SmarterDx

WHITE PAPER

The *clinical AI* RFP guide:

18 questions health system RCM leaders should ask before choosing a partner



Executive summary

Payers are adopting AI faster than health systems, and the impact is showing up in the denials queue. Health systems are racing to catch up, but most RFP processes weren't built to tell an impactful clinical AI partner from a good sales deck.

This guide gives leaders a shared framework for vetting AI partners within revenue cycle management: **the questions that separate genuine clinical intelligence from a repackaged model, and a way to score partners apples-to-apples**.

Why this guide, why now

Payers are outpacing health systems on AI, often employing layers of AI tools. Health systems feel the downstream impact every time a preventable denial lands in the queue.

So what's stopping health systems from catching up? According to research from SmarterDx and The Health Management Academy, health system leaders point to increasing denial complexity (70%), pressure to improve cash flow (68%), and payer AI adoption outpacing their own (58%).

The right AI platform can close that gap with greater review coverage, lower cost to collect, better documentation, and revenue that doesn't require adding headcount to find it. But many health systems are evaluating clinical AI partners for the first time, mid-RFP, without any shared internal playbook. That's how you end up choosing the best ***demo*** instead of the best ***partner***. And in a market where nearly every vendor now has an "AI story," that distinction matters more than ever.

The framework: Six categories of questions to ask every potential partner

Every question below came from **real health-system RFPs**, weighted by how often health systems actually asked it. Nothing here is hypothetical; these are the questions health systems are actually putting in front of vendors right now.

QUESTIONS 1-3
Platform

QUESTIONS 4-6
Data

QUESTIONS 7-9
Modeling

QUESTIONS 10-12
ROI and proven metrics

QUESTIONS 13-15
Safety, security, and compliance

QUESTIONS 16-18
Implementation and support

© Smarter Dx, Inc. 2026 All Rights Reserved

Platform

The best clinical AI doesn't just plug into your revenue cycle: It plugs into your EHR, your workflows, and your existing CDI and coding teams without adding friction, at whatever chart volume you're actually running. If a partner can't clearly explain how their platform fits inside the systems you already run, that's not a detail to sort out later. It's the whole ballgame.

KEY POINT

Of the 12 real health-system RFPs we reviewed, **EHR integration** and **CDI workflow fit** were the two most universal themes – both came up in all 12.

QUESTION 1

What is your experience integrating secondary CDI workflows into the EHR, including query workflows, reporting, and access controls?

Why it matters: *A strong answer shows the product working inside your existing operating model, not becoming one more disconnected queue your team has to check separately. A weak answer means hidden interface work you didn't budget for, duplicate processes nobody asked for, and CDI and coding teams who adopt slower because the tool doesn't actually fit how they work.*

QUESTION 2

How does the solution complement existing CDI and coding workflows?

Why it matters: *This is where you find out if the tool actually fits your process, or just sits next to it. A strong answer shows the technology working alongside your CDI and coding teams — not creating a second workflow they now have to babysit. A weak answer means unclear ownership, more review work than you had before, and opportunities that get flagged but never actually acted on.*

QUESTION 3

What does the data flow actually look like? What do you provide, and what do we get back?

Why it matters: *A strong answer lays out the interfaces, system boundaries, latency, writeback, and who owns what — before you're locked into implementation, not after. A weak answer hides custom feeds, manual handoffs, or outputs that never actually make it back into your EHR workflow.*

KEY TAKEAWAY

Good clinical AI should make your existing experts more effective, not force your organization to redesign itself around the technology. It should work the same way across every workflow it touches, not just the one it was piloted on.

© Smarter Dx, Inc. 2026 All Rights Reserved

Data

Every vendor will tell you their AI reviews “the full patient record.” The real question is whether it actually ingests enough of the chart to find what’s missing, and whether you can trust what it hands back.

QUESTION 4

What elements from the patient record do you ingest?

Why it matters: *A strong answer tells you whether the model actually sees enough of the chart to find something real, and flags which feeds will drive your implementation cost before you’re surprised by them. A weak answer leaves you guessing at completeness and latency, and hoping the performance numbers don’t secretly depend on data you can’t even provide.*

Get the details: Notes alone are not the full clinical picture. Labs, vitals, medications, orders, imaging, procedures, codes, and longitudinal context can all change the meaning of what appears in a note. And the question isn’t simply how much data a model ingests. It’s whether the model can connect it into a coherent clinical story.

