V6OPS Working Group Minutes — IETF 126

Date: 24 July 2026
Time: 09:00 – 11:00
Chairs: Nick Buraglio, Xipeng Xiao

Working Group Administration and Status Update

Nick Buraglio reported the following:

- draft-ietf-v6ops-rfc6146-bis completed Working Group Last Call and is being prepared for submission to the IESG as an Internet Standard.
- Working Group Last Call was issued for draft-ietf-v6ops-nat64-wkp-1918.
- draft-ietf-v6ops-ipv6-only was adopted as a working group document.
- Participants were encouraged to contribute to Brian Carpenter's open-source IPv6 textbook.

Jen Linkova suggested holding an interim meeting co-located with RIPE in May. Tom Hill and Tim Wicinski expressed support for exploring the idea.

Operational Presentations

How to Upgrade a SaaS Giant to IPv6

Presenter: Philipp Tiesel, SAP

Philipp described SAP's IPv6 deployment programme. The initial motivations included RFC 1918 address-space pressure across approximately 9,000 Kubernetes clusters and the United States government IPv6-only procurement mandate.

SAP created an internal product standard requiring its portfolio to operate without functional regressions in IPv4-only, dual-stack and IPv6-only environments. IPv6 deployment was incorporated into normal infrastructure and hardware refresh cycles rather than managed through a separate aggressive deadline.

Key observations included:

- IPv6 deployment in a large company is primarily an organisational and coordination challenge.
- IPv6 support differs significantly among cloud providers.
- Most application code worked after runtimes and environments were updated.
- Documentation and examples often remain IPv4-specific.
- Some runtimes have surprising defaults, such as Java preferring IPv4.
- Happy Eyeballs support varies among platforms.

SAP has delivered IPv6 support for relevant US government customers and is progressively extending it to additional products and public-cloud environments.

Draft Presentations

IPv6 Prefix Assignment to End Sites

Presenter: Jordi Palet Martínez
Draft: draft-ietf-v6ops-prefix-to-end-sites

The draft provides recommendations for assigning IPv6 prefixes to end sites and point-to-point links and is broadly aligned with RIPE BCOP-690.

Recent changes included:

- additional PPPoE and IPoE considerations;
- clarification of prefix persistence and privacy;
- removal or relaxation of some uppercase normative language;
- references concerning publication of prefix lengths;
- changing “end users” to “end sites” in the title; and
- additional discussion of temporary addresses and reputation concerns.

Jordi considered the document close to Working Group Last Call but noted that the latest version had received insufficient review.

Next step: Working group participants were asked to read the latest version and provide comments on the mailing list before the chairs consider a Working Group Last Call.

464XLAT/MAP-T Optimization

Presenter: Jordi Palet Martínez
Draft: draft-ietf-v6ops-464xlat-optimization

The draft addresses unnecessary double translation when IPv4-only devices behind a CE router access a local dual-stack CDN or cache in a 464XLAT or MAP-T deployment.

Jordi proposed focusing on the DNS/routing-based approach. The approach would allow traffic to reach a CDN or cache directly within the operator network rather than traversing a NAT64 translator. The availability of the NAT64 well-known prefix for RFC 1918 addresses could make this approach more practical.

Discussion included:

- Ondřej Caletka observed that AAAA records may not be required for the IPv4-only-device case.
- Ondřej suggested that explicit mapping provisioned to CE routers could be a preferable long-term solution.
- Jordi noted the practical difficulty of deploying CE firmware changes.
- Andriy Yurchenko said that existing MAP-T implementations and S46 provisioning mechanisms may already support static mappings.
- Jen Linkova requested evidence from a prototype or operational deployment.
- Éric Vyncke clarified that the proposed path applies to CDN infrastructure located within the ISP network, rather than general Internet destinations.

Jordi also asked whether the document should be split into an Informational document describing the alternatives and a BCP covering the preferred approach.

Next step: The authors will seek feedback from CDN, cache and operator implementers, investigate whether existing provisioning mechanisms are sufficient, and continue the architectural discussion on the mailing list. No decision was made during the session on splitting the document.

IPv6-Only and IPv6-Mostly Terminology

Presenter: Jordi Palet Martínez
Draft: draft-ietf-v6ops-ipv6-only

The first working group version included text restructuring, additional diagrams, further qualification of the part of a system or network being described, and clarification that APIs are generally expected to remain dual-stack for now.

Ondřej Caletka noted that the examples classify both a host using native IPv4 and IPv6 and a CLAT-enabled host generating only IPv6 packets on the network as "dual-stack." He suggested distinguishing the application-facing and network-interface operating modes more precisely.

Tim Chown questioned whether "scope" was the clearest term. "Context" was discussed as a possible alternative, but no terminology decision was reached.

Very few participants indicated that they had read the latest version.

