OPEN _FUTURE

# DIVERGENT MECHANISMS, ELUSIVE VOCABULARIES

*The state of machine-readable rights reservations*

OPEN FUTURE POLICY BRIEF #10

Author: Paul Keller

15 JANUARY 2026

# INTRODUCTION

This policy brief provides an overview of the state of machine-readable rights reservations, TDM opt-outs, and AI preference signalling at the end of 2025[1]. It follows our March 2025 policy brief proposing A Vocabulary for Opting Out from AI Training and Other Forms of TDM.

The brief takes a high-level view of the evolving landscape of preference-signalling mechanisms and standards, focusing on vocabularies used to express preferences about specific types of use of protected works. It does not provide an in-depth technical analysis of the individual standards or mechanisms discussed. The document consists of three parts. The first outlines key developments in the field over the past year. The second examines existing approaches to defining opt-out and preference vocabularies and analyzes their relationships. The final section provides an outlook and formulates recommendations in light of the European Commission's consultation on protocols for reserving rights from text and data mining under the AI Act and the GPAI Code of Practice.

At a high level, this paper argues that as machine-readable opt-outs, licensing signals, and compliance mechanisms proliferate, the absence of a shared vocabulary is becoming a structural risk for the information ecosystem. Without a common semantic layer, rightholders, users, and AI developers are pushed toward fragmented, mechanism-specific solutions that are difficult to interpret, prone to misapplication, and susceptible to over- or under-enforcement. The paper acknowledges that a shared vocabulary does not resolve underlying legal or economic tensions, but argues that it can provide common ground for balancing the interests of content producers, users, and AI developers across divergent technical and regulatory approaches. The design principles in the final section aim to clarify what such a vocabulary must—and must not—do to play that role.

Structurally, this brief assumes a separation between three layers: a shared vocabulary that defines concepts of use, a number of complementary attachment mechanisms (referred to as "opt-out solutions" in the remainder of this paper) through which preferences are expressed, and enforcement mechanisms, which depend on the legal context in which those attachment mechanisms operate.

[1] This document uses these three terms interchangeably. The terms machine-readable rights reservations and TDM opt-outs both refer to mechanisms used by rightholders to express opt-outs from the general Text and Data Mining exception in Article 4 of the 2019 Copyright in the Digital Single Market Directive. The term AI preferences — coined in the context of the eponymous IETF working group — conceptually overlaps with opt-outs / rights reservations but expands from rightsholders to contentholders and acknowledges the fact that globally there is no default state when it comes to the applicability of copyright to various types of uses of content in the context of AI training and use. It will depend on the specific legal context if expressed preferences have legal effect or not.

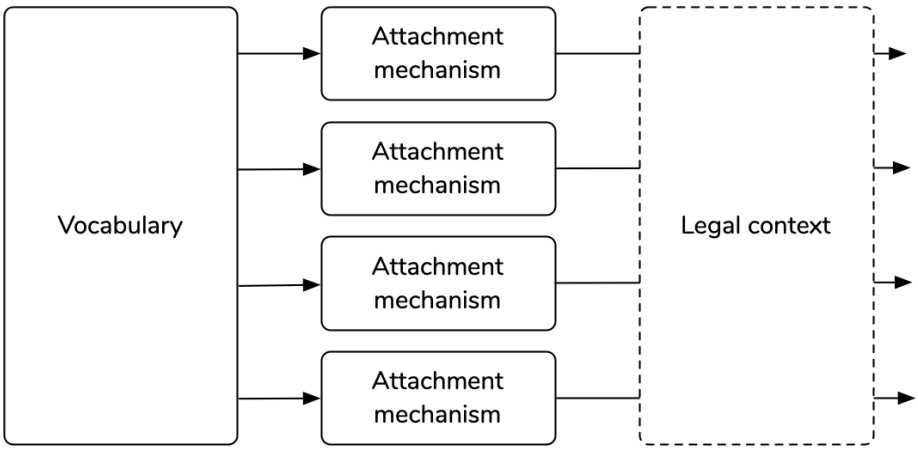


Figure 1: Separation between semantic definition (vocabulary), technical expression (attachment mechanisms), and legally contingent enforcement outcomes.

