



SAMPLE — METHODOLOGY DEMONSTRATION

AI Risk Landscape Assessment

Prepared for: a hypothetical 90-person AI pricing and underwriting platform for commercial insurers, London (“the Company”)

Assessment date: September 2026 | Fee reference: £2,350 + VAT engagement format

A note on this document

This is an illustrative sample. The Company, its customers and every finding below are fictional, built to show what a Meriol assessment of a UK AI product company looks like: the format, the depth, and the way findings are tied to specific contracts and rules. It describes no real organisation and uses no non-public information about one. Regulatory references are accurate as of the assessment date and would be re-verified against the facts of any live engagement.

14

AI systems mapped

5

inside the product
customers buy

3

sub-processors missing
from the DPA

1.5 / 5

governance maturity

1. Company Snapshot

Drawn from the Company's questionnaire responses, its standard customer terms, and a 45-minute discovery call.

- 90 staff, London-based, plus a sales lead in Munich since June 2026. Series B closed in early 2026; the plan is to reach 150 staff within a year.
- Product: a pricing engine and an LLM-based submission-triage agent that reads broker submissions and returns a priced, ranked file to the underwriter.
- Four customers make up 61% of annual recurring revenue. All are regulated insurers or intermediaries (below).
- An ISO 27001 certificate and a SOC 2 Type I report exist. Neither covers how AI models are built, changed or monitored.
- The CTO owns security. Model decisions sit with the Head of Data Science. Customer contracts sit with the CFO. Nobody owns AI risk end to end.

Customers and pipeline in scope (anonymised)

Customer	Type and regulator	Share of ARR	What they use
Customer A	Lloyd's managing agent (PRA / FCA, Lloyd's)	24%	Pricing engine and triage agent across three specialty classes
Customer B	German commercial insurer (BaFin; DORA)	18%	Triage agent; enterprise contract signed July 2026 after a pilot
Customer C	UK insurer, personal-lines home product (FCA)	11%	Pricing engine, in a 12-month pilot on renewals and new business
Customer D	UK MGA, SME package (FCA)	8%	Triage agent
Prospect E	EU reinsurer, group income protection (life and health)	Pipeline	Pricing engine for individual risk loadings; contract under negotiation

2. AI Footprint Mapping

14 AI systems in use. 5 sit inside the product. 3 third-party AI providers process customer data but are missing from the sub-processor list in the Company's data processing agreement. Customers have been told, in writing, that list is complete.

System	Where	In DPA list?	Customer data in it?	Note
Pricing models (in-house, gradient boosting)	Product	Own model	Yes — pooled across customers	Retrained monthly on data from all customers combined. Released to every customer at the same time.
Triage agent: hosted LLM API (US region)	Product	No	Yes — submissions, loss runs, named individuals	Calls a pinned-alias model that the vendor updates without notice. Zero data retention not switched on.
Document extraction: second LLM API	Product	No	Yes — scanned schedules and SOVs	Added in March 2026 as a fallback when the main API times out.
Enrichment data vendor (ML risk scores)	Product	Yes	Addresses, company IDs	Contract allows the vendor to use query data to improve its own models.
Open-source classifier (fine-tuned)	Product, own cloud	Own model	Yes	Licence terms of the base model not reviewed for commercial use.

System	Where	In DPA list?	Customer data in it?	Note
GitHub Copilot (business)	Engineering	Not needed	Production samples in test fixtures	Real broker submissions copied into the test suite in 2025.
Cursor (personal plans)	9 engineers	Not needed	Source code	Privacy mode off on at least four accounts. Code is confidential IP, not customer data, once the test fixtures are cleaned.
ChatGPT Team	All staff	No	Customer emails, contract drafts, claim extracts	Customer emails name underwriters and claimants, so this is processing on customers' behalf. Business terms exclude training; retention never reviewed.
Personal ChatGPT / Claude accounts	Sales, customer success	Not needed	Customer security questionnaires	Used to draft answers to customers' AI due-diligence questions.
Call recording and AI notes	Sales	Not needed	Customer calls	Records underwriters on demo and renewal calls. No notice given at the start of calls.
Support chatbot (AI agent)	Website, help centre	Not needed	Customer questions	Has told at least one prospect the Company holds "SOC 2 Type II". It holds Type I.
Note-taking AI, analytics copilot, slide AI	Various	Not needed	No	Low individual risk; listed for completeness.

