ITKNOWLEDGEBASES.COM

Security Intelligence & Advisory | itknowledgebases.com

AI VENDOR SECURITY QUESTIONNAIRE

FOR BUSINESS

Before your organization deploys ChatGPT, Claude, Copilot, Gemini, or any third-party AI tool — ask these questions first.

AUTHORED BY A SECURITY PROFESSIONAL · 2026 EDITION · ITKNOWLEDGEBASES.COM

David Wolfcale | Information Security Professional

. HOW TO USE THIS DOCUMENT

This questionnaire was built from real-world security experience — not from compliance checklists written by people who have never actually had to live with a bad vendor decision. Every question here exists because the gap it exposes is one that can cause real security incidents.

The document is organized into seven sections covering the full lifecycle of AI vendor risk — from data handling and access control through compliance certifications, model behavior, vendor security posture, contract risk, and AI-specific governance.

Who Should Complete This

- IT Managers evaluating a new AI tool for employee use
- Security teams vetting an AI-integrated SaaS product before onboarding
- Compliance officers reviewing third-party AI vendor contracts
- Procurement teams building AI vendor due diligence into purchasing workflows
- vCISOs standing up an AI governance program for client organizations

How to Score Responses

Each question includes a risk level — CRITICAL, HIGH, or MEDIUM. Use the scoring rubric in Section 8 to generate a vendor risk score after completing the assessment. A vendor that cannot answer CRITICAL questions should not be approved for use with sensitive data until those gaps are resolved.

> ⚠ **THE SUPPLY CHAIN BLIND SPOT**
>
> The most dangerous AI risk is not the tool your team chose — it is the AI embedded inside the tool they chose. Many SaaS products now include AI features powered by third-party models (OpenAI, Anthropic, Google, Cohere) without clearly disclosing this. A vendor saying their product is secure does not mean the AI dependency inside their product has been assessed. This questionnaire forces that visibility.

Instructions for Vendor Completion

- Send this document to your vendor contact with a completion deadline of 10 business days
- Request supporting documentation for any claim of certification or compliance
- A response of N/A requires a brief explanation of why the question does not apply
- Incomplete responses should be treated as unanswered for scoring purposes
- Request a follow-up call to clarify any answers that seem vague or evasive

1. DATA HANDLING & PRIVACY

The questions in this section expose the most common and dangerous gap in AI vendor due diligence — what actually happens to your data once it leaves your environment. From a security perspective, this is the section vendors are least prepared to answer honestly, because many of them have not done the internal audit to know.

REAL WORLD CONTEXT

A healthcare organization was questioning a customer support platform with AI-powered ticket summarization. When asked whether the AI component was HIPAA compliant and whether patient data could potentially get sent for model training, the vendor could not answer — because the AI was a third-party API integration they had never specifically audited. This is not an edge case. It is the norm.

Q1 Does your product use any third-party AI models, APIs, or machine learning components? If yes, list each one.

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Many vendors embed AI from OpenAI, Anthropic, Google, Cohere, or open-source models without disclosing this. Each of these creates a separate data flow that must be assessed independently. You cannot evaluate AI risk without knowing what AI is actually present.	A complete list of all AI components, including vendor name, model used, version, and data flow. If they use a custom in house, pay close attention to how it is trained. Vague answers like 'we use industry-standard AI' are not acceptable.	CRITICAL

Q2 For each AI component or API your product uses, where is customer data sent and processed? Specify country, cloud provider, and data center region.

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Data residency directly impacts your legal obligations under GDPR, HIPAA, state privacy laws, and industry regulations. Data processed in a foreign jurisdiction may fall under laws you did not account for in your compliance program.	Specific country and cloud region for each AI component. Bonus: documentation of data residency guarantees or contractual commitments on where data is stored.	CRITICAL

Q3 Do you maintain a software bill of materials (SBOM) or AI component inventory that documents all AI dependencies and their data flows?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Organizations that do not inventory their AI components cannot answer basic security questions about them. An SBOM for AI is the equivalent of knowing what packages are in your codebase — without it, supply chain risk is invisible.	A documented SBOM or AI inventory that is reviewed at least annually and updated when new components are added. Willingness to share this on request.	HIGH