QUESTION 5

Does the solution display the clinical evidence supporting each recommendation, with an auditable link back to the source record?

Why it matters: *A strong answer gives your clinicians and auditors a straight line from a recommendation back to the actual chart facts — something they can defend, not just accept. A weak answer hands you conclusions with nothing underneath them, which means more review time and more compliance risk.*

QUESTION 6

What’s your policy on data ownership, retention, deletion, and return upon contract termination?

Why it matters: *A strong answer means you actually keep control of your data — deletion that reaches backups, derived data, and subprocessors, not just the primary database. A weak answer means indefinite retention, data you can’t fully get back, or a model still learning from your patients long after the contract has ended.*

KEY TAKEAWAY

If a partner can’t tell you what they’re ingesting, where a finding came from, or what happens to your data when the contract ends, you don’t actually know what you agreed to.

© Smarter Dx, Inc. 2026 All Rights Reserved

Modeling

"It's AI" is not a methodology. Push past the buzzword and make them show you what's actually happening under the hood, because the difference between a rules engine, a fine-tuned model, and a general-purpose LLM that's never actually seen a real patient chart matters more than any demo will show you.

QUESTION 7

Can you describe your technology platform and AI/NLP methodology?

Why it matters: *A strong answer tells you which tasks run on rules, which run on statistical models, which run on language models, and why each one was picked for that job. A weak answer flattens the whole thing into "AI" — and if you can't tell what's actually making the decisions, you can't evaluate performance, risk, or change control.*

The goal isn't to reward the pitch with the longest list of model types. It's to understand whether the architecture was actually designed for the clinical problem you're asking it to solve.

QUESTION 8

Do you provide explainability or transparency into how AI-generated outputs are created?

Why it matters: *A strong answer shows the actual evidence, logic, confidence level, and model version behind a finding...enough for one of your own clinicians to push back on it if they disagree. A weak answer is "trust the model," and that's not something you can put in front of an auditor.*

QUESTION 9

What's your precision, recall, and score benchmark strategy?

Why it matters: *A good answer states the evaluation population, denominators, confidence intervals, and the tradeoff between missed findings and false positives. A weak answer offers an acceptance rate or a single headline metric that doesn't measure model performance end to end.*

This is also where sophisticated buyers should ask whether the model is simply reproducing what their own teams already find. For example, agreement with historical coding isn't the same thing as clinical accuracy, especially if the goal is to find opportunities the existing process misses. The stronger signal is whether the performance keeps climbing as more charts run through the model, not just where it started.

KEY TAKEAWAY

A partner who can't explain their own model in plain language probably can't defend it in an audit, either.

© Smarter Dx, Inc. 2026 All Rights Reserved

ROI and proven metrics

The ROI-methodology question came up in all 12 of the 12 health-system RFPs we reviewed — tied for the most universal question in this entire guide. Every AI partner has an ROI slide. Fewer can walk you through how they got there, and fewer still will put a number in writing they're willing to be held to.

QUESTION 10

Can you explain how your firm's solution calculates and tracks value added or ROI?

Why it matters: *A strong answer defines the unit of value, attribution logic, baseline, timing, and reconciliation to realized financial outcomes. A weak answer substitutes gross coded impact for cash, ignores reversals and denials, or counts opportunities you would have captured anyway.*

QUESTION 11

Can you provide references and case studies from comparable health systems, including documented financial and operational outcomes?

Why it matters: *A strong answer gives you comparable scale, EHR, workflow, baseline maturity, measurement period, and permission to validate the results directly. A weak answer offers logos or aggregate claims that can't be tied to a customer, a denominator, or a realized outcome.*

If there are no health systems willing to sing the praises of the AI partner in a case study and/or a private referral, it's a red flag... proceed with caution.

QUESTION 12

Can you provide calculation definitions, data sources, baselines, targets, and assumptions for each proposed KPI?

Why it matters: *A good answer makes the business case reproducible by your finance team and prevents changing definitions after launch. A weak answer presents percentages with no denominators, or benchmarks with no comparable population.*

KEY TAKEAWAY

If a partner can't show their ROI math, they don't have an ROI number – they have a marketing number.

© Smarter Dx, Inc. 2026 All Rights Reserved

Safety, security, and compliance

Your health system already treats patient data like the responsibility it is. A partner touching that data doesn't get a pass just because they're the AI company — hold them to the same bar, not a lower one.

