Multi-Addressing Considerations for IPv6 Prefix Delegation
(aka IPv6 Prefix Delegation Models)

Fred L. Templin (fltemplin@acm.org)

IETF102 v6ops Working Group

July 19-20, 2018

Why Multi-Addressing?

- RFC7934 (BCP204) discusses recommendations for providing hosts with multiple global IPv6 addresses
  - Benefits include:
    - Privacy addressing to prevent tracking
    - Support multiple processors inside the same device
    - Extending the network through "tethering"
    - Running virtual machines on hosts
    - Support translation technologies
    - Future applications such as per-application IPv6 addresses
    - Avoid continued dependence on Network Address Translation (NAT)
  - Methods include Stateless Address Autoconfiguration (SLAAC), DHCPv6 address assignment, unique IPv6 prefix per host (RFC8273), DHCPv6 prefix delegation
- **Of these, only prefix delegation allows multi-addressing without disturbing other nodes on the upstream link (no upstream-link MLD/DAD needed)**
- **Goal: document multi-addressing models for IPv6 prefix delegation**

IPv6 Prefix Delegation Models and Multi-Addressing

- Many different prefix delegation alternatives - DHCPv6 is primary example; others include PIO eXclusive, proprietary IPAMs, network management, static configuration, etc.
- **Classic routing model** is when the node provisions the delegated prefix to downstream-attached networks, e.g., a tethered Internet of Things, an internal network of virtual machines, etc.
- **Multi-addressing host model** is when the node uses the delegated prefix for its own internal multi-addressing purposes

Case 1: Classic Routing Model

- Requesting router receives delegated prefix and distributes it to downstream networks
- Useful for “Internet of Things (IoT)”
- Example 1: cellphone with tethered external network (e.g., bluetooth)
- Example 2: laptop with an internal virtual network of VMs
- Example 3: home network router

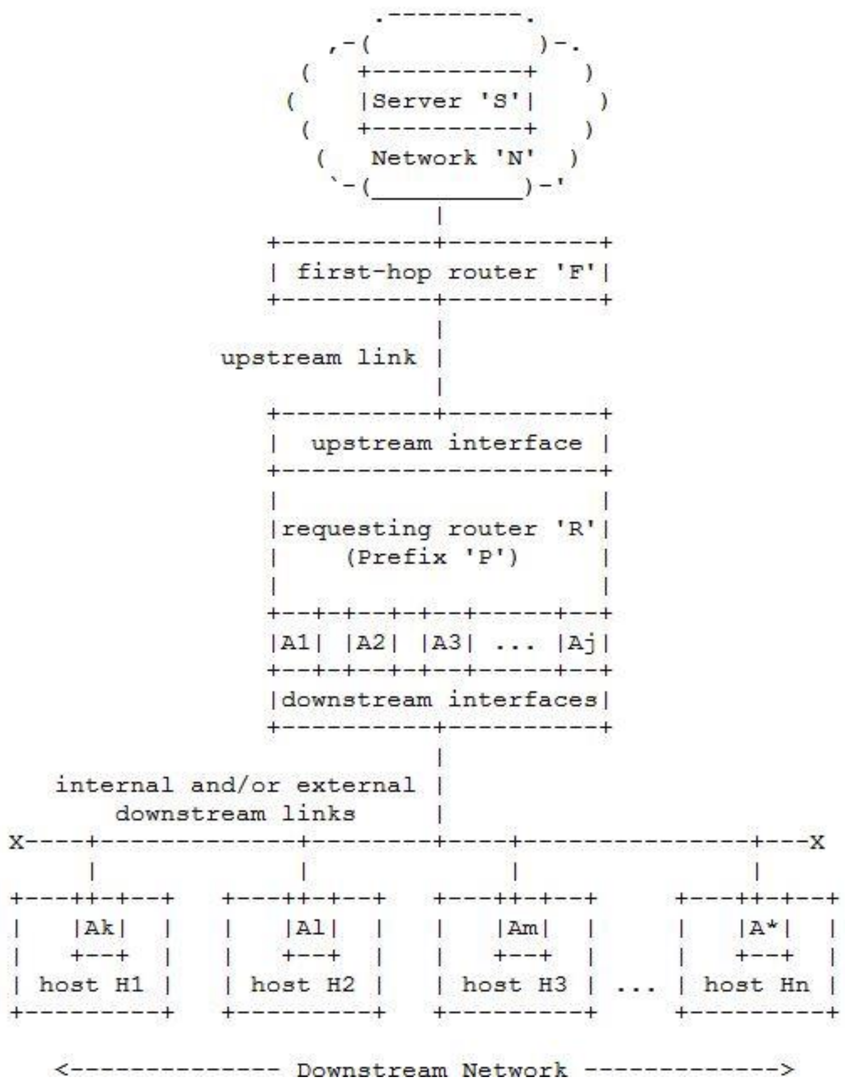

Case 2: Multi-addressing On Virtual Interfaces

- Requesting node assigns addresses from delegated prefix to internal virtual interface (e.g., a loopback) without invoking MLD/DAD on the upstream interface
- Unlimited numbers of addresses available
- Example: any host with an internal virtual interface on which addresses can be assigned

