Requirements for Extending IPv6 Addressing with Port Sets

draft-boucadair-softwire-stateless-requirements
Softwire WG Interim Meeting-BEIJING, September 2011

M. Boucadair, C. Bao, N. Skoberne, X. Li and X. Deng

Purpose & Scope

- Identify the list of requirements to be taken into consideration in the design of stateless 4/6 solutions
- These requirements cover the way IPv4-embedded IPv6 address and prefix are to be built when embedding the port information
- **Trivial requirements** such as the following **are not repeated** in the document
  - *Routing protocols should be kept the same, unaware of any A+P processing*
  - *Restoring as much end-to-end connectivity as possible*
  - *Leveraging existing mechanisms and protocols*
  - *Leveraging ISP's existing equipment and software systems (billing, AAA, etc.) as much as possible*
  - *Simple processing in the network*
  - *Allowing direct communication between A+P-aware customers for non IPv6-enabled applications*

Need to Harmonize Terminology

- **Need to adopt a common terminology**
  - Overall Solution: various terms are currently used
    - Stateless DS-Lite, Stateless A+P (RFC6346), 4RD, SMAP, dIVI, dIVI-PD,Stateless 4/6, S46T, etc.
  - Functional Elements or Nodes?
    - Some documents refers to nodes, e.g.,
      - BR (Border Router)
      - AFTR (Address Family Transition Router)
      - PRD (Port restricted Device)
    - Others define functions, e.g.,
      - ICXF (Stateless IPv4/IPv6 Interconnection Function)
      - SMAP (Stateless A+P Address Mapping Function)
  - Set of ports
    - Port Set, Contiguous Port Range, Non-Contiguous Port Range, port-range (RFC6052), Port Set Index, Port Range Mask, etc.

