RECORDS RETRIEVAL · CASE STUDY

Fragmentation Has a Number

What 34,808 Records Requests Reveal About Fragmentation in Medical and Billing Records Retrieval

October 17, 2026



the Records company™

Dr. Grady Marin, The Records Company



“We count delay, because delay is easy to count. What fragmentation actually does is quieter and considerably worse. **It destroys information.**

Ask anyone what is wrong with health information and you will get the same word back. Fragmented. It is the most agreed-upon diagnosis in American healthcare and nobody has ever put a number on it.

This case study uses the complete operating records of two retrieval engagements covering 34,808 requests to ask a harder question: what does fragmentation actually cost when the unit being fragmented is a medical history?

We count delay, because delay is easy to count. What fragmentation actually does is quieter and considerably worse. It destroys information. It does so at a rate nobody has established, in a pattern that is almost certainly not random, and it leaves no trace, because the thing it destroys is the evidence that anything was destroyed.

Your medical history does not exist

Not in one place. Not anywhere. There is no file with your name on it.

Break your wrist tonight and by morning you will have manufactured five documents. A record at the hospital. A bill from the hospital. A separate bill from the emergency physician, who almost certainly does not work for the hospital. A report from the radiologist who read your film. And a bill from the company that does that radiologist's billing, which is a different company again, in a different building, possibly in a different state.

Five organizations. One night. None of them talks to the others. None of them holds a list of what the others have. A complicated case runs to forty.

There is no index. No master file. No number you can call to ask what your own medical history contains. Reassembling it is not clerical work, whatever the price sheet says. It is detective work, and everyone in this industry, mine emphatically included, staffs and prices it as filing. The cost of that shows up with unnerving regularity.

Boston book: request types, mean calendar days

	Requests	Mean wait (calendar days)	Share that hit a problem
Clinical records	3,501	33.2 days	32.4%
The hospital's bill	3,411	44.9 days	38.0%
The ER physician's bill	410	43.9 days	42.2%
The radiologist's bill	308	54.9 days	54.9%

Same patient. Same night, frequently the same building, occasionally the same hour. Fragmentation is not a metaphor. It behaves like a dose.

The cost nobody counts

Everything above is a story about time, and time is what this industry measures because time is what a clock does. It is also the least interesting thing fragmentation does to a record.

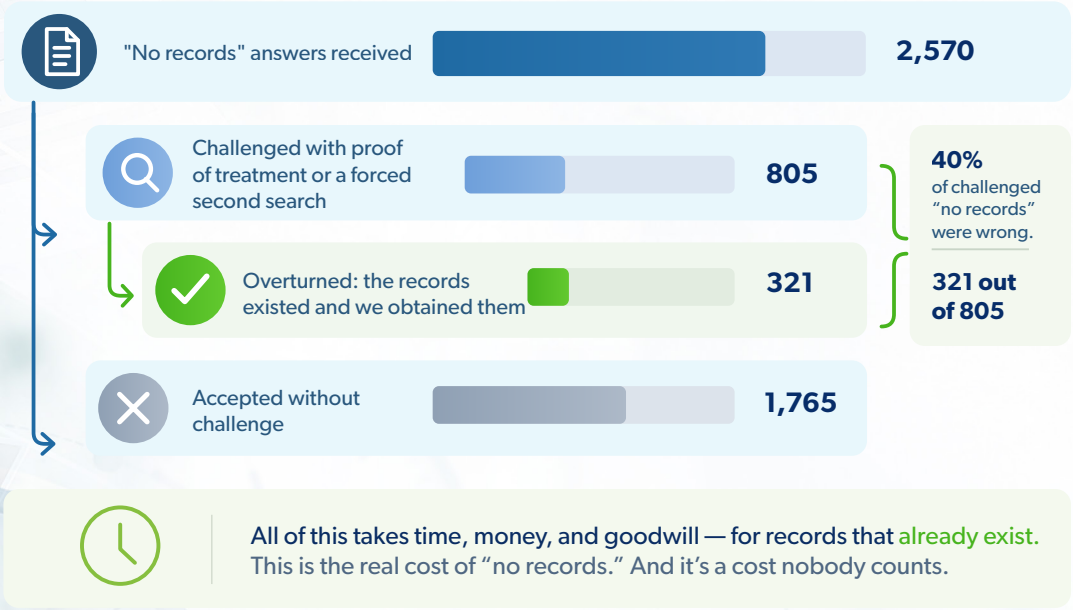
Here is the more serious consequence. Every organization standing between a question and an answer is an organization that can say no. And no is the one answer nobody can check.

Consider what a "no records" letter actually is. Somewhere, in a hospital or a billing office or a company you have never heard of, a person typed a name into a search box, did not see anything, and told you the thing you were looking for does not exist. That is the entire evidentiary basis.

You cannot inspect it. There is no document to review, nothing to hold up to the light. The only way to test it is to prove the treatment happened some other way and go back and make them look again, which costs time and goodwill and cannot be done for everything.

So it is done for almost nothing.

We did it anyway, across 27,111 requests placed for one national firm over seven months.



Two in five of the denials we tested were wrong.

The reasons, where a note recorded one, are gloriously mundane. Somebody had not read the request. Somebody searched a name spelled slightly differently. Somebody sent the wrong stack of paper and never went back. In one case our handler returned to clinical records the hospital had already produced, confirmed the imaging had in fact been performed, escalated, and the records that supposedly did not exist arrived.

Then there is the other column. We did not challenge 1,765 of them. Nobody did. Those files closed. Demand letters were written, claims were valued, cases settled or did not, and every one of those decisions rested on a medical history with a hole in it that no one involved knew was there.

How big is the hole

I cannot measure the false denial rate. I can bound it, and the bounds are worth putting on a table.

The floor is not arguable. Three hundred and twenty-one denials in this book are known false, because we went and got the records. That means at least 12.5 percent of every "no records" answer we received was wrong, and it is a floor rather than an estimate, because it counts only the ones we tested and won.

The ceiling needs an assumption, so here is the assumption in daylight. We challenged those 805 because something about them looked wrong, so their false rate of 39.9 percent should run higher than the rate among the 1,765 we let stand. How much higher is precisely what nobody knows. So take three cases.

False denial rate scenarios

IF THE FALSE RATE AMONG UNCHALLENGED DENIALS WERE…	ADDITIONAL FALSE DENIALS	TOTAL FALSE OF 2,570
A quarter of the challenged rate (10%)	~177	~498 (19%)
Half the challenged rate (20%)	~353	~674 (26%)
The same as the challenged rate (39.9%)	~704	~1,025 (40%)

A bounding exercise, not an estimate.
The middle column applies the assumed rate to 1,765 unchallenged denials.

Somewhere between one in five and two in five of the absences reported to us were probably not absences.

I would very much like someone to tell me that arithmetic is wrong. I have not managed it.

