

Update on (some) Standards

Bellagio

Mark Nottingham, April 2025

- AI Preferences
- Web Bot Auth
- Pay-Per-*
- Agent2Agent
- Other Thoughts

AI Preferences

Context

AI Preferences

- IETF Working Group, chartered January 2025
 - “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.”
- AI Vendors, Content, Civil Society
- Referenced in the EU AI Code of Practice, mid-2025
- Vocabulary
- Attachment

<https://ietf-wg-aipref.github.io/>

Vocabulary

AI Preferences

- **ai-training**

The act of using an asset in the production or refinement of an AI model that can generate content in one or more modalities (text, image, audio, etc...).

- **search**

Use of an asset in an application where the primary purpose of the application is to select assets and direct users to the location of those assets.

This category of use only applies under the following conditions:

- Where the presentation of an asset in search output -- if selected for presentation -- includes a direct reference or link to the original location from which the asset was retrieved.
- When excerpts from the asset are displayed they serve to assist users in evaluating the relevance of the result.
- This category does not include the use of assets to generate summaries.

- **ai-use / ai-inference ... ?**

Attachment

AI Preferences

- Initial focus is on robots.txt + HTTP headers

Content-Usage: train-ai=y

- Other attachment mechanisms possible
 - Especially in-content: IPTC, C2PA, ...
- Some discussion about registries, but out of scope for us

Challenges

AI Preferences

- Strong divergence of interests
 - **AI vendors** don't want future products constrained
 - **Content producers** want to tightly control how their products are used
 - **Civil society actors** don't want the Open Web enclosed
- Advisory only
 - **No enforcement** - left to legal or additional technical means
- Crawl/Inference gap
- **Training** horse has arguably bolted - "We have enough tokens"
- **Use** has massive scope creep - AI is an enabling technology

Status

AI Preferences

- It's been rough
- Productive meeting two weeks ago in Toronto
- If the proposals from there hold, we may achieve consensus on training and search terms.
- Use term is still an open question - but there is cautious engagement
- We don't want to be doing this in 2027

Web Bot Auth

Context

Web Bot Auth

- IETF Working Group, chartered October 2025
- Bots are currently identified using:
 - **User-Agent** (useless)
 - **IP Address** (problematic, favours big bots)
 - Custom auth mechanisms (not interoperable, requires foreknowledge)
- **Cryptographic identity for bots** can increase confidence, flexibility
- Use cases are primarily around crawling
- Some discussion of identifying agents
 - NOT about the user behind the agent

<https://datatracker.ietf.org/wg/webbotauth/about/>

Proposals

Web Bot Auth

- **HTTP Message Signatures** (Meunier)
 - Leverages RFC 9421 (2024)
 - One bot, one key
 - Key discovery is an exercise left for follow-on work
- **Anonymous Bot Authentication** (EKR)
 - Based upon PrivacyPass / ARC (draft)
 - Attestor role
 - unable to identify specific bots within anonymity set

Challenges

Web Bot Auth

- **Malicious actors** will still try to appear human
- Requiring authentication encloses the Web, risks centralization
- **Bot anonymity can be good or bad**
 - Embedding 'good' / 'bad' into architecture is... difficult
- **Complexity** - lots of cryptography
- **Key management** - always hard

Status

Web Bot Auth

- Very early days
 - Group still aligning on use cases
- Current focus is primarily on crawl use cases
- Anonymous bot auth may be workable, but lots to do

Pay-per-*

Context

Pay-per-*

- Many content folks want to replace advertising revenue streams with payment protocols
 - Micropayments, subscriptions, pay-per-click, pay-per-crawl
- You need a protocol for that
- TimBL defined the 402 status code with no semantics
- Lots of people are trying to fill that gap to create new markets

Challenges

Pay-per-*

- **Pricing**

- Not at all clear what's a good model. Many want a replacement for advertising revenue. Uncertainty about value measurement.
- Over-fixation on pay-per-click/crawl?

- **Equity**

- Advertising may be problematic, but is also very **progressive**. Primary concern for consumer cases, less so for business.

- **Trust**

- Browsers give publishers some confidence about how their content will be used. No such guarantees for agents.

- **Payment**

- Is already a hard problem - double spend, chargebacks (or not), KYC, etc.

Status

Pay-per-*

- Lots of semi-private, unilateral proposals
 - x402 (Coinbase et al)
 - Payment auth scheme (Stripe, Tempo)
 - Club Good Protocol (paywalls.net)
- No real standards activity yet

Agent2Agent

Summary

agent2agent

- Lots of heat, little light
- Discussion of
 - Discovery
 - "New" transports
 - Intent
 - Guardrails
 - Identity

Other Thoughts

User needs
come before the needs of web page authors,
which come before the needs of user agent implementors,
which come before the needs of specification writers,
which come before theoretical purity.

