

From ACAP to RightsML: the news industry and the future of ODRL

Brendan Quinn, International Press Telecommunications Council (IPTC)

mdirector@iptc.org

Position paper for the W3C Workshop on the Future of ODRL, London, 20–21 July 2026

Abstract

The IPTC is the global standards body for the news media and the steward of RightsML, the news industry’s rights expression language and an ODRL profile. This paper offers a practitioner’s perspective on ODRL: a short history of how RightsML developed from the earlier ACAP protocol and converged on ODRL; an account of IPTC’s role in the W3C work, including the two implementations that helped ODRL reach Recommendation and the extent to which RightsML’s vocabulary was incorporated into ODRL itself; and evidence from our members on why adoption remains difficult. We use examples from three news agencies — Reuters, Deutsche Presse-Agentur and AP — to illustrate both the gap between rights that are readable by humans and rights that are computable by machines, and the middle ground between ODRL’s complexity and simpler but less expressive alternatives. We conclude with what the news industry needs from the next phase of ODRL, particularly in the context of AI and text-and-data mining.

1. About IPTC and this contribution

The IPTC is a consortium of major news agencies, publishers and system vendors, and the body responsible for many of the standards used to exchange news internationally — NewsML-G2, ninjs, the NewsCodes vocabularies, the Photo and Video Metadata standards, and RightsML. It is therefore an implementer as well as a standards body. This contribution reflects that position: rather than proposing new formal semantics — addressed by other papers here — it offers an account of how ODRL is, and is not, being used in the news industry, and of what would help.

2. From ACAP to RightsML: a short history

The news industry’s first attempt at machine-readable rights was the **Automated Content Access Protocol (ACAP)**, proposed in 2006 by the European Publishers Council, the World Association of Newspapers and the International Publishers Association. ACAP extended the `robots.txt` Robots Exclusion Protocol so that publishers could indicate to search-engine crawlers what they were permitted to do. Its weakness was not technical but political: compliance was voluntary, and the major search engines declined to adopt it. A protocol of this kind has no effect unless crawlers observe it — a limitation directly relevant to current debates about AI training.

Faced with that impasse, the project changed direction. Rather than adding more to `robots.txt`, an effort known informally as “ACAP 2.0” set out to express rights as a proper profile and vocabulary built on **ODRL**, the Open Digital Rights Language created by Renato Iannella in 2000. In June 2011 the owners of ACAP handed the project to the IPTC, which renamed it **RightsML**. RightsML 1.0 was approved in March 2012 as a profile of ODRL 2.0; version 1.1 followed in October 2012. The significant change was that the news industry stopped developing its own rights mechanism and instead adopted and extended an existing, general-purpose one.

3. RightsML as an ODRL profile

3.1 Shaping the profile mechanism

When ODRL moved onto the W3C Recommendation track through the Permissions and Obligations Expression (POE) Working Group, the language was deliberately reworked into a *generic* information model: many actions, constraints and party functions were made non-normative so that any sector could extend the core through a formal new mechanism, the **ODRL Profile**. IPTC participated in this work directly and took on the role of defining the news-industry profile. When ODRL 2.2 became a W3C Recommendation in February 2018, IPTC published **RightsML 2.0** as that profile, one of the first industry profiles of ODRL.

3.2 Two implementations for the Recommendation

W3C process recommends implementation experience before a specification can advance: the model and vocabulary should be implemented into software by at least two parties and be tested against a list of criteria. Two of those implementations were written by IPTC participants: a **JavaScript library** (`rightsml-lib-js`, classes for the ODRL model and the RightsML profile) and a **Python library** (`licensed.py`, which produces ODRL and RightsML documents in XML or JSON with a schema-validated test suite). Both are now out of date; the IPTC would welcome collaborators to help with refreshing them.

3.3 The profile mechanism in practice

The relationship between RightsML and ODRL also illustrates how the profile mechanism is intended to work. As terms first developed for a particular sector prove more widely useful, they can be taken up into ODRL’s shared Common Vocabulary; the profile then need only declare which of those terms apply, together with any genuinely sector-specific additions. RightsML 2.0 shows this progression: of the roughly eighty terms its JSON-LD profile declares, all but one are drawn directly from the ODRL namespace, with a single term (`compensationValue`) unique to RightsML. Over successive versions the profile has therefore become thinner, its role shifting from defining a vocabulary towards selecting and constraining ODRL’s own. We see this as a natural evolution and a demonstration of the profile system in use.

4. RightsML in practice: adoption and the simplicity problem

ODRL and RightsML appear in a growing range of applications, among them TDMRep for text-and-data-mining reservations, and emerging licensing approaches such as Really Simple Licensing. However, uptake among ordinary publishers remains patchy. Ahead of this workshop we canvassed our members, and the recurring theme is not hostility but friction: the RDF and JSON-LD learning curve, the lack of tooling and validation, and the absence of any enforcement once a policy is expressed.

