

v6ops session – March 19, 2024 – IETF 119

94 participants

Chairs opening

IETF Note Well and WG Status

- Draft progress increasing.
- More participation from enterprises thanks to new chair Nick.
- Move DHCP-PD-per-device to last call + other 3 WG drafts to WGLC.
- Call for contribution: importance of Brian's book, please help with his book. Deployment of IPv6 is increasing, so people must know in advance what to do, especially for IPv6-mostly that unmasked many problems. Also, awareness on new drafts, measure IPv6 PLR, measure number of IoTs on IPv6.

Eric Vyncke asks for a scribe – Paolo confirms + Eduard online.

Chongfeng Xie

Framework of Multi-domain IPv6-only Underlay Network and IPv4-as-a-Service

<https://datatracker.ietf.org/doc/draft-ietf-v6ops-framework-md-ipv6only-underlay>

- Changes and review introduced in version 04.
- Acknowledges to contributors.
- Think draft ready for adoption.

No questions

Chongfeng Xie

EVN6: A Framework of Mapping of Ethernet Virtual Network to IPv6 Underlay

<https://datatracker.ietf.org/doc/draft-xie-v6ops-evn6>

- This is individual draft to promote EVN for IPv6, Ethernet VPN service as a hack to IPv6 header, no additional header. IP is used for MAC and VPN#. Only Data Plane, and Control Plane are put out of the scope.
- Overall architecture introduced.
- Asking for adoption.

Tommy Jensen: connect authors of "Ethernet connect", masque WG has a different approach. A: ok, will do.

Eric Vyncke: how to handle multicast traffic (BUM). A: in the next version.

Erik Kline: what is the advantage against 7442 (BGP) or SRv6?. Is the only motivation the lack of encapsulation? A: Yes.

EK: You are removing specification for the control plane support. A: Yes, it is easy.

EK: Are you considering traffic injection? What if somebody sends a packet directly to the particular MAC address? (security concern). A: need to think more. All MAC@ are put in IPv6.

EK: between ISPs, anybody can inject on your net. Is EVI same of attachment circuit? A: Yes.

Bob Hinden: Why is presented in v6ops. Xipeng Xiao: shall we move to IPv6Man? BH: No, routing area.

Tim Winters

IPv6 CE Routers LAN Prefix Delegation

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

- v6 in home gateways. Issue is ISP assign /64 with waste of prefixes unused in the home network as CPE needs to delegate more.
- Changes introduced since the last version.
- Question whether to document what internal routers should do (DHCP Relay).
- Ready for WG LC?

No questions

Tim Winters

Basic Requirements for IPv6 Customer Edge Routers

<https://datatracker.ietf.org/doc/html/draft-winters-v6ops-rfc7084bis>

- RFC updates shown and discussed, change of wording.
- V00 is the same as 7084.
- 6RD/DS-Lite removed, EUI-64 prohibited, NUST ingress filtering.
- Outstanding questions: BCP? Things to be included (port control protocol, NTP on the LAN, RFC 8925, RFC 9131), etc.

Eric Vyncke: asking for a feature to disable anti-spoofing filtering. A: good you mention. It is SHOULD now.

Jen Linkova: regarding ipv6-only, v6 preference option should be combined with PEF64. Need to be careful on how to enable it. A: expect that operators would provide it, we want to do it with operators in mind.

JL: So, you think that sometime operators will do, but do not forget Pref64 or things would be broken.

Dave Thaler (historical chair of PCP WG): could PCP be SHOULD? A: definitely in the CPE space. The original text included UPMP. PCP is ok to be included.

Erik Kline: What is planned for NTP? Looks like nobody using it, low value.

Answer: on the LAN, prefer to omit it but Jordi would like to add it. We end up with NTP server.

EK: Actually, are both things, but nobody uses them. A: yes.

EK: Can you ask NPT mail list if they have a preference?

Warren Kumari: shall we include NTP? Yes. But let's keep at the minimum the things otherwise no one will implement. A: agree, we should be careful. Let's remove the text.

David Lampaert: Relay NTP (fine) or run NTP in the router (horrible). Please, clarify. A: Yes, we should be clear.

Karsten Walther: NTP needs infrastructure, many networks are offline.

Discourage from requesting NTP. A: Point taken.

JL: Too many related drafts, asking for a reduction. A: Yes, the first is how a CPE is going to do these things, should implement. The other 4, let's see later.

Dave Thaler: My opinion, anything that is a MAY should be left out. Anything that is a SHOULD or a MUST should be mentioned.

Paolo Volpato

IPv6 Site Connection to Many Carriers

<https://datatracker.ietf.org/doc/draft-fbnvv-v6ops-site-multihoming>

Bob Hinden: The survey needs more responses to be valuable. A: For sure.

Jen Linkova: PVD with logical links may help, for example, SmartPhone. A: it was discussed.

JL: What kind of networks are you talking? What is CPE? A: Sites.

Tom Hill: NPT should not be encouraged, it has the same problems as NAT. A: Agreed, it is just to document.

TH: PI would drive BGP full-view growth. A: Let's discuss it offline.

TH: We should not always follow customers, we have to drive them. A: Agree. We are just documenting what is used.

David Lampaert: People in the survey may be puzzled: is it about 1 ISP or many?

Lorenzo Colitti: We need to say that NPT damages application developers.

Jen Linkova: People may be using it but are not sure that it works because DualStack would mask problems.

