Brain gain:

Why the secret to AI-powered RCM is orchestration

Pilots and point solutions won't necessarily bring true transformation. The true value in applying AI in RCM is when it's a modular, orchestrated layer that sits above existing systems, built up step by step and always aligned to a clear end-to-end vision.

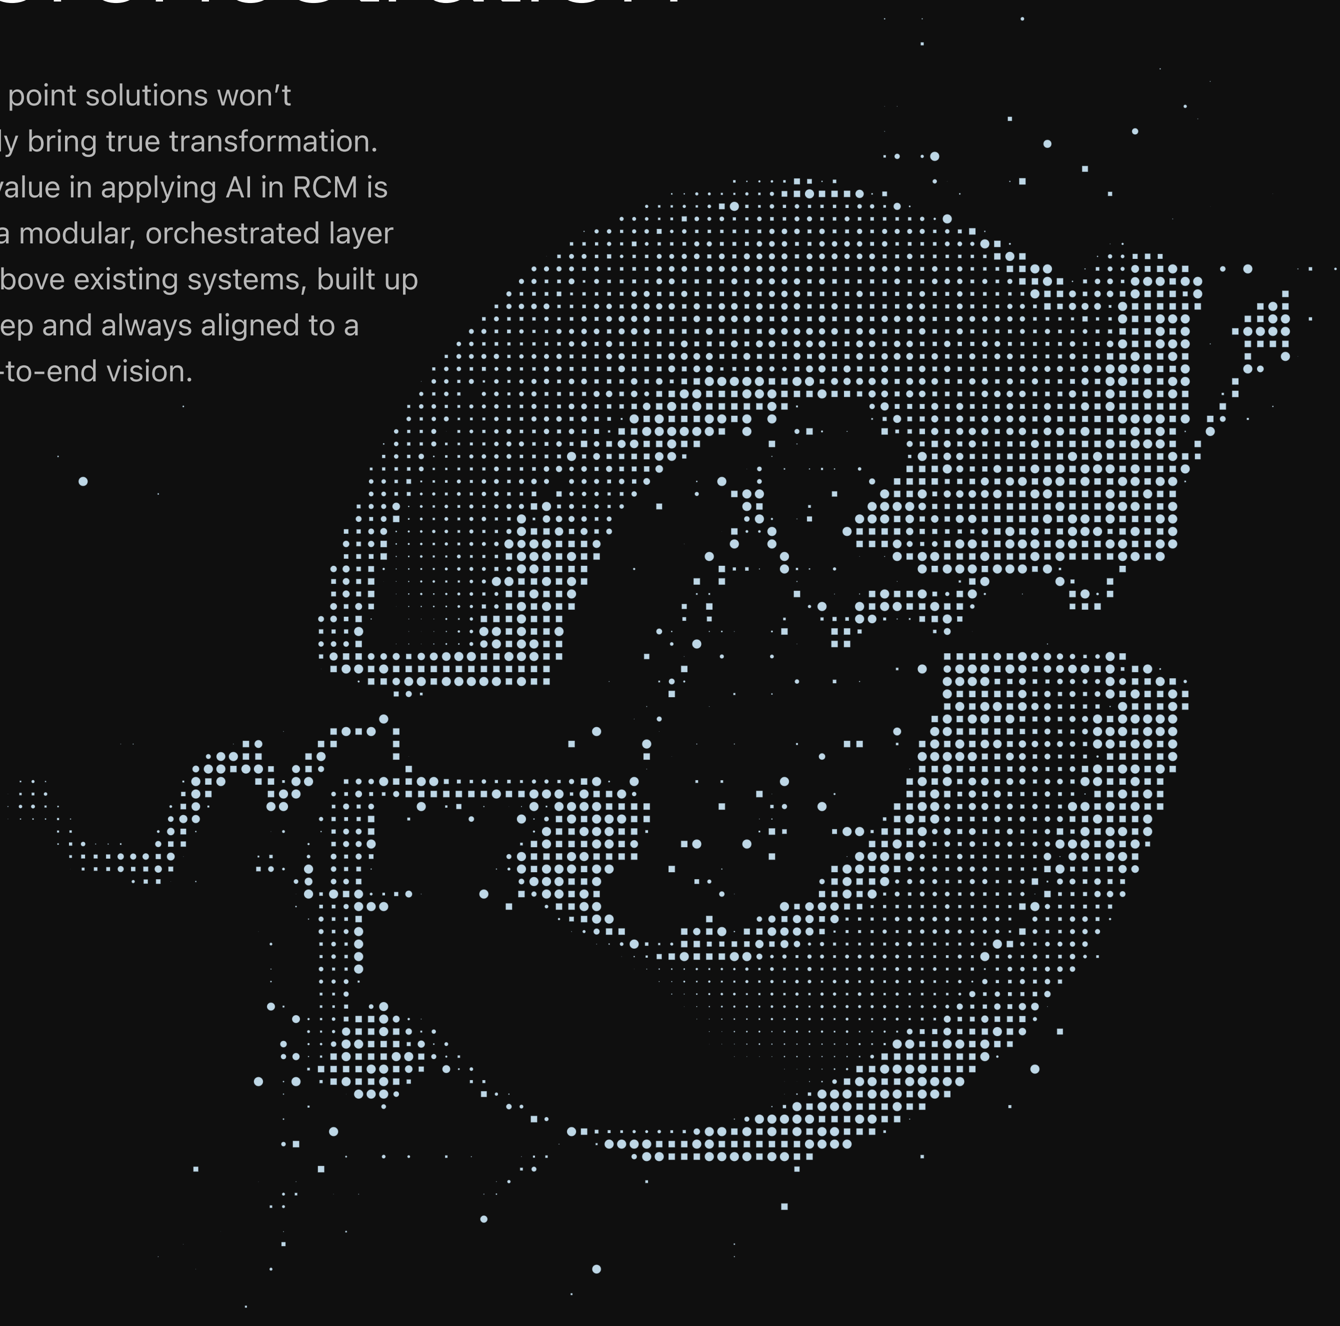

Table of contents

EXECUTIVE SUMMARY

AI is the operating system, not the app

Revenue cycle management in US healthcare systems is complex, fragmented, and under constant pressure. Leaders are asked to improve performance, manage rising costs, and cope with increasing operational complexity, often with limited resources and overstretched teams.

AI instinctively *feels* like the right move; who doesn't want better decisions, faster workflows, and more resilient economics? But when it comes to AI, many projects never move beyond pilot stages. With RCM in particular, we believe the disconnect comes from organizations still thinking about AI in terms of isolated trials for specific use cases like prior authorization or autonomous coding. These efforts may show promise in narrow areas, but they don't add up to a coherent change in how the revenue cycle process operates today.

Pilots and point solutions won't necessarily bring true transformation. We need a different way of thinking.

THE BRAIN, NOT THE LIMBS

Two steady hands are a great asset for any surgeon, but they need the brain to coordinate them effectively. Just as any intelligent system has a central nervous system, it's better to think of pilot AI tools and agents as limbs that do the work. For end-to-end complex workflows that we typically see in RCM, we need AI orchestration to guide the activity.

