TELCOR®

BOLT-ON OR BUILT-IN: AI IN LAB RCM

Where AI is paying off in laboratory billing today and where it isn't.

Sunjay Mouli

VICE PRESIDENT, AI SYSTEMS
TELCOR

BOLT-ON OR BUILT-IN: AI IN LAB RCM

Where AI is paying off in laboratory billing today and where it isn't.

A field perspective from Sunjay Mouli, Vice President of AI Systems, TELCOR. Written for laboratory billing leaders, RCM directors, and lab executives at independent and specialty laboratories.

IN THIS PAPER

Why most AI in lab billing stalls, the difference between bolt-on and built-in deployment, a crawl/walk/run maturity model, and the questions to ask a vendor.

READING TIME

About 12 minutes.

01 / THE REAL PROBLEM, NOT THE AI PITCH

Most labs already have automation. What they don't have is the context to make it useful. One question, three decisions.

Independent and specialty laboratories are being squeezed from three directions at once. Payer requirements keep tightening, per-claim documentation burden keeps rising, and the billing workforce that absorbs the difference keeps getting harder to hire.

The scale of the administrative drag is now well documented. The American Hospital Association estimates that hospitals spent $43 billion in 2025 simply trying to collect payments insurers owed for care already delivered, with the average hospital carrying roughly 64 dedicated administrative and billing staff to do it.[1] Availity's 2026 Abrasion Index puts a finer point on it: around **70% of denials** are eventually overturned and paid, which means most denial work is churn rather than recovery. In many cases the information was right and simply arrived too late to matter.[2] Laboratories carry a version of this problem that is structurally worse. A lab needs to articulate medical necessity while sitting outside the traditional provider-patient relationship. The hospital coder can go back to the chart. The lab usually cannot.

> **"Laboratory has all the complexity of having to articulate medical necessity, but sits removed from the traditional provider-patient relationship. That's what makes it especially tricky."**
>
> SUNJAY MOULI

That gap shows up as lost revenue in specific, predictable places. In laboratory billing, diagnosis drives the success of a claim, and in many cases the correct diagnosis was never available to

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

whoever entered the requisition. It may exist in the medical record, but if it never reached the requisition the claim goes out incomplete, and the money is gone before anyone has a chance to work it.

Molecular and genetic testing raise the stakes further, where high per-claim values make each documentation failure individually consequential rather than statistically absorbed.[3] Meanwhile the labor pool that would otherwise brute-force the problem is thinning; clinical laboratory vacancy rates run as high as 25% in parts of the country.[4]

That is why AI in lab RCM often disappoints. The tooling is aimed at the tasks, but the problem is the missing context.

02 / WHAT IS FORCING THE QUESTION NOW

Four capabilities, in plain language

Vendor language in this category has gotten loose enough that "AI" no longer tells a buyer anything. Almost everything sold into lab RCM today is a combination of four distinct capabilities. Knowing which ones a product has is most of the evaluation.

CAPABILITY	WHAT IT DOES	IN A LAB, THAT LOOKS LIKE
Classification	Identifies what a document is and who it belongs to	Sorting an inbound fax into denial letter, audit request, or prior auth determination, and attaching it to the right patient and claim
Extraction	Pulls structured facts out of unstructured clinical text	Finding the supporting diagnosis inside a 40-page record when the requisition arrived with an ICD-10 code that won't pay
Reasoning	Assembles an argument from assembled evidence	Constructing a payer- and patient-specific medical necessity argument for an appeal, citing the record it came from
Orchestration	Takes the next action, and hands off to the next process	An eligibility agent determining prior auth is required, then passing the case to a prior auth agent that builds and submits the packet

Most point solutions do the first two well. The value in lab RCM concentrates in the third and fourth, and those are the two that depend almost entirely on how much of your data the system can access.

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

03/ THREE WAYS LABS GET THIS WRONG

Automating actions instead of rethinking the process

The most common failure mode is a list of rote actions to automate (packet generation, denial reason interpretation, record summarization) with no view of the process as a whole. You can cut the time spent on a single action by 50, 80, 90%. If the process around it is broken, it stays broken.

