396 15444 29678 29678 29158 25489 25484 12731 5539 1930 9186
7660	2500 4697 10566 13944 25396 15444 29678 29678 29158 25489 25484 12731 5539 1930 9186
11537	7660 2500 4697 10566 13944 25396 15444 29678 29678 29158 25489 25484 12731 5539 1930 9186
1103	11537 7660 2500 4697 10566 13944 25396 15444 29678 29678 29158 25489 25484 12731 5539 1930 9186

- Ghost Route Buster flagged this as a ghost route coming through our AS 5539 (*gasp*), but actually it is quasi-legit.
- real problem is unlimited prefix distribution
(5539 → 12731 →)25484 → 25489 → 29158 → 29678...
- i.e.: leaf node ASes offering full BGP feed to both upstreams and (*worse*) upstreams accepting all of it

Observations (4) - 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	2001:608:0:3::7	1752 5511 4716 7668 i
* 2001:608:0:3::7	3FFE:1108:40A:FFFF::1:2	1930 20965 11537 7660 4716 7668 i
* 3FFE:1108:40A:FFFF::1:2	3FFE:C00:8023:19::1	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	3FFE:1108:40A:FFFF::1:2	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	::FFFF:203.14.5.73	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?
 ⇒ this seems to have been a hiccup around 8472 (Viag/BT DE)

Interesting Observations (5) - Invalid AS numbers

Network	Next Hop	Path
* 2001:568::/32		
	3FFE:1108:40A:FFFF::1:2	3274 790 790 6830 6175 6435 6342 64600 4787 64702 29155 3263 852 e
*>i	2001:650:F807::DE9:1	3561 145 6435 6342 64600 4787 64702 29155 3263 852 e
*	2001:608:0:3::7	1930 20965 3549 2500 4725 6175 6435 6342 64600 4787 64702 29155 3263 852 e
* 2001:D10::/32		
*	2001:478:FFFF::1	4555 13944 6435 6342 64600 4787 i
*	2001:7F8:2:8001::2	1752 4725 6175 10566 3758 4787 i
*>	3FFE:8150:0:1::17	9044 10566 3758 4787 i
* 2001:D10::/32		
*	2001:478:FFFF::1	4555 13944 25396 64702 4787 i
*>	2001:7F8:2:8001::2	1752 15444 25396 64702 4787 i
*	3FFE:8150:0:1::17	9044 5424 5609 4555 13944 25396 64702 4787 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
- private/unallocated AS numbers reappearing ?!
- ghost routes are reappearing ?!
- are “early” IPv6 networks already deteriorating?
- still quite a number of “unsolicited full transit” links
- but: 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/R48-v6-table/>

Questions?

gert@space.net