

The New Price Transparency Stack and the Very Investable Plumbing Hiding Inside It

JAN 03, 2026



Share

Table of Contents

Why This Rule Exists and Why It Looks the Way It Does

What Actually Changes in the Public Files

The Quiet but Massive Shift to Network Level Truth

Why CMS Is Obsessed With Context All of a Sudden

Utilization Data as the Missing Multiplier

Taxonomy Disclosure and the End of Contracting Theater

Out of Network Data Finally Becomes Usable

Quarterly Cadence and the Death of Panic ETL

Findability as a Regulatory Weapon

The Phone Requirement and the Operational Reckoning

Costs, Burden, and What the Government Is Accidentally Telling Investors

The Market Level Aggregation Gambit

The Percentage of Billed Charges Cleanup

What This Means for Founders

Where Venture Capital Fits Cleanly

Where Private Equity Has the Edge

The Service Provider Play and Why It Matters

The Compliance Software Stack Nobody Sees Coming

Deep Analysis of Investment Opportunities by Vertical

How to Underwrite This Wave Without Getting Burned

Closing Thoughts on What Actually Gets Built

Abstract

This essay breaks down the latest Transparency in Coverage proposed rule and why it matters far more to investors and operators than most people realize. The rule is about consumer shopping tools, and it is not about ideology. It is about converting a chaotic flood of pricing disclosures into a structured, contextualized, and operational dataset that can actually be used. The government is forcing plans and issuers to publish network level negotiated rates, utilization context, adjudication logic, and explicit change tracking, while also making the data easier to find and slower to change. This combination quietly unlocks a new generation of infrastructure, analytics, and compliance driven business models. The real opportunity sits with enterprise buyers, not patients, and favors companies that reduce operational costs, create negotiating leverage, and turn regulatory burden into recurring software and services revenue. Venture and private equity both have room to win here, but only if they underwrite the boring parts correctly.

The proposed rule also creates explicit permission structures for third party aggregation at unprecedented scale, legitimizes taxonomy management as a regulated business process, and forces phone based disclosure that will drive massive call volume.

software spend. The government estimates over nine hundred million dollars in time implementation costs and nearly seventy million in annual ongoing costs. That is not a compliance headache. That is total addressable market spelled out in Federal Register pages.

Why This Rule Exists and Why It Looks the Way It Does

The easiest way to misunderstand this proposed rule is to think it is a policy statement about transparency as a moral good. It is not. It is a technical correction to a system that technically worked but practically failed. The first wave of Transparency in Coverage did exactly what it was supposed to do in the narrowest sense. It forced the release of negotiated rates, allowed amounts, and drug pricing data that had been public before. What it also did was create files so large, duplicative, and confusing that only a handful of well capitalized data engineering teams could even open them without lighting money on fire.

CMS and the other agencies are not subtle about this in the proposal. They openly acknowledge that the in network rate files in particular became enormous because contracts enumerate every possible item and service for every provider regardless of whether that provider would ever be paid for that service. The result was petabyte scale data that mixed meaningful rates with nonsense combinations. Researchers complained. Engineers complained. Plans complained. Even the people publishing the data complained because storing and serving it was expensive and error prone.

The numbers tell the story. Some issuers are serving multiple terabytes per month for a single coverage entity. Individual files routinely exceed local storage and processing capabilities. The Departments conducted an internal analysis in 2024 sampling in network rate files market wide and found that eighty three percent of issuers were already using table of contents structures to reduce duplication, which means the industry has been screaming for relief through their implementation choices.

So this rule is not a philosophical pivot. It is an engineering intervention. The agencies are trying to shrink file size, reduce duplication, add context that

downstream users have been reverse engineering anyway, and align plan disclosure more closely with how hospital price transparency already works. The goal is not prettier files. The goal is to make the data usable enough that market pressure can actually happen.

That framing matters for investors because it tells you what kind of companies you want. This is not a moment for glossy front ends and consumer delight decks. This is a moment for plumbing, tooling, and operational software that assumes pricing data exists and focuses on making it reliable, explainable, and actionable.

What Actually Changes in the Public Financial

At a high level, the rule does five big things. It reorganizes in-network rates around provider networks instead of duplicating the same rates across every plan that uses them. It requires new contextual machine-readable files that explain what changed, what was actually used, and how adjudication logic works. It expands and stabilizes out-of-network allowed amount data so it is less sparse and more analyzable. It sets the update cadence for most files from monthly to quarterly. And it forces issuers to make the files easy to find by standardizing discovery.