KEY TAKEAWAY

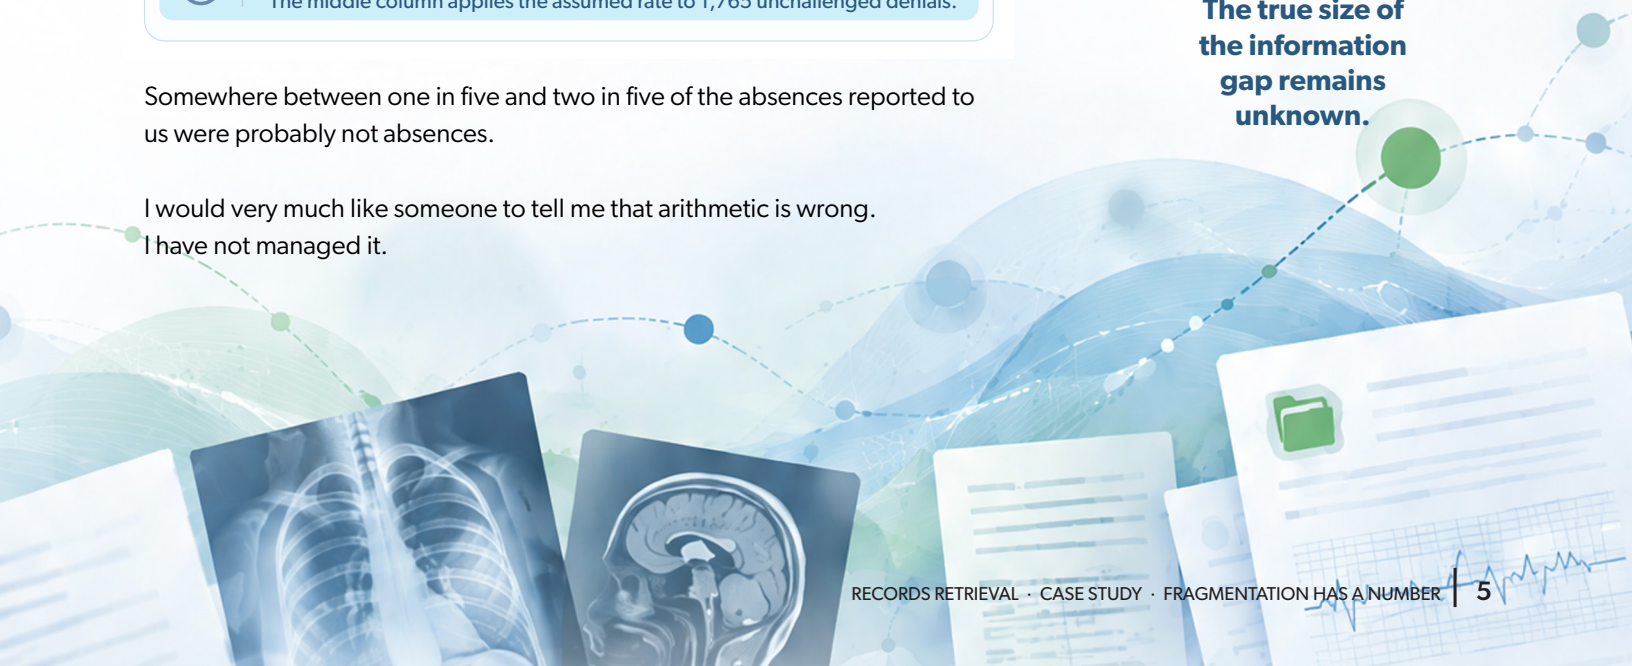
"No records" does not always mean no records.

When challenged,
321 of **805**
denials were overturned—showing that records reported as absent could, in fact, be found.



challenged denials were wrong.

The true size of the information gap remains unknown.



A prediction that would settle it

Put the two findings side by side and they imply something specific, which is the most useful thing a finding can do.

If difficulty rises with every organizational boundary a request crosses, and if denials are false partly because somebody searched carelessly or searched the wrong system, then falseness should rise with fragmentation too.

The prediction is not that distant custodians are slower. We have shown that. The prediction is that the further a record sits from the institution that delivered the care, the more likely a reported absence is to be false.

That is falsifiable, and it would matter enormously if it held, because it would mean the missing pieces are not scattered at random. They would cluster in exactly the corners of the record that are hardest to reach, and any dataset assembled from retrieved records would carry a truncation with a shape.

What would settle it is a cross-tabulation I have not run: false denial rate by record type and by custodian distance, with the challenge decision modeled rather than assumed away. The data sits in the operating record. The method for handling the selection problem is what I do not have. This is the question the industry still needs to answer.

Boston

The second kind of fragmentation happens within provider organizations. Boston offers a useful view of it: many custodians serving the same requesting firm within one regional market.

The first of the two books covers 7,697 requests placed between December 2022 and March 2026, reaching 1,979 patients and 1,334 provider organizations. Three quarters of the located volume sits in Massachusetts, and 3,351 requests went to Massachusetts custodians alone. Same firm, same kinds of records, same statute, same three and a half years.

Boston book: selected provider-name groups

Provider-name group	Requests	Mean calendar days	Met 30 days
Cambridge Health Alliance	245	24.3 days	74.7%
Boston Medical Center	600	23.3 days	75.7%
Boston Children's	79	26.2 days	68.4%
South Shore Health	132	35.4 days	68.9%
Signature Healthcare	90	38.4 days	48.9%
UMass Memorial	125	38.2 days	42.4%
Beth Israel Lahey Health	632	40.7 days	46.0%
Tufts Medicine	240	44.2 days	50.0%
Mass General Brigham	750	50.6 days	38.5%
Baystate Health	58	54.1 days	37.9%

Groups follow the source analysis's name-based classification, not a dated history of corporate ownership. The Boston Medical Center group includes 550 requests under the hospital's name and 50 under related BMC service labels. It excludes two records naming Berkshire Medical Center, a separate institution that also uses the abbreviation BMC.

Boston Medical Center and Mass General Brigham are two of the larger provider-name groups in this book. Across the 600 requests in the Boston Medical Center group, mean turnaround was 23.3 calendar days and 75.7 percent met the thirty-day benchmark. The Mass General Brigham group comprised 750 requests, with a 50.6-day mean and 38.5 percent meeting the benchmark. These are descriptive results for the groups shown, not a finding about which owner caused a delay.

Ownership changed during the study window. Carney Hospital closed on August 31, 2024. BMC Health System assumed operations of St. Elizabeth's Medical Center and Good Samaritan Medical Center on October 1, 2024. The Boston Medical Center group in this table contains no requests to either acquired hospital, before or after the handover. The table uses fixed provider-name groupings rather than ownership assigned at each request date; the former Steward group is not presented here.

The same institution, several strings

Hospital names in an export are not stable institution identifiers. A single hospital can appear under different spellings, abbreviations, punctuation, or encoded characters. The examples below show that data problem; they are not separate hospital performance scores.

Massachusetts General Hospital appears in our source data twice. Once as "Mass General Hospital," 118 requests averaging 48.9 days. Once as "Massachusetts General Hospital," 108 requests averaging 47.6 days. Same institution. Two rows.

Brigham and Women's main campus appears under multiple near-identical strings: "Brigham and Women's hospital," 54 requests at 60.7 days, and "Brigham and Womens Hospital," 59 requests at 52.5 days. Those two differ only by punctuation.

Beth Israel Deaconess Hospital in Milton appears twice, distinguished only by whether somebody typed an en dash or a hyphen. Forty-nine requests at 33.9 days, forty at 28.0.

The institution did not change; the string did. Until those variants are resolved, a scorecard can split one hospital into several apparently different performers. An abbreviation can create the opposite error by pulling two unrelated institutions into the same group.

Consolidating all of it reduced 1,368 raw provider names to 1,334 organizations across 74 clusters. Any scorecard built on the raw field, and most are, would be wrong about the largest teaching hospitals in this city.

KEY TAKEAWAY

- Provider fragmentation can exist within the same market
- Variation in naming can split one institution into multiple records.
- Neighborhood and cross-border differences affect request performance.
- The record path is inconsistent even when the institution is the same.

Neighborhoods

The variation runs below the institutional level too. Among locations with enough address data to examine, the slow end included Jamaica Plain (57 requests, 50.7-day mean), Lawrence (108, 48.2), Salem (50, 47.6), Newton (66, 44.2), Worcester (78, 41.2), and Somerville (62, 40.7). The faster end included Everett (57 requests, 22.6 days), Dorchester (231, 25.2), South Weymouth (65, 26.9), Cambridge (122, 30.8), Milton (61, 31.5), and Quincy (41, 32.4). These are locations, not all municipalities of the same type; Jamaica Plain and Dorchester are Boston neighborhoods.

The Boston location label covers 837 requests with a 38.2-day mean, close to the book average. The location figures describe where requests were sent; they do not isolate the effects of geography or provider size

.

Across the border

Massachusetts runs a 14.1 percent provider-caused problem rate on 3,351 requests. Connecticut runs 28.8 percent on 455, more than double, on an overall problem rate of 51.9 against Massachusetts' 35.2. New Hampshire sits between them at 19.7 percent on 365 requests, Maine at 23.3 on 163.

