v6ops Working Group

Meeting date / Time 2024-11-08 13:00-15:00 Europe / Dublin
Room: Liffey Hall 1
Participants: 79 (in meetecho)

Chairs: Ron Bonica, Nick Buraglio, XiPeng Xiao; Advisor: Fred Baker; AD: Warren Kumari

Minute taker: Eduard Vasilenko, Paolo Volpato

The AD invited people to join ietf-nat64 network, which may become the default net in Bangkok

Chair's opening - Note Well & WG status, Ron Bonica, Nick Buraglio, Xipeng Xiao, 10 minutes
Chairs discussed progress and current status of v6ops, and asked to help with Brian's IPv6 book, and other v6ops milestone tasks.

WG Drafts

1. IPv6-Mostly Networks: Deployment and Operations Considerations, Jen Linkova, 15m
   Main changes from the previous draft version presented.
   David Lamparter: Move ietf net to ipv6-mostly. ETSI event would have 15k users on IPv6-mostly. We'll do it at the congress in Chicago end of the year.
   Nick Buraglio: same at super-computing conference.
   Tim Chown: there is a UK university running an ipv6-mostly net of 5K devices
   Jen: let people know so that they are not scared of doing it.

2. 464XLAT Customer-side Translator (CLAT): Node Recommendations, Jen Linkova, 15m
   Updates to draft presented. ULA as the source would break the CLAT.
   Complex multi-homing with CLAT was discussed.
   Eric Vincke: for the IPv6 fragment, you need to add the new header again.
   Lorenzo Colitti: why slaac address as temporary? Why we should rotate them?
   Jen: because they need to rotate
   Lorenzo: is it for API use? why not pblic address?
   Tommy Jensen: no strong opinion whether to remove the normative

language.
?: “Should” is too much. Change to “Recommend”
Lorenzo: you talked about SA selection. Is SA selection aware of CLAT address?
Lorenzo: Multi PREF64 does not work in practice. Do you want to fix it?
Jen: Better fix, at least document.
Lorenzo: But you need more IPv4 address. Jen: It is not a problem to ask IANA
Jen: yes. We put text in draft, take a look. You need to do something when you have multiple CLAT addresses. We need more work.
Lorenzo: I don’t know how to implement the MUST.
Jen: use the same v4 source.
Philip Teisel: we don’t do SA before the packet entering the CLAT, but when leave CLAT.
Jen: text does not say rule 5.5 or 8028, it says send it to the next router.
Lorenzo: which way do you see your doc? design or operational?
Jen: provide guidance to implementers of CLAT.
Lorenzo: stateles translation does not do tuples. How many addresses do you use?
Jen: 8. we can ask for biggest networks if needed but the problem is to do with multiple CLATs.
Timothy Winters: is there a mechanism to renumber?
Jen: you can gracefully renumber or say don’t do it
David: the 5-tuple is confused, we need 1:1 map.
Jen: let’s assume the 5-tuple discussion is removed.
Eric: consider the pair of PREF64 and the local v6 address.
Jen: use PIOs.
Warren: the 5 tuple is solvable by hashing.
Lorenzo: PIOs is more complicated than the host can do. Can we use PREF64 configuration an implicit source substitute?

3. IPv6 CE Routers, Tim Winters, 15m
   LAN PD completed WGLC. 7084bis adopted. Asked for necessity of SLAAC renum and/or LAN-PD inclusion into the requirements.
   Ross Chandler:
   Lorenzo: do you believe that this is the last update?
   Timothy: let’s get it out by the November IETF.
   Lorenzo: Please include better interval for RAs. We could even propose the exact number.
   Timothy: yes.

Ted Lemon: we need SLAAC renum and LAN-PD.
Jen: it is not mandatory to include LAN-PD, depends on the 6man draft progress.
Timothy: will talk to the co-authors to have it included.
Tim: do we recommend timers for LAN-PD or other timers?
Timothy: no.
Lorenzo: hash …
Timothy: you should use it becuase you may no have flash memories.
Lorenzo: DHCP is based on big discs.
Timothy: it's about leases and bindings.
Lorenzo: DHCPv6 requirement should allow more IPv6 per host.
Lorenzo: ULA may create a problem for LAN-PD, especially if listed first.
Timothy: we are specifying the router that is plugged to ISP, subtending router is out of the scope.
David: the DHCP will give you any trouble.
Lorenzo: asking to keep DHCP states on the storage.