Each of these changes sounds incremental. Together they change what kinds of products can be built without heroic effort.

On top of that, the rule also tightens participant-facing disclosure by explicitly requiring that cost-sharing estimates be available by phone, not just online or on paper, and clarifies that doing so satisfies the No Surprises Act price-comparison requirement. That piece is less about data and more about operations, but it is where a lot of real money will move.

The Quiet but Massive Shift to Network-Level Truth

One of the most consequential changes in the rule is also one of the least flashy. Network rate files will now be organized by provider network, not by plan or coverage.

option. If ten plans all use the same PPO network, the rates for that network get published once, not ten times.

This matters because network is the real unit of negotiation and decision making in commercial health care. Employers choose networks. Brokers sell networks. Consultants benchmark networks. Contracting teams negotiate networks. Plans historically hidden behind plan proliferation as a way to make comparison harder. This rule strips that away.

By forcing network level files and requiring the common network name to be disclosed, the rule creates a cleaner identity layer for pricing data. It becomes much easier to say this network pays this much for these services, weighted by how often they are used. That is exactly the framing employers and advisors want when they are trying to understand spend.

The technical implementation matters here. Plans and issuers will be required to identify for each provider network every coverage option that uses that network. This maintains the plan to rate connection but removes the duplication. A researcher analyzing a specific network can grab one file instead of dozens. An employer comparing two networks for next year's renewal can diff two files instead of hunting through hundreds.

There is also a subtle but important pricing representation change. In network rates must be expressed as dollar amounts except in the narrow case where the contract is explicitly a percentage of billed charges and cannot be translated into a dollar amount ahead of time. That pushes the data toward computable reality. Percent of billed charge contracts still exist, but they are increasingly treated as an edge case rather than the default.

The rule text is explicit that plans and issuers should define what constitutes a separate provider network according to their current business practices. The Departments are not imposing a taxonomy of networks. They are forcing disclosure of the taxonomy plans already use internally. That is significant because it means that the data will reflect operational reality rather than a regulatory construct.

Why CMS Is Obsessed With Context All a Sudden

The biggest conceptual shift in the proposal is the idea that raw rates are not enough. CMS is now explicitly requiring context to travel alongside the numbers. That comes in three new machine readable files for each in network rate file.

The first is a change log file. This file identifies what changed since the last version of the in network file. From a developer perspective, this is a gift. Instead of re ingesting everything and computing diffs yourself, you get an official record of changes. From a compliance perspective, it creates a trail. Plans cannot quietly change rates without being obvious. That alone creates a market for monitoring and alerting.

The second is a utilization file. This file documents which covered items and services actually had claims submitted and reimbursed over a defined twelve month period ending six months before publication. It includes provider identifiers and place of service. This is the bridge between theoretical pricing and real spend. A negotiated rate that never gets used is trivia. A rate attached to high utilization services is a

The third is a taxonomy file. This file discloses the plan or issuer's internal provider taxonomy used in claims adjudication to decide whether a provider is appropriate for a given service. This is the logic that determines whether a claim gets paid or denied based on specialty. The rule then requires plans to use this same logic to exclude unlikely provider service combinations from the in network rate file.

This is a big deal. Plans already have this logic. It has just never been public. By forcing disclosure, the rule reduces garbage data and also exposes an internal data layer that has historically been opaque. That creates risk for plans, but it also creates an opportunity for vendors who can help manage, version, audit, and defend this logic.

The change log requirement becomes applicable on the first day of the calendar quarter following the date on which the first in network rate file is required to be posted, and updated and posted quarterly whether or not there are changes. The utilization file is required beginning on the first day of the calendar year quarter

following the applicability date and updated annually after the initial posting. The taxonomy file is required beginning on the first day of the calendar year quarter following the applicability date and updated and posted quarterly if changes to the internal provider taxonomy impact the information required in the in network rate file.

Utilization Data as the Missing Multiplier

If you talk to sophisticated employers or benefits consultants, the complaint about price transparency is always the same. Rates are interesting, but what matters is what people actually use. The utilization file directly addresses that.