Two Connecticut observations illustrate why small samples should be treated as flags rather than verdicts. Stamford, at 52 requests, posts a 48.1 percent provider-caused rate, the highest location-level rate in this book. A separate Connecticut health system, on only 40 requests, met the thirty-day benchmark on one request in forty and showed a 65 percent cancellation rate. Those samples are too small to support a settled conclusion. They are worth investigating because they sit inside a broader Connecticut pattern, not because either number should be quoted as a definitive score.

What I cannot explain

I can measure all of this to a decimal place. I cannot tell you what causes a single point of it. The candidate explanations are internal and invisible from outside. How release of information is staffed. Whether requests are triaged or worked in the order they arrive. What gets escalated and what sits. Whether the function reports to health information management or to counsel. Whether anyone is watching the queue at all.

These figures do not isolate a cause: record mix, request timing, routing, and ownership history can differ. Twenty-three calendar days for one provider group, fifty-four for another. Same state. Same requester. That spread is the question.

Two ways a request dies

Just over a third of the Boston book, 2,809 requests, hit at least one problem. Where the fault landed matters more than how often it happened, because the two failure modes are not interchangeable.

2,809
requests hit at least one problem Just over a third of the Boston book

74.6 days
Average time for provider-caused failures Nearly three times the 27.2-day clean baseline

Provider failures are slow. They average 74.6 days, nearly three times the 27.2-day clean baseline, and only 18 percent clear thirty days. But four in five eventually deliver. The cost is time and patience.

Client-side failures are faster and far more lethal. They average 52.6 days and then 63 percent are abandoned. Nine hundred and sixty-four requests died after an average of 56 days of work that produced nothing whatsoever. Across the book, roughly 98,000 excess days accumulated against the clean baseline, and about 75,800 elapsed days went into requests that delivered no document at all.
Rank the damage by cause rather than by drama and one provider behavior tops the list. Requests that are simply never answered, coded as backlogged or unreachable, number 227, average 103.2 days, and account for 17,232 lost days. That single failure mode is nearly a fifth of all the time lost in the book. Certification missing follows at 10,515 days, non-receipt at 6,807.

The worst outcome profile of anything in either book belongs to a category with only 45 requests in it: "no client response." Average 103.3 days, 93 percent cancelled. Those are requests where we asked the firm a question and never got an answer.

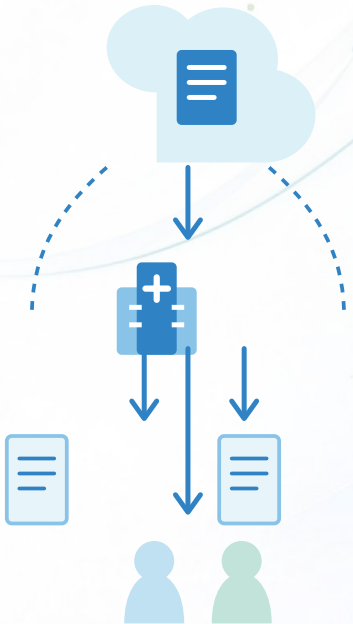
The client-side entries share a signature that is almost poignant. Wrong date of service. Wrong provider or campus. Demographic mismatch. Wrong date of birth. Misspelled name. Together, 720 requests and roughly 19,000 excess days, cancellation rates between 54 and 66 percent.

Every one of those is a field that could have been checked before the request went out. The reason none of them were is sitting in the data itself: in that system the patient's identity lived inside a free-text subject line rather than a structured field, so there was nothing to validate against.

One schema decision. Nineteen thousand days.

The middlemen

About a quarter of the time, a records request never reaches the hospital at all. It goes to a company that handles release of information on the hospital's behalf, and you have almost certainly never heard of any of them.



6,193
requests routed through a release-of-information platform

1,274
requests through Ciox, ChartSwap, and Datavant combined

1,057
requests bounced or misrouted by a platform

12,648
total resubmissions across roughly 6,601 requests

In the second book, 6,193 of 27,111

requests routed through a release-of-information platform. The platform table in the appendix records 2,021 named-platform requests; Ciox, ChartSwap, and Datavant together account for 1,274. The concentration raises a governance question: who can measure the exchange independently when one corporate family participates in more than one layer of it?

The failures at this layer have nothing to do with whether records exist. They are about handling. Our request vanished inside a platform's own system 218 times. It was bounced back with "we do not handle this" 113 times. Prepayment was demanded before any work began 36 times. Across the book, 1,057 requests were specifically bounced or misrouted by a platform and had to be sent again. The larger figure, 12,648, counts all recorded resubmission events across roughly 6,601 requests in the engagement, including repeated submissions arising from other causes. The two figures answer different questions and are not additive. We absorbed those resubmissions without billing an additional authorization fee for each repeat.

One note in the file captures the genre better than I can:

"DV rejected, due to them not handling this. They handle only hospital billing. Then we were told it is a clinic for physician billing and has to be submitted through the mail with a fifteen dollar check."

The first book shows the same thing from the other side. Requests routed through Ciox closed in 26.2 days at 73.6 percent on time. Chartswap, 42.0 days at 51.1. MRO, 41.6 at 42.9. The platform is chosen by the provider, not by us, so this is diagnostic rather than actionable, but the spread is real.

I want to be careful not to write a villain here, because the data does not support one. Independent platforms in this sample sometimes performed worse than companies inside the consolidated family. The family represented about 63 percent of named-platform volume and somewhat more than 70 percent of the overturned denials for which a platform was identified. That is modest overrepresentation, not proof of a disproportionate false-denial rate. A comparison using each platform's own denial count, challenge rate, and case mix would be needed to establish that. The descriptive results do not, by themselves, establish a reliability benefit or penalty from consolidation.

The structural point is about concentration layered on top of fragmentation. Consolidation can reduce the number of corporate owners without reducing the number of handoffs a record must survive. Publicly announced transactions in this market have increasingly placed provider-side release-of-information services, requester-side retrieval, and adjacent data products under common ownership. That may create efficiencies. It also makes independent measurement more important, because the same corporate family can have an economic interest in more than one part of the path being measured.

The closer measurement sits to the commercial interests being measured, the more important independent verification becomes.

The business of the missing record

The retrieval business is usually described as a service problem: somebody needs a record, somebody else has it, and a third party moves it from one side to the other. That description is accurate and incomplete. At scale, the record is also a revenue event. The patient creates

the underlying clinical event. The provider documents it. Then a commercial layer can earn money from releasing the record, locating it, moving it, delivering it, organizing it, summarizing it, and analyzing it.

There is nothing inherently wrong with that. Secure exchange costs money. People, technology, compliance, identity verification, payment processing, provider relationships, and follow-up are real work. The problem is structural: the people who bear the consequence of an incomplete answer are not necessarily the people who earn money from the transaction. A workflow can close, a fee can be charged, and a record can still be missing.

That separation between revenue and completeness is the part of this market that deserves more attention. The economic system knows when a request was submitted, when an invoice was paid, when a download occurred, and when a transaction closed. It often does not know whether the assembled history is complete. The easier things to count are monetized. The harder thing to count, whether the patient's story arrived intact, is left largely unmeasured.

The mergers changed the shape of the gate

The scale of the gatekeeper layer is no longer small. In 2021, Datavant and Ciox announced a $7.0 billion merger. The announcement said the combined company would have more than $700 million in annual revenue and described the result as the largest neutral and secure health-data ecosystem. Those are the companies' own public figures and words, not estimates from this study.

The important part for this paper is what happened next. In 2025, Datavant publicly announced the acquisition of Ontellus, a records-retrieval and claims-intelligence business, and said the transaction would create a Legal & Insurance vertical connecting providers and requesters. ChartSwap, which had been acquired by Ontellus years earlier, became part of that same legal-and-insurance ecosystem. Later in 2025, Datavant acquired DigitalOwl, adding automated medical-record summaries, chronology, de-duplication, and AI-enabled predictive analytics to the same end-to-end offering.

By 2026, Datavant's own public materials described a legal-and-insurance business spanning retrieval, provider access, payment, summaries, analysis, and related services, and said 90 percent of top carriers choose Datavant. In September 2026, KKR and Neuberger announced that they would acquire a significant minority stake while New Mountain Capital retained control. None of those transactions proves anything improper. Together, however, they show how much of the path from provider to requester can now sit inside one commercial ecosystem.