This paper is founded on the notion of AI as the operating system, not as apps alone. The true value in applying AI in RCM is when it's a modular, orchestrated layer that sits above existing systems, built up step by step and always aligned to a clear end-to-end vision.

At this point, some leaders may be thinking: does this mean I need to sign up to a long-term, organization-wide transformation project? Because we've seen that movie before. In healthcare, the road to modernizing systems is strewn with the wreckage of big promises and failed implementations. Projects that start with achievable deadlines have a habit of spinning into multi-year sagas.

Instead, what we propose in this paper is a modular approach that quickly builds self-standing components that work early and compound over time, building confidence and complexity towards arriving at the 'brain'. It's intentionally divided into a series of initiatives that start with a strong win. We're talking about validation within weeks rather than months or years, and value measured in impact that's felt in cost efficiency and productivity.

> The ultimate goal is an end-to-end solution but the key is not to start doing that all at once.

The ideas in this paper are presented in accessible terms for non-technical leaders, while going into enough depth to clarify what practical steps are needed.

PART I

Start small, think big

Why you shouldn't start with the whole enterprise.

Most health systems have data and processes spread across many systems and silos. The complexity of integrating into every workflow, harmonizing all data, and aligning all stakeholders can stall progress before any benefits appear. A more realistic approach is to treat the early stage like a controlled experiment rather than a full-scale transformation. Think of it like testing a treatment in a scientific environment before moving to clinical trials on real patients. Begin in a 'test tube', where conditions are simpler and easier to manage.

