aptem.
Intelligent technology®

Seven questions to ask your AI supplier

A minimum-viable AI assessment

A simplified starter for organisations that are not yet ready to run a full AI Suitability and Impact Assessment (AISIA). Use this as a floor, not a ceiling.

Provided by Aptem

Version: 1.0

Owner: [Assessment Owner]

Status: [Draft / Active]

Date: [Day Month Year]

Audience: customer-facing starter template

Follow us on LinkedIn and YouTube

www.aptem.co.uk

29 May 2026

aptem.
Intelligent technology®

Contents

Follow us on LinkedIn and YouTube

www.aptem.co.uk

29 May 2026

aptem.
Intelligent technology®

How to use this template

This is a deliberately simple document. Seven questions to put to any supplier offering an AI-supported tool or service, plus a few bonus items to bring the assessment into your wider governance.

Where this fits

The recommended path is a full AI Suitability and Impact Assessment (AISIA), paired with a Data Protection Impact Assessment (DPIA) where personal data is involved. The AISIA covers implementation risks, third-party AI system risks, user impact, operational risk, and sign-off. It is the standard your governance should aim for.

For less mature organisations, this seven-question version is the minimum. It is not a substitute for the full AISIA. As your governance matures, move to the full template, and treat this version as a fast-track screen for low-risk tools.

How to use it

1. Complete one assessment per supplier or use case. The same supplier with two different use cases needs two assessments, because the risk picture is different.
2. Put each question to the supplier in writing. Do not accept verbal answers for the consequential ones (questions 3, 4, 5, and 7).
3. Record the supplier's answer and link to any documentation they provide (terms of service, sub-processor list, data protection addendum, security white paper, model card).
4. Pair with a DPIA if personal data is processed. The seven questions inform the DPIA; they do not replace it.
5. Route for sign-off (see Section 9).

A note on audience

This template is written for training providers, with explicit reference to learners, learner data, and DfE expectations. The same questions apply to other regulated sectors, and the language can be generalised if you are not a training provider.

aptem.
Intelligent technology®

Overview

A lightweight metadata block. Replace each placeholder with your own content.

Field	Value
Supplier or tool	[Supplier name] - [tool URL]
Use case	[Brief one-line description of how you will use the tool]
Assessment date	[DD Month YYYY]
Assessed by	[Name, role]
AI risk tier (provisional)	[AI-Low / AI-High / AI-Prohibited]
Personal data processed?	[Yes / No. If Yes, a DPIA is required.]

aptem.
Intelligent technology®

Question 1. Whose model is this really?

Whose model is this really? Yours, the supplier's, or someone else's via an API?

Why it matters

Many AI products you encounter run on someone else's model. If your supplier's product uses OpenAI, Anthropic, Google, Microsoft, or any other foundation model under the hood, your data flow extends to that third party. Their terms, their data residency, their training policies become part of your data-protection picture.

If you cannot find out whose model sits underneath, treat that as a red flag. The supplier should be able to answer this in a sentence.

What a good answer looks like

- Names the underlying model provider (for example, OpenAI GPT-4, Anthropic Claude, Google Gemini, an in-house model, or a fine-tuned open-source model).
- References the contractual arrangement between the supplier and the underlying model provider (typically an enterprise API agreement).
- Confirms that the third-party model provider's terms are reflected in your supplier's terms with you, including data handling and training data restrictions.
- If the supplier uses a hosted enterprise tier of a foundation model (for example, Microsoft Azure OpenAI, Anthropic via AWS Bedrock), names that hosting environment explicitly.

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Question 2. What did you train it on, and can you prove it?

What did you train it on, and can you prove it?

Why it matters

This question tells you whether the training data is the supplier's, third-party, or scraped from the open web. It matters for IP, bias, copyright, and accuracy. It also matters because regulators, learners, and your own customers are increasingly asking the same question of you.

You are unlikely to get a full disclosure of training data. You should be able to get a credible high-level summary, supporting documentation, and a contractual position on copyright claims.

What a good answer looks like

- A documented description of training data sources at the category level (for example, licensed datasets, supplier-owned data, public web data, synthetic data).
- Auditable evidence available on request, such as model cards, system cards, third-party audits, or transparency reports.
- A stated position on copyright and IP claims arising from training, including any indemnity offered to customers.
- A clear statement on whether the model was fine-tuned on customer data and, if so, whose.

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Question 3. What happens to what I put in?

What happens to what I put in? Do you keep it? Train on it? Share it? Sell it?

Why it matters

This is the most consequential question by some distance. The answer determines whether your learner data, your IP, and your commercial information remain yours after they enter the supplier's tool.

A vague answer is a fail. The supplier needs to give you a specific, documented position on input retention, training use, sharing, and any commercial exploitation. "It is secure" is not an answer.

Worked example: Samsung, 2023

Within twenty days of ChatGPT being made available to staff, three Samsung engineers leaked confidential source code into it. They were not malicious. They were trying to get their work done. The tool processed what they typed exactly as it was designed to. The data was simply no longer Samsung's once it was typed in.

Apply the same logic to your learner records, your assessment data, your safeguarding notes, your commercial documents. Once it is in the supplier's system, the supplier's policy on inputs determines what happens next. Not your wish.

What a good answer looks like

- An explicit, contractual statement on input retention: none, a specified retention period, or indefinite.
- An explicit, contractual statement on training use: never, opt-in only, or default-in with the ability to opt out.
- An explicit, contractual statement on data sharing with sub-processors or third parties, including whether anonymised or aggregated data is shared.
- A clear position on whether your inputs can be sold or commercially exploited, and a contractual prohibition where applicable.
- Plan-tier or contract clause backing the above (free-tier consumer products often have very different terms from enterprise tiers).

