

Twelve operators, one by one

Each operator received on 31 August or 1 September 2026 its individual file: the sealed record of what its agent requested and received, a forensic report with the methodology and its limits, and in most cases a licence proposal. The dispatch went to the channel each operator publishes for such matters. What follows is what happened next, read from the mailbox, not recalled. This page contains no amounts and no commercial valuations.

12

notified

with proof of dispatch anchored on each message identifier

6

replied

2 remediated · 1 opened a commercial channel · 3 replied without addressing the question

6

have not replied

5 model providers and 1 search engine; from Amazon only an automatic reply arrived



answers to the question

how they identify and treat a machine-readable reservation that is not robots.txt

GENERAL-PURPOSE AI MODEL PROVIDERS

No reply as of 12 September

Five of the six operators without a reply are AI model providers. Three of them (OpenAI, Google and Amazon) signed the copyright chapter of the GPAI Code of Practice, checked against the Commission's list on 10 September, and thereby committed to identify and comply with machine-readable rights reservations (Measure 1.3) and to designate a point of contact for rightsholders (Measure 1.5). Meta and Apple are not on that list; Article 53(1)(c) of the AI Act applies to them all the same, without the presumption of conformity the Code provides.



OpenAI

NO REPLY

GPTBot · GPAI Code signatory · notified 31 August at dpo@openai.com and dsar@openai.com · follow-up 3 September

The strongest case in the file, and not by volume. GPTBot retrieved from OpenAI's network, with 200 responses, `tdm-policy.json` and `tdmrep.json` on 30 July 2026 at 21:54 and 21:55 UTC. OpenAI-SearchBot did the same on 23 August and ChatGPT-User on 4 and 6 September. GPTBot is the agent whose declared purpose, by the operator itself, includes training. OpenAI signed every chapter of the Code. It has not written a line.



Google

NO REPLY

GoogleOther · GPAI Code signatory · notified 31 August at partnerships@google.com and gemini-partnerships@google.com · follow-up 3 September

The file sent covers GoogleOther, the agent Google documents for uses other than search. Googlebot has its own sealed record and has deliberately not been contacted: ENTIA's machine-readable licence authorises indexing by search engines and names Google. From Google's network there is no retrieval of the reservation files or the legal pages, with the full filter. One of the addresses used on 3 September (search-console-team@google.com) bounced. Complete silence.

**Amazon****AUTOMATIC REPLY ONLY**

Amazonbot · GPAI Code signatory · notified 31 August at amazonbot@amazon.com, copying aws-EU-privacy@amazon.com · follow-ups 3 and 9 September

On 9 September ENTIA redirected the file to the single point of contact Amazon publishes for the European Union (amazon-dsa-compliance@amazon.com). The automatic reply states that the address serves only Member State authorities, the Commission and the European Board for Digital Services. Amazonbot retrieved only HTML legal pages, indistinguishable from ordinary indexing. No person has answered.

**Meta****NO REPLY**

Meta-ExternalAgent · not listed as signatory · notified 31 August at partnerships@meta.com and llama-partnerships@meta.com, copying webmasters@meta.com · follow-up 3 September

Meta-ExternalAgent is the agent with a declared dual purpose, including training, and with the most solid attribution of the whole record (99.98% from Meta's networks). It retrieved the HTML legal pages on 6, 7 and 9 September; there is no request to the machine-readable files. Nothing in writing.

**Apple****NO REPLY**

Applebot · not listed as signatory · notified 31 August at applebot@apple.com · follow-up 3 September

Apple documents that what Applebot crawls may be used to train its foundation models, with opt-out through Applebot-Extended. Applebot retrieved tdm-policy.json and two legal pages from Apple's network. As a non-signatory there is no Measure 1.3 to invoke, only Article 53(1) (c). Nothing in writing.

**Huawei****NO REPLY**

PetalBot · not listed as signatory · notified 31 August at petalbot@huawei.com, copying legal@aspiegel.com · follow-up 3 September

PetalBot feeds Petal Search; the operator declares search, not training. The file sent is an access record and a forensic report, without a licence proposal. Nothing in writing.

