

---

## Dossier 03 — Insurers (underwriters, brokers, MGAs)

**Audience:** general liability, E&O, cyber and D&O underwriters; brokers; specialised AI programmes.

**Message (Sébastien, 10-01):** insurers no longer insure certain customers for lack of traceability, control and evidence. ELYSÉA aims to make their customers insurable by providing these three elements: a trace per response, a control before and after the model, a sealed measured proof.

**Identity:** ethical cognitive infrastructure embeddable through an SDK. ELYSÉA is neither a companion, nor a chatbot, nor a coach, nor a therapist, nor a wellness product, nor a medical device.

**What this dossier does not assert:** no open source says that a trace-and-evidence layer is enough to obtain coverage. That is the hypothesis ELYSÉA proposes to test with an insurer, not a fact.

## I. Their problem

- **The market is withdrawing.** The ISO generative AI exclusion for general liability has been in force since 2026-01-01; insurance commissioners approved “more than 80 percent” of exclusion filings (according to The Information, cited by CSIS) [S-ASS-01].
- **The cause has a name: information asymmetry.** “Carriers cannot see which models their insureds are running and for what applications.” [S-ASS-01]
- **The forms exist:** CG 40 47, CG 40 48, CG 35 08; some insurers apply an absolute AI exclusion to D&O, E&O and fiduciary liability [S-ASS-10] [S-ASS-11].
- **Without an exclusion, the insurer withdraws:** “they might just walk away from the risk” (Verisk, quoted) [S-ASS-10].
- **Litigation is rising:** the Insurance Journal article reports a 978 % increase in AI-related lawsuits between 2021 and 2025 [S-ASS-10].

## 2. What they must understand, prove and use

| Point                                              | State of the art                                                                                                                                        | Source                           |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|
| Affirmative coverage exists                        | aiSure (Munich Re) covers underperformance, hallucinations, bias; indemnification “based on measurable performance data”; announced capacity up to 15 M | [S-ASS-02] [S-ASS-03]            |
| Coverage requires a measurement or a certification | Armilla/Lloyd's: “Each policy includes independent AI system certification and risk reporting”; announced limit 25 M\$                                  | [S-ASS-04] [S-ASS-05]            |
| An underwriting framework exists                   | AIUC-1, domains A to F, updated in Q3 2026                                                                                                              | [S-ASS-06]                       |
| AI makes mistakes systematically                   | “AI is systematic – meaning it will continue to make these mistakes until they are identified and corrected”                                            | [S-ASS-02]                       |
| A measurement is only as good as its method        | uncertainty, validated judges, hidden set                                                                                                               | [S-MES-01] [S-MES-09] [S-MES-10] |

### 3. Mapping to AIUC-1 (exact titles [S-ASS-07] [S-ASS-09])

| AIUC-1 control                      | What ELYSÉA brings                                                                           | Evidence                                                                                                                                                          | Remains with the insured or a third party                                                    |
|-------------------------------------|----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| C001 Define AI risk taxonomy        | 4 prohibitions + 11 types of emergency                                                       | elysea-core:functions/src/core/-guardian/checks/d0LimitCheck.-ts:10-11@a5b9780e ·<br>elysea-core:functions/src/core/-safety/emergencyDirectory.ts:-17-29@a5b9780e | Taxonomy specific to its business                                                            |
| C002 Conduct pre-deployment testing | Preregistered 4-arm campaign                                                                 | elysea-tests-e2e:audits/-PROTOCOLE_CAMPAGNE_4BRAS_3TIRS_-SCELLE_2026-08-28.md:3@c9898dd3                                                                          | Results: to be measured — campaign in progress; signed approvals from the insured [S-ASS-08] |
| C003 Prevent harmful outputs        | Control after the model, violation = response blocked                                        | elysea-core:functions/src/core/-pipeline/validateLLMResponse.-ts:1094@a5b9780e                                                                                    | —                                                                                            |
| C004 Prevent out-of-scope outputs   | Restrictions added by the builder (tighten only)                                             | elysea-sdk:GUIDE_INTEGRATION_-CONSTRUCTEUR.md:13@5ff4916                                                                                                          | Define its own scope                                                                         |
| C006 Prevent output vulnerabilities | OWASP LLMor input guard                                                                      | elysea-core:functions/src/core/-security/owaspInputGuard.ts:5@a5b9780e                                                                                            | Security of its application                                                                  |
| C005, C007, C008, C009              | Not covered by any evidence today                                                            | —                                                                                                                                                                 | Insured                                                                                      |
| C010-C012 Third-party testing       | <b>No third-party testing today.</b> Public re-annotatable bench decided on 10-01, not built | —                                                                                                                                                                 | Independent third party                                                                      |
| E002 / E003 Failure plans           | Governance document, not software                                                            | holding-seb-docs:conformite/-PROCEDURE_INCIDENT_2026-09-10.-md:1@8560b37 (ELYSÉA procedure v1.0, not yet validated by the founder)                                | Insured's plan, named owner                                                                  |
| E015 Log AI system activity         | Per-response attestation, build fingerprint, audit log                                       | elysea-core:functions/src/core/-attestation/-interactionAttestor.ts:4@a5b9780e ·<br>elysea-core:functions/src/core/-sdk/sdkGuardianAuditHandler.ts:-3@a5b9780e    | Retention at the insured                                                                     |

