How OpenEvidence compliantly and effectively scaled to $50M ARR witho BAAs or EHR integration: How their unique business model drove hyper-scaling metrics previously considered impossible in health.

DEC 16, 2025 • PAID

DISCLAIMER: The views and opinions expressed in this essay are solely my own do not reflect the views, opinions, or positions of my employer or any organizati with which I am affiliated.

If you are interested in joining my generalist healthcare angel syndicate, reach o trey@onhealthcare.tech or send me a DM. Accredited investors only.

ABSTRACT

OpenEvidence achieved what most healthcare software companies thought impo fifty million in annual recurring revenue within roughly two years of product lau capturing forty percent of US physicians without requiring business associate agreements or EHR integration during its critical growth phase. The company's trajectory from zero to half a billion in funding and a six billion dollar valuation happened because they understood a subtle regulatory distinction that let them like consumer software while everyone else was stuck in enterprise sales cycles. essay examines the specific mechanics of how they did it, why the approach work and why it probably won't work again.

TABLE OF CONTENTS

The Problem With Healthcare Software Distribution

Healthcare software has a distribution problem that's fundamentally different fr other enterprise software categories. The root issue isn't that healthcare buyers a conservative or slow, though they are. The issue is that clinical software usually requires two things that create massive friction: business associate agreements u HIPAA and integration with electronic health record systems. These requiremer turn what could be fast product-led growth into multi-year enterprise sales cycl

The typical path looks brutal. You build clinical decision support software that l doctors make better treatment decisions. To be useful, it needs patient data fron EHR, which means technical integration work. Epic's app orchard review proces takes six to nine months just to get listed. Then each health system needs to ena your app, which requires IT approval. Once you're dealing with patient data, you handling PHI, which means every health system needs their legal team to review sign your BAA. Their information security team does a vendor assessment. Thei clinical leadership needs to see evidence that doctors will actually use it. The wh

process takes twelve to twenty four months per health system, and at the end you might have a few hundred physicians using your software.

This explains why healthcare software companies typically raise a Series A to hir enterprise sales teams before they have real revenue. They spend millions on compliance infrastructure before they know if doctors actually want the product They modify features based on what hospital administrators think doctors need than what doctors actually need. The result is that most healthcare software companies are still trying to close their tenth customer three years after foundin ones that succeed usually do it by grinding through enterprise sales, not by achi viral adoption.

> **OpenEvidence just bypassed this entire machinery an got to fifty million in ARR in roughly two years.**

What HIPAA Actually Requires Versus What Everyone Thinks It Requires

The foundational insight that enabled OpenEvidence's growth was understandir exactly what HIPAA requires rather than assuming it applies to everything clinic Most people building healthcare software assume that any tool used by doctors f patient care needs HIPAA compliance and BAAs. This assumption is wrong but understandably so because the distinction is subtle.

HIPAA requires a BAA when a vendor creates, receives, maintains, or transmits protected health information on behalf of a covered entity. The key term is prote health information, which has a specific legal definition. PHI means individually identifiable health information. The identifiable part matters. If information doe identify a specific patient, it's not PHI under HIPAA, and HIPAA's requirements apply.

This creates a category of clinical software that provides value to doctors withou handling patient-specific data. Medical reference tools fall into this category. Cl calculators can fall into this category if designed carefully. Care coordination to

might fall here if they avoid patient identifiers. The key is that the software ansv clinical questions or provides clinical guidance without needing to know anythi about specific patients.

OpenEvidence launched as a medical search engine and evidence synthesis tool. doctor could ask "what's the recommended antibiotic for hospital-acquired pneumonia in a patient with penicillin allergy" and get an evidence-based answe citations to relevant studies and guidelines. That question doesn't contain PHI. There's no patient name, no medical record number, no date of birth. The fact t doctor is thinking about a specific patient when asking the question doesn't mak question itself PHI. OpenEvidence's system never received any patient data, so i wasn't a business associate under HIPAA, and no BAAs were required.

This wasn't some sketchy legal workaround. It's actually how HIPAA was design work. Medical textbooks don't need BAAs. UpToDate doesn't need individual B with every subscribing physician. PubMed doesn't need BAAs. These are inform resources, not systems that process patient data. OpenEvidence started in exactl category. It was a knowledge tool, not a workflow tool.

