



IPv6 Operations (v6ops) - IETF 118 Agenda

Tue. 13:00 - 15:00, Congress Hall 2

Meetecho link: <https://meetecho.ietf.org/client/?group=v6ops>

Chairs Opening, 5m

WG status and draft progression are presented.

Call for support/participation in different activities.

WG Drafts

Expanding the IPv6 Documentation Space, Geoff Huston, 10m

- <https://datatracker.ietf.org/doc/draft-ietf-v6ops-rfc3849-update/>

Shorter /32 for documentation.

Jen Linkova: Support

Antonios Moreiras: presented 10 years ago the same idea but you said it was not a good idea. We abandoned the idea. So what has changed?

Geoff Huston: thanks, between 2005/2008 we were concerned about address space. In that period we tended to be conservative. Since then I changed my opinion.

Bob Hinden: we see a lot of deployments, so we need it.

Warren Kumari: the new prefix is better to have a different one to see when it was documented.

Ray Bellis: you're right. It's a matter of address space.

Lorenzo Colitti: such a large allocation is not necessary. Networks may have multiple prefixes so you cannot support it anyway. Having smaller prefixes is good for documentation.

Geoff Huston: let's discuss it on the list, this block is not routable.

Individual Drafts

Use of IPREF as a tool for transitioning to IPv6, Waldemar Augustyn, 15m

- <https://datatracker.ietf.org/doc/draft-augustyn-intarea-ipref/>

IPref is presented as a transitional mechanism with a location address and reference to the node.



Warren Kumari: comparison with similar mechanisms to make it easier to understand

Jen Linkova: confused about why you don't need v4 on the Internet. Nice to compare it with other mechanisms to let people understand how it. If networks A and B are the same administration then fine. How does it scale globally? (full mesh of tunnels?)

Waldemar Augustyn: you carry v4 addresses on v6. As it can traverse NAT, v4 can be eliminated

Lorenzo Colitti: does it make sense from a commercial standpoint? The transition is already done, can it be commercialized? The cost of keeping v4 in DS is not high. If this was available 15 years ago then ok, but now it is too late

Waldemar Augustyn: I don't think current transition techniques are going to succeed. When the enterprises smell they cannot go through transition, they are not forced.

David Schinazi: You lack the analysis of why the other techniques failed, including MASQUE. My point is we don't need another transition mechanism.

Happy Eyeballs v3, Tommy Pauly, 15m

- <https://datatracker.ietf.org/doc/draft-pauly-v6ops-happy-eyeballs-v3/>

Basically, the same document as v2, but different due to different DNS types and transport since RFC8305 (SVCB RR, QUIC, Encrypted Client Hello.).

Jen Linkova: a good job, let's adopt. V6-only network support: how to detect such networks? If you detect DNS and DNS64, how to penalize 64?

Jan Zorz: support the HEv2 update. Many IPv6 implementations are broken, and HE hides problems.

Tommy Pauly: there is room for referring to other documents on that

Lorenzo Colitti: This does not belong here (v6ops). Probably this is for transport area folks. The draft is important. You cannot implement this on an IPv6 stack, you need higher layers (Chrome, ...).

Tommy Pauly: at least DNS+QUIC. Not sure where it belongs. But need IPv6 experts too.

Geoff Huston: Supports Lorenzo. Initially, it was just IPv6 against IPv4, it is now a comparison of many high-level protocols. Then this WG is not the best place for expertise. Support initiative in general.

Juliusz Chroboczek: priorities and other stuff is like a subset of ICE and I wonder if anyone with experience in ICE has reviewed it

Gorry Faihurst: from the transport working group: inviting to his WG.



Scalability of IPv6 Transition Technologies for IPv4aaS, Gabor Lencse, 10m

- <https://datatracker.ietf.org/doc/html/draft-lencse-v6ops-transition-scalability>

Tobias Fiebig: speaking for myself, OpenBSD is not the fastest when it comes to packet processing performance

Gabor Lencse: it is just a performance for stateful NIST 64

Lorenzo Colitti: RFC for benchmarking results is not very useful – it would be outdated very fast

Gabor Lencse: The methodology is close to WGLC in BMWG.

Tobias Fiebig: maybe move to IRTF

Gabor Lencse: methodology belongs to BMWG where there is a draft on that. It does not belong to this group, understand.

