

AI Preferences

Zürich

Q3 2025

Note Well

A reminder of IETF policies

This is a reminder of IETF policies in effect on various topics such as patents or code of conduct. It is only meant to point you in the right direction. Exceptions may apply. The IETF's patent policy and the definition of an IETF "contribution" and "participation" are set forth in BCP 79; please read it carefully.

As a reminder:

- By participating in the IETF, you agree to follow IETF processes and policies.
- If you are aware that any IETF contribution is covered by patents or patent applications that are owned or controlled by you or your sponsor, you must disclose that fact, or not participate in the discussion.
- As a participant in or attendee to any IETF activity you acknowledge that written, audio, video, and photographic records of meetings may be made public.
- Personal information that you provide to IETF will be handled in accordance with the IETF Privacy Statement.
- As a participant or attendee, you agree to work respectfully with other participants; please contact the ombudsteam (<https://www.ietf.org/contact/ombudsteam/>) if you have questions or concerns about this.

Definitive information is in the documents listed below and other IETF BCPs. For advice, please talk to WG chairs or ADs:

- [BCP 9](#) (Internet Standards Process)
- [BCP 25](#) (Working Group processes)
- [BCP 25](#) (Anti-Harassment Procedures)
- [BCP 54](#) (Code of Conduct)
- [BCP 78](#) (Copyright)
- [BCP 79](#) (Patents, Participation)
- <https://www.ietf.org/privacy-policy/> (Privacy Policy)

Logistics

- Local facilities
 - Toilets
 - Network: GoogleGuest
- We have **remote participants**
 - Use MeetEcho to queue
 - Speak clearly and loudly, one at a time, use the mic *closely*
 - Identify yourself

Scribe(s)

Agenda

1. Overview of where we're at
2. Discussion of potential paths forward
3. Depending on (2), issue discussion

Overview

- After Brussels, we appeared to have convergence and strong momentum towards shipping these documents on schedule.
- Since then, we have had many new participants join us.
- It was not apparent then how strong the disagreements in the group were, but now it is - WGLC has failed.

- Our task for this meeting is to identify the path forward with the best chance of success.
- Success is defined by our charter:
"standardize building blocks that allow for the expression of preferences about how content is collected and processed for Artificial Intelligence (AI) model development, deployment, and use."

There are distinct aspects of this work that seem to have different levels of agreement (these are impressions only):

1. The "ai training" preference -- seems to be possible to get to consensus, after considering definitional issues and the search carve-out
2. The "use / inference" preference -- little consensus currently, but discussion about 'purpose' may lead to something better
3. The "top-level / umbrella" preference -- little consensus currently
4. The overall model and the context around it (e.g. Section 3.2) -- open questions

- We started this work because it is difficult-to-impossible to express any preference regarding AI and content in a standard, widely recognized manner.
- In that regard, any improvement over the status quo is significant progress towards our chartered goal.
- Getting to consensus means that we need a level of support and acceptance of the outcomes. That means working through objections, and is likely to take considerable time.
- In Brussels, we seemed to be on track to produce our deliverables by their milestones. From recent discussions, however, it's not clear that we still have that sense of urgency driving us.

With all of that in mind, we'd like to discuss the following questions:

1. Is our charter still achievable? If not, can we come to consensus on the changes that are necessary to make it so?
2. Is it urgent that we deliver some -- potentially minimal -- capability to express a preference soon (e.g., in the next few months)? Or has that urgency now passed?
3. How can we stage our deliverables? E.g., can we deliver the "ai training" preference on its own? Or does it need to be delivered at the same time as the top-level and/or use preference?