HEALTHSAIL

Run healthcare integration discovery properly

The session agendas, the checklists, and the resources used in the sessions themselves.

HealthSail · Complimentary
Agendas, checklists and reference tables for pre-project discovery

HOW TO USE THIS

Run the three sessions, with us or without us

> This is the material we use to run discovery on healthcare integration projects. It is complimentary. Run the sessions yourself, with any vendor you are considering, or with nobody at all.

Most teams run 30 to 60 minute sessions and get through all three inside a week. Some take a few weeks, and that is fine — but a week is worth aiming for, because the momentum counts for something on its own. The tables in Part 3 onward are the parts worth keeping open on a second screen while you run them.

THE ONE THING TO DO TODAY

Turn to the access table in Part 3 and submit every request on it. Access is the longest pole on almost every project and it is almost never on the plan. A non-production environment at the customer takes two to eight weeks and nothing you do in the meantime shortens it.

PART 1

Why discovery is where projects are lost

Discovery is where integration projects are actually won or lost, and it is almost always compressed because it produces no visible artifact. The pattern is consistent: teams spend two weeks on discovery, six months on the build, and then discover in week twenty-two that the field they needed was never exposed, the customer's build does not include the module, or the rate limit makes the design impossible. Three things determine whether discovery worked. Did you get the right people in the room. Did you request access on day one. And did you write down field-level truth rather than system-level intentions.

PART 2

The three sessions

Each session has one job. Running them out of order, or collapsing them into a single meeting with everyone in it, is possible — but the second session depends on knowing who owns what, and the third depends on the first two having produced something to look at.

High level — the shape of the project

Session 1

Business people and technical leaders together. What the project is for, what it has to do, and the areas that will need access and decisions later. The output of this session is not a specification. It is a list of who else has to be in the room, and what has to be requested today rather than in week six.

- [ ] What the project is for, in the business's own words
- [ ] Which systems are in scope, and which are explicitly out
- [ ] Who owns each system named
- [ ] Which environments and credentials will be needed — start the requests now
- [ ] Who can make a final decision when two departments disagree

Technical — screens open, workflows walked

Session 2

The working session. Screens get opened and workflows walked as they actually happen, not as they are described. Record it where there is no PHI on screen: when people narrate a process they walk the happy path and skip the exceptions, and the exceptions are where the work is.

- [ ] Walk each workflow on the real screens, in the real system
- [ ] Capture field-level detail, not system-level intent
- [ ] Note every exception somebody demonstrates without naming it
- [ ] Confirm the API exposes each field you need — on that version, at that site
- [ ] Get rate limits in writing, or a written acknowledgement that none are published

Walkthrough — visualisations, testing, timeline and budget

Session 3

Everybody comes back to the diagrams. This is where somebody says the form does not capture the field the interface needs — on a diagram rather than in production. The testing plan, the timeline and the budget get agreed against what the first two sessions actually found.

- [ ] Walk the data flow end to end, with everyone looking at the same picture
- [ ] Confirm the exception handling: what retries, what alerts, who the human is
- [ ] Agree the test data, and that it exercises the edges rather than the happy path
- [ ] Timeline against real access lead times, not the build estimate
- [ ] Name the owners: business, technical, escalation, on-call

PART 3

Who needs to be in the room

Three groups have to be represented: the people who know how the work actually happens, the people who can grant access, and somebody who can make a final decision. Miss the second and it becomes a waiting problem — and waiting is expensive in a way that is easy to miss, because it is invisible on a status report.

ROLE	WHY THEY ARE ESSENTIAL	WHAT HAPPENS WITHOUT THEM
Interface analyst / integration engineer	Knows what already exists, which is usually more than anyone documented.	The project discovers a duplicate feed in month four.
Application owner for each system	Knows what the system actually does versus what the brochure says.	Requirements are written against marketing material.
Security	Will review it eventually. Reviewing it now costs days; reviewing it at the end costs weeks.	Go-live slips on a security review nobody scheduled.
Compliance / privacy	Decides whether the disclosure is permitted at all.	You build something you cannot turn on.
Clinical informaticist	Knows whether the workflow is real and whether anyone will use it.	Technically live, operationally ignored.
Revenue cycle	Knows which downstream financial process depends on the data.	A denial pattern appears three months after go-live and nobody connects it.
Vendor implementation lead	Knows their own product's undocumented behaviour.	Every question becomes a support ticket with a five-day turnaround.
One person who can make a decision	Resolves the field-definition dispute between two departments in the room.	Discovery produces an issues list instead of a specification.

