BASTION INTELLIGENCE · FREE TEMPLATE

Healthcare AI Vendor Questionnaire

A due-diligence questionnaire for assessing vendors whose products or services use artificial intelligence, built for US healthcare organizations. It covers PHI handling, training-data use, bias, validation, security, human oversight, monitoring, incident response, the supply chain, and exit terms, with a risk-tier screener, an evidence checklist, and a scoring rubric.

Version 1.0 · July 2026 · Published by Bastion Intelligence · **bastionintelligence.com**

Organization	
Vendor being assessed	
Product and version	
Risk tier assigned	
Date sent	
Response due	
Return completed form to	

Companion article, with best practices and sources: bastionintelligence.com/research/healthcare-ai-vendor-questionnaire

How to use this questionnaire

This questionnaire extends standard vendor security due diligence with the risk domains AI introduces: training-data use, bias, hallucination, drift, opaque fourth-party model dependencies, and fast-moving state AI laws. It reflects HIPAA requirements, the ONC HTI-1 transparency attributes, FDA lifecycle guidance, the Joint Commission and CHAI Responsible Use of AI in Healthcare guidance, NIST AI RMF, ISO/IEC 42001, and HITRUST AI controls.

1. Screen and tier the vendor. Run the nine-question screener on the next page. It assigns Tier 1 (high), Tier 2 (moderate), or Tier 3 (low) and tells you which form to send. A patient-facing clinical tool and an internal document summarizer should not get the same questionnaire.
2. Send the right form. Tier 1 and 2 vendors get all 60 questions. Tier 3 vendors get the 24 short-form questions marked ●. Set a response deadline (two to three weeks is typical) and name a return contact.
3. Require evidence, not assurances. Every question allows a narrative answer plus attached evidence. Treat unsupported answers as unanswered. The evidence checklist near the end of this document lists the artifacts to expect.
4. Score the responses. Use the 1-to-5 anchors and the weighted scoring sheet. Convert scores into a decision: approve, approve with conditions and a remediation plan, or reject. Any critical red flag stops the process regardless of score.
5. Push answers into the contract. Strong questionnaire answers become contract terms: BAA clauses (no PHI training, retention limits, subcontractor flow-down), performance thresholds, model-change notification, audit rights, breach timelines, insurance minimums, and data return at exit.
6. Reassess on a schedule. At least annually, at renewal, and on any material model change. High-risk vendors warrant quarterly touchpoints. Pair the questionnaire with continuous monitoring rather than treating it as a point-in-time exercise.

Route completed questionnaires through cross-functional review: compliance, security, privacy, legal, clinical leadership, and the business owner. Most US hospitals now assign this to a dedicated AI governance committee.

Section A: Vendor profile

Completed by the vendor.

Legal entity name	
Operating / d/b/a name	
Headquarters address	
Website	
Product(s) and version(s) in scope	
Deployment model (SaaS, on-premises, API, embedded feature)	
Primary contact (name, title, email)	
Security contact (name, title, email)	
Number of US healthcare customers	
Years serving healthcare customers	
Date this questionnaire was completed	

Section B: Risk-tier screener

Completed by the purchasing organization, with vendor input where needed. Answer all nine questions, then assign a tier using the table below.

#	Screening question	Yes	No	Notes
S1	Will the vendor create, receive, maintain, or transmit PHI on our behalf?	☐	☐	
S2	Is any AI output patient-facing (portal messages, chatbots, patient instructions)?	☐	☐	
S3	Does the AI inform or influence clinical decisions (diagnosis, treatment, triage, risk scoring, utilization review)?	☐	☐	
S4	Can the system take any action without a human approving each output?	☐	☐	
S5	Would a single AI failure affect many patients at once (population-level tools, automated outreach)?	☐	☐	
S6	Is the AI embedded in a system of record (EHR, billing, scheduling), or can it write to one?	☐	☐	
S7	Is any component a medical device, or has the vendor made FDA-related claims?	☐	☐	
S8	Does the vendor rely on third-party foundation models or subprocessors that would access our data?	☐	☐	
S9	Is this an AI feature added to an existing non-AI product or contract? (Existing BAAs may not cover it.)	☐	☐	

Tier assignment

Tier	Assign when	Form to send	Reassessment
Tier 1 · High	Yes to S1, and yes to any of S2 through S5.	Full questionnaire (all 60 questions), full evidence checklist.	Annual, plus quarterly touchpoints and on any material model change.
Tier 2 · Moderate	Yes to S1 with no to S2 through S5, or no PHI but yes to S3 or S6.	Full questionnaire, core evidence set.	Annual, and on material model change.
Tier 3 · Low	No PHI, internal-only, assistive use.	Short form: the 24 questions marked ●.	Annual attestation refresh.

