

Vendor questionnaire

Fifteen questions to ask any vendor whose software influences what a clinician orders. Published by ARKA Health, who answer all fifteen below.

Columns: Question | Your vendor's answer (blank) | ARKA's published answer | Where ARKA's answer is enforced

Q1. What proportion of your recommendations are supported by high-certainty evidence, and where is that published?

Your vendor's answer: [blank]

ARKA's published answer: Yes. Every rating carries a GRADE certainty and a COI class; the full distribution is published in the certainty census (including the thin end). Artefact: /evidence/methodology#how-certain.

Where ARKA's answer is enforced: CHECK-CERT-1

Q2. How many of your clinical appropriateness anchors are authored by a payer, and where is that count published?

Your vendor's answer: [blank]

ARKA's published answer: Zero. Payer-authored criteria never influence a clinical appropriateness verdict; the firewall count on that path is published. Artefact: /data-flow#appropriateness-firewall.

Where ARKA's answer is enforced: no-payer-criteria-in-appropriateness

Q3. Does your tool ever automatically deny an order or set a utilisation target, and where is that prohibition published?

Your vendor's answer: [blank]

ARKA's published answer: No. ARKA never automatically denies and never sets a utilisation target. Artefact: /docs/ai-boundary.

Where ARKA's answer is enforced: never-auto-deny

Q4. Do you fix the measurement period in writing before the intervention is switched on, and where is that analysis plan published?

Your vendor's answer: [blank]

ARKA's published answer: Yes. The pre-period is locked in writing and the analysis plan is published before the intervention is switched on. Artefact: /baseline.

Where ARKA's answer is enforced: CHECK-I5-LOCK

Q5. Do you publish corrections to your own public claims, dated with the commit that fixed them, and where?

Your vendor's answer: [blank]

ARKA's published answer: Yes. Corrections are dated, public, and tied to the fixing commit. Artefact: /evidence/methodology#corrections.

Where ARKA's answer is enforced: CHECK-CORR-2

Q6. Can anything a clinician writes into your tool be used for discipline, credentialing, or pay, and where is that prohibition published?

Your vendor's answer: [blank]

ARKA's published answer: No. Justification text is never used for discipline, credentialing, or pay. Artefact: /security#what-we-dont-hold.

Where ARKA's answer is enforced: no-justification-for-discipline

Q7. Can identical inputs regenerate identical outputs, and do language models write clinical or patient-facing prose in your product - where is that published?

Your vendor's answer: [blank]

ARKA's published answer: Yes - identical input produces byte-identical output; no language model writes clinical or patient-facing prose. Artefact: /evidence/methodology#determinism-asset.

Where ARKA's answer is enforced: CHECK-SDM-1

Q8. Have you published research artefacts under an open licence with a resolvable identifier before you have customers, and where?

Your vendor's answer: [blank]

ARKA's published answer: Yes. Research artefacts ship under an open licence with a resolvable identifier before revenue. Artefact: <https://doi.org/10.5281/zenodo.22051812>.

Where ARKA's answer is enforced: CHECK-BENCH-1

Q9. How often does your tool interrupt a clinician, against what published ceiling, and where is that published?

Your vendor's answer: [blank]

ARKA's published answer: Yes. Interrupt rate, acceptance decomposition, and sentiment are published against published fire-rate ceilings (Alert Burden Report). Artefact: /api/aj/burden-report.

Where ARKA's answer is enforced: CHECK-BURD-1

Q10. Who is on your evidence panel, what are their conflicts, and where are their decisions published?

Your vendor's answer: [blank]

ARKA's published answer: Yes. The independent evidence panel roster, conflicts, and decisions are public. Artefact: /governance/evidence-panel.

Where ARKA's answer is enforced: CHECK-CHARTER-PUBLISH

Q11. Do you monitor and publish your own calibration and drift, including when it gets worse, and where?

Your vendor's answer: [blank]

ARKA's published answer: Yes. Calibration and drift are monitored and published, including when they worsen. Artefact: /monitoring.

Where ARKA's answer is enforced: CHECK-CAL-1

Q12. What proportion of your recommendations does your system decline to make, and do you publish it?

Your vendor's answer: [blank]

ARKA's published answer: Yes. The volume-weighted abstention rate among the most-ordered scenarios is published in the GRADE census, with a dated independent-panel burn-down. Artefact: /evidence/methodology#how-certain.