That matters because the patient and the requester generally do not choose the release-of-information vendor sitting behind a hospital. The provider does. A requester can choose its retrieval partner, but not necessarily the platform through which a particular provider will release the record. When the same corporate family increasingly participates in provider-side release, requester-side retrieval, payment, and downstream analysis, the market has fewer truly independent vantage points from which to measure the exchange.

Consolidation can make a fragmented process easier to navigate. It can also turn the company operating the tollbooth into a company that participates in more of the road. Those two things can be true at the same time. The question is not whether consolidation is good or bad. The question is what independent measurement should look like when the party being measured is large enough to occupy several positions in the same transaction.

A gatekeeper can be paid even when the answer is nothing

Public litigation illustrates the incentive problem in unusually concrete terms. In Maryland, plaintiffs have challenged Ciox charges described in court filings as Electronic Data Archive, Digital Archive, Retrieval, Instant Download, Handling, and Basic fees. The cited opinions concern jurisdiction and allegations, not final findings that the fees were unlawful, and they should be read that way.

The numbers are still instructive. In Hoggard v. Ciox Health, the federal court's recitation of



the record says Ciox reported charging $5,039,050 in the challenged categories between February 2010 and February 2022. In the separate Ayers case, the court wrote that Ciox estimated it had collected more than $2.89 million over roughly twelve years from a Basic Fee charged when a request yielded no records, about $250,000 a year. Again, those cases do not establish liability in this paper. They establish something narrower and, for this analysis, more important: even an answer of nothing can sit inside a revenue-producing transaction.

Now put that public record beside our operating data. We received 2,570 no records answers in the California book. We challenged 805 because something else in the file suggested the answer might be wrong. Three hundred twenty-one were overturned and the records were produced. We cannot infer that the same error rate applies to every unchallenged denial, and this paper does not try to. But we can say that a closed transaction and a correct answer are not the same thing.

That is the core governance problem. A fee is easy to verify. Completeness is not. A platform can prove that it processed the request. A provider can prove that it responded. A requester can prove that it paid. None of those facts prove that the answer was complete. When the system rewards completion of the transaction but does not measure completeness of the record, it is optimized around the wrong finish line.

The patient is the source of the value and has the least leverage

There is an odd asymmetry at the center of the medical-record economy. The record exists because a patient was injured, became ill, had a child, underwent imaging, took a medication, had surgery, or otherwise needed care. The patient supplied the body, the history, the symptoms, the risk, and the event that made the data worth recording. Yet once the record exists, nearly everyone around it can have a clearer economic position than the person it describes.

The provider is paid for care. A release company can be paid to produce the record. A retrieval company can be paid to obtain it. A platform can be paid to transmit or provide access to it. An analytics company can be paid to summarize it. An insurer, law firm, life company, researcher, or other authorized user may derive substantial value from the information. The patient, meanwhile, is often the person waiting for the system to tell the truth about what happened to them.

This is not an argument that every fee is excessive or that commercial participation is illegitimate. It is an argument about alignment. In a well-designed market, the thing that creates value and the thing the system is rewarded for should point in roughly the same direction. Here they can diverge. The transaction can succeed economically while failing informationally.

That divergence is why language such as ‘consumer convenience’ or



> **“The easier things to count are monetized. The harder thing to count, whether the patient’s story arrived intact, is left largely unmeasured.”**

'administrative
efficiency' is not enough by itself. A fast portal is useful. A consolidated network is useful. AI summarization can be useful. But none of those features answers the prior question: did the complete record make it through the gate? If the answer is unknown, efficiency can simply move an incomplete file faster.

The danger is not monopoly as a slogan. It is dependency without an independent check

This paper is not an antitrust opinion, and market concentration by itself does not prove unlawful conduct. The concern is operational. A health system can outsource release. The release layer can be connected to the retrieval layer. Retrieval can be connected to claims intelligence and automated analysis. The result can be an elegant end-to-end product. It can also make the ecosystem increasingly dependent on the internal measurements of the same commercial family that operates it.

That is why independence matters more as the market consolidates, not less. A provider scorecard should not depend solely on the provider. A platform scorecard should not depend solely on the platform. A retrieval scorecard should not depend solely on the retrieval company. Every one of us, The Records Company included, should be measurable by somebody with access to the operating record and no economic reason to make the number look better than it is.

The answer to concentration is not to presume bad intent. It is to demand better instruments. Measure false no records responses. Measure bounce and resubmission rates. Measure fees by outcome, not merely by transaction. Resolve provider identities before publishing scorecards. Separate provider delay from requester error from platform failure. And most important, measure completeness. A market this consequential should not be allowed to define success as 'the request closed.'

The public merger announcements tell us where the market is going: more scale, more integration, more data products, and more of the workflow under common ownership. The operating record tells us what must travel with that growth: independent measurement. Otherwise consolidation risks making the gate easier to find while making the gatekeeper harder to audit.

What fragmentation makes people do

Underneath the timeline sits an operation that appears in no performance report I have ever seen. Seven months, one firm: 16,817 faxes. It is 2026. Twelve thousand and eighty-seven

Faxes **16,817**	Prepayment and check actions **12,087**	Items mailed **5,349**
Status inquiries **10,071**	Staff hours (mechanical actions) **3,256**	Cost of mechanical actions **$175,300**

99.8% of status inquiries arrived after a status had already been posted. One request was asked about eleven separate times.

prepayment and check actions. Five thousand three hundred and forty-nine items mailed, and 182 checks re-mailed because the first ones never cleared.

And 10,071 status inquiries, of which 10,052 arrived after a status had already been posted on that same request. That is 99.8 percent of the questions answering themselves before they were asked. One request was asked about eleven separate times. More than 2,100 were chased more than once. Those redundant questions alone consumed 671 hours, which is four months of somebody's working life spent telling people things they could already see.

The mechanical actions came to 3,256 staff hours and about $175,300, and that counts only the mechanical ones. It excludes the calls, the portals, the rejections worked and reworked, and the judgment every request demands. Meanwhile the provider fees were advanced and billed only on delivery, so roughly $903,000 of the firm's costs were being carried, interest free, at any given moment.

This is what a fragmented system produces when you look underneath the clock. Not just delay. Vast quantities of redundant communication whose only purpose is to compensate for the absence of a shared view.

What it costs the person the record is about

Everything so far is measured in requests, which is the unit my industry counts in and the wrong unit for this question. Nobody is a request.

Patients represented **1,979**	Requests **7,697**	Average requests per patient **3.9**
Patients with at least one request carrying a recorded issue **1,187** (60.0%)	Patients with at least one cancelled request **613** (31.0%)	Patients with at least one recorded request duration over 90 days **442** (22.3%)

Requests in repeat patient/provider/record-type clusters
937 (12.2%)

Repeat-cluster vs. single requests

What we can report directly is that repeat requests concentrate trouble. Nine hundred thirty-seven requests, 12.2 percent of the book, belong to 447 clusters where the same patient, provider, and record type were requested more than once.

Issue rate
Share of requests with a recorded issue

Repeat-cluster requests	53.5%
Single requests	34.1%

Cancellation rate
Share of requests that were cancelled

Repeat-cluster requests	28.7%
Single requests	14.4%

Patient percentages use all 1,979 distinct pseudonymous patient IDs as the denominator. Each patient is counted once per measure, so the rows can overlap. A duration over 90 days means a value recorded in the request-duration field; it is not an estimate of clinical delay. Repeat-cluster measures use requests, not patients, as their denominator.

The Boston book covers 1,979 patients across 7,697 requests, about 3.9 requests per person. Counting each patient once shows how often a problem reached the person behind the requests.