Q4 Is customer data used to train, fine-tune, or improve any AI models — either your own or your third-party AI vendors? Is this opt-in or opt-out?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
The Samsung ChatGPT incident demonstrated that data entered into AI systems can become training data. This question forces vendors to either confirm data isolation or admit that your sensitive data may be contributing to a shared model.	A clear written statement that customer data is not used for training, with contractual language to back it up. If training does occur, it should be opt-in only with explicit customer consent.	CRITICAL

Q5 What is your data retention policy for inputs, outputs, and prompts processed by AI components? How long is data stored and where?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
AI systems often retain conversation history, prompt data, and outputs for extended periods. Without a clear retention policy, sensitive data may persist indefinitely in systems outside your control.	A specific retention period (e.g., 30 days), a process for customer-requested deletion, and confirmation that retention policies apply equally to third-party AI components.	HIGH

Q6 Do you offer a Data Processing Agreement (DPA) that specifically covers AI processing of customer data?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
A generic DPA may not address AI-specific processing. Without explicit DPA coverage of AI components, you have no contractual protection if data is mishandled by an AI system.	A DPA that explicitly names all AI components, defines their data processing scope, and commits to GDPR Article 28 requirements or equivalent.	HIGH

Q7 For healthcare-related deployments: Will you sign a Business Associate Agreement (BAA) covering all AI components in your product?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
HIPAA requires a BAA with any vendor that handles Protected Health Information. Most AI platforms do not sign BAAs by default, and many cannot because their underlying AI APIs are not HIPAA compliant. Using AI tools with patient data without a signed BAA is a direct HIPAA violation regardless of intent.	A signed BAA that explicitly covers AI processing components. Vendors who cannot provide this should not be used with any PHI under any circumstances.	CRITICAL

Q8 Are customers notified — within your product UI or documentation — when their data is processed by AI? Is this disclosure clear and accessible?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Transparency about AI processing is increasingly required by regulation (EU AI Act, California AB 2013) and is a baseline expectation of ethical AI deployment. Vendors who hide AI processing create downstream compliance risk for your organization.	Clear in-product disclosure when AI is processing data, with accessible documentation of what AI is used and how. Not buried in terms of service.	MEDIUM

2. ACCESS CONTROL & AUTHENTICATION

AI systems that integrate into enterprise environments create new access control challenges. This section evaluates whether the vendor has implemented controls that meet enterprise security standards — and whether those controls extend to the AI components in their product, not just the application layer.

Q9 Does your platform support Single Sign-On (SSO) via SAML 2.0 or OIDC? Is SSO available on standard business tiers or only enterprise plans?

Why This Matters
SSO enforcement is a fundamental enterprise security control. AI platforms that gate SSO behind expensive enterprise tiers force organizations to choose between security and cost. From a SOC perspective, a platform without SSO creates orphaned accounts and bypasses your centralized identity controls.

What a Good Answer Looks Like
SSO available on business tiers or above, supporting SAML 2.0 or OIDC. Ability for admins to enforce SSO so users cannot authenticate with username/password as a fallback.

Risk If Unanswered
HIGH

Q10 Is multi-factor authentication (MFA) enforced or optional? Can administrators enforce MFA organization-wide?

Why This Matters
Optional MFA means some users will not use it. In a SOC environment, a single account without MFA is a single point of compromise for everything that account can access — including whatever AI system it connects to.

What a Good Answer Looks Like
MFA enforced at the organizational level with no admin bypass. Support for hardware tokens or authenticator apps, not just SMS.

Risk If Unanswered
HIGH

Q11 Does your platform implement role-based access control (RBAC)? Can administrators define granular permissions for what different users can do with AI features?

Why This Matters
AI features often grant broad capabilities — generating content, accessing data, making API calls. Without RBAC, every user in your organization has the same level of access to these capabilities regardless of their role or need.

What a Good Answer Looks Like
Granular RBAC that allows admins to restrict AI feature access by user, group, or role. Ability to prevent specific users from accessing certain AI capabilities entirely.

Risk If Unanswered
HIGH

Q12 Is there a full audit log of all AI interactions — including who accessed the AI, what prompts were submitted, and what responses were generated? How long are logs retained and are they exportable?

Why This Matters
Without audit logs, you have no ability to investigate a data incident involving AI. If an employee submits sensitive data to an AI tool and that data is later found in a breach, logs are the only way to establish what happened, when, and who was involved.

What a Good Answer Looks Like
Immutable audit logs retained for a minimum of 90 days (ideally 1 year), covering all AI interactions by user, timestamp, and content summary. Logs exportable to SIEM or downloadable in standard formats.

