



From Records to **Clinical Clarity**

A Buyer's Guide to Healthcare Data Interoperability
and Clinical Usefulness

How to evaluate vendors that acquire, structure, and
operationalize patient data, with key questions to ask each.

JULY 2026, Version 1

Contents

1. How to use this guide	03
2. The interoperability landscape	04
3. The core problem: connectivity is not clarity	05
4. The evaluation framework	06
5. The two data journeys	07
6. Acquisition and reach	09
7. Absorbing what comes back	10
8. From documents to structured data	11
9. From data to a decision	12
10. Trust and source traceability	13
11. Validation and evidence	14
12. Security, privacy, and compliance	15
13. Integration, scale, and commercial fit	16
14. The vendor landscape: who does what	17
15. Vendor scoring	18
16. Build your custom scorecard	19
17. Claims to pressure-test	21
18. Who should be in the room	22
19. Glossary	23
20. About xCures	24

1 How to use this guide

This guide explains how records are transacted today, the terms of participation, and where vendors fit into the interoperability journey. It walks you through a usability framework and a concrete set of questions for evaluating any vendor that acquires, structures, or operationalizes data from medical records and other healthcare data.

How it's organized:



Interoperability Landscape

How records move today and the core usability problem.



The Evaluation Framework

Establishing and evaluating the two data journeys that move raw records to answers you can use.



Eight Dimensions of Data Usability

one section each, with open-ended questions to ask every vendor.



Compare and Decide

The vendor landscape, the scorecard, claims to pressure-test, and what matters most by use case.



Reference

Glossary / Notes / Sources.

Who should read what:

Engineering / Data Science: Acquire, Absorb, Structure, Integrate.

Clinical / CMIO: Trust, Validate.

Operations / Workflow: Decide / What matters most by use case.

Commercial / Finance: Integration, the scorecard, and claims to pressure-test.

2 The interoperability landscape

Health records can now move freely across the country. This section covers how exchange works and what the rules require.

The three layers of nationwide exchange:

- **Regional and state HIEs.** Local networks connecting hospitals, labs, and clinics within a state or region. Services vary widely, from governance and record-locator matchmaking to full technical exchange. They solved local connectivity but were never built to link the whole country.
- **Carequality and TEFCA.** The governance layer. Both are rulebooks, not data acquisition pipes. They set the legal and policy terms that let independent networks and different record systems query each other, without running the technical exchange themselves. Together, they carry well over a billion document exchanges a month - and growing.
- **QHINs.** The technical layer under TEFCA specifically: record location, patient matching, and document retrieval across networks. QHINs are a TEFCA construct — Carequality runs through its own Implementers, not QHINs. There are 11 approved QHINs currently. TEFCA has operated this layer since December 2023, under one rulebook, run on the government's behalf by The Sequoia Project, the Recognized Coordinating Entity (RCE) that is designated to administer it.¹ Carequality is the older, higher-volume pathway today; TEFCA is the newer, federally established one.

The rules that make exchange work

Every request states an Exchange Purpose. Two Exchange Purposes require a response. That stated Exchange Purpose is tied to each request and is a representation to the other network participants of your legal basis for accessing records.

1. Treatment Exchange Purpose: lets healthcare providers and adjacent pull records for the purpose of treating their patients’.
2. Individual Access Services (IAS): lets a person gain access to their own medical records.

The regulatory backdrop

Interoperability is no longer optional. Sharing is legally required under the 21st Century Cures Act, which banned information blocking, but the law leaves flexibility on how.²

The document standards underneath

The nuances of how to parse information within documents depends on the document type. C-CDA is an older whole-document XML standard, so retrieving one value means parsing the whole file. FHIR is a modern API standard exposing discrete elements, but it does not guarantee that a value means the same thing across systems. USCDI defines which elements must travel. All three set a floor, not a ceiling, and most vendors build only to the floor.

¹ 11 designated QHINs under TEFCA, The Sequoia Project (RCE) – <https://rce.sequoiaproject.org/designated-qhins/>

² Information-blocking enforcement, HHS-OIG, 2025 – <https://oig.hhs.gov/documents/special-advisory-bulletins/10930/information-blocking-enforcement-alert-2025.pdf>

3 The core problem: connectivity is not clarity

Retrieving medical records is largely solved through connections with the network exchanges. However, connectivity is not clinical clarity - teams still need to gather relevant information from those records to make decisions. As the market matures, the true value is moving from simply acquiring the data to what happens after it arrives.

