MDA SOLUTIONS LLC

8-Question HIPAA AI Risk Checklist

A quick governance screen for AI tools that may touch PHI or influence care decisions.

Who it's for	What it covers
Healthcare executives, compliance and privacy officers, CISOs, EHR and AI implementation Directors/Managers, and consultants evaluating AI governance readiness.	Inventory, PHI flows, minimum necessary, risk analysis, vendor/BAA terms, transparency, validation and monitoring, and ongoing governance evidence.

6–8 Yes Green zone	3–5 Yes Yellow zone	0–2 Yes Red zone

Educational resource — not legal advice. | mdasolutionsllc.com | info@mdasolutionsllc.com

How to use this checklist

Walk through the eight questions for a single AI tool, vendor, or use case at a time. For each, mark **Yes**, **Partial**, **No**, or **Unsure**. **No** and **Unsure** indicate gaps to resolve — or evidence to gather — before deployment or expansion. Capture the responsible owner and remediation note in the space provided.

6–8 Yes	3–5 Yes	0–2 Yes
Green zone — reasonable readiness; continue periodic review.	Yellow zone — meaningful gaps; scope remediation before expansion.	Red zone — pause or restrict use until core controls are in place.

Zones describe governance readiness only. A high score does not establish HIPAA compliance, and a low score does not by itself indicate a violation.

01 Inventory

Have we identified every AI tool, chatbot, model, automation, or vendor system currently used or piloted?

WHY IT MATTERS
You cannot govern what you have not catalogued. Shadow AI — tools adopted by individual teams or embedded inside vendor products — is a primary source of unmanaged PHI exposure.

EVIDENCE TO LOOK FOR
A current inventory listing each AI feature or vendor, owner, business purpose, data inputs, deployment status (pilot vs. production), and whether it touches PHI/ePHI.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner: _____

02 PHI/ePHI Access

Do we know whether the tool creates, receives, maintains, transmits, stores, or learns from PHI/ePHI?

WHY IT MATTERS
HIPAA obligations are triggered by how a tool interacts with PHI. Training, fine-tuning, prompt logging, and vendor telemetry can all constitute PHI handling that needs safeguards.

EVIDENCE TO LOOK FOR
A documented data flow per tool covering inputs, outputs, storage locations, training/fine-tuning use, log retention, and any cross-border or third-party processing.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner: _____

03 Purpose & Minimum Necessary

Is the use case documented, and is PHI access limited to the minimum necessary for that purpose?

WHY IT MATTERS	EVIDENCE TO LOOK FOR
HIPAA's minimum necessary standard applies to most uses and disclosures. Generic 'all-data' access for AI tools rarely meets it and expands breach blast radius.	A written use case, defined data fields and record scope, role-based access controls, prompt/data filters, and rationale for why each PHI element is required.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner:

_____

04 Risk Analysis

Has the tool been included in a HIPAA Security Rule risk analysis with threats, vulnerabilities, likelihood, impact, and safeguards documented?

WHY IT MATTERS	EVIDENCE TO LOOK FOR
OCR has consistently identified inadequate or missing risk analysis as a top enforcement issue. AI introduces new threat vectors (prompt injection, model leakage, unsanctioned training).	Risk analysis artifacts: asset listing, threat/vulnerability pairs, likelihood and impact ratings, existing and planned administrative, physical, and technical safeguards, and a remediation tracker.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner:

_____

05 Vendor & BAA

If a vendor touches PHI, do we have an appropriate BAA plus AI-specific security, data-use, retention, and breach notification terms?

WHY IT MATTERS	EVIDENCE TO LOOK FOR
Standard BAAs predate generative AI. Without explicit terms, vendors may train on your PHI, retain prompts indefinitely, or use sub-processors you have not approved.	Executed BAA, plus contract terms covering: no training on PHI without authorization, data residency, retention/deletion, sub-processor list and notification, breach notice timing, and audit/SOC 2 evidence.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner:

_____

06 Transparency & Human Oversight

Do clinicians, staff, and/or patients receive appropriate information about AI use, limitations, and when humans must review or override outputs?

WHY IT MATTERS	EVIDENCE TO LOOK FOR
Clinical and operational decisions influenced by AI need clear accountability. Lack of disclosure and oversight raises patient safety, malpractice, and trust risks beyond HIPAA itself.	Disclosure language, clinician-facing model cards or use guidance, defined human-in-the-loop checkpoints, override workflows, and documented training for end users.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner:

_____

07 Bias, Validation & Performance Monitoring

Have we evaluated whether the tool performs safely and fairly across relevant patient populations and monitored for drift?

WHY IT MATTERS	EVIDENCE TO LOOK FOR
Models can degrade or perform unevenly across demographics, sites, or workflows. Unmonitored drift can produce systematically unsafe or inequitable recommendations.	Pre-deployment validation results, subgroup performance analysis, ongoing monitoring metrics and thresholds, incident logs, and a defined retraining or rollback trigger.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner:

_____

08 Governance Evidence

Is there a named owner, approval record, audit trail, training plan, and periodic review schedule?

WHY IT MATTERS

Durable governance — not a one-time review — is what regulators, boards, and partners expect. Each AI use should have a clear human accountable for it.

EVIDENCE TO LOOK FOR

Named accountable owner, approval memo or governance committee minutes, access and use audit logs, role-based training records, and a calendar for periodic re-review.

STATUS [ ] Yes [ ] Partial [ ] No [ ] Unsure

Notes / owner:

_____

Next steps

Five actions that turn checklist results into a defensible governance posture:

1. **Inventory your AI tools**
   Build (or refresh) a single source of truth that lists every AI feature, model, chatbot, automation, and embedded vendor capability — including pilots.

2. **Map PHI / ePHI flows**
   For each tool, document how PHI enters, where it goes, what is logged or retained, and whether any data is used to train or fine-tune models.

3. **Update your risk analysis**
   Bring AI assets into your HIPAA Security Rule risk analysis. Add AI-specific threats (e.g., prompt injection, model leakage, unsanctioned training) and document safeguards.

4. **Review vendors and BAAs**
   Confirm BAAs are in place where required and add AI-specific terms covering training restrictions, retention, sub-processors, and breach notification timing.

5. **Assign a governance owner and review cadence**
   Name an accountable owner per tool, define an approval workflow, and set a recurring review schedule (e.g., quarterly for high-risk uses).

Regulatory context

HIPAA Security Rule (HHS OCR). Regulated entities must implement reasonable and appropriate administrative, physical, and technical safeguards for ePHI, including risk analysis and risk management.

Risk analysis guidance (HHS OCR). Risk analysis is foundational and requires an accurate and thorough assessment of potential risks and vulnerabilities to the confidentiality, integrity, and availability of ePHI.

HIPAA Security Rule NPRM (HHS OCR, Dec. 27, 2024) — proposed rulemaking, not final. Would require written documentation of Security Rule policies, procedures, plans, and analyses, plus asset inventories and network maps and greater specificity in risk analysis.

FDA AI-enabled device software functions (draft guidance). Provides recommendations for marketing submissions for AI-enabled device software functions and reflects risk management throughout the device total product life cycle.

Ready to go deeper?
MDA Solutions offers a HIPAA AI Quick Scan starting at $1,500 — a 60-minute review, risk scorecard, top-5 gaps, and readout. Book a 30-minute strategy consultation at **mdasolutionsllc.com** or email **info@mdasolutionsllc.com**.