The sandbox environment

In practice, this means building a sandbox environment. Instead of plugging directly into live systems, you identify a single, clearly defined use case, export a subset of historical data relevant to that use case and then work in a controlled, static environment where the AI system can be developed and tested safely. The aim at this stage is not to deliver a production-ready solution but to answer a more basic question: Does the technology actually work on our data?

SANDBOX

A controlled, static environment built on exported historical data, where the technology can be developed and tested safely.

→

PRODUCTION

Live workflows on real patients. More complex architectural and workflow decisions come later, by design.

By focusing on a narrow use case and working with exported data in a sandbox, you avoid the heavy upfront work of full integration. You can move faster and generate early evidence of effectiveness, while postponing more complex architectural and workflow decisions to a later phase.

PART I · CONTINUED

Concrete proof-of-concept examples

Here are two examples of how such proofs of concept might look in the context of revenue cycle management.

EXAMPLE A

Denial simulation or prediction

In this scenario, the health system takes a set of historically submitted claims, where the final outcomes (paid or denied) are already known. The process then looks like this:

1. Mask the outcomes in the dataset
2. Train an AI model on the features of the claims to predict which ones are likely to be denied
3. Run the model on the masked data to generate predictions
4. Unmask the actual historical outcomes and compare them with the model's predictions

Because the real outcomes are known, the organization has a deterministic way to see how closely the model matches reality. This builds confidence that if the model performs well on historical data, it is likely to perform well on future data.

EXAMPLE B

AI-assisted coding

Another example involves coding. Here, the organization takes encounter data along with a batch of historically coded claims that include CPT and ICD-10 codes. Then it runs the encounter data through an AI coding model that proposes codes and compares those AI-generated codes with the actual codes that were submitted.

- Take encounter data with historically coded CPT and ICD-10 claims
- Run it through an AI coding model that proposes codes
- Compare the proposed codes with those actually submitted

Again, because the 'true' codes are already known, this becomes a controlled test of how accurate and reliable the model is. Fundamentally, the idea with both of these scenarios is the same: mask the outcome, make a prediction, and check how close the model was to the real-world decision. The goal is not a live deployment into production workflows within six to eight weeks, but to demonstrate within that timeframe, that the technology is capable of doing what it claims using your own historical data.

PART I · CONTINUED

Understanding compounding: beyond one-off cost savings

Compounding happens two ways: within a single use case, and across multiple use cases. Within a single use case, compounding comes from feedback and data. As humans interact with the system, correct or refine its outputs, and feed it more examples, the model can gradually improve. Over time, its decisions get better, and the amount of manual effort required decreases. Across multiple use cases, compounding arises when those use cases are designed to fit together. Instead of delivering isolated benefits, they feed into one another along the life cycle of an encounter becoming a claim. This is where the effect can become significantly greater than the sum of its parts.

A sequential example: from encounter to submitted claim

Here's a concrete sequence along the encounter-to-claim journey:

1. **Encounter routing**

   An AI system examines incoming encounters and classifies them into high, medium, and low complexity. Each category is routed to a different path or team. High-complexity cases might go to a highly specialized group, mid-level to another, and low-complexity to a third. Even before any automation, this specialization can save time. People become more expert in the types of work they handle, and routing reduces the friction of mismatched work.

2. **Auto-coding for low-complexity encounters**

   For simpler encounters, an AI model can handle the coding step autonomously. Many low-complexity cases can bypass human coders entirely, freeing them to focus on the more nuanced cases. This adds another layer of benefit. You begin with time savings through better routing and build on that by automating a significant portion of the volume.

3. **Agentic provider engagement for complex cases**

   High-complexity encounters often require additional information from providers to code them accurately. Traditionally, this involves back-and-forth communication that can be slow and inconsistent. An agentic AI system can automate much of this outreach, requesting clarifications and additional documentation, and following up as needed. This further reduces manual effort and accelerates resolution.

Each step adds its own savings, but because they apply sequentially to overlapping parts of the workload, the total effect grows in a compounding way. You may start with a modest percentage of time saved in one step, but as you add more aligned use cases along the journey, the combined impact becomes materially larger.