QUESTION 13

How do you ensure AI features comply with HIPAA and other applicable privacy and regulatory requirements?

Why it matters: *A strong answer connects the legal requirements to actual controls — how subprocessors are managed, how data flows, what's in the contract. A weak answer waves a certification or a signed BAA at you and calls it a day, without ever explaining how PHI actually moves through their product.*

QUESTION 14

Can you describe user authentication, role-based access, single sign-on, and access-audit capabilities?

Why it matters: *A strong answer means every user only sees what their job actually requires — backed by real identity management, fast offboarding when someone leaves, and an audit trail on every access to PHI. A weak answer runs on shared logins and broad roles, which means more administrative overhead for you and more ways for the wrong person to see the wrong chart.*

QUESTION 15

What safeguards are implemented to prevent errors in diagnosis recommendations or clinical interpretations?

Why it matters: *A strong answer layers technical controls, clinical review, confidence thresholds, monitoring, and a real incident response plan — so mistakes get caught before they matter. A weak answer is betting on your own team to catch every error after the system already influenced the chart or the code.*

KEY TAKEAWAY

Security isn't a checkbox item to close out at the end of the RFP – it's something you can see in how a partner handles your data every single day, not just on the day they sign the contract.

© Smarter Dx, Inc. 2026 All Rights Reserved

Implementation and support

Implementation timeline came up in all 12 of the 12 RFPs we reviewed — health systems know from experience that this is where good tools go to die. A brilliant model that takes a year to stand up isn't solving the problem you have today. The best partners can tell you exactly what go-live looks like, how many weeks it takes to get there, and exactly what it takes to keep it working after that.

Implementation is also where "platform" claims meet reality. If every new workflow requires another custom integration, consulting engagement, or manual queue, you're just buying complexity. The strongest partners can point to a defined implementation process, run the same way client to client, not one that gets reinvented every time.

QUESTION 16

Can you describe your onboarding approach, implementation timeline, key milestones, dependencies, and factors that can delay go-live?

Why it matters: *A strong answer lays out the critical path and tells you plainly what's on the partner and what's on you — your decisions, your data, your interfaces, your security review. A weak answer is a target date with nothing behind it: no assumptions, no evidence it's ever actually happened on that schedule before.*

QUESTION 17

What resources and level of effort are required during implementation and ongoing operations?

Why it matters: *A strong answer names the actual roles, hours, and decision rights — across IT, CDI, coding, compliance, finance, and whoever's signing off at the top. A weak answer undersells what your team will really be doing, and you find out the true number after the unplanned work has already landed on someone's desk.*

QUESTION 18

What training and resources are provided for each stakeholder group, including clinicians, operational teams, IT staff, and leadership?

Why it matters: *A strong answer is built for each role, timed around go-live, and reinforced after launch — with a real way to tell if it worked. A weak answer is one generic webinar and a hope that watching it once changes how people actually work.*

KEY TAKEAWAY

The best AI partner treats go-live as the beginning of the relationship, not the finish line of the sale.

© Smarter Dx, Inc. 2026 All Rights Reserved

PARTNER KPIS

KPIs and data inputs to request

Make every potential partner put real numbers behind these — their own client data, not an industry benchmark they found somewhere:

- **Realized net new revenue** per 10,000 discharges
- **Time to first measurable ROI**
- **Percentage of charts** actually reviewed
- **Denial prevention** and overturn rate
- **False-positive rate**
- **Number of health systems** live in production
- **Client expansion:** how many clients came back for more after year one

WATCH OUT

Common pitfalls in partner selection

- **Don't score the demo, score the data behind it.** A slick walkthrough tells you nothing about what the model was trained on.
- **Don't mistake agreement rate for accuracy;** high agreement with your own coding can just mean the model learned your blind spots and kept them.
- **Don't skip the reference call;** if a partner can't produce one happy client willing to talk, ask yourself why.
- **Don't leave security for the end;** retrofitting a compliance conversation after you've already picked a favorite is a much more complicated conversation than having it up front.
- **Don't walk into this without a shared scorecard;** without one, the loudest voice in the room wins, not the best partner.

ONE MORE

Don't evaluate AI by how autonomous it sounds. In clinical revenue cycle work, removing humans from the process is not automatically sophistication. The stronger question is whether the technology removes manual work while putting better evidence in front of the people whose judgment still matters.

