

The Health System Opportunity Stack A Builder's Guide to the Most Underserved Enterprise in America

MAR 05, 2026 • PAID



Share

Abstract

Health systems are operationally underwater and analytically blind across dozen high-value domains simultaneously. The gap between what exists commercially and what these organizations actually need is enormous, and the window for building category-defining companies inside this gap is open right now for a specific set of structural reasons. This essay maps the opportunity stack, explains why now is the right moment, and gives builders and investors a prioritized view of where the real opportunities are and why.

Key themes:

- The biggest health system software gaps are in financial operations, not clinical
- Data network effects compound faster in health system contexts than in any other enterprise vertical
- The entry point to most of these markets is a CFO conversation about quantifying revenue loss, not a clinical champion selling upward
- Several of these opportunities have natural exit buyers already circling the category
- Sequencing matters enormously, starting with the cheapest and fastest builds first, then the harder ones

Table of Contents

Why health systems are the most compelling enterprise software opportunity in market right now

The financial operations layer that nobody has built

Workforce intelligence and what \$29 billion in agency spend looks like when it's solvable

The operating room problem and why 65% block utilization is basically embezzled

Prior authorization as a regulatory forcing function

Clinical variation and what two surgeons doing the same knee replacement tells about margin

The data monetization play hiding in plain sight

Where to start and why sequencing is the whole game

Why health systems are the most compelling enterprise software opportunity in the market right now

Health systems are, by almost any analytical measure, the most complex operating organizations in the American economy. A single large health system manages tens of thousands of employees across dozens of employee classifications with wildly different compensation structures, benefits, union contracts, and scheduling requirements. It negotiates multi-year contracts with dozens of commercial payers where 60-70% of patient revenue depends on the outcome. It runs an OR generating 40-60% of its margin in a physical space that represents maybe 3-5% of its square footage. It processes hundreds of millions of dollars in accounts payable annually, manages pension obligations that would terrify a mid-size manufacturer, operates clinical

research programs, employs thousands of physicians with their own productivity compensation models, and does all of this while trying to keep people alive.

The software that exists to support most of these operations was not built for health systems. It was built for hospitals in the 1990s and 2000s, adapted for health systems as they consolidated, and patched together in ways that leave stunning amounts of operational intelligence on the floor. The CFO of a billion-dollar health system negotiates payer contracts using anecdote and historical rate data while UnitedHealthcare sits across the table with actuarial teams, claims data on millions of patients, and benchmarks across every market they operate. The OR director gets utilization reports on what happened last month, not predictions of what is going to happen next week. The nursing director finds out about a staffing gap when a shift starts, not days before when something could have been done about it.

This is not a problem of ambition or intelligence inside health systems. These are sophisticated organizations run by capable people. It is a problem of tools. The commercial software ecosystem that serves health systems has prioritized EHR data, billing compliance, and clinical documentation over operational intelligence, financial analytics, and the kind of real-time decision support that actually changes outcomes. Epic is extraordinary at what it does. What it does is not run a business.

The timing argument for building in this space right now has a few components stacked on each other. The financial pressure on health systems is at levels not seen since the post-ACA adjustment period, which means CFO conversations about quantifiable cost reduction and revenue recovery are shorter than they have ever been. The AI infrastructure available to builders today makes products feasible that would have required armies of data scientists three years ago – the extraction logic that would take a team of managed care attorneys can now run through an API call. And the regulatory environment, particularly around prior authorization and data interoperability, is creating forcing functions that compress enterprise sales cycles from 18 months to something actually manageable. When a mandate says you have to do something by January 2027, the procurement process gets motivated.

The counterintuitive insight is that the biggest opportunities are not in clinical. Clinical AI is genuinely important and the market for it is enormous, but the competitive intensity is also extraordinary and the sales cycles are brutal because clinical validation requirements are appropriately high. The most compelling near term opportunities are in financial operations, workforce intelligence, and data infrastructure – categories where the ROI is quantifiable, the buyer is a CFO or rather than a clinical committee, and the competitive landscape is weak because everyone with a healthcare AI thesis has been chasing clinical documentation and diagnostic support.

The financial operations layer that nobody has built

Start with the most egregious gap, which is payer contract intelligence. Health systems negotiate multi-year agreements with commercial payers that determine majority of their revenue, and they go into those negotiations almost entirely blind. The information asymmetry is total and structural. United Healthcare has actual teams modeling claims across every market they operate in. Aetna knows exactly every competing health system in a given geography is accepting for DRG 470. The health system managed care team has their current rates and maybe some consultant anecdotes from similar markets.