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Question 4. Where does the data actually live?

Where does the data actually live? UK, EU, US, somewhere else? Whose servers?

Why it matters

International transfers carry additional UK GDPR obligations. You need to know where the data is physically processed and stored, who hosts it, and what transfer mechanism applies if it leaves the UK or the EU.

The DfE explicitly requires that this is disclosed to users where learner data is processed by AI.

What a good answer looks like

- Named hosting countries and providers (for example, AWS eu-west-2, Azure UK South, Google Cloud europe-west2).
- A complete list of sub-processors and the regions in which they operate.
- A documented transfer mechanism for any international transfer (UK adequacy decision, Standard Contractual Clauses, International Data Transfer Agreement, supplementary measures).
- A change-control commitment so this cannot change quietly without notice.

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Question 5. How long do you keep it, and how do I get it deleted?

How long do you keep it, and how do I get it deleted?

Why it matters

Storage limitation is a UK GDPR principle. Data must not be kept for longer than necessary for the purpose.

The right to erasure belongs to your learner, not to your supplier. When a learner exercises that right, your supplier's retention policy and deletion process need to support you in honouring it.

What a good answer looks like

- A documented retention period for each category of data (inputs, outputs, logs, metadata, account information).
- A documented deletion process (in-product self-serve, support ticket, formal request) and the route to use.
- An SLA or commitment on time to honour an erasure request.
- An option to obtain a deletion certificate where required for audit or learner reassurance.
- Treatment of backups (often retained beyond the active deletion window) and any onward implication.

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Question 6. Can you explain a specific output to me?

Can you explain a specific output to me, if I have to justify it to a regulator?

Why it matters

Where AI output materially affects a learner (a recommendation, a flag, a score, an automated grading element, a route into safeguarding) you need to be able to explain how that output was reached. Not in theory. For the specific case in front of you.

This is a UK GDPR Article 22 question (rights related to automated decision-making) and a fairness question rolled into one. Ofsted and the DfE will increasingly expect the same.

What a good answer looks like

- A way to retrieve, for any specific output, the inputs used, the model version, and the parameters that produced it.
- A plain-English explanation route, in-product or via supplier support, that you can show to a learner or a regulator.
- A human-review route for contested outputs, with the ability to override the AI-driven outcome.
- An audit log that retains this evidence for as long as a learner could reasonably challenge the outcome.

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Question 7. What protects my IP and my learners' IP?

What protects my IP, and my learners' IP, from becoming your training data?

Why it matters

Your curriculum, your assessment materials, your internal training resources are all IP. They should not be ingested as training data without your written permission.

Your learners' submitted work is theirs (or yours under your terms with them); the same protection applies.

The DfE makes this explicit. Suppliers that have not addressed this question are not ready for procurement.

What a good answer looks like

- An explicit contractual provision excluding customer content (inputs, outputs, uploaded files) from training use.
- A technical implementation that backs the contract, such as enterprise tenant isolation, zero-retention API access, or segregated training environments.
- A change-control commitment, so the position cannot move without notice or consent.
- A position on indemnity if customer or learner IP appears in a downstream model output.

Your answer

[*Record the supplier's answer here, with links to supporting documentation where possible.*]

aptem.
Intelligent technology®

Bonus considerations

These items are not strictly part of the seven questions. They tie the assessment into your broader AI governance and make it useful long after sign-off.

AI risk tier

Confirm the provisional tier from the Overview section, now that you have the answers. Tiers align with the AI use and governance charter (Low / High / Prohibited).

- Low risk: little or no effect on individuals or consequential decisions.
- High risk: influences outcomes for individuals or processes personal data with potential impact on rights or eligibility. High-risk cases require additional sign-off.
- Prohibited: any use that contravenes UK law, regulation, or your AI charter. Stop here.

DPIA cross-reference

If personal data (including learner data) is processed, a DPIA is required and runs alongside this assessment. Record the link below.

DPIA reference: [Link or "Not required" with rationale]

Incident reporting

Where should suspected AI-related incidents (bias, data exposure, unexpected behaviour, supplier failure) be reported? Make sure staff using the tool know the route.

Incident reporting route: [Email or internal process link]

Review trigger

This assessment is not a one-off. Set a review trigger so it does not go stale.

- Annual review by default.
- Earlier review on any material supplier change (new sub-processor, new region, new model, change of underlying foundation model).
- Earlier review on any regulatory change (UK GDPR, DfE guidance, EU AI Act milestones).

aptem.
Intelligent technology®

Move to the full AISIA

This template is a floor. When the use case becomes consequential, or when your governance matures, move to the full AISIA template. The seven questions sit inside the full assessment; nothing you record here is wasted.

Follow us on LinkedIn and YouTube

www.aptem.co.uk

29 May 2026

aptem.
Intelligent technology®

Outcome and sign-off

Outcome

Tick the outcome that applies. Record any conditions or follow-up actions below.

Outcome	Notes
[ ] Proceed	[Conditions, if any]
[ ] Proceed with conditions	[Conditions: DPIA completion, contract amendments, technical controls, training, etc.]
[ ] Do not proceed	[Reason and what would change the decision]

Sign-off

For low-risk use cases, one named sign-off is sufficient. For high-risk cases, sign-off from [senior leadership / board] is required in addition.

Role	Name	Date	Signature
[Role, e.g. Service or department lead]	[Name]	[DD Month YYYY]	
[Role, if high-risk: senior leadership / board]	[Name]	[DD Month YYYY]	