A yes on S9 means the AI feature needs screening even though the base product is already under contract. AI features added mid-contract often fall outside the scope of the existing BAA and security review.

Section C: The questionnaire

Sixty questions across twelve domains. Questions marked ● form the Tier 3 short form. Answer in narrative form and attach evidence; check the Evidence box when an artifact accompanies the answer.

Domain 1: Product and AI Use Profile

Establish what the AI actually is and does before assessing how well it is governed.

ID	Question	Vendor response	Evidence
1.1 ●	Describe the AI system(s) in your product: model types (large language model, predictive machine learning, computer vision, other), whether each is customer-facing or internal, and every product function that uses AI.		☐ Attached ☐ N/A
1.2	Is AI used in the research, development, or production of your product or service, even where it is not customer-facing?		☐ Attached ☐ N/A
1.3 ●	What is the intended use, intended user, and target patient population for each AI function? What uses are explicitly out of scope?		☐ Attached ☐ N/A
1.4	Which functions produce a prediction, classification, recommendation, evaluation, or analysis relating to patients or their care? (This is the test for a predictive decision support intervention under ONC HTI-1.)		☐ Attached ☐ N/A
1.5	At what autonomy level does each AI function operate: assistive (a human reviews every output), conditional automation, or autonomous?		☐ Attached ☐ N/A

Domain 2: Regulatory Status and Compliance

Map the product to FDA, ONC, and state-law obligations, and confirm the vendor's compliance history.

ID	Question	Vendor response	Evidence
2.1	Is any component of the product a medical device (software as a medical device)? If yes, list the clearance pathway (510(k), De Novo, PMA), listing numbers, and whether an FDA-authorized Predetermined Change Control Plan (PCCP) governs post-market model updates.		☐ Attached ☐ N/A
2.2	Is your product certified under the ONC Health IT Certification Program? Do you support the decision support intervention criterion (45 CFR 170.315(b)(11)) and provide the 31 source attributes for predictive DSIs?		☐ Attached ☐ N/A

ID	Question	Vendor response	Evidence
2.3	How do you support customer compliance with state AI laws, including Texas TRAIGA and SB 1188 disclosures, California AB 3030 disclaimers, and Colorado's automated decision-making requirements? How do you monitor new legislation?		☐ Attached ☐ N/A
2.4 ●	Which AI governance frameworks does your AI risk program align to (NIST AI RMF, ISO/IEC 42001, HITRUST AI Security Certification)? Provide certificates and scope statements.		☐ Attached ☐ N/A
2.5 ●	Have you been subject to any regulatory inquiry, enforcement action, or litigation involving your AI claims, marketing, or data practices in the past five years?		☐ Attached ☐ N/A

Domain 3: Data Governance and PHI Handling

The HIPAA core. A vendor that cannot answer these cleanly should not touch PHI.

ID	Question	Vendor response	Evidence
3.1 ●	Will you create, receive, maintain, or transmit protected health information (PHI)? Will you sign our Business Associate Agreement? List which products, tiers, and features the BAA covers and which it excludes.		☐ Attached ☐ N/A
3.2 ●	Is customer data (including PHI, prompts, uploads, and audio) used to train, fine-tune, or improve any model, whether shared or customer-specific? Can this be contractually disabled?		☐ Attached ☐ N/A
3.3 ●	What customer data is stored, where, and for how long, including logs, caches, and human-review queues? Can retention be set to zero or to a customer-defined period?		☐ Attached ☐ N/A
3.4	In which regions and countries is customer data processed and stored? Does any personnel access occur from outside the United States?		☐ Attached ☐ N/A
3.5	How do you enforce minimum-necessary access to PHI? Describe any de-identification methods used and the controls against re-identification.		☐ Attached ☐ N/A
3.6	Describe encryption in transit and at rest, key management, and tenant isolation.		☐ Attached ☐ N/A
3.7 ●	What happens to our data at contract termination? State the return and destruction		☐ Attached ☐ N/A

ID	Question	Vendor response	Evidence
	timeline and whether you provide written certification of destruction.		

