Transparency Commitments for Software That Influences Clinical Ordering, v1.2 - published by ARKA Health

Version 1.2.0 Published 23 August 2026

When you evaluate software that influences clinical ordering, you cannot verify whether appropriateness ratings carry published certainty distributions, whether payer criteria sit on the clinical path, or whether a vendor's performance number separates intervention from secular trend. These fifteen transparency commitments name what must be published and checked - so any buyer can ask any vendor the same questions.

Every published result in imaging utilisation is uncontrolled. So nobody can tell your finance team how much of a vendor's number was the intervention. There are no FHIR scopes, because there is no integration.

Each clause names the artefact that proves it and the check that enforces it.

Clause 1 (live since 2026-08-16)
You publish every rating with a GRADE certainty and a conflict-of-interest class, and you publish the whole distribution - including the part that makes you look bad.
Why it matters: A vendor who will not publish how thin their evidence is is asking you to trust a black box.
Artefact: /evidence/methodology#how-certain
Enforced by: CHECK-CERT-1
Anchor: #standard-clause-1

Clause 2 (live since 2026-08-16)
You never let a payer-authored criterion influence a clinical appropriateness verdict - the count of payer-authored anchors on that path is zero.
Why it matters: If the clinical score is the payer's criteria with a white coat on, you bought utilisation management, not appropriateness.
Artefact: /data-flow#appropriateness-firewall
Enforced by: no-payer-criteria-in-appropriateness
Anchor: #standard-clause-2

Clause 3 (live since 2026-08-09)
You never automatically deny anything, and you never set a utilisation target.
Why it matters: An algorithm that can deny a scan or ratchet a target is a utilisation-pressure tool, not clinical decision support.
Artefact: /docs/ai-boundary
Enforced by: never-auto-deny
Anchor: #standard-clause-3

Clause 4 (live since 2026-08-09)
You fix the measurement period in writing before the intervention is switched on, and you publish the analysis plan first.
Why it matters: Without a locked pre-period, any later claim of effect is storytelling - your finance team cannot verify it.
Artefact: /baseline
Enforced by: CHECK-I5-LOCK
Anchor: #standard-clause-4

Clause 5 (live since 2026-08-16)
You publish your own corrections, dated, with the commit that fixed them.
Why it matters: A vendor that never corrects public claims is either perfect or opaque - and neither is credible.
Artefact: /evidence/methodology#corrections
Enforced by: CHECK-CORR-2
Anchor: #standard-clause-5

Clause 6 (live since 2026-08-09)
You never use anything a clinician writes for discipline, credentialing, or pay.
Why it matters: If free-text justification can reach HR or a bonus formula, clinicians will stop writing - and the mechanism dies.
Artefact: /security#what-we-dont-hold
Enforced by: no-justification-for-discipline
Anchor: #standard-clause-6

Clause 7 (live since 2026-08-21)

Clause 8 (live since 2026-08-18)
You publish research artefacts under an open licence with a resolvable identifier before you have customers.
Why it matters: A dataset that lives only in a private repo is not citable evidence - it is marketing copy with a spreadsheet attached.
Artefact: https://doi.org/10.5281/zenodo.22051812
Enforced by: CHECK-BENCH-1
Anchor: #standard-clause-8

Clause 9 (committed - target 2026-09-30)
You publish how often you interrupt clinicians, how often they accept, and what they said about it - including when the numbers are bad.
Why it matters: Every physician leader has been burned by alert fatigue; a published interrupt rate against a published ceiling is the scorecard they can hand a competitor.
Artefact: /api/aj/burden-report
Enforced by: CHECK-BURD-1
Anchor: #standard-clause-9

Clause 10 (committed - target 2026-11-15)
Ratings you are not confident about are reviewed by an independent multi-disciplinary panel constituted to the IOM standard, whose roster, conflicts, and decisions are public.
Why it matters: Who decides what is appropriate - and whether they are conflicted - is the question a medical director asks before they trust a rating.
Artefact: /governance/evidence-panel
Enforced by: CHECK-CHARTER-PUBLISH
Anchor: #standard-clause-10

Clause 11 (committed - target 2026-09-30)
You monitor your own calibration and drift and publish the result, including when it gets worse.
Why it matters: A model that only publishes when it looks good is not monitoring - it is advertising.
Artefact: /monitoring
Enforced by: CHECK-CAL-1
Anchor: #standard-clause-11

Clause 12 (live since 2026-08-23)
You publish how often you decline to give a verdict at all, weighted by how often those scenarios are actually ordered, and the dated schedule on which your independent panel is reducing it.
Why it matters: A vendor who will not say how often they refuse to answer is selling confidence they do not have.
Artefact: /evidence/methodology#how-certain
Enforced by: CHECK-ABST-1
Anchor: #standard-clause-12

Clause 13 (committed - target 2026-09-30)
You never let the protocol engine see the payment rate. No reimbursement amount, payer identity, site-of-service payment differential or contract rate is an input to protocol candidate generation, gate evaluation, or the autonomy tier - and a test fails the build if one becomes reachable from that code path.
Why it matters: Every other party in the imaging chain has a reason to want the payment rate visible at protocol selection. A written, enforced refusal is the answer a compliance officer can forward.
Artefact: docs/protocol-payment-firewall.json
Enforced by: CHECK-PROTO-2
Anchor: #standard-clause-13

Clause 14 (committed - target 2026-10-31)
You never protocol a study without a human in the loop unless a named licensed physician at that site has signed the specific indication class, at a recorded version and date, and any auto-protocoled study can be reversed by a radiologist at any point before acquisition.
Why it matters: Who is responsible when it is wrong stays legible: a named physician drew the boundary; the engine stayed inside it; here is the reversal rate.
Artefact: lib/siteconfig/attestation.ts
Enforced by: CHECK-AUT-1
Anchor: #standard-clause-14

Clause 15 (committed - target 2026-09-30)
You never let the protocol engine read an image. Its inputs are the chart, the report text and the device record - and a reachability test fails the build if a pixel path becomes callable from the protocol module.

Published 2026-08-23