COMPLIANCE BRIEF · 15 PAGES

EU AI Act, GDPR & ANSSI for production LLMs.

How a runtime LLM proxy maps to Article 15 of the EU AI Act, GDPR Chapter V and the 35 ANSSI-PA-102 recommendations — written for European B2B SaaS teams who shipped a chatbot, a copilot or an agent before the AI Office was watching.

NEXT ENFORCEMENT DATE

2 August 2026

GPAI fines · Art. 99

NEXT REPORTING DATE

11 September 2026

CRA incident & vuln. reporting

GAP ON THE GROUND

64 % vs 7 %

policy vs expertise · Wavestone 2025

How to read this brief

This document is the technical answer for one specific class of system: a B2B SaaS feature that calls a third-party LLM (OpenAI, Anthropic, Mistral, Gemini, OpenRouter) at runtime. It maps the obligations a European deployer owes under the AI Act, GDPR and the ANSSI-PA-102 guide to operational controls a runtime proxy can — and cannot — absorb. It is built to be forwarded to a DPO and cited in a procurement questionnaire.

WHAT YOU'LL FIND

FOR THE CTO READING ON A PHONE

The dates, the mapping table, and the self-assessment checklist (sections 1, 2, 3) are written to scan. Each table row carries its own meaning; the headings carry the gist.

FOR BOTH READERS

Section 4 lists seven questions every European LLM buyer should put to their model provider. Use it independently of any decision about Senthex.

FOR THE DPO READING CAREFULLY

Sections 5 (ANSSI mapping), 6 (out of scope), and 7 (certification honesty) are written to be defensible against a vendor questionnaire. Article numbers are cited verbatim and ANSSI titles are quoted from the official 29 April 2024 publication.

WHAT THIS BRIEF IS NOT

Not a substitute for legal advice. Not a conformity-assessment template. Not a survey of the AI Act in general — only of the obligations relevant to a deployer of a third-party GPAI model.

Table of contents

Cross-reference with the live page. This brief is consistent with `senthex.com/en/eu-ai-act` as of 2 May 2026. The page is canonical: regulatory dates, ANSSI mappings and certification status will move there first. The QR code on the back cover returns you to it.

SECTION 01

The four dates that actually trigger procurement

GPAI obligations have been applicable since **2 August 2025**, but the AI Office's supervision and enforcement powers — document requests, evaluations, fines up to **3 % of global turnover or €15 million** — only switch on **2 August 2026**. Two weeks later, the **Cyber Resilience Act** triggers reporting duties for actively exploited vulnerabilities and severe incidents on **11 September 2026**. For most deployers of a third-party GPAI model, those are the two procurement-shaping dates — and **2 February 2025** (Article 50 transparency for consumer-facing AI) is the one already past.

Timeline — the dates that bind a B2B SaaS shipping an LLM feature

2 Feb 2025 PAST — **Article 50 transparency obligations applicable.** End-users must be told when they interact with an AI system; deepfakes and AI-generated text on matters of public interest must be labelled. *In scope for any consumer-facing chatbot, copilot or voice agent.*

2 Aug 2025 PAST — **GPAI obligations applicable to model providers.** Technical documentation under Annex XI, copyright policy, transparency summary of training content. *In scope for OpenAI / Anthropic / Mistral / Google as upstream providers* — your downstream effect is that you can request the documentation that propagates through Article 53(1)(d).

2 Aug 2026 SOON — **AI Office enforcement powers + GPAI fines applicable.** Fines up to **€15 M or 3 % of world-wide annual turnover** under Article 101 for GPAI providers; up to **€35 M or 7 %** for prohibited practices under Article 99. *In scope for your deployer obligations under Articles 25–26 and, if you operate a high-risk Annex III system, Article 15.*

11 Sep 2026 SOON — **Cyber Resilience Act reporting duties applicable.** Manufacturers of products with digital elements must report actively exploited vulnerabilities within 24 h to ENISA, then a vulnerability notification within 72 h, and a final report within 14 days. *In scope for any SaaS shipped as a product with digital elements.*

2 Aug 2027 LATER — **Legacy GPAI compliance deadline.** GPAI models placed on the market before 2 August 2025 must be brought into compliance. *In scope for legacy fine-tuned deployments inherited from a predecessor.*

What every SaaS in scope needs by 2 August 2026

A defensible answer to four questions: **what does the model see**, **what does it return**, **who can prove it**, and **what happens when something is exploited**. The mapping table on the next page rewrites those four questions as concrete pieces of evidence — the kind that survive a procurement review.