- **71% can access, 42% use.** In 2023, 71% of US hospitals could access outside records at the point of care, but only 42% reported their clinicians often used them.³ Relevant information is buried in thousands of unstructured documents.
- **Most of what matters is unstructured.** By widely cited estimates, 80% of clinically relevant healthcare data lives in unstructured notes, PDFs, faxes, and scanned images, not in tidy coded fields.⁴ Teams spend large amounts of time, energy, and money finding insights and turning them into usable data.
- **The data comes back messy.** One peer-reviewed analysis collected 401 clinical documents and analyzed 346, finding about five data-quality errors in each.⁵ 47% of clinicians cannot quickly find outside information and must sift through duplicates.⁶
- **Reach and matching are incomplete.** Roughly 30% of US hospitals are not engaged in all four domains of interoperable exchange, with gaps concentrated among small, rural, and independent hospitals.⁷ With no national identifier, cross-organization match rates can fall to 50 to 60% when records are shared across separate organizations.⁸
- **The burden lands on people.** Physicians spend close to 2 hours on EHR and desk work for every 1 hour of direct patient care,⁹ much of it reconciling records by hand. When data arrives unusable, that manual work is the hidden cost.

The buyer's real question is no longer can you get the data. It is what happens after the data arrives: can you turn it into something my team can act on, and can I trust it?

³ Point-of-care access and use gap, ASTP/ONC, 2024 -

<https://www.healthit.gov/data/data-briefs/interoperable-exchange-patient-health-information-among-us-hospitals-2023>

⁴ 80% of clinically relevant data is unstructured, Deloitte Insights, "Dark analytics," 2017 -

<https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2017/dark-data-analyzing-unstructured-data.html>

⁵ Data-quality errors in clinical documents, D'Amore et al., AMIA, 2018 - <https://pmc.ncbi.nlm.nih.gov/articles/PMC6371305/>

⁶ Clinicians cannot quickly find outside data; duplicate burden, KLAS, 2024 - <https://klasresearch.com/archcollaborative/report/ehr-interoperability-2024/604>

⁷ 30% of US hospitals not engaged in all four interoperability domains, ASTP/ONC Data Brief, 2024 -

<https://healthit.gov/data/data-briefs/interoperable-exchange-patient-health-information-among-us-hospitals-2023/>

⁸ Cross-organization patient match rates fall to 50-60%, Pew Charitable Trusts, 2018 -

<https://www.pew.org/en/research-and-analysis/reports/2018/10/02/enhanced-patient-matching-critical-to-achieving-full-promise-of-digital-health-records>

⁹ Physician EHR and desk-work time, Sinsky et al., Annals of Internal Medicine, 2016 - <https://www.acpjournals.org/doi/10.7326/M16-0961>

4 The evaluation framework

Before you can compare vendors, you need to know what to expect from the data itself: how comprehensive the data is, how usable it arrives, and how much work is left to your team. Score every vendor on two sequential journeys:

JOURNEY ONE

The Data Usability Journey

Can the vendor get it, clean it, structure it in a way that is usable; trust it, and deliver it at scale?



JOURNEY TWO

Where The Data Gets Used

Once the data is transformed into something usable, does it reach your team at the moment they need it



5 The two data journeys

JOURNEY ONE: Can you find data, and can you make it usable?

Data climbs a ladder towards being usable, and each rung depends on the one before it. You cannot structure data you have not acquired, and you cannot deliver a decision-ready output from data that is not yet structured. Score a vendor on how far up this ladder they climb before handing the rest of the work back to you.

Stage	What the vendor delivers	What your team still does
Acquire	Raw records retrieved, but duplicated, out of order, and in a variety of formats (CCDA, fax, images, etc.)	Everything: OCR, parse, dedupe, structure, code, extract, decide.
Absorb	Duplicates removed, every format converted to one standard.	Read the free text, code the terminology, extract the specific facts.
Structure	Facts pulled from notes and coded to standard terms.	Confirm the extracted facts and assemble them for the decision.
Decide	The exact answer the workflow needs, checked and delivered.	Make the decision at hand.

The standing requirements: is it trusted and validated, is it secure, and can it be delivered at scale? A vendor can reach the top of the usability ladder and still fail. An output you cannot trust, verify, or connect to your systems at scale is not decision-ready.