By pairing rates with a standardized view of reimbursed services over a meaningful time window, the rule enables spend weighted analysis without stitching together external claims datasets. It does not replace full claims data, but it fills a gap that forced many analytics vendors to rely on proprietary or licensed data sources.

The utilization file must document for a twelve month period ending six months to publication all items and services for which a claim was submitted and reimbursed. It must include each in network provider identified by NPI, TIN, and place of service code who was reimbursed in whole or in part. This is not sample data. This is complete data for what actually happened.

For investors, the key point is that utilization context unlocks products that were previously too expensive or fragile to build. Network comparisons become more accurate. Contract optimization becomes less hypothetical. Benefit design models become more grounded. Even site of service analysis becomes easier when you know both the price and where care actually happened.

The proposal estimates the one time cost to build utilization files at over six hundred thirty eight million dollars across the industry. The annual ongoing cost is estimated at over nine million dollars. Those numbers represent the cost of compliance, but also represent the pain point that software can address. Any vendor who can reduce

the cost of generating, validating, or maintaining utilization files is selling into a market where the baseline cost is known and published.

Taxonomy Disclosure and the End of Contracting Theater

The taxonomy requirement deserves special attention because it will create friction inside plans and issuers. Internal provider taxonomies are messy. They evolve. They encode business rules that were never designed to be public. By forcing disclosure and alignment between adjudication logic and published rates, the rule collapses a long standing gap between what contracts say and what actually gets paid.

The rule requires plans and issuers to publish a taxonomy file that includes their internal provider taxonomy matching items and services represented by billing codes with provider specialties represented by specialty codes derived from the Health Provider Taxonomy code set established by NUCC. This taxonomy is used to determine if the plan or issuer should deny reimbursement for an item or service because it was not furnished by a provider in an appropriate specialty.

This is where a new class of tooling becomes inevitable. Taxonomy management is something most organizations treat as a product. It lives in spreadsheets, legacy systems, and institutional memory. Once it becomes a compliance artifact that must be published, updated, and defended, it turns into software.

Expect to see products that treat taxonomy like code. Version control. Impact analysis. Testing against claims history. Audit trails. These are not sexy features, but they are the kind of features compliance buyers pay for when regulators start asking questions.

The proposal also requires plans and issuers to exclude from in-network rate files provider rate combinations for items or services where the provider is unlikely to be reimbursed given that provider's area of specialty according to the plan's or issuer's internal provider taxonomy. The one-time cost estimate for this exclusion logic is forty-two million dollars. That is the cost of implementing the filtering. The ongoing maintenance is embedded in the taxonomy file updates.

Out of Network Data Finally Becomes Usable

Out of network allowed amount files have always been the weakest part of Transparency in Coverage. Too many codes never crossed the threshold to be reported. The windows were too short. The data was too noisy.

The proposal tackles this head on. It aggregates allowed amounts and billed charges at the health insurance market level rather than the plan level. It lowers the claim threshold from twenty to eleven. It expands the reporting period from ninety days to six months and the lookback window from six months to nine.

Each of these changes increases data density. Together they dramatically increase the likelihood that a given service appears in the file. That makes the data more useful for benchmarking, negotiation, and modeling out of network exposure.

This is not just academic. Out of network spend is still a meaningful driver of employer cost and member dissatisfaction. More stable benchmarks make it easier to design benefits, negotiate contracts, and evaluate vendor performance.

The market level aggregation is particularly clever. For self insured group health plans, health insurance market means all self insured group health plans maintained by the plan sponsor. For fully insured plans, it means the individual market, the group market, or the small group market as defined in existing regulations. This creates natural aggregation pools that are big enough to hit the eleven claim threshold more consistently but small enough to preserve price signal by market segment.

The rule also permits self insured group health plans under certain circumstances to allow another party such as a service provider to aggregate allowed amount files from more than one self insured group health plan including those offered by different sponsors. This is the explicit permission structure that turns TPAs and ASOs into disclosure platforms.

Quarterly Cadence and the Death of Panic ETL

Another underappreciated change is the move from monthly to quarterly update in network and allowed amount files. Monthly cadence sounded good in theory. In practice it created constant ingestion pressure and made it hard for downstream to finish analysis before the next update landed.

Quarterly cadence aligns better with how contracting and budgeting actually work. It reduces compute costs. It reduces storage and egress costs. It also reduces the likelihood of errors introduced by rushed updates.

