

IPv6Ops/Sunset4 Joint Meeting Minutes for Tuesday July 21st

=====

Minutes: Jason Weil

Jabber: Ole Troan

Recording at http://ietf93.conf.meetecho.com/index.php/Recorded_Sessions#V60PS

Slides at <https://datatracker.ietf.org/meeting/93/materials.html>

Opening remarks by Fred

Lots of drafts showed up on July 6th which is a problem because there is a rule that you need to discuss on mailing list before you can be added to agenda - so submit earlier and introduce your draft if you want to get on the agenda

Gap Analysis for IPv4 Sunset draft

Added legacy IPv4-only APIs, some editorial changes

Haven't added DHCP4o6, for v6o clients requiring DHCP(v4) configuration options.

Because it's not a general-purpose use; client must support DHCP4o6.

Ian Farrell: "yet"?

Peng: No, don't plan to incorporate

Ian: There's a mechanism to do it over DHCPv6, don't need it in the doc.

-IanF - there is already ways to disable DHCPv4

-needs a reviewer or two - no volunteers in the room

Also need more review of [Analysis of NAT64 Port Allocation Methods for Shared IPv4 Addresses](#)

IPv6 Design Choices draft

Victor presenting

Did WGLC, had comments

More cleanup still needed.

Overview of Diffs

Q: Does WG agree on inclusion of Enterprise and/or EIGRP?

Jan Zorz - yes to enterprise inclusion and no comment on EIGRP

Alain Durand - would be good to have info on enterprise but may be better to have separate document for that instead of including in this doc

Victor Kuarsingh: We tried to keep it high level so it applies to both.

Alain: Take it one level down, so it'll be more useful.

Q: Should we include discussion of NAT and NPTv6? Some had objected on list.

Brian Carpenter - We should include discussion. But you don't point out that ULA usually runs with GUA, NPT is odd, etc. Rewrite ULA text per his comments on list

Alain - knows NAT is cursed but NATv6 comes up in every one of his discussions with the enterprise

Erik Kline - what is the status of ULA Usage document - if moving along could be pointed to for the ULA portion of this draft (Victor - need some discussion in this draft)

Fred Baker - It keeps getting filibustered.

Erik - Let's bring it back.

Lorenzo Colitti - NPTv6 and NAT are presented together and they are very different - most large enterprises are not doing NPTv6 (Victor agrees edits need to be made there)

Erik Vyncke - an Enterprise Guidelines document has been written dealing with these issues

Tim Chown reiterates Erik's comment

Lee - The WG would like to hear from Enterprises about their deployments

Lorenzo - Like Google's enterprise?

John Brzozowski - Comcast can share what they did if interested

TimC - some University campuses may have looked at NPTv6. Not multimode, tend to trust the NREN.

Disposition: Wait for -09 draft, then need some good reviews.

Committed reviewers: Tim Chown, Jan Zorz

v6ops-siit-dc/siit-dc-2xlat

Tore presenting

Documents refer to deprecate-atomic-fragments, which has been moved to rfc6145-bis, so I need to update the docs.

TimC - good work and thanks for presenting

TuVo? from Chunghwa Telecom - virtual machine question on routing - no special consideration needed - should be consistent with virtualized datacenter architecture - presenter not sure what the problem is there

JJB - looking at this for their datacenters - We're trying to make sure we don't require provisioning IPv4 addresses. would like to see this go in the direction of OpenStack - agrees this is useful work

Fred - going to start WGLC on at least one of these on Sunday, 2-week WGLC on one followed by 2-week WGLC on the other.

IPv6-Only for Wired Thin-Clients

Erik Vynke presenting

enterprise desktop issues when trying to deploy with v6-only

-Wake-On LAN - uses magic ether type packet

-One way in IPv6 is to send a global IPv6 addy statically mapped to ff-ff-ff...

Alex Petrescu - many solutions - some only Ethernet others IP - what solution do you need? drawbacks some cards don't support the ffff...

