

ARKA

ARKA - Implementation Dossier

=====

v1.0.0 21 August 2026

IF YOU DISLIKE IT, YOU TURN IT OFF

If you dislike it, you turn off one hook. There is no data migration, no schema change, and no residue.

ARKA attaches as a CDS Hooks service registration and a SMART on FHIR app record. Full deactivation is deregistering those entries in EHR configuration - one analyst, no interface-engine rebuild, no order-set rollback. Scope it to one department or one hook, or turn the whole thing off.

After you turn the hook off, the justification text stays in the chart because it is a clinical record entry, and that is correct. ARKA keeps only the receipt - hash, length, timestamp - and that receipt is unused once the hook is gone.

IF ARKA IS UNREACHABLE AT ORDER TIME

If ARKA is unreachable at order time, the order proceeds.

The EHR does exactly what it did before ARKA existed: no card renders, no order is delayed or blocked. ARKA cannot place, cancel, modify, or hold orders. A rules-only fallback also covers ML-service outages so guideline-anchored cards continue without the optional model.

NO FHIR SCOPES FOR THE FIRST PRODUCT

Volume VIII ?9.2 ranks third among all your differentiators: "For the first product there are no FHIR scopes, because there is no integration."

There are no FHIR scopes, because there is no integration.

Modules 1 and 2 need none of these scopes. The land motion is a claims extract.

WHICH COMMITTEES - AND WHICH THE LAND MOTION SKIPS

Buyers know their own committees. We do not publish a single blended go-live number. Each row is one gate: what it decides, a duration range we can source (or unknown), the artefact ARKA hands it, and whether Modules 1 and 2 - claims extract, no EHR - require it at all. The column that says not required is the argument.

Columns: M1 / M2 / M3 / M4 / M5

Security review [1-3 weeks] M1=Required; M2=Required; M3=Required; M4=Required; M5=Required

Whether ARKA's trust boundary, architecture, and control posture are acceptable for the data that will actually move - a batch claims extract for Modules 1 and 2; a live CDS Hooks / SMART connection for Modules 3 and 4.

Artefact: Security & compliance package

Privacy / HIPAA [1-3 weeks] M1=Required; M2=Required; M3=Required; M4=Required; M5=Required

Whether the BAA, field list, hashing boundary, retention, deletion, and subprocessors are acceptable before any member-level extract or live FHIR session leaves the institution.

Artefact: Data-flow & PHI statement

Clinical informatics governance [unknown] M1=Not required; M2=Not required; M3=Required; M4=Required; M5=Not required

EHR app review, FHIR scope negotiation, CDS Hooks registration, SMART launch, and interface-engine change control. Modules 1 and 2 never attach to the EHR, so this seat does not vote on the land motion. Module 5's gold-card export is a file, not a hook. For Module 3, this seat also decides whether to enable the considered-and-deferred note write-back: a named physician leader must approve enablement in writing (site, date, scope, Clause 10.5) before the feature may turn on - recorded in docs/governance/WRITBACK_ENABLEMENT_LOG.md.

Artefact: AI governance packet

Value analysis [unknown] M1=Required; M2=Required; M3=Required; M4=Required; M5=Required
Whether the spend, evidence posture, and intended use are acceptable for a departmental or enterprise purchase. This is a procurement seat, not an EHR-app board - it still meets for a claims-only deal.

Artefact: Value Analysis Committee packet

Contracting / legal [2-6 weeks] M1=Required; M2=Required; M3=Required; M4=Required; M5=Required

MSA, BAA, baseline-lock addendum, exit language, and the written commitments that justification text is never used for discipline, credentialing, or compensation.

Artefact: Vendor portal

IRB [unknown] M1=Not required; M2=Not required; M3=If a study; M4=If a study; M5=If a study
Human-subjects determination when a study is involved - QI / quality-assessment, expedited minimal-risk, or full board. The commercial land of Modules 1 and 2 (claims extract, no study) does not require this seat.

Artefact: Evidence methodology (study status)

Data-governance committee [1-2 weeks] M1=Required; M2=Required; M3=Not required; M4=Not required; M5=Required

Who may see whose rate. Small-cell suppression at $n < 11$, no ranking at any transparency level, clinician vs. medical-director surfaces, retention, and the dispute path. This seat owns the risk thesis for Modules 1, 2, and the Module 5 gold-card export.