Web Platform Design Principles,
Priority of Constituencies
<https://www.w3.org/TR/design-principles/>

Openness

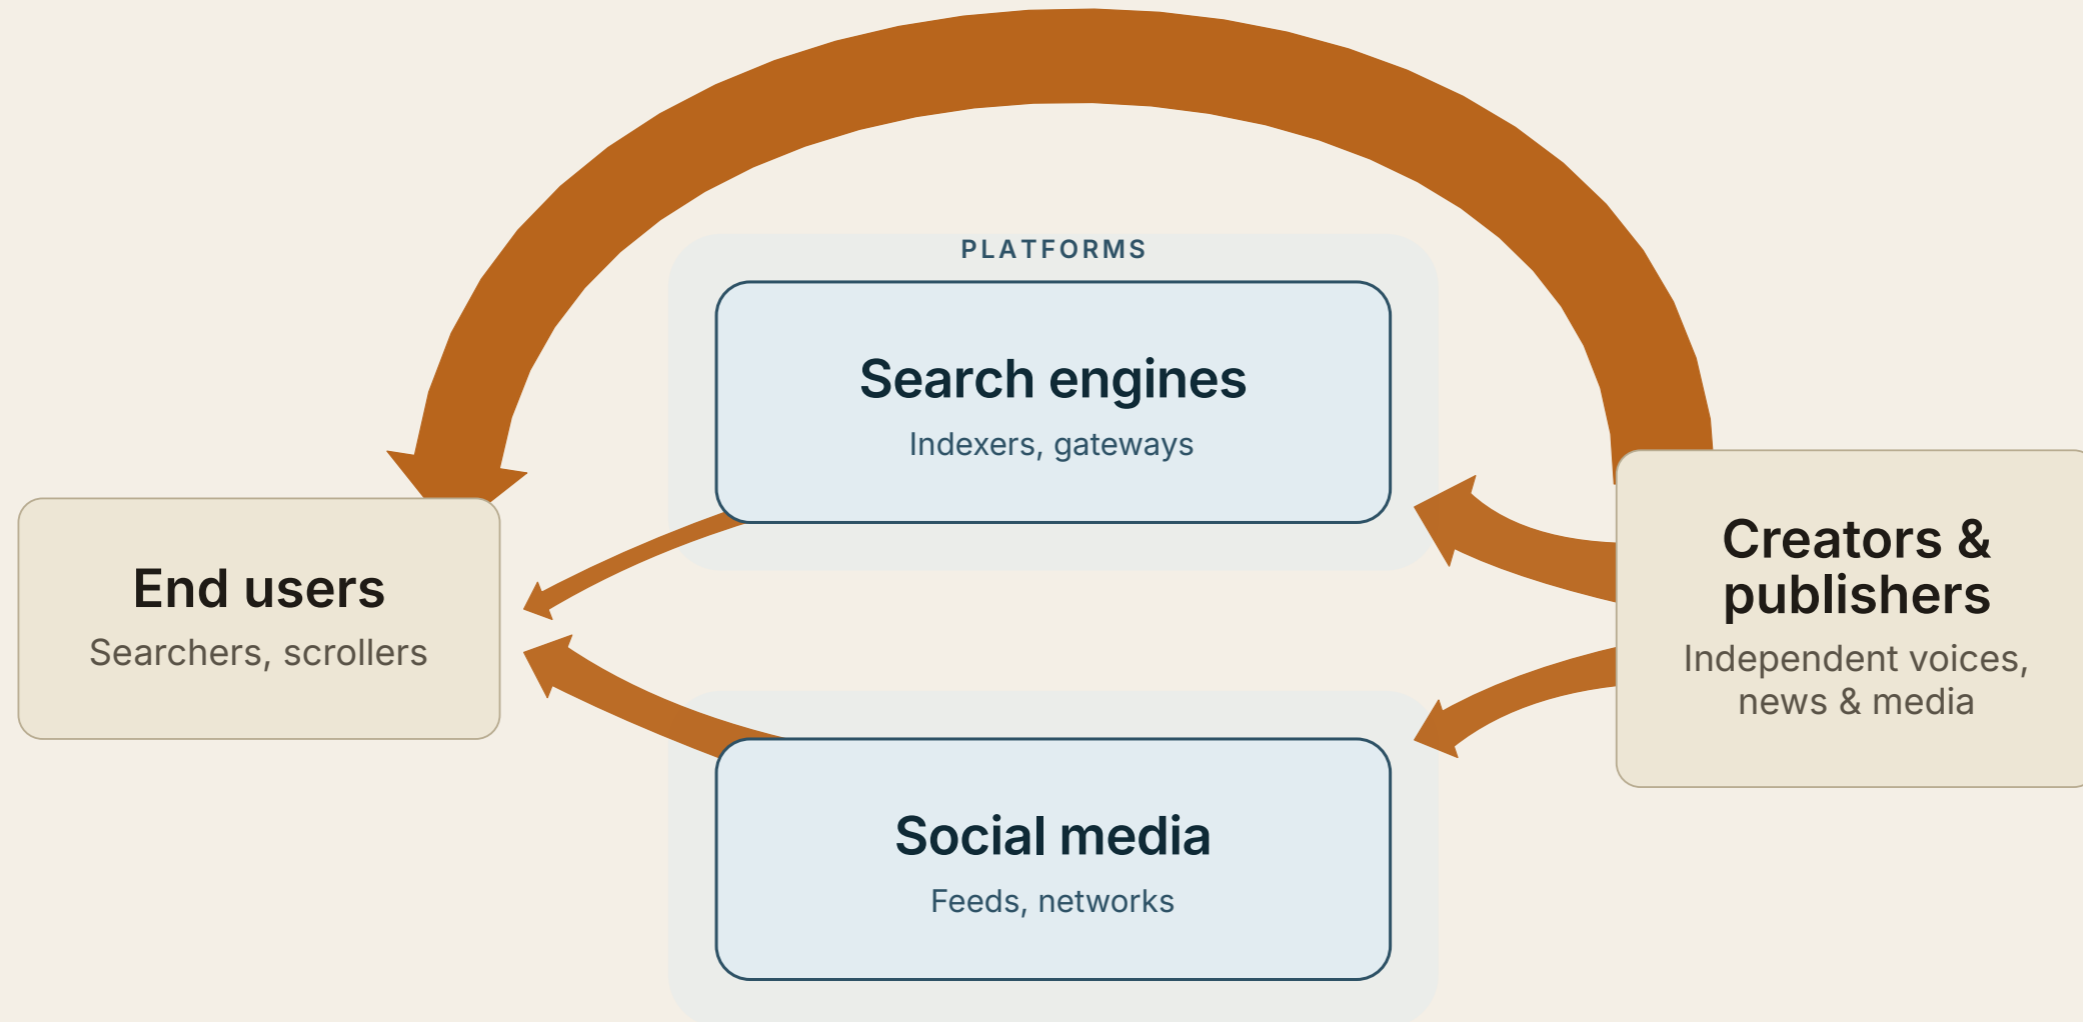
as a system property

- **Openness** is a key attribute of Internet and Web infrastructure
 - Ensures low entry costs ("permissionless innovation")
 - Helps avoid centralisation
 - Closed systems necessitate governance
- Designing open systems brings unique challenges
 - Adapting closed systems to be open is even harder
- The **Open Web** has changed - will the replacement be better?

