

DIPA Vendor Capability Guide

Describe. Interpret. Predict. Act.

Evaluation and RFI Guide

Version 3.0 · September 2026 · Working draft for review and comment

Matthew Ayinde, MBA, BSN, RN

Emergency Nurse · Clinical Documentation Integrity · AI Workflow and Prompt Architecture

Published openly. No license or fee. Attribution requested for derivative or published work. The name and mark DIPA: Describe. Interpret. Predict. Act. are retained by the author. Comments and corrections: dipa@reveiqai.com

DIPA Vendor Capability Guide

Describe. Interpret. Predict. Act.

Evaluation questions for EHR, clinical AI, CDI, and documentation partners

Version 3.0 | September 2026 Aligned to DIPA Framework v3.0

About this version

Version 3.0 makes three structural changes.

Predict is scoped to clinical trajectory. Version 1.0 mixed clinical prediction with denial risk, underpayment, and coding rework forecasting under a single heading. Those are different constructs with different data, owners, and users. Section 4 now covers clinical trajectory only.

Financial and patient-access capabilities move to an appendix. They remain worth evaluating. They are no longer presented as DIPA capabilities, because they are not.

Two sections are new and carry disqualifying weight. Section 6 covers trigger monitoring and closed-loop escalation. Section 7 covers regulatory status and data architecture. A vendor that fails the disqualifying criteria in §12 should not advance regardless of how it scores elsewhere.

Purpose

This guide helps healthcare organizations evaluate whether a technology partner can support DIPA, a clinical reasoning framework intended to make the clinical story explicit while it can still change the outcome, and to keep documented commitments active until they are met.

DIPA is not a request for an automated diagnosis engine, a deterioration detection system, a coding checklist, or a revenue-maximization tool.

The goal: make clinical reasoning visible, ensure documented plans are executed, and preserve human judgment for the decisions that matter.

Use as: vendor demonstration script, RFI appendix, proof-of-concept evaluation tool, multidisciplinary due-diligence checklist, or pilot-design companion.

The governing distinction

Before any question below, establish this:

DIPA technology tracks execution of a clinician-authored plan. It does not detect deterioration.

A vendor that does not understand this distinction, or that markets its product on the other side of it, is not a DIPA partner regardless of product quality. The distinction determines clinical trust, medico-legal posture, and regulatory status simultaneously.

Capability model

Element	Required capability	Illustrative use case
Describe	Aggregate and organize patient-specific clinical evidence across the EHR	Show trends in oxygen requirement, vitals, labs, imaging, medications, procedures, and change from baseline
Interpret	Surface clinically relevant evidence and possible documentation-story gaps for clinician review	Identify evidence that may warrant clarification of acuity, etiology, linkage, severity, or clinical significance where supported by clinician judgment
Predict	Support clinician articulation of expected trajectory, deterioration signals, monitoring needs, and escalation criteria	Surface trends and candidate risks for clinician acceptance or rejection; never author the prediction
Act	Monitor clinician-authored triggers, notify on unclosed commitments, and route work within defined governance rules	Notify when a documented escalation criterion is met and no response is documented; escalate on non-response

Vendor response instructions

For each question, vendors should:

1. Provide a yes / no / partial response
2. Describe current production capability, not roadmap
3. Identify whether the capability is native, partner-delivered, custom-built, or client-configured
4. Demonstrate the workflow using patient-level de-identified example evidence
5. State limitations, required data sources, integration dependencies, and implementation effort
6. Identify the role of human review and any confidence thresholds used

7. Provide evidence of performance in comparable settings, including setting, scale, EHR, use case, and measured outcome

1. Clinical evidence and data foundation

Evaluation question	Response	Evidence
Can the solution ingest structured and unstructured clinical data, including notes, laboratory results, imaging reports, medication administration, flowsheets, orders, diagnoses, procedures, and problem lists?		
Can it distinguish current encounter data from historical information and identify change from baseline?		
Can it reconcile conflicting or duplicated information across notes, problem lists, orders, and other sources?		
Can it identify stale, missing, delayed, or incomplete data before generating a prompt or recommendation?		
Does it provide source-level traceability for every finding, including author, date and time, note section, laboratory result, imaging report, medication, or flowsheet value?		
Can users see the evidence without navigating to a separate portal or conducting a second search?		
What EHR data standards and integration methods are supported, FHIR, HL7, SMART on FHIR, APIs, CCD/C-CDA, direct feeds?		
Which EHR platforms have live production integrations? Specify Epic, Oracle Health, Meditech, athenahealth, eClinicalWorks, or others.		
Can the organization configure role-specific views for clinician, nursing, CDI, coding, case management, physician advisor, and utilization review?		

