

Signalling DHCPv6 Prefix Delegation Availability

[draft-collink-6man-pio-pflag-02](#)

Lorenzo Colitti, Jen Linkova, Xiao Ma
IETF 117, July 2023

Recap: Problem Statement

- IPv6 devices almost always have multiple addresses
 - Link-local, stable, privacy, 464xlat, multiple prefixes/renumbering, ...
 - Often, share connectivity to other IPv6 stacks (VMs, containers, tethered devices, ...)
- On some networks, it's advantageous to hand a prefix to every device
 - Tracking individual /128s is a scaling problem
 - Some enterprise APs drop packets after X (=6, 8, ...) addresses per device
- Other networks, (e.g., home network with a /60) have no problem with lots of addresses, but don't have enough prefixes to hand out a /64 per device
- This draft defines a way to tell the device which prefix (PIO or PD) to use

Recap: Proposed Solution

- Add a new P flag to the PIO
 - “If you understand this flag, please DHCPv6 PD to get a unique prefix, and assign addresses from that prefix, instead of using SLAAC on the on-link prefix”
- Why in the PIO?
 - Must be available to the device before it does SLAAC => must be in RA
 - Specific to the particular prefix
 - Allows, for example, to SLAAC for ULA and PD for global space
 - In a multihoming situation, not different upstreams might support different mechanisms

Update since IETF 116

- Requested last call for [draft-ietf-v6ops-dhcp-pd-per-device](#) in v6ops this week
- Changes since -00:
 - Added Security Considerations and Privacy Considerations
 - Spec changes:
 - On routers, P flag SHOULD be configurable and SHOULD be set to 0 by default
 - Device MUST perform a REBIND if it sees set of P-flag PIOs change
 - Device MUST ignore prefix if it's not of suitable size to support SLAAC (longer prefixes remove the device's ability to extend connectivity downstream via ND proxying)
 - Device MAY fall back to SLAAC if it does not receive any prefixes

Changes since -01 (this week)

- [#20](#) Clarify that devices that is configured to run PD, or devices whose primary purpose is to provide connectivity to other devices (e.g., CE routers as described by [RFC 7084](#)), SHOULD run PD even if P=0
- Minor changes:
 - [#18](#) Remove mentions of DECLINE
 - [#19](#) Change “host” to “device” to match v6ops document

Next steps

- Requesting WG adoption

Questions?

Appendix: Using the delegated prefix

- Device MAY use as many addresses as it wants
- Device MAY use prefix to assign IPv6 addresses to internal components such as VMs and containers
- If permitted by device policy, device MAY use prefix to extend the network
 - => device MUST use DHCPv6 PD hint for prefix size sufficient to use SLAAC
 - Extending the network is already always possible via NAT44 / NAT66. PD allows extension of end-to-end connectivity as well

Appendix: Renumbering

- Device tracks every (unexpired) PIO with P=1
 - Keep DHCPv6 PD running as long as at least one such prefix exists
 - Start PD SOLICITs or REBIND when such a prefix appears or is deprecated
 - Routers that require DHCPv6 PD to work can still request PD regardless of P flag
- Why not RECONFIGURE?
 - Not widely implemented, difficult to use (requires authentication)

Appendix: Multihoming

- If multiple PVDs on link, every packet's source address must match next-hop
- Device shall maintain the mapping between delegated prefixes and routers (relay) link-local addresses so Rule 5.5 can be used
 - (yet unclear) what if relay is not collocated with the router
- Why not PVD option in DHCPv6?
 - Previous work in this area was blocked by an IPR claim
- Why not ICMPv6 redirects?
 - Redirects not specific to source address