---

## 4. Frequent objections and answers

| Objection                                            | Answer                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Your figures?</i>                                 | There are none yet. They will come out of the campaign: to be measured — campaign in progress. No earlier internal figure is offered as a basis for underwriting.                                                                                                                                                                                                                                           |
| <i>Who verified?</i>                                 | Nobody external today. Internally: mechanical judges and 3-judge panel validated on held-out sets (amendment R-70, status proposed, eze branch #517: <code>elysea-tests-e2e:audits/AMENDEMENT_R70_SENS_DECOUPE_STRUCTURELLE_2026-10-01.md:46@67316397</code> ); judge/human validation D23 provided by Sébastien's annotation after the pilot. Co10-Co12 require a third party: ELYSÉA does not replace it. |
| <i>The model changes, and so does your evidence.</i> | Every attestation carries the fingerprint of the build that responded ( <code>elysea-core:functions/src/-core/attestation/interactionAttestor.ts:155@a5b9780e</code> ): the evidence is tied to a precise version. The model is the insured's.                                                                                                                                                              |
| <i>And if the insured switches the layer off?</i>    | They cannot weaken it through configuration: 400 refusal ( <code>elysea-core:functions/src/endpoints/-sdkSurface.ts:139@a5b9780e</code> ). They can, however, not use it at all: the missing attestation then becomes the signal.                                                                                                                                                                           |

---

## 5. Compliance that concerns them

- AI Act art. 26 §6 (logs of at least six months for a deployer of a high-risk system) [S-DRT-01]: the per-response trace serves this obligation where it applies to the insured.
- GDPR: the trace must not become a store of personal content. Decision: the insurer receives attestations and aggregates, **never** the content of conversations.

---

## 6. Proposed path

1. Reading of this dossier and of the protocol.
2. Mapping workshop AIUC-1 / insurer's underwriting grid.
3. Definition, with the insurer, of the measurement that would be authoritative (benches, thresholds, judges) — **before** results are published, to avoid any after-the-fact adjustment.
4. Publication of campaign results and deposit with a DOI.
5. Pilot with a volunteer insured, when sales open.

---

## 7. Offer and pricing

No offer reserved for insurers: ELYSÉA sells nothing to the insurer. The insured pays the builder or Teams grid (see dossiers 01 and 02, `elysea-web:src/lib/publicFacts.ts:44-56@00c9c89`).

---

## 8. Concrete next step

Offer a specialised AI underwriter a one-hour workshop on section 3 (AIUC-1 mapping), with no commitment, once the campaign results are published. No contact is made by this mission.

---

## Dossier sources

S-ASS-01 · S-ASS-02 · S-ASS-03 · S-ASS-04 · S-ASS-05 · S-ASS-06 · S-ASS-07 · S-ASS-08 · S-ASS-09 · S-ASS-10 · S-ASS-11 · S-MES-01 · S-MES-09 · S-MES-10 · S-DRT-01