

# IPv6 Routing Recommendations

Rob Evans  
...and YOU!

# Why are we doing this?

- Previous IPv6 allocation policy was very strict
  - Allocation must be announced as a single prefix
  - $>X\%$  usage before an LIR can get another prefix
- Problems
  - Traffic engineering
  - LIR made up of disconnected networks

# What was the solution?

- Step 1
  - Remove the requirement to announce an allocation as a single prefix
    - The alternative was to allocate /32s for routing reasons
- Step 2
  - Suggest that in some circumstances, breaking an allocation into smaller chunks was acceptable
- Step 3
  - Allow those chunks to be routed

# How are we doing this?

- 2009-06 removed the routing requirements from the allocation policy.
- Draft document to this group on routing recommendations
  - Follows on from RIPE-399
  - Admits that in some cases it is desirable to break up a PA block.
    - Recommends that the scope of more specific prefixes is limited where possible
  - Suggests an (arbitrary) filtering limit

# Reasons for the limit.

- Nobody wants to see /32s broken up into 65,536 /48s.
- The limit of routability can ... no ... **WILL** be determined by the operator community.
  - Of which we're part.
  - We pay for the memory.
  - We want BGP to converge.

# What should the limit be?

- Document currently says 'prudent subdivision' of a PA allocation.
  - Suggests quantifying that as four bits more than the minimum PA allocation size.
  - /36 with current practices

# What should the limit be?

- RENATER have said this is not enough for their network
  - Should it be more?
    - If it should, will there really be many operators that will simply announce e.g. 256 x /40s?
  - If we don't want to split an *aggregatable* block further, is this back to Address Policy?
    - Where is the boundary between PA and PI?
    - Does this require a more liberal PI policy instead?

# Other comments

- Ruediger has said the document uses the term 'deaggregation' too much.
  - I can reword that to mean breaking up a provider aggregatable allocation.



# What do you think?

- How should IPv6 routing recommendations be documented?
  - We are but one forum in one region.
- What should those recommendations be?
  - If the current document is reasonable for most cases, how do we handle the exceptions?
  - As Shane Kerr pointed out, routability cannot be completely divorced from allocation policy.