Stage	What the vendor delivers	Why it matters
Trust	Every output traced back to its source in the record, ready for your experts to verify.	Clinical facts need sources. If your team reopens the record to check it, the vendor saves them nothing.
Validate	Proof that each output matches the source record.	Accuracy needs to be audited, because errors no one caught becomes denials.
Secure	Audited, demonstrable data stewardship.	Handling PHI with the right security measures in place is essential.
Integrate	Output that fits your systems and flexes across workflows at scale.	A demo on one record isn't enough, it needs to scale across your entire patient population and exist in outputs your team can actually leverage.

Once you establish that a vendor can clear the 8 different usability metrics above, you can actually use the data.

JOURNEY TWO: Does the data reach the right team at the time of the decision?

Once the data is deemed usable, the question shifts. Where inside your specific organization does that data land? Does the data show up as the specific answer each workflow needs, and does it reach your team at the right time? One record has to serve very different people at very different moments, and a single generic summary forces every one of them to dig out their own piece.

Example: Maria, 58, Type 2 diabetes, A1c 8.9%, eGFR 52 and falling, referred for a prognostic test to gauge her risk of kidney disease progression. One longitudinal patient history, but five teams need five different things from it:

- **Order and intake** needs her full outside history the moment the requisition opens, not a week of chasing faxes.
- **Clinical review** needs only the clinical facts that prove the test is warranted, not a 60-page document to read through.
- **Prior authorization** needs the fields that the payer requires to approve the test before it runs.
- **Billing** needs the right codes, each tied to the source it came from.
- **Appeals** need the full defensible record ready the day a denial lands.

Same patient with the same records, but five different answers for five different teams are needed. You need an checklist output built for the exact question each team is asking; so the test gets authorized, the claim gets paid, or the denial gets defended.

The eight dimensions of usability

The next sections cover each stage in detail, with the questions to ask each vendor, to determine what work is left for your teams.

Stage	The question it answers	Where many vendors stop
Acquire	Can you reach my patients' records?	A single network connection
Absorb	Can you handle the mess that returns?	Raw, messy documents handed back
Structure	Can you turn text into usable data?	Using already structured feeds only
Decide	Can you deliver what a workflow needs?	A generic summary you still need to read
Trust	Can you verify every output?	Unsourced answers
Validate	Can you prove your output matches the source record?	A marketing number with no tracked evidence
Secure	Can you be a trustworthy data steward?	A HIPAA self-attestation
Integrate	Can you fit my systems at scale?	One output format, for only one workflow

6 Acquisition and reach

Strong vendors should use several data acquisition methods to gather the maximum amount of medical data they can on each patient:

- Network queries across multiple exchange frameworks (Carequality, TEFCA via QHINs, etc.)
- Including Individual Access Services via TEFCA, the individual-directed exchange purpose
- Direct connections to EHRs
- Bringing your own data (faxes, PDFs, scans, images, uploads)

The treatment trap: if you are not a Covered Entity with a treatment relationship with your patients, the treatment purpose is closed to you. Other companies may use other use cases or Individual Access services to acquire data from the networks.



Questions to ask every vendor

- ☒ Which networks do you query, and are you a QHIN or a participant? If a participant, through whom?
- ☒ What Exchange purposes do you support today? What is your plan for future exchange purposes? Do you support individual access services?
- ☒ Can I bring my own faxes, PDFs, scans, and images, processed the same way as network data?
- ☒ For a typical patient, what share of their records do you actually retrieve? And how do you know what is missing?
- ☒ How do you handle partial responses and sources that are not connected?
- ☒ What does onboarding entail, and what legal agreements will I need to sign before we can exchange?

7 Absorbing what comes back

A real query returns a pile of raw documents: duplicates, conflicting values, and many formats, which is not an actionable patient record. A capable vendor absorbs that complexity so you never see it. If a vendor hands you raw documents, your team becomes the integration layer, a real recurring cost that is easy to underestimate in a clean demo.



Questions to ask every vendor

- ☒ Do you deduplicate at both the patient and document level, and how?
- ☒ How do you reconcile conflicting values from different sources?
- ☒ Do I receive one normalized record, or many raw documents I must combine?
- ☒ Which input formats do you accept, and what exactly do you return?
- ☒ When two sources disagree, how do you represent which value you chose and why?