Risk If Unanswered
HIGH

Q 13 **If your product exposes an API, how are API keys scoped and managed? Can keys be restricted by permission level, IP address, or usage rate?**

Why This Matters

Poorly scoped API keys are one of the most common AI security failures. A single leaked API key with broad permissions can expose your entire data environment. Rate limiting prevents abuse and reduces the blast radius of a compromised key.

What a Good Answer Looks Like

API keys with scoped permissions, per-key usage limits, IP allowlisting capability, and the ability to revoke individual keys without affecting others. Key rotation supported.

Risk If Unanswered

MEDIUM

3. COMPLIANCE CERTIFICATIONS

Certifications are not a substitute for understanding — a vendor can be SOC 2 certified and still have AI components that fall completely outside the scope of that certification. Always ask whether the certification specifically covers the AI features you are evaluating, and always request the actual report rather than accepting a badge on a website.

Q 14 Do you have a SOC 2 Type II certification? Does the scope of the report explicitly include your AI processing components and any third-party AI APIs?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
SOC 2 Type II is the baseline enterprise security certification. However, many vendors have SOC 2 certifications that were written before they added AI features, meaning the AI components are outside the audit scope. A certification that doesn't cover the AI is not useful for evaluating AI risk.	A current SOC 2 Type II report (within the last 12 months) with explicit scope coverage of AI features and data flows. Willingness to share under NDA.	HIGH

Q 15 Are you ISO 27001 certified? What is the scope of the certification and when does it expire?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
ISO 27001 demonstrates a systematic approach to information security management. Like SOC 2, the scope matters — a certification that covers headquarters but not the cloud infrastructure running your AI is limited in value.	Current ISO 27001 certification with scope that includes AI processing infrastructure. Certificate available on request.	MEDIUM

Q 16 Are you FedRAMP authorized? At what impact level (Low, Moderate, High)? Does FedRAMP authorization cover your AI components?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
For any organization working with federal agencies, state agencies with federal funding, or defense contractors, FedRAMP is often a hard requirement. AI tools that are not FedRAMP authorized may not be legally permissible in these environments.	Active FedRAMP authorization at the appropriate impact level with AI features within scope. Authorization status verifiable through the FedRAMP marketplace.	HIGH

Q 17 For organizations with EU customers or operations: Do you offer Standard Contractual Clauses (SCCs) for data transfers? Are these specific to AI processing?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
GDPR requires specific legal mechanisms for transferring EU personal data outside the EU. SCCs are the most common mechanism, but they must cover the actual data flows — including AI processing — to be effective.	Current SCCs available that explicitly cover AI data processing flows, not just core product data. Legal team available to review on request.	HIGH

Q 18 **Have you conducted a GDPR Data Protection Impact Assessment (DPIA) for your AI features? Is documentation available?**

Why This Matters

The EU AI Act and GDPR both require DPIAs for high-risk processing. An AI tool that processes personal data at scale almost certainly requires one. A vendor who has not done a DPIA has not seriously assessed the privacy risk of their AI features.

What a Good Answer Looks Like

Completed DPIA documentation for AI features that process personal data, reviewed by a qualified Data Protection Officer.

Risk If Unanswered
MEDIUM

Q 19 **Under the EU AI Act, how do you classify your AI systems by risk level? What compliance steps have you taken for any high-risk classifications?**

Why This Matters

The EU AI Act (partially in force as of 2024) classifies AI systems by risk level and imposes significant obligations on high-risk systems. HR screening, credit scoring, medical AI, and biometric systems are high-risk by default. Vendors who cannot answer this question have not assessed their EU AI Act obligations.

What a Good Answer Looks Like

A clear risk classification for each AI feature with documented compliance steps for any high-risk systems. Transparency about systems that may require conformity assessments.

Risk If Unanswered
MEDIUM

4. MODEL & OUTPUT RISK

AI outputs carry risks that traditional software does not — hallucination, bias, inconsistency, and the potential for model behavior to be influenced by adversarial inputs. This section evaluates whether the vendor has thought seriously about what their AI actually does and what happens when it does it wrong.