PART 4

Access — the real critical path

Access is always the longest pole and is almost never on the plan. Submit every one of these on the day the project is approved, not when the build is ready.

WHAT	REALISTIC LEAD TIME	NOTE
Vendor developer account / sandbox	Same day to a week	Do it before the kickoff call. It is free at most vendors and costs nothing to have early.
Non-production EHR environment at the customer	2-8 weeks	This is usually the longest pole. Request it the day the project is approved.
Production client ID / app registration	Days to weeks, gated by governance	At Epic, get scopes and redirect URIs right the first time — app records are immutable once marked production-ready.
Network path (VPN, private link, IP allowlist)	2-6 weeks	Involves two security teams and a firewall change window. Never fits in a two-week sprint.
Interface engine access or a channel build	1-4 weeks	Depends entirely on whether the customer has an analyst with capacity.
Test data	Immediate if synthetic	Build a Synthea-based library once. It removes the PHI question from every future project permanently.
Named security and compliance contacts	Immediate	The single cheapest thing on this list and the most frequently skipped.

PART 5

Deciding what to integrate

Score every candidate workflow on four axes before anyone estimates anything.

AXIS	AUTOMATE WHEN	LEAVE MANUAL WHEN
Volume	High and repeating — daily or hourly	A handful of times a month
Variance	The process is the same every time	Every case needs judgement
Clinical risk	An error is recoverable and visible	An error reaches a patient before anyone notices
Reversibility	You can turn it off and fall back	Turning it off means a manual backlog nobody can clear

Sometimes the right answer is a better report, not an interface. Saying so is the fastest way to earn a technical buyer's trust, and it is the thing no vendor content will ever say.

THE FIVE FEASIBILITY CHECKS

Before anyone estimates anything

1. **Does the API expose the field at all?** Not the resource — the specific field, on that version.
2. **Does it support write, if you need write?** Read is common; write is where programmes stop.

3. **Does the vendor permit this use?** Custom scopes at Oracle require sales approval and partner certification. Epic app records cannot be changed after production registration.
4. **Does *this customer's* build include the module?** The capability existing at the vendor and existing at the site are different questions.
5. **What is the rate limit?** eClinicalWorks publishes 25 calls per minute per base URL. Most publish nothing — get a written answer before you design.

PART 6

Who controls the connection, by vendor

Approval for a production connection sits in a different place with every vendor. This decides your timeline more than your architecture does.

VENDOR	WHO APPROVES PRODUCTION	PRACTICAL CONSEQUENCE
Epic	The health system. Epic bills the health system a quarterly per-production-directory subscription.	Your sales cycle is 200 hospital IT governance committees, not one Epic contract.
Oracle Health	Split — customer files a Service Request with your app ID and client ID; PECA form where applicable; custom scopes need Oracle sales plus OPN certification.	CommunityWorks, PowerWorks and Continuum customers are excluded from self-service entirely.
athenahealth	athenahealth, centrally.	One integration reaches every practice. Structural advantage worth naming.
eClinicalWorks	Statement of work, per interface.	No public developer programme. 25 FHIR calls/min per base URL.
Veradigm	Paid tier gate — certification sits behind a Gold+ subscription.	Six named tiers, zero published prices.
Altera	Sales plus information-governance review.	No public self-serve developer portal found at all.
MEDITECH	Greenfield is free to enter; the production path is not documented.	Sandbox has three test patients.

PART 7

What discovery has to produce

Discovery is finished when the remaining unknowns are isolated and small enough to say so. In practice that means eight things exist.

Discovery is done when all eight exist. Not before.

1. A field-level mapping document with a named source of truth for every field, signed by both sides.
2. Confirmed API capability: the field is exposed, write is supported if needed, and the customer's build includes the module.
3. Documented rate limits, or a written acknowledgement that the vendor does not publish them.
4. An agreed message pattern — batch, queue, pub/sub or real-time — with the latency requirement written down.
5. Error handling defined: what retries, what alerts, what goes to a human, and who that human is.

6. Volume baselines for normal and peak, so monitoring has something to compare against.
7. Named owners: business owner, technical owner, escalation path, on-call rotation.
8. Environment plan with dates, and confirmation that access requests are already submitted.