© Smarter Dx, Inc. 2026 All Rights Reserved

SCORECARD

Scoring partners consistently

Once you've asked every partner the same questions, score them on the same scale. For each of the six categories, rate them 1 to 5 based on how they answered the three questions in that section:

1. **Couldn't answer any of the three questions with real specifics**
3. **Answered some questions with real detail, others vague or incomplete**
5. **Answered all three with concrete detail, evidence, and comparable references**

Score Platform, Data, Modeling, ROI, Safety/Security, and Implementation, then add up all six for a total out of 30. Partners clustering toward the top of the scale are worth serious consideration. Partners clustering toward the bottom have real gaps worth pushing on before you move forward.

CATEGORY	1	2	3	4	5	SCORE
Platform	☐	☐	☐	☐	☐	_____ / 5
Data	☐	☐	☐	☐	☐	_____ / 5
Modeling	☐	☐	☐	☐	☐	_____ / 5
ROI and proven metrics	☐	☐	☐	☐	☐	_____ / 5
Safety, security, and compliance	☐	☐	☐	☐	☐	_____ / 5
Implementation and support	☐	☐	☐	☐	☐	_____ / 5
Total						_____ / 30

This gives you a fast, in-the-room read on where they stand as you work through the RFP.

© Smarter Dx, Inc. 2026 All Rights Reserved

FINAL WORD

Before you sign anything

The right questions do more than help you pick a clinical AI partner. They get finance, IT, and clinical leadership aligned before anyone signs anything. The score isn't meant to make the decision for you. It makes the tradeoffs visible and makes it much harder for a polished demo to outrun weak evidence.

Clinical AI is moving too quickly for health systems to evaluate it on brand recognition, feature lists, or demo quality alone. The partners that matter will be able to show what their technology sees, how it reaches a recommendation, where the evidence came from, how humans stay in control, and exactly how the financial impact is measured.

The goal isn't to buy the best AI story. It's to choose the clinical intelligence, evidence, and economics that still hold up when the demo is over and the real work begins – what do the team members executing the workflow actually think of it? Does it meaningfully help them, or simply check an AI box for someone above them? For success, any tool has to help the people doing the work.



LEARN MORE

Learn more about clinical AI built for revenue cycle performance: **smarterdx.com**

© Smarter Dx, Inc. 2026 All Rights Reserved

Bonus *checklist*

All 18 questions, by category

PLATFORM

- [ ] Describe your experience integrating Secondary CDI workflows into the EHR, including query workflows, reporting, and access controls.
- [ ] How does the solution complement your existing CDI and coding workflows?
- [ ] What does the data flow actually look like? What do you provide, and what do we get back?

DATA

- [ ] What elements from the medical record do you ingest?
- [ ] Does the solution display the clinical evidence supporting each recommendation, with an auditable link back to the source record?
- [ ] What's your policy on data ownership, retention, deletion, and return upon contract termination?

MODELING

- [ ] Describe your technology platform and AI/NLP methodology.
- [ ] Do you provide explainability or transparency into how AI-generated outputs are created?
- [ ] What's your precision, recall, and score benchmark strategy?

ROI AND PROVEN METRICS

- [ ] How does your solution calculate, track, and report value added or ROI?
- [ ] Provide references and case studies from comparable health systems, including documented financial and operational outcomes.
- [ ] Provide calculation definitions, data sources, baselines, targets, and assumptions for each proposed KPI.

SAFETY, SECURITY, AND COMPLIANCE

- [ ] How do you ensure AI features comply with HIPAA and other applicable privacy and regulatory requirements?
- [ ] Describe user authentication, role-based access, single sign-on, and access-audit capabilities.
- [ ] What safeguards are implemented to prevent errors in diagnosis recommendations or clinical interpretations?

IMPLEMENTATION AND SUPPORT

- [ ] Describe your onboarding approach, implementation timeline, key milestones, dependencies, and factors that can delay go-live.
- [ ] What resources and level of effort are required from your team during implementation and ongoing operations?
- [ ] What training and resources are provided for each stakeholder group, including clinicians, operational teams, IT staff, and leadership?

See what clinically grounded AI looks like in practice. Learn more about the SmarterDx clinical intelligence platform for revenue cycle performance: **smarterdx.com**

© Smarter Dx, Inc. 2026 All Rights Reserved