IPv6 CE Routers LAN Prefix Delegation, Tim Winters, 10m

- <https://datatracker.ietf.org/doc/draft-winters-v6ops-cpe-lan-pd/>

Ted Lemon: does the spec speak about how to get prefix delegation? The reason is that having a doc to explain bad ideas, is a bad idea.

Lorenzo Colitti: do whatever you can to have this doc adopted (cancel whatever disappoints people). Chairs, please adopt.

Xipeng Xiao: we'll take it to the list

Lorenzo Colitti: another change, PD on the LAN

Tim Hill: hope the doc will be adopted

Lorenzo Colitti: the other document is required. For EUI, can I get a persistent ID used as a proxy for location info?

Tim Hill: Yes

Tom Hill: more prescriptive about transition technologies rather than descriptive. Never seen a CE doing them all.

Tim Winters: scared of the dispute that may happen

Lorenzo Colitti: It is RFC, not an RFP

Extending the Network - Host Requirements, Ole Trøan, 15m

- <https://datatracker.ietf.org/doc/html/draft-troan-v6ops-extending-network-reqs-00>



/64 may run out.

Jen Linkova: I am horrified of NAT. how to define what we want to support, options are too many. It is not possible to solve generic problems.

Ole Trøan: Let's draw a picture of what can be done

Lorenzo Colitti: Share your interest in network extension. There is room in the work presented tomorrow in 6man.

The problem is that the amount of addresses upstream is at the mercy of other organizations.

Please, consider the E2E requirement.

Ole Trøan: Agree. Consequence of failure on extension

Timothy Winters: willing to help. If we don't need more guidance from people.

Philip Tiesel: would solve lots of problems having this sort of extension from a security point of view

Martin Hunek : we need to define the behavior of the hosts which is missing in the PD draft. The device should be capable of asking extension. Not every network should be extended.

Markus ?: interested in address extensions for containers. Would love to have a spec with less than /64

Operational Presentation: Issues with Happy Eyeballs in various OSs, Nalini Elkins -> Vanessa Maria Fernandes, 15m

45k University terminals. /36 block. Simple and flat architecture.

Tommy Jensen: talk to me later and share some information (solution for the problems claimed)

Xipeng Xiao: interested in finding the traffic % level. Please give an update on this. Why does Windows use more IPv4?

David Schinazi: thanks for bringing the research to us. Talk to Tommy to see if he can solve your issues

Cheng Li: thanks, hope to see more students at the IETF

Greg Shuls: IESG: If you want to explore DNS/DHCP more, come to us

Operational Presentation: IPv6-mostly Enterprise Network - a Google Deployment Experience, Jen Linkova, 20m (for more info please attend the side meeting on Thur.



Motivation is a scarcity of IPv4. SLAAC-only. NAT64/DNS64. Dual-stack SSID because people may move to IPv4 and stick. IPv6 hosts signal that they would be only IPv6-only (DHCPv4 hint). Migration did happen during this year, 90% now. 300k IPv4 released. HY hides problems – need to turn IPv4. Some nodes use NAT44 inside – it is visible after IPv4 is disabled. IPv6 on laptops was disabled by NOC. EH was disabled in some places.

Lorenzo Colitti: forgot to mention the RFC you had to write in preparation for this. I did not believe it happened so quickly. Need to use

network protocols to change the configuration. Write a draft to explain how you did it. Encourage everybody to go SLAAC-only.

Jen Linkova: My boss asked to do it only when you are 100% sure

Ondrej Caletka: we implemented the DHCPv4 hint on 2 RIPE meetings. It was easy. Only VPNs are broken.

Jen Linkova: Probably VPNs have broken because of EH fragmentation.

Ondrej Caletka: DNS 64 on Windows is now optional

Tommy Jensen: XLAT should be activated. Gratitude to Jen's recent drafts.

Xipeng Xiao: Jen has a great presentation, but disagrees on point #0. End users don't care it's IPv6 or IPv4. If we can make IPv6 deployment easy, then users will move to IPv6. For example, Akamai is making it easy for their customers to turn on IPv6. Similarly, Mythic Beasts is providing IPv6-only hosting and using high IPv4 address price as incentive to move their customers to IPv6.