What we can report directly is that repeat requests concentrate trouble. Nine hundred thirty-seven requests, 12.2 percent of the book, belong to 447 clusters where the same patient, provider, and record type were requested more than once. Those repeat-cluster requests carry a 53.5 percent issue rate versus 34.1 percent for single requests, and a 28.7 percent cancellation rate versus 14.4 percent. Some repeats are legitimate scope changes, but the outcome profile shows that problems are not distributed independently across a patient's record.
The person waits on the slowest fragment, not the average one. This is what gets lost when performance is reported as a mean. A 40-day average across a book tells you almost nothing about what a human being experienced, because their case was gated by whichever of their four requests went worst. Averages describe the vendor. Maximums describe the patient.

Fragmentation makes people do this

12.2%
of the book, belong to 447 clusters where the same patient, provider, and record type were requested more than once.

53.5%
issue rate for repeat-cluster requests vs 34.1% for single requests.

28.7%
cancellation rate for repeat-cluster requests vs 14.4% for single requests.

40
day average across a book tells you almost nothing about what a human being experienced.

75,800
elapsed days consumed by requests that ultimately produced nothing.

38
days per patient, if those 75,800 days were spread evenly across the 1,979 patients.

And in a contingency case nothing moves until the records arrive. The demand letter cannot be written. The negotiation cannot open. The settlement cannot close. Someone who was injured, who may not be working and may already be receiving bills for the treatment in question, waits on the slowest fragment. Across this one book, requests that ultimately produced nothing consumed roughly 75,800 elapsed days. Spread evenly that would equal about 38 days per patient, but the burden is not evenly distributed; repeat-request clusters show that trouble concentrates.

The false denials are the worst of it, because nobody involved ever learns. A delay is visible and somebody eventually complains. An absence that was never real produces no complaint at all, because the file simply proceeds as though the treatment did not happen. The imaging that documents an injury. The specialist bill that establishes what it cost. The record from the facility that saw the patient first. When one of those is reported not to exist and the report is

wrong, the case is valued without it, and the person is compensated for a smaller injury than the one they have.

That is the cost worth sitting with. Not days, not dollars, not turnaround. A file that quietly understates what happened to somebody, closed by people acting in good faith on an answer nobody could check.

The stakes do not stop at a claim file

This dataset cannot tell us how many people were harmed because a record was missing, and it certainly cannot count deaths caused by a failed retrieval. It would be irresponsible to pretend otherwise. But it would be equally irresponsible to treat the consequence as administrative simply because mortality is harder to observe than turnaround time.

Health-services research has repeatedly linked fragmented care and poor continuity with worse clinical outcomes. A large nationwide cohort study published in 2023 found that multiple measures of care fragmentation were associated with potentially inappropriate medication and increased mortality after adjustment for major demographic, socioeconomic, and morbidity factors. That literature studies fragmentation of care rather than record retrieval, so it does not prove that a missing chart caused a particular death. It does establish that fragmentation is not a paperwork problem with paperwork-sized consequences.

Two firms. Opposite coasts. Different years.

Essentially no shared provider network.

Boston book

In the Boston book, 36.5 percent of requests hit at least one problem.

Hit at least one problem

36.5%

Client-caused problem-rate range

6.0%–31.8%

California book

In the California book, about 36 percent. The requester spread reproduces too.

Hit at least one problem

~36%

Overall requester problem-rate range

25%–68%

Repeat requests behave identically in both books.

A missing record can hide a prior diagnosis, a medication, an imaging result, a specialist's recommendation, or the fact that a procedure happened at all. The causal chain from one missing document to one bad outcome is rarely visible after the fact. That is precisely why the risk is easy to discount. Lives can be affected, and in some circumstances lives can be lost, without the information failure ever being recorded as the reason.



The same number, twice

Two firms. Opposite coasts. Different years. Essentially no shared provider network. In the Boston book, 36.5 percent of requests hit at least one problem. In the California book, about 36 percent.

The requester spread reproduces too. In Boston, 29 people submitting 40 or more requests ranged from a 6.0 percent client-caused problem rate to 31.8, a factor of five, on the same work in the same period. In California, 25 requesters ranged from 25 to 68 percent overall. Repeat requests behave identically in both.

The similarity is not well explained by ordinary sampling error. At a 95 percent level, the approximate margin around the Boston book's 36.5 percent rate is about plus or minus 1.1 percentage points; for the California book's roughly 36 percent rate it is about plus or minus 0.6 points. The more important limitation is comparability: the two coding schemes were built independently and the turnaround clocks differ, so the rates should not be treated as a formal replication until they are rebuilt on common definitions.

But if the definitions survive scrutiny, a rate that reproduces across books with this little in common is more plausibly a property of the system than of either firm.

These were TRC-managed engagements, not audits of other retrieval vendors. The operating record lets us separate our own work from requester input, provider response, and platform handling. That distinction matters whenever a delay is attributed to a vendor without first tracing what happened.

What a useful scorecard would measure

Most retrieval scorecards collapse performance into one number: turnaround. The data in these two books argues for something more discriminating. A useful scorecard should separate at least five things.

A useful scorecard should separate at least five things

Cause, not just age.

A sixty-day request caused by a provider backlog is not the same event as a sixty-day request waiting on a wrong date of service.

Completeness, not just closure.

A request closed as "no records" may still be wrong. A program that measures only whether a request closed and how long it took can look efficient.

Provider and platform identity, after entity resolution.

The Boston data reduced 1,368 raw provider names to 1,334 organizations across 74 clusters.

Rework and administrative load.

Resubmissions, redundant status inquiries, checks, prepayments, mailings, and repeated routing attempts are measurable.

Case-level exposure.

A patient with several requests waits on the slowest fragment. Maximum delay often matters more than the mean delay across a book of business.

Cause, not just age. A sixty-day request caused by a provider backlog is not the same event as a sixty-day request waiting on a wrong date of service, an unanswered client question, or a platform that lost the submission. Age tells you that something is late. Attribution tells you what to fix.

Completeness, not just closure. A request closed as "no records" may still be wrong. A program that measures only whether a request closed and how long it took can look efficient while quietly producing incomplete files. The 321 overturned denials in the second book are enough to show why closure cannot be the final quality measure.

Provider and platform identity, after entity resolution. The Boston data reduced 1,368 raw provider names to 1,334 organizations across 74 clusters. Until names are reconciled, a scorecard can split one hospital into several rows and manufacture differences that do not exist.

Rework and administrative load. Resubmissions, redundant status inquiries, checks, prepayments, mailings, and repeated routing attempts are part of the cost of fragmentation even when the final record eventually arrives. They are operational waste, and they are measurable.

Case-level exposure, not only request averages. A patient with several requests waits on the slowest fragment. That means the maximum delay inside a case often matters more than the mean delay across a book of business.

Once those measures are separated, accountability becomes more useful. The question stops being simply "Who was slow?" and becomes "Where did the failure begin, what did it cost, did the file arrive complete, and what can be changed before the next request?"

What I cannot tell you

Four limits bound all of it, and I would rather hand them over than have them found.

Fault is assigned, not audited. Whose problem a delay was is an internal classification made by an operations team reading a file. It is a judgment. In the Boston book, "no record found" appears 230 times coded as a client error and 67 times as a provider failure, and the underlying observation is identical in all 297 cases: the request went out and nothing came back.

Every figure is a floor, because causes were counted only where somebody wrote one down in plain language at the time.

Accountability becomes more useful when we can answer:

- Who was slow?
- Where did the failure begin?
- What did it cost?
- Did the file arrive complete?
- What can be changed before the next request?

Two in five is not a rate, for the selection reason above.

And the two books do not share definitions. One measures turnaround in calendar days from creation to solve, the other in median business days excluding weekends and holidays. Any comparison across them, including the 36 percent I just made much of, needs rebuilding on common definitions first. That work has not been done.

There is a fifth limit, and it is the one that tells you the most about the state of measurement in this field. Before any of these numbers existed, the data had to be repaired.

The hospital naming problem described in the Boston section is only part of it, and one string in that data differed from another only by an HTML-escaped apostrophe. Four categories of problem were spelled six different ways, including two spellings of "cancelled" and three of "no record found." Fifty-one requests appeared twice with conflicting outcomes, one row marked complete and its twin marked cancelled. Six columns were empty in all 7,697 rows. Any scorecard built on the raw fields would have been wrong. That is not a footnote. It is a finding.

