IETF 124 – v6ops Working Group

Date: 4 November 2025
Chairs: Nick Buraglio, XiPeng Xiao
Tech Advisor: Ron Bonica

---

1. Chairs' Opening, Note Well & WG Status (Nick Buraglio, XiPeng Xiao)

- Chairs reviewed the Note Well, IETF processes, and policies, including IPR disclosure obligations and treatment of meeting recordings and personal data.
- Participants were reminded of the IETF code of conduct, to treat each other with respect and to contact the Ombudsteam in case of concerns.
- Housekeeping: attendees were asked to sign the blue sheet (QR code / Meetecho) and to join the queue system via mobile before approaching a mic.
- The chairs thanked Ron Bonica, welcomed him formally as the technical advisor to v6ops, and noted his past IESG experience as especially helpful for evaluating work and understanding IESG expectations.
- WG status:
    - v6ops was recently re-chartered, with Ron leading much of that work.
    - Since IETF 123, two drafts were published as RFCs, and three more are in WG Last Call or the RFC Editor queue.
    - The number of attendees has gone up; the chairs view this as part of the overall IETF trend, not a unique anomaly.
- New milestones were highlighted, including work items related to IPv6 deployment guidance and input to Happy Eyeballs.
- The chairs mentioned planned input to the Happy Eyeballs WG about logging aspects; after discussion with Jordi and advice from AD Med, they expect to send input directly to HAPPY for faster processing rather than via a separate v6ops document.

---

2. WG Draft: 7084bis – Basic Requirements for IPv6 CE Routers (Tim Winters)

- Tim presented updates to 7084bis, clarifying behavior for IPv4/IPv6 customer-edge routers and cleaning up accumulated issues from RFC 7084 experience.
- A specific operational issue was discussed where DHCPv6 and SLAAC interactions can lead to implementations wrongly stopping DHCPv6 when certain RA flags change; this can break connectivity even if a valid lease exists.

- Tim proposed text (adapted from existing host-related RFCs) clarifying that if a CE router has a current valid DHCPv6 lease, it must not drop or stop using it merely because the RA flags changed.
- Chairs and participants agreed that it is important to place this guidance in a v6ops document where CE-router implementers will see it, rather than leaving it only in host-focused RFCs.
- Tim reported no remaining open technical issues; only some editorial cleanups (e.g., fixing “obsoletes” metadata and typos) remain.
- Tim requested to move 7084bis to Working Group Last Call (WGLC).
- Poll (WGLC for 7084bis):
    - Question: *“wglc?”*
    - 14 Yes, 0 No, 6 No Opinion (62 present when poll closed).

---

3. WG Draft: IPv6-Mostly Networks – Deployment and Operations Considerations (Jen Linkova)

- Jen recapped the definition of IPv6-mostly networks: environments where IPv6-only, dual-stack, and IPv4-only hosts can coexist on the same L2 segment, with “ready” hosts signalling that they no longer need IPv4 (e.g., via DHCPv4 option 108) so operators can reduce IPv4 usage.
- She noted real deployments, including the IETF Wi-Fi network, her corporate network, and RIPE, where IPv6-mostly is already in use.
- NAT64 / NAT44 pool discussion:
    - Eric Vyncke asked about using the same address pool for both NAT64 and NAT44 and shared experience where the NAT64 pool quickly exhausted while NAT44 was rarely used in practice, because almost all hosts used IPv6-mostly.
    - Discussion concluded that sharing or separating pools is an operational choice, but operators must consider remote ACLs that may need updating if address pools change.
- “Sticky” NAT44 behavior: Warren Kumari clarified that NAT44 should be sticky (i.e., maintain bindings for a connection as long as possible), and that this is not especially complex; the draft text should reflect this.
- Jen described issues seen in deployment:
    - Some firewalls, especially on laptops, dislike atomic fragments, and some NAT64 implementations still behave per RFC6145 instead of RFC7915, causing atomic fragments to be dropped, leading to user-visible breakage.

- UDP/IPv4 with zero checksum:
    - The draft was updated to discuss UDP/IPv4 packets with zero checksum seen in the wild (e.g., some mobile gateways), which broke Wi-Fi calling when translated.