Erik - open

Dave Thaler - Default was changed in IPv4, as you said, for security reasons, so let's not try to create it for IPv6. In MS enterprise, even when the machine is asleep, the card stays awake, is responsible to look for ARP/ND. That's enough to get magic packet to work. "ND offload." Widely implemented—all cards installed in machines with Windows pre-installed since Win7 support this. - if you can get card to respond to arp and ND then much better solution (ND Offload seems to be fairly well implemented per Dave)

-Lorenzo - how does this get advertised - Erik - per VLAN

-Barbara - solution for in the home may be much different - needed for power saving devices but interesting problem

-Benedikt S - could also do with a raspberry pi

-Stuart Cheshire - Bonjour Sleep Proxy does something similar circa 2009 - NIC listens to ND and ARP and looks for magic packet

-Stig Venaas - multicast could be used but could be some security issues

-James Woodyatt - put thought into constrained devices

Pre Boot Execution (PXE) Environment

-RFC5970 not widely supported (needed to support DHCPv6)

-UEFI will fix but not available yet (wait?)

-Wes - should be captured in the gap analysis - if important then should get people to work on it

Disposition: unclear

v6ops-solicited-ra-unicast draft

Lorenzo Colitti presenting

network with thousands of nodes - send RS and all devices wake up - rate-limited to once every 3 secs

-some experience bad battery, drop IPv6 packets or just multicast v6 packets

-Unicast RAs are an option - this draft asks for a knob to configure RA response to RS
JJB - what do you do for solicited node multicast - solicited node multicast scales with size of network so not many collisions - 3 secs undesirable -

-Andrew Yourtchenko - solicited node multicast can be fixed by increased timers - RAs can't

-Fred Templin - ISATAP has been doing unicast RS/RA for years -

-Ole Troan - Just an operational configurable option, that wouldn't change the spec, or update the algorithm? i.e., does this go to 6man or not

Lorenzo - configuration option; 6man won't pass it.

Kyle Rose (jabber) Drop Box LAN has the same effect - might be a window in 4861bis - Lorenzo says could add text

Andrew - maybe do 2 docs - one in ops and another in 6man

Suresh - add reason for why and point to 4861 wording

Mikael - reference ErikV document from two IETFs ago that has a recommendation relevant to this

Mark Townsley - presented the findings of the problem 2-4 IETFs ago - this is recommendation for knob to fix it

-AlexP - "Networks serving lots of devices SHOULD" is too strong. E.g., if doing DNA. So, specify when SHOULD does and doesn't apply

Lorenzo: DNA has recommendation on Alex's use case (I believe)

-qualified with 2 things - fast moving mobiles that don't do DNA and other is Suresh's (like with 802.11p networks)

Christian Huitema - WiFi has a solution - for WiFi device is always polling so why should a mobile wait for a RA

TimC - in Univ networks have multiple VLANs with multiple subnets - vendors implemented a broadcast type service that made this worse - VLAN pooling fixes this in some cases on WiFi too - Tim suggesting some text to address this similar problem

Tim Winters - are people not including link layer options -

Victor - sports working on this

Jen Linkova - keep both Shoulds

Philip Homburg - not sure what he was suggesting
Steven Barth - should we also make recommendations on interval as well - maybe wording like networks that support mobile devices should increase their interval length to improve battery life
StuartC - Bonjour now does something similar - added capability that says reply with unicast for example when laptops wake up
**Call for Hums showed were all for making this a WG Document
Disposition: Resubmit as WG doc

Friday

Charter discussion

Job 1: Find operators' issues and workarounds

Wes: Change "Ipv4/Ipv6 Internet" to "IPv6 Internet"

Joel: We don't have to publish everything. Sometimes just a draft and related discussion is useful. Allow discussion to occur here even if it doesn't result in publication.

Tore: I have an operational issue that requires a protocol change, but the right WG no longer exists.

