The Free Lunch Is Over, Except Now It Not: What Near-Zero Software Costs Mean for Every Player in Healthcare

FEB 24, 2026 • PAID

Abstract

This essay argues that the collapse of software development costs, driven by AI c tools, will be one of the most disruptive forces in healthcare over the next five to years, arguably more disruptive than any single regulatory change or clinical breakthrough. The implications cut differently across hospitals, payers, pharma, vendors, but the common thread is that software as a moat is largely dead, and tl winners will be those who figured that out early.

Key claims:

- Software development costs are falling 80-90% for many use cases, with agentic coding tools like Cursor, Devin, and GitHub Copilot dramatically compressing b timelines

- For hospitals and health systems, this means internal IT teams become credible builders again, threatening incumbent EHR and middleware vendors

- For payers, it means utilization management, prior auth, and claims adjudicati logic can be rebuilt internally at a fraction of historical cost, destabilizing a generation of point solutions

- For pharma, clinical trial software, regulatory submission tooling, and commer analytics platforms become commoditized, shifting value to data and relationshi

- For health tech vendors, any company whose core defensibility was “we built tl thing and you can’t” is in serious trouble

- The real winners are those sitting on proprietary data, clinical workflows, and regulatory relationships that software alone cannot replicate

> This is not a five-year story, it’s a two-year story

Table of Contents

The Actual Premise: Software is Becoming a Commodity Input

There’s a useful analogy buried somewhere in the history of electricity. Before widespread electrical grids, manufacturers built their own power generation on- It was expensive, it required specialized expertise, and it was a legitimate compe differentiator to have reliable power when your competitor didn’t. Then the grid happened, and power became a utility, and overnight the differentiator evaporat Nobody today builds a factory and considers their access to electricity a competi moat.

Software in healthcare has operated for the last thirty years roughly like private generation. Building it was expensive and slow. A mid-sized health system tryin custom-build a care management platform was looking at multi-year timelines, figure budgets, and a constant risk of the whole thing collapsing when three key engineers left for Google. So instead, everyone bought. They bought Epic. They bought Salesforce. They bought a hundred point solutions for a hundred specifi workflows. And the vendors who built those things had real moats, because the switching costs were brutal and the alternative was attempting to rebuild intern which was essentially impossible at reasonable cost.

That equation is breaking down fast. Tools like GitHub Copilot, Cursor, and the newer agentic coding platforms are compressing development timelines by anyw from 50 to 90 percent depending on the use case. Some enterprise teams are rep that work that used to take a senior engineer six weeks is getting done in three d The models are not perfect, the output requires review, and complex distributed systems still require serious human architecture decisions. But for the enormou category of healthcare software that is essentially business logic wrapped in a UI some integrations, the cost structure is collapsing.

This matters more in healthcare than almost anywhere else because healthcare is uniquely full of software that is essentially business logic wrapped in a UI with integrations. Prior auth platforms. Care gap identification tools. Claims repricin engines. Quality reporting dashboards. Population health analytics. Contract modeling tools for value-based care. Virtually every category of health tech poin solution that raised a Series A in the last decade is, if you strip away the brandin set of rules encoded in software running against healthcare data. And encoding in software is exactly what these new tools are extraordinarily good at.

The people who built those companies are smart and they know this. The honest will tell you privately that they are terrified. The less honest ones are writing blo posts about how AI will create more demand for their platform, which is technic true in some narrow sense and deeply misleading about the structural shift happ underneath them.

What This Means for Hospitals

Hospital IT departments have historically been places where software goes to di slowly. The organizational dynamic is well understood: clinical operations dema something, IT scopes it, procurement negotiates for eighteen months, implemen takes another year, and by the time the thing is live the clinical workflow it was designed for has already changed. The internal build option gets floated approximately once per initiative, immediately dies when someone calculates the cost and timeline, and everyone agrees to just buy the vendor product and live w seventy percent of their requirements being met.

That calculus changes materially when the build timeline compresses by eighty percent and the required expertise becomes substantially less specialized. A hos system that used to need a team of twelve engineers and eighteen months to buil custom prior auth workflow tool now potentially needs three engineers and six v The total cost of internal build for a specific use case drops from, say, four milli dollars over two years to somewhere around three hundred thousand dollars. At price point, a lot more use cases become worth building internally, especially wh the vendor alternative comes with integration headaches, contract lock-in, and a implementation process that requires hiring a consulting firm to manage.

