



RIPE Address Policy Working Group

Nov 3, 2011

RIPE 63, Vienna

WG Chairs: Gert Döring, & Sander Steffann

please remember: this session is webcast

APWG Agenda - Thursday

- R. on IPv6 PI/PA Unification - Gert Döring, WG Chair
 - recap on the genesis of the current IPv6 PA/PI policies
 - proposal how to re-do the IPv6 allocation framework
 - discussion!
 - (please read the concept document sent to the AP WG list)
- Y. Open Policy Hour
 - “The Open Policy Hour (OPH) is a showcase for your policy ideas. If you have a policy proposal you’d like to debut, prior to formally submitting it, here is your opportunity.”
- Z. AOB

Let's enter the discussions

- No decisions are made here(!). This is to get feedback to the proposers and to get a feel for the Working Group's opinions.
- Consensus based process based on the *open* mailing list.
- please remember to speak into the microphone
- please speak your name, so the scribes can properly attribute what you said
- the session is webcast, so people that couldn't come to Vienna can still be participate
- remote feedback can be provided by webchat

returning to IPv6 PI discussion

- **this is only about IPv6**
- IPv4 is different, and we take this into account
- looking into the future

why is there a difference between PA and PI?

- in the end, it's "just some numbers" given out by the RIPE NCC to "consumers" of these numbers
- difference comes from intended use:
- PA
 - intended to aggregate (A) thousands or millions of end users into a single block, single routing table slot
 - assumed that "ISP" would be RIPE member anyway
 - liberal sizing, no strings attached
- PI
 - intended for a single independent (I) end-user network
 - not intended as "cheap replacement for RIPE membership"
 - specific purpose (BGP multihoming) \Rightarrow strings attached

history of IPv6 allocation/assignment


- initial IETF model was very strict on aggregation
 - “Top Level Aggregator” ISPs get a /13
 - default-free zone hard bounded to 8192 routes
 - question “who is worthy?” not answerable \Rightarrow abandoned
- initial RIR IPv6 policy (1999) gave LIRs a /35 (minimum)
 - avoided TLA problem, but a bit on the small side
 - changed to /32 in 2002
 - still strong focus on aggregation \Rightarrow no PI
- since then, detail tuning of policy for allocation to LIRs
 - HD ratio and end user assignment size adjusted
 - removal of the requirements to announce as an aggregate (only) from the policy (2009-06), deferring to routing WG

history of IPv6 allocation/assignment (2)


- IPv6 PI proposal introduced by Jordi Palet in 2006-01
 - *strong* resistance from the “aggregation!” camp
 - experience with IPv4 PI caused quite some opposition
 - argument that finally got accepted: multihoming proposals from IETF are not going anywhere (in reasonable time), and solution needed for *enterprise end-users* that want to do *BGP-based multihoming with IPv6*
 - proposal accepted in April 2009
 - lots of strings attached (multihoming, no sub-assignment)
- over time, emphasis shifted from “maximum aggregation” to “find workable compromise, encourage *use* of IPv6”