PART 8

Timeline reality

Where the time actually goes, and which stages are outside your control:

STAGE	TYPICAL	WHO CONTROLS IT
Vendor sandbox and developer account	Same day to a week	You
Customer non-production environment	2–8 weeks	The customer
IT governance / security approval	2–8 weeks	The customer
Spec agreement and field-level mapping	2–4 weeks	Shared
Build	Varies	You
Integration testing	2–4 weeks	Shared
Pre-production validation	1–2 weeks	Shared
Cutover window	Scheduled, often months out	The customer

Anyone quoting "six weeks end to end" is quoting the build stage and calling it the project.

APPENDIX

Role briefs

What each role cares about, what they will ask, and what goes wrong when they are not in the room. Use these to brief people before the first session rather than discovering their concerns in it.

Chief Information Security Officer

Personally exposed. Every OCR settlement since 2024 rests on the same finding — failure to conduct an accurate and thorough risk analysis — and integration infrastructure is systematically excluded from those analyses.

What they ask

- Where does PHI actually rest in this integration, and for how long?
- Is the interface engine's message archive in our asset inventory? (It usually is not.)
- What are the OAuth scopes, and are any of them wildcards?
- Is this vendor a business associate, and what is the in-scope service list on the BAA?
- Who holds the signing keys, and what happens when the person who created them leaves?
- Can we prove minimum necessary at the field level, not just the message level?
- What is in the dead-letter queue right now, and who has exported it to a spreadsheet?

What changes their answer

- The interface engine message store is usually the largest concentration of PHI outside the EHR and is routinely absent from the risk analysis.
- The HIPAA Security Rule update is not arriving in 2026 — final action moved to a July 2027 projection. Vendors selling that deadline are selling a deadline that does not exist.
- Build to the NPRM's substance anyway: OCR enforces asset inventory, risk analysis, MFA and patching today under the existing rule.
- 36% of breaches occur at business associates by count, and far more by record volume. A BAA is a promise, not evidence, and it does not transfer liability — a 2026 ruling rejected the vendor-only liability defence.
- Unpurged SFTP landing zones and error queues are quiet PHI repositories nobody classified.
- Recovery is the weakest NIST CSF function in the sector: only 22% of health-sector CISOs reach maturity level 4-5 on Recover.

CIO / VP or Director of IT

Owns the outcome and controls almost none of the inputs. Access, vendor governance, EHR release cycles and customer IT committees all sit outside their authority.

What they ask

- Which systems, which versions, which volumes — and who owns each one?
- Do we have an interface inventory, or a list of things people remember?

- How long does access actually take, and did we start on day one?
- What breaks when the EHR upgrades in six weeks?
- Build, buy, or aggregator — and what does year three cost?
- Who is on call for this at 2am, and how would we even know it broke?

What changes their answer

- Access is always the critical path and is never on the project plan. Start every sandbox, credential and VPN request the day the project is approved.
- The health system controls the connection at Epic; Oracle splits it between a service request and a customer PECA form; athenahealth controls it centrally; eCW runs on statements of work.
- Sandboxes lie. Shared test patients, thin data, different scopes and different rate limits are the recurring cause of go-live slippage across every vendor.
- eClinicalWorks publishes 25 FHIR calls per minute per base URL — roughly 0.4 per second. Bulk patterns are simply off the table. Most vendors publish no rate limit at all.
- Community Connect and shared-instance customers frequently cannot approve an integration; the parent organisation can, and often does not.
- If you run Mirth Connect 4.5.2 or earlier you are on frozen, unpatched code. 4.6+ went closed commercial in March 2025.

Integration Engineer / Interface Analyst

The person who will actually be paged. Technical, allergic to condescension, and the internal champion who eventually brings the vendor conversation to the CIO.

What they ask

- Does the API expose this field, and does it support write?
- What are the real rate limits and token lifetimes?
- How do I replay three days of missing messages?
- What happens to my mapping when the lab changes a code?
- Which environment is telling me the truth?

What changes their answer

- An HL7 ACK of AA means received, not processed. Most "we sent it" disputes die on this distinction.
- Epic app records are immutable once marked ready for production. Wrong scope or redirect URI means registering a new app.
- Half-open MLLP sockets, delimiter injection, A40 merges and terminology drift account for the majority of silent production failures.
- Exactly-once delivery is mostly a lie; idempotency keys are the real answer.
- FHIR Subscriptions can set status=error and stop, with no heartbeat and no replay — a permanent gap with no alert.
- Build a synthetic patient library once with Synthea and the PHI risk leaves every future test and demo permanently.