For software companies, this changes product design. You can build around quarterly change instead of constantly chasing the latest file. That makes analytics more stable and products easier to explain to buyers.

The proposal explicitly notes that plans, issuers, and researchers have indicated since provider networks and rates do not change significantly from month to month, switching to a quarterly reporting cadence would not lead to a significant reduction in meaningful data. This reduced reporting cadence may also provide more time to analyze the data, as some file users have informed the Departments that they have difficulty keeping up with the pace of downloading and ingesting the file data monthly.

The benefits section of the proposal estimates annual cost savings of over two hundred fifty seven million dollars from reduced data cleaning, storage, and network egress costs driven by the combination of file size reduction and quarterly cadence. That number is split between plans, issuers, third party developers, and other file users. Those savings become margin expansion for compliance vendors and efficiency gains for analytics platforms.

Findability as a Regulatory Weapon

One of the most practical parts of the rule is also one of the most impactful. Plan issuers will be required to publish a simple text file in the root of their website that points to the location of the machine readable files and names a contact person. They will also be required to include a standardized footer link labeled Price Transparency or Transparency in Coverage.

This does two things. It makes it trivial for crawlers to find the files without custom scraping logic. And it assigns responsibility. When there is a named contact, errors get reported and fixed faster.

The text file must be in a dot txt format located in the root folder of the plan or issuer's website with information on the specific location of the machine readable files as well as contact information including a name and email address for those responsible for the files. The footer link must route directly to the publicly available web page that hosts the machine readable files.

From an investor perspective, this commoditizes basic ingestion. If your moat is that you can find the files, your moat is gone. If your value is what you do once you have them, this rule helps you.

The text file and footer requirements apply to all machine readable files including the prescription drug file. The text file must be posted beginning on the first day of each calendar year quarter following the applicability date and updated and posted as soon as practicable but no later than seven calendar days following a change in any of the required information.

The Phone Requirement and the Operational Reckoning

The participant facing side of the rule is where operations meet regulation. Plan issuers will be required to provide cost sharing estimates and related disclosures by phone, using the same customer assistance number that appears on ID cards. This information must be accurate and provided at the time of the request.

This is not a small ask. Many existing tools are estimation engines with disclaimers. Phone delivery implies real time quoting and accountability. It also satisfies the Surprise Act price comparison requirement, including for grandfathered plans that were otherwise exempt.

The rule proposes to allow plans and issuers to limit the number of providers with respect to which cost sharing information for covered items and services is provided to no fewer than twenty providers per day and to require plans and issuers to disclose the applicable provider per day limit to the participant, beneficiary, or enrollee when the request for information is made. This mirrors the existing limitation for paper requests.

This requirement will drive spending on call center tooling, agent assist software, workflow automation, and quality assurance. It creates an opportunity for vendors who can reduce handle time and error rates while creating compliance logs.

The proposal estimates the one time training cost for customer service representatives and supervisors at over thirteen million dollars. That is just training. The ongoing annual cost of providing phone based disclosure is estimated at over fifty two million dollars. Those are operational costs that software can compress.

The applicability date for the phone requirement is for plan years beginning on or after January 1, 2027. That gives the industry roughly two years from when the rule is finalized to build out the infrastructure. For vendors selling into this space, the clock is already ticking.

Costs, Burden, and What the Government Is Accidentally Telling Investors

The proposal includes detailed cost and benefit estimates. One time costs are dominated by building utilization files and change logs. Ongoing annual costs are driven by participant disclosures, utilization file maintenance, and responding to inquiries.

The total one time cost across the industry is estimated at over nine hundred thousand million dollars. The total annual ongoing cost is estimated at over sixty eight million dollars. These are not rough guesses. These are line item estimates based on labor hours, systems development, and operational assumptions.

This is essentially a TAM estimate for vendors who can reduce these burdens. When the government says this will cost the industry hundreds of millions of dollars to implement, it is implicitly saying there is money to be saved by doing it better.

The benefits estimates focus on reduced data cleaning, storage, and discovery costs. That tells you where CMS thinks inefficiency lives today. Products that attack those inefficiencies directly are aligned with regulatory intent.

The cost estimates also reveal assumptions about how plans and issuers will compete. For example, the utilization file build cost assumes plans and issuers will need to develop new data pipelines to extract and format claims history. Any vendor who sells a pre built pipeline or a managed service is competing against that baseline cost.