## DIVERGENCE INSTEAD OF CONVERGENCE

2025 was widely expected to bring greater clarity on machine-readable rights reservations, AI preferences, and AI opt-outs. In practice, these expectations converged around two parallel processes:

In early January, the IETF approved the charter for the AI Preferences Working Group[2]. The approved charter includes two particularly relevant work items:

> *A standards-track document defining a vocabulary for expressing AI-related preferences, independent of how those preferences are associated with content.*
>
> *Standards-track document(s) describing mechanisms for attaching or associating those preferences with content in IETF-defined protocols and formats, including—but not limited to—the use of Well-Known URIs (RFC 8615), such as the Robots Exclusion Protocol (RFC 9309), and HTTP response header fields.*

At the same time, discussions were ongoing in the European Union on the voluntary Code of Practice for providers of general-purpose AI (GPAI) models. The final version of the Code was anticipated to include language on how GPAI model providers could demonstrate compliance with Article 53(1)(c) of the AI Act, which requires them to "put in place a policy to comply with Union law on copyright and related rights, and in particular to identify and comply with, including through state-of-the-art technologies, a reservation of rights expressed pursuant to Article 4(3) of [the CDSM] Directive".

These two processes, while independent from each other, were expected to inform one another. The original target date for the IETF AI Preferences Working Group to deliver its output—August 2025—was formulated against the backdrop of the EU AI Act's timeline, under which the

[2] The author is one of the editors of the vocabulary draft discussed by the working group.

provisions for general-purpose AI models, including the obligation in Article 53(1)(c) to comply with machine readable rights reservations, would become applicable at the beginning of August 2025. At the same time, the GPAI Code of Practice was widely expected by stakeholders to incorporate or reference emerging preference-signalling standards developed by the IETF working group.

Rather than converging on a shared technical and conceptual framework, these processes exposed divergent assumptions about scope, legal grounding, and the intended role of preference-signaling mechanisms. These divergences became most visible in the work of the IETF AI Preferences Working Group.

## *IETF AI preferences working group*

In early March, against this background, we published Policy Brief #8 proposing A Vocabulary for Opting Out from AI Training and Other Forms of TDM. In May 2025, the IETF AI Preferences Working Group adopted the proposed vocabulary as the working-group draft for its vocabulary work item[3].

The AI Preferences Working Group initially appeared to be on track to deliver both the vocabulary and the related—but distinct—mechanism for attaching preferences to content. However, progress on the vocabulary draft stalled in the aftermath of a July interim meeting in London[4]. Subsequent discussions on the mailing list exposed fundamental disagreements among stakeholder groups.

While a full analysis of the dynamics within the working group is beyond the scope of this paper, three sets of issues can be identified that have so far stood in the way of reaching consensus on the vocabulary draft:

1. The first set of issues concerned proposals to expand the scope of the vocabulary beyond AI training. While the original draft was limited to training-related use cases, publishers and other rightsholders advocated for the inclusion of additional categories that would allow the expression of preferences related to inference-time uses of content by AI models and systems—the use of content as input to already trained and deployed models.

   The primary concern articulated by these stakeholders was a perceived need to opt out of substitutive uses of content, such as the generation of summaries. Proposals to expand the vocabulary in this direction were met with opposition from AI model developers and also raised concerns among stakeholders representing public interest perspectives. These stakeholders cautioned that expressing preferences relating to the use of content by deployed AI systems would have implications for user rights, including freedom of

[3] Adoption as a working-group draft indicates that the document was selected as the starting point for the group's work on this topic. It does not imply consensus on its content, and drafts are generally expected to evolve substantially through discussion, revision, and the incorporation of alternative approaches and stakeholder input.

[4] For more details on the state of discussion and the main points of divergence see this unofficial progress update from August 2025.

expression, and could enable content owners to assert control over the use of publicly available content in ways that go beyond the scope traditionally afforded by copyright law.

2. Closely related to this, stakeholders representing online publishers sought the inclusion of a category that would allow them to remain indexed by search engines while opting out of AI-related uses of their content. This demand was articulated in response to the status quo, in which a dominant search provider does not offer a mechanism to opt out of AI-generated summaries—perceived by publishers as substitutive—without also opting out of inclusion in the search index altogether.