8 From documents to structured data

Most clinically important information sits in unstructured text and images; physician notes, by widely cited estimates, this is more than 80% of healthcare data.¹⁰ Turning that into structured, coded data is the hard work, and where vendors most often stop.

Decode “we use AI.” It means three different things: operational AI that automates steps, general-purpose NLP that parses text broadly, and domain-specific validated extraction that is checked and source-linked. Only the last produces data you can trust.



Questions to ask every vendor

- ☒ Do you extract structured data from free text, PDFs, and scanned images, or only structured feeds?
- ☒ How is your extraction built, and how do you keep it current as record types change?
- ☒ Do you map to standard terminologies such as SNOMED, LOINC, RxNorm, and ICD?
- ☒ How do you handle negation, uncertainty, and patient-versus-family attribution?
- ☒ What is your measured extraction accuracy, broken out by field? And how many fields have you measured, out of how many total?

¹⁰ 80% of clinically relevant data is unstructured, Deloitte Insights, “Dark analytics,” 2017 – <https://www.deloitte.com/us/en/insights/topics/technology-management/tech-trends/2017/dark-data-analyzing-unstructured-data.html>

9 From data to a decision

Structured data is still not an answer. The last mile is delivering the specific information a workflow needs, inside that workflow, at the moment of decision. Too much data is as useless as too little.



Questions to ask every vendor

- ☒ Can I define the exact data elements each workflow requires?
- ☒ Do you return decision-specific outputs, or a single general summary?
- ☒ Can outputs reach my team through both a user interface and an API?
- ☒ How do you avoid burying my team in irrelevant records?
- ☒ Will you show a completed example for my highest-volume workflow during evaluation?

10 Trust and source traceability

An output a clinician cannot verify is an output a clinician will not use. Every value should link to its exact source document and location. Clinical summarization shows measurable error rates, including hallucinated and omitted facts,¹¹ so vendors need to show the source for every claim and label inference clearly.



Questions to ask every vendor

- ☒ Does every output link to its exact source document and the location within it?
- ☒ Can a reviewer open the original record in one click?
- ☒ Where do you use automated inference, do you label it and show the basis for it?
- ☒ Can you verify if the test set was clinician reviewed ground truth, or self reported?
- ☒ Do you keep an audit trail of what was extracted, from where, and when?

¹¹ Clinical summarization error and hallucination rates, Asgari et al., npj Digital Medicine, 2025 – <https://www.nature.com/articles/s41746-025-01670-7>

11 Validation and evidence

Accuracy claims are easy to make and hard to verify. A 2024 analysis reviewed ten certified developers' real-world testing and found none presented their results in a meaningful way.¹² There is a difference between a vendor that extracts and one that extracts and then validates against ground truth.



Questions to ask every vendor

- ☒ Will you share your validation methodology, sample size, and ground truth definition?
- ☒ Is your extraction accuracy validated independently or externally, and can you show it?
- ☒ What is your accuracy by field and by document type, not only overall?
- ☒ How do you monitor accuracy in production over time?
- ☒ What happens, and who is accountable, when an output is wrong?

¹² Real-world testing among certified developers, Handley et al., JAMIA Open, 2025 – <https://academic.oup.com/jamiaopen/article/8/3/ooaf044/8160436>

12 Security, privacy, and compliance

Patient data carries heavy obligations. Confirm them before a contract, not after a breach. Look for an independent framework such as HITRUST rather than a HIPAA self-attestation alone, and confirm how consent and sensitive categories are handled.



Questions to ask every vendor

- ☒ Do you hold an independent certification (such as HITRUST) and at what level?
- ☒ Will information shared through TEFCA be collected (or kept) in your systems?
- ☒ Is the data you retrieve retained in your systems, and what secondary uses do you make of it within the Common Agreement — reuse, de-identify, or resell?
- ☒ Can you identify and segment sensitive records if needed (like 42 CFR Part 2 etc.) or does it all come back in one undifferentiated stream?
- ☒ Where is data stored, who can access it, and how is access logged?
- ☒ When you sign a business associate agreement, what are your breach notification terms?

13 Integration, scale, and commercial fit

The best data is worthless if it cannot reach your systems at your volume. Two traps live here. The good-enough trap: accepting seventy percent of what you need to avoid managing another vendor is a real tradeoff, so you should price the thirty percent gap that your team would absorb. And the build versus buy trap: you can build basic extraction, but not years of domain-trained, validated extraction and a quality-control process.



