

✦ GUIDE · COMPLIANCE

Healthcare AI Compliance Guide

Consent, data minimization, de-identification, audit trails, human oversight, vendor contracts, data residency and model change control for AI that touches patient data.

FORMAT	LEVEL	READ TIME	UPDATED
Guide	Intermediate	16 min	Apr 2026

OVERVIEW

What this guide covers

Healthcare AI projects rarely fail on model quality. They stall in security review, get paused after an incident, or never leave the pilot because nobody can show where patient data went. This guide sets out the compliance controls we design in from day one, so that an assistant can be explained to a privacy officer, a clinical safety lead and an auditor before it reaches a single patient.

KEY TAKEAWAYS

- Map every data flow, including prompts, retrieved context, logs and evaluation sets, before choosing a model or vendor.
- Consent and purpose limitation decide what the AI may see; enforce them in the data layer, not in the prompt.
- De-identify for analytics and evaluation, but treat de-identified data as re-identifiable until proven otherwise.
- AI assists clinicians — a qualified human reviews and signs off anything that affects diagnosis or treatment.
- Every model, prompt and retrieval change is a controlled change with an evaluation record and a rollback path.

WHO IT'S FOR

Hospital and clinic CIOs / CTOs

Privacy officers and DPOs

Clinical safety and quality leads

Health-tech product and engineering teams

HOW TO USE IT

Use section 01 to see which regimes apply, then treat 02–06 as design controls. The checklist in 07 doubles as an evidence pack. Not legal advice; confirm obligations with counsel.

CONTENTS

01 Which rules apply to you
02 Consent, notice and purpose limitation
03 Data minimization and de-identification
04 Clinical safety and human oversight
05 Audit trails and access control
06 Vendors, contracts and data residency
07 Model change control and the evidence pack

01 / SCOPE

Which rules apply to you

Obligations follow the patient, the data and the function of the software — not where your servers happen to be. Most healthcare AI programs sit under more than one regime at once.

Common regimes for healthcare AI

REGIME	APPLIES WHEN	WHAT IT MEANS FOR AI
HIPAA (US)	You are a covered entity or business associate handling PHI	Privacy, Security and Breach Notification Rules; minimum necessary use; BAAs with every vendor touching PHI
DPDP Act 2023 (India)	You process digital personal data of individuals in India	Notice and consent, purpose limitation, security safeguards, breach reporting, verifiable parental consent for children
ABDM (India)	You integrate with Ayushman Bharat Digital Mission rails	Consent-managed exchange of health records via ABHA and the consent manager; follow the Health Data Management Policy
GDPR (EU / UK GDPR)	You process data of people in the EU or UK	Health data is special category (Art. 9); DPIA for high-risk processing; limits on solely automated decisions (Art. 22)
Medical device rules	The software informs diagnosis or treatment decisions	May qualify as software as a medical device under FDA, EU MDR or India's Medical Devices Rules 2017 — assess intended use early

INTENDED USE

Decide what the AI is for before you build it

An assistant that drafts discharge summaries for a clinician to edit is a different regulatory object from one that suggests a diagnosis. Write the intended use down, keep the product inside it, and re-assess when features expand.

02 / CONSENT

Consent, notice and purpose limitation

Consent is not a checkbox at sign-up. It is a set of purposes, each of which the system must be able to check at the moment data is used.

01

Purpose-specific consent

Record consent per purpose — care delivery, AI-assisted documentation, analytics, model improvement — rather than one blanket agreement.

02

Plain-language notice

Tell patients that AI is used, what it does, and that a clinician remains responsible. Offer the notice in the languages your patients read.

03

Withdrawal that works

Withdrawing consent must stop future processing for that purpose, and the system should be able to prove it did.

04

Children and guardians

Under DPDP, processing a child's data needs verifiable parental consent; plan the guardian flow and its records explicitly.

Enforcing consent in the data layer

01 **Model consent as data**

Store purpose, scope, grant time, expiry and withdrawal per patient in a consent service, not in application flags.

02 **Tag records with purpose eligibility**

Index and retrieval filters check consent for the current purpose before a record can reach a prompt.

03 **Fail closed**

If the consent service is unavailable or ambiguous, the AI path is skipped and the clinician works without it.

