



IPv6 (hot) topics in the IETF

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

HEPiX IPv6 WG meeting, 3rd June 2020

IETF Who?

So what is the IETF?

- Standards body that defines a broad range of standards, esp. transport and network protocols.
- Over 8,000 RFCs have been published - <https://www.ietf.org/>
- Has over 100 WGs across multiple Areas, including the Internet and Operations Areas
- Meets three times a year, ~1,200 people, though IETF 108 in Madrid in July is on-line only
- Everything is discussed and agreed on mail lists; these are **open, no membership is required**
- Anyone can write a draft, present it to a WG, and potentially have it adopted.
- Once adopted, the WG then seeks to refine and publish the document as an RFC.
- Key WGs relevant to IPv6?
 - **6man** – “IPv6 maintenance” and **v6ops** – “IPv6 operations”
 - Plus others, such as dhc, intarea, opsec, spring, ...
 - In theory all WGs include IPv6(!), and no new IPv4-specific work is undertaken

So what's going on in 6man/v6ops?

Current drafts – some hotter than others

- All WG and personal drafts are listed in the datatracker:
 - <https://datatracker.ietf.org/wg/6man/documents/> and <https://datatracker.ietf.org/wg/v6ops/documents/>
 - To see the personal ones (not yet adopted) scroll right down
- Selected for your delectation in this talk today:
 - “The Packet Steering Wars” – many, many, many emails discussing segment routing
 - Better handling of flash renumbering – “oh, my prefix just changed and my address doesn't work”
 - An IPv6-only networking hint – “don't bother using IPv4”
 - Marking traffic – RFC 8321 for IPv6 – relevant to Shawn's talk
- Plus a few other bits of news
- And there remain the old chestnuts, which are always hotly debated
 - “Wait, there's no DHCPv6 default gateway option?”
 - Though there is now <https://datatracker.ietf.org/doc/html/draft-troan-6man-universal-ra-option-02>

The Packet Steering Wars

aka “why are there a million emails in my IETF inbox?”

- The Source Packet Routing in Networking (spring) WG is defining mechanisms to steer traffic
- Their work led to RFC 8402, the “Segment Routing Architecture”
 - Defines segment routing for MPLS or IPv6
 - “Segment Routing (SR) leverages the source routing paradigm. A node steers a packet through an ordered list of instructions, called segments. **A segment can represent any instruction, topological or service based.** A segment can have a semantic local to an SR node or global within an SR domain. SR provides a mechanism that allows a flow to be restricted to a specific topological path, while maintaining per-flow state **only at the ingress node(s) to the SR domain.**”
- Complemented by RFC 8754 “Segment Routing Header for IPv6 (SRHv6)”
 - The new IPv6 header is essentially an ordered list of Segment IDs (SIDs) that manifest themselves as IPv6 addresses. The packet is routed to each SID in turn, where the next SID in the SID list in the header is taken as the next destination

So what is controversial about SRv6?

It's down to how SR is applied on path

- From RFC 8754 section 3.1:
 - "A SR source node is any node that originates an IPv6 packet with a segment (i.e., SRv6 SID) in the destination address of the IPv6 header. The packet leaving the SR source node may or may not contain an SRH. This includes either:
 - A host originating an IPv6 packet, or
 - **An SR domain ingress router** encapsulating a received packet in an outer IPv6 header, followed by an optional SRH.
- The issue is when an existing packet **enters an SR domain**.
 - It says “encapsulate” to be RFC 8200 compliant, but implementations are **inserting** headers.
- Even with encapsulation, there are gotchas, e.g., you need to ensure the MTU in the SR domain is large enough to support the encapsulation.

What are the various camps doing?

A few further drafts that are emerging...