The Market Level Aggregation Gambit

The rule creates explicit permission for third party aggregation at a scale that did not previously exist. Self insured group health plans can allow another party such as a service provider with which they have an agreement to aggregate allowed amounts for more than one self insured group health plan including those offered by different plan sponsors.

This is not just a technical detail. This is the Departments saying that TPAs, ASOs, and other service providers can act as disclosure platforms. They can aggregate claims across multiple employers, publish it once, and point all the individual plans to that aggregated file.

The same logic applies to in network rate files. Plans and issuers can allow another party to make available in a single in network rate file the information required

more than one plan, insurance policy, or contract including those offered by different plan sponsors across different health insurance markets.

This creates a natural consolidation point. Whoever operates the largest aggregated platform has the cleanest data, the most leverage with plans, and the best position to upsell adjacent services. This is where private equity can build roll up plays and where venture can fund platforms that aim to become the definitive source of truth.

The Percentage of Billed Charges Cleanup

The rule tightens the representation of negotiated rates. In network rates must be reflected as a dollar amount except for contractual arrangements under which plan and issuers agree to pay an in network provider a percentage of billed charges and are not able to assign a dollar amount to an item or service prior to a bill being generated.

This is a narrowing of when percentage of billed charges can be used. Previously percentage of billed charges was treated as an acceptable rate representation. Now there is explicitly an exception that only applies when a dollar amount cannot be assigned in advance.

This matters because percentage of billed charges is inherently less transparent. The final payment amount depends on what the provider bills, which is not disclosed in the transparency file. By limiting when this representation can be used, the rule pushes the industry toward more concrete pricing.

The one time cost estimate for implementing this requirement is over seven million dollars. That cost is driven by the need to convert existing percentage of billed charges contracts into dollar amounts where possible and to document why conversion is not possible where it is not.

What This Means for Founders

The biggest mistake founders can make in this space is to chase consumer behavior change. The smarter move is to sell to organizations that already feel the pain and have budgets.

There is room for pricing data observability platforms that validate files, track changes, and flag anomalies. There is room for utilization weighted network analytics that inform contracting and benefit design. There is room for taxonomy governance tools that turn a compliance headache into a managed process. There is room for center automation that turns a phone mandate into a cost reduction story.

The common thread is that these are enterprise products tied to operations, not consumer engagement.

The buyers are plans, issuers, TPAs, ASOs, benefits consultants, large employers and state regulators. These are entities with procurement processes, multi year contracts and willingness to pay for risk reduction. They are not consumers trying to save dollars on an MRI.

The wedge is compliance. The expansion is optimization. The endgame is becoming embedded infrastructure.

Where Venture Capital Fits Cleanly

Venture capital fits best where there is horizontal software leverage and the potential for platform expansion. Data infrastructure, observability, and analytics layers that can be reused across many customers are good candidates.

Products that ingest, validate, and normalize transparency files at scale can sell to plans, issuers, consultants, and regulators. Products that automate change detection and impact analysis can sell to anyone who needs to monitor pricing or compliance. Products that layer utilization data on top of negotiated rates can sell to anyone to model total cost of care.

The key is distribution. Products that wedge into benefits consultants, TPAs, or employers can scale. Products that require selling plan by plan with heavy

customization will struggle unless they price accordingly.

The venture opportunity is also in tooling for the service providers who are aggregating data. If TPAs and ASOs become disclosure platforms, they need software to manage ingestion, aggregation, versioning, and distribution. That software can be sold as a recurring subscription with expansion revenue tied to the number of plans and volume of data.

The timing matters. The rule has a twelve month implementation period from when it is finalized. Assuming finalization in late 2025, the first network level files will be due in late 2026 or early 2027. The first utilization files will be due shortly after. Vendors who can get to market before the compliance deadline will capture the initial wave of spend.

Where Private Equity Has the Edge

Private equity shines where services and software intersect. Managed compliance, transparency operations, and TPA adjacent platforms are ripe for roll up. The rule explicitly allows third parties to publish files on behalf of plans, including aggregating across multiple self insured plans. That legitimizes service provider disclosure operators.

Once a provider owns that role, upselling software becomes easier. Margins improve with automation. Contracts get sticky.

