The 340B Software Stack: The Next Healthcare SaaS Vertical

MAR 15, 2026 • PAID

Table of Contents

Abstract

- 340B Drug Pricing Program generates an estimated $44-54B in covered entity savings annually (HRSA 2023 data, various estimates)

- Program complexity has quietly produced a dedicated SaaS vertical analogous t early revenue cycle management software

- Six discrete software categories have emerged, each representing a standalone venture opportunity

- Manufacturer restrictions since 2020 have massively accelerated software dema across the stack

- The vertical is fragmented, underleveraged on AI/ML, and early in consolidatic

How 340B Became a Software Problem

The 340B Drug Pricing Program was signed into law in 1992 as part of the Vete Health Care Act. The original intent was simple enough – qualifying covered en mostly safety-net hospitals and federally qualified health centers, could purchas outpatient drugs at significantly reduced prices, with the spread theoretically subsidizing care for low-income and uninsured patients. For about the first fifte years of the program's life, the operational complexity was manageable. Covered entities had relatively limited formularies, contract pharmacy arrangements wer and the program's administrative footprint was small enough that a decent phar and a spreadsheet could handle most of the compliance work.

That era ended somewhere around 2010, and the catalysts were structural. The Affordable Care Act dramatically expanded the universe of eligible covered enti and, more importantly, exploded the volume of eligible patients moving through qualifying facilities. At the same time, the commercial specialty pharmacy mark exploding. Biologic therapies, oncology agents, and specialty injectables – drugs carrying average wholesale prices measured in thousands of dollars per unit – w becoming a larger share of total drug spend. The intersection of those two dyna turned 340B from a modest safety-net benefit into one of the largest drug pricin mechanisms in the U.S. system. By the mid-2010s, covered entities were purchas somewhere between 5% and 6% of all outpatient drugs in the country through 3 channels. That share has since grown.

The contract pharmacy expansion was the decisive inflection point from a softwa perspective. HRSA's 2010 guidance allowed covered entities to use an essentially unlimited number of contract pharmacy arrangements, meaning a qualifying hos could register retail pharmacies across a geographic footprint as dispensing site 340B-purchased drugs. That ruling turned the program into a distributed financ network almost overnight. A mid-sized academic medical center might now man hundreds of contract pharmacy locations, each requiring patient eligibility

verification, prescription attribution, drug replenishment tracking, and split bil reconciliation. The data flows involved in running that network properly at scale genuinely complex – on par with the transaction processing infrastructure behir regional health plan.

The analogy to revenue cycle is not cosmetic. Revenue cycle management softwa emerged because billing and collections in healthcare became too operationally complex to handle through manual processes or generic accounting tools. The sa dynamic is now playing out in 340B. The program's rules are byzantine, the data requirements are substantial, and the financial stakes are high enough that erro expensive in both directions. An eligibility misclassification that allows a non-qualifying prescription to accumulate 340B savings is a compliance liability. A reconciliation failure that lets contract pharmacy claims go unmatched means le real money on the table. The covered entity operating a serious 340B program to needs purpose-built tooling, and a generation of startups has been building it.

The Eligibility Engine Layer

Patient eligibility is the foundational technical problem in 340B, and it is harde it looks. The basic rule is that a 340B-purchased drug must be dispensed to a pa of the covered entity, meaning the patient must have received a qualifying clinic encounter at a registered outpatient facility. That sounds straightforward until y to implement it at scale across a health system with dozens of outpatient sites, multiple EHR instances, high patient volume, and a contract pharmacy network spread across a metro area.

