Visit-Summary Review-Gate & Plain-Language Checklist

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

GATE 1 · DRAFT, NEVER FILE (the review gate)

- [ ] Model output is labeled 'AI draft — unverified' until a clinician signs
- [ ] Clinician sees the draft next to the source and can edit any fact
- [ ] Nothing posts to the record or the portal without a human signature
- [ ] The signature, not the model, makes the document the record
- [ ] Draft, edits, and sign-off are logged for an audit trail

GATE 2 · PLAIN LANGUAGE (the patient summary)

- [ ] Patient text written at ~6th-grade reading level, short sentences (AHRQ/CDC)
- [ ] A reading-level check runs on patient-facing text before release
- [ ] Medical shorthand translated to everyday words (no 'SOB on exertion')
- [ ] Summary offered in the patient's language; no raw MT of safety text (1557)
- [ ] Clinician note and patient summary built as two documents, not one

GATE 3 · PHI BOUNDARY (the model is a business associate)

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

GATE 4 · THE RECORD + RED FLAGS

- [ ] Signed summary treated as the designated record set (45 CFR 164.501)
- [ ] Patient access (164.524) and amendment (164.526) paths exist and are tested
- [ ] Design as if the patient reads the note (info blocking, 45 CFR Part 171)
- [ ] RED FLAG: auto-filing an unreviewed summary straight to the chart
- [ ] RED FLAG: a factually perfect summary written at a clinical reading level

THE ONE-LINE RULE

After a telemedicine visit, AI can draft two documents — a clinical summary for the provider and a plain-language after-visit summary the patient reads at home — and they have different readers, reading levels, and failure modes. The one rule that keeps both safe is the review gate: the model drafts, a clinician reads, edits, and signs, and nothing is filed to the record or shown to the patient until a human owns it, because summarization models still invent and omit facts at measurable rates. The patient summary owes two extra duties — a roughly sixth-grade reading level and the patient's language — and since the 2021 information-blocking rule, patients can usually read the clinical note too, so design as if they do. Every word of the visit is Protected Health Information, so the model is a business associate that needs a signed BAA with a no-training clause before it sees a transcript; a tool that recaps the visit is administrative support, while one that issues a new decision drifts toward an FDA-regulated device.