    - Richard Patterson asked whether the text is recommending a specific behavior or merely listing options.
    - Jen clarified it is primarily a warning: implementations can drop such packets or compute a valid checksum, but operators should be aware they exist and may cause trouble in v6-mostly environments.
    - Richard suggested mentioning a third option: translate and send with zero checksum in UDPv6 in tightly controlled domains, noting some equipment today is unhappy with UDPv6 zero checksums, so this may be risky.
- New text was also added on using the well-known prefix (WKP) in enterprise v6-mostly environments when communicating with non-global address space; this ties into the later NAT64 WKP / 6052-update draft.
- Jen emphasized that many issues documented are not unique to IPv6-mostly: they already exist in dual-stack networks but are masked by Happy Eyeballs; in v6-mostly they become more visible when IPv4 is turned off.
- Elliot Lear provided an extensive review; Jen plans another revision to address comments and believes the draft is ready for WGLC afterwards.
- Poll (WGLC for “IPv6-mostly Networks” draft):
    - Question: *“WGLC?”*
    - 23 Yes, 0 No, 4 No Opinion (62 present when poll closed).

---

4. WG-Related (6man) Draft: YANG Data Model for IPv6 Neighbor Discovery (Fan Zhang)

- Fan Zhang (China Telecom) presented the YANG Data Model for IPv6 ND draft, now a 6man document, seeking operational input from v6ops.
- The model adds configuration and management for IPv6 Neighbor Discovery and related functions that are not covered by existing IPv6 unicast routing YANG models.
- Functions covered include address resolution, redirect, proxy ND, enhanced DAD, ND statistics, and other ND-related features.
- The work originally lived in an RTGWG document, alongside an ARP YANG model, but feedback suggested moving the IPv6 ND piece to 6man where ND expertise resides.
- Since the last meeting and mailing list feedback, the module name was aligned with ietf-ipv6-nd, and notes about SEND/ND-proxy were removed because those features were dropped from IPv6 node requirements due to limited use.

- An appendix was added summarizing existing ND-related nodes scattered across different YANG models/RFCs (e.g., from RFC 4861, 4862, and others) to clarify coverage.
- Eric Vyncke asked a technical question about whether certain timers (e.g., reachable time) are per-interface or per-RA, pointing to the need for precise modeling; chairs suggested conferring with 6man experts.
- The authors invited more operational feedback on whether the model matches real-world ND configuration practices.

---

5. Active Individual Drafts:

5.1 RFC6146-bis, RFC7915-bis, SIIT-DC to STD, SIIT-DC-DTM to STD (Jordi Palet Martínez)

- Jordi presented his set of NAT64/SIIT-related drafts:
    - RFC6146-bis (NAT64)
    - RFC7915-bis
    - SIIT-DC to Internet Standard
    - SIIT-DC-DTM to Internet Standard.
- Jordi reported reaching out to the original NAT64-spec authors to ask if they can contribute or help validate updates and maturity.
- He asked the WG about the process to move -bis documents to Internet Standard (IS) and to move some existing STD documents to PS or vice versa.
- Med (Med Boucadair) clarified that the -bis drafts must follow the full process (Call for Adoption, WGLC, IESG evaluation), while advancing the others to full Standard is partly a WG chair judgement call, depending on maturity and deployment evidence. Chairs will ask in mailing list whether there is objection to send documents that has no change to IESG for publication as standards.
- Lorenzo raised concerns about RFC7915's "stateful mode":
    - He could not find any document that clearly explains how this mode works.
    - Without a clear specification, the WG cannot claim interoperability and thus cannot safely advance to full Standard.
- Ted Lemon supported this concern, noting that advancing to full Standard without a concrete description of stateful mode would effectively standardize behavior that is not fully specified and possibly unused.
- The conclusion was that the authors and WG need to dig into mailing list history and existing docs, and possibly create additional text or a separate document, before progressing the status of RFC7915.

