Clinical-AI Compliance & Safety Checklist

Run any telemedicine AI feature through the four gates before you build it. Engineering guidance, not legal advice — confirm specifics with counsel.

GATE 1 · THE CONTRACT (the model is a business associate)

- ☐ Signed BAA with the model provider before any PHI is sent (164.502(e))
- ☐ BAA names the exact enterprise service/tier — not 'the vendor offers one'
- ☐ No-training clause: patient data never trains the vendor's future models
- ☐ BAA flows down to every sub-processor in the chain
- ☐ Free consumer AI endpoints blocked for any identifiable patient data

GATE 2 · THE PHI BOUNDARY (keep it in, send the minimum)

- ☐ Model runs inside the boundary (self-host or enterprise API + BAA)
- ☐ Send only the minimum-necessary fields, not the whole record (164.502(b))
- ☐ Analytics outside the clinical loop run on de-identified data (164.514(b))
- ☐ De-identification is Safe Harbor (18 categories) or Expert Determination
- ☐ 'Removed the name' is NOT de-identification — a DOB + ZIP can re-identify

GATE 3 · THE FDA LINE (support vs diagnosis)

- ☐ Feature organizes info for a clinician — it does not diagnose or direct care
- ☐ Recommendation goes to a clinician who can independently review the basis
- ☐ Not patient-facing: a symptom checker telling a patient a diagnosis is a device
- ☐ Output is explainable — 'the AI said so' fails the CDS carve-out
- ☐ If it IS a device: plan the FDA pathway + a PCCP for model updates

GATE 4 · HUMAN + VALIDATION + RED FLAGS

- ☐ A qualified human reviews and owns every output that affects care
- ☐ Safe default: escalate uncertain cases, never silently act
- ☐ Validated by patient subgroup — not one average that hides the worst group
- ☐ Bias identified and mitigated per Section 1557 (45 CFR 92.210); AI use disclosed
- ☐ RED FLAG: a patient-facing feature wired to a free, no-BAA AI endpoint

THE ONE-LINE RULE

Every clinical AI feature — the scribe, the chatbot, the vital-sign flag — passes the same four gates, and this is the checklist that defines them for the whole product. Gate one is the contract: any model that sees Protected Health Information is a business associate that needs a signed BAA with a no-training clause before it touches a single record, which rules out the free consumer endpoints teams reach for first. Gate two is the boundary: keep PHI inside that contracted perimeter, send the model only the minimum it needs, and de-identify (Safe Harbor or Expert Determination) anything used outside the live clinical loop. Gate three is the FDA line between software that supports a clinician who can review it and software that diagnoses, speaks a conclusion directly to a patient, or cannot be independently reviewed — cross it and the feature is a regulated medical device. Gate four is the human, the validation, and the bias test: a clinician owns the output, performance is measured by subgroup rather than hidden behind a single average, and discrimination risk is identified and mitigated as Section 1557 now requires. Fail any one gate and the feature is not ready.