   While this request seems straightforward in principle, defining such a category proved difficult in practice. AI model developers that also operate search engines argued that what is commonly understood as "traditional search" is now so reliant on AI technologies that it would be infeasible to provide search functionality in cases where content owners opt out of AI-related uses, including training.

3. Finally, the working group was divided over the desirability of an overarching category that could function as a catch-all. The original draft included such a category, aligned with the EU Copyright Directive's definition of Text and Data Mining.

   One group of participants criticized this approach as overly broad, arguing that a catch-all category would impose wide-ranging restrictions on legitimate uses of content and could have negative implications for innovation and user rights more broadly. Another group argued that a catch-all opt-out was necessary for rightholders to retain meaningful control over their content, pointing to the rapid and unexpected emergence of generative AI systems, which in their view had made it difficult or impossible to opt out "before the damage was done".

   In addition, some participants emphasized that a category covering the concept of TDM was necessary for the vocabulary to be internally coherent and comprehensive, given that broad TDM reservations are already expressed through existing legal and technical mechanisms.

This last set of issues also points to an additional source of divergence within the working group: Some participants urged caution, expressing concern that technical standards developed by the working group could, over time, acquire legal force if legislators were to mandate their adoption as part of regulatory compliance frameworks.

This line of argument gained traction after the European Commission published the GPAI Code of Practice in July 2025, which included the following commitment:

> *Signatories commit: […] to employ web crawlers that read and follow instructions expressed in accordance with the Robots Exclusion Protocol (robots.txt), as specified in Internet Engineering Task Force (IETF) Request for Comments No. 9309, and any subsequent version of this Protocol for which the IETF demonstrates that it is*

*technically feasible and implementable by AI providers and content providers, including rightholders.*

The reference to "any subsequent version" of RFC 9309 was interpreted by some participants as an indication that the European Union might seek to mandate broad compliance with standards emerging from the IETF process. These participants argued that this could amount to an attempt to extend a specific regulatory approach beyond the EU by leveraging a global technical standardization process.

While this interpretation is not supported by the Code of Practice text—and, at a minimum, the vocabulary draft under discussion does not qualify as a "subsequent version" of RFC 9309—the perceived linkage between the Code of Practice and the work of the AI Preferences Working Group nevertheless further complicated the group's internal dynamics.

By year's end, after two contentious meetings in Zurich and Montreal failed to produce consensus around several core concepts[5], the IETF working group chairs decided to remove a number of the more controversial categories from the vocabulary draft. As a result, the current version includes only two categories: one relating to foundation model training, and one intended as an approximation of "traditional search".

Neither category reflects a settled consensus within the working group. Rather, their inclusion represents an attempt to narrow the scope of the draft in order to allow work to continue in the absence of agreement on more expansive or contentious elements.

## *Pay per crawl to the rescue?*

As progress within the IETF working group slowed, alternative approaches to signaling AI preferences began to emerge. In late September, Cloudflare announced contentsignals.org, a proposed extension of the Robots Exclusion Protocol. It allows site owners to communicate preferences using three predefined "content signals": search, ai-input, and ai-train. Cloudflare presents this as its own implementation of a mechanism for allowing website publishers to declare how automated systems should use their content.

This initiative builds on Cloudflare's earlier introduction of bot management and AI crawler controls, which already allow site operators to selectively permit or block automated access at the network level. By combining technical enforcement with a vendor-controlled[6] signaling vocabulary deployed at scale, Cloudflare's approach is positioned to become a de facto standard. This is largely due to structural asymmetries in which infrastructure providers are able to shape preference-signaling practices more rapidly and decisively than standards bodies or individual publishers. Cloudflare has also introduced a pay-per-crawl mechanism, currently in private beta,

[5] For more details on the evolution of the vocabulary during the Zürich meeting and the main points of contention see this unofficial progress update from October 2025 and this blog post by Martin Thomson (co-editor of the vocabulary draft)

[6] The term "vendor-controlled" is used here to indicate that in its current form Content Signals is fully controlled by Cloudflare. Cloudflare has released the Content Signals policy under CC0 to ensure that anyone can implement and freely use it.

which would allow publishers to condition automated access not only on declared preferences but also on remuneration. This further reinforces the role of infrastructure providers as gatekeepers capable of coupling signaling, enforcement, and monetization within a single technical stack.