> **Reading note for the DPO.** The €15 M / 3 % cap (Article 101) applies to GPAI providers, not deployers. Deployer fines (Article 99) are tied to specific obligations — prohibited practices, transparency, high-risk duties — and graduated by Article 99(7) proportionality. The €15 M figure on the cover is the headline; the operative number for most B2B SaaS deployers comes from Article 99(4).

Sources. EU AI Act Implementation Timeline, AI Office, `artificialintelligenceact.eu/implementation-timeline`. AI Act Service Desk FAQ, European Commission, `ai-act-service-desk.ec.europa.eu/en/faq`. Cyber Resilience Act, European Commission DG CONNECT, `digital-strategy.ec.europa.eu/en/policies/cyber-resilience-act`. ENISA single-reporting platform notice (CRA Art. 14).

SECTION 02

Feature × regulation mapping

The table below is the row a procurement officer asks for: **which technical control** answers **which obligation**, and **which evidence** you can hand over to a DPO. Coverage means the proxy provides the technical control; it does **not** mean you are exempt from documenting the control in your risk register or DPIA. Use this table as the spine of your vendor questionnaire response.

SENTHEX CONTROL	EU AI ACT	GDPR	ANSSI-PA-102	EVIDENCE TO PRODUCE
Zero-data-retention by default	Recital 69 (privacy by design)	Art. 5(1)(c) Art. 25	R7	Architecture diagram showing request/response bodies are forwarded in-flight, not persisted; retention schedule for metadata only (timestamp, latency, shield verdicts, token counts).
EU-only inference path (Hetzner Falkenstein, DE)	Recital 134	Chapter V (Art. 44-50)	R11	DC location attestation (Hetzner ISO 27001 certificate); routing diagram showing the Senthex hop never leaves the EU; downstream provider region selected by you (this is your decision, not ours).
26 input/output shields (prompt injection, jailbreak, PII, secrets, toxicity, off-topic, hallucination markers, etc.)	Art. 15(5) cybersecurity	—	R25, R27, R33	Per-call shield verdicts in the audit log, including rule fired, score and rule version. Map each shield to the OWASP LLM Top 10 risk it addresses (LLM01 prompt injection, LLM02 sensitive information disclosure, etc.).
Per-request audit log (Business and Enterprise tiers)	Art. 12 record-keeping	Art. 30 records of processing	R23, R29	Tamper-evident log: timestamp, route, shield verdicts, token cost, response status. Raw input/output excluded by design — the controller-to-data-subject link stays in your systems. Exportable to your SIEM via API or signed S3 archive.
Real-time dashboard + alerts (Slack, webhook, email)	Art. 14 human oversight	—	R26	Operational signal feeding the human-in-the-loop your high-risk system is required to have. Documented oversight procedures remain your responsibility — the dashboard is a feed, not a process.
OWASP LLM Top 10 posture report (Business and Enterprise)	Art. 9 risk management	—	Annex residual risks	Monthly per-app posture summary mapped to the OWASP LLM Top 10 (2025). Helps populate the risk register; does not write it for you.
Multi-provider routing	Art. 25 value chain	Art. 28 processor	R26	DPA documenting Senthex as a controller-aware processor. Pass-through routing — we never re-purpose model traffic. Audit log unified across providers (one place to query OpenAI, Anthropic, Mistral and Gemini calls).

How a DPO reads this table. Coverage in a single column does not equal compliance with the article. Article 15, Article 9 and Article 12 are **system-level** duties; the proxy provides controls and evidence the broader system needs to cite. The conformity assessment, the technical documentation under Article 11, and the residual-risk acceptance remain on you.