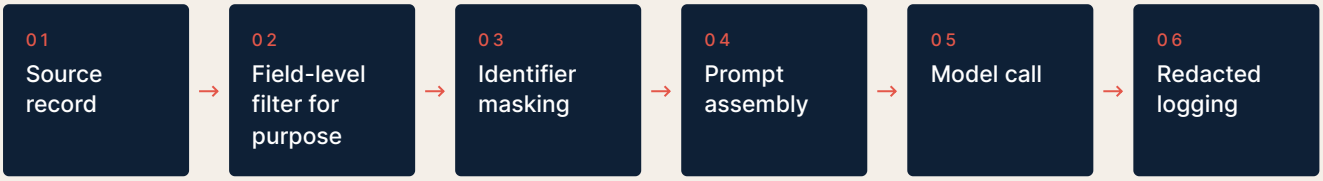
04 **Log the consent decision**

Each AI request records which consent it relied on, so audits can trace use back to permission.

03 / DATA

Data minimization and de-identification

The safest patient data is the data the model never sees. Minimize what goes into prompts, logs and evaluation sets, and de-identify wherever identity adds nothing.



COMMON SHORTCUTS
Full chart pasted into every prompt
Raw prompts and outputs kept in vendor logs
Production data copied into test environments
Evaluation sets built from identifiable notes

MINIMIZED BY DESIGN
Only fields the task needs are retrieved
Logs redact identifiers and expire on a schedule
Synthetic or de-identified data outside production
Golden sets de-identified and access-controlled

De-identification in practice

- Under HIPAA, data is de-identified by **Safe Harbor** (removing the 18 listed identifiers) or **Expert Determination** — document which method you used.
- Free-text notes hide identifiers in unexpected places: names in quotes, dates in narratives, rare conditions, small locations. Test redaction on real samples.
- Pseudonymized data is still personal data under GDPR and DPDP; keep the re-identification key separate and access-logged.
- Re-assess re-identification risk when datasets are combined — joining two safe tables can produce an unsafe one.

✦ 04 / SAFETY

Clinical safety and human oversight

AI in a clinical setting assists qualified people. It does not diagnose, prescribe or decide on its own, and the workflow must make that true in practice, not only in the policy document.

DESIGN

Draft, never final

AI output enters the record only after a clinician reviews, edits and signs it.

DESIGN

Show the evidence

Summaries and suggestions cite the source notes, labs or guidelines they came from.

DESIGN

Visible AI labeling

Clinicians can always tell which text was AI-drafted and which they wrote.

GUARDRAIL

Out-of-scope refusal

Requests outside the intended use — dosing, diagnosis, triage — are declined or routed to a human.

GUARDRAIL

Escalation paths

Red-flag phrases in patient-facing chat hand off to staff with defined response times.

MONITORING

Clinical incident loop

AI-related near misses feed the existing incident reporting system, not a separate inbox.

Hazard log starter

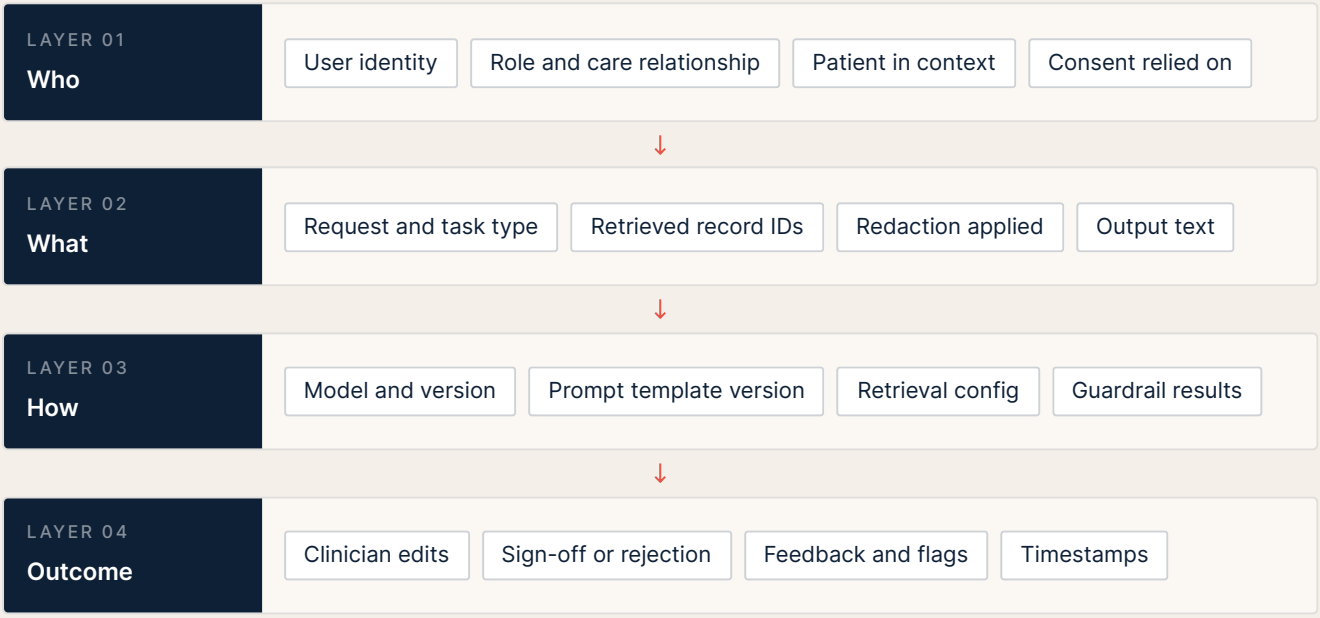
HAZARD	POSSIBLE HARM	CONTROL	OWNER
Hallucinated finding in a summary	Wrong clinical picture	Source citation + clinician sign-off	Clinical safety lead
Missed allergy or contraindication	Adverse event	Structured data shown alongside AI text	Product owner
Wrong patient context	Privacy breach and wrong advice	Patient ID check before prompt assembly	Engineering lead
Over-reliance by staff	Reduced clinical scrutiny	Training, sampling audits, edit-rate monitoring	Department head