IPv6 PA and PI given out by RIRs

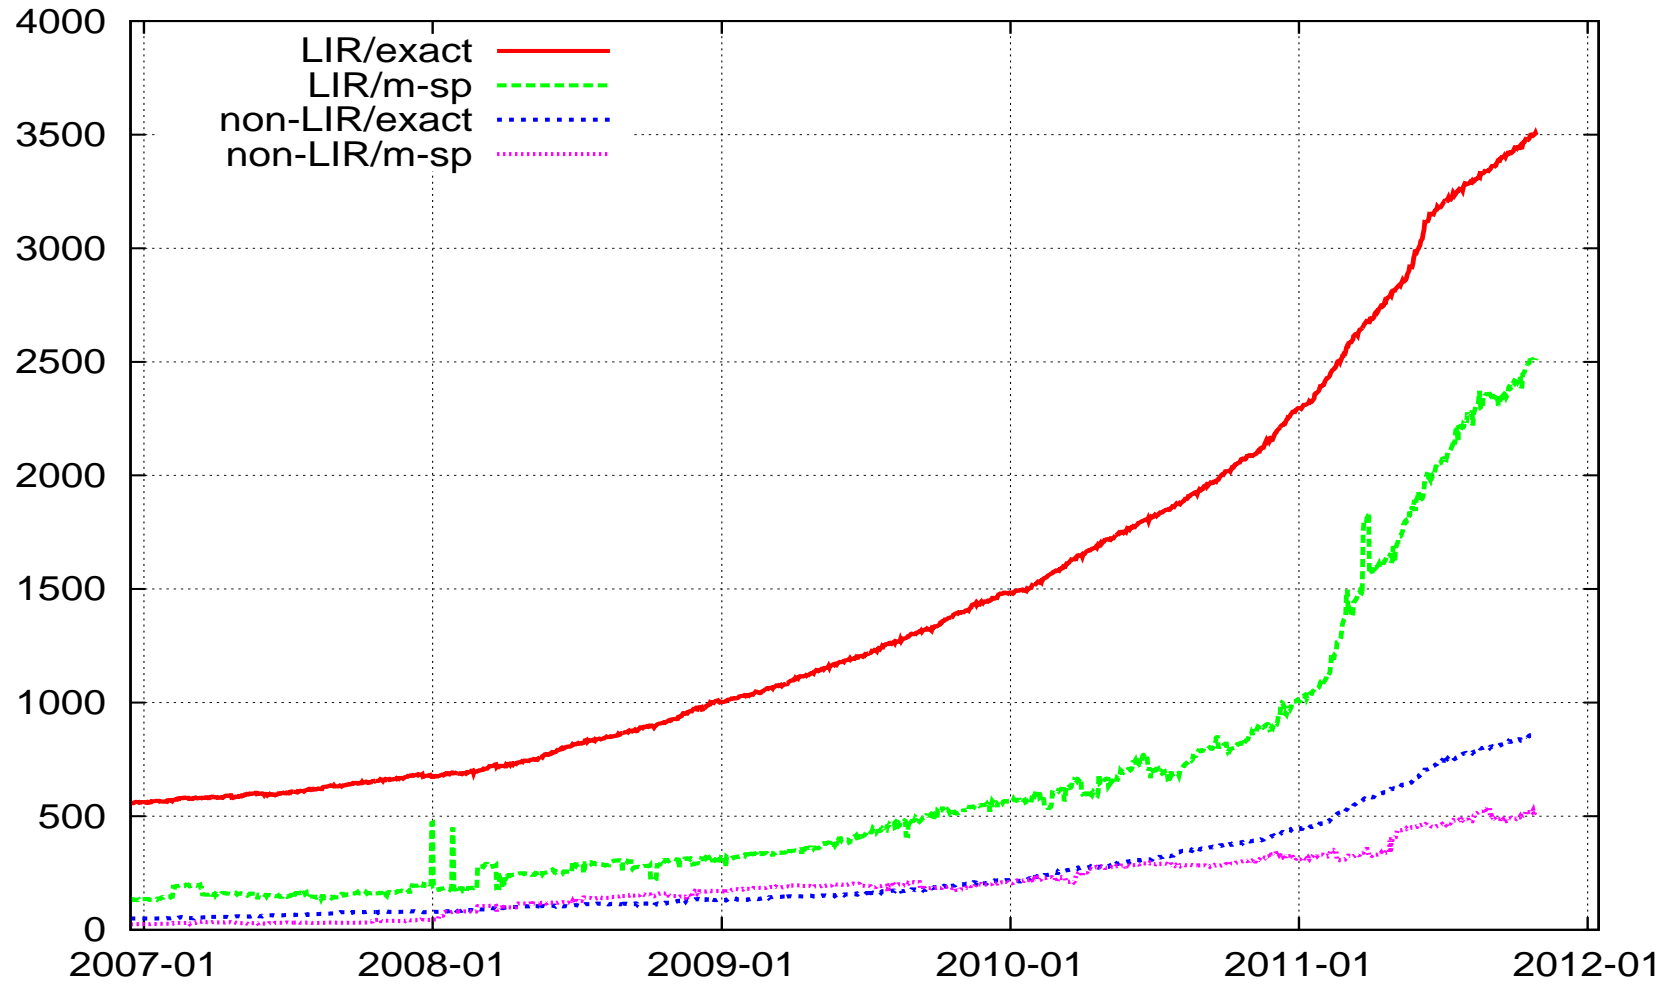


prefix size in IPv6 routing table





IPv6 routing table by prefix class



pillars of IP address management

- registration
 - clear documentation who “owns” a certain number
 - goes along with verifiability in the routing system
- aggregation
 - keeping the routing table size bounded
 - trying to balance business needs/wants and global cost
- conservation
 - making sure we don’t run out of addresses
 - for v6, we can be more liberal, but still finite resource

address policy needs to balance...

