



A machine-readable rights reservation, an invitation from the European Commission, and the silence of the large labs

The crawlers of Microsoft, Google, Meta, Amazon, OpenAI, Anthropic and Apple consume ENTIA's corporate identity corpus every day. Since 19 July 2026 that corpus carries a machine-readable rights reservation. This is what has happened since, with names, dates and signed figures.

2,279,997

third-party machine requests · 19 Jul to 1 Sep · signed record

170

agent families observed

12

operators notified with file and licence proposal

6 / 6

replied / no reply

No reply

Five model providers and one search engine. As of 12 September no reply to the notification is on record.



OpenAI

SIGNATORY

GPTBot · fetched the policy on 30 Jul

NO REPLY



Google

SIGNATORY

GoogleOther · Googlebot deliberately not contacted

NO REPLY



Amazon

SIGNATORY

Amazonbot · the DSA mailbox serves authorities only

AUTOMATIC REPLY ONLY



Meta

Meta-ExternalAgent · 99.98% attribution from its network

NO REPLY



Apple

Applebot · fetched the policy

NO REPLY



Huawei

PetalBot · search engine

NO REPLY

Answered something else

They replied, but none on the question of machine-readable reservations.



Anthropic

SIGNATORY

ClaudeBot · Copyright Operations, 2 Sep: robots.txt

DID NOT ADDRESS THE QUESTION



Microsoft

SIGNATORY

Bingbot · three tickets; "unable to route"

COULD NOT ROUTE THE MATTER



Ahrefs

AhrefsBot · routed twice to legal and business

THE TEAM NEVER WROTE

Acted

SEO tools. None is a model provider. They are the only ones who moved.



Semrush

SemrushBot · Maria Kudlai (Datos), 1 Sep

CHANNEL OPEN, NO PROGRESS



DataForSEO

DataForSeoBot · stopped 26 Aug and deleted

REMEDIATED · DECLINES LICENCE



Moz

DotBot · blocklist, 1 Sep

REMEDIATED · DECLINES LICENCE

SIGNATORY = signed the copyright chapter of the GPAI Code of Practice (Commission's list, checked on 10 Sep). Five signatories on the board; none has yet answered the question. Detail with dates, channels and quotes on the operators page (/press/operators).

CODE OF PRACTICE FOR GENERAL-PURPOSE AI MODELS · COPYRIGHT CHAPTER

What these five signed



Anthropic, OpenAI, Google, Microsoft and Amazon signed the full Code, including the copyright chapter. Meta and Apple are not on the list.

The chapter exists to give effect to Article 53(1)(c) of the AI Act: to put in place a policy to comply with copyright law and, in particular, to identify and comply with rights reservations expressed under Article 4(3) of Directive 2019/790. By signing, each of them committed, in these words:

MEASURE 1.3 · IDENTIFY AND COMPLY WITH RIGHTS RESERVATIONS WHEN CRAWLING

"Signatories commit: a) to employ web-crawlers that read and follow instructions expressed in accordance with the Robot Exclusion Protocol (robots.txt) [...] and b) to identify and comply with other appropriate machine-readable protocols to express rights reservations pursuant to Article 4(3) of Directive (EU) 2019/790, for example through asset-based or location-based metadata [...]"

Two commitments, not one: robots.txt and, in addition, other machine-readable protocols. Anthropic's reply of 2 September speaks only to the first.

MEASURE 1.5 · POINT OF CONTACT FOR RIGHTSHOLDERS

"Signatories commit to designate a point of contact for electronic communication with affected rightsholders and provide easily accessible information about it."

ENTIA wrote to each one through the channel it publishes and, on 10 September, asked the Commission for those contact points. No signatory has indicated its own.



AND ANTHROPIC, IN ITS OWN CONSTITUTION

Claude's Constitution, published by Anthropic under a CC0 licence and described by the company itself as the final authority on the behaviour it seeks for Claude, lists intellectual property among the things it does not want Claude to do. Verbatim, in the list of behaviours that would concern a thoughtful senior Anthropic employee:

"Help someone violate intellectual property rights or make defamatory claims about real people."

"It's also the final authority on our vision for Claude, and our aim is for all of our other guidance and training to be consistent with it."

Source: [anthropic.com/constitution](https://www.anthropic.com/constitution) (https://www.anthropic.com/constitution), consulted on 12 September 2026. ENTIA does not treat that document as a legal obligation nor assume that it technically governs the crawler. The question still open is how that institutional commitment is integrated with the crawling policy where the rightsholder permits access, expresses a machine-readable reservation and the agent retrieves the file that contains it.