5.2 IPv6 Application Testing (Philipp S. Tiesel)

- Philipp introduced a new draft on IPv6 application testing, motivated by increased pressure (e.g., from US government mandate 21-07) to support IPv6-only deployments properly in commercial products.
- The draft proposes a structured testing framework that organizes tests by environment (e.g., enterprise, mobile, cloud, etc.) rather than using a single monolithic test suite.
- It collects common sources of IPv6-related failures and misbehaviors observed in real deployments, aiming to help app developers and operators proactively test and fix them.
- Philipp asked for feedback on missing scenarios, especially from perspectives beyond software vendors and testbed operators.
- Tim Winters supported adoption, saying the draft reflects real-world testing experience and that organizing by deployment environment is a useful structure. He suggested distinguishing more clearly between pure IPv6-only and environments needing additional services (e.g., NTP, certificates) that might require DHCP rather than only SLAAC.
- Tommy Jensen also supported the work and suggested:
    - Broadening the justification beyond the US mandate, mentioning other national mandates and market forces.
    - Framing the document as a reference that app developers can consult to understand *why* IPv6 testing matters, not only *how* to test.
- Philipp agreed to add references to other mandates (e.g., in China, Germany, Czech Republic) and to refine the introductory motivation.
- The chairs asked the room about WG adoption.
- Poll (Adoption for “IPv6 Application Testing” draft):
    - Question: *“adoption?”*
    - 20 Yes, 1 No, 4 No Opinion (72 present when poll closed).

---

6. Discussion Drafts (as time allowed)

6.1 BCP IPv6 Prefix Assignment (Jordi Palet)

- Jordi presented the BCP IPv6 Prefix Assignment draft, which aggregates operator experience and proposes recommendations on prefix sizes and assignment models.
- Jen Linkova questioned the reasoning for recommending certain prefix lengths with normative “MUST” language, suggesting the draft must clearly justify strong requirements.
- Lorenzo argued that any MUST-level requirement should be checked carefully against existing deployment practices; otherwise, it risks being ignored or contradicting real networks.
- Tom Hill noted the draft’s recommendation on persistent prefixes (MUST keep the same prefix) might be problematic in scenarios where operators need to change prefixes (e.g., renumbering, different access technology).

- Nick Buraglio echoed concerns, pointing out that many broadband deployments use unnumbered or link-local-only WAN interfaces, so some recommendations may not reflect the diversity of real-world deployments.
- Ron Bonica suggested the draft enumerate candidate approaches for WAN addressing (e.g., link-local-only, /127, GUA) and provide an ordered preference list; he personally favored:
    1. Link-local-only
    2. /127
    3. GUA.
- Eric Vyncke suggested that the document may be better as Informational rather than BCP, and that using “RECOMMENDED” instead of “MUST” would be more appropriate for deployment guidance where no one will “inspect” ISP compliance.
- Richard Patterson disagreed with the statement that GUA is “most common,” again stressing that many operators use link-local or unnumbered WANs, often historically driven by hardware scale constraints which may or may not still apply today.
- Jordi replied that recommendations were based on surveys of ISPs worldwide, but acknowledged the need to refine wording and possibly document the survey results more explicitly.
- The chairs suggested continuing detailed text changes on the mailing list before considering WG adoption.

6.2 Considerations of Gradual IPv6-only Deployment in 5G Mobile Networks (Chenhao Ma)

- Chenhao (China Telecom) presented a draft on gradual IPv6-only deployment in 5G mobile networks, focusing on 464XLAT/9964-based approaches on the user plane.
- The draft discusses challenges such as roaming UEs, DNS64 configuration, and how to signal DNS64 behavior to UEs.
- One proposal is to use RA options to deliver prefix and DNS64 server information specifically for IPv6-only UEs.
- Lorenzo questioned whether a new DNS64-related option is really needed, given that DNS64 is fundamentally a network function and hosts do not need to know if synthesis is happening.
- Chenhao described a use case where mobile and broadband/wireless users share a common DNS infrastructure, but operators want only IPv6-only mobile users to reach IPv6-only services via DNS64, not all users.
- Lorenzo and Nick Buraglio suggested that this sounds like a network architecture choice, and that operators could:
    - Use RDNSS options selectively, or

    - Use different PDP contexts / APNs and allocation policies (IPv4-only, dual-stack, IPv6-only) to control exposure, rather than creating a new option.
- Jordi noted that in his deployment experience, DNS64 usage and IPv6-only vs dual-stack behavior can be controlled through existing mechanisms without new signaling.
- Chairs recommended Chenhao discuss further with Jordi and others offline to refine the problem statement and determine whether new protocol work is truly needed.