Next step: The working group was asked to review the latest version. The authors will consider how to distinguish host, application, interface and network operating modes and whether a term other than "scope" would be clearer.

IPv6 Application Testing

Presenter: Philipp Tiesel
Draft: draft-ietf-v6ops-ipv6-app-testing

The draft aims to become a BCP describing how to test applications for IPv6-only readiness. It covers conventional applications, cloud and service applications, connectivity scenarios, intermediaries, name resolution, partially broken connectivity and common IPv6-related failures.

Recent changes included:

- a section on name-resolution testing;
- expanded treatment of partially broken connectivity;
- testing without IPv4 loopback addresses;
- considerations for addresses stored or handled as data; and
- replacing "true IPv6-only" with "IPv6-only strict."

Discussion included:

- Tim Winters asked whether performance testing was in scope. Philipp said the draft currently focuses on functional testing because performance problems are more commonly associated with infrastructure deployment than application implementation.
- Tim Winters suggested explicitly testing management or control traffic separately from application data traffic because the two may use different paths.
- A NIC.br participant strongly supported the document and mentioned a related testing tool under development.
- Tim Chown suggested separating guidance for teams constructing test environments from guidance for application developers.
- Tim Chown also mentioned the possible value of making an IPv6-only test environment available as a service.
- Alex suggested testing applications that were initially installed in dual-stack environments and later reconfigured for IPv6-only operation.

Session poll: Is the application-testing draft ready for Working Group Last Call?

- Yes: 7
- No: 1
- No opinion: 13 (total: 108 participants registered)

The poll showed some support but limited participation.

Next step: The authors will consider the comments concerning control and data paths, test-environment guidance, transition from dual stack to IPv6-only operation, and additional common failure scenarios. Further review is needed before the chairs decide whether to start Working Group Last Call.

Filtering Address Records in Stub DNS Resolvers

Presenter: Ondřej Caletka
Draft: draft-ietf-v6ops-aaaa-filtering

The draft addresses hosts and applications that suppress AAAA queries when they believe IPv6 connectivity is unavailable. This behaviour can break access to IPv6 destinations reachable through VPNs or non-default interfaces.

The proposed principle is that a resolver should issue AAAA queries when the host has a route to IPv6 address space that may legitimately appear in DNS. Link-local connectivity alone would not be sufficient, but ULA connectivity could be.

The draft also explains why suppressing unusable queries can sometimes reduce delays caused by packet loss or broken recursive resolvers. It discusses interaction with Happy Eyeballs v3 and multi-question DNS mechanisms.

Discussion included:

- Jen Linkova cautioned that resolving through one interface and sending traffic through another can produce unexpected behaviour on multihomed hosts. She suggested considering Provisioning Domains.
- Nick Buraglio confirmed that VPN-related DNS failures are a real operational problem.
- A Cloudflare participant requested clarification of what address space should trigger AAAA queries.
- Ondřej clarified that ULA routes may be relevant because ULA addresses can appear in globally accessible DNS, while link-local addresses should not trigger such queries.

Next step: The author will continue aligning the proposal with Happy Eyeballs v3 and will consider whether VPN and Provisioning Domain behaviour needs more explicit treatment. The document was not considered ready for Working Group Last Call.

IPv6 CE Router Requirements

Presenter: Tim Winters
Draft: draft-ietf-v6ops-rfc7084-bis

Tim presented updates to the IPv6 CE router requirements, including:

- incorporating the LAN-side prefix-delegation requirements from RFC 9018 and obsoleting that RFC;
- clarifying SLAAC and RA-Guard requirements;
- expanding the explanation of why CE routers should generate ULAs;
- adding requirements concerning DNS proxying, DNS over HTTPS and DNS over TLS;
- preventing temporary M/O-flag changes from causing a CE router to discard active DHCPv6 state;
- adding DNS packet-size and fragmentation-avoidance requirements; and
- updating references, including the DHCPv6 specification.

Discussion included:

- Stefan Ubbink requested that the requirement for a DNS proxy to support IPv6 be stated as a normative “MUST.” Tim agreed.
- Tim Chown asked whether multicast requirements should be expanded. Tim said multicast behaviour at the residential edge remains operator-specific and should not delay the document.
- A participant noted that the DNS-over-HTTPS text should make clear whether the requirement applies only when DNS proxy functionality is implemented. Tim agreed to clarify the wording.

Next step: The authors will apply the agreed DNS-proxy wording corrections and other editorial updates. Tim indicated that the document should then be ready for Working Group Last Call.

Enhanced Dual Stack

Presenter: Xipeng Xiao
Draft: draft-xiao-v6ops-eds

Shuping presented Enhanced Dual Stack, which aims to reduce the operational risk of enterprise IPv6 deployment by allowing systems to prefer IPv6 when it performs adequately and to fall back to IPv4 when IPv6 fails or performs poorly.

