

# IPv6 Prefix Length Recommendation for Forwarding

**6MAN WG**

IETF 91-Honolulu, November 2014

M. Boucadair ([mohamed.boucadair@orange.com](mailto:mohamed.boucadair@orange.com))

A. Petrescu ([Alexandru.Petrescu@cea.fr](mailto:Alexandru.Petrescu@cea.fr)) speaker

# Context & Goal

- Discussions about the prefixes length is recurrent in IPv6-related mailing lists (including 6man and operational mailing list)
  - e.g., Recent discussions on the 64-bit boundary in IPv6 addressing revealed a need for a clear recommendation on which bits must be used by forwarding decision-making processes
- A clear statement is missing
  - A simple statement making it unequivocally clear is helpful for the operational community

# A Simple Recommendation

- Forwarding decision-making processes **MUST NOT** restrict by design the length of IPv6 prefixes. In particular, forwarding processes **MUST** be designed to process prefixes of any length up to /128, by increments of 1.

# Action Point for the Next Iteration of the I-D

- Comments from F. Gont
  - Suggest that (either):
    - The track is changed to “Informational”, such that this document serves as a kind of “implementation guidance”, or,
    - It clearly states which document it is updating.
- Any option about how to proceed?

# More comments

- Example TCAM performance
- Faster for  $\leq 64$ bit prefixes, but price
- Some performance ratios are the same since decades
- CIDR like
- Is an RFC going to help?
- Relate to src-dependent routing work
- Useful to a wider audience
- Helps interpret 64bit in architecture

# Next Step

- Fair support in the mailing list
- Request adoption