Earlier in the summer, the first version of the Really Simple Licensing (RSL) protocol was published by the RSL collective. RSL positions the protocol as "a standardized XML vocabulary and associated discovery and authorization mechanisms for expressing machine-readable usage, licensing, payment, and legal terms that govern how digital assets may be accessed or licensed by AI systems and automated agents."

The initial version, published in July, included a limited list of permitted uses based on what it described as "a superset of the proposed IETF AI Usage Preference," drawing on the 02 version of the IETF vocabulary draft. In a subsequent version released in September, the explicit reference to the IETF vocabulary was removed, and the list of permitted uses was revised to align with the terminology used in Cloudflare's Content Signals, with the addition of a catch-all category.

The current 1.0 version of the RSL specification, published in November 2025, further expands the list of permitted uses and includes the following note:

> *This attribute includes the Cloudflare Content Signals vocabulary and MAY include terms from additional standardized vocabularies as they become available (e.g., IETF AI Preferences).*

What is particularly notable in this context is the way in which references to the IETF vocabulary are progressively replaced by Cloudflare's proprietary Content Signals vocabulary. This shift is underscored by the fact that Cloudflare—alongside other major CDN providers—is listed as a supporter of the otherwise largely US-publisher-driven RSL initiative, which positions itself as an open pay-per-crawl solution.

At present, RSL is best understood as a parallel and complementary effort to Cloudflare's signaling vocabulary. In combination, these initiatives complicate the emerging landscape, blurring the boundaries between infrastructure-led, standards-based, and licensing-centric approaches to signaling and enforcing AI-related content preferences.

Taken together, these developments highlight a fundamental difference between licensing-centric approaches such as RSL and opt-out or preference-signaling mechanisms. The latter primarily express permissions or reservations that must be interpreted against an existing legal default. Licensing-centric approaches, by contrast, reframe automated access to publicly available content as a transactional activity that is conditional on payment. This shifts the centre of gravity from signaling preferences that operate in the context of legally permitted uses to negotiating access on contractual terms—with implications for who can realistically participate. In practice, such approaches tend to favor actors with the technical and organizational capacity to manage licensing and payment flows and risk, structurally constraining access to and use of publicly available content.

Rather than converging on a single opt-out standard, the developments described above point to an increasingly fragmented landscape of preference-signaling and licensing mechanisms. In this context, the central challenge returns to one already identified in earlier work: the need for a shared vocabulary that enables translation between different approaches. Such a vocabulary does not constitute an opt-out standard in itself, but rather an intermediary layer that allows preferences expressed in one system to be meaningfully interpreted in another.

## MAPPING THE VOCABULARY LANDSCAPE

As we enter 2026, where does this leave the various vocabularies that have been proposed or incorporated into existing standards and specifications?

While the emergence of contentsignals.org and the RSL specification represents notable recent developments, the overall vocabulary landscape has otherwise remained relatively stable. The following overview focuses on standards, specifications, and related proposals that include at least one usage category that can be targeted for the expression of preferences.

This includes the aforementioned RSL specification, contentsignals.org, C2PA/CWAG TDM Assertions, the TDM Reservation Protocol (TDMRep), two versions of the IETF AI Preferences vocabulary draft (with the earlier version included for reference), the IPCT Plus extension of the IPTC Photo Metadata Standard, and Open Future's vocabulary proposal, which provided the basis for the initial IETF draft.

While some of these initiatives appear in the May 2025 EUIPO's study The Development of Generative Artificial Intelligence from a Copyright Perspective, others referenced in that study have been excluded here for the following reasons:

- Spawning.ai's Do Not Train registry and ai.txt have been discontinued following the apparent cessation of Spawning's operations in summer 2025.
- Valuenode's proposed Open Rights Data Exchange and the JPEG Trust Core Foundation v2 governance framework do not currently function as opt-out mechanisms and do not define vocabularies for expressing usage-level preferences.[7][8]
- TDM•AI, while under active development, does not seek to define its own vocabulary, but instead aims to implement an existing standard.