2. Describe: clinical situation awareness

Evaluation question	Response	Evidence
Can the solution summarize relevant current clinical facts without omitting material contradictory evidence?		
Can it display time trends for vital signs, oxygen support, laboratory values, medication escalation, interventions, and treatment response?		
Can it identify a clinically meaningful change from baseline, new oxygen requirement, acute kidney function change, mental status change, functional decline?		
Can it identify relevant data not yet available, such as pending cultures, imaging, consult recommendations, or incomplete medication reconciliation?		
Can clinicians control which data are presented and suppress low-value or duplicative information?		
Describe alert-suppression and user-preference capabilities. How does the solution avoid overwhelming clinicians with generic summarization?		

3. Interpret: clinical reasoning support

Evaluation question	Response	Evidence
Does the solution distinguish between presenting evidence and establishing a diagnosis?		
Can it identify evidence that may warrant clinician assessment of etiology, acuity, severity, temporal relationship, complication status, or linkage between conditions?		
Can it support clinically compliant, non-leading prompts rather than suggesting a diagnosis based on reimbursement opportunity?		
Can the prompt explicitly state that the clinician should document only what is clinically supported and independently assessed?		
Can CDI, coding, physician-advisor, quality, and compliance teams review and govern prompt logic before deployment?		
Can the solution account for conflicting documentation, alternative explanations, and uncertainty?		
Does it support condition- or specialty-specific logic while allowing organization-specific governance and evidence thresholds?		
How does the vendor detect and mitigate hallucination, unsupported inference, omitted contradictory evidence, and bias in generated content?		

4. Predict: clinical trajectory

Scope note. This section addresses clinical trajectory only. Denial risk, coding rework, authorization exposure, and financial forecasting are evaluated in Appendix A and are not DIPA capabilities.

4.1 The two-lever structure

DIPA requires that clinical intuition govern and that AI predictive support inform rather than author. These questions establish whether the vendor’s architecture respects that hierarchy.

Evaluation question	Response	Evidence
Does the solution author predictions, or propose candidate risks for clinician acceptance, modification, or rejection? Describe the exact interaction.		
Can a prediction or escalation trigger enter the record without affirmative clinician action?		
Are rejected candidate risks retained anywhere in the record or used to generate downstream activity?		
Is every Predict statement attributable to a named clinician in the record?		
Does the solution function if the predictive capability is disabled entirely?		
Can the organization configure which clinician roles may author, modify, or retire predictions and triggers?		

4.2 Trajectory support

Evaluation question	Response	Evidence
Can the solution surface patient-specific trends relevant to expected clinical course, deterioration, or treatment failure?		
Can it distinguish patient-specific trend observation from population-derived risk scoring, and does it present them differently?		
Can it propose escalation criteria as observable conditions with thresholds and time bounds, for clinician review?		
Can it identify anticipated discharge barriers for clinician consideration?		
Does the product report sensitivity, specificity, positive predictive value, false-positive rate, false-negative rate, and calibration for each predictive use case?		
How frequently are models retrained or recalibrated, and how is drift monitored?		
Can the organization review predictive performance by race, ethnicity, language, age, disability, payer, and geography where data are available and appropriate?		

5. Act: workflow orchestration and human escalation

Evaluation question	Response	Evidence
Can the solution route work to the appropriate role based on case type, confidence, clinical complexity, or time sensitivity?		
Can it distinguish low-risk tasks appropriate for automation from high-risk tasks requiring explicit human review?		
Can it create or prepopulate a task, work item, documentation prompt, or escalation within the EHR or approved workflow system?		
Can it assemble evidence packets for CDI review, coding review, utilization review, or physician-advisor escalation?		
Can it track task ownership, response time, completion, override, escalation, and final resolution?		
Can it support closed-loop communication across clinicians, nursing, CDI, coding, case management, physician advisors, and utilization review?		
Can it prevent duplicate work when multiple teams review the same patient?		
Can it prioritize work by clinical urgency, patient impact, compliance risk, and time remaining before discharge?		

6. Trigger monitoring and closed-loop escalation

This section is new in v3.0 and carries disqualifying weight. See §12.