Domain 4: Model Transparency and Documentation

If the vendor cannot document the model, you cannot govern it.

ID	Question	Vendor response	Evidence
4.1 ●	Provide a completed CHAI Applied Model Card (or an equivalent model card) for each model or AI solution in scope.		☐ Attached ☐ N/A
4.2 ●	Which foundation or base models does the product use (provider, model name, and version)? Are they proprietary, open-source, or third-party APIs?		☐ Attached ☐ N/A
4.3	How are model versions identified, and how are customers told which version they are running?		☐ Attached ☐ N/A
4.4	What documentation covers known limitations, failure modes, and contraindicated uses?		☐ Attached ☐ N/A
4.5	Explain, at a level appropriate for clinicians, how the system produces its outputs (your explainability or interpretability approach).		☐ Attached ☐ N/A

Domain 5: Training Data and Bias

Bias that enters through training data surfaces in patient care. Require evidence, not assurances.

ID	Question	Vendor response	Evidence
5.1	Describe training-data sources, provenance, licensing and intellectual property status, and the time period covered.		☐ Attached ☐ N/A
5.2	What is the demographic composition of training and validation data (race, ethnicity, language, age, sex and gender, payer type, geography)? How representative is it of our patient population?		☐ Attached ☐ N/A
5.3 ●	Describe your bias detection and mitigation methodology, the fairness metrics you use, and your testing frequency. Attach your most recent bias audit results.		☐ Attached ☐ N/A
5.4	How do you evaluate performance across demographic subgroups? What disparities have been identified, and how were they addressed?		☐ Attached ☐ N/A

ID	Question	Vendor response	Evidence
5.5	How does the product support our own FAVES assessment (fair, appropriate, valid, effective, safe)?		☐ Attached ☐ N/A

Domain 6: Performance and Validation

Marketing claims are not evidence. Tie every claim to a metric, a dataset, and test conditions.

ID	Question	Vendor response	Evidence
6.1 ●	Report performance metrics tied to each claimed benefit (for example AUROC, sensitivity and specificity, precision and recall, error and hallucination rates), with the datasets and conditions used to produce them.		☐ Attached ☐ N/A
6.2	Provide external validation evidence: peer-reviewed studies, independent evaluations, or real-world performance data.		☐ Attached ☐ N/A
6.3	What support do you provide for local, site-specific validation before go-live (test environments, validation datasets, statistical support)?		☐ Attached ☐ N/A
6.4 ●	For generative AI: how do you measure and reduce hallucination or confabulation, and what guardrails constrain outputs?		☐ Attached ☐ N/A
6.5	Which benchmarks or acceptance thresholds will you commit to contractually?		☐ Attached ☐ N/A

Domain 7: Security

General security attestations plus the AI-specific attack surface.

ID	Question	Vendor response	Evidence
7.1 ●	Which attestations and certifications do you hold: SOC 2 Type II, HITRUST (e1, i1, or r2, and the AI Security Certification), ISO/IEC 27001, ISO/IEC 42001? Attach reports and scope.		☐ Attached ☐ N/A
7.2 ●	How do you defend against the OWASP Top 10 risks for LLM applications, including prompt injection, insecure output handling, training-data poisoning, model theft, and sensitive information disclosure?		☐ Attached ☐ N/A
7.3	Describe access controls for models and data: multi-factor authentication, least privilege,		☐ Attached ☐ N/A

ID	Question	Vendor response	Evidence
	privileged access management, and audit logging.		
7.4	Describe your penetration testing and red-teaming program for AI components. Attach the latest summary and remediation status.		☐ Attached ☐ N/A
7.5	How is customer data segregated from other tenants and from model-improvement pipelines?		☐ Attached ☐ N/A
7.6	Describe vulnerability management and secure development practices, including model supply chain integrity.		☐ Attached ☐ N/A

Domain 8: Human Oversight and Clinical Safety

The boundary between supporting a clinician and replacing one. State laws increasingly draw it in statute.