Q 20 If your product allows fine-tuning or custom model training on customer data, is that data strictly isolated per customer? Can one customer's data influence another customer's model outputs?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Shared model training is one of the most serious data privacy risks in AI. If customer A's data can influence the outputs generated for customer B, you have a direct data leakage pathway. This is not theoretical — it has happened in production AI systems.	Strict per-customer data isolation with technical controls (not just policy) preventing cross-customer data influence. Independent security validation of isolation controls.	CRITICAL

Q 21 What content filtering and output guardrails does your AI implement? How are these tested and how frequently are they updated?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
AI systems without content guardrails can generate harmful, biased, legally problematic, or simply incorrect outputs. In a business context, an AI tool that generates confidently wrong information and presents it as fact creates liability.	Documented content filtering policies with specific categories of restricted output. Regular red team testing of guardrails. A process for customers to report guardrail failures.	HIGH

Q 22 Does your AI system log prompts and responses at the account level? Can administrators review AI interaction history for compliance and incident investigation purposes?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
From a SOC perspective, an AI tool that does not log interactions is a black box. When something goes wrong — a data leak, a compliance violation, an inappropriate output — you need to be able to reconstruct exactly what happened.	Per-account logging of all AI interactions, accessible to administrators, with a minimum 90-day retention period. Logs should include prompt text, response text, user identity, and timestamp.	HIGH

Q 23 How does your AI system handle hallucinations — cases where the model generates false or fabricated information presented as fact? Is there any mechanism to flag uncertain outputs?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
AI hallucination is not a bug that will be fixed — it is a fundamental characteristic of large language models. A vendor who claims their AI does not hallucinate does not understand their own product. The right answer is a vendor who has built transparency mechanisms around uncertainty.	Explicit acknowledgment that hallucinations occur, with mechanisms to flag low-confidence outputs, cite sources where possible, and recommend human review for high-stakes decisions.	MEDIUM

Q 24 **Has your AI system been tested for prompt injection vulnerabilities? What controls prevent malicious inputs from manipulating AI behavior or extracting sensitive data?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Prompt injection is the AI equivalent of SQL injection — an attacker crafts an input that causes the AI to ignore its instructions and behave in ways the vendor did not intend. In an enterprise context, a successful prompt injection attack can extract data, bypass guardrails, or cause the AI to take unauthorized actions.	Documented prompt injection testing as part of security assessments, with specific controls (input validation, output filtering, instruction hierarchy enforcement) in place. Ongoing monitoring for prompt injection attempts.	HIGH

Q 25 **Who owns the intellectual property rights to outputs generated by your AI using our data? Does your platform make any claims to content generated during our use of the service?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
AI output ownership is legally unsettled in many jurisdictions, and vendor terms of service vary significantly. Some vendors claim rights to outputs, which creates complications for organizations using AI-generated content commercially.	Clear contractual statement that the customer owns all outputs generated using their data. No vendor claims to AI-generated content. Legal review of ToS recommended regardless.	MEDIUM

5. VENDOR SECURITY POSTURE

This section evaluates the vendor's fundamental security program — independent of their AI features. A vendor with a weak overall security posture cannot be trusted to secure their AI systems, regardless of what their marketing materials claim. Ask for evidence, not assertions.

Q 26 **Has your organization undergone a third-party penetration test within the last 12 months? Does the scope include your AI features and integrations? Can you share an executive summary under NDA?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Self-assessed security is not security. Third-party penetration testing by qualified firms is the baseline for taking vendor security claims seriously. The willingness to share results is itself a trust signal.	Annual third-party penetration test with AI features in scope. Executive summary available under NDA. Remediation of critical and high findings documented.	HIGH

Q 27 **Do you have a published vulnerability disclosure program (VDP) or bug bounty program? What is your average time to remediate critical vulnerabilities?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Organizations with no VDP are telling security researchers there is nowhere to report vulnerabilities responsibly. This does not mean fewer vulnerabilities exist — it means they go unreported until they are exploited.	Active VDP with a clear process for external researchers to report vulnerabilities. Target remediation SLAs for critical (24-48 hours), high (7 days), and medium (30 days) findings.	MEDIUM

Q 28 **What is your security incident notification SLA? How will you notify customers if there is a breach affecting their data, and within what timeframe?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
GDPR requires breach notification within 72 hours. Many state laws have similar requirements. A vendor with a vague incident notification policy will not help you meet your own regulatory obligations when something goes wrong.	Written incident response policy with customer notification within 72 hours of confirmed breach, specific communication channels defined, and a dedicated security contact.	HIGH

Q 29 **Do you carry cyber liability insurance? What is the coverage limit? Does the policy cover AI-related incidents including data breaches caused by AI components?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Cyber insurance for vendors is a signal of risk maturity — insurers require organizations to meet minimum security standards before issuing coverage. More practically, it determines whether you have any financial recourse if a vendor breach causes you losses.	Active cyber liability insurance with coverage limits appropriate to the vendor's size and data exposure. Policy covers AI-related incidents. Certificate of insurance available on request.	MEDIUM

Q 30

What is your uptime SLA for AI features specifically? How do you handle AI service degradation — for example, if a third-party AI API you depend on experiences an outage?