6.1 Architecture and constraint

Evaluation question	Response	Evidence
Does the solution monitor clinician-authored criteria, or does it independently assess patient state? Describe precisely.		
Can the solution generate a notification from criteria no clinician authored? If yes, describe when and why.		
What data does the monitoring capability consume? Specify whether it includes continuous physiologic waveforms or signal-acquisition output.		
Does every notification display the clinician's own trigger language alongside the data that met it?		
Can the solution characterize a patient's clinical state, assign a risk score, or make a deterioration assessment? If yes, is that function separable from trigger monitoring?		

6.2 Trigger lifecycle

Evaluation question	Response	Evidence
How are triggers authored? Can they be expressed as observable conditions with thresholds and time bounds?		
Can triggers be modified, suppressed, or retired by a clinician at any time?		
Do triggers expire? Can expiry be tied to a reassessment cycle, and are unreaaffirmed triggers retired automatically?		
What close paths are supported? Confirm that documented reassessment with the plan unchanged closes a trigger exactly as intervention does.		
How many interactions are required to close a trigger by documented reassessment?		
Can the organization review and report close-path distribution?		

6.3 Routing and notification burden

Evaluation question	Response	Evidence
Can notification be tiered, with escalation on non-response rather than on trigger?		
Can the organization configure recipients, intervals, and escalation thresholds by unit, role, shift, and condition?		
Can notification volume be capped or throttled per clinician per shift?		
What reporting exists on notification volume, close rate, time to close, and tier escalation frequency?		
Can the solution measure trigger specificity over time; that is, detect clinicians writing vaguer criteria to avoid notifications?		
What evidence does the vendor have on alert burden and workaround behavior in production deployments?		

7. Regulatory status and data architecture

This section is new in v3.0 and carries disqualifying weight. See §12.

7.1 Regulatory

Evaluation question	Response	Evidence
Has the vendor determined whether the monitoring capability is a medical device under FD&C Act §520(o)(1)(E)? Provide the written analysis.		
If non-device, address each of the four criteria explicitly, including the independent-review criterion and any time-critical considerations under the January 2026 FDA guidance.		
If the capability is a device, what is its regulatory status, cleared, exempt, under enforcement discretion, in submission?		
Has the vendor obtained outside regulatory counsel opinion? Will it be shared under appropriate terms?		
How does the vendor handle a product change that alters regulatory status? Describe change control.		
Does the vendor market this capability as deterioration detection, early warning, or patient risk prediction in any materials?		

7.2 Data layer separation

Evaluation question	Response	Evidence
Can the solution architecturally separate clinical record content from aggregate quality-improvement analysis? Describe the mechanism, not the policy.		
Can the client determine what monitoring data enters the medical record and what does not?		
Can aggregate analysis: close rates, response times, unit patterns, override analysis, be segregated into a defined patient safety evaluation system?		
Can that segregated layer be exported to a listed Patient Safety Organization in a usable format?		
Can retention and disposition policy be configured separately for each layer?		
Does the vendor retain, access, or analyze either layer for its own purposes? Under what terms?		
Is client data used to train shared models? Describe opt-in/opt-out, de-identification, contractual terms, and governance.		

8. CDI, coding, and query governance

Evaluation question	Response	Evidence
Does the solution support concurrent CDI prioritization rather than only retrospective review?		
Can it identify cases appropriate for clinical validation, physician-advisor review, coding review, or compliant query consideration?		
Can users configure policies to avoid duplicate or conflicting queries?		
Can it support query templates approved by CDI, HIM, compliance, and physician leadership?		
Does the system preserve the distinction among clinical evidence, clinician diagnosis, coded data, and reimbursement impact?		
Can professional coders audit, correct, or reject automation outputs, particularly in complex or high-risk cases?		
Does it measure query quality, clinician response, coding rework, final DRG agreement, and clinical-validation outcomes?		

9. Human oversight, safety, and compliance

Evaluation question	Response	Evidence
Which recommendations can be automated, which require human approval, and which are never automated?		
Can the client configure automation thresholds by workflow, role, setting, or risk class?		
Does every recommendation and notification retain an audit trail showing source evidence, model and version, trigger, prompt, user action, close path, override, and outcome?		
Are audit logs exportable for internal compliance, quality, coding, and legal review?		
How are privacy, HIPAA, minimum-necessary access, role-based access, data retention, and data-use restrictions managed?		
How does the vendor address clinical-safety incident reporting, error remediation, model rollback, and change control?		
Can the organization test prompts, triggers, and models in a sandbox before production release?		
Can the vendor support independent validation, ongoing audit, and client-defined quality thresholds?		