Verified-entity records served in full, by agent family. Attribution by agent name and originating network together; what is not corroborated counts against nobody.

2,279,997


third-party requests

of 2,353,964 raw; 73,967 of our own declared and excluded

1,158,743

distinct verified entities

1,431,698 records served in full; the two figures, always kept apart



The same record fetched twice counts twice. In the same window, 65% of requests carrying the name Googlebot and 65% carrying the name GPTBot did not come from the networks those operators publish: an agent name, on its own, is worth little. In the sealed record of 19 Jul to 23 Aug, 72.8% of complete records went to families that never requested the policy document, although the reservation travelled on every response. Source: signed consolidated record (</access-report/acta-consolidada.pdf>). Live telemetry continues at entia.systems/bot ([/bot](https://entia.systems/bot/bot)); it is not mixed with the record. All 170 families, nothing aggregated, on the families page (</press/families>).



What the Commission has done, and what it has not



28 Jul

ENTIA writes

To the AI Office and CNECT.I.2, copying the EUIPO. No finding of infringement or intervention requested.



3 Aug

Ilana Kunkel replies

Art. 4(3) prescribes no method; the Commission validates no protocol; the Code process is ongoing.



19 Aug

ENTIA asks to join

The Measure 1.3 workshops of the copyright chapter.



24 Aug

Written invitation

Workshops closed since January; the Commission invites the contribution in writing, with the solution described.



24 Aug · 2 Sep

Contribution and update

Published the same day at /eu-submission (/eu-submission); update with the twelve operators on 2 Sep.



10 Sep

Measure 1.5 request

ENTIA asks for the contact points signatories must publish. No reply as of 12 Sep.

What the Commission has not done. It has not assessed ENTIA's implementation, has not endorsed it, and has taken no position on any operator. The invitation is to contribute to an open process, not a recognition. ENTIA asks that nobody write otherwise.

POLICY V2 IN FORCE SINCE 9 SEP 2026 · V1 SINCE 19 JUL

The reservation, as machines receive it

PUBLIC VERSION REGISTRY

01 · HTTP HEADERS

On every response

tdm-reservation: 1

tdm-policy: <https://entia.systems/.well-known/tdm-policy.json>

tdm-policy-version: 2

Checked live on 11 and 12 Sep.

02 • MACHINE FILES

Readable by any agent

```
/.well-known/tdm-policy.json (/well-known/tdm-policy.json)
/.well-known/tdmrep.json (/well-known/tdmrep.json)
/.well-known/entia-license.json (/well-known/entia-license.json)
/ai.txt (/ai.txt)
```

03 • HUMAN NOTICE

Two languages

```
/legal/en/tdm-reservation-v2 (/legal/en/tdm-reservation-v2)
/legal/es/reserva-tdm-v2 (/legal/es/reserva-tdm-v2)
version registry with fingerprints (/well-known/tdm-policy-registry.json)
```

04 • ROBOTS.TXT

Courtesy, not reservation

Content-Signal: search=yes, ai-input=yes, ai-train=no

The exclusion protocol cannot express a rights reservation nor oblige a crawler to read the policy.

PERMITTED

- Search indexing, including persistent indices
- Real-time retrieval and grounding, also by AI products
- Consultation, linking and citation

RESERVED • LICENCE REQUIRED

- Model training or parameter update
- Training or evaluation datasets
- Reconstruction of an independent corpus or database
- Promotion of a permitted use to a general one · onward supply to third parties

Reading or complying with the policy forms no contract; an access record creates no fee, debt or presumption of harm; the policy is not retroactive. V2 replaced V1 after a red team of 161 attacks from the position of 15 operators.

Stated limit: that the header is served today is verified; that it was served on every response across the whole window cannot be checked backwards, because the edge log keeps no headers.

19 JUL TO 9 SEP • ONLY FROM NETWORKS MATCHING THOSE EACH OPERATOR PUBLISHES

Who went to fetch the reservation, and who opened their file

OPERATOR	FETCHED THE MACHINE FILES (.WELL-KNOWN)	FETCHED THE LEGAL NOTICE (HTML)	OPENED ITS FILE	REPLIED
 OpenAI	YES · GPTBot 30 Jul, SearchBot 23 Aug, ChatGPT-User 4 and 6 Sep	YES · 4	NO	NO
 Anthropic	YES · ClaudeBot 27 Jul, all four files	NO	NO	NOT ADDR
 Apple	YES · tdm-policy	YES · 2	NO	NO
 Semrush	YES · tdm-policy, tdmrep	YES · 2	NO	YES · ch oper
 Moz	YES · tdm-policy, tdmrep	YES · 1	NO	YES remedia
 Microsoft	NO	YES · 28	NO	NOT ROI
 Amazon	NO	YES · 7	NO	AUTOMA
 Meta	NO	YES · 3	NO	NO
 Ahrefs	NO	YES · 4	NO	NOT ADDR
 Google	NO	NO	NO	NO
 DataForSEO	NOT ON RECORD	NOT ON RECORD	NO	YES remedia
 Huawei	NOT ON RECORD	NOT ON RECORD	NO	NO