SECTION 02 · CONTINUED

How to use this table in a vendor questionnaire

A typical European vendor security questionnaire from a buyer in scope of the AI Act will ask, in some order:

Where is the data processed and under which legal mechanism?	`Row 2 (EU-only path)`
What is your data retention policy for the prompts, the responses and the meta-data?	`Row 1 (zero retention)`
Which Article 15 cybersecurity controls do you operate against prompt injection, jailbreak and output filtering?	`Row 3 (shields)`
Can you provide a per-request audit trail for our DPO and our SOC?	`Row 4 (audit log)`
What human-oversight signals do you expose under Article 14 of the AI Act?	`Row 5 (dashboard)`
How does your OWASP LLM Top 10 posture feed our Article 9 risk-management system?	`Row 6 (posture report)`
Are you a processor under Article 28 GDPR — and what happens to our model traffic across multiple providers?	`Row 7 (routing + DPA)`

A concrete worked example — one Annex III row

Consider a recruitment SaaS using a GPT-class model to score CV-to-job matches. This system falls in Annex III §4 (employment) and is therefore **high-risk**. The deployer owes:

- **Article 9 (risk management)** — a documented, lifecycle risk-management system. *The proxy contributes* the OWASP posture report (row 6) and the per-shield verdicts (row 3). *The proxy does not write* the risk register itself.
- **Article 12 (record-keeping)** — automatic event logs throughout the system's lifetime. *The proxy contributes* the per-request audit log (row 4) for the inference hop. *The proxy does not log* user-facing recruiter actions or candidate-side events.
- **Article 14 (human oversight)** — a recruiter must be in the loop on every adverse decision. *The proxy contributes* alerts on shield-blocked calls (row 5). *The proxy does not implement* the human-in-the-loop UI.
- **Article 15 (cybersecurity)** — resilience against attempts to alter inputs, outputs or performance. *The proxy contributes* the 26 shields (row 3). *The proxy is not* a cybersecurity certificate; it is a control among others.
- **GDPR Article 22** — automated decision-making affecting a data subject is restricted unless one of three legal bases applies. *The proxy does not address this* — it is a controller-side legal exercise.

> **The reading the brief is built to support.** The proxy is **evidence-producing**. It is not, on its own, **conformity-producing**. A DPO who reads this brief and then asks *"so we still need to write our Annex IV technical documentation, right?"* — has read it correctly.

SECTION 03

Self-assessment — 18 questions to score your own setup

Score each row honestly. The point is not the total — it is the rows you score **No** or **Partial**. A reader who scores poorly should want to fix something; a reader who scores high should be able to defend each row in front of an auditor.

A. DATA FLOW & RESIDENCY (ROWS 1–4)

01	We can produce, on request, a one-page diagram showing every hop a prompt takes between the user, our app, our LLM provider and any third party — including geographic location of each hop.	☐☐☐☐
02	Our DPA with the LLM provider names a specific data-processing region and we can demonstrate that the request never transits another region under normal operation.	☐☐☐☐
03	We have written down what is **retained** and what is **not retained** at each hop, with a defined retention schedule for what is retained.	☐☐☐☐
04	If we use a runtime proxy or gateway, its hosting jurisdiction is covered by an appropriate Chapter V GDPR transfer mechanism (or it sits in the EEA).	☐☐☐☐

B. CYBERSECURITY & ARTICLE 15 CONTROLS (ROWS 5–10)

05	We can name the controls we operate against prompt injection (LLM01) — input filtering, output filtering, system-prompt hardening, or a combination — and produce evidence for each.	☐☐☐☐
06	We can name the controls we operate against sensitive information disclosure (LLM02), including PII detection and secrets redaction.	☐☐☐☐
07	We can produce, for any single user-facing call from the last 30 days, a complete trace: which model was called, with which system prompt version, which shields fired, what the verdict was.	☐☐☐☐
08	We have a documented response procedure for a confirmed prompt-injection incident — who is notified, how the affected users are identified, what is logged.	☐☐☐☐
09	Our LLM-feature error budget is monitored alongside our application SLOs, with an explicit threshold for shield false-positive rate.	☐☐☐☐
10	We have run, in the last 6 months, at least one red-team exercise specifically targeting the LLM feature (not the surrounding app).	☐☐☐☐