The core challenge is real-time or near-real-time eligibility determination at the of prescription. When a patient leaves a qualifying outpatient encounter and tak prescription to one of the covered entity's contract pharmacy locations, the phar needs to know whether that prescription qualifies for 340B accumulation before dispensing. The window is tight – the determination ideally happens during the adjudication workflow, before the pharmacist processes the claim. Getting that determination right requires pulling the relevant encounter data out of the cove

entity's clinical systems, confirming the prescribing provider is a qualifying prov at a registered site, validating the prescription is for a covered outpatient drug, a passing that determination back to the pharmacy in time to influence the transa

Early 340B eligibility solutions were mostly rules-based batch processors. A nigl extract from the covered entity's EHR would generate a list of qualifying encoun that list would be pushed to the pharmacy network, and pharmacists would chec patients against it. The latency was tolerable when contract pharmacy volumes w modest, but at scale it creates meaningful attribution failures. Patients who had qualifying encounters that morning but weren't in the last batch export get miss Prescriptions written at the encounter itself are particularly vulnerable to this la

The more sophisticated eligibility engines now operating in the market are doin event-driven eligibility streaming – pulling ADT feeds, appointment completion events, and clinical documentation signals in near real-time to maintain a continuously updated eligibility population. The better platforms are also doing probabilistic attribution, using historical patient behavior and prescribing patte improve match rates for prescriptions that arrive at contract pharmacies withou clean encounter-to-prescription linkage. That starts to look more like a machine learning problem than a rules engine problem, and the vendors who have investe that direction are pulling away from the legacy batch processors.

The startup opportunity here is not fully captured. Most of the current eligibilit platforms were built before health systems had anything like their current EHR integration surface area. FHIR API standardization, which was still largely aspirational five years ago, is now creating a meaningfully different integration landscape. A modern eligibility engine built natively on FHIR event streams, wit robust ML attribution layer and real-time pharmacy API integration, would be architecturally superior to most of what is currently deployed. The market is fragmented enough that a well-funded seed-stage company with strong health sy integration experience could compete effectively.

Split Billing and the Reconciliation Sta

Split billing is where the financial complexity of 340B really concentrates, and it almost certainly the largest software opportunity in the vertical. The term refers the operational requirement to correctly separate 340B-eligible claims from non claims within a shared pharmacy dispensing environment. In a contract pharmac arrangement, the same physical pharmacy location is dispensing drugs for multi payors and multiple covered entities, some of which have 340B arrangements an some of which do not. The pharmacy needs to know, for each individual prescrip whether to replenish the dispensed drug through the covered entity's 340B acco through normal commercial channels. Get that wrong at scale and you either hav 340B inventory mixing violations or you are missing eligible replenishments.

The underlying data problem is a claim-level matching exercise across multiple systems that do not natively talk to each other. The pharmacy's dispensing syste generates a claim record. The covered entity's EHR has the encounter record. Th 340B administrator or third-party claims processor has the eligibility determina The covered entity's pharmacy benefit structure may involve a PBM layer with it claims data. Correctly attributing each dispensed prescription to the right buck 340B or non-340B, which covered entity if there are multiple, which cost center within the health system – requires joining data across all of those systems at th transaction level. For a health system running meaningful contract pharmacy vo that is millions of claim-level reconciliations per year.

The reconciliation software market grew up largely around a handful of third-pa administrators and pharmacy service administrators who were doing this work manually or semi-manually, and then building software to systematize their own processes. Companies like Sentry Data Systems, TPA+ (now part of various consolidation plays), and a handful of regional players accumulated large covere entity customer bases by being early to market with workable reconciliation too The limitation of most of that first-generation software is that it was built aroun reconciliation problem as it existed in 2012 – relatively limited contract pharma networks, less complex specialty drug volumes, and before manufacturers starte actively trying to disrupt the contract pharmacy channel.

The manufacturer restrictions that began rolling out starting around 2020 turne reconciliation from a difficult operational problem into a genuinely adversarial c problem. When drug manufacturers – starting with AstraZeneca, Eli Lilly, and t rapidly expanding list of others – began limiting 340B pricing to a single contra pharmacy per covered entity or imposing data reporting requirements as a condi of providing 340B pricing on contract pharmacy claims, they created a new layer system complexity. The covered entity now needs to track, manufacturer by manufacturer, which drugs are available at which contract pharmacy locations, u what conditions, and with what documentation requirements. The reconciliation software needs to incorporate a manufacturer restriction data model that is cons changing as manufacturers update their policies and as legal challenges work th the courts and HRSA administrative process.