The practical implication was enormous for distribution. Without BAAs require OpenEvidence could offer their product directly to individual physicians with ze institutional approval needed. A doctor could verify their NPI, create an accoun start using it in under three minutes. No legal review. No IT security assessment procurement process. No vendor evaluation. Just sign up and go.

Compare this to what happens when you do need BAAs. Even if you have a stanc BAA ready to go, each health system's legal team wants to review it. They have questions about data retention, breach notification procedures, audit rights, liab limits. The negotiations take weeks or months. Information security does their assessment, which generates more questions. By the time you're done, six month passed and you haven't added a single user. OpenEvidence completely avoided t structuring their initial product so it wasn't handling PHI at all.

The EHR Integration Tax

The other major distribution barrier OpenEvidence avoided was EHR integratio
This deserves its own analysis because the conventional wisdom in healthcare so
is that you must be in the EHR workflow to get physician adoption. This belief i
mostly true for certain types of software but not universally true, and the except
create opportunities.

EHR integration makes sense for workflow automation tools. If you're automatic
generating orders or pre-filling documentation or flagging drug interactions in r
time, you need to be embedded in the EHR because that's where those workflow
happen. But not all clinical software is workflow automation. Some clinical softv
supports decision-making that happens before or alongside the EHR workflow.

Think about the actual process of clinical decision-making. A doctor sees a patie
takes a history, does an exam, and thinks about diagnosis and treatment. During
process, questions come up. What does the latest evidence say about this treatm
approach? Are there newer therapies I should consider? What are the current
guidelines for this scenario? These questions often occur before the doctor ever
the patient's chart to document anything. They might happen at the bedside dur
the encounter or immediately after while formulating the plan.

For these information lookup moments, EHR integration doesn't actually help. T
doctor doesn't need the software to read data from the chart. They just need a fa
to get an answer to a clinical question. A standalone app on a phone or laptop ac
works better than navigating through EHR menus. This was OpenEvidence's pro
insight: don't compete with the EHR, compete with Google and UpToDate for th
moments when doctors need information to make decisions.

The strategic advantage of avoiding EHR integration early on was eliminating th
enterprise dependency entirely. Getting into Epic's app orchard takes months.
Cerner's marketplace has its own requirements. Even after you're listed, each he
system's IT team has to enable your app. Then you need clinical champions to tr
people on where to find it in the interface. The whole process easily takes a year
health system.

OpenEvidence skipped all this by being a standalone tool that doctors could disc through word of mouth. A hospitalist in Seattle tries it and finds it helpful. She mentions it to colleagues at a conference. They sign up and start using it at their hospitals in Boston and Atlanta. None of them needed IT permission. None need vendor assessments. None needed training beyond maybe a two minute product The product spread through professional networks the way consumer apps spre not through enterprise procurement.

This distribution approach also created different competitive dynamics. Traditic healthcare software moats come from integration complexity and institutional contracts. Once you've integrated with Epic and signed agreements with fifty he systems, competitors face the same barriers. But you also face those barriers for adoption. OpenEvidence's moat came from network effects among physicians. O forty percent of US doctors were using it daily, that became the standard. Medic residents learned about it from their attendings. Doctors changing hospitals kep using it because it wasn't tied to any institution. The moat built from user habit peer adoption rather than from technical lock-in.

There's also a channel conflict advantage to staying out of the EHR initially. EHI vendors view third-party apps with suspicion, especially successful ones. If OpenEvidence had needed EHR integration from day one, they would have been competing with Epic's and Cerner's own clinical content tools. They might have pressure to revenue share or encountered subtle distribution barriers. As a stanc tool, they avoided that conflict until they were too big to ignore.

How OpenEvidence Structured Their Product to Avoid Both Barriers

The sequencing of OpenEvidence's product evolution shows sophisticated think about how to achieve scale before adding complexity. They didn't try to do every at once. They started with the minimum viable product that could provide enorr value without touching PHI or requiring integration, then layered in more advar capabilities once they had leverage.