The problem worth solving

Everything here points at one missing instrument.

We can tell you how long a record took. We can tell you what it cost. We cannot tell you whether it is complete, and completeness is the only property that matters to anyone downstream. A file assembled in nine days with a piece missing is worse than a file assembled in ninety with nothing missing, and today there is no way to tell those two files apart.

So here is the problem I want to work on: give an assembled medical record a completeness score.

Not a guarantee. A probability, with the reasoning visible. For a given episode of care, what should exist, what do we hold, and how likely is it that a reported absence is a real one.

I think it is buildable, in four steps.

Learn what an episode should produce. An emergency visit with imaging generates a predictable constellation: facility record, facility bill, emergency physician bill, radiology report, radiologist bill. A surgical admission generates a different one. Across millions of observed episodes that mapping is learnable, because we watch it resolve every day. Compare held against expected. The residual is the candidate gap. This alone would change practice, because it converts "we did not get everything" into "the anesthesia bill for this date is unaccounted for."

Model the denial rather than accept it. Here the data we already have becomes unusually valuable. Those 805 challenges are labeled. We asked again, and 321 times we learned the answer had been wrong. That is a training set for the one thing nobody has ever modeled: the probability that a reported absence is false, conditional on custodian, record type, channel, and distance from the institution that produced the care.

Publish the score with the record. An assembled file arrives carrying its own coverage estimate, the way an estimate arrives carrying a confidence interval.

The hard part is not the first two steps. It is the third, because the challenge decision was made by people using judgment, which makes this a selection problem rather than a clean sample. Any model built from the challenge data has to account for that before its output can be trusted.

What it would unlock

Research-ready records instead of merely linked ones. The progression from fragmented records to linked data to research-ready assets to AI-enabled discovery is attractive, but linkage alone is not enough. Linked records with unknown coverage are still unquantified. A cohort carrying a coverage distribution is a different scientific object than one that does not, and every downstream analysis inherits whichever one it was built from.

Accountability that survives contact with incentives. Once completeness is measurable, a custodian, a platform, or a vendor can be scored on it. Today every performance standard in this field measures speed, which means every party is optimized for what is easy to observe rather than what matters. That is not a moral failure. It is a measurement failure, and measurement failures are fixable.

A different basis for policy. State statutes set deadlines for producing records. None set a standard for producing all of them, because there has been no way to check.

And this is where it stops being an operations problem. Fragmentation is not distributed evenly across people. The records most likely to scatter belong to patients who moved between facilities, who were transferred, who were treated in safety-net settings, who saw many specialists, whose care was interrupted or paid for in unusual ways. Those are not random patients.

If false absences rise with fragmentation, as the prediction above holds they should, then the information loss falls hardest on exactly the populations whose care is already the most fragmented. And it falls silently, because nobody is counting.
A patient whose care was continuous and local gets a record that is nearly complete. A patient whose care was scattered gets a record with holes in it, and every

KEY TAKEAWAY

Completeness is the missing instrument.

321 of 805
challenged denials were proven wrong.

BUILDING A COMPLETENESS SCORE

1. Learn expected episode components
2. Compare expected vs. held record
3. Model false absences
4. Publish the score with the record

Better data.
Fairer decisions.
Healthier futures.

decision made about them afterward, clinical, financial, legal or scientific, is made on the worse file without anyone knowing that is what happened.

What the operating record makes possible

The operating record creates an unusual test environment: de-identified retrieval histories with outcomes attached across multiple client books and jurisdictions, a provider database exceeding 90,000 records, processing notes at the granularity where cause is actually written down, and a live environment where a hypothesis can be tested against real requests rather than simulated ones. Most retrieval operations generate this data. Very few retain it in a form that can be analyzed.

It also captures a part of the system that rarely appears in published datasets: the operational layer between a request and the record that eventually arrives, or does not.

That record can show how raw operating history becomes a usable dataset, and why a provider scorecard built on unresolved names and inconsistent fields will be wrong before the first metric is calculated.
It can support an attribution scheme that separates client-side, provider-side, platform-side, and vendor-side failures while making the uncertainty in those assignments explicit.

It can show how absence is actually produced: who says no, why, through which channel, and what it takes to overturn the answer.

Boston also provides a useful institutional comparison: major hospitals and health systems in the same regional market, operating under the same external rules, show twofold and sometimes larger differences in observed retrieval performance.

And the honest version of applying AI to this layer. We built a system that reads the full history behind every request, identifies what is blocking it, names who is responsible including when that is us, and carries what it learns about a custodian forward. I think it makes the work better. I cannot tell you what it does to the number of false denials that slip through, because nobody has established what that number is. We are automating a layer whose error characteristics have never been measured, and the automation will be graded on speed and volume, which is precisely how the layer became invisible in the first place.

Everything in this paper came out of ordinary operating records. No survey, no instrumentation, no special study. Two engagements, 34,808 requests, 270,150 processing notes, read line by line. The same data exists in every retrieval operation in the country. Almost nobody reads it.

That is the opportunity. The data already exists inside ordinary retrieval operations. The missing piece is the discipline to read it as an operating record rather than a queue of transactions.

> **The data already exists. The missing piece is the discipline to read it as an operating record rather than a queue of transactions.**

Six questions the industry still cannot answer

These are the questions the data leaves open:

1. How do you prove something is missing, when the only evidence that it existed is the thing you could not get?
2. We only argued with the denials that looked suspicious. How would you find out how often the ordinary-looking ones were wrong too, without asking everyone to do twice the work?
3. If the records that vanish belong mostly to patients whose care was already scattered, what happens to a study that never knew any of it was gone?
4. Major Boston institutions operate under the same state and federal rules, receive similar requests from the same counterparty, and still show sharply different retrieval performance. What internal operating differences would explain that spread, and how would you measure them?
5. Everyone in a position to measure this layer either owns a piece of it, depends on it, or competes with it. So who is left to do the measuring?
6. We are automating this work right now, mine included, without knowing how often it gets absence wrong. What should anyone have to show you before you trust what comes out?

Public-source note

All market-structure discussion in this case study is based on publicly announced transactions, public company materials, and public court filings. The transaction history cited includes the 2021 Datavant-Ciox merger; the 2025 Ontellus/ChartSwap and DigitalOwl acquisitions; and the 2026 announced minority investment by KKR and Neuberger. The Maryland fee discussion summarizes allegations and amounts described in federal court opinions addressing Hoggard v. Ciox Health and Ayers v. Ciox Health; those opinions are not cited here as findings that the challenged fees were unlawful. The patient-safety discussion cites the 2023 nationwide cohort study, 'Healthcare fragmentation, multimorbidity, potentially inappropriate medication, and mortality.' Hospital transitions are documented in Massachusetts Department of Public Health notices and BMC Health System's May 1, 2025 announcement: Carney Hospital closed August 31, 2024, and BMC assumed operations of St. Elizabeth's and Good Samaritan on October 1, 2024. Berkshire Health Systems identifies Berkshire Medical Center as a separate hospital in its own system, also abbreviated BMC.

Public sources: https://www.datavant.com/press-release/datavant-and-ciox-health-announce-merger-creating-the-largest-neutral-and-secure-health-data-ecosystem ; https://www.datavant.com/press-release/datavant-completes-acquisition-of-ontellus ; https://chartswap.com/about-us/ ; https://www.datavant.com/press-release/datavant-completes-acquisition-of-digitalowl ; https://www.datavant.com/legal-and-insurance ; https://www.datavant.com/press-release/neuberger-and-kkr-to-invest-in-datavant-alongside-new-mountain-capital ; https://law.justia.com/cases/federal/district-courts/maryland/mddce/8%3A2023cv03457/550980/26/ ; https://law.justia.com/cases/federal/district-courts/maryland/mddce/8%3A2023cv03079/551187/25/ ; https://pubmed.ncbi.nlm.nih.gov/37580711/ ; https://www.mass.gov/info-details/steward-health-care-transition-to-the-new-operators ; https://www.bmc.org/news/boston-medical-center-health-system-unveils-new-names-good-samaritan-and-st-elizabeths-medical ; https://www.berkshirehealthsystems.org/location/berkshire-medical-center/

Fewer delays.
Clearer records.
A healthier tomorrow.

