IPv6 Prefix Delegation for Hosts

Fred L. Templin (fltemplin@acm.org)

IETF100 v6ops Working Group

November 16, 2017

Draft History

- Draft -00 posted 11/06/2015 and announced to v6ops
- Draft -01 resolved list comments on MLD/DAD
- Draft -02 published 6/27/2016 and was reviewed by Internet Draft Review Team July/August 2016; resulted in publication of -03
- List discussion August 2016 resulted in publication of -04
- -04 expired and was replaced by -05 March 2017
- September 2017 – present there have been 10 updates (now at -15) In parallel with considerable list discussion
- https://datatracker.ietf.org/doc/html/draft-templin-v6ops-pdhost

IPv6 Prefix Delegation for Hosts

- IPv6 Prefix Delegation entails:
  1) the communication of a prefix from a delegating router to a requesting router,
  2) a representation of the prefix in the delegating router's routing table, and
  3) a control messaging service between the delegating and requesting routers to maintain prefix lifetimes.
- Example service is DHCPv6 Prefix Delegation (DHCPv6 PD)
- Document considers the case where the “requesting router” is a host that obtains a delegated prefix for its own internal multi-addressing purpose or to attach a tethered “Internet of Things”

Case 1: Classic Routing Model

- ‘D’ delegates prefix ‘P’ to ‘R’
- ‘R’ can delegate sub-prefixes from ‘P’ to downstream networks and/or assign addresses ‘A(i)’ taken from ‘P’ to a downstream interface
- Hosts ‘H(j)’ assign addresses ‘A(i)’ taken from ‘P’, and may also further delegate sub-prefixes from ‘P’ on their own downstream interfaces
- Example 1: cellphone with tethered external network (e.g., bluetooth)
- Example 2: laptop with an internal virtual network of VMs

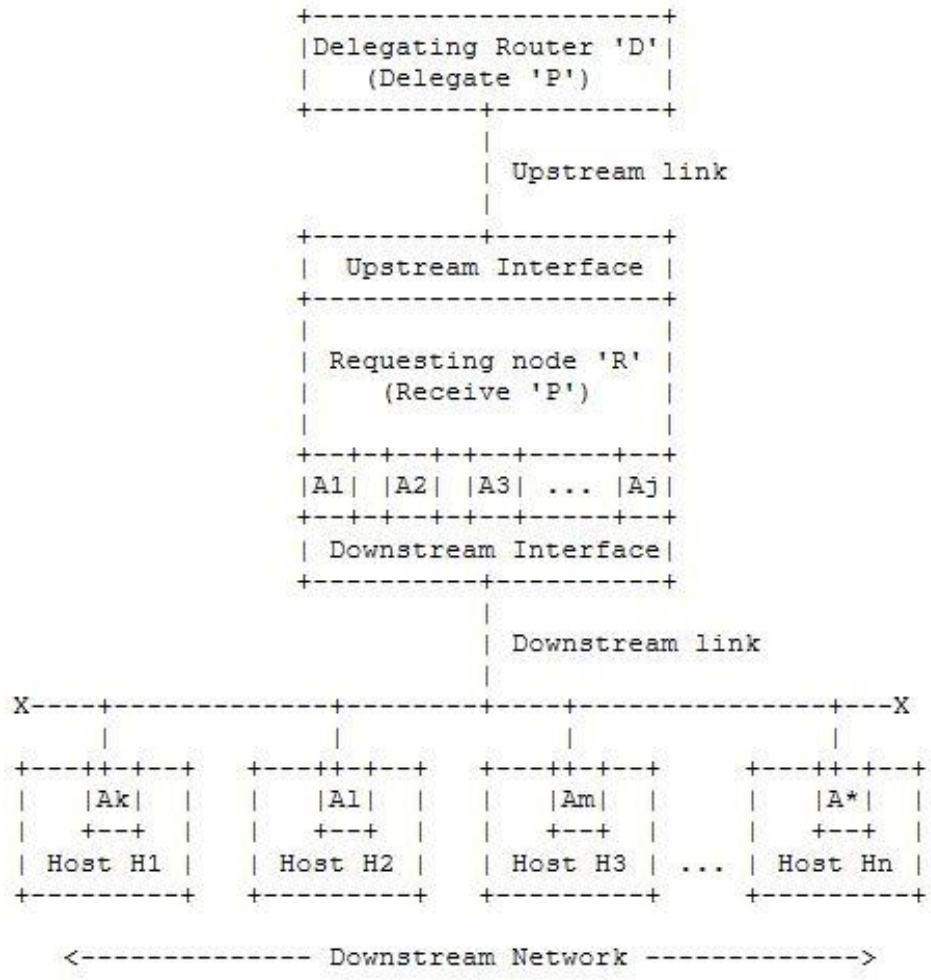

Case 2: Weak End System Model

- ‘D’ delegates ‘P’ to ‘R’
- ‘R’ can assign addresses ‘A(i)’ to an internal virtual interface (e.g., a loopback) without invoking MLD/DAD on the upstream interface
- Example: any host with an internal virtual interface on which addresses can be assigned

