V6ops session – Monday 27, 9.30-11.30

1. Chongfeng Xie – Framework of Multi-domain IPv6-only Underlay Networks and IPv4-as-a-Service (draft-ietf-v6ops-framework-md-ipv6only-underlay)

Provide e2e v4 service delivery over multi-domain v6 underlay.

- Xing Li (coauthor, CERNET Center/Tsinghua University) - There is another draft in IDR related to implementation of this one.

2. Chongfeng – Operational Issues with Processing of the Hop-by-Hop Options Header (draft-ietf-v6ops-hbh)

Discuss the processing of the HBH option header and issues currently found in networks. Request for WG adoption.

- Xipeng Xiao (chair) – We take the request for adoption to the mailing list

3. Jordi Palet – 464XLAT/MAT-T Optimization (draft-ietf-v6ops-464xlat-optimization)

Doc was in WG's last call 1.5y ago but expired. Started again. With v6-only networks, there is an issue of unwanted v4 translation. The goal of the draft is to do translation only once at the CPE. 3 ways to resolve the problem: routing to local CDN, create translaton table on CPE or centrilized.

- Jared Mauch (Akamai) – Proposal #1 is unlikely, CDNs don't touch the packet – it is very expensive. We and ISPs don't want to cooperate on this. Operational performance would be compromised. We would probably not deploy solution#1.
  Jordi – That is my point. With IPv4 CGNAT we could download private address space to CDN, that's not posssible in IPv6. It is why we propose #2.
- Warren Kumari – Doc expired for 2 years, chairs should confirm WG is working on it because there were issues during WG's last call.
- Erik Nygren – Doc was actively discussed on solution #2, some issues were resolved, but issues still remain, we need experiment with this, may be a lot of coner cases. If option 2 is deployed at scale, issues of support may arise. If would be operational issues traffic may switch back to IPv4-only. It may be that IPv4-only clients would have small enough market share to this draft not needed.
  Content provider support will also suffer, since this CPE bugs will be really hard to chase down.
- Jen Linkova – Looked at the diff, concerned there are not so many corrections, the text appears the same.
  Wonders what problem we are trying to solve. Who is interested in solving it? CDN or ISP?
  If this would be done wrong then users and ISP support would suffer. It is interesting to see of how much problem we solve here?
  CPEs are already complicated and bug-prone. This solution complicates them even more.

Please address these concerns.
Jordi – Doc stuck because 2 CPE vendors promised to come up with implementations, but it did not happened.
We have addresses many comments in previous versions.
Any new protocol change takes time in the market even when we work for improvements, this is true even in the US/Europe where the replacement of CPE is faster than in other regions.
The advantage is for ISPs. They like that CGNAT is a natual solution for IPv4. Hence, less motivation for IPv6 migration.
The thing is to try to get v6 to work.

- Masanobu Kawashima (NEC) – Support the draft. Prefer approach 3, not 2. Because 2 uses a public DNS resolver. Do you know CDN providers supporting 3? We need to find more use cases.
  Jordi – Yes, but we need to talk with CDN providers again.
- Ron – could we sense in the room of much interest on this? (quick pool) 16:14
- Anthony Somerset - CPE vendors are remarkably slow in this space... i mean look how long it took for CPE's to have reliable native IPv6 support....
- Erik Nygren - I wonder if this should be Experimental rather than Proposed Standard, regardless? and/or is this something where we should keep as a draft until we can get CPE implementations and experience there?
  Éric Vyncke - Without any hat, this seems experimental indeed and should then include a metric/test to check whether the experiment was a success

4. Timothy Winters – IPv6 CE Routers LAN Prefix Delegation (draft-winters-v6ops-cpe-lan-pd)

Requirements for IPv6 CE Routers to support DHCPv6 PD for distributing prefixes to the IPv6 CE Router. The draft proposes flat mode (ask for multiple /64) assigned from the v6 CE router to the LAN, other CEs are DHCPv6 relays.
RFC 7084 update – do we need “bis”?

- Jen Linkova – like the draft, not sure about “bis”. It looks like ti should be a default already.
  Tim – We have 30 CPEs – no one supporting this.
  Jen Linkova – Do you need to cover scenarios such as reboot or recovery? For DualStack it is not so critical, but for IPv6-only it would.
  Tim – It should work. Put more information around that
- Jen – Not acceptable to have reboot-related issues longer than 5 minutes.
  Tim- Will put some text on it.
- Lorenzo Colitti – Second that. Corner cases should be better considered.
  Better to have open source code for that.
  “bis” is not a good idea.
  What reason do you have to work away from hierarchical PD?
  Tim – I haven’t been thinking for a while. See if flat is adopted.