The roll up thesis is straightforward. Acquire regional TPAs or benefits administrators who are already managing transparency files for clients. Standardize the technology stack. Centralize data operations. Cross sell adjacent services like network analytics, contract benchmarking, and taxonomy management.

The rule also creates opportunity in the taxonomy and change log space. Plans and issuers will need help building, maintaining, and defending their internal taxonomies. That is a services business that can be productized over time. Start with consulting

taxonomy design. Sell managed taxonomy as a service. Build software to automate taxonomy updates and impact analysis. Roll it into the next acquisition.

Private equity can also target the call center software and services market. Plans issuers will need to stand up or expand phone based disclosure capabilities. This creates demand for agent assist tools, IVR systems, quality assurance platforms, outsourced call center services. Acquire a call center operator focused on health plans. Layer in software to reduce handle time and improve accuracy. Expand to other compliance disclosure requirements.

The Service Provider Play and Why It Matters

The rule creates explicit permission for service providers to act as disclosure platforms. Self insured group health plans can allow another party to aggregate and publish files on their behalf. That party can be a TPA, an ASO, a benefits consultant or any other entity with whom the plan has an agreement.

This is significant because it shifts the unit economics of compliance. Instead of plan building its own infrastructure, they can outsource to a platform that amortizes the cost across many clients. The platform gets recurring revenue. The plans get reduced operational burden. The regulators get better data quality because platforms have more resources to invest in validation and error correction.

The platform play is especially attractive in the self insured market. There are millions of self insured plans sponsored by employers of all sizes. Most of them use TPAs to handle claims administration. Those TPAs already have access to the claim data needed to build utilization files and the contracts needed to populate in-network rate files. They just need the software and processes to turn that data into compliance disclosures.

Whoever builds the best platform for TPAs wins. That platform needs to ingest claim history, map it to the required file schemas, generate change logs, apply taxonomy exclusions, and serve the files with the required metadata. It needs to handle

versioning, error correction, and inquiry response. It needs to scale to hundreds thousands of plans without requiring custom work for each one.

The Compliance Software Stack Nobody Sees Coming

The rule will create demand for an entire stack of compliance software that does not yet exist or exists only in fragmented form. At the bottom of the stack is file generation. Plans and issuers need tools to extract data from their systems, apply required transformations, and output compliant files. This is ETL for regulatory disclosure.

Above that is validation. Files need to be checked for schema compliance, logical consistency, and completeness. Errors need to be flagged and corrected before publication. This is where data quality platforms fit.

Above that is monitoring. Files need to be tracked over time. Changes need to be detected and explained. Anomalies need to be investigated. This is where observability platforms fit.

Above that is analytics. The data in the files needs to be turned into insights. Network performance needs to be benchmarked. Utilization trends need to be tracked. Contract opportunities need to be identified. This is where BI and analytics platforms fit.

And at the top of the stack is governance. Taxonomies need to be managed. Change logs need to be audited. Inquiries need to be responded to. This is where workflow and compliance management platforms fit.

Each layer creates a business. Each layer can be sold separately or bundled. Each has different buyers and different competitive dynamics. But they all share a common foundation, which is that the rule creates demand that did not exist before.

Deep Analysis of Investment Opportunities by Vertical

The taxonomy management vertical is particularly interesting because it sits at the intersection of compliance and operations. Plans and issuers have always had internal taxonomies, but they have never had to publish them or defend them. Now they have. That creates demand for tools that make taxonomies auditable, versionable, and testable.

A taxonomy management platform needs to import existing taxonomies from whatever format they currently live in, map them to the NUCC code set, identify gaps and inconsistencies, simulate the impact of changes, and generate compliant taxonomies. It also needs to track changes over time and provide an audit trail.

The market for this is every plan and issuer. The urgency is the compliance deadline. The expansion is helping plans optimize their taxonomies to reduce denials, improve provider satisfaction, and defend against audits.

The network analytics vertical is interesting because utilization data makes network analysis more accurate. Right now, most network analytics are based on contract rates without weighting for actual use. The utilization file changes that. Analytics platforms that can join negotiated rates with utilization data can produce spend-weighted benchmarks that are much more useful for contract negotiation and benefit design.

The market for this is employers, benefits consultants, and brokers. The pitch is better data leads to better decisions leads to lower costs. The pricing model is likely subscription with expansion based on number of employees or total spend under management.