IPTC’s recent work in this area targets the ease-of-use problem. The **RightsML Generator** lets users build multi-constraint policies without hand-writing JSON-LD. The prototype **simplerights.iptc.org** service goes further, expressing common restrictions as short URIs. Inspired by Creative Commons and `rightsstatements.org`, a URI can issue a simple, broad rights permission or restriction, for example `https://simplerights.iptc.org/geo/exceptfor/CN` for “usable anywhere except China” — which resolve, via HTTP content negotiation, either to a human-readable page or to a full ODRL policy in JSON-LD, RDF/XML or Turtle. A single identifier therefore serves both a human reader and a machine.

5. RightsML/ODRL in other IPTC standards

RightsML and ODRL can also be carried within IPTC’s other standards. Rights expressions may be included in NewsML-G2 and ninjs, and the Photo Metadata Standard and Video Metadata Hub provide dedicated Embedded and Linked Encoded Rights Expression properties: a policy can be embedded directly in the asset’s metadata (via XMP for images and video) or linked by reference to an external URL that holds the ODRL.

The linked form corresponds to ODRL’s own `hasPolicy` mechanism, by which an asset references a policy and is then inferred to be the target of its rules. This is similar to the SimpleRights approach: because ODRL requires every permission and prohibition to name a target asset, a policy served independently of any asset — as SimpleRights does — is not strictly conformant, but is in effect a reusable template, completed only when an asset links to it. ODRL currently has no formal construct for such asset-independent templates, a point we return to below.

6. Member evidence: how three agencies signal rights today

A useful indicator of ODRL’s prospects in the news industry is the current practice of the major agencies. Ahead of the workshop we examined three, via direct member feedback and consulting public API documentation. None uses ODRL in production; each has addressed the same problem differently; and together they indicate the gap the standard needs to close.

6.1 Reuters: compact codes over free text

Reuters’ published *Restrictions Guide* defines a compact, near-controlled convention: a media use type (*Broadcast* or *Digital*) before a colon, and after it a restriction from a small set — `None`, `No Use [Country]`, `No Use [Company]`, `No New Uses After [date]`, `No Resales`, `No Archive`. In distributed content, these codes are conveyed as free text within a single element.

```
<usageTerms xml:lang="en">Broadcast: No Use Brazil. Digital: No Use Brazil.</usageTerms>
```

The restriction is therefore free text, not a machine-readable rights expression: there is no ODRL — or other structured — policy here. To an editor on a picture desk, “No Use Brazil” is unambiguous; to a syndication pipeline, an embargo-enforcing asset-management system, or an AI-training filter, it is opaque prose that must be parsed, and — as the `xml:lang="en"` attribute indicates — parsed in English only. (Reuters has historically carried structured “outs” as NewsML-G2 `exclAudience` filters, but these are a coarse distribution filter — “do not deliver to this audience” — rather than a rights model.)

6.2 Deutsche Presse-Agentur: “too complex, too simple”

Deutsche Presse-Agentur (dpa) reached a similar conclusion by a different route. Needing machine-readable spatial restrictions on third-party photos it republishes — now required per item, as those photos flow into general multimedia wires rather than separate feeds — dpa weighed its contracts against both ODRL and SimpleRights and judged ODRL too complex and SimpleRights too simple. Its real need is narrow: express either “only these few countries” or “the world except these few countries”. SimpleRights cannot state a set of several countries at once, and neither scheme names the base “world” from which exceptions are subtracted. So dpa, too, invented a compact code: `dnlrights:area-AT-CH-DE` for “Austria, Switzerland and Germany only”, and `dnlrights:area-world--GB--IR` for “the world except the UK and Ireland”.

6.3 Associated Press: rights spread across fields

AP’s public Media API documentation shows a third approach. AP carries its restrictions as free text across several NewsML-G2 fields — `usageTerms` for client-facing restrictions, `edNote` for internal editorial instructions not to be passed on, and a `RESTRICTION SUMMARY` block inside the video script — rather than as RightsML or ODRL policies in this feed. The phrasings echo Reuters closely — “No Use in Japan”, “No Archiving Permitted”, “AP Clients Only”, “No re-sale, re-use or archive” — while a few elements, such as an `expires` timestamp, are already machine-actionable. AP’s approach spreads the same information across more fields, but the substance of the restrictions still travels as prose.

6.4 The common thread — and how RightsML/ODRL can help

The three agencies show a common pattern: rights carried almost entirely as human-readable text in the same IPTC container; two agencies using a compact code and one using prose; and a de facto shared vocabulary — No Use [Country], No Archive, clients-only — that is never formalised. That two of the three agencies, and IPTC’s own SimpleRights service, each independently developed a compact code indicates a specific gap: there is no standard shorthand at the level of abstraction these providers require.