Phase one from 2023 through April 2025 was purely the medical search and evid synthesis product. Doctors typed in clinical questions and got evidence-based ar with citations in five to ten seconds. The system searched through thirty five mil peer-reviewed papers and clinical guidelines. It synthesized information across multiple studies rather than just returning a list of abstracts. It worked on mobi doctors could use it anywhere. And critically, it never needed any patient data to function.

The product was immediately useful. Instead of spending twenty minutes search PubMed and reading abstracts, a doctor could get a synthesized answer with rel citations in seconds. The value proposition was obvious and the barrier to trying was zero. This combination drove viral adoption through medical professional networks. Doctors who tried it told other doctors about it. No marketing spend required.

During this phase, OpenEvidence scaled to over one hundred thousand verified physician users in the first year after launch. By July 2024, they were handling th hundred fifty eight thousand physician consultations per month. All of this grov was organic. They weren't spending money on ads or enterprise sales. The unit economics were remarkable: essentially zero customer acquisition cost, minutes first value, minimal churn because it was free and useful. The only real costs wer infrastructure and content partnerships with journals.

In February 2025, OpenEvidence raised a seventy five million dollar Series A led Sequoia at a one billion dollar valuation. This happened roughly two years after product launched. For context, most healthcare software companies at the two y mark are still trying to sign their tenth enterprise customer. OpenEvidence had hundreds of thousands of active users and was embedded in the daily workflow substantial percentage of US physicians.

Then in April 2025, OpenEvidence announced HIPAA compliance and the availa of BAAs. This marked phase two. They added features that did require handling document upload, encounter transcription through their Visits feature, prior authorization letter generation with patient details. For these advanced features

now functioned as a business associate and required BAAs with healthcare organizations.

The sequencing matters enormously. By April 2025, OpenEvidence already had massive adoption among individual physicians. When they approached health sy about BAAs, the conversation was completely inverted from a normal vendor pit They weren't asking IT departments to take a risk on an unproven product. The responding to demand from physicians who were already using the basic version wanted the advanced features. The sales motion flipped from push to pull.

This also changed leverage in BAA negotiations. Normally when a small softwar company asks a major health system to sign a BAA, legal picks apart every claus the process drags on for months. But when that health system's physicians are al using your product extensively and demanding more features, legal has pressure move quickly and accept more standard terms. OpenEvidence could also point t months of usage by physicians at that institution without issues, which reduced perceived risk.

By July 2025, just three months after enabling BAAs, OpenEvidence was handlir eight and a half million physician consultations per month. That represents roug twenty four hundred percent growth from July 2024. They announced a two hun ten million dollar Series B led by Google Ventures and Kleiner Perkins at a thre a half billion dollar valuation. In October 2025, they raised a Series C at six billi dollars. The entire arc from product launch to six billion dollar valuation took roughly two to two and a half years.

The Growth Numbers and What They Mean

The scale OpenEvidence achieved is genuinely unprecedented for healthcare sof As of mid-2025, they had over four hundred thirty thousand registered US physi representing about forty percent of all physicians in the country. They were processing over eight million clinical consultations per month. Over one hundre million Americans were treated by a doctor who had used OpenEvidence that ye

The growth rate is what makes this exceptional. From July 2024 to July 2025, mo consultations grew from three hundred fifty eight thousand to eight and a half million. That's twenty four hundred percent growth in a single year. Consumer s apps sometimes hit these growth rates but healthcare tools just don't, specificall because of the distribution barriers discussed earlier.

The revenue progression tells a similar story. Based on available information, OpenEvidence reached approximately fifty million dollars in annualized revenue mid-2025, growing from an estimated eight million dollars at the end of 2024. T kind of growth trajectory is more typical of consumer internet companies than healthcare software businesses. The primary revenue model was pharmaceutical advertising shown to physicians in contextually relevant ways. When a doctor qu about heart failure treatment, showing information about relevant cardiac medications is actually useful rather than intrusive.

The advertising model worked because of the scale and targeting. With eight mi monthly consultations and forty percent of US physicians using the platform, OpenEvidence had better reach among doctors than almost any other channel. A they had superior targeting because they knew exactly what clinical questions ea doctor was asking in real time. A pharmaceutical company launching a new diab drug would pay premium rates to reach endocrinologists and primary care docto the moment they're researching diabetes treatments. Reports suggest OpenEvide was achieving CPMs of seventy to one hundred fifty dollars compared to five to f dollars for typical social media advertising.