[7] The proposed JPEG Trust Core Foundation v2 operates at the level of trust, provenance, and governance for media assets and identities, and could—at most—carry or reference opt-out assertions defined elsewhere in the future. As such, JPEG Trust Core Foundation v2 is neither an opt-out mechanism nor does it provide an opt-out vocabulary.

[8] According to the first report of the Copyright Infrastructure Task Force, ORDX may eventually support or reference TDM opt-outs, but it is not itself an opt-out mechanism, nor does it currently define the vocabulary, semantics, or technical behaviour required for AI Act–relevant TDM opt-out compliance. This places ORDX in the same conceptual category as JPEG Trust: potentially complementary infrastructure, but not a candidate opt-out standard.

The diagram below visualizes the current vocabulary landscape by mapping relationships[9] between terms across different initiatives, making visible both areas of convergence and points of divergence without implying a single, comprehensive taxonomy.

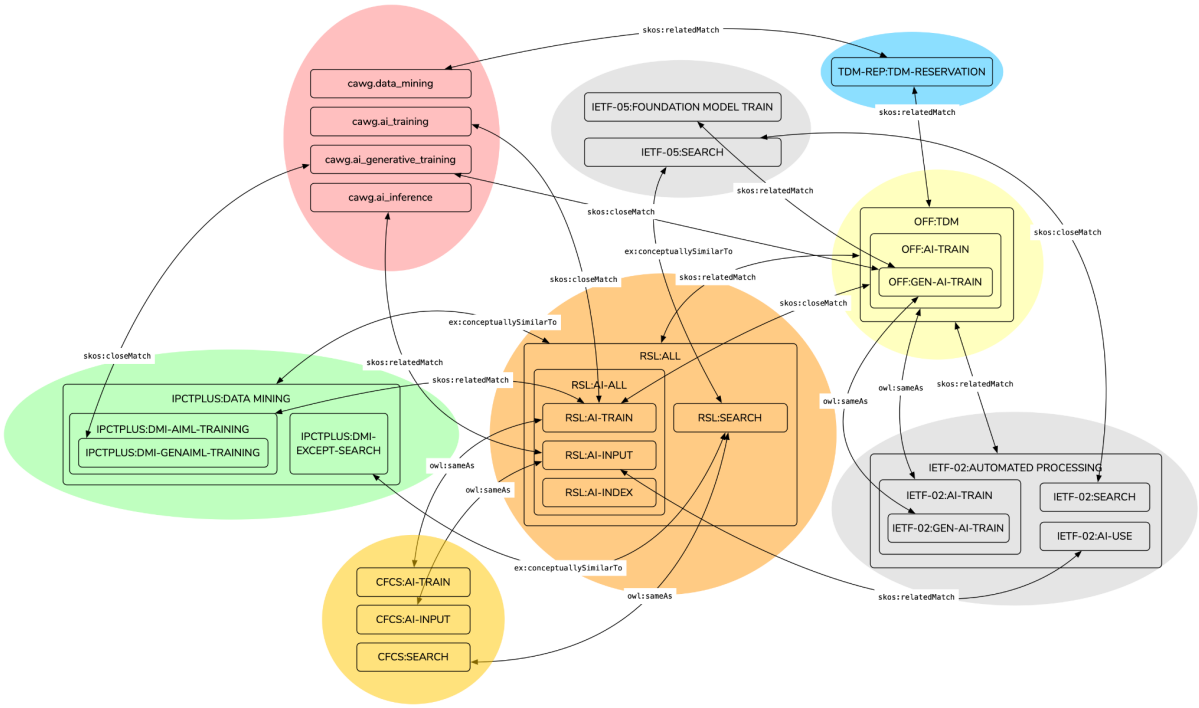


Figure 2: Conceptual relationships between opt-out vocabularies. Click to enlarge.

The diagram reveals several important insights into the current state of efforts to define and align vocabularies for AI preferences. At a very high level, it clearly illustrates that by the end of 2025, no single shared vocabulary has emerged. Instead, multiple partially overlapping vocabularies coexist, addressing similar concerns while differing in scope, terminology, and underlying assumptions. Several observations stand out:

- **Training-related concepts are the most consistently represented:** With the exception of TDMRep[10], all vocabularies include at least one category covering AI model training.

