

Questions to ask AI suppliers

Prepare what to ask before choosing an artificial intelligence solution.

Use this guide to support your research on AI products or services before committing.

You don't need to ask every question. Focus on the ones most relevant to the task or process you plan to use AI for, especially where the potential risk or impact is higher.

Each section includes:

- questions to ask suppliers
- what to check for in a good answer
- what to capture in writing.

Managing risks and potential harms

Questions to ask

- What legal, operational or reputational risks have you identified for this product or service?
- How do you monitor and respond to new or emerging risks over time?
- What safeguards or back-up processes exist for high impact failures?
- How are incidents reported, investigated and fixed, and who is responsible?

What to check for

- The supplier can explain the main risks in plain language.
- They have a clear process for incidents, escalation and remediation.
- They can show how risk controls change based on business impact.

What to capture

- Incident reporting timeframes and escalation path.
- Shared responsibilities between you and the supplier.
- Any back-up or safemode processes.

Transparency and explainability

Questions to ask

- How can we understand why the system produced a particular output or recommendation?
- What documentation do you provide on capabilities and limitations?
- How do you help users avoid overreliance or misuse?
- If someone is affected by an AI-supported outcome, what review or complaint process is available?

What to check for

- Clear documentation on limitations and known failure points.
- A practical explanation (not 'black box' or "trust us" responses).
- Clear ways for people to question and override outputs.

What to capture

- What information must be shared with staff or customers (where relevant).
- What documentation is provided and how often it is updated.

Accountability, roles and integration

Questions to ask

- Who is accountable once the system is live (your organisation, the supplier or subcontractors)?
- What roles do you provide during setup, integration and training?
- Who built the AI model or components (proprietary, open source or third party)?
- How will this change our workflows, staff tasks and decision points?

What to check for

- A clear split of responsibilities.
- Transparency about subcontractors and their roles.
- Integration and support match your capability and systems.

What to capture

- Named accountability contacts
- Supplier responsibilities across the lifecycle (setup, run, fix, and exit)

Data training, use and security

Questions to ask

- Where was this AI trained and what data was used?
- Is the training data relevant to my industry and task?
- What data does the AI system collect, store or process?
- Will any of our data be used to train AI models now or in the future?
- Is any data shared with third parties, such as cloud providers or subcontractors?
- Where is data stored or processed, including offshore locations?
- Can AI features be limited, configured or turned off to restrict data use?
- Who owns the data we put into the system and the outputs it generates?
- How is personal or sensitive data protected throughout its lifecycle?
- Can I review what data the AI is using to make decisions?
- What happens to my data if I stop using this solution?
- Are there any data licensing or ownership agreements I need to know about?

What to check for

- Clear explanations of data flows and data use.
- The ability to control, limit or opt out of certain data uses.
- Security measures that are appropriate for the sensitivity of your data.

What to capture

- Data ownership, retention and deletion terms.
- Clear statements about whether data is used for training.
- Security responsibilities and breach notification processes.

Testing, monitoring and performance

Questions to ask

- How has the system been tested for this type of use?
- How do you monitor accuracy, reliability and performance once the system is live?
- How do you detect when performance degrades or behaviour changes over time?
- How are updates managed and communicated?
- Can we test or review changes before they are applied?

What to check for

- Evidence of testing, not just assurances.
- Ongoing monitoring that matches the importance of the use.
- Clear processes for updates and change management.

What to capture

- Expectations for testing before and after deployment.
- How performance issues are identified and handled.
- How and when updates will occur.

Human oversight and control

Questions to ask

- Where do people review or approve AI outputs?
- Can the system be paused, overridden or switched off?
- What training or guidance is provided for staff?
- What happens if the AI fails or produces unexpected results?

What to check for

- Clear points where people have control for higher-impact uses.
- Practical override or 'stop' controls.
- Training that matches your team's roles and responsibilities.

What to capture

- When review by people is required.
- Who can override or shut down the system.
- Backup processes if the AI is unavailable.

Fairness, inclusion and broader impacts

Questions to ask

- How do you assess and reduce bias in the system?
- How is fairness monitored over time?
- Who could be negatively affected if the system makes errors?
- What feedback or complaints processes are available?

What to check for

- Evidence that fairness and impacts have been considered.
- Clear processes for monitoring and responding to issues.
- Willingness to discuss limitations and tradeoffs.

What to capture

- Any fairness or impact assessments.
- How issues are escalated and addressed.
- Responsibilities for ongoing monitoring.

Exit, decommissioning and continuity

Questions to ask

- What happens if we decide to stop using the system?
- Can we export our data in a usable format?
- How is data deleted after exit?
- What happens if the product is retired or stops working unexpectedly?
- Can we continue operating without the AI if needed?

What to check for

- Clear exit and data return processes.
- Reasonable transition or shutdown timeframes.
- Plans that support business continuity for critical systems.

What to capture

- Exit and termination terms.
- Data export and deletion commitments.
- Decommissioning responsibilities.

Environmental sustainability

These questions are optional but are becoming more relevant for businesses.

Questions to ask

- Is the system sized appropriately for our needs?
- Do you have information on energy use or efficiency?
- How often is the model retrained and on what infrastructure?
- Is sustainability considered as part of the system design?

What to check for

- A fit-for-purpose approach rather than unnecessary complexity.
- Willingness to discuss environmental impacts at a practical level.

What to capture

- Any sustainability metrics or commitments.
- Alignment with your organisation's environmental, social and governance priorities.