IN HIS WORDS

"What AI actually gives you is the space to rethink the process. A lot of people are doing things the same way just because they've always done it that way. We'll ask someone, why are you doing this? And the only real answer is: this is the process we've always had."

This is the difference between a technology purchase and an operating change, and it is visible in the industry data. An HFMA and FinThrive survey found that 63% of organizations are using AI or automation in the revenue cycle, while only 15% report positive ROI.[5]

REACHING FOR AI WHEN A PLAIN FIX WOULD DO

Sometimes the process change itself delivers the result. An integration, a removed handoff, an eliminated step. Introducing a model where none is needed adds cost, adds a thing to govern, and obscures the fact that the underlying step should not have existed.

STARTING TOO BIG

"Automate my entire denials process" is not an instruction an agent can execute if the organization has no documented SOP for denials. Bad inputs produce bad outputs, and model capability is no substitute for knowing what good looks like. The same applies to analytics: AI can write the queries and structure the datasets, but it cannot supply the business understanding of why your billing looks the way it does.

04 / LAYER ONE: THE QUICK WIN

Correspondence processing, before and after

Bolt-on AI, a tool layered over whatever systems you already run, is real and worth doing. For most labs it is the right place to start, and correspondence processing is the usual first candidate.

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

BEFORE

A fax, PDF, or lockbox document arrives. Someone opens it, identifies the patient, identifies the claim, drags it into another system, types in the data, and then decides what to do about it. Most of the time it gets indexed and never used again. Teams of people spend their days on this.

AFTER

The document is read, classified, indexed to the right patient and claim, and reasoned over. Is this an audit request, a denial letter, a prior auth determination? The downstream action is triggered from the answer. The cost center starts producing feedback on how the process is performing.

The second-order effect matters more than the time saved. Correspondence contains the instructions for winning: the submission modality, the correct address or portal, the response deadline. Labs lose appeals they should win because nobody read the letter closely enough to submit correctly, an error that disappears once the document is machine-read every time.

THE FILTER: WHAT MAKES A GOOD FIRST WORKFLOW

Clear inputs and outputs

You know where the work comes from and exactly where it needs to go.

Real volume

Enough hours are going into it that redirected effort is worth something.

Measurable performance

You can evaluate accuracy and speed, not just assert improvement.

Historical baseline

Past data to compare against, so the win is provable to leadership.

The last criterion is the one most often skipped, and it is the one that decides whether you get funded for a second workflow. If you cannot objectively show the organization that this worked, the program stops here regardless of the result.

The ceiling on this layer is context. A bolt-on tool sees the documents it was pointed at. It does not see the denial in the billing system, the demographics in the relational database, the records in the document store, or the contract terms. Beyond well-bounded work it has to fill in the gaps, and that is where accuracy drops.

05/ LAYER TWO: NATIVELY EMBEDDED AI

Everything depends on context

The most expensive failure in lab RCM is not a broken task but a decision made without enough context around the claim. That description covers prior authorization, denials, appeals, and most processes underneath a lab's lost revenue.

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

The information exists. It is just scattered: medical records in one document store, the denial in the billing system, correspondence sitting in a repository untranscribed, patient data in a database that was never linked to the claim. Few teams can consolidate that at pace, and no model can reason over data it cannot reach.

Asked what changes once that data is in one place, Mouli does not hedge:

> **"Once you do have that full context, the ability to do something with the data becomes exponentially larger. The usage of a language model depends entirely on the source data you have."**
>
> SUNJAY MOULI

None of this is a new argument. Pre-AI, you could buy the best LIS or the best CRM on the market and still be stuck if it wasn't connected to anything, with staff downloading reports from one system and retyping them into another. AI does not change that logic. It raises the price of ignoring it, because the model's output quality now scales directly with how much of your operation it can see.

Native access to the data model and workflow engine changes two things at once: the context available for a decision, and the actions available afterward. Assembling a medical necessity argument, generating a payer- and patient-specific appeal, writing the result back into the system of record. All of these are downstream of integration, not of model choice.