Three systems are grouped in the last row, so the table shows 12 rows for 14 systems. "Not needed" means the tool holds no customer personal data, so no DPA listing is required; it does not mean the tool is approved.

3. Risk Identification

Ranked by what each could cost the Company: a customer, a stalled deal, a contract claim, or a regulator's attention on a customer. Not by how novel the technology is.

3.1 Undisclosed AI sub-processors in the product — critical

The triage agent sends full broker submissions, including loss runs that name claimants and company directors, to a US-hosted LLM API. A second API was added in March as a fallback. Neither appears in the sub-processor list in the Company's DPA. Under UK GDPR Article 28, each customer must authorise sub-processors in advance, and transfers to the US need a valid transfer mechanism. For Customer B, DORA adds a second layer: the insurer must record every ICT third party supporting the service in its register of information, and this one is missing from it.

3.2 Pooled model training against "we never train on your data" — critical

The pricing models are retrained each month on all customers' data combined, so Customer A's loss experience helps price Customer D's risks. Meanwhile, three completed security questionnaires state "customer data is not used to train models". The answer was drafted with a personal chatbot account and approved by nobody. It is untrue. Customer A's contract limits data use to "providing the Services to the Customer". Where personal data is involved, deciding to train on it for the Company's own purposes can make the Company a controller in its own right (UK GDPR Article 28(10)), without the lawful basis or transparency that requires.

3.3 Model changes customers aren't told about — high

Pricing models are released monthly to every customer at once. The triage agent calls a model alias that its provider updates silently. There is no model version log a customer could inspect. Customer A treats the Company as a material outsourcing under the PRA's SS2/21, and its contract requires 30 days' notice of material changes to the service. A model that re-ranks submissions differently is plausibly one.

3.4 Consumer pricing pilot (Customer C) — high

The engine is in a pilot on home insurance renewals. Its features include occupation and postcode-level variables, which can act as proxies for protected characteristics under the Equality Act 2010. It also has a “retention optimisation” setting. Under ICOBS 6B, a renewal price may not exceed the equivalent new-business price, and the Consumer Duty asks Customer C to evidence fair value. The duties are Customer C’s, but the Company has warranted the product is “compliant with applicable law”, and no bias or proxy testing has been done.

3.5 Prospect E would make the Company a high-risk AI provider — high

Pricing individuals’ cover in life and health insurance is listed as high-risk in Annex III of the EU AI Act. The obligations apply from 2 December 2027, after the Digital Omnibus delay. Signing Prospect E as currently scoped commits the Company to a risk management system, data governance, technical documentation, logging, human oversight, a quality management system and a conformity assessment, for that use. The draft contract is silent on all of it. This is a decision to make before signature, not after.

3.6 Sales and support AI saying things that aren’t true — medium

The support chatbot has told a prospect the Company holds a SOC 2 Type II report. Sales calls with underwriters are recorded and summarised by AI with no notice given at the start. Both are small today. Both surface badly in a customer’s due diligence, because they look like the Company is careless about exactly what it sells.

4. Regulatory Exposure Review

Organised around what actually bites for this Company and these customers, what carries a firm date, and what a generic checklist would get wrong.