Nalini Elkins/Amogh Umesh

Implementation and Performance Evaluation of PDM using eBPF

<https://datatracker.ietf.org/doc/draft-elkins-ebpf-pdm-ebpf>

- Use of EH using eBPF. Reason for using eBPF explained.
- Packet processing results shown and compared.

Eric Vyncke: congrats on presenting something on your first IETF meeting.

Jen Linkova

Using Subnet-Specific Link-Local Addresses to Improve SLAAC Robustness

<https://datatracker.ietf.org/doc/draft-link-v6ops-gulla>

- After IPv6-mostly deployment, found some issues, e.g. renumbering.
- Virtual Router per PIO (multihoming).

Karsten Walther: Good point. Simpler to say if you use ULA, don't use LLA? A: No, ND works on LLA, RA does not work on ULA.

KW: Controversial. We have routable addresses. First phase, you need LLA, but then GUA available. A: this is 6man discussion because you are talking about changing ND

Timothy Winters: Couple of things. Lots of CPE make configurable LLA. You want to set them differently? A: recommended to use algorithms so that LLA are unique.

TW: But most of them are not automatic. Tested many Enterprise routers. Never LLA auto-generated. You are asking too much – it would not be successful. A: Think about them.

David Lampaert: prefer LLA per PIO than LLA per all PIO, but the draft should choose one. A: I'll work on it.

Warren Kumari: Should we always advertise the global, maybe we need to change ND. Propose to revisit architecture. A: You may not have global. It is 6man agenda.

Thomas Jensen

464 Customer-side Translator (CLAT): Node Recommendations

<https://datatracker.ietf.org/doc/draft-link-v6ops-claton>

- Why CLAT.

- Asking for adoption.

Nick Buraglio: important work. CLAT progress is instrumental to v6 implementations and makes easier to go v6-only

Lorenzo Colitti: good to document even if with gotchas learnt by experience. Should be adopted because even if there are gotchas it's production ready.

Jen Linkova: Prefer CLAT or native IPv4 when you have both? It is a fundamental question that happens in many places (for example HYv3).

Should be interesting to discuss. A: Expect to be independent, solution may not exist, we just need to disclose the pros and cons.

LC: From deployment perspective Android prefers native v4 over CLAT but at that time code was terrible. But today if you have option 108 works.

JL: If there are not configurational errors, Pref64 faster than DNS64. Question is where we want to go.

LC: Does that matter to network operators?

JL: network operators interested to DNS64, interested to know what is expected behavior.

LC These days CLAT preferred over native v4.

Jen Linkova

IPv6-Mostly Networks: Deployment and Operations Considerations

<https://datatracker.ietf.org/doc/draft-link-v6ops-6mops>

- Nice to have a raft on IPv6-mostly.
- Asking for Adoption.

Jared Mauch: Some android devices that seem to ignore resolvers. DNS address may be given in many ways. Which one to prefer? Answer: IPv6 migration needs IPv6.

LC: you can provide a stub record ipv4only@arpa. the rest may use IPv6 DNS.

Warren Kumari: IETF SSID moved to ipv6-mostly, will be presented in plenary.

Jared Mauch: DNS servers and v6 to get through: RAs, DHCPv6. Is there a recommendation on what we should we be doing? A: I didn't think even in v4.

LC: great support for the draft. It is the only way to transition to IPv6 only

While I like the name v6-mostly, it's not ambitious. It works, but propose to rename the solution to "transition mechanism to IPv6-only". A: The world needs to distinguish when you don't have v6-only.

Stuart Cheshire: Corporate policy may ignore local DNS (IPv6-mostly would not work), we could ask exception for Option 108 to be asked anyway. A: Ok, talk to you later.

Geoff Huston

Operational issues

- It's a talk without a draft.
- How well IPv6 is supported in DNS today?
- 3901 (BCP91): DNS needs IPv4. 3901bis: opposite.
- How much is DNS support in IPv6 now? How fast is it?
- 43% IPv4-only, 11% IPv6-only, 46% both. 56% of DNS fragmentations failed.
- Conclusion: DNS is not ready for IPv6!

Stuart Cheshire: trying to understand. Green bar there is a recursive resolver. In DNS we don't know what a resolver is. Each AD is a single name. now it may be 6 then 4, but resolver may flip the answers.

David Lampaert: Reason to believe that DNS mask in theory does not work.

SC: HE meant to prefer v6 but this is not what HE prefers. I send first request, then wait for a little time. It's not all at once and not the old sequential way, it's middle ground. A: Remember your presentation in Seoul. Today DNS is the protocol to win. Is v6 slower? That's a question dependent on density of population where things get fed.

SC: If v4/6 are similar than go w/ 6.

Tobias Fiebig: it is about DNS. What about E2E? A: Google Ads would not permit this.

Bob Hinden: Thanks for the presentation. Discourage to use fragmentation when v6 was designed. Unfortunate DNS over UDP. A: Next to that, fragmentation was a compromise.

BH: DNS over TCP. A. there are more talks on it. Telstra does not work with TCP and v6 with DNS (rest of world works).

LC: Use of fragmentation is correct behavior. It does not work when applications use fragments. A: A server can't say no, I don't answer, so next step is fragmentation. Industry has deployed resolvers. It's not interactive problem.

Warren Kumari: asking again to look at the architecture. What if prohibits fragmentation? A: Jen's draft said fragmentation is high. Ron Bonica proposed to deprecate fragmentation.

LC: Let's drop fragmentation.