Why This Matters

AI features often depend on external APIs that are outside the vendor's direct control. An outage at OpenAI or Google affects every product built on their APIs simultaneously. Understanding how vendors handle this reveals their resilience planning.

What a Good Answer Looks Like

Documented uptime SLA for AI features (99.9% or higher). Fallback behavior defined for third-party AI API outages. Historical uptime data available.

Risk If Unanswered

MEDIUM

6. CONTRACT & BUSINESS RISK

The questions in this section are the ones that matter when things go wrong — when a vendor gets acquired, shuts down, or breaches the contract. These are rarely asked during vendor evaluation because they are uncomfortable. They are the most important questions to ask precisely because of that.

Q 31 What happens to our data if your company is acquired, merges with another organization, or ceases operations? Is there a data return or deletion guarantee?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Vendor acquisitions happen. When they do, your data may be transferred to a new entity with different privacy practices, different jurisdictions, and potentially different competitive interests. Without contractual protections, you have no control over this.	Contractual data return or deletion guarantee within 30 days of contract termination, regardless of the reason. Acquisition scenarios explicitly addressed in contract terms.	HIGH

Q 32 Can we export all of our data — including AI-generated content, conversation history, and model outputs — in a standard, portable format at any time?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Vendor lock-in is a business risk that becomes a security risk when you cannot leave. If you cannot export your data, you cannot migrate to a more secure alternative if the vendor's security posture degrades.	Full data export in a standard format (JSON, CSV, etc.) available at any time without requiring vendor assistance. Export includes all AI interaction history and generated content.	MEDIUM

Q 33 Does your platform indemnify customers against intellectual property claims arising from AI-generated outputs? Have you faced any copyright or IP litigation related to your AI features?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
The legal status of AI-generated content is unsettled, and several major lawsuits are working through courts. If your organization uses AI-generated content commercially and the vendor's model was trained on copyrighted material, you may face downstream liability.	Contractual indemnification for IP claims related to AI outputs. Transparent disclosure of any ongoing litigation related to training data or AI outputs.	MEDIUM

Q 34 What are the terms for API rate limiting and fair use? Can your usage be throttled or suspended without notice? What are the SLAs for API availability?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Organizations that integrate AI APIs into critical workflows can be severely disrupted if the vendor unilaterally changes rate limits, introduces pricing changes, or suspends access. This is not theoretical — it has happened to businesses dependent on AI APIs.	Clear documented rate limits, advance notice requirements for limit changes, and SLA commitments for API availability. Emergency contact process for unexpected throttling.	MEDIUM

Q 35 **What is your process for notifying customers of material changes to your AI features, data handling practices, or terms of service? What advance notice is provided?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Vendors frequently update their AI features and terms of service. Without adequate notice, you may find your organization suddenly non-compliant because a vendor changed how they handle data without telling you.	Minimum 30-day advance notice for material changes to data handling or AI features. Email notification directly to account administrators, not just a terms of service banner update.	HIGH

7. AI-SPECIFIC GOVERNANCE

This final section addresses the emerging regulatory and governance landscape around AI. These questions were largely irrelevant two years ago. They are increasingly essential today, and will be mandatory requirements within the next 12–18 months for many regulated industries.

Q 36 Do you have a published responsible AI or AI ethics policy? Who in your organization is accountable for AI governance decisions?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Organizations that have not thought seriously about responsible AI have not thought seriously about the risks of irresponsible AI. A published policy and named accountability signal organizational maturity.	Published responsible AI policy available on the vendor website. Named Chief AI Officer, VP of AI Safety, or equivalent role with clear governance accountability.	MEDIUM

Q 37 Can you provide documentation of your AI model training data sources? Are there known biases in your models and how are they documented and mitigated?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
AI models trained on biased data produce biased outputs. In HR, lending, healthcare, and law enforcement applications, biased AI outputs can create legal liability under discrimination laws. A vendor who cannot discuss their training data cannot discuss their bias risk.	Documentation of training data sources and categories. Published bias evaluation results from third-party audits. Documented mitigation steps for known biases.	HIGH