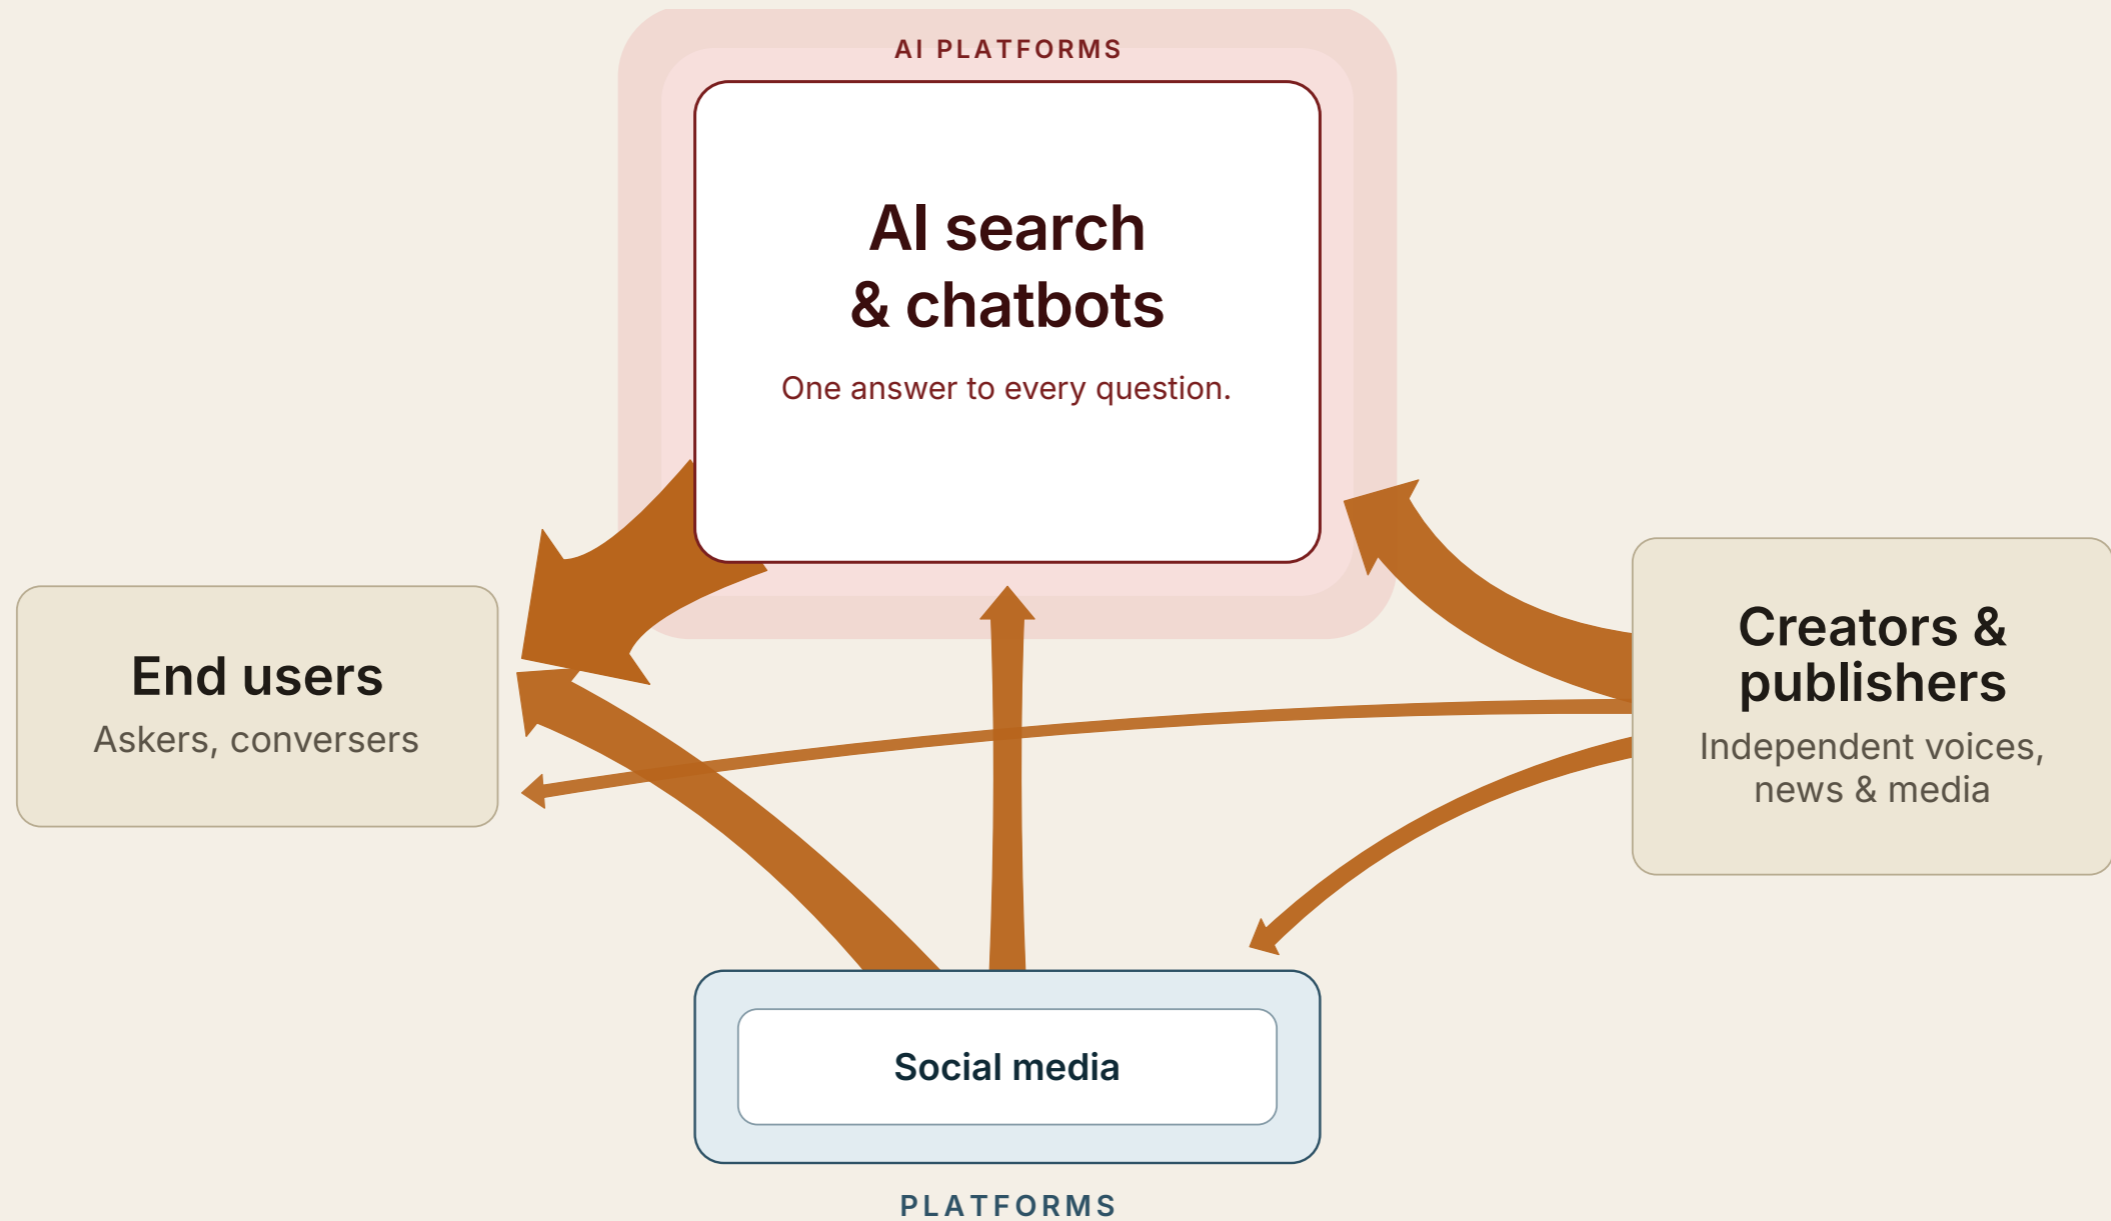
https://mnot.net/blog/2026/open_systems

https://mnot.net/blog/2026/open_web

The platform information economy



The AI information economy?



Centralisation

Infrastructure for Humanity's Access to Information


- Access to and control over information flows can lead to:
 - Reduced incentives for creators and publishers
 - Decreased infrastructure **resilience**
 - Loss of **provenance** and **branding**
 - Opportunity for **tracking** and **profiling** (digital twins)
 - Ability to **reshape perception** of the world
 - ... leading to geopolitical tension and likely further **fragmentation**
- ... and **abuses are deniable** thanks AI's opacity

<https://www.rfc-editor.org/rfc/rfc9518.html>

Promoting Incentive Alignment

Some Strategies to Establish Trust

- “I have good intention / We’re the good guys”
- “The market will punish us if we abuse our position”
- “Sue us if we abuse our position”
- “\$Authority will punish us if we abuse our position”
- **“We’re not in a position to abuse our position”**



Crypto

(no, not that)

- Zero Knowledge Proofs (ZKPs)
- Merkle Trees
- Multi-Party Computation (MPC)
- Mix Nets
- Private Information Retrieval (PIR)
- ...

Partitioning

Not Just for Privacy

- Protocols often create roles to separate concerns
 - MASQUE
 - Oblivious HTTP
 - PrivacyPass
- RFC 9614 - Partitioning as an Architecture for Privacy
- "How do I compose an ecosystem where actors' incentives are aligned?"
 - Strong links to economic perspectives expressed so far

User Agency

Possibly the best part of the Web

- Web browsers are “User Agents” - they represent user interests
 - Privacy, security, sandboxing, predictable operation
- But they *also* represent the interests of Web sites
 - Predictable display of content, ad presentation*
- Both sides' interests are balanced
 - Browser is effectively an intermediary / two-sided platform
 - **Allows a market to form**
- **Definition venue** should be intentional - consider mobile vs web

https://mnot.net/blog/2026/agents_as_collective_bargains