The benchmarking database that would fix this has never existed because no single health system has data across markets. The opportunity is to build it collectively anonymized rate data from a network of participating systems. The product itself is not particularly technically complex once you have the data – contract term extraction using modern LLMs, a rate benchmarking layer that answers questions like what United pay for a knee replacement in Dallas versus Denver, and an underpayer monitoring dashboard that flags where actual remittances are running below contracted rates. That last feature alone is worth the price of admission for most health systems, because underpayment rates of 2-4% of net revenue are essential and most systems lack the tooling to detect them systematically.

The business case closes itself. A billion-dollar health system that is being underpaid at 3% of net revenue is leaving \$30 million on the table annually. A platform that takes half of that and charges \$1 million per year is a no-brainer conversation. The negotiation intelligence feature is additive – showing a managed care team that other systems in similar markets are receiving rates 15% higher for a specific DRG is worth far more than a million dollars in the next contract cycle. The strategic acquirers in this category are obvious: Kaufman Hall, Vizient, Premier, any major consulting firm that wants to own the data layer of managed care advisory, or private equity building a healthcare analytics platform. Five to eight times ARR multiple at meaningful ARR is a very credible exit.

AP automation is a different kind of financial operations play but equally underdeveloped. A large health system processes \$500 million to \$2 billion in accounts payable to vendors annually – medical supplies, pharma, staffing agencies, software, facilities, contractors – and almost all of it moves by check or ACH. The interchange revenue that could be generated by converting those payments to virtual card, typically 1% per transaction, goes entirely to banks who have done almost nothing to earn it. A virtual card program that shares the majority of interchange back to the health system as a cash rebate and wraps it in an AP automation layer generates what is essentially found money for a CFO who is looking everywhere for non-clinical revenue.

The reason this has not been done well is vendor enrollment. Getting vendors to accept card payment requires outreach at scale, and most fintechs doing vendor enrollment have no leverage with healthcare suppliers. The leverage comes from representing enough combined purchasing volume that vendors have a genuine incentive to change their payment acceptance infrastructure. Represent a network of health systems paying a specific medical supply company \$100 million annually and you have an entirely different conversation than a generic AP fintech that can only get access to maybe \$10 million in volume. The business model generates interchange revenue from the first transaction and scales with every additional vendor that converts, which means the economics compound without proportional cost increase.

Treasury management is adjacent and similarly underdeveloped. A large health system has dozens of bank accounts across multiple entities – operating accounts, restricted

funds, foundation accounts, bond proceeds, self-insurance reserves, pension assets managed by a small treasury team using bank portals and spreadsheets. Cash visibility across the enterprise is minimal. Idle cash optimization is essentially nonexistent. Interest income left on the table through poor cash positioning is millions annually at current rates. The product that Kyriba or Broadridge provides for large corporations does not map cleanly to health system needs because the nonprofit status, bond covenants, and restricted fund rules create compliance requirements that general treasury platforms handle poorly. Purpose-built means native handling of those constraints, which means the product actually works rather than requiring a treasury analyst to manage the edge cases manually.

Workforce intelligence and what \$29 billion in agency spend looks like when it is solvable

Health systems spent \$29 billion on contract and agency nursing labor in 2023. At the pandemic peak it was \$49 billion. Even at normalized rates, this is an extraordinary cost that is largely preventable, and the prevention mechanism does not require eliminating agency nurses – it requires predicting demand accurately enough to act proactively rather than reactively. The gap between those two things is the product.

The core insight that sells this category is the distinction between preventable and structural agency spend. Every CFO knows their agency nursing line item. Almost none of them know how much of it was genuinely unavoidable versus the result of poor demand visibility, inadequate float pool management, or scheduling decisions made on the basis of last month's patterns rather than a forward-looking forecast. Show a CFO that \$20 million in agency spend contains \$7 million that was structurally unavoidable and \$13 million that a 30-day demand forecast would have allowed them to cover with internal resources, and you have created urgency that a vendor pitch deck could manufacture.

The technical build for this is not exotic. Historical census data by unit, scheduling system exports, and agency invoice files are the inputs. Time series demand prediction

at the unit-shift level is a solved problem technically. The complexity is in the institutional data aggregation – getting clean, consistent data out of health system scheduling systems requires integration work that is not glamorous but is very real. The output is a 30-90 day staffing demand forecast by unit and shift, a float pool optimization recommendation layer that matches predicted gaps with available internal resources before the gap becomes a crisis, and a CFO dashboard that shows preventable versus structural agency spend with enough granularity to act on it.