REPLIED WITHOUT ADDRESSING THE QUESTION

Those who answered something else



Anthropic

REPLIED WITHOUT ADDRESSING THE QUESTION

ClaudeBot · GPAI Code signatory · notified 31 August · Copyright Operations Team reply on 2 September · ENTIA's questions on 3 September, unanswered

On 27 July 2026, between 15:28 and 15:30 UTC, ClaudeBot retrieved from address space registered to Anthropic PBC the four reservation files, with 200 responses on all of them. On 2 September Anthropic's Copyright Operations Team replied that its crawlers are controlled through robots.txt, that a site can add a Disallow directive and that a form exists to request inclusion on its blocklist. ENTIA had asked for no block: its robots.txt has a group dedicated to ClaudeBot with Allow and an ai-train=no signal, sealed on 19 July with an independent timestamp.

On 3 September ENTIA asked whether that reply represents Anthropic's position on Article 4(3) reservations, which machine-readable protocols it recognises beyond robots.txt, what operational effect the 27 July retrieval had and, if its position is that no reservation has effect without a Disallow, that it confirm this and reconcile it with Measure 1.3 of the Code it signed. It asked for escalation to the team responsible for Article 53(1)(c). As of 12 September there is no reply. Measured on 10 September: zero requests from Anthropic's network to the dossiers and zero requests to the reservation files since Policy V2 was published on 9 September.



Microsoft

REPLIED; COULD NOT ROUTE THE MATTER

Bingbot · GPAI Code signatory · notified 31 August at eureg@microsoft.com, copying bingbot@, bwtsupport@ and bingwb@microsoft.com · three Bing Webmaster Support tickets

Bingbot is the family with the most verified-entity records served in the record (453,243 between 19 July and 1 September) and the only one where the registered network holder matches without a single contradiction. Microsoft acknowledged receipt and replied seven times across three tickets. On 20 August it asked ENTIA to verify that the traffic was really Bingbot and to share the logs, evidence the file already contained. On 3 September Bing Webmaster Support (ticket UCM000007475380) stated that it handles organic search and webmaster tool matters. On 9 September it declared intellectual property, database rights, text and data mining and data licensing out of scope and said it was unable to route the request to the appropriate team. EUReg@microsoft.com replies automatically that it handles only product safety and regulatory authority queries; bingwb@microsoft.com bounced the message because it accepts only internal senders. It arrived; there is nobody to take it.



Ahrefs

ROUTED TWICE; THE TEAM NEVER WROTE

AhrefsBot · not a model provider · notified 31 August at support@ahrefs.com · five messages received

On 22 August support indicated that the matter needed the attention of its legal and business development team; on 25 August, that it had forwarded the details and that team would reach out if interested. On 3 September, that it cannot take a position on an external record or a licensed-access proposal, and pointed to its API plans. On 9 September ENTIA clarified that it is not seeking access to the Ahrefs API, and support again passed the details to the appropriate team. That team has never written.

Those who remediated or opened a channel

None of these three is a general-purpose AI model provider, so the GPAI Code does not reach them by definition. Their route is contractual and bilateral, and they are the only ones who have acted. Ceasing to crawl is not a bilateral resolution: future cessation and the historical record are different things, and that has been put to them in writing.

D

DataForSEO

CEASED AND DELETED; DECLINES LICENCE

DataForSeoBot · formal notice 19 August · file 31 August · legal position 4 September · thirteen messages

On 21 August it declined the licence. On 28 August it confirmed that it had stopped crawling entia.systems and removed the related data from its database, that its bot retains only hyperlinks and that it does not use the content for training. On 3 September it specified that crawling stopped on 26 August and passed the case to its legal representative. On 4 September Andrii, at legal@dataforseo.com, set out its position: the activity complied with robots.txt, it rejects ENTIA's legal characterisation and any licensing obligation, and considers the matter closed. Still outstanding: the signed confirmation ENTIA requested on 28 August and the single question ENTIA put on 10 September: what is its policy when faced with a TDM reservation expressed by machine-readable means.

MOZ

Moz

BLOCKLIST; DECLINES LICENCE

DotBot · file 1 September · reply from Andrew Hignett (Help Team, ticket 228066) on 1 September · eleven messages

Moz added entia.systems and its subdomains to its blocklist, stated that it collects only public link-graph data and anchor text for ranking metrics, that it does not store or commercialise the underlying records, declined the licence and considered the matter resolved. What ENTIA measured: DotBot's last access was on 24 August at 18:33 UTC, eight days before that email, with a positive control of other crawlers active every day. It is recorded as a declared cessation, technically corroborated; eight days of absence do not prove it cannot return. On 3 September Moz referred to its public documentation and declined to disclose its internal engineering.



