

## Preamble

At RIPE83, the new APWG chairs suggested an overhaul of the IPv6 address policy - both policy text and policy itself, if needed. Some WG members volunteered to look into the part “why is the current policy the way it is? Are the fundamental motivations still valid?”... and this is a starting document for further discussions in the WG.

This document is not a formal WG document, nor a product of a formal task force, but a personal view of the authors.

## Authors

Gert Döring, Kurt Kayser, Sander Steffann

## Scope

We have focussed exclusively on IPv6 policy. This does not mean that policies about other resources, such as ASNs, won't benefit from a closer look. But that is out of scope for *this* document.

## Underlying assumptions and motivations

The following underlying assumptions, goals and drivers went into the IPv6 policy, as it is now - not commenting on the merits of either point here. Discussion follows, grouped in a somewhat topical way, in the next sessions.

- It should be easy to get IPv6 address space.
- RIPE IPv6 address policy should encourage IPv6 rollout, not get in the way.
- Aggregation is very important, both inside an ISP network and for the global routing system
- Conservation of IPv6 address space is a desirable criteria, but at the same time it is acknowledged that the IPv6 address space is very large, so allocation and assignments can be more liberal (“wastive”) than for IPv4
- Implementation of allocation and assignment policies should be easy, and require less paperwork and bureaucracy than for IPv4
- Use of N:1 NAT (ISPs assigning single IPv6 addresses to customers) is seen as undesirable
- “Special Networks” exist that need stable IPv6 addresses which must not be tied to any particular uplink provider - namely, Anycast root DNS servers and IXP exchange fabrics (some need to be seen in the global routing system, others just need to be unique)
- “Renumbering networks is hard”

- “Multihoming needs to be done with BGP” (due to lack of widely accepted IETF alternatives)
- IETF failed to deliver on “new” multihoming solutions, so a certain - limited - number of companies exist that need address space not tied to a specific upstream provider, but can be announced over multiple providers or taken from upstream A to upstream B when changing contracts.

## Evaluation

Looking into each of these - partially conflicting - goals and assumptions in more detail:

### “Simplicity, encouragement, aggregation, discourage N:1 NAT”

The end result of forming policy for these goals is:

- ISPs can give end users a whole network block, up to a /48 per end site, without needing justifications (generally a /56 is recommended for SoHo end users, which is still 256 subnets of IETF-recommended size /64) - and most ISPs that provide IPv6 connectivity to end users seem to assign a /56, so the policy works.
- Unlike IPv4, end customers are not encouraged to use NAT to work around address shortage, and also do not need to come back regularly because they need more addresses due to network growth.
- RIPE LIRs can get a /29 address block from the RIPE NCC without needing any justification, so receive sufficient address space for 0.5 Million /48 assignments, or about 134 Million /56 assignments. Less, if internal aggregation is needed, but still sufficient for all but the largest service providers.
- Most RIPE LIRs will not come back to the RIPE NCC to ask for more address space for a long time - or not at all.
- If a RIPE LIR needs more than a /29, the need for internal network aggregation is taken into account when judging allocation size based on number of /56s needed (numerically, the algorithm used is “HD ratio”)

**Most general assumptions leading there are still valid and the existing IPv6 policy seems to reasonably balance these needs. Especially the “easy to operate, little bureaucracy, encourage IPv6 usage” part is still important.**

Items under discussion:

- 2.7 Utilization & 2.8 HD-Ratio are seen as “too complicated”
  - the wording (“bits to the left of the efficiency measurement unit”) could be less scientific, and easier to understand
  - it does take into account that “larger networks with multiple levels of aggregations” can not be as efficient as “smaller networks”, so there is a mathematical justification for doing so, and not just asking for “10% utilization” - but are we overdoing it, for the common cases?