The commercial model can include performance pricing – a share of documented agency spend reduction above a baseline – which turns the product into an easier self-justifying investment. A health system spending \$20 million on agency labor; the platform reduces by 20% has saved \$4 million. A \$750 thousand annual contract against that outcome is an easy conversation, and the performance fee on top of it aligns incentives in a way that makes renewal automatic rather than a negotiation.

Physician workforce intelligence is a related but distinct problem. Health systems spent the last decade acquiring physician practices at extraordinary valuations and then destroying value during integration. The failure mode is consistent: compensation model misalignment, productivity drops, physician attrition, cultural friction, and a general absence of systematic integration infrastructure. Every acquisition is essentially a bespoke consulting project that ends when the consultants leave, taking the institutional knowledge with them and leaving behind unhappy physicians and data that does not flow coherently between the acquired practice and the health system.

The platform that manages the full lifecycle of an employed physician enterprise – compensation benchmarking, productivity analytics, satisfaction monitoring, credential management, cultural integration workflows – does not exist as a cohesive product. The pieces exist in various EHR modules and point solutions but the integrated system of record for the physician enterprise is a gap. And the data network effect is real: compensation benchmarking requires peer data, which makes the product gets better for every participant as the network grows. Workday and Oracle are the logical acquirers, alongside PE platforms building physician management infrastructure.

The operating room problem and why 65% block utilization is basically embezzlement

The OR is where health system finances live. Forty to sixty percent of margin, the five percent of physical space. And it is run with tools that were designed for documentation and scheduling, not optimization. Block utilization nationally averages 65-70%. First case on-time start rates average below 60%. Implant costs vary by 40% between surgeons doing identical procedures with equivalent outcomes. The average health system leaves \$10-30 million in OR margin on the table annually through inefficiency that is predictable and therefore preventable.

The intelligence layer that would actually drive OR performance is a prediction problem. If you can accurately forecast which blocks are going to be underutilized before the day, which cases are likely to run long, and which implant choices are adding cost without improving outcomes, you can intervene proactively rather than generating reports about what went wrong. The difference between an analytics product and an operations product is that timing – analytics tells you what happens, operations changes what is going to happen.

OR data is also the most structured and accessible clinical data in any health system. Case times, block schedules, surgeon preference cards, and supply utilization are clean exports from OR management systems. An MVP does not require EHR integration. Historical case data from multiple systems trains prediction models in weeks, not months. The business case is immediate: a 400-bed hospital that improves block utilization from 67% to 75% adds roughly \$8-12 million in OR margin annually on cases that were already scheduled. The product pays for itself in weeks at almost any reasonable pricing.

The strategic acquirers in perioperative intelligence are interesting – Vizient, GE Healthcare, Intuitive Surgical expanding beyond robotics, Epic rounding out perioperative workflow. What makes this category particularly attractive is that the data moat compounds quickly. Training prediction models across multiple OR

environments simultaneously generates accuracy that a cold-start competitor would need years to replicate, because surgical preference patterns and case time variations are highly specific to the combination of surgeon, procedure type, facility, and patient population in ways that synthetic data cannot approximate.

Prior authorization as a regulatory forcing function

CMS has mandated prior authorization API compliance from payers by January 2026. This is one of those rare regulatory moments where the forcing function is real enough to compress enterprise sales cycles dramatically. You do not need to convince a health system leader that prior authorization is a problem – the AMA estimates costs health systems \$13 billion annually in administrative burden, physicians spend 16 hours per week per physician on PA, and denial rates have increased 56% over the last decade. The problem is universally understood. The question is whether the deadline is real enough to change procurement behavior, and it is.

The product architecture for prior auth automation has two sides. The provider-facing product automates submission, predicts denial probability before submission, generates supporting clinical documentation automatically, and manages appeals with minimal human involvement. The denial probability scoring feature deserves particular emphasis because it changes the workflow in a way that nothing else can. If a physician can see before submission that a specific request has a 73% predicted denial probability and here is the additional clinical documentation that would reduce that to 18%, you have intervened at the point where intervention actually matters rather than managing appeals after the fact. That feature requires training data good enough to build a reliable prediction model, which is why health system partnerships are the moat.

The payer-facing product is where the business gets dramatically larger, but it requires the provider-side credibility first. Payers pay for AI-driven clinical review that applies evidence-based criteria consistently rather than relying on offshore reviewers following rigid scripts. The strategic acquirers are Availity, Optum, E1

and major payers that want to own the PA infrastructure rather than buy access. Cohere Health raised \$50 million and sold for \$400 million. Build this with better training data, built-in distribution through founding health system relationships, and a regulatory tailwind they did not have.