[9] The relationship labels used in the diagram describe how terms from different vocabularies relate to one another at a conceptual level. "ExactMatch" denotes terms that are intended to represent the same concept and can be treated as functionally equivalent across vocabularies. "CloseMatch" is used where terms are similar in scope and meaning but are not fully interchangeable, for example because they differ in granularity or implicit assumptions. "RelatedMatch" indicates a looser association between terms that address adjacent or overlapping concerns but are defined for different purposes or usage contexts. "ConceptuallySimilarTo" is used where terms are not directly mappable but express comparable underlying ideas or policy objectives. These relationships are descriptive and comparative in nature and do not imply formal interoperability, technical substitutability, or legal equivalence.

[10] The TDMRep specification was drawn up before the emergence of Generative AI models as a tool to express rights reservations that cover the full scope of the TDM exception in Article 4 of the CDSMD directive.

However, the distinction between general AI training and generative AI models training—a key feature of earlier vocabularies (including C2PA, IPTCPlus, Open Future's proposal, and early versions of the IETF draft)—is no longer present in the more recent vocabularies, such as RSL and Content Signals. The current IETF vocabulary stands out in this regard, as it introduces a more targeted category—foundation model training—that is conceptually similar to generative AI training, but does not include a broader category covering AI training more generally.

- **Search-related concepts are widely included but remain difficult to define consistently:** Almost all vocabularies include a category intended to capture "traditional search". However, the diagram illustrates that this category is among the least stable conceptually. While often treated as a carve-out from AI-related uses, "search" is defined differently across vocabularies. As a result, relationships between search-related terms are frequently associative rather than equivalent, reflecting ongoing disagreement over whether "search" can still be meaningfully defined as a distinct, non-AI category.

- **Inference- and deployment-related uses are gaining prominence but remain highly fragmented:** Across vocabularies, there is increasing attention to AI input, use, or inference-time interactions with content, reflecting growing concern over downstream and potentially substitutive uses. At the same time, the diagram shows significant variation in how these uses are defined and scoped, with weak or ambiguous relationships between corresponding terms. Notably, the current IETF AI Preferences vocabulary draft stands out as an outlier—it does not include categories addressing inference- or deployment-time uses, underscoring the lack of consensus on how such uses should be represented.

- **The role of an overarching or top-level category remains contested:** The diagram shows differing approaches to how relationships between broad concepts like TDM and more granular usage categories are modeled. Some vocabularies rely on an explicit top-level or umbrella category, while others avoid such constructs and instead establish relationships through more specific or contextual terms[11]. This divergence is reflected in ongoing disagreement over whether umbrella concepts (such as general TDM or Automated Use) should coexist with more narrowly defined categories, or whether vocabularies should rely exclusively on the latter.

Overall, the diagram shows that RSL currently provides the vocabulary covering the widest set of concepts discussed in this space. It includes categories that are not present in any other vocabulary, such as ai-index—defined as the inclusion of an asset in an AI system's internal index or retrieval database—and an ai-all container category encompassing all AI-related uses.

This breadth contrasts with the current IETF vocabulary draft, which has narrowed to only two categories. These categories exhibit relatively weak conceptual alignment with those found in

[11] Both the RSL specification and Cloudflare's Content Signals address the relationship with the EU TDM exception by including—or, in the case of RSL, allowing reference via its optional Default Access Terms—the following statement: "ANY RESTRICTIONS EXPRESSED VIA [CONTENT SIGNALS / THE ASSOCIATED RSL LICENSE] ARE EXPRESS RESERVATIONS OF RIGHTS UNDER ARTICLE 4 OF THE EUROPEAN UNION DIRECTIVE 2019/790 ON COPYRIGHT AND RELATED RIGHTS IN THE DIGITAL SINGLE MARKET."

other vocabularies, placing the current IETF draft at the periphery of the emerging vocabulary landscape.

This situation is problematic for several reasons. The RSL specification—and the closely related Content Signals approach—have been developed outside formal standards bodies[12]. From the outside, these initiatives appear to be supported by a relatively narrow and homogeneous set of stakeholders with aligned commercial interests, namely a mix of leading global Content Delivery Network (CDN) providers and US-based online publishers[13].