PART II

What this means for leaders

Thinking in terms of compounding pushes leaders beyond the mindset of 'one successful AI pilot'. It encourages them to ask how each use case will connect to the next, and how the overall journey from encounter to claim can be redesigned so that each improvement amplifies the others. It also reinforces the idea that meaningful results can appear in months rather than years, provided that use cases are selected and sequenced thoughtfully.

Purposeful connectedness

ONE BRAIN, DENSELY CONNECTED

Imagine an adult human brain, which has around 86 billion neurons. The neurons are densely interconnected. They form loops and networks that continuously feed into and out of each other. This structure makes complex reasoning and behavior possible.

86 BILLION ISOLATED NEURONS

Imagine 86 billion snails, each with one neuron. Even though the total number of neurons may be similar, billions of separate snails, each acting in isolation, just don't have that same capacity.

The same concept applies to AI in revenue cycle management. A health system can have many small, disconnected AI tools: one for prior authorization, one for coding, another for denial management, and so on. Each may work reasonably well on its own, but if they're not designed to work together, the overall effect is limited. Purposeful connectedness means designing AI use cases, data flows, and workflows so that they form a coherent system, more like a brain than a collection of isolated parts.

PART III

The AI operating system

A modular AI layer that sits above your existing tools.

For what we're proposing with RCM, it helps to think of AI as an operating system that sits as an intelligent layer above existing tools and systems. Rather than replacing every underlying application, the AI layer:

THE AI OPERATING SYSTEM

Understands	Routes	Orchestrates
the data that flows through the revenue cycle	work to the right paths or teams	the use of various automation tools and models

SITS ABOVE ↓

EXISTING TOOLS & SYSTEMS

EHR | Billing | Coding tools | Clearinghouse | Point solutions

This isn't something that appears all at once. It's built up over time through a series of modules, each corresponding to a use case or a set of related capabilities.

Why modularity matters

Modularity is important for two reasons. First, it avoids a false choice between 'do nothing' and 'attempt a complete end-to-end AI operating system from day one'. Health systems can hold a clear vision of where they want to go, while still starting with a narrow, practical first step. Second, it allows for a phased approach in which each new piece is added only when there is sufficient confidence and when it clearly supports the overall strategy. Over time, these modules together form a coherent, orchestrated layer.

PART IV

The journey to unified data and intelligence

A central challenge in revenue cycle management is that data is dispersed across many systems. Different applications, departments, and vendors hold pieces of the picture. A key first step in the maturation journey is to unify the data, at least conceptually. This doesn't necessarily mean physically moving everything into a single database, but it does mean being able to see and work with the data as a coherent whole for the purposes of AI and orchestration.

Standardizing data

Once data is unified, it still needs to be standardized. Different systems may record similar information in different formats, with different field names, or with inconsistent usage. Standardisation involves aligning these representations so that the AI systems can interpret and work with them consistently. This reduces errors, simplifies model development, and makes it easier to extend solutions across new areas.

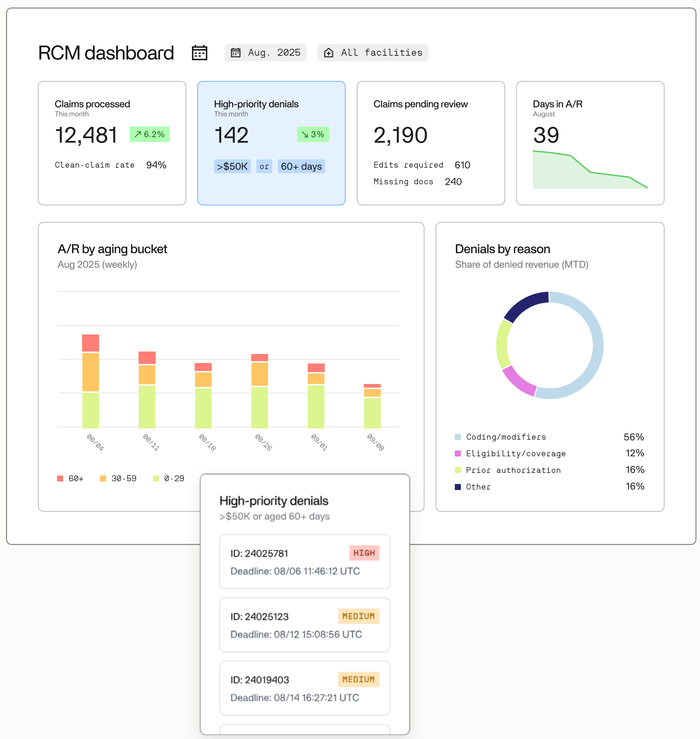

Developing an ontology

The next step is to make the data "intelligent" by developing an ontology or structured way of representing the entities and relationships within the revenue cycle. This might include encounters, claims, codes, providers, statuses, and so on, along with how they relate to one another. An ontology helps the AI layer to understand what different pieces of data mean and how they connect, which in turn supports smarter routing, coding, denial prediction, and other decisions.