Clinical variation and what two surgeons doing the same knee replacement tells you about margin

Clinical variation intelligence is one of the most politically sensitive and financially significant opportunities in the health system stack. Two orthopedic surgeons doing the same knee replacement might use implants that cost \$3,000 versus \$8,000 with identical outcomes. One hospitalist discharges pneumonia patients in 3 days on average, another in 5 days, with the same readmission rates. These variations are known anecdotally inside health systems and almost never addressed systematically because the data is not organized in a way that enables productive conversations between physicians.

The product design insight here is that punitive dashboards do not work. Physicians are not hourly workers who respond to performance management pressure in predictable ways – they are highly trained professionals with significant institutional leverage who will disengage from, work around, or actively resist tools that feel like surveillance. What works is peer comparison with outcome context in a format designed for physician-to-physician conversation. The department chair who can show a colleague that the top quartile of orthopedic surgeons in the group achieve equivalent outcomes at 22% lower supply cost, and here is specifically what is different about their preference cards, is having a productive clinical conversation. The administrator who publishes a cost ranking is creating a grievance.

Building this correctly requires understanding physician psychology as well as data architecture, which means the design partners have to be physician leaders rather than IT or revenue cycle leadership. The data network effect compounds with every health system that participates because compensation benchmarking and outcomes

comparison require peer data at scale to be statistically meaningful. Vizient and Premier are the natural acquirers, alongside healthcare analytics platforms that outcome-contextualized variation data as a differentiator.

Referral leakage is adjacent and equally under-addressed. Health systems lose 20% of their referral revenue to out-of-network providers – patients referred to specialists who are not in the employed or affiliated network. For a billion-dollar health system, that is \$50-100 million in lost downstream revenue annually. Every health system has this problem. Almost none of them have software to solve it.

The data to solve it already exists in the EHR and in claims data – referral orders, specialist scheduling, downstream claims. The gap is that nobody has organized an actionable layer. The MVP is a leakage attribution dashboard and a PCP-level score that answers two questions: where is referral revenue going and which primary care physicians are the highest leakage drivers. The intervention that actually changes behavior – surfacing the right in-network specialist at the moment of referral, with wait times, outcomes data, and patient satisfaction scores, so that the in-network choice is the easiest choice rather than requiring directory navigation – is the second version. You need the attribution intelligence first to create urgency for the intervention.

The data monetization play hiding in plain sight

Real-world clinical data – structured EHR data, imaging, pathology, genomics – is the most valuable commodity in life sciences and the least efficiently monetized asset most health systems own. Pharma companies spend billions on synthetic data, clinical proxies, and limited real-world datasets because the actual clinical data is fragmented across thousands of institutions with no consistent access model. AI foundation model developers need high-quality clinical data at scale and are paying for it wherever they can access it.

The business model for clinical data monetization is licensing – structured access to de-identified longitudinal clinical records for pharma research, clinical trial

recruitment, pharmacovigilance, and AI model training. The governance complexities: patient consent frameworks, de-identification validation, BAA structures, and IRB considerations all require legal infrastructure that takes time to build correctly. But the revenue potential is extraordinary. Data licensing revenue from pharmaceutical large enough dataset runs into the hundreds of millions annually, and AI model training contracts add to that.

The execution insight is that the data catalog and access portal are the actual product, not the data itself. Pharma companies and AI developers need to understand what is available before they can structure a research contract – patient volume, condition mix, data completeness by field, longitudinal depth, geographic diversity. Building that catalog with enough specificity that a research team at a pharma company can structure a meaningful dataset request, and then handling the de-identification and delivery pipeline in a standardized way that passes data governance scrutiny, is where most of the technical work lives. The first revenue comes from direct research contracts, probably \$2-5 million each, before anything resembling a self-service platform exists.

The federated compute model is an interesting alternative structure that avoids some of the regulatory friction around raw data commercialization. Rather than exporting datasets, you build secure enclaves where external entities bring algorithms to the data rather than the reverse. Pharma companies and AI developers run approved analytics inside the compute environment and pay for the results, not the underlying data. This positions the platform as a neutral governance-backed compute utility rather than a data seller, which changes the political dynamics inside health systems significantly and may be the right architecture for the long term even if direct licensing generates faster initial revenue.