Job 2:

Victor: Can we bring in work from NOGs and other groups? I'd like to focus on v6o, but too much of the glue of v6 is in IPv4.

Dave Thaler: Discuss things necessary to facilitate deployment of IPv6. That would include IPv4-related IPv6 issues (or vice versa).

Lorenzo: Remove IPv4.

Dave: "With IPv6 Internet, including operational issues depending on IPv4."

Lorenzo: A v4o topic is off topic. Agree with #2.

Dave Apple: v4 is out of scope. Dual-stack is in scope.

Fred: So maybe remove IPv4, and add a sub-bullet saying, "Issues in the interaction between IPv4 and IPv6"

Some nodding and thumbs up in the room.

Mikail Abrahamasson: What I want to avoid is that any problem I have is out of charter for both v6ops and sunset4.

Fred: If it's an Ipv4-specific problem, sunset4. If IPv6-, v6ops. If both, v6ops.

Marc: If it's in scope for IETF, IETF should find a place to discuss. But I disagree with "if this is IPv4, go to sunset4."

Mikhail: I'm okay with v6ops being v6o, but it means expanding sunset4 charter.

Victor: Operators don't want to run around to multiple working groups. Need a default route between WGs.

JJB: On the cusp of v6o. Doing more with v6 than v4, talking about how to trun down with v4.

Dan York: Yes, keep #1. Lots of BCPs at NOGs, etc.

Job 3: Describe roadmap to turn off IPv4

Wes: This overlaps sunset4 significantly. There's not enough work for two WGs. Is

this work important, and where does it belong? There may be protocol work and interesting discussion, and we don't want to split participants. We're going to try to complete our existing documents, then go dormant (sunset4). That might inform what we do and where we put things. So, yes, we should work on how to make v6o work.

Dave Thaler: Wes answered some of my question. V6ops has been how to get from v4o to dual-stack, and sunset4's job was to go from DS to v6o. This bullet looks like we're going to combine that into one document.

Lorenzo: I don't like this bullet. I think we should do it. We don't need to describe a roadmap, per se, but do bullet #1 with the numerous transition technologies already defined. This is just a special case of #1.

JJB: "when IPv4 is turned off" is now, in some cases. Sooner than 5-10 years some think. If sunset4 is going to close for lack of interest, that concerns me.

Dan York: Yes, we need this kind of document. We've written some of these already; maybe we could summarize. I get questions from people wanting to move to v6o now, need more of that. I think the IETF is the right place to do it; RFCs are canonical and easily consumed.

Mikhail: If sunset4 is going dormant, I think it's a good idea to just merge these. V6ops has energy sunset4 never got it. People here have become interested in turning off v4.

Matt: Brilliant to have 2 WGs and have them always meet jointly. How people think of the problem alters how they enter the process.

Tsia from China Telecom: v6o. I've used the term, but it's confusing, is it access, backbone, or data center? It's better to design several metrics to measure Ipv6 development. Ipv6 penetration, traffic, those metrics would be useful. Better to keep v4 and v6 together for discussion.

Brian Haberman: Terry couldn't be here. What Matt said about joint meetings makes sense. Then if there's protocol work, sunset4 is already here.

Jen Linkova: I'm excited. Need volunteers to help document v6o deployment.

Fred: Okay, we'll send updated draft charter and send to list. Please send suggested wording changes.

Apple Report

Stuart Cheshire

Question: in the NAT64 in MacOS for developers to use, we assign 2001::/64. What should we use? 2001:2:a:bb1e::/64

Kurt Warner? Re 464xlat vs NAT64. If you do NAT64 it comes DNS64. 464xlat doesn't need DNS64, which is important if you want to do DNSSEC local validation.

James Woodyatt: only for A records

[laughter]

Stuart: If they're never running NAT64, there will never be a record synthesized. It's just where in the stack you're lying to the app.