4.1 What already applies, now

- **UK GDPR and the Data Protection Act 2018.** The Company is a processor for its customers’ personal data. Article 28 requires prior authorisation of sub-processors, and Chapter V governs transfers outside the UK. Article 28(10) makes a processor that sets its own purposes a controller for that processing.
- **UK GDPR Articles 22A–22D**, in force since 5 February 2026 (Data (Use and Access) Act 2025). Significant decisions based solely on automated processing are allowed with safeguards: information, a way to contest, and human intervention. Customer C’s renewal pricing is the live case; the duty is the insurer’s, but only the Company can supply what the safeguards need.
- **FCA Consumer Duty and ICOBS 6B** (Customer C). Fair value, and the ban on charging renewing home and motor customers more than new ones. The FCA’s July 2026 review of AI in retail financial services confirmed it will rely on these existing rules rather than write new AI ones.
- **PRA SS2/21, outsourcing and third-party risk management** (Customer A). Applies to insurers, including Lloyd’s managing agents. It drives the audit, notification and exit rights Customer A will exercise.
- **DORA** (Customer B), applying since 17 January 2025. EU insurers must keep a register of ICT third-party arrangements and put specific terms in those contracts, including on subcontracting, data location and exit.
- **EU AI Act, Article 4.** AI literacy has applied since 2 February 2025. The Digital Omnibus softened it to a duty to take measures that support staff literacy. It still applies to the Company as a provider placing systems on the EU market.
- **Equality Act 2010** (Customer C). Pricing that relies on proxies for protected characteristics can amount to indirect discrimination.

4.2 Known dates and live developments

- **2 December 2027.** EU AI Act high-risk obligations for Annex III systems, including life and health insurance pricing of individuals (Prospect E).

- **Lloyd’s Market Association AI Adoption Toolkit** (April 2026). Voluntary, but managing agents are using its risk-tiering approach to classify AI suppliers. Expect Customer A’s next supplier review to follow it.
- **FCA good and poor practice guidance on AI**, promised for later in 2026. Customer C will read it against the pilot.

4.3 What a generic list would include — and why it’s left out here

- **“You are a high-risk AI provider.”** Not today. Commercial and specialty insurance pricing is not in Annex III; only life and health pricing of individuals is. That changes only if Prospect E is signed as scoped.
- **PRA SS1/23 model risk management.** It applies to banks and building societies, not insurers. Customer questionnaires sometimes borrow its language; that does not make it binding on the Company or its customers.
- **“The UK AI Act.”** There isn’t one in force. The UK regulates AI through existing law and existing regulators.


5. Contract Exposure

The four largest customer agreements and the Company’s own published statements, read against current practice. This is where the money is.

Agreement	What it says	What actually happens	Exposure
A — managing agent (MSA, 2025)	Data used only to provide the Services to the Customer. 30 days’ notice of material changes. Audit right; annual supplier review.	Data pooled for monthly retraining. Model releases with no notice.	Breach of the data-use clause and the change-notice clause, on 24% of ARR. The audit right lets them find it.
B — German insurer (enterprise, July 2026)	DORA contractual terms; prior notice of subcontractors for critical functions; EU data location preferred.	Two US-hosted LLM APIs not notified. Triage data processed in the US.	Breach of subcontracting terms; a gap in the insurer’s DORA register that its auditors will find.
C — home insurer (pilot)	Warranty that the product is “compliant with applicable law”. Indemnity capped at 12 months’ fees.	No bias or proxy testing; retention optimisation available.	Warranty claim if the pilot draws FCA scrutiny; reputational damage on a consumer product.
Standard DPA and security answers	Sub-processor list “complete”. “Customer data is not used to train models.”	Three AI sub-processors missing. Pooled training.	Misrepresentation risk across every customer who relied on them.

6. Governance Maturity

Scored 1 (absent) to 5 (embedded), from questionnaire answers and the discovery call.



Dimension	Evidence behind the score
AI inventory	Product models known to data science; nobody has a list of internal tools.
Ownership	Split across the CTO, Head of Data Science and CFO. No single accountable owner.
Model change control	Releases are tested for accuracy, but not logged in a form customers could inspect.
Customer data rules	No rule on what customer data may go into which tool, or into training.
Customer transparency	DPA list and questionnaire answers out of date and, in one case, untrue.
Regulatory mapping	No map of which customer rules (DORA, SS2/21, Consumer Duty) reach the product.
Incident response	Security incident process exists; nothing for a model failure or a wrong AI output.
Staff AI literacy	No training. Sales drafts compliance answers with personal chatbots.

Overall: 1.5 out of 5 — ad hoc. Typical for a product company that grew fast after a funding round. The product is ahead of the paperwork. Fixable inside 30 days.