Q 38 Have your AI systems undergone third-party AI safety audits or red team assessments specifically for AI risks (adversarial inputs, data poisoning, model extraction)?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
Traditional security assessments do not cover AI-specific attack vectors. MITRE ATLAS and OWASP LLM Top 10 define a distinct set of AI threats that require specialized evaluation. Vendors without AI-specific security testing have untested attack surfaces.	Third-party AI safety assessments using MITRE ATLAS or OWASP LLM Top 10 frameworks. Red team results available under NDA. Ongoing monitoring for AI-specific threats.	HIGH

Q 39 How does your AI system handle requests for information about your own security controls or internal systems? Have you tested whether your AI can be manipulated into disclosing sensitive vendor information?

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
AI systems that have access to internal documentation, code, or configurations can sometimes be manipulated into disclosing that information through carefully crafted prompts. This is an AI-specific risk that traditional access controls do not address.	Documented testing of AI behavior when probed for internal information. Technical controls preventing AI from accessing or disclosing sensitive internal data. Incident history of any data disclosure via AI manipulation.	HIGH

Q 40 **What is your process for a customer to request deletion of all data — including data used in AI model training, fine-tuning, or evaluation — from your systems?**

Why This Matters	What a Good Answer Looks Like	Risk If Unanswered
GDPR Article 17 right to erasure applies to AI training data where personal data is involved. Many vendors have deletion processes for application data but no process for removing data from model training. If your customer data was used in training, deletion may require model retraining — a significant operational undertaking most vendors have not planned for.	Documented data deletion process that explicitly covers AI training and fine-tuning data. Deletion confirmation within 30 days. Legal commitment to deletion verifiable.	HIGH

8. SCORING RUBRIC & RISK ASSESSMENT

Use this rubric to calculate a vendor risk score after receiving completed responses. Score each question, total by section, and use the overall score to inform your approval decision.

Individual Question Scoring

Response Quality	Score	Description
Fully answered with documentation	**3 points**	Complete answer with supporting evidence provided
Answered without documentation	**2 points**	Credible answer but no supporting documentation
Partial or vague answer	**1 point**	Incomplete, evasive, or unclear response
Unanswered or N/A without explanation	**0 points**	No response or N/A with no justification provided

Risk Level Weighting

Risk Level	Weight Multiplier	Implication
CRITICAL	**x3**	Any score of 0 on a CRITICAL question = automatic vendor review hold
HIGH	**x2**	Score of 0 on HIGH questions should trigger compensating controls
MEDIUM	**x1**	Document gaps and plan for remediation within 90 days

Overall Vendor Score Interpretation

Score Range	Risk Rating	Recommended Action
85–100%	**LOW RISK**	Approved for use. Annual reassessment recommended.
70–84%	**MEDIUM RISK**	Conditional approval. Document gaps. Compensating controls required.
55–69%	**HIGH RISK**	Do not approve for sensitive data. Escalate to security leadership.
Below 55%	**CRITICAL RISK**	Reject vendor. Do not deploy. Escalate immediately.

> ⚠ **AUTOMATIC HOLD CONDITIONS**
>
> Regardless of overall score, place the vendor on hold if: (1) Any CRITICAL question scores 0. (2) The vendor cannot confirm whether a BAA applies to their AI components. (3) The vendor cannot identify all third-party AI APIs in their product. (4) The vendor's SOC 2 scope does not cover AI features. These are not negotiable.

9. IMPLEMENTATION GUIDE

Completing this questionnaire is the beginning of vendor assessment, not the end. This section provides guidance on building the assessment into your vendor management workflow and what to do with the results.

Before You Send

- Identify your internal stakeholder who owns the vendor relationship — they should be involved in reviewing responses
- Set a response deadline of 10 business days — longer deadlines result in lower quality responses
- Prepare a list of the specific data types the vendor will process — this context helps vendors give accurate answers
- Review the vendor's public privacy policy and terms of service before sending — this gives you a baseline to identify inconsistencies

When You Receive Responses

- Score each question immediately rather than waiting until the full document is returned
- Flag any response that seems inconsistent with the vendor's public documentation
- Request supporting documentation for any claim of certification within 5 business days
- Schedule a 30-minute call to review any vague or incomplete responses before scoring them as partial
- Never accept a SOC 2 badge, logo, or checkbox on a security page as evidence — request the actual report