Appendix

A1. Where the fault landed (Boston book, n = 7,697)

Distribution of requests, average turnaround time, and outcomes by where the problem originated.

Category	Requests	Average wait	Met 30 days	Cancelled
No problem	4,888	27.2 days	68.0%	1.6%
Client-caused	1,522	52.6 days	34.6%	63.3%
Provider-caused	1,195	74.6 days	18.3%	16.2%
Our own error	92	59.8 days	25.0%	6.5%

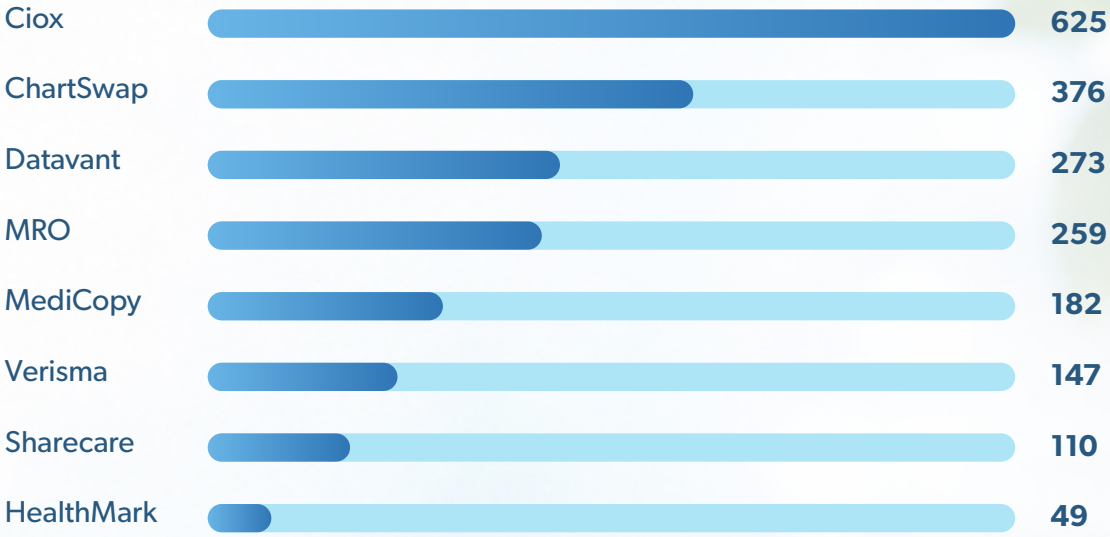
A2. Damage ranked by cause (Boston book)

The most common specific problems, who they came from, and their operational impact.

#	Cause	Whose	Requests	Average wait	Days lost
1	Backlogged or unreachable	Provider	227	103.2 days	17,232
2	Certification missing	Provider	346	57.6 days	10,515
3	Non-receipt	Provider	123	82.6 days	6,807
4	Wrong date of service	Client	238	54.1 days	6,381
5	Wrong provider or campus	Client	180	58.6 days	5,643
6	No client response	Client	45	103.3 days	3,424

A3. Requests routed through each named platform (California book)