05 / EVIDENCE

Audit trails and access control

When a regulator, patient or clinician asks why the AI said something, you need to reconstruct the answer exactly: who asked, what data was used, which model and prompt version produced it, and who approved it.

What an AI audit record captures



- Use role-based access tied to the care relationship: an AI request should only retrieve records the requesting clinician could open directly.
- Store audit records append-only, separate from application data, with access to the logs themselves logged.
- Agree a retention period with your privacy and legal teams and apply it consistently to prompts, outputs and traces.
- Run periodic access reviews and sample AI outputs for quality alongside your usual clinical audits.

06 / VENDORS

Vendors, contracts and data residency

Model providers, vector databases, transcription services and observability tools may all process patient data. Each one needs a contract, a residency answer and a deletion story.

QUESTION	WHY IT MATTERS	EVIDENCE TO REQUEST
Will you sign a BAA / DPA?	Required for HIPAA business associates and GDPR processors	Signed agreement covering the AI services used
Is our data used for training?	Training on PHI can breach purpose limits	Contractual no-training term and settings
Where is data processed and stored?	Residency and cross-border transfer rules	Region list for inference, storage and logs
How long are prompts retained?	Hidden retention creates breach exposure	Retention config; zero-retention options
Which sub-processors are involved?	Your obligations flow down the chain	Current sub-processor list and change notice

RESIDENCY

Check every hop, not just the database

Teams often pin the primary database to a region and then send prompts to a model endpoint, embeddings to another service and traces to a third — each in a different country. Map the whole path. India's DPDP Act permits transfers except to countries the government restricts, while sector rules and contracts may require stricter local processing.

✦ 07 / CHANGE CONTROL

Model change control and the evidence pack

Models, prompts and retrieval settings change more often than any other part of a healthcare system. Treat each change like a clinical software release.

01 Version everything

Model, prompt templates, retrieval config and guardrails each carry a version recorded on every output.

02 Evaluate before release

Run the de-identified golden set, safety prompts and refusal tests; record results against the previous version.

03 Clinical sign-off

A clinical safety lead approves changes that could alter clinical content, with the hazard log updated.

04 Staged rollout

Release to a small group first, monitor edit rates and flags, then widen.

05 Rollback ready

The previous version stays deployable, and provider-side model updates are pinned rather than accepted silently.

GOVERNANCE

- ☐ Intended use written and approved
- ☐ Applicable regimes identified with counsel
- ☐ DPIA or risk assessment completed

DATA

- ☐ Consent checked per purpose at retrieval
- ☐ Prompts and logs minimized and redacted
- ☐ De-identification method documented

SAFETY

- [ ] Clinician sign-off on all clinical output
- [ ] Hazard log owned and current
- [ ] Escalation path for patient-facing chat

OPERATIONS

- [ ] Audit trail reconstructs any output
- [ ] BAA / DPA signed with every processor
- [ ] Model changes versioned, evaluated and reversible

We apply the same controls to our own platforms. BloomSenz, our therapy-centre platform, handles sensitive child and family records, so consent, role-scoped access, therapist sign-off and audit trails sit in the core data model rather than in individual AI features.

✦ GET IN TOUCH

Let's Talk

We design and build healthcare AI with compliance in the architecture — consent-aware retrieval, audit trails, clinician-in-the-loop workflows and vendor-ready evidence packs.

- Compliance-by-design architecture review
- Consent and audit layer implementation
- Clinical safety and evaluation setup

BOOK AN AI STRATEGY SESSION →

EMAIL
hello@splendenslabs.com

WEB
splendenslabs.com

OFFICE
Sona Incubation Foundation, Salem