C. RECORDS, OVERSIGHT & RISK (ROWS 11–14)

11	Our Article 30 GDPR records of processing list the LLM-feature processing as a distinct activity, with a named purpose, retention, and recipients.	☐☐☐☐
12	If our system is high-risk under Annex III, we have an Article 9 risk-management system that is updated at least annually and that names the LLM-specific risks (hallucination, prompt injection, training-data lineage).	☐☐☐☐
13	We have a documented human-oversight procedure under Article 14 — and the human in the loop has the technical means to intervene before an adverse decision is finalised.	☐☐☐☐
14	Our internal audit, on the LLM feature, can be performed without depending on the cooperation of the LLM provider.	☐☐☐☐

D. VENDOR MANAGEMENT & SUPPLY CHAIN (ROWS 15–18)

15 For each LLM provider we use, we have a current DPA that covers the LLM-specific processing — not just a generic SaaS DPA from 2022. ☐☐☐☐

16 For each LLM provider we use, we have a documented position on the joint-controller question (Articles 4(7) and 26 GDPR) — even if the conclusion is "we treat them as a processor and accept the residual risk". ☐☐☐☐

17 We have a documented downgrade path: what happens to the user-facing feature if our primary LLM provider is unavailable for 24 h, 7 days, indefinitely (also addresses ANSSI R15). ☐☐☐☐

18 We have a position, on file, regarding the AI Act conformity assessment status — even if that position is "we do not place a high-risk system on the market and therefore Article 43 does not apply". ☐☐☐☐

SCORING — EACH ROW ☐ ☐ ☐ ☐ = YES · PARTIAL · NO · N/A

How to read your score

15+ Yes

Well ahead of the median European B2B SaaS deploying an LLM feature. Use the brief to compare your evidence pack against a procurement reviewer.

10–14 Yes

Typical for a team that has shipped an LLM feature in the last 18 months. The Partial and No rows are your runway to 2 August 2026.

5–9 Yes

Common for teams whose LLM feature was added by product without DPO involvement. The next 90 days are the pragmatic window.

0–4 Yes

Your LLM feature is exposed. Start with rows 1, 5, 7, 11, 15 — they give you the most defensibility per hour of work.

On the 64 % vs 7 % gap. The Wavestone AI Cyber Benchmark 2025 found that **64 %** of European enterprises have defined an AI security policy, while only **7 %** have the in-house expertise to operationalise it (`wavestone.com/fr/insight/ai-cyber-benchmark-2025`). The 18 questions above are written to make that gap concrete — most teams discover the policy answer is "yes" while the operational answer is "we'll need to check". A runtime proxy narrows the gap; it does not, on its own, close it.

SECTION 04

Seven questions to ask your LLM provider

Whichever proxy you do or do not use, these are the seven questions a European deployer should put — in writing — to OpenAI, Anthropic, Mistral, Google, OpenRouter or whichever GPAI provider is on the other end of the line. Keep the answers; they go in your Article 30 records and they pre-empt half of a vendor questionnaire.

01 In which region(s) is our inference traffic processed by default, and what is the legal mechanism for any cross-border transfer?

Why — Defines your Chapter V GDPR posture. "In the EU on a best-effort basis" means the burden of proof stays with you.
What "good" looks like — Named region (e.g. EU-Ireland), with a contractual commitment that traffic does not transit other regions, plus the SCC version or other Article 46 transfer tool in force.

02 Are our prompts, completions, embeddings or fine-tuning datasets used to train, evaluate or improve any model — opt-in or opt-out, by default and after we change the setting?

Why — The opt-out default for OpenAI API and the opt-in default for Anthropic API are different. Your DPA needs to reflect the actual setting on your account.
What "good" looks like — Explicit reference to the API setting (not the consumer-product setting), the effective date, and a statement of what is retained "for safety review" and for how long.

03 What is the Article 53 GPAI provider documentation summary you publish, and where can our DPO retrieve it?

Why — GPAI providers owe transparency about training-data summary and copyright policy under Article 53(1)(c)–(d). A deployer is entitled to receive this.
What "good" looks like — A URL or attached document, dated after 2 August 2025, naming the model series in scope and the date of the most recent update.

04 What is your incident-notification SLA for a confirmed model-side incident affecting our traffic — model regression, data-handling incident, confirmed jailbreak class?

Why — An Article 28(3)(f) GDPR processor obligation, and a CRA consideration from 11 September 2026.
What "good" looks like — A named SLA in hours, a named contact channel, and a commitment that the notification reaches you whether or not the incident is publicly disclosed.