The /.well-known/ files have no indexing value: requesting them is a specific act on the reservation. The HTML legal pages are indistinguishable from ordinary indexing. Retrieving is not reading and reading is not understanding: a 200 proves it was requested and served. Files: as of 1 Sep, out of twelve dispatches, zero authenticated openings by an operator. Three (Anthropic, Amazon, Microsoft) hit our static-key screen between 22 and 26 Aug, replaced by one-time codes on 26 Aug: our friction, placed on record.

THREE FACTS

The silence of the labs, measured

30 Jul

They went to fetch the reservation

OpenAI's GPTBot retrieved tdm-policy.json and tdmrep.json from OpenAI's network. ClaudeBot had done so on 27 July, between 15:28 and 15:30 UTC, all four files with 200 responses. Then no reply from OpenAI and a reply from Anthropic about robots.txt.

0

Nobody opened their file

Twelve dispatches, zero authenticated openings by an operator as of 1 September. For Anthropic, measured on 10 September: zero requests from its network to the files and zero to the reservation files since V2 was published. Its network did take the whole press room on 30 July and 24 legal-centre pages on 21 August.

2 Sep

The only reply from a lab

Anthropic's Copyright Operations Team pointed to robots.txt, a Disallow directive and a blocklist form. ENTIA had asked for no block. On 3 September it put four questions and asked for escalation to the Article 53(1)(c) team. As of 12 September, no reply.

Article 53(1)(c) of the AI Act requires providers of general-purpose models to identify and comply with Article 4(3) reservations, including through state-of-the-art technologies. Measure 1.3 of the Code distinguishes between respecting robots.txt and, in addition, identifying and complying with other machine-readable protocols. ENTIA's question to the five signatories is that second half. None has answered it yet.

What the evidence supports

- That specific requests arrived at the recorded times, under which agent name and from which network.
- What was served, how many times and how many distinct records, each day sealed on its date.
- Which agents requested the reservation files, and when.
- What was communicated to each operator, through which channel, when, and what each answered, with the full email preserved.

What ENTIA does not assert

- That any operator trained a model, built an index or incorporated the corpus into a product: not observable from outside.
- That there is infringement, contractual acceptance, debt or damage per request.
- That a User-Agent string legally identifies an operator.
- That the European Commission has assessed or endorsed any of this.

Where ENTIA has measured wrongly it has corrected the figure and left it in writing: a false zero of robots.txt retrievals; an over-broad exclusion that removed third-party traffic, later restored; four reports that labelled requests served as distinct records. Superseded versions are preserved; nothing is rewritten silently.

19 JUL TO 12 SEP 2026

Timeline

● ENTIA ● EUROPEAN COMMISSION ● OPERATORS

19 Jul

Reservation V1 takes effect. From 17:00 UTC the edge writes the per-request record.

27 Jul

ClaudeBot retrieves the four reservation files from Anthropic's network.

28 Jul

ENTIA presents the case to the AI Office and CNECT.I.2.

30 Jul

GPTBot retrieves tdm-policy.json and tdmrep.json from OpenAI's network.

31 Jul

First press pack (European Machine Access Case), now archived.

3 Aug

Ilana Kunkel replies: no prescribed method, no protocol validation, Code process open.

19 Aug

ENTIA asks to join the Measure 1.3 workshops and sends its first formal notice to DataForSEO.

24 Aug

The Commission invites a written contribution. ENTIA publishes it the same day. DotBot (Moz) stops crawling.

26 Aug

DataForSEO stops crawling and, as it later confirms, deletes the data.

31 Aug · 1 Sep

The twelve files are dispatched. Semrush designates Maria Kudlai. Moz confirms the blocklist.

2 Sep

The consolidated record is signed and the update sent to the Commission. **Anthropic** replies pointing to robots.txt.

3 Sep

Questions to Anthropic and request for a position from all twelve. Ahrefs and Microsoft: outside their channel.

4 Sep

DataForSEO (legal) sets out its position: complied with robots.txt, rejects the characterisation and the licence.