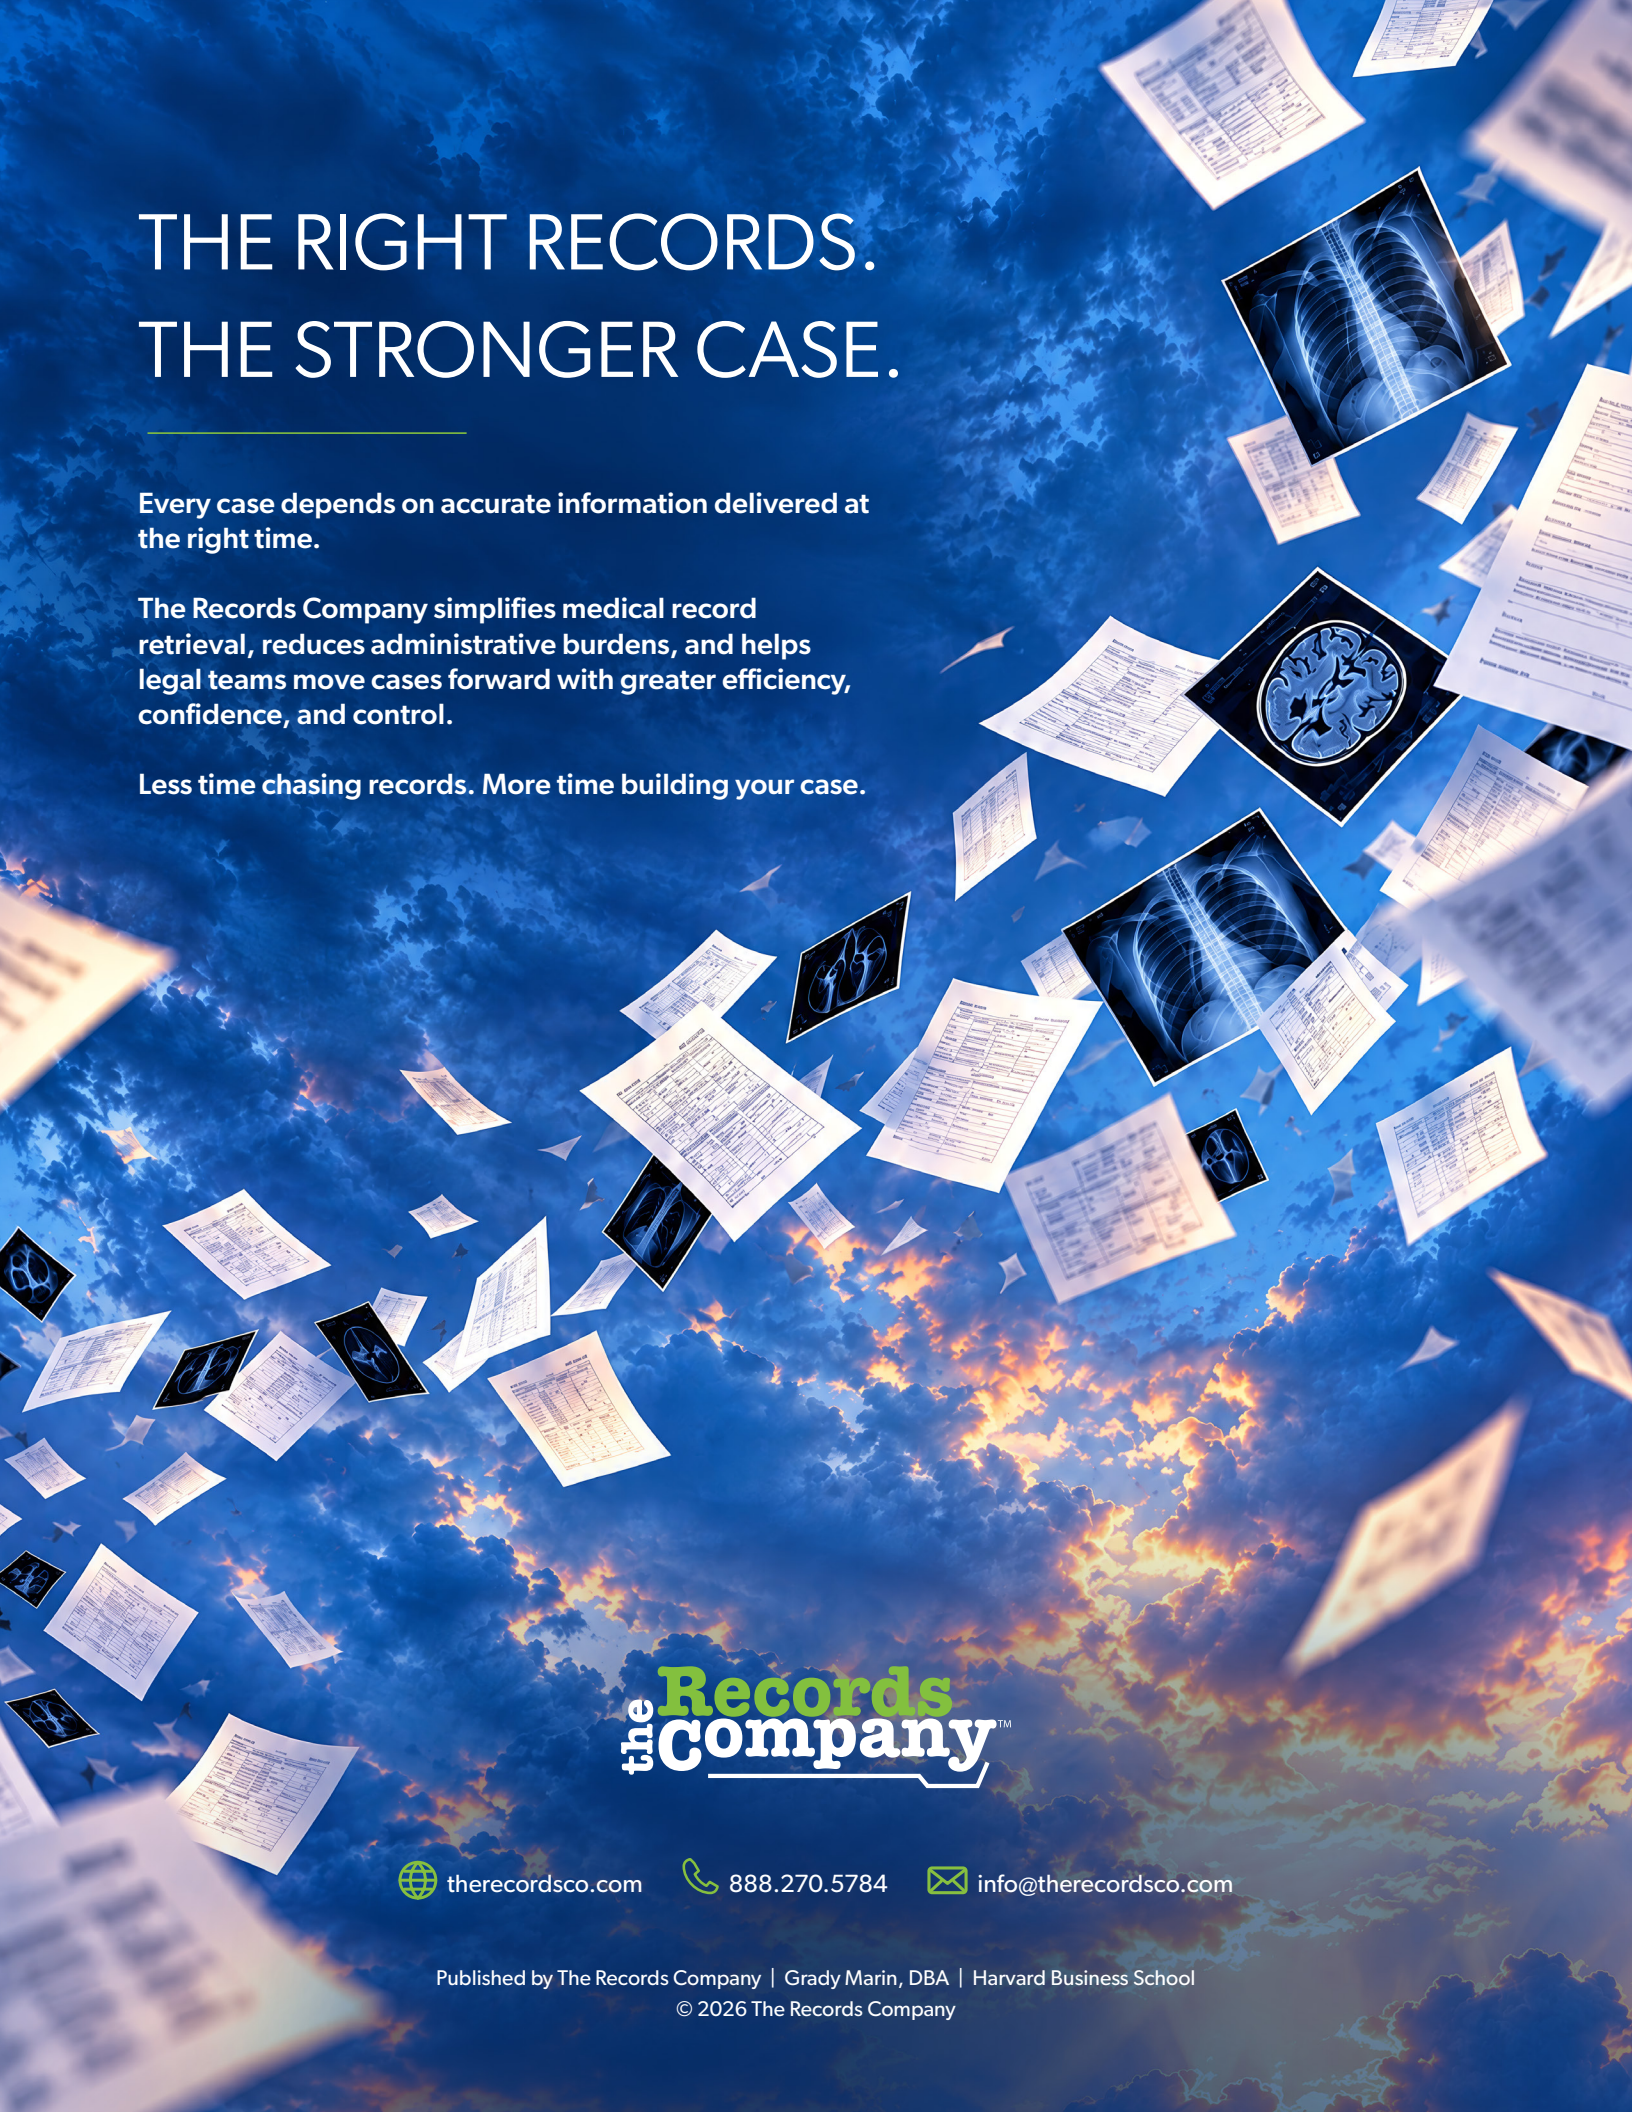
Number of requests sent through each named records retrieval platform.



Note on sources

The Records Company was the retrieval vendor in both engagements. These are analyses of TRC's own operating records, not independent audits of two outside vendors; "Our own error" denotes an internal TRC issue. Book A: 7,697 unique requests placed by a national plaintiff firm from December 2022 through March 2026, covering 1,979 pseudonymous patients and 1,334 resolved provider organizations across nineteen states, concentrated in greater Boston. Turnaround is measured in calendar days. Book B: 27,111 requests and 270,150 processing notes from a second national firm over seven months, weighted toward California; turnaround is reported in business days excluding weekends and provider holidays. Cause is read from contemporaneous processing notes and counted only where the note states a cause, so cause-based figures are conservative floors.
Book A was normalized before analysis: 51 duplicate ticket IDs were resolved, provider-name variants were consolidated, inconsistent category labels were normalized, and impossible turnaround values were blanked. Figures in this case study are reported as computed from the source records after those documented cleaning and entity-resolution steps. Client identities, patient identities, and individual identities are withheld.

Patient measures count each distinct patient once per measure. An issue is any request with a recorded problem; a cancellation is a request recorded as cancelled; a long duration is more than 90 calendar days. A patient can meet more than one condition. The Boston Medical Center group excludes requests to St. Elizabeth's and Good Samaritan. Provider-name groupings are not ownership assignments at the date of each request.


THE RIGHT RECORDS.
THE STRONGER CASE.
Every case depends on accurate information delivered at the right time.
The Records Company simplifies medical record retrieval, reduces administrative burdens, and helps legal teams move cases forward with greater efficiency, confidence, and control.
Less time chasing records. More time building your case.
the Records company™
therecordsco.com
888.270.5784
info@therecordsco.com
Published by The Records Company | Grady Marin, DBA | Harvard Business School
© 2026 The Records Company