05 What controls do you operate against prompt injection and jailbreak in your current model, and what is your 12-month roadmap — including version-pinning options?

Why — Article 15(5) asks for resilience against alteration of inputs / outputs. Your provider's controls compose with yours.
What "good" looks like — Specific controls named (input classifiers, output filters, training-time mitigations), plus a commitment that you can pin a model version for at least 6 months.

06 Can we obtain a sub-processor list with their geographic locations, and a notification mechanism when it changes?

Why — Article 28(2) GDPR. The list often surfaces a US sub-processor the headline "EU-hosted" claim does not mention.
What "good" looks like — A current list with locations, plus a commitment to email you ahead of a material change with a window long enough to object.

07 Will you sign our DPA, and if not, which changes to your standard DPA are non-negotiable from your side?

Why — "We only sign our standard DPA" is itself an answer. The non-negotiable clauses tell you which risks the provider declines to hold.
What "good" looks like — Written confirmation of which clauses are accepted and which are not, with reasons. A flat refusal to discuss is a procurement signal.

SECTION 05

Where a runtime proxy fits in the data path

Print-friendly schematic. Every arrow is labelled; nothing is conveyed by colour alone. The bottom block is the audit hop — it does not carry user data, only the metadata listed.

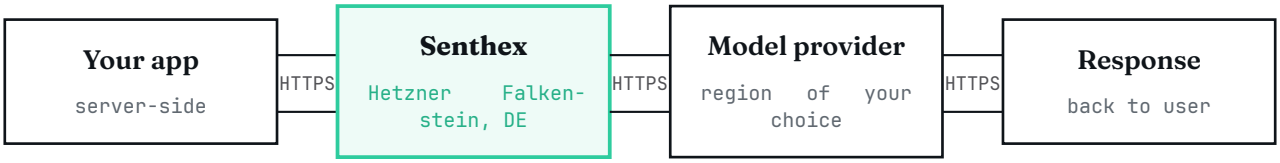

Senthex applies, in this single hop: input shields → forward to provider → output shields → return to your app. The request and response bodies are **never** persisted. Only the metadata below is written.

What is logged, what is not

LOGGED (PER REQUEST)

- Timestamp (ISO 8601, ms precision)
- Route identifier (your app's internal name)
- Provider and model name (e.g. `anthropic/claude-3-5-sonnet`)
- Token count (input, output)
- Latency (Senthex hop, provider hop, total)
- Shield verdicts (rule fired, score, rule version)
- HTTP status and any error code
- Request/response **hash** (for tamper-evidence)

NOT LOGGED

- Request body (the prompt itself)
- Response body (the model output)
- End-user identifiers (unless you pass them as metadata you control)
- Document attachments forwarded to a multimodal model
- Tool-call payloads in agentic flows
- Any field you mark as `x-senthex-redact` in the request header

On the secrets-redaction shield (opt-in, off by default), the original-vs-rewritten diff **is** retrievable by you for legal-basis reasons; otherwise outputs are not stored.

The architectural claim that survives a procurement review. "Our runtime proxy is in the EU." — incomplete. "Our runtime proxy is in the EU and the DPA names the region, the sub-processor list is the Hetzner Falkenstein DC, and the audit log holds metadata under a 90-day retention with a tamper-evident hash." — defensible.

SECTION 06

ANSSI-PA-102 — the seven recommendations a runtime proxy touches

The ANSSI guide *Security recommendations for a generative AI system* (ANSSI-PA-102, 29 April 2024) lists 35 recommendations across the lifecycle of an AI system. A runtime proxy can contribute to seven of them. Titles below are quoted verbatim from the official English publication; the commentary is ours.

R7 — "Manage data confidentiality issues from the AI system design phase"

A zero-data-retention runtime proxy is a load-bearing answer here: prompts and completions are not persisted at the proxy hop. The recommendation extends to the model provider — the proxy does not absorb provider-side caching, which remains a DPA-side question between you and OpenAI / Anthropic / Mistral / Google.

R11 — "Host the AI system in trusted environments consistent with security needs"

The Senthex hop is hosted on Hetzner Falkenstein (DE), an ISO 27001–certified data centre. ANSSI explicitly elevates SecNumCloud (R14) for sensitive use cases — Senthex is **not** SecNumCloud-qualified; deployers needing SecNumCloud should route to a SecNumCloud-qualified provider end to end.