The specific categories where hospitals are going to start building internally, or dramatically expanding internal capabilities, are predictable. Anything touching management workflows is high on the list, because the gap between what a vend product does and what a specific health system's clinical model requires is almos always enormous, and most of the value of care management is in the workflow specificity. Patient communication and engagement tooling is another obvious o since the major vendors in this space have built generic solutions and hospitals increasingly have specific needs around populations, languages, and care modali that generic doesn't serve well. Quality reporting and regulatory submission log third, which is basically a tax on hospitals today because the vendors who built t for this have permanent repricing power as long as switching costs remain high.

The harder question is what happens to the relationship with the big EHR playe Epic and Oracle Health have spent decades building something that is genuinely to replicate: clinical workflow integration, deep regulatory compliance baked in product, decades of implementation experience, and a network of trained staff a every health system in the country. Near-zero software costs do not immediately threaten that core. What they do threaten is the vast ecosystem of bolt-on tools t exist only because the EHR doesn't do something and building it internally was expensive. A large chunk of the health tech market is basically "Epic doesn't do yet," and that whole category is about to get very interesting.

Hospital venture funds and innovation arms are also worth watching here. For tl several years these have been primarily check-writing and ecosystem-building vehicles, because the health system genuinely could not build most of what they investing in. If internal build costs drop dramatically, a health system's build-ve buy math shifts on basically every category, and their innovation arm becomes something closer to an actual product studio. Some health systems are already m in this direction.

What This Means for Payers

Payers are sitting on a particularly interesting problem. On one hand, they have enormous internal engineering capabilities already, especially the large national On the other hand, they have spent thirty years outsourcing logic to vendors pre because building and maintaining that logic internally was expensive and the ve specialized in it. The prior auth vendors, the utilization management platforms, network adequacy tools, the fraud waste and abuse detection systems: all of thes in their current form partly because payers decided at some point that buying w more efficient than building.

The prior auth category deserves special attention because it is both enormous a genuinely vulnerable. The vendor companies that built prior auth platforms did encoding clinical criteria, building payer-specific rule sets, and developing the integrations into payer claims systems and provider EHRs. The business logic it

not particularly complex. The defensibility came from implementation inertia ar cost of rebuilding. Both of those defenses weaken significantly when a payer's in team can replicate the core functionality in weeks rather than years.

There is a regulatory dimension here too. The CMS prior auth transparency rule are rolling out over the next couple of years are creating pressure on payers to re their prior auth infrastructure anyway. A payer that was going to have to rebuild integrations for regulatory compliance has a natural moment to ask whether the rebuild on the vendor platform or rebuild internally. As internal build costs fall, of them are going to choose internally.

The population health and care management vendor space faces a similar dynam Large payers already have clinical informatics teams that understand the logic b than the vendors do in many cases. What they lacked was cheap build capacity. I constraint loosens, the sophisticated national plans start internalizing much mo what they currently buy from point solution vendors, and the market for popula health platforms starts bifurcating sharply between plans that can build and sm regional plans that can't.

Pharmacy benefit management and drug pricing analytics is another category wh the build economics are shifting. The complexity of PBM logic is real, but a significant portion of it is business rules, contract math, and formulary manager that does not require specialized software engineering. Companies that have bee selling analytics and decision support tools to payer pharmacy teams are going t meaningful internal build competition within the next few years.

What does not get disrupted as easily on the payer side is anything requiring regulatory approval or actuarial certification, and anything touching claims adjudication at the core infrastructure level, where the legacy systems are so dee embedded that even cheap build capacity doesn't solve the problem of replacing something processing a billion claims a year. Payers will continue buying in thos areas, and the vendors serving those areas have more durable positions.