Artefact: What we will never do

WHO MAY SEE WHOSE RATE

For a risk-bearing buyer this section matters more than the FHIR scopes. Each statement is generated from the never-do catalog or a firewall artefact and names the check that proves it. A claim with a linter behind it is a different kind of statement from one without.

Attributed clinician MAY SEE: monthly peer packets comparing a clinician to top performers (not the mean, not a rank). The ledger query returns that clinician's own risk- and reliability-adjusted rows plus organisation aggregates (filterCellsForCaller in lib/tcoc/ledger/query.ts).

MAY NOT SEE: No leaderboard, no rank, no bottom-N. No personally attributed dollar figure in a clinician-facing artefact. Peer names, other clinicians' case lists, and ranks are out of the response.

Check: CHECK-I3-RANK; CHECK-I4-DOLLARS

Medical director / quality lead MAY SEE: The organisation distribution and organisation-level dollar model. filterCellsForCaller returns the full tenant set, unsorted by rate.

MAY NOT SEE: No leaderboard, no rank, no bottom-N. An ordered leaderboard with names is forbidden on every surface, including this one.

Check: CHECK-I3-RANK

Who may see whose rate is the ledger role split in filterCellsForCaller: a clinician sees their own risk- and reliability-adjusted rate; a medical director sees the distribution, never sorted by rate. No leaderboard, no rank, no bottom-N. No personally attributed dollar figure in a clinician-facing artefact. Enforced by CHECK-I3-RANK and CHECK-I4-DOLLARS.

[who-may-see] CHECK-I3-RANK (linter) never-do=no-leaderboard
firewall=docs/arka-tcoc/firewall.json

The default transparency level is blinded-distribution - The distribution is shown without names; each clinician sees only their own position on their packet. Ranking is not available at any of aggregate-only, blinded-distribution, named-to-leadership, named-to-group. No leaderboard, no rank, no bottom-N. Enforced by CHECK-I3-RANK and CHECK-GOAL-3. [default-transparency-no-rank] CHECK-I3-RANK (linter) never-do=no-leaderboard firewall=docs/arka-tcoc/firewall.json

Locked baselines, ledger aggregates, and delivery audit rows are retained for the contract term plus 6 years (or longer if law requires), aligned with security-evidence retention (45 C.F.R. ? 164.316). baseline locks that are append-only and hash-chained. Aggregates and lock rows remain as long as the analysis must remain replayable. Small cells ($n < 11$) are suppressed and the suppression is stated. Lock immutability is enforced by CHECK-I5-LOCK; the 6-year floor is the contractual term in docs/arka-tcoc/BASELINE_ADDENDUM.md clause 7, and CHECK-DOSSIER-2 fails if that file and this sentence drift. [retention] CHECK-I5-LOCK (linter) never-do=none firewall=docs/arka-tcoc/firewall.json

Justification text is never used for discipline, credentialing, or compensation. justification text and its hash cannot be joined to peer, ledger, stats, group, or incentive aggregates. Enforced by CHECK-I10-JUSTIF in scripts/lint-tcoc.ts; CHECK-J5-JOIN / CHECK-J5-BRIDGE in scripts/lint-aj-firewall.ts; aj_ schema forbid. [justification-not-for-discipline] CHECK-I10-JUSTIF (linter) never-do=no-justification-for-discipline firewall=docs/arka-tcoc/firewall.json

ARKA measures which clinician's ordering behaviour changed against a locked pre-period, and exports that measurement. ARKA never uses, alters or influences the assignment of beneficiaries to clinicians, and never derives a contribution figure from a field that determines attribution or risk score. The two are computed from disjoint inputs and the build proves it. deriving a contribution figure from a beneficiary-attribution or risk-score field (contribution-attribution-is-not-beneficiary-attribution). ARKA does not compute, influence, or recommend changes to beneficiary attribution or to risk-adjustment coding. Attribution and risk scores are read-only inputs supplied by your organisation. We measure avoided utilisation against the attribution and risk model you give us, and we publish the specification of how we use them. Enforced by CHECK-INC-1 and CHECK-INC-2 in scripts/lint-tcoc.ts; the disjoint-input proof in docs/attribution-firewall.json; Clause 8 in docs/arka-aj/DEPLOYMENT_AGREEMENT_CLAUSES.md. [contribution-carve-out] CHECK-INC-1 (linter) never-do=contribution-attribution-is-not-beneficiary-attribution firewall=docs/attribution-firewall.json

