Signalling DHCPv6 Prefix Per Client Availability

draft-ietf-6man-pio-pflag-05

Lorenzo Colitti, Jen Linkova, Xiao Ma, David Lamparter
IETF 120, July 2024

Recap

Recap: goal of this work

- Signal that network supports draft-ietf-v6ops-dhcp-pd-per-device
  - Similarly to how A=1 signals SLAAC and M=1 signals DHCPv6

- Why do we need a flag?
  - Device needs to know **before** running SLAAC if it should do PD instead
  - Some networks (e.g., home network with a /60) are happy to hand out a /64 to devices that ask, but will run out of space if too many devices ask
    - Device **needs to know if it is safe to ask** for a SLAAC-sized prefix

- Why SLAAC-sized prefix? See draft-ietf-v6ops-dhcp-pd-per-device

Update since IETF 119

- Companion draft-ietf-v6ops-dhcp-pd-per-device is now in RFC editor queue
- Draft has been in WGLC since April
- Discussion since WGLC focused on deployment model, not protocol
- No changes to protocol specification: see diff from -02 to -05

Summary of discussion since WGLC

Signalling arbitrary prefix length

Q. Why not instead define P as “PD is available, but prefix length is unspecified”?

A: Clarified text and renamed draft to say that P is specific to pd-per-device

Rationale:

- Delegating long prefixes (e.g., /80) will not cause prefix exhaustion, so “PD is available” does not need to be signalled - device can just ask
- pd-per-device implementations ask for a SLAAC-sized prefix. Many networks today will hand out /64s but will out of space if many devices get them
- DHCPv6 server cannot easily know if DHCPv6 client **needs** a /64 (e.g., RFC 7084 router) or **could use** a /64 (pd-per-device)
    - Even if server detects this and returns NoPrefixAvail, this slows down pd-per-device clients

Signalling other PD use cases as well

Q. Can we use P to signal other deployment models? How about tethering?

A. Out of scope for this draft, left for future work

Rationale:

- Deployment models with longer prefixes (/80, /96) do not need signalling
- Might make sense to signal “network supports extension via PD, but does not have enough space to support pd-per-device”
    - Seems to be consensus that this should be done in a separate document

Prefix exhaustion

Q. Setting P=1 could cause networks to run out of prefixes

A. Administrator is responsible for configuring network correctly

Rationale:

- Network administrators today are accustomed to sizing IPv4 address pools and PD address pools
- Planning for /64s in IPv6 is easier than planning for IPv4 /32s because there are more bits to subnet
- P flag SHOULD be 0 by default

Excessive address space usage

Q. /64 per device uses too much address space

A. It is acceptable in many environments, but not others - that’s why this flag is needed

Rationale:

- Many networks today (e.g., mobile) already assign /64 to every device
- See pd-per-device draft for considerations about address space usage
    - Example: enterprise that uses all of 10.0.0.0/8 today needs a /40 to do pd-per-device
    - There are 32 times more /40s in 2000::/3 than there are IPv4 addresses

Creating a new addressing model

Q. Why are we creating a new addressing model and burning a flag for it?

A. There seems to be rough consensus that this is justified

Rationale:

- Both of our current addressing models are signalled by flags
  - DHCPv6 signalled by M=1, SLAAC signalled by A=1
- pd-per-device will soon be an operational model published by the IETF
- pd-per-device is a meaningful improvement **over both** models
  - Combines “infinite addressing” benefits of SLAAC with accountability benefits of DHCPv6

Next steps

Next steps

- Close WGLC
- Request early allocation

Questions?