ID	Question	Vendor response	Evidence
8.1 ●	What human-in-the-loop controls exist for high-stakes outputs? Can any output execute an action without human review?		☐ Attached ☐ N/A
8.2	How does the workflow ensure a licensed practitioner reviews AI-generated content before it informs clinical decisions (as Texas SB 1188 and similar laws require)?		☐ Attached ☐ N/A
8.3	Can users override, correct, or flag outputs? How is that feedback captured and acted on?		☐ Attached ☐ N/A
8.4 ●	What in-product transparency tells users that content is AI-generated (supporting California AB 3030, Texas TRAIGA, and similar disclosure laws)?		☐ Attached ☐ N/A
8.5	Describe safety guardrails for harmful, out-of-scope, or crisis-related outputs, especially in patient-facing tools.		☐ Attached ☐ N/A

Domain 9: Monitoring, Drift, and Change Management

AI systems change faster than traditional software. Assess the vendor's discipline around change.

ID	Question	Vendor response	Evidence
9.1	Describe post-deployment monitoring: the metrics tracked, drift-detection methods, alert thresholds, and any dashboards available to customers.		☐ Attached ☐ N/A

ID	Question	Vendor response	Evidence
9.2 ●	How and when are customers notified of model updates, retraining, or material behavior changes? Is there a customer-facing changelog?		☐ Attached ☐ N/A
9.3	Can customers pin a model version or opt out of updates? What is your rollback process?		☐ Attached ☐ N/A
9.4	For FDA-regulated functions: how do updates map to the authorized Predetermined Change Control Plan?		☐ Attached ☐ N/A
9.5	Who at your company owns model performance in production, and what is the escalation path when performance degrades?		☐ Attached ☐ N/A

Domain 10: Incident Response and AI Safety Events

Covers both security incidents and AI-specific safety events, such as harmful outputs and systematic errors.

ID	Question	Vendor response	Evidence
10.1 ●	Describe your incident response plan for security incidents and for AI-specific safety events (harmful outputs, systematic errors). When was it last tested?		☐ Attached ☐ N/A
10.2 ●	What breach and incident notification timelines will you commit to contractually? (Our target: initial notice within 24 to 72 hours of discovery.)		☐ Attached ☐ N/A
10.3	Have you experienced a reportable breach or AI safety event in the past 36 months? Describe the scope and remediation.		☐ Attached ☐ N/A
10.4	Do you participate in blinded AI safety-event reporting, such as the CHAI Health AI Registry?		☐ Attached ☐ N/A

Domain 11: Fourth Parties and Supply Chain

Your vendor's vendors are your risk. Foundation-model providers and cloud subprocessors need downstream agreements.

ID	Question	Vendor response	Evidence
11.1 ●	List every subprocessor that touches our data: foundation-model APIs, cloud providers, vector databases, moderation and analytics tooling.		☐ Attached ☐ N/A
11.2 ●	Do downstream Business Associate Agreements or data processing agreements exist with every subprocessor that handles PHI? Provide evidence.		☐ Attached ☐ N/A

ID	Question	Vendor response	Evidence
11.3	How are we notified of subprocessor changes, and can we object?		☐ Attached ☐ N/A
11.4	What happens to service continuity and our data if a critical fourth party (for example your foundation-model provider) changes terms, deprecates a model, or fails?		☐ Attached ☐ N/A

Domain 12: Business Terms, Viability, and Exit

Plan the exit before the entrance. AI products fail, pivot, and get acquired.

ID	Question	Vendor response	Evidence
12.1	Describe your financial stability: years in operation, funding or profitability posture, healthcare customer references, and your experience maintaining AI systems over multiple years.		☐ Attached ☐ N/A
12.2 ●	Which contract protections do you offer: performance guarantees, indemnification (including intellectual property and AI-output claims), cyber liability and errors-and-omissions insurance minimums, and audit rights?		☐ Attached ☐ N/A
12.3 ●	Describe data portability and export formats at exit, decommissioning support, and timelines.		☐ Attached ☐ N/A
12.4	Describe your pricing model and any usage-based cost drivers tied to AI, such as token or consumption exposure.		☐ Attached ☐ N/A

Section D: Evidence checklist

Artifacts the vendor should provide with its responses. Absence of an artifact is information: score accordingly.