R20 — "Protect the integrity of AI system files"

The proxy contributes the audit-log integrity dimension: per-request hash, append-only retention, signed export. Integrity of the **model files** themselves remains the model provider's responsibility under their own change-management process.

R23 — "Conduct security audits of AI systems before deployment to production"

The proxy makes audits practical: per-request shield verdicts and posture reports give an auditor something to inspect without depending on the model provider's logs (which most enterprise plans do not expose at request granularity).

R26 — "Manage and secure the interactions of the AI system with other business applications"

Multi-provider routing and tool-call traceability are the operational form of this recommendation for a runtime hop. As MCP-based tool calls become standard in 2026, this row will widen to cover the audit of tool invocations alongside completions.

R27 — "Limit automatic actions performed by an AI system handling uncontrolled inputs"

The shields that block before forward (prompt injection, jailbreak, off-topic) are a runtime contribution to this recommendation. The proxy cannot stop your application from acting on a passed-through completion — the action-side mitigation belongs in your code.

R33 — "Strengthen security measures for AI services hosted on the Internet"

Internet-exposed LLM features are explicitly singled out by ANSSI for hardened controls. The shield catalogue (input/output filtering, secrets redaction, jailbreak detection) and the EU-resident audit trail are the proxy's contribution. Other R33 measures — WAF, rate limiting, authentication — remain in your edge stack.

> **What ANSSI-PA-102 is and is not.** It is a **guide**, not a regulation. It is non-normative — ANSSI states explicitly that the recommendations are delivered "as is" and that adaptation to the target IS is expected (see *Avertissement*, page 1 of the official document). French RSSI teams treat it as a de-facto check-list because no equivalent national guide existed before 29 April 2024; treat your mapping as evidence of due diligence, not as a certificate.

Source: ANSSI-PA-102, *Security recommendations for a generative AI system*, 29 April 2024, English version, 35 recommendations. Available at `cyber.gouv.fr/en/publications/security-recommendations-generative-ai-system`.

SECTION 07

Out of scope — what a runtime proxy cannot absorb

Three classes of duty are not absorbable by any runtime hop. Pretending otherwise is how vendors lose buyers in the first technical review. This page is what we say **to** a procurement reviewer who asks "and?".

01 · Model documentation and conformity assessments

If your system falls under Annex III (recruitment, credit scoring, education ranking, critical infrastructure, justice, etc.), you owe **technical documentation under Article 11** and **conformity assessment under Article 43**. A runtime proxy cannot fill those forms. We can hand you operational evidence — logs, posture reports, shield definitions, version history — that you cite in section 4 of Annex IV. The document itself, signed by the provider and notified to the conformity-assessment body, is on you.

Example. A health-tech SaaS deploying an LLM-based triage assistant under Annex III §5 will need a notified-body assessment under Article 43(1). The proxy contributes evidence; the dossier is the SaaS's own work and typically an outside specialist's engagement.

02 · Training-data lineage and bias testing

Article 10 (data governance) applies to **providers** of high-risk systems, not to a runtime hop. If your team fine-tunes a model — even a small LoRA over a public base — you own that lineage end-to-end, including the data subjects' rights under Article 13(2)(f) GDPR and the training-set documentation. A proxy that observes only inference traffic has no way to reconstruct training-set composition.

Example. If a recruitment SaaS fine-tunes a base model on its own historical hiring data to score candidates, the bias evidence required by Article 10(2)(f) is **not** derivable from runtime logs. It comes from the training pipeline and from periodic adversarial testing of the fine-tuned model.

03 · Joint-controller analysis with your model provider

GDPR Articles 4(7) and 26 are a legal exercise between you and the model provider. The exercise is a written position your DPO maintains, supported by the DPA in force, the actual data flow (which the proxy can document), and the residual-risk acceptance from your controller. Senthex's DPA covers our **processor relationship to you**. It does not adjudicate your relationship with OpenAI, Anthropic, Mistral or Google.

Example. If a fintech SaaS uses an LLM for AML alert triage and the provider performs safety classification on the prompts (a common provider-side step), a serious DPO will ask whether that classification creates a joint-controllership situation. The answer is provider-specific, depends on the DPA in force, and is not absolved by a proxy.