10. Implementation and operating model

Evaluation question	Response	Evidence
What data, interfaces, configuration, policies, staffing, and governance are required before go-live?		
Can the solution support a phased deployment, documentation framework first, monitoring capability second?		
What is the typical timeline from contracting to first production use?		
What client roles are required during implementation, and how many hours per week?		
What workflow redesign is required versus optional?		
How are local facility variations accommodated without losing systemwide definitions and comparability?		
What training, adoption support, change management, and ongoing optimization are included?		
What support exists specifically for trigger-writing education?		
What support model, uptime commitment, incident response, and service levels are offered?		

11. Outcomes, measurement, and references

Evaluation question	Response	Evidence
Provide outcomes from comparable organizations, including baseline, intervention, measurement period, denominator, and methodology.		
Separate outcomes for clinical and documentation quality, operational efficiency, compliance, and user experience.		
Report false-positive and false-negative performance, not only gross financial return.		
Can outcomes be stratified by facility, service line, payer, specialty, patient population, and user role?		
Can the solution demonstrate reduction in avoidable rework rather than merely faster downstream processing?		
For monitoring deployments: report close rate, time to close, notification volume per clinician, and any observed trigger-quality degradation.		
Provide at least two reference clients using a similar EHR, workflow, size, or use case.		
What metrics does the vendor recommend for a phased DIPA pilot?		

12. Disqualifying criteria

A vendor should not advance if any of the following is true, regardless of scores elsewhere.

- 1. The solution authors predictions or escalation criteria that enter the record without affirmative clinician action.**
- 2. The solution generates clinical notifications from criteria no clinician authored**, outside a separately governed early warning function the client has independently selected.
- 3. The monitoring capability consumes continuous physiologic waveforms or signal-acquisition output** without a corresponding device regulatory position.
- 4. The vendor has no written regulatory analysis** of the monitoring capability's device status.
- 5. The solution cannot separate clinical record content from aggregate analysis architecturally.**
- 6. Documented reassessment is not supported as a close path:** or requires materially more effort than intervention.
- 7. The vendor markets the capability as deterioration detection or patient risk prediction** while presenting it as a DIPA capability.

Items 1, 2, and 7 are the same failure in three forms: the AI displacing clinical judgment rather than supporting it. Item 6 appears minor and is not, a system without an easy non-action close path forces clinicians to either intervene or ignore, and they will ignore.

13. Required demonstration scenarios

Using de-identified data, a realistic sandbox, or structured simulation.

Scenario 1: Clinical reasoning support

A patient with pneumonia, hypotension, elevated lactate, increasing oxygen requirement, and acute kidney injury is admitted.

Demonstrate how the solution: - Presents relevant clinical evidence with source traceability - Distinguishes evidence from diagnosis - Supports clinician assessment of etiology, acuity, severity, trajectory, and monitoring needs - Proposes escalation criteria **for clinician acceptance**, showing the accept, modify, and reject paths - Avoids leading or reimbursement-oriented prompts - Routes to CDI or physician-advisor review where appropriate - Records the clinician's decision with audit trail

Scenario 2: Trigger monitoring, full lifecycle

Using the triggers authored in Scenario 1, demonstrate: - Trigger authoring and the record entry it produces - A trigger firing on discrete data, with the clinician's own language displayed - Tier-1 notification to nursing, then escalation on non-response - Close by intervention - **Close by documented reassessment with the plan unchanged**: count the interactions required - Trigger revision and trigger retirement - Automatic expiry at the reassessment cycle - What enters the medical record versus the aggregate layer, shown explicitly

Scenario 3: Adverse conditions

Demonstrate behavior when: - A trigger fires repeatedly on a patient whose clinician has judged the plan appropriate - Multiple triggers fire simultaneously across several patients on one nurse's assignment - A trigger fires on data later corrected or invalidated - A clinician writes a vague, unmonitored trigger, what does the system do, and can the organization see it?

Scenario 4: Pre-final-coding documentation review

A complex medical admission with documentation ambiguity affecting severity and clinical validity.

Demonstrate how the solution: - Identifies the case before final coding - Aggregates clinical evidence relevant to the ambiguity - Prioritizes by clinical and compliance risk - Routes to the appropriate reviewer - Supports a compliant, non-leading query - Tracks resolution

14. Scoring

Score each dimension 1 to 5.