Questions to ask every vendor

- ☒ Do you offer both an API and a usable interface for non-technical staff?
- ☒ How do you perform at my volume, and what are your latency and uptime commitments?
- ☒ Can you verify if it is per-patient/per-query pricing vs platform/seat pricing? How does it behave as we scale?
- ☒ What is hard-to-replicate depth, and what is an API I could swap?
- ☒ Do you cover all conditions and use cases, or only a subset?
- ☒ If I leave, can I export everything in a standard format? Am I locked in by proprietary formats or re-integration cost?

14 The vendor landscape: who does what

Interoperability vendor-types cluster into categories. Each category does real work, but the work tends to stop at a point. Use the map to see what a vendor is doing vs not doing.

Dimension	Category	What it does well	Where it typically stops
Acquire	Connectivity and transport	Moves messages between systems and formats	Hands you a stream of data, not a patient record
Acquire	EHR-native exchange	Deep sharing within one vendor's network	Strongest inside its own network, weaker across others
Acquire	Network and aggregation	Queries networks and returns combined records	Returns normalized documents or FHIR, rarely captures free text
Acquire	Developer data platforms	Gives engineers FHIR data and components	Customer builds the workflows and supplies the meaning
Acquire	Real-world data and analytics	Builds de-identified datasets for research	Not for real-time, patient-level care decisions
Absorb /Structure	Enrichment and NLP	Extracts from text and adds codes (bought structuring)	A layer you must feed with records first, weak provenance
Absorb /Structure	In-house AI or custom AI builds	Internal builds on general models – no vendor involved	Can be difficult to validate and maintain over time
Absorb /Structure	Ambient documentation	Turns the current visit into a note	Does not retrieve or structure prior records
Absorb /Structure	Manual abstraction	People reading charts, with judgment	Slow, inconsistent, and hard to scale

15 Vendor scoring

Score each vendor on the eight data usability dimensions. Weight them for your use case, then rate one to five.

Worked example below: a diagnostics lab evaluating vendor A and vendor B.

A top priority for labs is medical necessity review: extracting the right clinical facts accurately enough to justify a test, and proving that accuracy if a payer questions it later. That means Structure, Decide, and Validate might matter more to this buyer than Acquire or Integrate. Vendor B seems to be a better fit for your specific needs.

Dimension	Weight (1-3)	Vendor A score (1-5)	Vendor A weighted	Vendor B score (1-5)	Vendor B weighted
Acquire	1	5	5	2	2
Absorb	1	4	4	2	2
Structure	3	2	6	5	15
Decide	3	3	9	4	12
Trust	2	2	4	5	10
Validate	3	2	6	4	12
Secure	2	4	8	2	4
Integrate	1	5	5	2	2
Unweighted total		27		26	
Weighted total			47		59

16 Build your custom Scorecard

The eight dimensions do not change, but their weights can. Determine which market vertical you fit into, and which portions of the data journey are more important than others for your specific use case.

Market Segment	Primary Use Case
Diagnostics	Test requisition completion. The retrieved records needed to support a medical-necessity determination.
Providers (in person)	Quicker access to a consolidated patient history for intake. Pre-treatment review and better documentation to support reimbursement.
TeleHealth	Complete patient history at the point of care.
Value-based care	Longitudinal records assembled across connected systems, condition documentation, and support for risk-adjustment workflows.
Payers	API and SDK depth, breadth across conditions, white-label fit, and predictable scale.
Consumer Health	A consolidated health history to drive product or engagement insights.
Specialty pharmacy	Medication reconciliation and payment. Speed to therapy, prior-authorization packages assembled from retrieved records, and reliable fax and document intake.

*Starting weights by market segment based on preliminary market research
(1–3 scale, high - low importance)

Market Segment	Acquire	Absorb	Structure	Decide	Trust	Validate	Secure	Integrate
Diagnostics	1	1	3	3	2	3	2	1
Providers (in person)	2	2	2	3	2	2	1	2
Telehealth	3	2	2	3	2	1	1	2
Value-based care	3	3	3	2	2	3	1	2
Payers	2	2	2	2	2	3	2	3
Consumer health	2	2	2	2	3	2	3	2
Specialty pharmacy	2	2	3	3	2	2	1	3

Example Scorecard Below:

Dimension	What a top score looks like	Weight	Score
Acquire	Multiple methods with measured reach		
Absorb	One normalized record from many messy inputs		
Structure	Validated extraction from text, PDFs, and images		
Decide	Configurable, workflow-specific outputs		
Trust	Every value links to its exact source		
Validate	Published or independent, per-field accuracy		
Secure	Independent certification		
Integrate	API and interface, scale, and breadth of use cases		

17 Claims to pressure-test

The claim

The Reality

"We get all the data."