> **The honest framing.** A runtime proxy is a **defence-in-depth control** for the inference hop and an **evidence layer** for what you can show an auditor. It is not a **legal layer**. The question "what does a runtime proxy not give us" is a feature of the brief, not a hole in it.

SECTION 08

Honest about limits — what we are not yet certified for

Senthex is an early-stage company shipping a production product. The infrastructure is built to compliance standards; the certification paperwork is not all in place yet. Stating this in the brief — and not only in a footnote — avoids the awkward call later. A buyer who notices is reassured, not deterred.

SOC 2 Type II	NOT HELD	Not certified at the corporate level. Roadmap target — late 2026 / early 2027 once paid customer count justifies the audit cost. A buyer for whom SOC 2 is a hard procurement gate today should choose a SOC 2-certified vendor; we will not pretend otherwise.
ISO 27001	NOT HELD	Not certified at the corporate level. Same window. Hetzner Falkenstein, our hosting layer, is ISO 27001-certified independently — but corporate ISO is a separate scope and we do not conflate the two.
HDS (French health data hosting)	NOT HELD	Senthex is not appropriate for clinical-data flows that require HDS hosting. Use a HDS-qualified provider for those routes; the runtime-proxy layer cannot rescue the certification gap.
SecNumCloud (ANSSI)	NOT HELD	Not qualified. Deployers needing SecNumCloud (e.g. public-sector OIVs, sensitive ANSSI R14 use cases) should route end to end on a SecNumCloud-qualified stack — this is not a proxy-resolvable gap.
GPAI provider obligations (AI Act Art. 53)	NOT APPLICABLE	Senthex does not train, distribute or place a general-purpose AI model on the market. Article 53 obligations apply to model providers (OpenAI, Anthropic, Mistral, Google, Meta, etc.), not to a downstream runtime hop.
AI Act conformity assessment	NOT APPLICABLE	Not applicable to a runtime proxy at the time of writing. The relevant harmonised standards (CEN-CENELEC JTC 21) are still in draft. We monitor them and will publish a position when finalised.
DORA (financial-sector resilience)	OUT OF SCOPE	Senthex does not provide financial services. DORA's third-party-ICT register may capture us if you are a regulated financial entity using us — handle that registration on your side; we will provide the DORA-relevant information on request.

> **The procurement-honest reading.** "100 % AI Act compliant" is a marketing claim no vendor in this category can substantiate today — the harmonised standards are still in draft and no conformity-assessment body issues certificates for runtime LLM controls. If you hear that phrase from any vendor, ask which body issued the certificate. There is not one yet.

SECTION 09

Glossary

For the DPO walking into this brief without an LLM-engineering background, and for the CTO who has not had to write the acronyms down before. One line each.

AI Act	Regulation (EU) 2024/1689 on artificial intelligence, in force since 1 August 2024.
AI Office	The Commission body within DG CNECT supervising GPAI providers; enforcement powers applicable from 2 August 2026.
Annex III	The list of high-risk AI systems in the AI Act — recruitment, credit scoring, education, critical infrastructure, law enforcement, migration, justice.
Annex IV	The technical-documentation contents required for high-risk systems (cited via Article 11(1)).
Article 5(1)(c) GDPR	Data minimisation principle. The reference for "store no more than necessary".
Article 12 / 14 / 15 (AI Act)	Record-keeping (12), human oversight (14), accuracy/robustness/cybersecurity (15) for high-risk systems.
Article 28(3)(g) GDPR	The processor's obligation to delete or return personal data at the end of the service.
Article 30 GDPR	Records of processing activities — the controller's register.
Article 50 (AI Act)	Transparency obligations — informing end-users of interactions with an AI system. Applicable since 2 February 2025.
CEN-CENELEC JTC 21	The European joint technical committee drafting harmonised AI standards under the AI Act.
CRA	Cyber Resilience Act — Regulation (EU) 2024/2847 on horizontal cybersecurity for products with digital elements. Reporting duties from 11 September 2026.
DORA	Digital Operational Resilience Act — financial-sector ICT-risk regulation, in application since 17 January 2025.
DPA / DPIA	Data Processing Agreement (Art. 28 GDPR) and Data Protection Impact Assessment (Art. 35 GDPR — mandatory for high-risk processing).
FRIA	Fundamental Rights Impact Assessment — Article 27 AI Act, for certain deployers of high-risk systems.
GPAI	General-Purpose AI model — defined in Article 3(63) AI Act; obligations under Articles 53–55.
HDS	Hébergement de Données de Santé — French certification for health-data hosting.
ISO 27001 / SOC 2 II	International ISMS standard / audited US operational-controls report. Common procurement gates.
MCP	Model Context Protocol — open protocol for connecting AI models to tools, contributed to the Linux Foundation in late 2025.
NIS2	Directive (EU) 2022/2555 on a high common level of cybersecurity. Member-state transposition since 17 October 2024.
OWASP LLM Top 10	Reference list of the most critical LLM-application risks; current version 2025.
RAG	Retrieval-Augmented Generation — pattern of retrieving documents at query time and feeding them to the LLM as context.
SCC	Standard Contractual Clauses — the European Commission's tool for cross-border GDPR transfers under Art. 46(2)(c).
SecNumCloud	ANSSI's French qualification for cloud providers handling sensitive data.