COO / VP Clinical or Ambulatory Operations

Feels the problem daily as staffing, throughput and staff frustration, and rarely connects it to an integration decision.

What they ask

- Where are people retyping data between screens, and how many hours is that?
- What is our procedure when the interface is down but the EHR is up?
- Which workflow should we automate first?
- Are we standardised enough to automate at all?
- How would we know an integration went live and nobody used it?

What changes their answer

- Run a rekeying audit before buying anything: one week, a ranked list, hours and dollars attached. It becomes both the roadmap and the ROI input.
- Automate in order of volume, low variance and low clinical risk — not in order of how interesting the problem is.
- Automating a broken process makes it broken faster and harder to see. Standardise first.
- One health system removed 748 documentation groups, rows and options over 24 months and returned roughly 30,000 hours a year to direct patient care. Deleting fields beat buying software.
- Downtime plans usually assume the EHR fails. Partial failure — labs flowing, orders not — is the more common and less planned scenario.
- 40.3% of provider directory listings were still inaccurate eighteen months after the No Surprises Act verification requirement. That is a data pipeline failure, not a clerical one.

CFO / CEO

Approves or kills the project. Will not read a technical argument and will not fund an efficiency claim without a number attached to a line item.

What they ask

- What does this replace, and what is the second-year cost?
- What breaks if it fails, and what does that cost per day?
- How will we know it is working?
- Is this tied to a revenue line or a regulatory deadline, or is it a nice-to-have?
- What is our exposure if a vendor in this chain is breached?

What changes their answer

- For every dollar US hospitals spent on direct patient care in 2023 they spent nearly two on administrative and operational costs — $687B against $346B.
- Providers spent $25.7B adjudicating claims in 2023 and roughly $18B of it was potentially wasted, because about 70% of denials are overturned and paid anyway.

- Claim submission is 98% electronic. Eligibility is 94%. Prior authorization is 31% and attachments are 29% — the manual work is concentrated in exactly two transactions.
- 63% of healthcare organisations deployed AI or automation in the revenue cycle; only 15% achieved positive ROI, and the third-ranked barrier was integration with existing systems.
- Payers must expose four FHIR APIs on 1 January 2027. As of March 2026, 33% of providers had not begun and provider confidence had fallen from 69% to 47%.
- Median hospital operating margin was 1.9% year-to-date through February 2026, down 13% year over year. Every proposal now competes against that.

Revenue Cycle Director

Sits exactly where integration failures become unpaid claims, and is usually the fastest internal champion once the connection is drawn.

What they ask

- Why do claims disappear between the clearinghouse and the payer?
- Which denial reasons trace back to a data problem upstream?
- Why is my team still logging into payer portals when 270/271 exists?
- What changes for us on 1 January 2027?

What changes their answer

- The 837 → 999 → 277CA → 835 acknowledgement chain is rarely reconciled end to end. Claims die silently in the gap.
- Initial denial rates run around 11% while final denial rates run around 2.7%. The gap between them is the rework cost, and that gap is the entire business case.
- Prior-auth denials rose from 3.2% of claims in 2022 to 10.4% in 2023.
- Provider organisations lost $48B to final denials and bad debt in 2025, up 25% year over year.
- The cost to rework a single denied claim rose from $43.84 to $57.23 in one year — a 30.5% increase.

Compliance / Privacy Officer

Can stop any project and is rarely consulted early enough to say yes cheaply.

What they ask

- Is this disclosure permitted, and under which basis?
- What is the minimum necessary here?
- Does any of this data fall under 42 CFR Part 2 or a state law that overlays HIPAA?
- Are we creating an information blocking exposure by saying no?

What changes their answer

- Information blocking enforcement became operational in 2026 — nonconformity notices to certified developers, with decertification on the table. But no OIG settlement or imposed provider disincentive has been publicly announced. Do not overstate this.
- Median hospital disincentive exposure is $394,353, with a documented range from $30,406 to $2,430,766.
- 42 CFR Part 2 alignment changed what substance use disorder data can travel through general-purpose interfaces. Segmentation is an engineering problem.
- Twenty comprehensive state privacy laws are in effect, several with private rights of action. Washington's My Health My Data Act and California's AB 45 geofencing ban are the sharpest.
- The OCR web tracking bulletin was partially vacated — but that ruling touched none of the CMIA, state wiretap or CIPA theories driving $100M+ in settlements.