That dynamic created a genuine greenfield opportunity in what might be called manufacturer restriction intelligence – software that maintains a live database o manufacturer-specific contract pharmacy policies, maps those policies against a covered entity's specific network configuration, and flags compliance exposure o missed replenishment opportunities. A few startups have moved into this space, remains underbuilt relative to the operational need. The manufacturers are not making their restriction policies easy to track programmatically, which means th data aggregation work required to build a comprehensive restriction intelligenc is substantial – but that also makes it a durable competitive moat for whoever bu it well.

Manufacturer Compliance and the Dispute Layer

The adversarial dynamic between manufacturers and covered entities has create distinct software category that did not meaningfully exist five years ago. Call it manufacturer compliance and dispute management. The basic problem is that manufacturers who have imposed contract pharmacy restrictions or data submis requirements have also, in many cases, been refusing to honor 340B pricing on c that they allege do not comply with their stated policies. That creates a formal d

process – the covered entity believes it has a valid 340B claim, the manufacturer disputes it, and there is a resolution mechanism that involves documented evide HRSA involvement in some cases, and legal proceedings in others.

The scale of this dispute activity has grown substantially. HRSA's 340B dispute resolution process, which was designed as a backstop for the occasional disagree is now handling a volume of manufacturer-covered entity disputes that it was ne resourced for. The administrative dispute workload for a large covered entity wi sophisticated contract pharmacy network and exposure to multiple manufacture with restriction policies can be significant – tracking claim-level disputes, assem the supporting documentation, managing the timelines, and pursuing resolution through the appropriate channels.

The software opportunity here is essentially a legal operations and dispute management platform purpose-built for 340B. The commercial legal tech marke built sophisticated tools for tracking litigation, managing document collections, monitoring regulatory proceedings, but none of those generic tools are configure the specific workflow of a 340B compliance dispute. A purpose-built platform w need to integrate with the covered entity's claims data to identify potentially dis claims, pull the relevant supporting documentation, track manufacturer-specific dispute policies and resolution pathways, and give compliance teams a single workbench for managing the full dispute portfolio.

There is also a related opportunity in what might be called proactive compliance documentation – software that helps covered entities build and maintain the evidentiary record to defend 340B claims before a dispute arises, rather than scrambling to assemble documentation after a manufacturer challenges a claim. analogy here is to the audit trail and documentation tools that evolved in the rev cycle space around RAC audit defense. The covered entities that built good documentation practices ahead of audit pressure consistently outperformed thos did not, and the 340B compliance environment is moving in a similar direction.

Specialty Pharmacy Optimization Software

Hospital-owned specialty pharmacy is one of the larger strategic moves a health system can make inside the 340B framework, and it has its own distinct software stack. When a covered entity establishes an owned specialty pharmacy rather tha relying entirely on contract pharmacy arrangements, it internalizes the margin t would otherwise be split with a retail chain or pharmacy service administrator. I gets much cleaner data – the covered entity controls both the prescribing encou and the dispensing transaction, which simplifies eligibility determination and reconciliation substantially. The tradeoff is operational complexity. Running a specialty pharmacy is a different business than running a hospital, and the softw infrastructure required to operate one well – formulary management, prior authorization, specialty drug logistics, patient adherence programs, payer contra – is substantial.

The specialty pharmacy software market is not new, but the 340B context create specific optimization opportunities that generic specialty pharmacy platforms d address well. The most important is formulary and channel optimization – determining, for each specialty drug and each patient, which dispensing channel maximizes the covered entity's 340B capture while also satisfying payer requiren and patient access needs. That optimization problem involves the drug's manufa restriction status, the patient's insurance coverage and prior authorization requirements, the health system's payer contracts, and the available dispensing infrastructure. It is genuinely a multi-variable optimization that most health sys are solving manually or with crude tools.