ARKA can auto-approve; it can never auto-deny. adverse determination (deny / modify / downgrade) only with a verified licensed-physician reviewer identity. adverse determinations require a human licensed-physician reviewer identity with a verified NPI; automation may only auto-approve or route to clinical review; ARKA-IP may recommend a second review and may never cancel, convert, hold or set the status of an order. Enforced by evaluateReviewerAuthorization in lib/ins/reviewer-queue.ts; lint:never-deny (CHECK-ND-1?CHECK-ND-8); lib/arka-ip/never-deny.ts; IRE decision-surface test. [never-auto-deny] CHECK-ND-1 (linter) never-do=never-auto-deny firewall=docs/never-deny-firewall.json

A payer-authored or utilisation-management source may inform a coverage determination and may never inform clinical appropriateness. The two field sets are disjoint and the build proves it. a payer-authored or utilisation-management source informing clinical appropriateness (no-payer-criteria-in-appropriateness). No payer-authored or utilisation-management source may inform a clinical appropriateness rating. Verdict A (score.ts and the knowledge matrix) and Verdict B (lib/coverage) are computed from disjoint field sets in docs/appropriateness-firewall.json; the clinical engine cannot reach

lib/coverage, lib/pa, lib/davinci or lib/ins even transitively; no ModalityRating carries a payer_um / payer_utilization anchor.. Enforced by CHECK-APPR-1 in scripts/lint-tcoc.ts; CHECK-FW-1..3 in lint:sources / lint:coverage / lint:scope; the disjoint-input proof in docs/appropriateness-firewall.json; Clause 9 in docs/arka-aj/DEPLOYMENT_AGREEMENT_CLAUSES.md. [appropriateness-coverage-firewall] CHECK-APPR-1 (linter) never-do=no-payer-criteria-in-appropriateness firewall=docs/appropriateness-firewall.json

No ARKA surface, artefact, metric, export, contract clause or piece of marketing may reference evaluation-and-management coding, billing level, code selection or reimbursement in connection with the considered-and-deferred note - and the build fails if one does. The one-page write-back brief at /api/compliance/writeback-brief (also on /compliance/regulatory-standing#writeback-brief) states what the note is, who authors it, what ARKA never sees, counts or is paid for, the five rules with the check that enforces each, the Clause 10 mirror of each rule, and the enablement record. Every claim in that brief is traced under CHECK-WBB-1. [writeback-brief] CHECK-WBB-1 (linter) never-do=no-billing-language-in-clinical-writeback firewall=none

WHAT LEAVES THE INSTITUTION - AND WHAT DOES NOT

Generated from the live flow graph, not drawn. Identifiable data crosses only the CDS Hooks session and the Module 3 chart write; federated outcomes and calibration aggregate without moving identifiable data.

Outcomes and calibration aggregate without moving identifiable data. The coordinator receives masked partials or noisy counts - never row-level lake rows.

CDS card [allowed] [not identifiable]

Appropriateness card with FDA Non-Device CDS disclosure. Not a chart extract, not pixels.

Hashed claims extract [allowed] [not identifiable]

Flat claims extract for Modules 1 and 2. Member identifiers SHA-256 hashed with a per-tenant salt at the ingest boundary and never stored raw. No FHIR scopes.

Aggregates only [allowed] [not identifiable]

Outcomes and calibration aggregates. Row-level lake rows never leave the institution (lib/federated/, docs/FEDERATED_PRIVACY.md).

Structured indication [allowed] [not identifiable]

Structured indication the clinical engine produced. Coverage may read it. This is the 27.4 allowed arrow.

Payer criteria [forbidden] [not identifiable]

Payer criteria. The clinical engine cannot read coverage criteria. This arrow is forbidden - the 27.4 firewall.

Justification text (stays) [stays] [identifiable]

Module 3 justification text written to the chart under the clinician's token. It stays in the institution. It is a clinical record entry.

{hash, length, timestamp} [allowed] [not identifiable]

Receipt only - hash, character length, timestamp. Not the justification text. Not identifiable narrative.

Structured FHIR (session) [allowed] [identifiable]

CDS Hooks prefetch - structured FHIR in the EHR's BA session at order time. Identifiers hashed before any persistence.

The same generated SVG is embedded in this PDF (stream /EmbeddedFile).

EVERY FHIR SCOPE, JUSTIFIED

Volume VIII ?9.2 ranks third among all your differentiators: "For the first product there are no FHIR scopes, because there is no integration."

Modules 1 and 2 need none of these scopes. The land motion is a claims extract.

AllergyIntolerance patient/AllergyIntolerance.rs M1/M2=Not required

Contrast and other allergies support premedication or modality-change context in the chart.