The change detection and monitoring vertical is interesting because it turns a compliance requirement into an operational capability. Plans and issuers will publish change log files every quarter. Someone needs to ingest those files, compare them

prior versions, flag meaningful changes, and alert stakeholders. That someone could be internal staff, or it could be a monitoring platform.

The market for this is broad. Employers want to know when their network's rate change. Consultants want to know when their clients' competitors change rates. Regulators want to know when outlier changes happen. Providers want to know their contracted rates get published correctly. A monitoring platform that serves these constituencies at once has serious revenue potential.

The call center automation vertical is interesting because the phone requirement creates an immediate operational problem with a known cost. Plans and issuers estimate they will spend over thirteen million dollars on training and over fifty million annually on ongoing phone operations. Any vendor who can reduce hand time, improve accuracy, or automate responses is competing against that baseline.

The opportunity is agent assist software that sits alongside the existing cost sharing estimation tools and helps representatives navigate phone requests in real time. The software needs to pull the same data that powers the online tool, format it for voice delivery, handle the twenty provider per day limit, log the interaction for compliance, and do it fast enough that call times stay reasonable.

The more ambitious play is full automation through conversational AI. If the cost sharing estimation logic is already codified for the online tool, it can be wrapped in a voice interface. The challenge is handling edge cases, maintaining accuracy requirements, and building trust with callers who expect a human. But the economics are compelling if it works.

The utilization file generation vertical is interesting because the one time cost is six hundred million dollars and the annual maintenance cost is over nine million. The challenge is the cost of building custom pipelines to extract twelve months of claims history, filter to reimbursed claims only, group by item or service code, attach provider identifiers, and format for disclosure.

Most plans and issuers do not have this capability today. Claims data lives in adjudication systems that were not built for this kind of batch extraction. The d

needs to be joined across multiple tables, deduplicated, validated, and mapped to the required schema. That is classic data engineering work that can be productized.

A utilization file platform needs to connect to common claims platforms, handle extraction and transformation logic, apply the six month lag requirement, generate the files on the required annual cadence, and version them appropriately. The business is owned by anyone responsible for transparency compliance at a plan or issuer. The pricing is either per plan, per file, or per claim volume.

The file hosting and distribution vertical is less obvious but still real. Plans and issuers need to serve potentially massive files from publicly accessible URLs with high availability and low latency. The files need to be discoverable via the text file and footer link requirements. The hosting environment needs to handle traffic spikes when researchers or vendors scrape all the files at once.

Most plans and issuers are not set up for this. They are used to serving member portals and broker sites, not public data repositories. Standing up the infrastructure is a one time cost, but maintaining it is ongoing. A managed hosting service that handles file storage, distribution, metadata management, and access logging could be sold into plans and issuers who do not want to build this themselves.

The aggregation platform vertical is where the biggest outcome potential lives. The rule explicitly allows third parties to aggregate files across multiple plans. Whoever builds the dominant aggregation platform becomes the de facto source of truth for transparency data. They can monetize through subscriptions to plans for the compliance service, subscriptions to data users for access and analytics, and transaction fees for any marketplace features they layer on top.

The aggregation platform needs to ingest files from multiple sources, normalize to a common schema, deduplicate rates across plans that use the same networks, quality checks, version everything, and expose it through APIs and bulk downloads. It also needs to handle the change log aggregation problem, where changes from multiple underlying plans need to be rolled up into a coherent view.

The moat is data quality and coverage. The platform that has the cleanest data and most complete coverage becomes the default choice. Network effects kick in as more plans join because data users prefer platforms with broader coverage, and more users join because the platform has better data.

How to Underwrite This Wave Without Getting Burned

Underwriting in this space requires discipline. The first rule is to look for companies that treat the public data as an input, not gospel. The transparency files will have errors. Rates will be stale. Coverage will be incomplete. Companies that assume the data is perfect will build fragile products. Companies that assume the data is messy and build validation and correction into their workflows will build resilient products.

The second rule is to look for alignment with buyer incentives. The buyers here are not consumers. They are enterprises trying to reduce cost, manage risk, or comply with regulations. Products need to map to budget line items that already exist. Compliance software competes against legal and regulatory risk. Analytics compete against consulting fees. Automation competes against internal labor costs.