This limited supporter base is reflected in the licensing and contract-centric framing adopted by the initiative. **The extensive RSL documentation focuses almost exclusively on copyright-based permissions and contractual access terms, while making little or no reference to user rights, limitations and exceptions to copyright, or the broader need to balance the interests of rightholders and publishers with those of other stakeholders.** In doing so, the approach effectively conflates the exclusive rights afforded by copyright law with technical access-control mechanisms operated by online publishers and site owners, who in many cases are not themselves rightsholders.

Taken together, these characteristics suggest the RSL approach risks extending the scope and reach of publisher control over the use of publicly available content beyond what copyright law alone affords.

At the same time, many concerns about overreach and the need to account for limits, exceptions, and competing interests were raised within the IETF working group by participants representing user rights advocates, public-interest technologists, and some AI developers. While these concerns have contributed to the group's difficulty reaching consensus and making progress, the IETF AI Preferences vocabulary draft remains the only approach discussed in this paper that explicitly acknowledges the limits of copyright and other forms of control over access to content.

By the end of 2025, this has resulted in a situation where the open, multi-stakeholder approach pursued within the IETF framework risks being crowded out by initiatives operating under different governance models, with a stronger focus on publisher-facing tooling than on balancing the interests of all affected stakeholders, including individual users, researchers, non-profit organizations, and AI developers.

---

[12] That said, it is worth noting that Cloudflare released its Content Signals policy to the public domain and has expressed intent to continue working with neutral standards bodies, like the IETF. In line with this in October the Content Signals specification was introduced as a vocabulary draft in the AI Preferences Working Group for further consideration.

[13] The RSL standard is also supported by Creative Commons. This support indicates engagement from Creative Commons, although it should not be read as implying endorsement of specific implementation choices or broader policy positions reflected in the RSL specification.

# WHAT LIES AHEAD IN 2026?

In early December 2025, the European Commission launched a public consultation on protocols for reserving rights from text and data mining under the AI Act and the GPAI Code of Practice. The consultation is part of a Commission-led process aimed at "the identification and general agreement on opt-out protocols that can be considered as state-of-the-art, technically implementable, and widely adopted by rightsholders across different cultural and creative sectors"

The consultation builds on the EUIPO study on development of generative artificial intelligence from a copyright perspective, published in May 2025. It consists of two components. First, a public questionnaire—open until 23 January 2026—seeks stakeholder views on the text and data mining opt-out solutions identified in the EUIPO study. Second, the Commission has announced a series of follow-up workshops intended to discuss opt-out protocols that Code of Practice signatories and other providers would be expected to respect in the context of compliance with the AI Act.

Because it is largely based on the EUIPO study's findings, the consultation risks having limited practical relevance. As discussed above, several of the opt-out solutions identified by the EUIPO in the first half of 2025 are no longer available—most notably ai.txt and Spawning.ai's Do Not Train registry. Others, such as the proposed Open Rights Data Exchange (ORDX) and JPEG Trust Core Foundation, do not currently function as opt-out mechanisms.

At the same time, two of the most consequential recent developments in this space—Cloudflare's Content Signals and the RSL specification—emerged only after the EUIPO study was published and are therefore not reflected in its analysis. As a result, the consultation is anchored in a snapshot of the landscape that no longer fully captures the range of approaches currently shaping practice.

Against this background, the consultation appears to reflect an expectation that it will be possible to identify a set of opt-out protocols that are technically implementable and widely adopted by rightsholders across different cultural and creative sectors, and that these protocols can serve as a stable reference point for compliance with the AI Act and the GPAI Code of Practice. Implicit in this framing is the assumption that convergence at the level of opt-out mechanisms is both achievable and desirable.

However, as the preceding analysis shows, the current landscape is characterized less by convergence on specific mechanisms than by the coexistence of multiple, partially incompatible approaches reflecting different legal interpretations, technical architectures, and commercial incentives. In this context, the question is not which opt-out mechanism should be preferred, but how preferences expressed through different mechanisms can be meaningfully understood and interpreted in a consistent way.

This shifts the focus from opt-out mechanisms themselves to the vocabularies they rely on: the conceptual layer that determines how different forms of signaling, reservation, or licensing are interpreted across systems. It is at this level—not at the level of individual protocols or

enforcement mechanisms—that alignment through standardization processes is both most feasible and most impactful.