Patient financial navigation deserves mention in the context of data-driven operations because the moat is similar – proprietary training data from billing and collection history that produces propensity-to-pay models more accurate than anything built from purchased credit data. A patient who just received a cancer diagnosis has completely different payment behavior than their FICO score would suggest. A patient who has been coming to the same health system for 15 years has a payment history that is

more predictive than generic credit bureau data. The product is a segmentation that routes each patient to the appropriate pathway – charity care, payment plan financing, or collections – automatically rather than running everyone through the same process. The commercial model on percentage of incremental collections makes this an extremely easy internal sale because the alternative is leaving money on the floor.

Where to start and why sequencing is the whole game

The sequencing logic for building in the health system stack is governed by two variables: how quickly does the product generate revenue, and how much does it strengthen the data foundation for the companies that come after it. These are not always the same product.

The fastest path to cash is the cheapest technical build with the most quantifiable ROI and the shortest sales cycles. Payer contract intelligence and purchased services benchmarking are the canonical examples. The MVP for payer contract intelligence is a contract term extraction engine, a rate benchmarking database, and an underpayment monitoring dashboard. The technical build is 10-12 weeks at \$140 thousand. The pilot pricing is \$500 thousand to \$1 million per health system and the product at launch says here is the \$8-15 million you are currently leaving on the table, and here is the negotiation intelligence that will help you recover it at your renewal. That conversation is extremely short.

Purchased services benchmarking is even simpler technically. Hospitals spend 30% of non-labor costs on purchased services – food service, equipment leases, IT consulting, facilities contractors – with almost no visibility into whether they are paying market rates. The product is a spend audit tool that ingests AP data, categorizes it against a purchased services taxonomy, and benchmarks each category against peer spend. The first deliverable to each participating health system is a report showing the categories where they are paying above median and the estimated savings opportunity. That report is the sales motion. It often closes before a dash

has been built. Cost to MVP is \$80-120 thousand, timeline 6-8 weeks, and it is genuinely the fastest path to studio cash flow and proof of concept for a broader portfolio.

Workforce intelligence and perioperative intelligence are the next wave – slightly more complex data collection requirements but equally quantifiable ROI and accessible data. OR data is already structured and available as clean exports. Historical census and scheduling data requires 24 months of exports from each system, which is an operational ask but not a technical one. These build to \$10 million plus ARR businesses within 12 months if execution is clean.

Virtual card AP automation is the highest-revenue Wave 1 opportunity by a significant margin if vendor enrollment executes. The economics are straightforward at \$500 million in converted AP spend per health system at 1.5% interchange with 25% revenue share, a network of health systems generates tens of millions in annual revenue from interchange alone before adding a single outside customer. The vendor enrollment operation – convincing suppliers to convert from check or ACH to card acceptance – is the real product, not the technology. The leverage comes from representing enough combined purchasing volume that the conversation with any given vendor is categorically different from what a generic fintech can have.

Prior auth, clinical documentation improvement, referral intelligence, and patient financial navigation come next because they require EHR integration that adds 60 days to timelines regardless of technical readiness – health system IT review cycles are what they are. Starting those IT relationship conversations during the first wave so the technical builds run in parallel is how you compress the overall timeline.

Data monetization, cybersecurity, and marketplace businesses come last because they require legal infrastructure, institutional credibility, or operational scale that the earlier companies build. Data monetization specifically needs BAA structures and identification validation across multiple health systems that are expensive and time-consuming to build correctly but not technically complex. The institutional credibility that comes from having deployed multiple products inside a health system

network makes those legal conversations significantly easier than approaching the cold.

The aggregate math on this sequencing is compelling. The cheapest Wave 1 build payer contract intelligence, purchased services, workforce intelligence, perioperative intelligence, and the AP automation business development operation – can be capitalized for under a million dollars and generate \$60-100 million in ARR across a portfolio within 18 months if execution is disciplined. The Wave 2 builds add ceiling, particularly prior auth and data monetization, which are the businesses most likely to generate eight- and nine-figure exits.

The common thread across all of these is that the entry point is a CFO conversation about quantifiable revenue loss or cost reduction, not a clinical champion selling upward through a committee. Health systems are under enough financial pressure right now that a product demonstrating \$10 million in recoverable revenue gets an audience quickly. The data network effects that make these businesses defensible compound faster than in most enterprise verticals because health system data is scarce and extraordinarily valuable, meaning the benchmarking advantage of a network system is very difficult for a single-system point solution to replicate. Build the network first, then build the products on top of it. The sequencing is the strategy.



1 Like

← Previous

Next →

Discussion about this post

Comments

Restacks



Write a comment...