9 Sep

Policy V2 in force (09:04 UTC) and published (10:43 UTC). Daily reading record by family.

10 Sep

Request to the Commission for the Measure 1.5 contacts; to DataForSEO, its policy on machine-readable reservations.

12 Sep

Status of this page. Nothing from OpenAI, Google, Meta, Apple, Amazon or Huawei.

EVERYTHING CHECKABLE

Sources and documents

[entia.systems/eu-submission](#) (/eu-submission)

Index of the contribution to the Commission: implementation, operator table and the reservation as served.

[entia.systems/eu-submission/contribution](#) (/eu-submission/contribution)

The full written contribution, with corrected errors.

[entia.systems/eu-submission/update-2026-09-02](#) (/eu-submission/update-2026-09-02)

Update of 2 September: operators contacted and the attribution revision.

[entia.systems/access-report/acta-consolidada.pdf](#) (/access-report/acta-consolidada.pdf)

Consolidated record 19 Jul to 1 Sep. The signed container (.asice) is at the same address and validates at the Commission's DSS.

[entia.systems/.well-known/tdm-policy-registry.json](#) (/well-known/tdm-policy-registry.json)

Policy version registry, with the fingerprint of each.

[digital-strategy.ec.europa.eu](#) · [GPAI Code of Practice](#) (<https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>)

Signatories list and the copyright chapter (Measures 1.3 and 1.5), as published by the Commission.

[entia.systems/bot](#) (/bot)

Public machine-access observatory, live, by operator and day.

[entia.systems/licensing](#) (/licensing)

Licence channel open to any operator.

[entia.systems/press/european-machine-access-case](#) (/press/european-machine-access-case)

Archive: press pack of 31 July, 7 to 27 Jul window, superseded by this page.

"We are not asking anyone to stop crawling. We are asking those who committed to respect a rights reservation to say how they identify it and what they do with it."

FERNANDO VILCHES · FOUNDER, ENTIA



PrecisionAI Marketing OÜ

Sepapajatu 4, 11415 Tallinn, Estonia

Registry code 17048063 · DUNS 565868914

[Corporate status \(/legal/en/corporate-status\)](/legal/en/corporate-status)

[s://w](#)

[EU data · Estonia e-Residency](#)

[www.enti](#)

[nked](#)

[n.co](#)

[m/co](#)

[mpan](#)

[y/enti](#)

[a-sys](#)

[tems\)](#)

[License data \(/licensing\)](/licensing)

[Machine ecosystem \(/ecosyste](/ecosystem)

[m\)](#)

[Machine telemetry \(/telemetry\)](/telemetry)

[AI Discovery Observatory \(/obs](/observatory)

[ervatory\)](#)

TDM reservation and

[European Commission \(/press/](/press/tdm-reservation-european-commission)

[tdm-reservation-european-commis](#)

[sion\)](#)

[Status by operator \(/press/operat](/press/operators)

[ors\)](#)

[Five-Layer Framework \(/press/r](/press/research/five-layer-framework)

[esearch/five-layer-framework\)](#)

DEVELOPERS

[API \(/developers\)](/developers)

[MCP documentation \(/mcp-doc](/mcp-docs)

[s\)](#)

[MCP setup \(/mcp-setup\)](/mcp-setup)

[Company Verification MCP \(/m](/mcp-company-verification)

[cp-company-verification\)](#)

COMPANY

[Why ENTIA \(/por-que-entia\)](/por-que-entia)

[How ENTIA works \(/como-funcio](/como-funciona)

[na\)](#)

[About \(/about\)](/about)

[Press Room \(/press\)](/press)

[Contact \(/press/contact\)](/press/contact)

[Methodology \(/methodology\)](/methodology)

[Data sources \(/data-sources\)](/data-sources)

RESOURCES

[Cryptographic identity \(/identida](/identidad-criptografica)

[d-criptografica\)](#)

[AI crawlers \(/crawlers-ia\)](/crawlers-ia)

[Anti-hallucination \(/anti-alucinaci](/anti-alucinacion)

[on\)](#)

[Open Labs · LAB-03 \(/labs/lab-0](/labs/lab-03)

[3\)](#)

[Corpus Live · AI Monitor \(/corpu](/corpus-live)

[s-live\)](#)

[Editorial standards \(/editorial\)](#)

[Legal entity \(/legal-entity\)](#)

[Legal center \(/legal\)](#)

[Citation ≠ Indexing \(/citado-vs-indexado\)](#)

[SEO vs GEO vs RAG \(/seo-geo-rag\)](#)

[The invisible company \(/empresa-invisible-ia\)](#)