Active Individual Drafts
4. Using Dummy IPv4 Address and Node Identification Extensions for IP/ICMP translators (XLATs), David Lamparter, 15m
Tobias Fiebig: concerned that fragmentation may happen that would be difficult to translate.
David: not in the draft now.
Jen: Do not understand.
Tobias: a packet translated into 2 fragments and ends up into 2 v6 packets.
Jen: currently PATHMTU is dropped. We need to decide. let's take it offline.
Lorenzo: not looked much at the details. Definitely useful draft. there is value in translating of packet too big. You need a dummy addess in Linux that when you do an update, it will find the route looked for.
Jen: do we need a route for this fake address?
Lorenzo: let's be sure that solution works.
David: it's stored in the v6 cache, because it's v6.
Lorenzo: not sure that dummy addresses are needed.
Jen: if we are overspecifying, we can remove text. It's in the draft because we did work on ICMPv6 to v4 translation.
Lorenzo: if an implementation does not know what to do, is there a simple part that says what to do?
Jen: just use he dummy address 1:1.
Lorenzo: if use use multiple clats for all interface is okay.

5. IPv6 Address Accountability Considerations, Tim Chown, 15m
   Jen: useful. we hear people confused in address assignment mechanism. Valid to document what the you could do in this cenario. Warren and me developing a solution for ND telemetry – a very limited subset of messages is needed.
   ND is very chatty even in the test environment, it may not be useful. Layer 2 addresses may be not enough, better to log users.
   Tim: what is missing here is speaking to the peole that want accountability.
   Timothy: support the work. Link locals may be needed too for some governments. ULAs may be needed even more.
   Lorenzo: if there is still text on prefix per host, it does not belong here.
   Tim Chown: No.
   Tim: if you are in v4 net with everything v6-enabled, there will be v6 local traffic.
   Nick: useful work. Accountability want to tie flow to a user, typically done w/ address, MAC, switchport, office machine user.
   David: EVPN seemed to have figured out hot to map IP to MAC using BGP. For SAVI you have DHCPLEASE query. Several methods available to learn from.
   Tim: thanks
   Jen: respond to David. PD per device introduced. if you have bgp that a scary argument.
   Thomas: we have awesome methods that can be used.
   Tommy: don't use BGP. Support the work.
   Lorenzo: say at the beinning what we do today (for IPv4) than explains parity.
   David: everything listed here have privacy considerations.

6. IPv6-only Definitions, Jordi Palet, 15m
   All terminology is scope-dependent. Asking for adoption. Many WGs have terminology.
   Jen: disagree on IPv6-mostly. Delete NAT reference. when you give definitions, people start disagreeing because would be always corner cases. risk of limiting other documents defining specific scenarios. better to leave them to the documents specifying there in their specific scenarios?
   Jordi: important here is not the definition, but set the scope.
   Lorenzo: echo Jen, we don't need definitions. It was not a problem for a long time. it's context dependent.

Jordi: what you call context, it's scope for me.
Paolo: value in work, you can take text from RFC 9386.
Tommy: useful to have a set of central definitions. could be co-author.
Timothy: IPv6 definition for US government already exists. It would be difficult to agree on everything, consensus on something is enough.
Nick: having definitions of profiles would be respected documents. context is key, having an idea of v6-only varies based on the device sits or the network. I see value here.
Jordi: effort to define scopes.
David: value is the lists all the different aspects.
Bob: lot of docs have terms defined, pay intention not to redefine the existing names.
Philip: terminology is painful, understand reluctancy. better to have draft or an IETF page (instead of RFC)
Tobias: DNS people are doing authorative document on the names. Define common usages to avoid getting consensus every time the doc updates.
Tommy: it is useful for the future documents. A central repository may be valuable too. good to put it into an RFC.
Richard: agree on having definitions.

7. IPv6 Network Monitoring Deployment Analysis, Cao Chang, 10m
   Nick: invite everyone to read the draft and comment.

Discussion (as time allows)
8. 464XLAT to STD (discussion only), Jordi Palet, 10m
Lorenzo: before moving a STD we should explain what we are doing. don't object but there is work to do
Jordi: I'm aware
Lorenzo Colitti: what is deployed is different from the 464XLAT specification. We need to re-write it for standard. Moreover, the specification is not complete. DNS64 is not used anymore.
Lorenzo: a lot of thing are not described in 6877. the other two are ok.
Jen: 7915 is proposed std. don't understand what is going to change if we move fom proposed std to full std. Speaking of DNS64 i don't understand why move it to STD. I would oppose. For the others, I don't see value but don't oppose.
Tommy: support for 4x4XLAT. 6877 don't see it referenced.Deprecate 7050.
Jordi Palet: agree on the last one
Jen: 4x4XLAT is not normative, it simply puts thing together.
tommy: deprecate 7050

Bob: generally support moving things to std. pay attention to the doc that will be moved, leave the others apart.
Warren: unclear the work done to deprecate 7050, where it is going to happen.

9. SLAAC Prefixes with Variable Interface ID (IID) Problem Statement, Gyan Mishra, 10m
   Not time to discuss