



Current IPv6 topics in the IETF

Tim Chown, Jisc, tim.chown@jisc.ac.uk

UK IPv6 Council annual meeting, 21 November 2023

Agenda

A quick tour of current IETF IPv6 activity

- The Internet Engineering Task Force (IETF) is the de facto standards body for much of IP and transport networking
- Meets in person three times a year
- Most recently IETF118 in Prague two weeks ago
 - <https://www.ietf.org/how/meetings/118/>
 - <https://datatracker.ietf.org/meeting/118/agenda>
 - Included a very active hackathon the weekend before
- I'll try to cover the hot topics off the WG mail lists

The main IPv6 WGs

One for standardisation, one for operations

- *6man – IPv6 maintenance – new standards or updates to standards*
 - <https://datatracker.ietf.org/wg/6man/documents/>
 - 55 RFCs since WG formed (replacing the ipv6 WG in 2007)
 - 8 active drafts adopted by the WG, 6 expired drafts
 - 16 personal drafts (not yet adopted – many may never be)

- *v6ops – IPv6 operations – usually best practices, informational*
- See <https://datatracker.ietf.org/wg/v6ops/documents/>
 - 83 RFCs since WG formed
 - 8 active drafts adopted by the WG, 15 expired drafts
 - 18 personal drafts

Other IPv6-related WGs

Of particular interest:

- *dhc* – dynamic host configuration
- *snac* – stub network autoconfiguration (new) – Tom talking on this next
- *6lo* – IPv6 over networks of resource-constrained nodes

• **Note: all WGs now include IPv6 as standard**

• No new IPv4-specific work should be started

Hot 6man/v6ops mail list topics?

Some of the interesting topics and areas of debate...

- “IPv6 mostly” (turning off IPv4 on hosts)
- Prefix delegation to hosts
- Variable length IIDs
- RFC8415-bis (DHCPv6 to Internet Standard)
- Registering self-generated addresses
- RFC 6724 update (ULAs)
- Extension header processing

“IPv6 mostly”

Turning off IPv4 on hosts where you can

- Jen will be talking in detail about this
- Sites want to streamline operations – run one protocol not two – dual-stack NOT an end game - something of a theme here today
- Jen’s approach - use DHCPv4 option 108 and CLATs
- Related CLAT draft – draft-link-v6ops-claton-00
 - With RA-based PREF64 prefix discovery (RFC 8781)
- Is there appetite here from dual-stack campuses/enterprises?

Prefix delegation to hosts

A cleaner model for multiple addresses per host?

- Allows a host to obtain a prefix via DHCP-PD
 - *draft-ietf-v6ops-dhcp-pd-per-device-05*
 - Noting a /64 is not “hard coded”, rather a prefix length that allows SLAAC
- Driven by hosts on large L2 networks generating lots of addresses
 - Problematic for wireless controllers and accountability
- May use this to connect other devices or host VMs/containers
 - Alternatives are described in *draft-troan-v6ops-extending-network-reqs-00*
 - e.g., ChromeOS uses SLAAC and ND proxying
- Related - *draft-collink-6man-pio-pflag-03* to signal PD availability to hosts
 - And that they should use that and not SLAAC

Variable length IIDs

Alternatives to /64

- SLAAC is currently defined to use /64
- The “Why 64?” topic is discussed in RFC 7421
- DHCPv6, and DHCP-PD to hosts raises the question of whether host subnets could be longer than /64
- Perhaps make length completely variable, or pick say /80?
 - But then how would you transition?
 - And would there be a “race to the bottom” and use of /128?
- SLAAC at least ensures every network gets at least a /64
- Currently discussion but no WG-adopted draft on this topic

DHCPv6 to Internet Standard

Signaling maturity of DHCPv6

- RFC8415-bis
- See *draft-ietf-dhc-rfc8415bis-03*
- Not adding anything new, tidying up and clarifying from experience, obsoleting some options, applying errata
- Similar to RFC 8200 which saw IPv6 core spec move to Internet Standard in 2017

Registering self-generated addresses

Improving accountability / logging

- See *draft-ietf-dhc-addr-notifaction-06*
- Allows device to inform a DHCPv6 server that it has a self-generated or given statically configured address
 - e.g., host forms address via SLAAC (RFC7217) and then informs DHCPv6 server of that address
- Aim is to facilitate address accountability
 - In a similar-ish approach to DHCPv4 today, without using DHCPv6 to assign addresses to hosts
 - Alternative to querying switch/router ND tables or syslogging
- Though a “bad” host is unlikely to register

RFC 6724 update

Precedence of ULAs in address selection

- Currently the default address selection policy table in RFC 6724 means a host will prefer to use IPv4-to-IPv4 (RFC1918 or global) than ULA-to-ULA
- Not aligned to preference for IPv6 globals over IPv4 globals
- The update would change this
 - See *draft-ietf-6man-rfc6724-update-03*
- Some concern though where ULAs may be put in global DNS (even if not generally good practice)
- Topic arose from US government deployments

IPv6 Extension Header (EH) processing

Recommendations for handling EHs

- Two WG drafts on the topic
- One setting minimum expectations for processing:
 - *draft-ietf-6man-hbh-processing-11*
- The other setting maximum limits
 - *draft-ietf-6man-eh-limits-08*
 - Includes 128 byte parsing buffer, or thus 104 byte EH chain
- The WG seems more comfortable with the former than the latter
- A new personal draft on inflight hop-by-hop or routing header removal
 - *draft-herbert-eh-inflight-removal-01*

Comments or questions?