- routing table
 - 1 million routes will break it for everybody
- NCC costs
 - we need the NCC to have a solid financial basis
- end user costs
 - too expensive RIR cost will lead to creative workarounds
- usefulness
 - address space acquired must be useful for the purpose
- address space efficiency
- good stewardship: encourage /48.../64 to end users

ongoing discussions about PI

- multihoming requirements
 - what *is* multihoming and how to prove it?
 - 2011-02 aims to remove this requirement
- costs
 - PI seen as “cheaper way to number my ISP business”
 - PI isn’t meant as such, and that causes frictions
- usage restrictions (no sub-assignment)
 - “why can’t I number my datacenter customers from my PI?”
 - some proposals in the discussions, nothing tangible yet
- is this detailed fine-tuning of PI policies the right approach?

more radical approach

- abandon distinction between PA and PI completely
- RIPE members (LIRs) go to RIPE NCC and ask for “numbers”
- numbers are then used to “number things”
- sounds easy...?
- the tricky bit is to get the details right
- draft at <http://www.ripe.net/ripe/mail/archives/address-policy-wg/2011-October/006496.html>

1. WHO gets address space?

- keep LIR (= RIPE member) and “sponsoring LIR” model
- all number blocks go from the RIPE NCC to a LIR
- then
 - either LIR uses it “for its own network”
 - or LIR passes on to customer that has signed appropriate contracts (keeping the requirements of 2007-01)

2. HOW BIG should a single block be?

- /48 “by default”
- larger than /48 for “end-sites with large networks”, if justified
- /32 (.../29) when planning to assign /48.../64 to 3rd parties
- larger than that: when documented need (as now)
- automatic consequence: “multiple blocks of numbers” to a single LIR will have to be accepted as “frequently seen usage case” (we’ll come back to this)

3. SPECIAL CASE networks

- currently: special case PI exists for IXP, Root DNS, Anycast DNS
- proposed implementation:
 - checkbox on the template “I want to use that for IXP/Root/Anycast DNS”
 - not used for evaluation (this is just numbers!) but used for selection of the address range to pick numbers from
 - people want to have the option to treat “special numbers” differently in their routing policy, and that’s easier if they are easily recognized
- also proposed to have a well-documented range for /33.../48 number blocks (smaller “minimum assignment size” range)

4. COSTS

- we don't decide costs
- but we can send recommendations to the AGM (and then vote)
- one possible model could be:
 - yearly base fee for LIR
 - per-piece yearly fee for each number block held
 - /48 = 50 EUR, /32 = 100 EUR, /31+ = 200 EUR /year?
 - (per-piece installation fee for each new number block?)

5. MULTIPLE BLOCKS per LIR

- “more than one block for a single LIR” is not possible today for PA (unless full), but would be needed for new “number blocks”
- “get any number of blocks you ask for”
 - not likely to get consensus
- proposal for compromise:
one “block of numbers” per “network”
- definition of “network”?
 - interconnected nodes
 - operated by same entity
 - operated with a common routing policy
 - operated as a *layer 3* network
- goal is to be *reasonably* flexible here

problem cases regarding multiple blocks

- ISPs providing L3 services on top of other ISPs' Layer 2 infrastructure
- single LIRs providing addresses to two independent Layer 3 networks that are not directly connected - e.g. due to political (commercial/NREN), legal or geographical reasons
- “classical PI” type connections - end users with independent numbers, having multiple non-connected end sites
- multiple L3 providers providing address space to a single user must be allowed (multi-homing, special-purpose networks, etc.)

what do YOU think?

- feedback from the room, please
- next steps: take feedback, form policy text, propose to the list

Y. Open Policy Hour

- Y.1 Kurtis Lindqvist - on AS numbers, anycast, and the IETF
- Z. AOB - anything else?

Thanks!

- thanks for your input
- thanks for your help in forming policies in the RIPE region
- ...enjoy your coffee break!
- ...and see you on the mailing list...