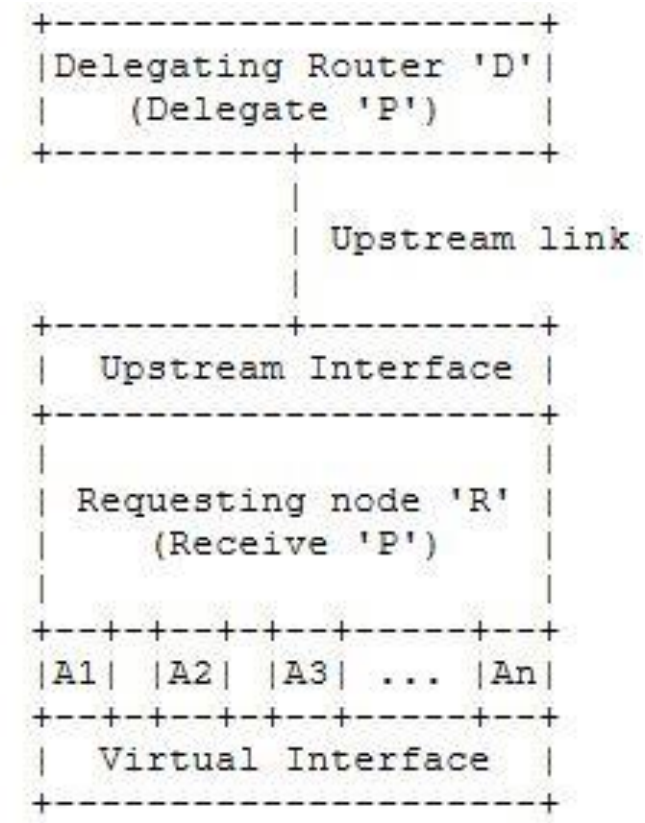

Case 3: Strong End System Model

- ‘D’ delegates ‘P’ to ‘R’
- ‘R’ can assign addresses ‘A(i)’ to an upstream interface without invoking MLD/DAD
- Example: any host that cannot assign addresses to any other interfaces besides the upstream

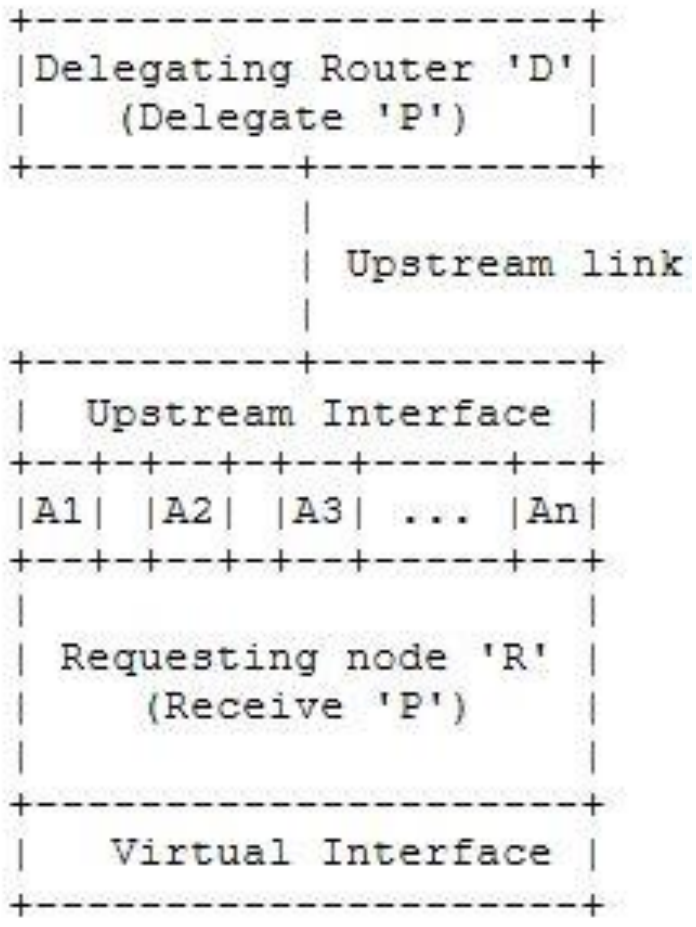

Additional Considerations

- MLD/DAD Implications
  - The host does not use MLD to join the solicited-node multicast group and does not perform DAD over the upstream interface in any of the three cases
  - Acceptable because PD guarantees that no other node will receive the prefix
- Dynamic Routing Protocol Implications
  - Host can be configured to participate or not participate in a routing protocol over the upstream interface
  - Nodes that do not participate in routing protocol send all outbound packets to the delegating router as their default router
  - Future Redirects may inform host of a better first hop

Additional Considerations (cont'd)

- IPv6 ND Implications
  - Node acts as a simple host to send RS messages over the upstream interface
  - Node sets "Router" flag to TRUE in NA messages
  - Node does not send RA messages over upstream interface (upstream interface is not an advertising interface)
  - Delegating router may return a Redirect informing node of a better first hop via the upstream interface
- ICMPv6 Implications
  - Routers send Destination Unreachable (DU) "No route to destination" and "Address unreachable" to remote sources as necessary
  - Hosts send DU "Address unreachable" to local sources and "Port unreachable" to remote sources as necessary
  - Hosts that maintain prefix delegations per this document observe the ICMPv6 specifications for both hosts and routers

Implications for Vendors and Operators

- Hosts that require prefix delegations act like routers from the standpoint of prefix delegation but act as hosts from the standpoint of their local applications (but the network doesn’t care!)
- Allows for unlimited multi-addressing in the spirit of RFC7934
- Multi-addressing does not cause any MLD/DAD messaging over the upstream interface
- Opens new possibilities (e.g., a different and unique IPv6 address for each of the host’s local applications)

Next Steps

V6ops Working Group item?

Backups