Escalation Triggers

- Vendor cannot identify all AI components in their product — escalate to CISO before proceeding
- Vendor cannot confirm BAA availability for healthcare data — stop evaluation immediately
- Vendor scores below 55% overall — do not approve, document rejection rationale
- Vendor refuses to provide SOC 2 report under NDA — treat as a red flag equivalent to failed audit

Annual Reassessment

AI vendor risk is not static. Models change, terms of service change, compliance certifications expire, and vendor security postures evolve. Build annual reassessment into your vendor management calendar for any AI vendor with access to sensitive data. At minimum, request an updated questionnaire response and verify that compliance certifications are current.

Keeping This Document Current

The AI regulatory landscape is changing rapidly. EU AI Act compliance deadlines are rolling out through 2026. US state AI legislation is accelerating. New attack vectors against AI systems are being discovered regularly. This questionnaire should be reviewed and updated at least annually to reflect current requirements.

For updates to this document and additional AI security resources, visit itknowledgebases.com.

10. GLOSSARY

BAA (Business Associate Agreement)

A legally required contract under HIPAA between a covered entity and any vendor handling Protected Health Information on their behalf. Required before any healthcare data can be processed by an AI system.

CAIQ (Consensus Assessments Initiative Questionnaire)

A Cloud Security Alliance questionnaire designed for cloud consumers and auditors to assess cloud service provider security capabilities.

Data Poisoning

An AI-specific attack where an adversary manipulates training data to cause a model to behave incorrectly or produce biased outputs.

DPA (Data Processing Agreement)

A contract governing how a data processor handles personal data on behalf of a data controller. Required under GDPR for any vendor processing EU personal data.

EU AI Act

European Union regulation classifying AI systems by risk level and imposing obligations on developers and deployers of AI. Partially in force as of 2024 with rolling compliance deadlines through 2026.

FedRAMP

Federal Risk and Authorization Management Program. US government framework cloud providers must meet before federal agencies can use their services.

Fine-Tuning

The process of further training a pre-trained AI model on a specific dataset to specialize its behavior. Creates risk if customer data is used in shared fine-tuning processes.

Hallucination

When an AI model generates false, fabricated, or nonsensical information and presents it as factual. A fundamental characteristic of large language models, not a bug.

MITRE ATLAS

Adversarial Threat Landscape for Artificial-Intelligence Systems. A knowledge base of adversarial tactics and techniques targeting machine learning systems.

Model Extraction

An attack where an adversary queries an AI model repeatedly to reconstruct its behavior or training data.

OWASP LLM Top 10

The Open Web Application Security Project's list of the top 10 security risks specific to Large Language Model applications.

PHI (Protected Health Information)

Any health information linked to an individual that is created, received, stored, or transmitted by a HIPAA-covered entity.

Prompt Injection

An attack where malicious content in an AI input causes the model to ignore its instructions and perform unauthorized actions.

SBOM (Software Bill of Materials)

A formal record of all components in a software product. An AI SBOM specifically documents all AI models, APIs, and dependencies.

SCCs (Standard Contractual Clauses)

EU-approved contract clauses enabling lawful transfer of personal data outside the EU to countries without adequacy decisions.

SIG (Standardized Information Gathering)

A comprehensive vendor risk questionnaire from Shared Assessments covering cybersecurity, IT, privacy, data security, and business resiliency.

SOC 2 Type II

A security audit certification covering security, availability, processing integrity, confidentiality, and privacy over a defined period. Type II covers actual operating effectiveness, not just design.

TPRM (Third-Party Risk Management)

The process of identifying, assessing, and managing risks introduced by vendors, suppliers, and other third parties.

More Resources from itknowledgebases.com

→ AI Acceptable Use Policy Template — Tell your team what is and is not allowed

→ AI Tool Intake & Approval Form — Process for reviewing new AI tool requests

→ Data Sovereignty Quick-Start Guide — Free resource on AI data compliance

CVE writeups · Security tool reviews · Certification guides · Free security tools

www.itknowledgebases.com

DISCLAIMER

This document is provided for general informational purposes only and does not constitute legal, regulatory, or compliance advice. Laws and regulations vary by jurisdiction and change frequently. Consult qualified legal counsel for guidance specific to your organization. itknowledgebases.com makes no warranties as to accuracy or completeness of this material.

© 2026 itknowledgebases.com | David Wolfcale | All Rights Reserved