More data is not the goal, completeness plus relevance is. Ask what share of your sources they reach, and what they do with what returns.

"We are the system of record."

Systems of record are built to store data for their own setting, not to reconcile it against outside sources.

"We use AI to read records."

Ask which models, how they are validated, and whether every output links to a source. General models vary between runs and state wrong facts confidently.

"We are connected to the network."

Connection is not a match, not completeness, and not clean data. Ask what actually arrives, and in what shape.

"Our accuracy is 99 percent."

Ask 99 percent of what, measured how, on how many records, by whom, and broken out by field.

"We will just add more people."

People are the most expensive and least scalable layer, and inconsistent across reviewers and days.

"It is good enough, and it is bundled."

Decide it on purpose. Accepting 70 percent to avoid managing a vendor is a real tradeoff. Price the gap your team closes by hand.

18 Who should be in the room

Evaluating a data vendor is not one person's job. Bring four perspectives, and give each its own questions.

Stakeholder	What they should press on
Engineering and data	Formats, APIs, latency, and scale; and whether this is real depth or a swappable API.
Operations and workflow	Whether output lands inside the workflow, how much manual review remains, and how the highest-volume case performs.
Clinical and CMIO	Whether experts can verify every output, whether inference is labeled, and that it stays clear of unsanctioned clinical decision support.
Commercial and finance	How pricing behaves as you scale, the cost of the work your team still absorbs, and the ROI on the gap it closes.

19 Glossary

Evaluating a data vendor is not one person's job. Bring four perspectives, and give each its own questions.

Term	Plain definition
Regional HIEs	A regional network that shares records among local hospitals, labs, and clinics.
Carequality/TEFCA	Provide governance/ policy/ frameworks rather than technical services (e.g., Record location).
QHIN	A certified backbone hub under TEFCA (11 QHINs established currently), others can connect through them. • CommonWell: original network for all non-epic EHRs and was a counter to Epic's CareEverywhere rollout (which initially only allowed all Epic organizations to share data with one another)
Exchange Purpose	The stated, attested reason for a request (Treatment, Payment, Operations).
IAS	Individual Access Services through TEFCA only. The pathway to pull records on behalf of the individuals with whom you already have a direct relationship.
Bring your own data	Records you already hold, like faxes, PDFs, and uploads, added to the mix.
FHIR	A modern API standard exposing discrete data elements as resources.
C-CDA	An older whole-document XML standard you must parse to find one value.
USCDI	The federal list of data elements that must be interoperable. A floor, not a ceiling.
Provenance	The traceable link from a data value back to its exact source.
HITRUST	An independent security certification, stronger than a HIPAA self-attestation.
Patient matching	Linking records to the right person without a national identifier.
Information blocking	Interfering with the access, exchange, or use of records. Now actively enforced.

20 About xCures

Optimize your tech stack. xCures® takes you all the way to Clinical Clarity. We don't drop you off along the way.

- 1 Find and gather. Acquires records through multiple intake pathways; via treatment-purpose network retrieval or individual access services, or by bringing your own data or EHR connections.
- 2 Combine and clean. Deduplicates and normalizes the records we retrieve into a single consolidated record.
- 3 Structure and extract. Pulls the facts into structured data: the Automated Patient History.
- 4 Verifiable, source-linked, trusted. Runs every value through our validated extractors and links it to its source: source-linked clinical data. Our extractors are validated against clinician review, with reported extraction performance of 97-99% precision and recall scores across checklists on a defined validation cohort.^{13*}
- 5 Deliver decision-ready answers. Delivers Decision-Ready Checklists straight into your workflow.



**We don't make clinical decisions.
We make them possible.**

To discuss your specific use case, book a demo at:

xcures.com

*Results are based on a retrospective analysis of a defined historical dataset and do not constitute a guarantee or warranty of future or comparable performance.

¹³ Validated extraction accuracy, Stuhlmiller et al., medRxiv, 2025 – <https://www.medrxiv.org/content/10.1101/2025.02.25.25322898v2>