## *Design principles for a shared vocabulary layer*

If convergence at the level of opt-out mechanisms is neither realistic nor desirable, attention should shift to designing a shared vocabulary layer that can support translation and interpretation across different signaling, licensing, and enforcement approaches. Such a vocabulary does not determine legal outcomes or technical behavior on its own, but it does shape how preferences are expressed, understood, and acted upon across systems. The following principles emerge from the analysis above as particularly important.

- **Legal minimalism:** A shared vocabulary should be designed to operate within existing legal frameworks rather than to restate, reinterpret, or expand them. Its role is to provide a stable and shared semantic reference for expressing preferences—not to encode legal conclusions or to create new normative effects through standardization.

  Vocabulary terms should therefore not be defined or framed in ways that presume the existence, scope, or enforceability of exclusive rights, permissions, or prohibitions beyond what is determined under applicable law. While the legal effect of acting on a given preference may vary across jurisdictions, the underlying meaning of the preference itself should remain consistent and should not be used to bypass, pre-empt, or narrow the operation of exceptions, limitations, or other user rights.

- **Explicit acknowledgement of limits and exceptions:** A shared vocabulary should make explicit that expressed preferences operate within legal limits, including exceptions and limitations to copyright and other user rights. This is particularly important to avoid misinterpretation of preferences as absolute prohibitions and to ensure that lawful uses—such as those by researchers, libraries, or individual users—are not inadvertently chilled.

- **Separation of signaling and enforcement:** The vocabulary layer should remain clearly distinct from mechanisms for technical enforcement or contractual licensing. Its function is to express preferences in a structured and interpretable way, not to determine how those preferences are implemented or enforced. Maintaining this separation helps prevent vocabularies from becoming de facto control mechanisms and preserves flexibility for different implementation contexts.

- **Interoperability before completeness:** Rather than aiming to capture every conceivable use case or business model, a shared vocabulary should prioritize interoperability and conceptual clarity. A smaller set of well-defined, widely understood terms that can be mapped across systems is more valuable than a comprehensive but idiosyncratic taxonomy that resists translation.

Taken together, these principles point toward a vocabulary layer that functions as shared semantic infrastructure—enabling coordination across diverse technical and legal approaches,

while remaining neutral with respect to enforcement models, commercial arrangements, and regulatory compliance strategies.

## OUTLOOK: A SHARED VOCABULARY AS INSTITUTIONAL INFRASTRUCTURE

The analysis in this brief points to a growing urgency for stakeholders to address the vocabulary problem directly. As preference-signaling mechanisms proliferate and harden into practice, the absence of a shared, interoperable vocabulary increasingly risks fragmentation, inconsistent interpretation, and over-enforcement. This issue cannot be left to regulators alone, nor can it be solved by initiatives driven by narrow or predominantly commercial interests.

At the same time, the choice confronting policymakers and stakeholders is not a technical one. Seen in this light, the divergence of technical protocols reinforces the case for moving coordination efforts up the stack—from enforcement mechanisms and signaling formats to the vocabulary and semantic layer. As protocols increasingly reflect divergent legal interpretations, technical architectures, and commercial incentives, attempts to impose convergence at the mechanism level become both less realistic and more brittle. This brings the analysis back to an earlier observation made in our 2023 policy brief on opt-out compliance considerations: shared vocabularies, rather than shared protocols, represent the most viable point of coordination across heterogeneous systems.

At this stage, choices about how vocabularies are defined and aligned are beginning to harden into practice. Decisions made today—whether through standards processes, infrastructure-led initiatives, or policy-adjacent coordination—will have lasting effects on how lawful access, reuse, and innovation are balanced against control and exclusion. Open, multi-stakeholder forums remain essential for maintaining that balance. Whether the IETF AI Preferences Working Group can reclaim and sustain this role remains an open question, but the need for such a forum has become increasingly clear.

# ABOUT OPEN FUTURE

Open Future is a European think tank that develops new approaches to an open internet that maximize societal benefits of shared data, knowledge, and culture.

Paul Keller is a co-founder and director of policy at Open Future. His work focuses on the intersection of copyright policy and emerging technologies. He works on policies and systems that improve access to knowledge and culture and protect the digital public sphere.