ODRL, and by extension RightsML, can make that intent computable. The compound notice “not in China, not by CNN, not resold” becomes, in RightsML/ODRL (JSON-LD, abridged):

```
"prohibition": [
  { "target": "...item", "action": "use",
    "constraint": [ { "leftOperand": "spatial", "operator": "eq",
      "rightOperand": "http://cvx.iptc.org/iso3166-1a3/CHN" } ] },
  { "target": "...item", "action": "use",
    "assignee": "https://example.com/party/CNN" },
  { "target": "...item", "action": "sell" }
]
```

Each clause now has a defined meaning that software can act on: the country is a controlled-vocabulary identifier rather than the bare word “China”, and the rules are queryable, enforceable and language-independent. This does, however, assume agreed identifiers for the values involved — a standard that does not yet exist for countries, and still less for organisations such as CNN (hence the placeholder identifier above); we return to this below.

Would `simplerights.iptc.org` close the gap? Only partly. It covers the common single-dimension cases — geography, dates and embargoes, platform — well, and there the “why re-encode?” objection largely disappears, since emitting a URI costs little yet yields a machine-actionable signal. But in its current beta form it cannot express party or company exclusions, resale or archive rules, or a set of several countries, and as dpa found, it cannot combine restrictions. A useful feature to help with adoption is the `seeterms` URI: a machine-readable pointer indicating “consult the usage terms”, which wraps the text already published by the agencies and provides a low-cost first step towards adoption.

7. What the news industry needs from the future of ODRL

A consistent theme in our members’ feedback is that the tool they want sits *between* ODRL’s full expressiveness and the minimalism of simpler schemes. On this basis we offer six concrete, publisher-driven priorities for the next phase of ODRL:

1. **A URI-shortcut or named-policy mechanism** in the core, so that common restrictions

can be referenced as compactly as a code while still resolving to full semantics — generalising what SimpleRights, and agencies’ own codes, already prototype.

2. **Composability of simple expressions**, so that a set of territories, and party and action restrictions, can be combined without dropping to hand-authored RDF.
3. **First-class, lightweight vocabularies for the most common real-world restrictions** — territory, embargo and expiry, platform, party exclusion, resale and archive — shaped by what providers actually emit today.
4. **An explicit notion of scope**. Provide a standard way to name the base set — “all countries”, “the world” — from which exceptions are drawn, so that “all except {...}” is unambiguous rather than merely implied.
5. **Standard identifier schemes for constraint values and parties**. A policy is only interoperable if different systems resolve its identifiers identically. There is no agreed standard for country names as URIs (which is why IPTC minted its own scheme using ISO-3166, for example <http://cvx.iptc.org/iso3166-1a2/CN>). Schemes for identifying news brands such as “CNN” strike problems because “CNN” is a brand name and doesn’t have an organisational identifier. (Initiatives such as the Global Media Registry’s Global Media Identifier (ISO IWA 44:2025) are beginning to address this problem.) So whether an industry profile should specify identifier schemes is a worthwhile topic for discussion at this workshop.
6. **Support for reusable, asset-independent policies**. ODRL requires every permission and prohibition to name a target asset, so a policy published independently of an asset, as SimpleRights does, is not strictly conformant. A standard notion of reusable policy templates, applied to assets by reference (via `hasPolicy`) or instantiation, would let common restrictions be defined once and reused across many assets. Proposals elsewhere at this workshop address the same need.

Underlying all six is a central question: signalling versus enforcement. Publishers can express rights; they cannot compel anyone, least of all an AI crawler, to observe them. ACAP failed for this reason two decades ago. If ODRL is to be useful to the news industry in the age of AI crawling and data mining, the task is not only to refine the semantics but to make machine-readable rights simple to author, inexpensive to consume, and visible on the open web, where crawlers operate.

References

- [1] IPTC, “IPTC to Update RightsML to Align With W3C’s ODRL”, October 2017. <https://iptc.org/news/iptc-update-rightsml-align-w3cs-odrl/>
- [2] IPTC, “RightsML 2.0 is now published”, July 2018. <https://iptc.org/news/rightsml-2-0-is-now-published/>
- [3] IPTC, “Beta release of simplerights.iptc.org”, May 2025. <https://iptc.org/news/iptc-simplerights-beta-announcement/>
- [4] RightsML Generator. <https://iptc.org/std/RightsML/generator/>
- [5] TDM Reservation Protocol (TDMRep), W3C Community Group Final Report, May 2024. <https://www.w3.org/community/reports/tdmrep/CG-FINAL-tdmrep-20240510/>
- [6] Really Simple Licensing (RSL), RSL Collective. <https://rslstandard.org/>
- [7] Global Media Identifier (GMI) for distribution channels and brands, ISO IWA 44:2025. <https://www.iso.org/standard/88469.html>
- [8] Rights Statements. <https://rightsstatements.org/>
- [9] ACAP — Automated Content Access Protocol. <https://the-acap.org/>
- [10] AP Media API 2.2 Help Center, “Developer guide: Usage instructions and restrictions”. https://api.ap.org/media/v/docs/Usage_Instructions_and_Restrictions.htm
- [11] Reuters Liaison, “Reuters Restrictions Guide”. <https://liaison.reuters.com/tools/restrictions>