- Toerless Eckert – The redundancy scenario (many uplinks to many ISPs) is not covered it would delay market adoption.
  Tim – worried that we would be lost in this complicated discussion
- David Lampaert – to be explicit: CE should not rely anything to the ISP. it is not clear in the draft.
  Tim – I need to put more words about this. The leak should not demage anything.
- Eric Vincke – Draft is important for SNAC, keep it simple, we need to avoid Homenet thing.
- Michael Richardson - I'd rather do a few more updates to 7084, before we do a -bis.

5. Nalini Eskins – Deep Dive into IPv6 Extension Header Testing: Behind a CDN (draft-elkins-v6ops-eh-deepdive-cdn)

Look at CDN in detail. CDN prefers v4. If it is v6-only, CDN is not handing EH. It works when the server is just directly connected.

- Anthony Somerset (Liquid Telecom) – The main issue is proxy or CDN are L7, this is why EH don't go through always. Signalling to CDN provier is needed to support EH. But how to get respose back because it is transport connection not L7.
  The proxy is a problem because it terminate connection first, you could get the response to the CDN of the end user but not to the origin CDN.
  Nalini – Yes, I want to bring in the problem even if I'm not convinced is L7, but CDN is probably needed in the conversation.
- Erik Nygren – Many issues are conflated here: some implementations get dropping packets with EHs is a different problem from getting the EH to the original server. CDN is L7 proxy is the separate issue from EH drop.
  Does it makes sense to send EH in a separate connection to the Origion?
  Nalini – Yes, let's keep it short. Defenetely multiple issues here.
- Jared Mauch (Akamai) – if two requests come one w/ and the other w/out EH. What is actually the problem?
  Not always preserve behavior on transport and application.
  Nalini – Would be great to have a response for PDM management from CDN. discuss off-line
- Warren Kumari – proxy connect 2 connections with diferent parameters (MTU, etc). We do not propogate TTL or checksum why we need propogate EH though proxy? Proxy works above L3.
  If you like to propagate EH then TCP proxy or IP proxy would be better not a web proxy.
  Nalini – CDN prioritize IPv6 over IPv4, we need to think: do we need this.
- Andrew Mcgregor (Fastly) – You can't pass EH through the proxy, but you can run custom code in the proxy and pass EH to the API of the proxy.
  Some EHs (like performance measurement) should have strict behaviour.

But you should make the customer choose the behavior they want for other EHs.
Nalini – discuss off-line

- Erik Nygren - should CDNs be encouraged to do IPv6 to Origin Server\", at least at Akamai we finally have configurable support for this (and do prefer IPv6 to Origin when its enabled). I'm happy to provide details offline. But CDNs tend to operate as \"surrogate origins\" above L3 and it is often a feature to isolate the origin from clients from a security/etc perspective. Most of our customers likely wouldn't want EH to be passed through blindly --- although making this visible in customer configuration does make sense if there's a demand / use-case.
- Anthony Somerset - i think CDN's could do a feature where they could be configured to do specific EH's such as PDMv2 to the origin when making requests to origin, they would have to work out capturing the response data and presenting to the customer

6. Jen Linkova – Using DHCP-PD to Allocate /64 per Host in Broadcast Networks (draft-collink-v6ops-ent64pd)

V6 hosts are expected to have multiple addresses. The proposed solution is a prefix delegation assignment. Asking for adoption.

- Jared Mauch (Akamai) – Small ISP may give only /60.
  One of the challenges with v6 is that there are many ways to do the same thing. Not clear how this protocol should be used. /64 per device is great, but things get more complicated, in particular in enterprises where you want device-level accountability.
  There is concern that it will make IPv6 even more comlicated.
  Privacy is affected by this solution.
  I can’t keep an SSH session open for too long because privacy address rotation.
  So, how many ways do we want to select an IP address?
  Jen – It is not a small problem. We have chance to close DHCP vs SLAAC discussion.
  The proposed change is small – all mechanism are available, it is just a different combination.
- David Lampaert – a big shift in the prefix management from router to host, is it a good thing?
- Tim Winters (QA cafe) – Put text on the host to host, hair pinning text otherwise not clear.
- Comment – SSH is supported in terms of keeping a long connection (life time is possible to provision).
  Jen – we did mention in the draft: it is possible to rotate prefixes
- Comment – Strongly support draft, in our lab, we are doing with /64 assignment.

If we use IPv6 as IPv4 we would lose a lot.