Beyond advertising, the enterprise BAA model created a higher-value revenue st Health systems that wanted their physicians to access advanced features needed enterprise agreements. These were priced per physician and represented signific higher revenue per user than the ad-supported model. The enterprise motion als created switching costs because once physicians got used to advanced features li document upload and encounter transcription, they'd resist losing access.

The capital efficiency deserves attention. Through their Series C, OpenEvidence raised roughly five hundred million dollars. That sounds like a lot in absolute te

but it's actually quite efficient given their scale and valuation. Most healthcare software companies raising similar amounts might have ten to twenty million in OpenEvidence was doing fifty million in ARR on comparable capital raised. The efficiency came from avoiding enterprise sales costs, extended sales cycles, and c integration work during their growth phase.

The valuation progression also tells the story. One billion dollars in February 20 three and a half billion in July 2025, six billion in October 2025. The company's valuation increased six fold in eight months. The investors weren't naive capital. Sequoia led the Series A. Google Ventures and Kleiner Perkins co-led the Series These firms have deep healthcare investing experience and saw something genui different in OpenEvidence's scaling dynamics.

Why This Window Has Closed

The obvious question is whether this playbook can be replicated. The answer is mostly no for fundamental structural reasons, and definitely no for the specific medical search use case because OpenEvidence has already captured it. Understanding why helps clarify what was actually special versus what was circumstantial.

The core requirement for avoiding BAAs is that your product genuinely doesn't PHI to provide value. This is true for medical reference and education tools. It n be true for some clinical calculators and risk scores if designed carefully. It coul work for certain care coordination tools if structured to avoid patient identifiers it's fundamentally not true for most clinical workflow software.

If you're building clinical documentation tools, you need the patient's chart data you're optimizing medication orders, you need to know current medications. If y identifying care gaps, you need to identify patients. All of these require PHI fro one and therefore require BAAs. There's no way around it.

The medical search and evidence synthesis use case OpenEvidence addressed wa uniquely suited to the PHI-free model. Doctors constantly have clinical questior

These questions are often generic enough that they don't require patient-specifi to answer. The answers come from medical literature and guidelines that are the regardless of which patient prompted the question. This let OpenEvidence provi massive value without knowing anything about specific patients.

Even within the constraints of PHI-free products, OpenEvidence had circumstar advantages. The medical literature is largely publicly accessible through PubMe journal partnerships, so they could build a comprehensive knowledge base witho proprietary data access. Clinical questions have commonality across specialties a practice settings, so the product worked for all physicians rather than being nicl The value proposition was immediate and obvious, get answers in seconds instea searching for twenty minutes.

The timing also mattered significantly. OpenEvidence launched in 2023 right as language models became capable enough to do sophisticated literature synthesis Earlier attempts at AI-powered medical search struggled because the underlying wasn't good enough. Later attempts will struggle because OpenEvidence already captured the market. There was a narrow window where the technology enabled product and the market was still open.

The distribution dynamics also benefited from timing. Physicians were primed t AI tools by 2023 because ChatGPT had created mainstream awareness. But ther weren't yet many clinical AI tools that doctors trusted, which created an opening OpenEvidence also benefited from being genuinely better than existing alternati like UpToDate for the specific use case of quick evidence lookups. That perform advantage won't last forever because incumbents will adapt.

The regulatory environment might also shift in ways that close this approach. As clinical tools become more common, there's increasing scrutiny about whether t should be regulated as medical devices or whether HIPAA's definitions of PHI sh expand to cover de-identified clinical scenarios. If regulations change to require even for tools that don't technically handle identifiable patient data, the distribu advantage disappears.

There are probably a few other healthcare software categories where a similar approach could work. Medical education and continuing education tools might a PHI requirements. Clinical trial matching could potentially work with de-identi patient criteria. Second opinion platforms might structure workflows to avoid identifiable data in certain scenarios. But the number of high-value clinical use that can provide significant utility without any patient-specific data is quite limi

The meta lesson is that there are occasionally windows in healthcare technology new capabilities enable distribution models that bypass traditional gatekeepers. Mobile health apps in the early 2010s were one such window. Patients and docto could download apps without institutional approval, creating opportunities that didn't exist before. AI tools in 2023 through 2024 were another window, specific for products that could provide clinical value without data integration. These windows tend to be temporary because once a category gets established, the dyn shift toward more regulation and more competition.