- ☐ Signed Business Associate Agreement, or the vendor's BAA template with covered products and exclusions
- ☐ SOC 2 Type II report (most recent period)
- ☐ HITRUST certification letter and scope (e1, i1, or r2; AI Security Certification if held)
- ☐ ISO/IEC 42001 certificate and scope statement
- ☐ ISO/IEC 27001 certificate and scope statement
- ☐ CHAI Applied Model Card, or an equivalent model card, for each model in scope
- ☐ Most recent bias and fairness audit summary, with the metrics used
- ☐ Penetration test or AI red-team summary and remediation status
- ☐ Incident response plan attestation and date of last test
- ☐ Complete subprocessor list with downstream BAA or DPA status for each
- ☐ Certificates of insurance: cyber liability and errors and omissions, with limits
- ☐ Sample state-disclosure language or in-product screenshots (Texas, California, and other applicable states)
- ☐ FDA clearance letters and Predetermined Change Control Plan, if any function is a regulated device
- ☐ Data flow diagram covering PHI, prompts, outputs, logs, and human-review queues
- ☐ Data retention schedule and destruction certification template
- ☐ External validation evidence: peer-reviewed studies or independent evaluations

Section E: Scoring rubric

Score anchors

Score	Anchor
5	Complete, specific answer supported by independent third-party evidence (certification, audit, external validation).
4	Complete, specific answer supported by internal evidence (policies, test results, documentation).
3	Partial answer; evidence promised but not yet provided, or provided with material gaps.
2	Vague or evasive answer; no evidence.
1	No answer, refusal, or an answer contradicted by other responses or public information.

Domain scores

Score each domain 1 to 5 using the anchors, multiply by the weight, and compute the weighted average. Weights shown are Tier 1 defaults; for Tier 2, set Domain 3 and Domain 7 to 1.5 and all others to 1.0; for Tier 3, weight all domains equally.

Domain	Weight (Tier 1 default)	Score (1–5)	Weighted score
Domain 1: Product and AI Use Profile	1.0		
Domain 2: Regulatory Status and Compliance	1.0		
Domain 3: Data Governance and PHI Handling	2.0		
Domain 4: Model Transparency and Documentation	1.0		
Domain 5: Training Data and Bias	1.5		
Domain 6: Performance and Validation	1.5		
Domain 7: Security	1.5		
Domain 8: Human Oversight and Clinical Safety	1.5		
Domain 9: Monitoring, Drift, and Change Management	1.0		
Domain 10: Incident Response and AI Safety Events	1.0		
Domain 11: Fourth Parties and Supply Chain	1.0		
Domain 12: Business Terms, Viability, and Exit	1.0		
Weighted average (sum of weighted scores ÷ sum of weights)			

Decision thresholds

- 4.0 or higher: approve. Convert commitments into contract terms before signature.
- 3.0 to 3.9: approve with conditions. Document findings in the remediation plan below, assign owners and dates, and re-verify before go-live.
- Below 3.0: do not proceed. Require remediation and a full re-score, or reject.
- Any critical red flag (next page): stop regardless of score until resolved.

Remediation plan

#	Finding	Severity	Owner	Due date	Status

Section F: Critical red flags

Each of these stops the assessment until resolved, regardless of the weighted score.

- PHI or customer data is used to train shared models by default, or answers about training-data use stay vague after follow-up.
- The vendor will not sign a BAA, or the BAA excludes the AI features you actually plan to use.
- No model card, and no clear account of training-data provenance, limitations, or intended use.
- No bias testing, or the vendor refuses to share fairness metrics and audit results.
- No third-party attestations (SOC 2, HITRUST, ISO 27001 or 42001) for a vendor handling sensitive data.
- No documented, tested incident response plan, or vague breach-notification timelines.
- The vendor cannot name its foundation-model providers or subprocessors, or confirm downstream BAAs.
- No human-oversight design for high-stakes outputs, and no override or escalation path.
- No process for notifying customers of model updates or performance drift.
- Accuracy or hallucination claims with no supporting evidence. (Unsubstantiated accuracy claims drew the first state AI enforcement action in healthcare: Texas AG v. Pieces Technologies, 2024.)

Section G: Vendor attestation

An authorized representative of the vendor certifies that the responses in this questionnaire are accurate and complete as of the date below, and agrees to notify the purchasing organization in writing within 30 days of any material change to these responses, including material model changes, new subprocessors, changes to data-handling practices, and reportable security or AI safety events. This attestation anticipates the annual written verification of business associate safeguards proposed in the January 2025 HIPAA Security Rule update.

Vendor legal entity	
Authorized representative (name)	
Title	
Signature	
Date	