RPM Anomaly-Detection Safety & Compliance Checklist

Run a remote-patient-monitoring anomaly-detection feature through the four gates before you ship it. Engineering guidance, not legal advice — confirm specifics with counsel.

GATE 1 · THE CLAIM LINE (wellness vs device)

- [ ] Alert describes wellness/a trend, not a named disease (FDA General Wellness)
- [ ] No directive 'do X now' to the patient unless built as a regulated device
- [ ] Clinician-facing trends are support; patient-facing alerts are a device by default
- [ ] Time-critical 'deterioration' alarms are treated as device software (FDA CDS)
- [ ] Intended use and claims, not the algorithm, decide device status (SaMD)

GATE 2 · ALERT FATIGUE (suppress the noise)

- [ ] Personalized baseline per patient, not one fixed threshold for all
- [ ] Trend detection catches slow drift; a noise filter drops obvious non-events
- [ ] Alerts are ranked and routed to a clinician who triages, not auto-fired
- [ ] Alert volume is sized to what the care team can actually act on
- [ ] False-alarm rate is measured and tracked, not assumed

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

- [ ] Signed BAA with the model provider before any reading is sent (164.502(e))
- [ ] BAA covers the exact enterprise service/tier — not 'the vendor offers one'
- [ ] No-training clause: patient readings never train the vendor's future models
- [ ] Free consumer AI endpoints blocked for any identifiable reading
- [ ] Analytics outside the clinical loop run on de-identified data (164.514(b))

GATE 4 · BIAS + THE RECORD + RED FLAGS

- [ ] Alert performance tested across patient subgroups (skin tone, age, body type)
- [ ] Known device bias handled (e.g. pulse oximetry reads high on darker skin)
- [ ] Bias risks identified and mitigated per Section 1557 (45 CFR 92.210)
- [ ] RED FLAG: a patient-facing auto-alert that names or predicts a disease
- [ ] RED FLAG: fixed-threshold flood that buries the one true alert in false ones

THE ONE-LINE RULE

Remote patient monitoring streams home readings to a care team, and AI anomaly detection finds the ones that matter — but the same alert is a harmless wellness nudge or a regulated medical device depending on what it claims and who it speaks to. Keep the claim on the wellness or clinician-support side of the FDA line: a trend a clinician reviews is support, while a patient-facing, time-critical, or disease-naming alert is a device. The bigger day-to-day problem is alert fatigue — monitoring alarms are false 80–99% of the time in studied settings — so suppress noise with personalized baselines, trend detection, and a clinician triage step instead of adding sensitivity. Every reading is Protected Health Information, so the model is a business associate that needs a signed BAA with a no-training clause before it sees data, and every anomaly model must be tested for bias across patient groups because it is only as fair as its data and its devices.