6.3 IPv6 Network Deployment Monitoring and Analysis (Jing Zhao)

- Jing Zhao (China Unicom) presented a draft on IPv6 network deployment monitoring and analysis, describing an end-to-end monitoring architecture to understand why IPv6 deployment may be lagging in some segments.
- The draft identifies several key problems:
    - Monitoring focused mainly on backbone; lack of visibility into access networks, terminals, and applications.
    - Metrics often limited to traffic ratios, without insight into quality, active connections, or element states.
    - Data silos between network domains hinder end-to-end correlation and troubleshooting.
    - Insufficient depth to detect incomplete IPv6 transformation in private apps/CDN paths.
    - Lack of dynamic prediction models that factor in external influences and policy changes.
- The proposed architecture has three layers:
    - Data collection layer aggregating metrics from home broadband, mobile, backbone, and application networks.
    - Intelligent analysis layer, including dynamic traffic attribution and correlation analysis to locate problematic sub-networks or services.
    - Presentation and decision-support components (described in the draft).
- Jing showed a use case (e.g., cloud gaming) where end-to-end flow reconstruction across three network segments (N1, N2, N3) identified abnormal latency in the home segment as the root cause, tracing issues to poor home IPv6 network quality and device/maintenance problems.
- Eric Vyncke asked if the data all comes from a single AS and whether the system uses IPFIX; Jing confirmed it is deployed in their own operator network and leverages such mechanisms.
- Erik Nygren suggested aligning with discussions in the Happy Eyeballs WG about measuring IPv6 “brokenness” and also recommended focusing on residual IPv4 (rather than only IPv6 adoption) since IPv4 remnants often drive CGN pressure and operational cost.

- XiPeng cautioned that ranking different regions' IPv6 deployment is not an IETF responsibility; the draft should focus on monitoring methodology, not on providing deployment scorecards.
- The authors welcomed suggestions and indicated they plan to expand use cases and may demonstrate the system at a future meeting.

6.4 NAT64 WKP / 6052-update – Using the Well-Known Prefix for Non-Global IPv4 (Jen Linkova)

- Jen presented the NAT64 WKP / RFC6052-update draft, motivated by operational issues in IPv6-mostly enterprise networks that use public DNS64 and want to reach non-global IPv4 addresses (e.g., RFC1918) via the well-known prefix (WKP).
- Problem: RFC 6052 Section 3.1 says translators "MUST NOT" use the WKP to represent non-global IPv4 addresses and "MUST drop" such packets; this conflicts with enterprise needs and with v6-mostly deployment scenarios where WKP-based synthesis is convenient.
- Historically, when 6052 was written, DNS64 was assumed to be under local administrative control, making per-prefix behavior easy (different views for global vs private addresses). Today, more deployments rely on public DNS64 plus local prefix discovery, making 6052's restrictions painful.
- The draft proposes relaxing the "must drop" requirement so that WKP may be used safely for private/non-global IPv4, enabling IPv6-only hosts using public DNS64 to reach those endpoints.
- Discussion points:
    - Some participants worried about making currently compliant implementations non-compliant if requirements become stricter or flip.
    - Jen's initial text used "should translate and must not drop"; Ron Bonica pointed out that this wording is confusing.
    - Suresh Krishnan suggested making both sides SHOULD so that existing implementations that still drop such packets can remain compliant while guidance nudges the ecosystem toward translation.
    - Ted Lemon argued that any update should state clearly what is desired and that being fuzzy about normative requirements is not helpful.
- The group agreed that the problem is real and worth solving; exact normative language will be refined on the mailing list.
- Poll (Adoption for 6052-update / NAT64 WKP draft):
    - Question: *"Adoption for 6052-update?"*
    - 7 Yes, 1 No, 7 No Opinion (76 present when poll closed).