THE ROLE OF COMPLIANCE IN AI

There is real and warranted caution about AI touching clinical data. These systems must be designed for the sensitivity of the material, since mishandling patient data is not a recoverable error. But the assumption underneath most of the fear-mongering, that AI weakens auditability, is backwards. On this point he is blunt:

> **"Agents are far better at compliance than humans are. If you tell an agent to time-stamp an action, it does that every single time. A human being rarely wants to."**
>
> SUNJAY MOULI

ICD coding is a clear example. Diagnosis drives claim success in laboratory billing, and labs frequently receive requisitions with a diagnosis that won't support the test, often because whoever

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

ordered it didn't have the full record. Some labs now staff certified coders authorized to make those edits, substantiated against the medical record. Experienced CPCs are good at logging their reasoning. But "good at logging" is not the same as complete.

In an AI-assisted version of that workflow, every input is captured: which passage of which record supported the edit, cited. The human stays in the loop. What changes is that you now have an accurate representation of what information produced the decision, which is exactly what a payer audit asks for and exactly what humans, designing for their own tolerance, tend to under-record.

This gets more consequential as laboratories gain access to more patient records through health information exchanges and similar initiatives. The more external clinical data enters lab decisionmaking, the more traceability stops being a compliance formality and starts being the operational record of how revenue was earned.

06 / WHERE YOU ARE: MATURITY MODEL

Crawl, walk, run

Every mature deployment worth pointing at started in the crawl stage. Not one of them started by automating a whole department.

STAGE	WHAT IT LOOKS LIKE	HUMAN ROLE	TYPICAL HORIZON
Crawl	One point-to-point workflow with clear inputs and outputs, running largely off-the-shelf. Correspondence processing is a typical example.	Confirming results case by case	Days to weeks to deploy
Walk	A workflow that needs real reasoning, such as appeals or prior authorization, fed with full patient context, writing results back into your systems.	Reviewing exceptions only	Weeks to a quarter
Run	Multiple workflows talking to each other, with evals running on your own data. An eligibility agent determines prior auth necessity and hands the case to a prior auth agent that builds and submits the package.	Monitoring activity and the feedback loop	Multiple quarters

WHAT A REALISTIC TIMELINE LOOKS LIKE

Using TELCOR AI Solutions, a pre-built workflow addressing a known lab RCM pain point can be running inside an organization in a couple of days to a week or two. From there, tuning the workflow to your process and getting real value out of it usually takes another six to eight weeks.

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

Then there is a third phase, measured in months, where the metrics actually move. Some of that is out of anyone's hands: a denials or appeals workflow waits eight to sixteen weeks on payer responses before results are legible. The reasonable expectation is positive movement within the first quarter or two.

If you are two quarters in and nothing has moved, that is a poor prognosis. Either the solution isn't working or you picked the wrong problem, and both are yours to fix.

07 / SIX QUESTIONS TO ASK A VENDOR

01	**How many workflows do you have in production and with whom?** A vendor who cannot answer this specifically is selling a roadmap. Production count is the single most diagnostic question in the category.
02	**Show me one decision's audit trail, end to end.** Every input, every citation, every timestamp. If the trail is thinner than what your coders currently document by hand, the system was not designed for your compliance posture.
03	**What data can you actually reach in my environment?** Billing system, order management, document stores, records retrieval. Not what's on the integration roadmap, but what is live, and how it was done at a comparable lab.
04	**How will we evaluate this and against what baseline?** Ask for the eval methodology, in writing, before deployment. A vendor confident in the result will define the scoreboard first.
05	**What breaks if I scale this five times?** Bolt-ons have a ceiling and it is usually context. Ask where theirs is instead of discovering it in month seven.
06	**Who from your team sits with mine and how often?** Workflows that stay off-the-shelf stay generic. Iteration requires someone with both the technical ability and the business context to do it.

A self-check that needs no vendor at all: pick your last 20 overturned denials and ask what information the win depended on, and where that information lived. If most of it came from a document somebody had to go find, you have a context problem, not a productivity problem.