Mikhail [missed it]

Stuart: Yes, if you use literals, we do the synthesis locally. You can do bump in the stack or bump in the API. We do it in the API.

Marc: A few mainstream apps are v4o.

Dave Thaler: IETF NAT64 has been working better than IETF SSID this week. Yes, use the new prefix. 2001:2:: is the benchmarking prefix, ask BMWG if it's okay with them.

Fred: rfc6052 gives a prefix for this purpose. Section 2.1

Lorenzo: 64:f99b? Some things say you can't use that with any private addresses.

Stuart: We clearly want to make sure the software doesn't recognize it as a special address.

Lorenzo: With 464xlat, there's a v6 client as well. It's required. It's a way to provide v4 when you already have v6. Still, I agree NAT64 is better if you can make it stick, and I hope you can.

Stuart: Yes. We want developers to do v6 apps, and we don't think it's hard.

James Woodyatt: Why in 2015 using ios 8.1 I still can't configure it on Tmo or Orange or any European carrier with IPv6 or NAT64. I want to test on a real network.

Stuart: Some carriers will let you do that.

Andrew Yourtchenko: Thanks for making my job easier. I was trying to convince developers to do v6, and one of them said they'd started because of Apple. When you can ensure app has IPv6, you're in a better position. We do v6o+NAT64 as default, and just less than 40% of clients stick to that one. IPv4 native is called "-legacy".

Tore Andreson: Change prefix from /64 to /96. Would let you test IPv4-embedded. See rc6052.

David: It's not the adres of the NAT64 server, it's the NAT64 prefix.

Tore: Assign a prefix, document it in WHOIS.

Joel applauds

Edwin Protego: Using Teredo or 6o4 prefix will mess up enterprises who filter it.

Stuart: These will never be in enterprise network, it's one hop, between your phone and your Mac. Nobody else in the world will ever see these address. Agree-ask IANA to assign an address.

IPv6 Deployment at OTE

Yannis Nikolopoulos

1.3M DSL subs, 380k TV (IP/SAT)

No issues with BRAS

Bugs in BNG and in CPE

Lots of bugs in CPE

Default security policy is "default deny," but will allow users to opt for "default permit"

IPv6 bugs go unnoticed in dual-stack environment

Chose LW4o6, launching Q1-2016

Still nervous that their IPv4 exhaustion may happen before their v6o service

20% of users have IPv6. 10% of traffic.

Host address availability

Lorenzo Colitti

20-30 ppl have read it

Shin Miyakawa: Support. Consider NAPT66 vs NPT66. Prefix NAT might battery.

Lorenzo: Didn't include NPT because host isn't getting a prefix. We'll add that to the doc.

Shin: VMs need more addresses than you say.

Wes: Good, important document. This need was new information to me. Limitations of asking for more on the fly section is vague. Lots of things require waiting for something to happen. Lots of cases where you know ahead of time you'll need more than one. Launching a VM will require on the fly, but enumerate these cases better. I'll send notes to the list. No point debating how many is "enough," we'll never get consensus. Discuss by way of example, but not to limit the problem.

Fred Baker: Need to cut this short—out of time.

Fred Templin: -PD. It's in AERO. Get PD over wireless, assign to virtual interface.

Tore: If each endpoint gets a /64, from a /56, not enough space.

Lorenzo: Those networks aren't ones that need administrator interference. In a home, you can use SLAAC.

Adopt as a WG item? Strong hum yes, silent no.

Disposition: Adopt as WG item

IP/ICMP Translation Algorithm (SIIT) (rfc6145bis)

Xing Li

Where to discuss?

BEHAVE doesn't exist.

Discuss in sunset4? V6ops?

Joel has agreed to let us discuss on v6ops, and then take as an AD-sponsored draft.

Fred had thought to fold the EAM draft into this, but due to the hairpinning question, it should be separate. We'll publish the EAM draft, and this one can refer to it.

Disposition: Discuss, then send as AD-sponsored draft.