- Fernando – What if a host needs multiple /64?
  Jen - Don't see problems. Multiple requests. PD is mandatory for any migration anyway.
- Fernando – What about a router that does not support PD or runs out of prefixes? Fallback to SLAAC?
  Jen – Not worse than we have today, expand the address pool. If not PD, it may have SLAAC fallback to error.
- Xipeng – adopton would be on the mailing list
- Éric Vyncke - Just curious about the host to host connection, i.e., from one /64 to another one... ICMP redirect the router ? or hair pining ?
- Daniel Bernier - agree on useful for DCs and/or large consumers of IP addresses (think PodCIDR)
- Erik Nygren - This would be useful for hosting providers. Giving a /64 per VM/host as a default seems like a cleaner default world.
- Philipp Tiesel - I did the calculations for our K8 infrastructure and /64 per host will hurt me by limiting the aggregation levels above (though I have no problems with using a /80 from software)
- Lorenzo Colitti - We should really just think of this as a replacement for IA_NA that gives the host as many addresses as it needs
- Yasunobu Toyota - Basically support this draft, but some hosts require several separate L2 segments; you may need to run both docker containers and VMs.
\nIs there any solution for hosts that need multiple /64s

7. Nick Buraglio – IPv6 Site connection to many Carriers (draft-fbnvv-v6ops-site-multihoming)

Informational draft, we don't want to solve anything new. What you can do with multi-homing.

- Tim winters QA café – 2 things. In the picture you have ULAs, but in the router you have to agree on what those addresses are.
  Second, Multihoming is a hard problem – supports this draft.
- Eric Vincke – Asking to mention PDV solution.
- Toerless Eckert – Tanks for the work. There is almost nothing positive we get from PA addresses. There are fewer "+" but there are several positive criteria. PI is more positive – I agree.
  How do we come up with criteria and agree on it? Other prople may other opinion.
  Nick – PI works, but PA we have to work too.
- Lorenzo – NPT is not recommended/experimental. ULA and private addresses are not endorse by IETF – experimental or discoraged.
  Nick - we don't endorse these technologies. We just list what exists.
  The problem is that it dumps the cost on app developers, and the

consequences (brittle connectivity, bad battery life) end up being borne by users.

- Sheng Jiang – Thanks for the work, good start. Not only for PA addressing important but there are still operations issues also for PI with multiple providers. ISP filters packets from other providers. Please include this analysis in new versions.
  Nick – Yes, we can expand.
- Erik Nygren - very supportive of us having a draft here
- Fernando Gont - I support the document. I would say that the document should analyze the topic from an operational and pragmatic perspective, and not exclusively based on how IETF would like things to be.
- Fernando Gont - aside from working NPT to PS, its document status doesn't mean that the thing is not used in the operational world. i.e., this document should consider that as an option -- no matter whether \"we\" like it or not.
- Erik Nygren - Better might be to focus first on \"draft-fbnvv-v6ops-site-multihoming\" which seems like a good start comparing the options. PI isn't an option for residential, so then it becomes a question of which another option is least bad. If PA with failover can be made robust enough and handle various scaling and failover issues, great. But if not then perhaps NPTv6 is the next least bad option.
- Jared Mauch - if the Internet is for end-users, then that v6ops-site-multihoming and giving ULA equal v6 access for a NAT66/NAT-PT behavior would be of good value for the end-user networks that have the option to multi-home.
- Juliusz Chroboczek - I'm claiming that address translation (whether stateful or stateless) is user-hostile, and that we need to develop better host behaviors for dealing with multiple prefixes being announced on the local link.
- Juliusz Chroboczek - if you're not afraid of unclean hacks, Proxy-ND might be your friend. And if you're not afraid of the IETF's wrath, subnet longer than /64 and use DHCPv6.
- Lorenzo Colitti - The problem with ULA is that it dumps the cost on app developers, and the consequences (brittle connectivity, bad battery life) end up being borne by users

8. Momoka Yamamoto – IPv6 only iterative resolver utilising NAT64 (draft-momoka-v6ops-ipv6-only-resolver)

Why we want to do it. With v4 to v6 translation, IPv6-only iterative resolvers can operate in an IPv6-only environment. When the DNS zone is served by a v4-only authoritative server, the iterative resolver translates the v4 address to v6 to access the authoritative server's IPv4 address via stateful NAT64.
Question to WG: Is this worth documenting?

- Eric Vincke – Presented to DNSops? Simple stuff, but needs to be documented. Document is needed.
- Masanobu Kawashima (NEC) – Could you explain your expected use case (DC, home net, Enterprise, ...)
  Momoka – we operate v6-only AS for experimental use – we have no knowledge as ISP.
  Home networks can ask their ISP public resolver, but they can move the resolver at home.
  Masanobu Kawashima – why not use CLAT instead? Please, clarify use caes better.
- Ralf Weber (Akamai) – Would it work w/ CLAT?
  Momoka – Yes, it works. It just use IPv4 packets.

9. Rei Shimizu (LINE, software engineer for the Line Cloud) – IPv4/v6 dual-stack migration for in-house load balancers in LINE data centers
   Operational case and experience discussed on the use of Load Balancers. LINE proposes a private cloud. CLOS on eBGP. Some pods are IPv4-only and some DualStack. IPv4 may be encapsulated into IPv6 for LB, hence MTU increased to 1650.