AI Triage Safety & Compliance Checklist

Run a telehealth AI-triage feature through the four gates before you ship it. Engineering guidance, not legal advice — confirm specifics with counsel.

GATE 1 · SUPPORT, NOT DIAGNOSIS (the FDA line)

- [ ] Output describes and routes; never 'you have X, take Y' to a patient
- [ ] Clinician-facing recommendations show their basis to review (CDS criteria 3-4)
- [ ] Single recommendation only where one option is clinically appropriate
- [ ] Patient-facing checker treated as possibly a device — assessed, not assumed exempt
- [ ] Does not analyze a medical image or device signal (non-device criterion 1)

GATE 2 · PHI BOUNDARY (the model is a business associate)

- [ ] Signed BAA with the model provider before any symptom is sent (45 CFR 164.502(e))
- [ ] BAA covers the exact enterprise service/tier — not 'the vendor offers one'
- [ ] No-training clause: patient data never trains the vendor's future models
- [ ] Free consumer chatbots and LLM endpoints blocked for any symptom data
- [ ] Symptoms encrypted in transit and at rest; correct region and logging on

GATE 3 · BIAS + GATE 4 · SAFE DEFAULT

- [ ] Reasonable effort to identify race/age/disability inputs or proxies (45 CFR 92.210)
- [ ] Reasonable effort to mitigate bias; outcomes tested across groups, on a schedule
- [ ] Predictive-DSI transparency: what data trained it, on whom, how it performs
- [ ] Red-flag rules escalate immediately and never down-route (chest pain, stroke, SI)
- [ ] Human confirms routing for anything above self-care; model may only escalate the default
- [ ] Every step logged: suggestion, override, final route — for audit and bias review

RED FLAGS (stop before you ship)

- [ ] Patient-facing chatbot that diagnoses and directs instead of routing
- [ ] Symptoms piped to a consumer LLM with no BAA — a HIPAA breach
- [ ] Tool tuned for 'efficiency' that silently down-routes borderline cases
- [ ] No bias check, and a proxy (ZIP, language) doing protected-class work invisibly

THE ONE-LINE RULE

AI triage is software that asks a patient about symptoms before a visit and helps decide how urgently they need care and who should see them. The line that governs every choice is decision support versus diagnosis: the FDA's January 2026 Clinical Decision Support guidance keeps its 'not a device' carve-out for tools that support a clinician who can review the reasoning, and says plainly that device rules still apply to software for patients and caregivers — exactly where a symptom checker sits. Every typed symptom is Protected Health Information, so the model behind the chatbot is a business associate that needs a signed BAA, with a no-training clause, before it sees a word; and a routing algorithm that uses race, age, or a proxy must be checked and de-biased under 45 CFR 92.210. The safe design keeps AI as support with a human and a conservative default in the loop: triage suggests, a clinician or a red-flag rule decides, and the system always errs toward more care, never less.