Dimension	1	3	5
Clinical relevance	Generic or retrospective outputs	Useful but limited clinical context	Evidence-based, patient-specific, clinician-usable insight
Clinician authorship integrity	AI authors; clinician reviews	Mixed authorship, unclear boundaries	Clinician authors; AI proposes only, with clear accept/reject/modify
Workflow fit	Separate portal, manual handoffs	Partial integration	Embedded workflow, closed-loop routing, minimal duplicate work
Safety and governance	Limited transparency and auditing	Standard controls	Source traceability, configurable oversight, layer separation, robust audit, client control
Regulatory and legal posture	No analysis performed	Analysis exists, gaps remain	Documented position, counsel-reviewed, architecture supports it
Operational impact	Vague claims	Demonstrated point solution	Measurable reduction in avoidable rework
Scalability	Site-by-site customization	Some repeatable configuration	Enterprise standards with controlled local flexibility

Weighting

Dimension	Weight
Clinical relevance	20%
Clinician authorship integrity	20%
Safety and governance	20%
Regulatory and legal posture	15%
Workflow fit	15%
Operational impact	7%
Scalability	3%

Clinician authorship integrity is weighted equally with clinical relevance because it is the constraint on which the medico-legal and regulatory positions both rest. A vendor may be excellent at everything else and still be unusable if it authors clinical judgment.

A solution should not advance on financial return. It must meet minimum standards for clinical usability, clinician authorship, evidence traceability, human oversight, regulatory posture, and measurable reduction of avoidable rework.

Appendix A: Adjacent capabilities outside the DIPA framework

The capabilities below are legitimate and often valuable. **They are not DIPA capabilities**, and a vendor's performance here should be scored separately from the DIPA evaluation above. They are retained because organizations frequently evaluate these alongside documentation capability, and separating them prevents a strong revenue-cycle product from scoring well on a clinical framework it does not serve.

A.1 Payer, authorization, and revenue integrity

- Can the solution identify authorization requirements before service delivery or an inpatient procedure?
- Can it surface payer-specific coverage criteria, documentation requirements, and authorization status in workflow?
- Does it support payer API, FHIR, clearinghouse, portal, or EDI methods? Specify production capability.
- Can it identify missed carve-out payment opportunities, implants, drugs, technology add-ons, high-cost supplies, contract-specific provisions?
- Can it compare expected reimbursement with remittance to identify contractual underpayment?
- Can it classify denials by root cause and identify upstream prevention opportunities?
- Can it prioritize appeals using expected recovery value, payer behavior, filing deadlines, clinical support, and likelihood of overturn?
- Can it report denial trends by payer, facility, service line, procedure, diagnosis, clinician, authorization status, and root cause?
- Can it estimate likelihood of denial, underpayment, coding rework, or delayed payment at the encounter level?

A.2 Patient access and financial experience

- Can the solution support eligibility verification, coverage discovery, benefit interpretation, network validation, and estimate generation before service?
- Can it identify patients who may benefit from financial counseling, coverage screening, charity-care review, or payment-plan outreach?
- Can it automate routine reminders, digital intake, payment links, and standard communications while routing sensitive or complex cases to trained staff?
- Can it document and respect patient communication preferences, language needs, accessibility needs, and consent requirements?
- How does the vendor ensure financial predictions do not result in discriminatory, coercive, or inequitable patient treatment?

A.3 Demonstration scenarios for adjacent capabilities

Denial prevention and payer intelligence. A payer begins denying a defined procedure and diagnosis combination across facilities. Demonstrate detection, exposure quantification, root-cause separation, identification of in-process encounters at risk, preventive workflow recommendation, and appeal or contract-management support.

Patient financial navigation. A scheduled patient has incomplete coverage information, significant estimated responsibility, and hardship indicators. Demonstrate coverage verification or discovery, understandable estimate generation, routing to a person for counseling or assistance screening, and protection of privacy, equity, and patient choice.

Closing

The right DIPA partner will not generate more documentation, create another queue, or automate a downstream task. It will help clinicians make their reasoning explicit, keep the commitments in that reasoning active, and return judgment to the people whose judgment it is.

A vendor that cannot articulate the difference between tracking a clinician's plan and assessing a clinician's patient has not understood the framework, whatever else it does well.

Version 3.0 | September 2026 | Working draft for review and comment