

AI Vendor Risk Assessment Questionnaire

30 questions across huby's six evaluation categories, for assessing an AI vendor before you buy. Ask for evidence rather than yes/no answers. Copy or adapt these freely.

#	CATEGORY	QUESTION	EVIDENCE TO REQUEST
1	Quality	What accuracy or performance benchmarks do you publish, and were any run independently?	Benchmark reports with method and test date
2		How do you test for errors and regressions before releasing a model update?	QA process, release notes
3		How will we be notified of model or behavior changes?	Changelog, notification policy
4		What are the known failure modes and limitations for our use case?	Model card or limitations doc
5		What defect and resolution-time data can you share?	Support SLAs, incident metrics
6	Privacy	Are our inputs, outputs or files used to train or improve any model? Can we opt out by contract?	DPA, plan terms
7		How long are inputs and outputs retained, and can we delete them on request?	Retention schedule, deletion procedure
8		Which sub-processors, including model providers, process our data, and where?	Sub-processor list, data residency
9		Which certifications and regulations apply (SOC 2, GDPR, CCPA/CPRA, HIPAA)?	Audit reports, signed DPA, BAA if needed
10		How do users exercise access, export and deletion rights?	Privacy policy, admin docs
11	Security	When was your last independent penetration test, and what was remediated?	Pen test summary, remediation log
12		Do you support SSO, MFA and role-based access?	Admin documentation
13		How is data encrypted in transit and at rest, and who holds the keys?	Security whitepaper
14		How do you defend against prompt injection, cross-customer data leakage and model abuse?	AI security controls, red-team results
15		Which security incidents or CVEs affected the product in the past 24 months?	Incident history, CVE records
16	Use Cases & Pricing	Which of our use cases are supported today, and which are roadmap?	Product docs, dated roadmap
17		Which integrations and APIs are included on the tier we would buy?	Tier comparison
18		How is usage metered, and what happens at limits?	Pricing page, overage terms
19		Which data-protection features are gated to higher tiers?	Tier feature matrix
20	Sustainability & Reliability	What does exit look like: export format and transition support?	Exit and termination terms
21		What uptime do you commit to, and what was actual uptime over the past 12 months?	SLA, status page history
22		What happens to our service if your model provider changes terms or fails?	Continuity plan

#	CATEGORY	QUESTION	EVIDENCE TO REQUEST
23	Sustainability & Reliability	What evidence of financial stability can you share?	Funding status, references
24		What is your incident response and customer notification timeline?	Incident response policy
25		Which customers of similar size use the product today?	Named references, case studies
26	Impact, Ethics & Safety	Do you publish an ethics or acceptable-use policy, and how is it enforced?	Policy, enforcement examples
27		How do you test for and reduce bias in outputs that affect people?	Bias testing results
28		What human oversight do you recommend for consequential outputs?	Usage guidance
29		How do you handle copyright in generated content, and do you offer IP indemnity?	Terms, indemnity clause
30		How have you responded to past safety or ethics incidents?	Incident reports, public statements