- A draft on how it is being implemented (by Cisco, Huawei, Barefoot and others):
 - <https://tools.ietf.org/html/draft-voyer-6man-extension-header-insertion-09>
- A draft on why EH insertion is considered bad by many people:
 - <https://tools.ietf.org/html/draft-smith-6man-in-flight-eh-insertion-harmful-02>
- Juniper seem to be defining an alternative, the Compact Routing Header (CRH)
 - <https://tools.ietf.org/html/draft-bonica-6man-comp-rtg-hdr-22>
- Important to note that the spring WG is only using SRHv6, and is opposed to CRH

- **Question:** does WLCG have use cases to deploy SR, and if so using SRv6?
- Some NRENs are deploying MPLS-based SR; SURFnet is one, for TI-LFA
 - See <https://wiki.geant.org/display/APM/18th+STF+-+Copenhagen%2C+October+2019>

Better handling of flash renumbering

When your IPv6 prefix changes suddenly

- Two drafts have been written to describe the issue of hosts having their IPv6 prefixes changed from under them, and possible host and CPE-oriented mitigations
- Appears driven from residential ISP scenarios where an ISP changes a customer's prefix
- Host/SLAAC based view
 - <https://tools.ietf.org/html/draft-ietf-v6ops-slaac-renum-02>
- CPE based view
 - <https://tools.ietf.org/html/draft-ietf-v6ops-cpe-slaac-renum-03>
- Generally recommending better heuristics, including lower default prefix information timers
- The problem may be masked somewhat by Happy Eyeballs for browsers

- **Is this an issue for WLCG?** Probably not as renumbering is rare?

An IPv6-only networking hint

For host networks where you want to turn off IPv4

- A new draft in the dhc WG
- <https://tools.ietf.org/html/draft-ietf-dhc-v6only-01>
- “It allows IPv6-only capable hosts to turn off IPv4 only upon receiving an explicit signal from the network and operate in dual-stack or IPv4-only mode otherwise. “
- **NB. This is a DHCPv4 option**
- A fair number of comments have been made that it is not needed, nor useful. An argument against is that IPv4 stacks need to be updated for it to be useful, so legacy devices will never use it. But the draft has a longer-term view.
- The idea is that IPv4 chatter on the network is then minimised.
- Not clear that anyone has implemented it yet.

Marking IPv6 traffic

Using the “Alternative Marking Method” with IPv6

- Recently adopted as a 6man WG document
 - <https://tools.ietf.org/html/draft-ietf-6man-ipv6-alt-mark-00>
- Based on RFC 8321, which is a passive packet “colouring” specification
- Designed to measure loss, but can also support observations of latency and jitter
- Basic idea is to mark “blocks” of traffic with the same value, and observe their reception (or non reception!) at the desired destination
- The alt-mark draft proposes both an IPv6 HBH and an IPv6 Destination option
 - It had considered other methods, such as the IPv6 Flow Label
 - More discussion on this in Shawn’s talk...
- NB. the HBH/DO is added by the source, not on path
- NB2. RFC 7872 notes that packets with HBH options are more likely to be dropped (in 2016)

Other changes you may have missed

New IPv6 stuff in the last two or three years

- IPv6 became an Internet Standard with the publication of RFC 8200
 - Much gnashing of teeth over certain parts, esp. insertion and processing of headers on path
 - But a strong, positive signal of the maturity of IPv6
- Old school SLAAC with embedded MAC addresses is (in theory) no more – RFC 8064
 - Says use stable, per-prefix host identifiers as per RFC 7217
- Happy Eyeballs got improved with v2 defined in RFC 8305
 - Includes better support for IPv6-only networks using NAT64/DNS64
- New IPv6 Node Requirements RFC; updated from the 2011 version – RFC 8504
- NAT64 / DNS64 RA option to inform hosts of the PREF64 prefix – RFC 8781
 - Potentially useful for WLCG IPv6-only networks
- And finally DHCPv6 got updated, obsoleting the version from 2003(!) – RFC 8415

Any questions?

Feel free to email me at tim.chown@jisc.ac.uk