SECTION 10

What Senthex is and is not

Senthex is

A **runtime proxy** between your application and a third-party LLM provider, hosted in a **single EU data centre** (Hetzner Falkenstein, Germany), operating **26 input/output shields** and producing a **per-request audit log** you can export to your SIEM.

A **processor** under Article 28 GDPR, with a DPA available on request before signature, naming our processor relationship to you and committing to the standard processor obligations including Article 28(3)(g).

An **evidence layer**. We give you the operational artefacts — shield verdicts, posture reports, signed audit exports — that you cite in the documents your system owes (Annex IV technical documentation, Article 30 records, Article 9 risk-management evidence).

Built and maintained by a small team in France. Solo-founder origin, Cyber Booster cohort, currently shipping. Team size and certification roadmap are public on the live page.

Senthex is not

Not a **certification**. Not a substitute for a conformity-assessment body, an Annex IV dossier, or a notified-body engagement.

Not a **legal opinion**. The joint-controller analysis with your model provider, the residual-risk acceptance, and the FRIA where applicable remain a controller-side exercise.

Not a **guarantee** against prompt injection or jailbreak. The OWASP LLM Top 10 lists LLM01 precisely because no filter is exhaustive. We expose per-shield verdicts, scores and rule versions so the residual risk is **measurable**, not absent.

Not **SOC 2 Type II**, not **ISO 27001**, not **HDS-qualified**, not **SecNumCloud-qualified** at the corporate level today. A buyer for whom one of those is a hard gate today should choose a vendor that holds it.

Not the **only** defence-in-depth layer your LLM feature should have. Provider-side guardrails (Bedrock, Azure Content Safety, Vertex), edge controls (WAF, rate limiting) and your application's own checks are complementary.

How to test the claim before talking to anyone

The Senthex **Free plan** (1,000 requests/month, no credit card) gives you the same 26 shields and the dashboard. The €49/month **Pro** plan adds 90-day log retention; **Business** at €199 adds the AI Act audit trail and the OWASP LLM Top 10 report; **Enterprise** covers signed S3 archives, custom retention, and the high-risk-deployment scoping. For high-risk Annex III or HDS-adjacent flows, a scoped DPIA conversation precedes the proxy decision — reach out via the contact form on the live page.

> **The reader this brief was written for.** A CTO of a 50–200 person scale-up reading this on a phone Tuesday morning, and a DPO of a 500–2000 person company reading it carefully Thursday afternoon. If a row above does not survive your review, tell us — we will rewrite the row before the next version. The page at `senthex.com/en/eu-ai-act` is canonical.

ONE LAST PAGE

The regulatory landscape will keep moving. The page won't.

Scan the code to return to the live mapping at senthex.com/en/eu-ai-act. Dates, ANSSI references and certification status will surface there first. This PDF is a snapshot dated 2 May 2026 — the page is the source of truth.

OPEN ON A PHONE

senthex.com/en/eu-ai-act

TALK TO US

senthex.com/en/contact

VERSION

v2026-05-02 · EN · 15 pages