What This Means for Pharma

The pharma industry has a slightly different relationship with software than hos and payers do. Pharma doesn't primarily run on software in the operational sens a health system does. Software for pharma is a set of tools for running clinical tr managing regulatory submissions, doing commercial analytics, and optimizing n access. These are all important, but they are not the core value-creating activitie the enterprise, which remain chemistry, biology, and clinical development exper

The practical implication is that software cost compression hits pharma somewh differently. The clinical trial technology space, which has attracted substantial venture investment over the last several years on the thesis that decentralized an hybrid trial models would transform drug development, is looking at a world wh the software layer becomes dramatically less differentiated. Companies that rais substantial rounds on the premise of proprietary technology enabling decentrali trials are going to find that their technology advantage erodes faster than expect The underlying market need, which is genuinely reducing trial costs and timelin remains real. But the defensibility increasingly has to come from regulatory relationships, site networks, patient recruitment infrastructure, and operational expertise rather than from software per se.

Regulatory submission tooling is worth specific attention. The major platforms serving regulatory affairs teams at large pharma companies are genuinely compl deeply integrated into regulatory workflows, and subject to constant change as F guidance evolves. But they are also fundamentally software products encoding ru and formats, and pharma companies with large regulatory affairs operations are to start asking hard questions about build versus buy as the build cost drops. Th medium-sized pharma companies and specialty pharma operators, who historica lacked the internal capabilities to build anything, may actually be the more inter opportunity for existing vendors, since those customers will continue to buy eve the largest accounts start internalizing.

The real estate of highest value in pharma software is anything touching real-wo evidence and data linkage. Pharma's ability to use real-world data for label expa outcomes research, and market access decisions is increasingly central to comme strategy. The software and data infrastructure enabling that is valuable not prim

because the software is hard to build but because the data relationships and regulatory standing are hard to replicate. Companies sitting at the intersection c longitudinal patient data and pharma's real-world evidence needs are actually reasonably well-protected from the software cost compression story, because the value is the data and the relationships, not the query engine sitting on top.

Commercial analytics for pharma, which is a massive market covering market ac modeling, payer coverage analytics, and sales force optimization, is more vulnera Much of this is sophisticated but ultimately replicable analytics logic running ag data that pharma companies already have. As build costs fall, expect pharma ana teams to insource more of what they currently buy from specialized vendors, and expect that market to consolidate significantly as the easy use cases get handled internally.

What This Means for Health Tech Vend

This is where things get genuinely uncomfortable. The health tech vendor mark been built on a set of assumptions about the relative cost of building software, tl specialization required to do so, and therefore the defensibility of having built a thing. Those assumptions are breaking down, and the implications for a large pc of the venture-backed health tech market are not good.

The categories most at risk are those where the core product is essentially a wor tool with healthcare-specific configuration, no proprietary data asset, and a cust relationship that is primarily about the software rather than about services, relationships, or regulatory positioning. This describes a lot of companies. Patie engagement platforms with no unique data. Prior auth tools with no proprietary clinical criteria or payer relationships. Care gap identification platforms sitting data that customers could access directly. Quality reporting tools encoding CMS that are publicly available.

The argument that some of these companies make, and it is not entirely wrong, i even if the software cost falls dramatically, customers still lack the organizationa capability to execute an internal build. This is true for small and mid-sized cust

and it is increasingly less true for large sophisticated customers who have been c
building engineering capacity. The health systems that will start building intern
are the ones with revenue above a couple billion, mature IT operations, and exist
data and analytics capabilities. That is exactly the customer segment that these
vendors have been most aggressively targeting, because those are the customers
can pay real money.

What survives and potentially thrives in this environment is anything where the
software is inseparable from something that is genuinely hard to replicate. The b
version of this is proprietary data that only exists because the vendor has built u
unique longitudinal asset through years of customer relationships and data
aggregation. A vendor sitting on connected data across thousands of providers a
payers has something that a health system building internally cannot easily repli
because the value of that data comes from the network, not from the software
processing it. Scale network effects on data are actually getting more valuable as
software gets cheaper, because the data becomes the moat rather than the techno

Regulatory positioning is similarly durable. A vendor that has spent years naviga
FDA clearance, CMS certification, or HIPAA compliance programs has built
something that does not become less valuable when AI coding tools get better. If
anything, the compliance infrastructure becomes more valuable as a differentiat
when everyone can build the software quickly but not everyone can build the sof
quickly and get it certified.