- 5.1.2 for “allocations larger than a /29”, and 5.2.1 “subsequent allocation criteria” - it is (sometimes) very hard to bring the necessary documentation
  - especially the bump “no documentation for /29” to “full documentation for /28” is steep(!).
  - Some cases are easy, like “traditional ISP networks”, where a LIR can just count customers, POPs, etc., and document their network aggregation hierarchy.
  - For other large networks, more situated in the enterprise or government arena, it is considered a huge effort to prepare this detailed level of documentation.
  - increasing the general allocation size (“a /24 for everybody”) would not solve the problem for very large networks, and needlessly use more address space for “smaller LIRs”
- 5.7 “Existing IPv6 address holders” is seen as contradicting 5.2.1 - does a LIR need to bring documentation for an extension of its address allocation, or not? (Clarification in 5.7 that this refers to LIRs having been allocated a smaller-than-usual network, i.e. a /32 or /35, should help reducing the confusion)
- 4.2 Routability is not guaranteed, but there is little guidance on routing aggregation in the address policy (currently 3.4: “*should be distributed in a hierarchical manner*”) - this is a situation outside the direct responsibility of address policy, but if we could give advice here, or point to a BCOP document, this would be helpful for newcomers that are looking for guidance on commonly accepted levels of deaggregation, vs. what could be called “network vandalism”.
- Unbounded deaggregation of LIR PA blocks (/29s) into “more specifics announced by other entities” might create lots of demand for AS numbers as well.
- 4.4 “Consideration of IPv4 Infrastructure” - is this still a relevant criteria? Does it create an unfair situation for “old networks that have lots of IPv4” vs. “new networks that have great plans, but no v4 to back their numbers”? Is it just moot and should go?

### “Special Networks”

The needs of IXP fabrics (unique, not tied to any particular ISP, possibly not even routed) and anycast root DNS servers (unique, tied to root DNS instance and not to any particular operator) have not changed.

With the general availability of IPv6 PI address space, it is unclear whether these special-case blocks are still needed or whether “plain” IPv6 PI could serve the needs.

### “Renumbering, Multihoming, Provider-Independence”

For non-trivial networks, the whole aspects of “renumbering on ISP change” and also “multihoming to multiple different providers” still does not have good solutions in IETF. Existing solutions are “use BGP routing and multihoming” - which has obvious scalability limitations in regards to the global routing table - or “use IPv6 NAT/Prefix translation” - which has implications on end-to-end address transparency.

Solving these is outside the RIPE APWG's scope.

If we acknowledge that the requirements for "BGP based multihoming" for a certain subset of networks exist, the consequences are "owners of such networks need to have portable address space". With the current address policy, LIRs can use a PA /29 for that (splitting up the /29 among multiple independent locations, if needed), while non-LIR end users can receive IPv6 PI space, with a /48 per end site, non-aggregatable.

Some problems in this space are

- routing table size (global impact)
- yearly recurring costs for address space holders (IPv6 PI is cheaper than a full RIPE NCC membership)
  - discouraging IPv6 PI use conflicts with encouraging IPv6 deployment
- aggregability of multiple IPv6 PI assignments to PI holders that have multiple independent end sites (e.g. 5 sites = 5x /48, not 1x /45)
- IPv6 PI assignment to LIRs holding IPv6 PA is something the policy permits (7.2), but only by demonstrating a "unique routing requirement", the interpretation of which can create quite a bit of friction between LIRs and the NCC registration services.

## Transfer Policy and Stockpiling

IPv6 address blocks can be transferred under the "normal" Transfer Policy (RIPE-682), just like IPv4 address blocks or AS numbers.

This is not something intrinsic to IPv6 needs, but came out of the desire to have a unified Transfer Policy for all number resources managed by the RIPE NCC. That said, there are legitimate business cases for transfers of IPv6 blocks, like mergers and acquisitions, and having identical policies for all resource types makes this much easier for all parties involved.

We have observed that some entities have acquired multiple /29s, in a few cases up to "over 50" /29s. This is not exactly the intended outcome, but we consider this as not critical yet (50 /29s are about a /22 worth of total space, of which there are many). The common assumption is that this is a side effect of certain actors opening multiple LIRs to collect multiple IPv4 /22 address blocks, and then merging the LIRs with their v4 and v6 space - so the amount of /29s involved should become somewhat bounded. Monitoring of future developments is still seen as necessary.

## Recommendations: items where revisiting policy might be worth considering

(not calling this "Conclusions" as we want the work to start here, not to end)

- simplify language
- simplify requirements for /28, /27, /26 ... allocations - explicit numbers, not HD ratio, and maybe start with "less drastic" documentation requirements for "just +1 bit"?
- change default allocation size? (4.3)

- revisit “special case” assignments for IXP fabrics and root DNS operators (6.) <-> loosen up general IPv6 PI policy (for LIRs!) or keep special cases?
- possibly fold IPv6 for IXP policy into main IPv6 policy document (<https://www.ripe.net/publications/docs/ripe-451>)
- revisit IPv6 PI (7.x)
  - cost
  - aggregability
  - multihoming options
  - criteria for assignment
- revisit aggregation requirements and expectations (BCOP)