Kailash Consult Free resource · kailashconsult.com/resources

AI Vendor Review Pack

The 25 AI questions enterprise security teams now ask, an example answer for each, and where each question maps in ISO 42001, AIUC-1, the EU AI Act and the NIST AI RMF.

How to use

- Use this when a buyer's security or legal team adds an AI section to their review.
- Example answers are starting points. Replace every [bracket] and only state what your product does today.
- The mapping columns show which standard each question comes from. They are typical references, not legal advice; EU AI Act duties depend on how your system is classified.

Example answers and mappings are templates, not facts about your company. Edit every bracket and validate each statement before sharing it with a customer or auditor.

What AI you use and where data goes

ID	Question buyers ask	Example answer (edit before use)	ISO 42001	AIUC-1 area	EU AI Act	NIST AI RMF
C01	Which AI features are in the product and which models power them?	Features [X, Y] use [provider, model]. Feature [Z] uses our own model [type]. The full list is kept in our AI system inventory and updated with each release.	A.6	Accountability	Art. 11 (if high-risk)	Map
C02	Which third-party model providers process our data, and in which regions?	[Provider] processes prompts in [region] under our enterprise agreement. AI providers are listed as subprocessors at [link].	A.10	Data & Privacy	Art. 25	Govern
C03	Is our data used to train or fine-tune any model?	No. Customer data is not used to train or fine-tune any model, ours or our providers'. This is written into our provider contracts and our customer terms.	A.7	Data & Privacy	Art. 10 (if high-risk)	Map

Free to use inside your company. Not for resale. kailashconsult.com/resources

ID	Question buyers ask	Example answer (edit before use)	ISO 42001	AIUC-1 area	EU AI Act	NIST AI RMF
C04	Do model providers keep prompts or outputs? For how long?	Our provider settings are [zero data retention / retention limited to [X] days for abuse monitoring]. We store prompts and outputs for [X days] for [support/quality] and then delete them.	A.10	Data & Privacy	-	Manage
C05	Can we turn AI features off or limit them for our workspace?	Yes. Workspace admins can [disable AI features / restrict them by role] in settings. [Or: available on request for enterprise plans.]	A.9	Accountability	-	Manage
C06	How do you stop one customer's data from appearing in another customer's AI output?	Each request only includes data from the requesting customer's workspace. Retrieval is filtered by tenant ID at query time, and we test for cross-tenant leakage before release.	A.7	Security	Art. 15	Manage

Security of AI features

ID	Question buyers ask	Example answer (edit before use)	ISO 42001	AIUC-1 area	EU AI Act	NIST AI RMF
C07	How do you protect against prompt injection and data leaking through AI features?	We separate system instructions from user content, filter inputs and outputs, limit what data the model can reach, and test against known injection patterns before release.	A.6	Security	Art. 15	Measur
C08	What can your AI do on its own, and with what permissions?	[Our AI only suggests; a user confirms every action.] [Or: the agent can [actions] using the user's own permissions, and [actions] always need user confirmation.]	A.9	Safety	Art. 14 (if high-risk)	Manage
C09	How do you control changes to models and prompts?	Model and prompt changes follow our change process: review, evaluation against a test set, staged rollout, and the ability to roll back.	A.6	Reliability	-	Manage
C10	Do you remove personal data before sending it to a model?	[We redact [data types] before model calls.] [Or: personal data is sent only when the feature needs it, to providers under a DPA, with no training and limited retention.]	A.7	Data & Privacy	-	Manage

Quality, safety and human oversight

ID	Question buyers ask	Example answer (edit before use)	ISO 42001	AIUC-1 area	EU AI Act	NIST AI RMF
C11	Where does a person review AI output before it affects a customer or a decision?	AI output is presented as a draft or suggestion. [Decisions about [X] always require human approval.] Users can edit or reject every output.	A.9	Accountability	Art. 14 (if high-risk)	Manage
C12	How do you test AI quality before release and monitor it after?	Each AI feature has an evaluation set of [N] cases scored before release. In production we track [user feedback, error reports, sampled reviews] and investigate drops.	A.6	Reliability	Art. 15	Measur
C13	How do you handle wrong or made-up AI output?	Outputs [cite their sources / are limited to the customer's own data]. Users can flag bad output in one click, and flagged cases feed our evaluation set.	A.6	Reliability	Art. 15	Manage
C14	Have you assessed bias or unfair outcomes?	[For features that affect people, we test outputs across [groups/scenarios] and review results before release.] [Or: our AI features do not make decisions about individuals.]	A.5	Society	Art. 10 (if high-risk)	Measur
C15	Do you log AI inputs, outputs and actions? For how long?	We log [requests, outputs, actions, user and time] for [X days], access-controlled and used only for security, support and quality.	A.6	Accountability	Art. 12 (if high-risk)	Measur
C16	How would you detect and tell us about an AI incident or misuse?	AI incidents follow our incident response plan: detection through monitoring and user reports, severity rating, and customer notification within [X hours] when their data or decisions are affected.	A.8	Safety	Art. 73 (if high-risk)	Manage

Governance and transparency

ID	Question buyers ask	Example answer (edit before use)	ISO 42001	AIUC-1 area	EU AI Act	NIST AI RMF
C17	Do you have an AI policy approved by leadership?	Yes. Our AI policy, approved by [CEO/CTO] on [date], sets rules for building, buying and using AI, including data use, human oversight and acceptable uses.	A.2	Accountability	-	Govern
C18	Who owns AI risk in your company?	[Name, title] owns AI governance. Product and engineering leads own risks for their AI features, reviewed [quarterly].	A.3	Accountability	-	Govern

ID	Question buyers ask	Example answer (edit before use)	ISO 42001	AIUC-1 area	EU AI Act	NIST AI RMF
C19	Do employees get training on responsible AI use?	Yes. Staff who build or use AI complete training on our AI policy, data handling and limits of AI output at onboarding and annually.	Clause 7.2	Accountability	Art. 4	Govern
C20	Have you done an AI risk or impact assessment for your AI features?	Yes. Each AI feature has a risk and impact assessment covering intended use, misuse, affected people and controls, reviewed when the feature changes.	Clause 6.1.4, A.5	Accountability	Art. 9 (if high-risk)	Map
C21	Do you tell end users when they are dealing with AI or AI-generated content?	Yes. AI features are labeled in the interface, and AI-generated content is marked [as AI-generated / with an indicator].	A.8	Society	Art. 50	Govern
C22	Can you give us documentation of how the AI works and its limits?	Yes. We provide feature documentation covering purpose, data used, known limits and how to report problems at [link].	A.8	Accountability	Art. 13 (if high-risk)	Govern
C23	Where did your training or fine-tuning data come from, and do you have rights to it?	[We do not train our own models; we use [provider] base models.] [Or: training data comes from [sources] that we are licensed or permitted to use; we keep records of sources.]	A.7	Data & Privacy	Art. 10, Art. 53	Map
C24	Have you classified your AI features under the EU AI Act?	Yes. We assessed each AI feature against the Act's categories (prohibited, high-risk, transparency, minimal). Our current features are [classification], with the obligations that follow.	A.5	Accountability	Art. 5, 6, 50	Map
C25	Do you follow, or plan to certify against, ISO 42001 or AIUC-1?	[We are certified to [standard] since [date].] [Or: we align our AI controls to ISO 42001 and plan [certification/assessment] by [month, year].]	Whole standard	Whole standard	-	Govern

Free resource from Kailash Consult. Free to use inside your company. Not for resale or redistribution. kailashconsult.com/resources