Clinical expertise embedded in the product is the third durable category. The re
Epic is not going to be disrupted by near-zero software costs in any short-term
timeframe is not really the software itself, it is the thirty years of clinical workflo
logic, regulatory compliance, and implementation methodology embedded in the
product. That takes time to accumulate and cannot be replicated cheaply regard
what the AI coding tools do.

Where the Real Moats Are

The pattern here is clear enough to be worth stating directly. Software is becomi commodity input, and the winners in health tech going forward are those who h built moats in things that are not software. The shorthand is data, relationships, regulatory positioning, and the companies that stack multiple of these have genu durable positions.

The most interesting class of company to own right now is one where near-zero software costs actually help rather than hurt. This exists. Companies that have historically been constrained by build costs in their ability to deliver customized implementations, expand into adjacent workflows, or build integrations with ne sources, stand to benefit significantly from the same cost compression that threa pure software players. A company that has strong clinical workflow expertise an customer relationships but has always struggled to build fast enough gets dramatically better as an organization when internal build velocity increases. Th are often the services-heavy or implementation-intensive companies that venture investors have historically underweighted because the gross margin profile was The margin profile improves as build costs fall while the clinical and relationshi assets appreciate.

The other interesting class is pure data plays with software that was previously necessary but not central. There are companies that built software tools essentia a means to an end, where the end was accumulating a data asset through clinical claims integration, and the software has always been a cost center rather than a driver in the business model. As the cost of maintaining and improving that soft layer drops dramatically, the economics of those businesses improve significantl data asset remains valuable. The tax of maintaining the software wrapper aroun decreases.

What Investors Should Actually Be Doi About This

The straightforward implication for health tech investors is to be very skeptical company whose defensibility narrative centers primarily on having built a thing.

build is replicable now, faster than it used to be, and in two to three years it will faster still. The diligence question that matters is not "is this technically impres but "what happens to this company when every mid-sized health system can reb this in six weeks."

The companies worth backing are those that can answer that question with some other than "we'll move faster." Moving faster is not a moat when the underlying technology is available to everyone. The answer has to be something like "we hav years of longitudinal data that cannot be recreated," or "we have regulatory appr that took five years to obtain," or "we are deeply embedded in clinical workflows three hundred health systems in a way that creates actual switching costs indepe of the software itself."

The portfolio construction implication is to look more carefully at companies th have been historically undervalued because of gross margin concerns. If a compa has sixty percent gross margins because it is services-intensive around deep clini expertise, and the reason competitors haven't been able to replicate it is the expe rather than the software, that company may be dramatically better positioned th pure-play software company at ninety percent gross margins whose product enc publicly available clinical criteria.

For founders specifically, the message is to run the thought experiment honestly someone handed a competent three-person engineering team a hundred thousan dollars and agentic coding tools, could they replicate your product's core functio in six months? If the answer is yes, you do not have a defensible technology posi and that is fine as long as you are honest about it and building defensibility elsev The founders who are in trouble are those who have convinced themselves that t technology is harder to replicate than it actually is, and are counting on that as t primary competitive advantage.

The macroeconomic implication for the overall market is a meaningful compress the valuation multiples being applied to undifferentiated health tech software. T last several years have seen revenue multiples for health tech SaaS that implicitly assumed strong pricing power and high switching costs. Both of those assumptio

weaken as build costs fall and customers develop credible internal alternatives. T compression is not uniform, and companies with genuine data or regulatory moa may actually see their relative valuations improve as investors distinguish betwe real and illusory defensibility. But the average multiple paid for average health t software in 2021 was pricing in a world that is no longer coming.

None of this plays out overnight. Healthcare is famously slow to change, procure cycles are long, and organizational inertia is real. The health system that theoret could build internally will still spend twelve to eighteen months deciding wheth do so, scoping the project, and getting internal alignment. Vendors have time to adapt. The ones who will fail are those who use that time to argue the disruptior real rather than repositioning toward assets that actually last. The electricity gri analogy holds here too: the manufacturers who went out of business weren't tho who failed to generate their own power. They were the ones who couldn't figure anything else to sell.

3 Likes • 1 Restack

← Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...