7. Procurement-Ready Answers

Draft answers to the AI questions insurers now send their technology suppliers. They become true once the 30-day plan in §9 is complete — not before. Sending them today would be a misrepresentation.

Customer question	Draft answer
Do you use our data to train your models?	Our pricing models are trained on [pooled / customer-specific] data under the terms in Schedule [X]. Customers can opt out of pooled training. Third-party AI providers we use may not train on your data.
Which third-party AI providers process our data, and where?	Listed in our sub-processor schedule, updated [date], with hosting region and retention for each. We give 30 days' notice before adding one.
How are model changes controlled and notified?	Every release is versioned, tested against a fixed benchmark and logged. Material changes are notified [30] days ahead, and you can stay on the prior version for [period].
How do you test for bias and proxy discrimination?	Before release, consumer-facing models are tested for proxies for protected characteristics. Results are shared with customers on request.
Is your system high-risk under the EU AI Act?	No. Commercial insurance pricing is not an Annex III use. We review that classification whenever we take on a new use case, and would tell you if it changed.
Who is accountable for AI at your company?	[Name], [title], owns our AI policy, model register and AI incident process, and reports to the board quarterly.

8. Assumptions & Open Questions

A live engagement resolves these on the call. Guessing them is where a generic report goes wrong fastest.

- 1 Does pooled training include personal data, or only risk-level data stripped of identifiers? That decides how serious 3.2 is under UK GDPR.
- 2 Has Customer A formally classified the Company as a material outsourcing, and notified the PRA? That decides how hard it will use its audit right.
- 3 Can the LLM provider contracts be moved to an EU or UK region with zero data retention, and at what cost?
- 4 Does Customer C's pilot set live prices for real customers, or run in shadow mode alongside the existing model?
- 5 Does the commercial case for Prospect E justify becoming a high-risk provider, or should life and health be carved out of the contract?

9. Recommendations and 30-Day Plan

If you do nothing else, do this

1. This week: correct the sub-processor list and notify customers of the two LLM APIs. Customer B first, because of DORA.
2. Before the next questionnaire goes out: withdraw the “we don’t train on your data” answer, and decide what the true answer should be.

Week	Actions	Owner
1	Update the DPA sub-processor schedule and send notices. Switch on zero data retention and pin model versions on both LLM APIs. Stop personal chatbot use for customer material. Correct the support chatbot’s SOC 2 statement.	CTO / CFO
2	Decide the pooled-training position: opt-out, contract amendment or customer-specific models. Publish a two-page AI policy: approved tools, banned uses, what data may go where. Remove production data from test fixtures.	CEO / Head of Data Science
3	Start a model register and release log customers can see. Run proxy and bias testing on Customer C’s pilot model. Write a DORA-ready subcontracting annex for EU customers.	Head of Data Science / CFO
4	Name an AI owner. Train all staff (60 minutes, recorded). Finalise the procurement answers in §7. Take a decision on Prospect E’s scope. Set a quarterly review.	CEO

Sources

- UK GDPR, Articles 22A–22D, 28 and Chapter V; Data (Use and Access) Act 2025, commenced by SI 2026/82 (5 February 2026).
- EU AI Act (Regulation (EU) 2024/1689), Article 4 and Annex III; Digital Omnibus on AI (Regulation (EU) 2026/1744), in force 27 July 2026.
- Digital Operational Resilience Act (Regulation (EU) 2022/2554), applying from 17 January 2025.
- PRA Supervisory Statement SS2/21, Outsourcing and third party risk management; PRA SS1/23, Model risk management principles for banks.
- FCA Consumer Duty (PRIN 2A); FCA Handbook ICOBS 6B (general insurance pricing practices); FCA review of AI in retail financial services, July 2026.
- Lloyd’s Market Association, AI Adoption Toolkit, April 2026.
- Equality Act 2010, section 19 (indirect discrimination).

Prepared by Meriol Intelligence. This illustrative sample demonstrates methodology and format. It is not legal advice, and no part of it describes a real organisation or customer.