A few startups have built point solutions around pieces of this optimization pro – specialty pharmacy prior authorization automation, 340B formulary manageme patient assistance program coordination – but the integrated optimization layer mostly unbuilt. The vision here would be a platform that sits across a health syst specialty pharmacy operations and continuously optimizes channel and formula decisions at the patient level, incorporating real-time data on manufacturer

restrictions, payer rules, and the health system's own financial objectives. That platform would look architecturally similar to some of the more sophisticated re integrity tools that have emerged in the revenue cycle space – decision support systems that help health systems maximize legitimate revenue within a complex environment.

The Platform Play and What Comes Ne

The five software categories described above – eligibility, split billing reconcilia manufacturer compliance, dispute management, and specialty pharmacy optimiz – are currently served by a fragmented market of point solution vendors, some division of a handful of larger pharmacy services companies, and a meaningful a of homegrown tooling inside larger health systems. The consolidation trajectory predictable. Revenue cycle started with the same fragmentation, and the market about fifteen years consolidating into a handful of scaled platforms before the m EHR vendors started absorbing pieces of the stack. 340B software is probably ei ten years behind that consolidation curve, which means the window for building platform-level company is open.

The platform thesis goes something like this: a company that builds excellent eligibility infrastructure has a natural expansion path into reconciliation, becaus eligibility data and the encounter-to-prescription matching logic are inputs to tl reconciliation workflow. A company that owns the reconciliation layer has natur leverage into manufacturer compliance and dispute management, because the cl level data it is already processing is the foundation of the compliance documenta A company that owns compliance documentation has a natural path into specialt pharmacy optimization, because it has the most complete picture of a covered er drug economics across channels. Each layer builds on the previous one, and the switching costs compound.

The AI/ML angle is real and currently underexploited. Most of the existing 340B software stack is rules-based – it is implementing HRSA guidance and covered e policy as logic trees, not learning from patterns in data. The opportunities for

machine learning are concentrated in eligibility attribution (probabilistic patien matching), reconciliation anomaly detection (identifying systematic misattributi patterns before they become compliance issues), and optimization (the multi-var channel and formulary optimization problem in specialty pharmacy). The vendo move first on AI-native architectures will likely find themselves with both bette products and meaningfully better unit economics, because the marginal cost of scaling an ML-based eligibility engine is much lower than scaling a rules-based

The regulatory risk is real and worth discussing plainly. 340B has been under sustained political and legal pressure for several years. Manufacturers have been winning some of their legal challenges to HRSA's contract pharmacy guidance. Congressional interest in program reform is biennial at minimum. A startup bui in this space needs a thesis for why the program's core economics survive meani regulatory change, which most sophisticated operators have – the covered entity is politically durable, the safety-net hospitals that depend on 340B revenue have strong congressional relationships, and the program has survived multiple refor attempts. But the specific contract pharmacy economics that drive much of the software demand are more vulnerable than the program itself. Building software whose value proposition survives a significant contraction in contract pharmacy arrangements – by focusing on owned pharmacy optimization, compliance infrastructure, or eligibility accuracy rather than contract pharmacy volume maximization – is the more durable bet.

The investor framework here should look familiar to anyone who watched the re cycle software market develop. Large, recurring financial flows. Complex regulat environment that creates demand for specialized tooling. Fragmented vendor landscape with no clear platform leader. Health system customers who are sophisticated enough to pay for good software but often still using inferior tools because switching costs feel high and procurement cycles are long. The category real, the timing is right, and the margin physics of 340B – a program generating of billions in covered entity savings annually – mean the software spend justified even modest improvements in capture rates or compliance risk reduction is substantial.

1 Like • 1 Restack

← Previous

Nex

Discussion about this post

Comments | Restacks

Write a comment...