Implications for Healthcare Software Strategy

OpenEvidence's trajectory offers several strategic lessons for anyone building or evaluating healthcare software, though the specific playbook probably won't rep The high-level insight is that challenging assumptions about necessary distribut barriers can create enormous value if you find the right category and execution.

The first implication is about product design and regulatory strategy. Most healt software teams assume they need full HIPAA compliance from day one because t product will eventually handle PHI. This leads to building compliance infrastru before achieving product-market fit. A better approach might be to ask whether core value proposition actually requires PHI or whether there's a useful PHI-fre version that could drive initial adoption. If you can deliver value without PHI, yo achieve fast distribution and validate product-market fit before adding compliar complexity.

This doesn't mean ignoring compliance or cutting corners. It means understand exactly what regulations require for your specific product design rather than assuming maximum regulation applies. The key question isn't "is this a clinical but rather "does the core functionality require creating, receiving, maintaining, transmitting identifiable patient data." If the answer is no, you might have a distribution advantage.

The second implication is about distribution strategy and unit economics. Produ that can scale through individual adoption rather than institutional procuremen fundamentally different economics. Customer acquisition cost should be near ze you're achieving organic viral growth. Time to value should be minutes, not mor Implementation costs should be minimal. If you're building a direct-to-clinician product but need significant marketing spend or complex onboarding, somethin wrong with product-market fit.

The challenge is that direct-to-clinician products have different monetization constraints. OpenEvidence could monetize through pharmaceutical advertising because they had scale and the engagement patterns worked for that model. Mos clinical tools won't have that option. Many will eventually need to monetize thro institutional subscriptions, which brings back the need for enterprise sales capa The strategic question is whether you can achieve enough individual adoption fi that institutions feel pull to purchase rather than needing push sales.

The third implication is about competitive dynamics and moats. Traditional healthcare software moats come from integration complexity and institutional relationships. These are real but they also slow initial adoption. Direct-to-clinic products have weaker moats in some ways because competitors can build simila standalone tools. But they have stronger moats in other ways because user habits network effects among clinicians create stickiness that's hard to displace. The ke reaching critical mass quickly enough that you become the standard before competitors can copy you.

The fourth implication is about sequencing and timing. OpenEvidence didn't try everything at once. They started with the minimum product that could provide v

without PHI, achieved scale, then added advanced features that required BAAs o they had leverage. This sequencing only works if the minimal version truly solve painful problem. Too many products try to start minimal without ensuring the minimal version is actually valuable enough to drive adoption.

The timing element is also crucial. Healthcare tends to adopt new technologies s but occasionally there are windows where new capabilities enable bypass strateg around traditional gatekeepers. The teams that win are the ones who recognize t windows early and execute fast enough to capture them before the window close challenge for strategists is identifying what the next such window might be and whether they can move quickly enough to exploit it.

The final implication is about risk and variance. Traditional healthcare software predictable but slow growth and long sales cycles but also relatively predictable outcomes if you execute on enterprise sales. Direct-to-clinician products are hig variance. If they achieve product-market fit with individuals, they can scale incr fast with great economics. But if they don't achieve that fit, there's no enterprise motion to fall back on. You either get viral adoption or you get nothing. For som teams and some strategies, this risk profile is attractive because the upside is dramatically higher.

OpenEvidence represents a rare but important case study in healthcare software distribution strategy. By finding a clinical use case that provided significant valu without requiring PHI handling or EHR integration, they achieved consumer-gr scaling velocity in what's normally an enterprise software category. The specific playbook won't work for most healthcare software because most clinical use case require PHI and integration by their nature. But the meta lesson about carefully analyzing which barriers are actually necessary versus which are assumed to be necessary should inform how healthcare software gets built. The next breakthro healthcare software company probably won't look like OpenEvidence specifically it might share the characteristic of finding creative ways around traditional distribution barriers through careful product design and regulatory understandi rather than just accepting that healthcare software has to be slow.

4 Likes

Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...

Zahid Butt Dec 16

Excellent analysis. Enjoyed reading it.

LIKE REPLY