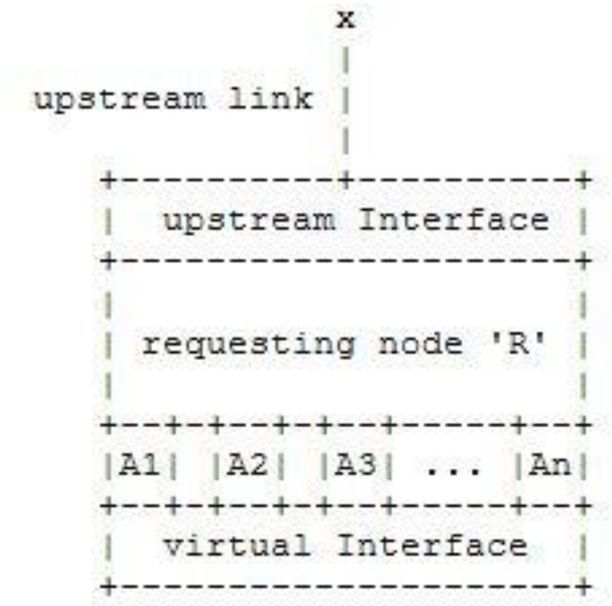

Case 3: Multi-Addressing on Upstream Interface

- Requesting node assigns addresses from delegated prefix to an upstream interface without invoking MLD/DAD
- Unlimited number of addresses available
- Example: any host that cannot assign addresses to any other interfaces besides the upstream

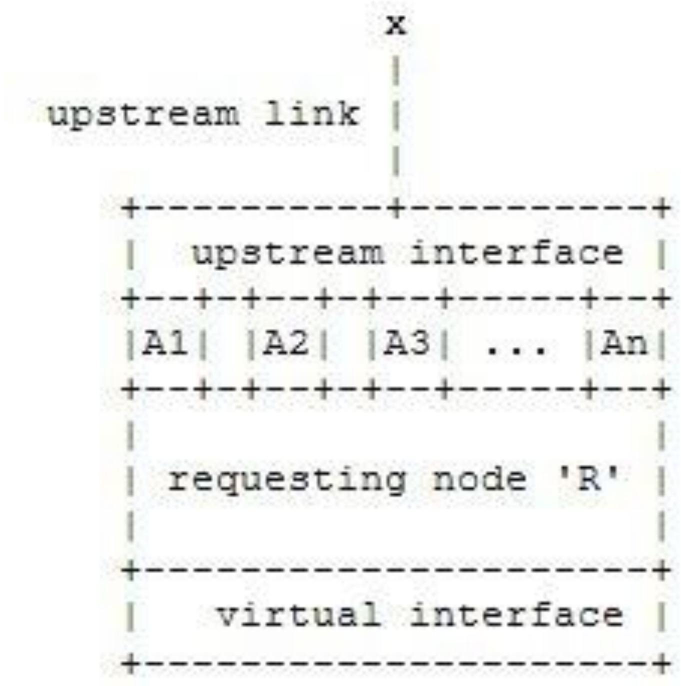

Case 4: Application Addressing

- Requesting node assigns addresses from delegated prefix to its local applications; each application is a “virtual interface” for address assignment per RFC4291, Section 2.1
- Unlimited number of addresses available
- **New model with distinct benefit that a unique address per application may obviate the need for port numbers in the future**

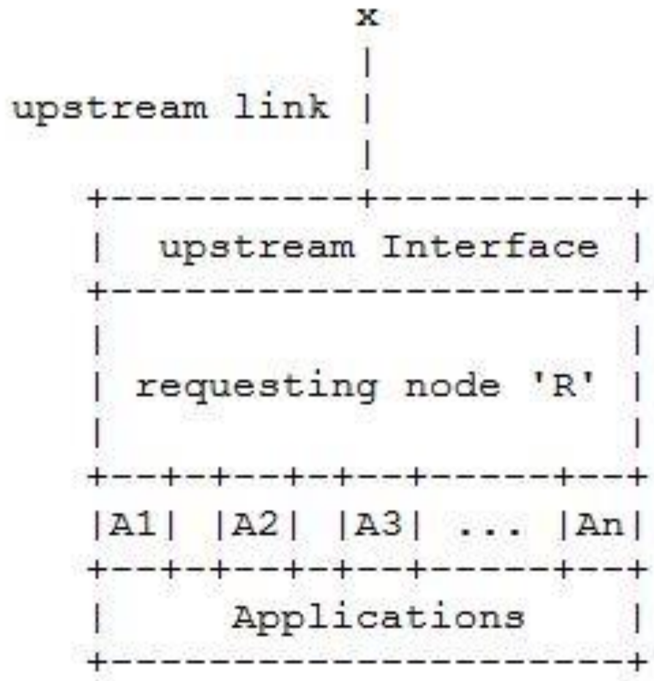

Changes since IETF101

- Clarified that the classic routing model applies to both external (physical) networks and internal (virtual) networks
- Dropped discussion of weak vs strong end system
- Added discussion on “address per application”
- Added discussion on relation to RFC7934 and RFC8273
- Received list comments on 6/15/2018 and posted proposed resolutions on 6/18/2018. Resolutions to be folded into next version.

Draft History

- Draft -00 posted 11/06/2015 and announced to v6ops
- Draft -15 presented at IETF100. Significant comments received at wg session and on the list afterwards.
- Draft-19 presented at IETF101. No time for questions at wg session, but significant comments received on list afterwards
- Now at Draft -21 (includes version-by-version changelog)
- https://datatracker.ietf.org/doc/html/draft-templin-v6ops-pdhost

  - **Working Group Item?**

Backups