Terminology

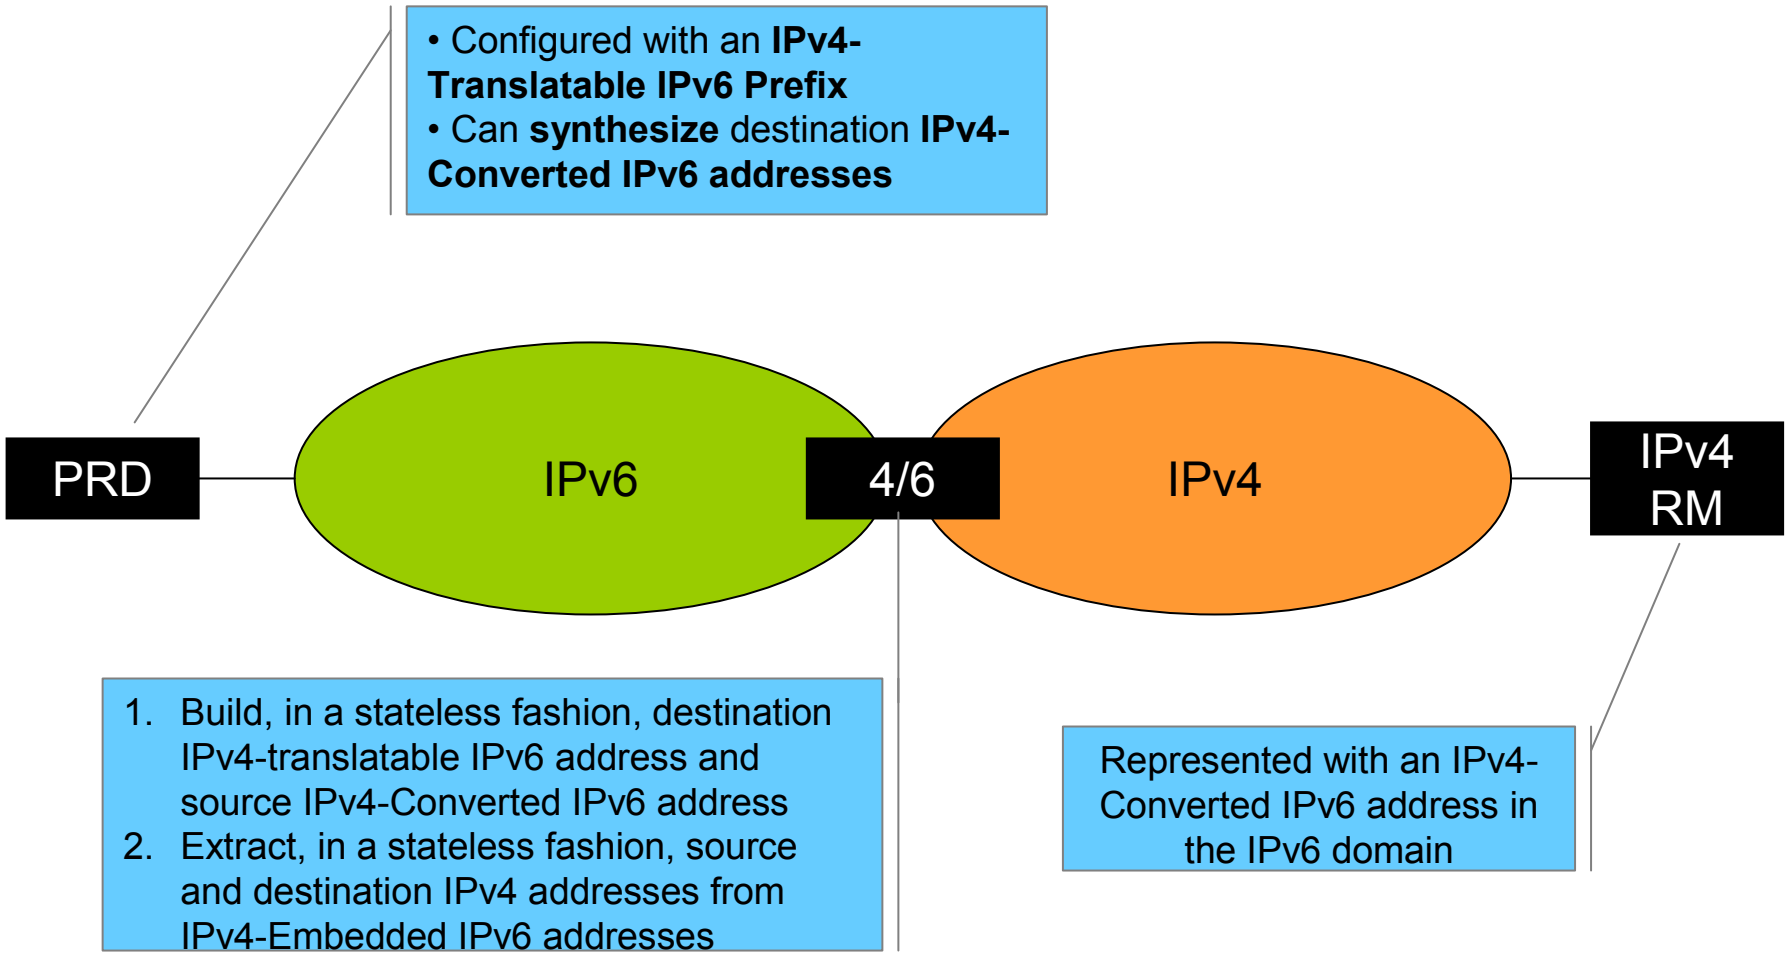

REQ#1

- **The administrative entity operating the stateless solution MUST be able to select the length of the prefix to be used to build IPv4-translatable IPv6 addresses/prefixes**
- Discussion

REQ#2

- **When extending the IPv6 address with the port, the same format SHOULD be used to build both IPv4-translatable IPv6 prefixes/addresses and IPv4-converted IPv6 addresses**
- Discussion

REQ#3

- **Some service providers may require the ability to unambiguously distinguish IPv4 traffic from native IPv6 traffic**
  - **e.g., multi-topology contexts where IPv4 and IPv6 traffic may be conveyed over different paths**
  - **Accounting purposes**
  - **Dedicated per-subscriber policies**
- Discussion

REQ#4

- **When only one single IPv6 prefix is assigned for both native IPv6 communications and the transport of IPv4 packets, the IPv4-translatable IPv6 prefix MUST have a length < /64**
- Discussion

REQ#5

- **The algorithm that computes how port information is conveyed in IPv4-embedded IPv6 addresses MUST be standardized for the sake of interoperability.**
  - **Do we allow the support of multiple algorithms a la RFC6056?**
- Discussion

REQ#6

- **The allocation policy of IPv4-translatable IPv6 prefixes embedding the port information MUST preserve proper prefix aggregation**

- Discussion

REQ#7

- **Service Providers SHOULD be able to support different classes of customers:**
  - **i.e., be able to assign port ranges of different sizes to customers without requiring any per-customer state to be instantiated in network elements involved in data transfer**
- Discussion

REQ#8

- **Applications requiring even/odd and port contiguity (e.g., RTP/RTCP) SHOULD NOT be broken due to the port set assignment scheme**
- Discussion

REQ#9

- **The ability to assign or not the 0-1023 port range should be left to each Service Provider and not excluded by default**

- Discussion

IGD.1

- **As discussed in RFC6269, IGD.1 is broken**
  - **Is there any need to spent effort on designing algorithms which are IGD.1-friendly?**
  - **This induces complexity with no guarantee IGD.1 will succeed**
  - **Our take is NO**
- Discussion