Degrade - contrast and other allergy context is omitted and the miss is recorded.

Condition patient/Condition.rs M1/M2=Not required

Active conditions define the differential the order is working - core structured indication material.

Degrade - completeness is recorded as degraded and the missing class is written to the degradation array; scoring continues on remaining structured context.

Device patient/Device.rs M1/M2=Not required

Implanted devices and MRI conditionality matter to indication completeness and modality choice context.

Degrade - implanted-device and MRI-conditionality context is omitted and the miss is recorded.

DiagnosticReport patient/DiagnosticReport.rs M1/M2=Not required

Pathology reports (known malignancy) change staging-protocol framing for the indication.

Prior radiology reports change protocol context (e.g. delayed-phase follow-up) for the reconstructed indication.

Degrade - prior radiology and pathology reports are omitted from Phase B; the miss is recorded.

DocumentReference patient/DocumentReference.rs M1/M2=Not required

Clinical notes are the richest narrative source for Phase B (mechanism, red flags, functional impact).

Degrade - Phase B narrative notes are skipped; reconstruction stays structured_only and the miss is recorded.

Encounter patient/Encounter.rs M1/M2=Not required

ED vs inpatient vs ambulatory changes the decision frame for the reconstructed indication.

Degrade - setting (ED / inpatient / ambulatory) is omitted from the reconstructed indication and the miss is recorded.

ImagingStudy patient/ImagingStudy.rs M1/M2=Not required

Recent imaging-study metadata (identifiers and timestamps, not pixels) lets the engine see prior exams in the episode without a PACS pull.

Degrade - prior-study metadata is an empty bundle; scoring continues without that context and the miss is logged, not silent.

MedicationRequest patient/MedicationRequest.rs M1/M2=Not required

Active meds (anticoagulation, metformin, nephrotoxics) enrich mechanism and prior-treatment context.

Degrade - active-medication context (anticoagulation, metformin, nephrotoxics) is omitted and the miss is recorded.

Observation patient/Observation.rs M1/M2=Not required

Recent labs (WBC, lactate, CRP, troponin, D-dimer, LFTs, lipase, INR, hCG, eGFR) change

study territory and urgency framing. Vitals (HR, BP, SpO₂, temp, RR, FiO₂) ground acuity and exam-finding elements for Phase A.

Degrade - labs and vitals are omitted from the indication; the miss is recorded. The engine does not invent values.

Patient patient/Patient.rs M1/M2=Not required

Age and administrative sex gate reconstruction and every card; the Patient resource is the session context token, not a chart dump.

Abstain - reconstruction does not start and the EHR receives no card. The order still proceeds.

Procedure patient/Procedure.rs M1/M2=Not required

Recent procedures and surgeries frame implant/post-op status relevant to the indication.

Degrade - recent procedure and implant/post-op context is omitted and the miss is recorded.

ServiceRequest patient/ServiceRequest.rs M1/M2=Not required

Prior completed imaging orders are the duplicate and redundancy window; the in-flight draft order arrives from hook context, not this prefetch key.

Degrade - prior completed ServiceRequests become an empty bundle; the in-flight order still comes from CDS Hooks context. The miss is recorded.

The write-back path is one narrow, named case: Module 3 accountable justification, posting the clinician's words to the chart as a ServiceRequest annotation under the clinician's SMART token ('user/ServiceRequest.u', with 'user/ServiceRequest.write' as the v1 fallback). That is not a prefetch scope, and it is not requested for Modules 1 and 2.

SECURITY POSTURE, PENETRATION TEST, AND SUBPROCESSORS

HIPAA - Privacy, Security & Breach Notification Rules (Program in force). Full privacy and security policy suite adopted and operating; business associate agreement ready for execution; annual NIST 800-30 risk analysis complete.

SOC 2 - Security, Availability, Confidentiality (In progress). Readiness assessment complete against the 2017 Trust Services Criteria (2022 points of focus). Type I examination targeted December 2026; Type II report mid-2027.

HITRUST e1 - Essentials, 1-Year Validated (Roadmap active). 44-requirement e1 tier selected as the appropriate entry certification; MyCSF self-assessment Q1 2027, validated assessment H1 2027. ~80-90% control overlap with the SOC 2 program lets one evidence body serve both.

2025 HIPAA Security Rule NPRM (Adopted as baseline). The December 2024 OCR proposal's heightened specifications are ARKA's internal baseline today: mandatory MFA, universal encryption, asset inventory and network map, six-month vulnerability scans, annual penetration testing, segmentation, 72-hour restore capability.