08 / WHAT SEPARATES THE LABS THAT GOT VALUE

One reference point: a national-scale molecular diagnostics laboratory running high-volume NGS oncology testing, which spent significant effort redesigning its workflows around automation

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

rather than layering automation onto the workflows it already had. That redesign, not the technology it bought, is what separated it from labs that saw nothing.

> **"The best processes win. AI lets you throw more resources at a problem without finding more human beings, but if you're not redesigning your process to take advantage of that, you're not going to move the needle."**
>
> SUNJAY MOULI

The labs that layer AI on top of broken processes will look back at the spend and find the needle never moved. The trap is easy to fall into because the spending itself feels like progress.

The second group of winners is doing something less discussed: redesigning their context. Removing the walled gardens around patient data, policy data, and contracts, and getting it into one place where models can use it. Results depend on what data is available. Most of that work is not AI work at all, and none of the rest of it pays off until it is done.

09 / WHY THIS IS TELCOR'S PROBLEM TO SOLVE

There is no shortage of healthcare AI vendors. There is a shortage of vendors who understand laboratory.

The nuances that determine whether a lab claim gets paid are specific and unusual, and largely invisible from a general healthcare RCM vantage point. TELCOR has spent decades inside that detail, with more laboratory customers and more accumulated knowledge of those complexities than anyone else in the category. When the difference between a paid claim and a write-off is a diagnosis buried in a record the ordering physician never sent, that depth is what decides the outcome.

SHARED CONTEXT

If your billing system is already TELCOR, the context problem described above is largely solved by architecture rather than integration work. If it isn't, the integrations into other laboratory RCM platforms and adjacent data stores already exist and have been run in production.

WORKFLOWS IN PRODUCTION

Ask us the first of those six questions. The answer is a list of named TELCOR AI Systems workflows running against live claims at real laboratories, along with the deployed engineers who tuned them, not a roadmap and not a demo.

The deployed engineering model matters more than it sounds. Off-the-shelf workflows get a lab to first value quickly, but the value compounds only when someone who understands both the technology and your operation keeps iterating on it. That is a staffing choice rather than a feature,

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

and it is why implementations move at the pace described earlier.

Being AI-native rather than AI-adjacent is ultimately a statement about where the investment goes. Those six questions cut both ways, and they are the right ones to put to every vendor in this market, including this one. There is no shortage of healthcare AI vendors. There is a shortage of vendors who understand laboratory.

ABOUT THE AUTHOR



Sunjay Mouli

VICE PRESIDENT, AI SYSTEMS, TELCOR

Sunjay Mouli leads applied AI at TELCOR, where he works with independent and specialty laboratories on deploying agentic workflows across the revenue cycle. Prior to joining Sample Healthcare, now part of TELCOR, he led revenue cycle operations and analytics at Adaptive Biotechnologies, where he oversaw reimbursement, payer operations, analytics, and technology-enabled transformation for a $200M+ molecular diagnostics portfolio.

His work focuses on applying large language models and agentic systems to complex, high-context operational processes with some of the largest diagnostic laboratories in the U.S. Sunjay holds a B.S. in Bioengineering from the University of Washington, with a focus on Medical Instruments & Devices, and is based in Seattle, WA.

SOURCES

1. American Hospital Association, Costs of Caring, March 2026. aha.org/costsofcaring
2. Availity, 2026 Abrasion Index, as reported by HIT Consultant, April 2026. hitconsultant.net
3. 3Gen Consulting, “Laboratory Revenue Cycle Management for Medicare in 2026,” July 2026. 3genconsulting.com
4. Office of Rep. Jen Kiggans, “Medical Laboratory Personnel Shortage Relief Act of 2025” announcement. kiggans.house.gov
5. HFMA and FinThrive survey of 101 healthcare organizations, as summarized in “State of AI in Revenue Cycle Management: 2026 Benchmarks,” July 2026. aiassemblylines.com

Quotations are drawn from a recorded conversation with Sunjay Mouli, 2026, lightly edited for length and readability.

©2026 Confidential and proprietary to TELCOR Inc. All rights reserved.