Semrush

COMMERCIAL CHANNEL OPEN, NO PROGRESS

SemrushBot · file 31 August · Maria Kudlai, Director of Data Partnerships, Datas (a Semrush company), on 1 September

The only commercial opening of the whole series. On 19 August support had said that Semrush does not purchase or license third-party web data. On 1 September Maria Kudlai wrote that she was interested in learning more about the proposal, suggested a call that week and asked to be the single point of contact. ENTIA replied the same day accepting the designation, keeping the procedure in writing, as with every operator, and offering to correct any figure that does not match Semrush's own records. ENTIA followed up on 3 and 9 September. No reply since 1 September. SemrushBot is the second family by verified-entity records served (273,113) and retrieved tdm-policy.json and tdmrep.json from Semrush's network.

How this page was measured

Dispatches are read from ENTIA's sent mail, anchored on each message identifier. Replies were read in full on 11 September 2026 across all messages received since 19 August, searching by sender domain across the whole mailbox and not only the dispatch thread: one reply from Anthropic arrived in an old support thread and was found that way. A message in sent mail proves it left, not that it arrived; where an acknowledgement or a reply exists, it is stated.

What this page does not say

It does not assert that any operator trained on the corpus or that there is infringement. It does not turn silence into guilt: it records it as absence of reply as of a stated date. It publishes nothing about any negotiation beyond what the operator itself wrote. When an operator provides its copy and it differs from what is described here, it will be corrected with a new dated version and the previous one preserved.



ENTIA

DEFINING TRUST IN AI

PrecisionAI Marketing OÜ

Sepapaja tn 4, 11415 Tallinn, Estonia

Registry code 17048063 · DUNS 565868914

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

s://w

EU data · Estonia e-Residency

www.ifi

nkedi

s://x.

s://gi

License data (/licensing)

Machine ecosystem (/ecosyste

m)

Machine telemetry (/telemetry)

AI Discovery Observatory (/obs

ervatory)

TDM reservation and

European Commission (/press/

tdm-reservation-european-commis

sion)

Status by operator (/press/operat

ors)

Five-Layer Framework (/press/r

esearch/five-layer-framework)

DEVELOPERS

API (/developers)

MCP documentation (/mcp-doc

s)

MCP setup (/mcp-setup)

Company Verification MCP (/m

cp-company-verification)

PLATFORM

Entia Home (/products/entia-hom

e)

Score (/auditorias/)

Sector Intelligence (/sectores)

Knowledge (/knowledge/es/)

SOLUTIONS

For Advertisers (/advertisers)

For Agencies (/partners/)

AI search visibility (/funnel?fresh

=1)

COMPANY

Why ENTIA (/por-que-entia)

How ENTIA works (/como-funcio

na)

About (/about)

Press Room (/press)

Contact (/press/contact)

Methodology (/methodology)

Data sources (/data-sources)

Editorial standards (/editorial)

Legal entity (/legal-entity)

Legal center (/legal)

RESOURCES

Cryptographic identity (/identida

d-criptografica)

AI crawlers (/crawlers-ia)

Anti-hallucination (/anti-alucinaci

on)

Open Labs · LAB-03 (/labs/lab-0

3)

Corpus Live · AI Monitor (/corpu

s-live)

Citation ≠ Indexing (/citado-vs-in

dexado)

SEO vs GEO vs RAG (/seo-geo-r

ag)

The invisible company (/empres

a-invisible-ia)

[COOKIES \(/LEGAL/ES/POLITICA-COOKIES\)](/LEGAL/ES/POLITICA-COOKIES) [TERMS \(/LEGAL/ES/TERMINOS-SERVICIO\)](/LEGAL/ES/TERMINOS-SERVICIO)

[DPA \(/LEGAL/ES/PLANTILLA-DPA\)](/LEGAL/ES/PLANTILLA-DPA) [CORPORATE STATUS \(/LEGAL/EN/CORPORATE-STATUS\)](/LEGAL/EN/CORPORATE-STATUS)

[CENTRO LEGAL \(/LEGAL\)](/LEGAL)