INSIGHT · AI & CONTRACTS

Reading the AI vendor DPA: ten clauses your standard template doesn't cover

By Shalini Jhunjhunwala · Aequita Advisory · 9 min read

Most enterprise procurement teams use the same Data Processing Agreement they have used for years. It works fine for ordinary SaaS. It does not work for AI, and the gap shows up six months into the engagement, when the vendor has trained on your data, your outputs sit in a model you do not own, and your DPO discovers a sub-processor nobody disclosed.

The hard part is not that AI vendor contracts need different clauses. The hard part is that they need clauses that are *not asked for* in the standard procurement checklist. A vendor first-draft will satisfy the checklist completely and still leave the buyer exposed.

Below are the ten clauses I look for in any AI vendor contract, and the negotiating trap each one tends to come with.

Why the standard template fails

A 2024 analysis of AI vendor contracts found that 92% of vendors claim broad data-usage rights, only 33% offer indemnification for third-party IP claims, and only 17% commit to full regulatory compliance. These numbers are materially different from established SaaS contract norms. The EU AI Act's Article 25 compounds the issue by establishing a shared-responsibility model between providers and deployers, meaning that an underspecified contract sits on both parties' books, not just the supplier's.

The combined effect is that "standard practice" no longer protects either side adequately. The contract has to specifically address what makes AI different.

The ten clauses

Listed in rough order of how often they are missing from vendor first drafts.

1. Training-data restrictions. Explicit prohibition on using customer data to train, fine-tune, or improve the vendor's models, including aggregated, anonymised, or derived forms. Default vendor terms permit training. You must opt out in writing. The negotiating trap: vendors offer "no training on identifiable data", which leaves a wide opening.

2. Data ownership and derived artefacts. Clear ownership of inputs, outputs, embeddings, fine-tuning data, logs, and any model weights derived from your data. Vendor licence-back should be narrow and time-bounded. The trap: vendors retain broad rights over "feedback and usage data", which often covers everything you put through the model.

3. Model lineage and sub-processor transparency. Disclosure of base models, foundation-model providers, and any transitive AI sub-processors. Required for accurate Records of Processing Activities under GDPR Article 30. The trap: vendors claim "trade secret" protection over model architecture, which is reasonable up to a point, but not for the identity of the underlying foundation models.

4. Third-party IP indemnification. Vendor indemnifies you for third-party claims arising from training data or model outputs. Uncapped or super-capped for IP infringement. The trap: only one in three AI vendors offers this today. If they refuse, that itself is the answer to your risk question.

5. EU AI Act role allocation (Article 25). Explicit allocation of provider, deployer, importer, or distributor roles per system. Identification of any high-risk classification under Annex III. Commitment to share technical documentation needed for your role. The trap: vendors push role allocation downstream by "deeming" the customer to be the provider whenever convenient.

6. Transparency and documentation rights. Access to model cards, system cards, technical documentation, evaluation results, and specifically the Annex IV documentation required for high-risk systems. The trap: documentation is provided "on reasonable request" with no committed turnaround. Useless during a regulator query.

7. Bias, audit, and testing rights. Right to request bias and fairness testing results. Right to commission independent audits of the system as deployed for you. The trap: audit rights restricted to security audits only, not algorithmic or bias auditing.

8. Human oversight commitments. Where the system supports decisions affecting individuals, the vendor commits to features and configurations that enable meaningful human review, not just a notional override button. The trap: human oversight defined as the customer's sole responsibility, even when the system architecture prevents it.

9. International transfers and data residency. GDPR Chapter V compliance: identified transfer mechanism, named destination jurisdictions, special handling for inference traffic to non-EU model endpoints. The trap: vendors disclose "hosting" locations but not inference routing. Data may leave the EU for processing even with EU storage.

10. Termination, data deletion, and model decommissioning. On termination: return or deletion of all customer data, deletion certificate within a stated period, and critically, decommissioning of any models or fine-tunes derived from customer data. The trap most often missed: deletion clauses cover data but not derived models. Your data may be gone, but the model trained on it may persist indefinitely.

What to do with this in practice

Run the vendor's first draft against the ten clauses. Each missing clause is a redline. Each trap phrase is a negotiation point.

The goal is not to win every clause. It is to know which ones you are knowingly conceding, and which ones the vendor refused to acknowledge at all. The second category is where the real risk sits, because if a vendor will not negotiate on training data restrictions or IP indemnification, that tells you something concrete about what their business model depends on.

> *"The clauses your standard template doesn't cover are the ones that determine whether the engagement protects you or quietly transfers risk to you."*

Closing thought

AI vendor contracts are still moving territory. The vendors are improving their first drafts. The standard templates will catch up eventually. Until they do, the buyer who reads carefully and pushes for the right ten clauses ends up with a contract that actually reflects the risks the engagement creates.

A useful test before signing: if a regulator asked you, six months from now, what protections you have around the vendor's use of your data, the training of their models, and the obligations they have under the EU AI Act, could you point to specific paragraphs in the contract? If yes, the contract is doing its job. If not, the homework is not done.

Aequita Advisory
Privacy & AI Compliance · EU practice

hello@aequita.eu
Shalini Jhunjhunwala, Founder