The third rule is to be wary of moats based on ingestion alone. The text file and requirements commoditize file discovery. The network level organization reduces the complexity of joining data across plans. The change log files reduce the need for custom diff logic. Ingestion used to be a meaningful technical challenge. Now it's a table stakes. The value has to be in what you do with the data once you have it.

The fourth rule is to pay attention to regulatory timing but not bet everything on the rules. This is a proposed rule with a sixty day comment period. The final rule could change. The implementation timeline could shift. The enforcement approach could evolve. Companies that are only viable if the rule lands exactly as proposed are taking on regulatory risk. Companies that solve problems that exist regardless of the specific rule text are more robust.

The fifth rule is to favor companies that reduce operational cost or create negoti leverage. Those are the budgets that renew. Employers will pay for tools that hel them negotiate better network rates. Plans will pay for tools that reduce the cost compliance. Consultants will pay for tools that make their analysis faster and m accurate. Those are recurring revenue streams tied to ongoing pain points.

The sixth rule is to watch for concentration risk. If a company's entire revenue n depends on selling to health plans and there are only a few hundred meaningful buyers in the country, customer concentration becomes an existential risk. Prodi that can sell to plans, issuers, employers, consultants, and service providers have diversified revenue streams.

The seventh rule is to be realistic about sales cycles and implementation timelin Enterprise software sales to health plans are slow. Compliance deadlines create urgency, but procurement processes do not move faster just because the deadlin tight. Companies need enough runway to survive twelve to eighteen month sales cycles and another six to twelve months of implementation before they see mean recurring revenue.

The eighth rule is to understand that this wave will have multiple phases. The fi phase is compliance. Companies sell software and services that help plans and is meet the basic requirements. The second phase is optimization. Companies sell analytics and insights that help buyers use the data to make better decisions. Th third phase is transformation. Companies sell platforms that fundamentally cha how healthcare pricing and contracting work. Most companies will only ever get phase one or two. The ones that get to phase three are the ones that become generational outcomes.

Closing Thoughts on What Actually Get Built

This proposed rule will not make patients perfect shoppers. It will not magically healthcare pricing. What it will do is force the industry to emit cleaner, contextu pricing telemetry that can be used by people who already care deeply about cost.

For investors and entrepreneurs, that is enough. The opportunity is not to change human behavior. It is to sell better tools to the humans already tasked with managing billions of dollars in spend.

The rule creates explicit permission structures for third party aggregation, mandates the disclosure of internal business logic that has never been public, establishes new operational requirements that will cost hundreds of millions to implement, and all of this on a timeline that gives vendors roughly eighteen months to get to market before the compliance deadline hits.

The government has essentially published a detailed TAM estimate, identified the specific pain points, estimated the baseline costs, and told the industry that third parties are allowed to build infrastructure to solve these problems at scale. That is a regulatory headache. That is an investor roadmap.

The companies that win will be the ones that understand this is fundamentally about reducing operational burden and creating decision leverage for sophisticated buyers, not about empowering consumers to shop for care. They will sell into existing business lines, integrate with existing workflows, and solve problems that renew annually.

The products that matter will be boring. File validation tools. Taxonomy management platforms. Change detection services. Utilization weighted network analytics. Content center automation. Aggregation infrastructure. These are not exciting demos. But they are what enterprises will actually buy.

The exits that work will be to buyers who care about infrastructure and data quality, not consumer engagement metrics. Strategic buyers will be larger health tech platforms that need pricing data as a feature. Financial buyers will be private equity firms that understand how to roll up fragmented service providers and cross sell software. The IPO path exists but requires getting to meaningful scale in a market with concentrated buyers and long sales cycles.

The timing is now. The rule is proposed but the direction is clear. The compliance deadline is coming. The industry is unprepared. The technology does not exist in productized form. The winners will be the companies that launch in the next two

months, sign their first customers in the next eighteen months, and have revenue scale before the compliance deadline hits.

This is not a consumer transparency play. This is an enterprise infrastructure play disguised as a transparency rule. The investors who see it clearly will fund the companies that build the plumbing. The ones who chase the consumer narrative fund companies that struggle to find buyers.

The price transparency stack is real. The investment opportunity is real. The question is who builds it and who funds them.



5 Likes

← Previous

Next

Discussion about this post

Comments

Restacks



Write a comment...