ACCESS CONTROL

- MFA on all production, code, and data systems; FIDO2 hardware keys for admins
- Role-based least privilege from a documented access matrix; quarterly reviews
- 24-hour deprovisioning; break-glass access alarmed, logged, and reviewed

ARKA processes structured FHIR resources in the EHR's context at the moment of order as your Business Associate - transiently in the CDS Hooks / SMART request path, not bulk-exported or warehoused as raw charts. ARKA does NOT ingest image pixels. Only de-identified order features and scoring outcomes persist (SHA-256-hashed order and patient keys, age buckets, ICD-10/CPT codes, AIIE scores, and scrubbed factor metadata), redacted CDS decision-log fingerprints, federated aggregate query audit metadata, and standard security event logs -

encrypted in transit and at rest, with append-only audit logging. Model improvement uses federated analytics on de-identified or aggregated signals - row-level patient records never leave the institution in federated queries.

Penetration test: first independent test scheduled Q4 2026; engagement letter shareable under NDA; executive summary shared with customers post-remediation.
Customers are notified at least 30 days before ARKA adds a PHI-touching subprocessor - aligned with BAA subcontractor flow-down and customary objection rights.

Vercel, Inc. - Edge / application hosting (Next.js, CDS Hooks & API routes) - United States (primary: iad1) - Required before production PHI
Structured FHIR prefetch and CDS context processed transiently in the request path - no PHI persisted at the edge

Supabase, Inc. - PostgreSQL data tier (tenant-isolated databases, row-level security) - United States - Required before production PHI
De-identified audit rows, hashed identifiers, AIIE scores, security event logs - encrypted at rest

Render Services, Inc. - Optional ML inference service (when ML_SERVICE_URL is configured) - United States - Required before production PHI
Structured order features for AIIE scoring in the CDS request path - rule-based fallback when unset or unreachable

ANALYST HOURS

The integrated modules (CDS Hooks / SMART on FHIR) cost 40 analyst hours.
The claims-only baseline motion costs 0 analyst hours.
A "free" pilot that consumes 200 hours of a customer's analyst time is not free.

REFERENCE TIMELINE

Phase 0 - Paperwork (Week 0)

ARKA: Execute MSA and BAA, share the 21-document security package from our Trust Center, provision sandbox credentials, and seed your security questionnaire responses from the published controls on /security.

You: Procurement and security reviewers to execute agreements, complete the vendor questionnaire, and approve sandbox access for the integration team.

Phase 1 - Connect (Weeks 1-2)

ARKA: Provide CDS Hooks discovery URL, SMART on FHIR launch parameters, and registration checklists. ARKA consumes structured FHIR order-select context from the EHR prefetch bundle - no image pixels, no bulk PHI export.

You: Your EHR/integration team registers ARKA as a CDS Hooks service and SMART on FHIR app in non-production - designed for Epic, Oracle Health (Cerner), and athenahealth after customer-specific validation. Typical lift: one analyst, light ARKA support.

Phase 2 - Configure (Weeks 2-4)

ARKA: Map service lines and payer mix, narrow pilot scope to one department and one payer, tune auto-clear thresholds toward 35-40%, and align card copy with our CDS Card Language Style Guide (supportive, non-coercive, evidence-first).

You: Clinical and revenue-cycle sponsors to confirm pilot boundaries, approve threshold targets, and sign off on in-flow card wording before shadow mode.

Phase 3 - Validate (Weeks 4-6)

ARKA: Run shadow mode alongside live traffic, publish validation metrics and drill-down

rows, and complete conformance checks against CDS Hooks and Da Vinci CRD expectations.
You: Clinical champion reviews shadow cards against local practice patterns, logs sign-off,
and confirms the pilot KPI baseline before production enablement.

Phase 4 - Go live (Weeks 6-8)

ARKA: Enable ARKA in production for the agreed pilot scope, monitor API latency and card delivery, and report ROI and utilization metrics through the observability dashboards.

You: Change-control approval for production CDS registration, a named on-call contact for the first two weeks, weekly readouts with your champion and rev-cycle lead, and confirmation that the rollback runbook is delivered and walked through before production enablement.

CITE THIS DOCUMENT

ARKA - Implementation Dossier (v1.0.0, 21 August 2026). /api/implementation/dossier.
Generated; no customer name.

Ungated. No form. Zero customer data. No customer name.

- end of dossier -