The proposal contains three complementary elements:

- an OS or platform Happy Eyeballs implementation profile;
- a compatibility path for existing applications using APIs such as getaddrinfo() and connect(); and
- selective reporting of IPv6 fallback and failure information for operational diagnosis.

The draft also proposes incremental deployment, beginning with early-adopter hosts while retaining IPv4-only operation for critical hosts until sufficient confidence has been established.

Discussion included:

- Philipp Tiesel said the draft combined several distinct mechanisms that might be better considered separately. He strongly questioned changing the behaviour expected from existing socket APIs and considered adoption by OS implementers unlikely.
- Tom Hill warned that the proposal could reinforce the perception that IPv6 is inherently difficult and could encourage enterprises to treat IPv6 deployment as a minimal compliance exercise.
- Xipeng argued that reducing deployment difficulty could improve motivation, particularly for small and medium-sized enterprises.
- Philipp said that the hardest part in large enterprises is often deciding to begin, rather than the technical deployment itself.
- Warren Kumari observed that many small organisations lack IPv6 expertise and that reducing operational pitfalls could still be valuable.
- Jen Linkova argued that enterprises routinely operate other complex technologies and often avoid IPv6 primarily because their existing networks continue to work without it.

Next step: The authors requested technical review, especially from people with enterprise IPv6 deployment or software-development experience. The proposed mechanisms and whether they should remain in one document require further discussion. No adoption request or working group decision was made.

Unicast Transmission of Unsolicited Router Advertisements

Presenter: Jen Linkova

Jen described a proposal to improve the reliability of unsolicited Router Advertisements, particularly over Wi-Fi during renumbering. Multicast RAs can be

lost or delayed, so a router could additionally send copies to individual hosts using L3 unicast or L2 unicast delivery.

Discussion included:

- Andriy Yurchenko questioned some L3-unicast implications and suggested that L2-unicast transmission of an L3-multicast packet may be cleaner.
- Ondřej Caletka considered one packet per client a reasonable cost when important RA information changes.
- Tim Winters noted that CE router implementations often keep RA generation separate from neighbour-cache processing, making per-neighbour transmission difficult.
- Éric Vyncke suggested sending the normal multicast RA as well.
- Lorenzo Colitti argued that the preferable long-term solution may be an IEEE 802.11 requirement for access points to convert critical multicast control packets to unicast.
- Tim Chown noted DHCPv6 Reconfigure as another mechanism for notifying hosts of changes.

Next step: Jen and Tim Winters agreed to continue discussing implementation considerations. Participants also suggested engaging IEEE 802.11 experts. The work is associated with 6MAN, and no V6OPS working group decision was requested.

Observations on the Reachability and Evasion of Packets with IPv6 Extension Headers

Presenter: Le Gai

The presentation reported measurements of IPv6 packets containing extension headers, including both reachability and the possibility of bypassing firewall access controls.

The study observed comparatively good reachability for some Destination Options and Fragment Header cases, while Hop-by-Hop Options experienced substantially more filtering. It also identified networks in which filtering devices inspected only the first header and therefore failed to apply policy to the encapsulated transport header.

Discussion included:

- Éric Vyncke recommended comparing the results with previous studies and RFC 9511.
- Éric noted that a single geographic vantage point limits the conclusions and recommended using vantage points on multiple continents.
- He also warned that SSH and SNMP probes may introduce bias because those services are intentionally filtered in many networks.
- Justin noted that the study used relatively small extension headers and therefore measured a best-case scenario; larger headers may have lower reachability.

Next step: The researchers agreed to add vantage points and compare the work with existing research.

Considerations for IPv6-Only Deployment in 5G Mobile Networks

Presenter: Chenhao Ma
Draft: draft-ma-v6ops-ipv6-only-deployment-in-5g

The draft describes operational considerations for IPv6-only deployment in 5G networks and aims to bridge gaps between existing standards and operator deployment practice.

The update added:

- introductory material on 5G concepts;
- discussion of PCO-based and RA-based mechanisms;
- testing isolation;
- organisational and operational considerations; and
- clarification of IPv4 service support and CLAT availability.

The presentation compared 3GPP Protocol Configuration Options with Router Advertisement mechanisms such as PREF64. The authors preferred the RA-based approach because IP-layer information would remain in the IP layer, the mechanism would be access-independent, and changes to the NAT64 prefix could be propagated dynamically.

Discussion included:

- Jordi Palet Martínez supported continued work and suggested coordination or liaison with 3GPP concerning PCO and RA behaviour.
- Jen Linkova welcomed the focus on PREF64 and encouraged operators to request PREF64 support from equipment vendors.

Next step: The authors will continue refining the deployment guidance and seek further operator and implementation feedback. The draft remains an individual Informational document; no adoption decision was made.