PART IV · CONTINUED

Building the action layer

On top of unified, standardized, and semantically understood data, organizations can then build the action layer. This is where visible change happens for frontline teams:

- Routing encounters by complexity
- Automating coding for appropriate subsets
- Orchestrating communication with providers for complex cases
- Coordinating multiple tools and agents across longer workflows

This combination of data foundation and action layer is what brings the idea of an AI operating system into practical reality.

Timelines and expectations

Every health system's environment is different, and rigid timelines are rarely helpful. A more flexible way of communicating expectations:

Weeks

Technology validation in weeks, using sandbox proofs of concept on historical data.

Months

Meaningful results in months, not years, as selected use cases move into production and begin to compound.

This framing sets ambition without overpromising and recognizes the complexity of each organization's ecosystem.

PART V

Mindset, impact, and the platform model

Find a design partner, not a vendor

Because AI for revenue cycle management is still a relatively new and evolving area, many organizations are looking not just for tools, but for design partners. These are collaborators who will help shape a roadmap of use cases, share responsibility for experimentation, and work through challenges that inevitably arise when theory meets real-world operations. This approach acknowledges that no-one has a finished playbook. Instead, leaders and partners work together to design and refine the system over time.

Frame impact beyond pure cost savings

Cost savings are important, but focusing solely on them can be limiting. Some use cases, like denial simulation, are primarily about revenue upside rather than cost reduction. A more balanced way to frame impact is across three dimensions:

UPSIDE POTENTIAL

For example, identifying and preventing avoidable denials.

COST EFFICIENCY

Reducing manual work and reallocating people to higher-value tasks.

PRODUCTIVITY

Increasing the volume of work that can be handled without a proportional increase in resources.

This broader framing can help secure support from across the organization, not just from those tasked with cutting budgets.

PART V · CONTINUED

Why a platform matters

Healthcare operates at the "speed of trust", and trust depends on more than model accuracy in a sandbox. It depends on the reliability and safety of the entire system in production. A platform model provides shared capabilities that every AI use case needs, such as:

- Auditability of actions and decisions
- Robust data standardization and ontology
- Human-in-the-loop controls
- Fallback mechanisms when systems are uncertain or fail
- Monitoring and performance tracking of AI workflows
- Operational support and uptime

Without such a platform, organizations risk duplicating effort across pilots, and they face higher chances of failure when moving from controlled experiments into live workflows. Many AI experiments fail when they leave the lab and enter real-world settings, especially if there's no mature infrastructure and governance behind them. A platform helps mitigate these risks and ensures that as new modules are added, they sit on a stable, shared foundation.

CONCLUSION

A practical path forward

01 **Clarify the vision**

Define the long-term ambition for AI in RCM. What does an orchestrated, AI-enabled revenue cycle look like for your organization?

02 **Choose a first wedge use case**

Select a narrow, well-bounded problem such as denial simulation or auto-coding where historical data is available and outcomes are already known, which allows for more accurate testing of the results.

03 **Run a sandbox proof of concept**

Export a defined historical dataset, mask the outcomes, run predictions or coding, and then compare with the real results. Use this to validate the technology and build confidence.

04 **Map a compounding roadmap**

Identify sequential use cases along the encounter-to-claim journey. Think explicitly about how each will build on the previous ones and contribute to a coherent system.

05 **Lay the data and platform foundations**

Begin the work of unifying and standardising data, defining an ontology, and establishing the monitoring, audit, and human-in-the-loop capabilities that will support multiple use cases over time.

06 **Communicate clearly and accessibly**

Use simple analogies such as test tubes vs. clinical trials, or the example of neurons in the human brain to explain the approach to boards, clinicians, and operational leaders.

Taken together, these steps shift from scattered, disparate pilots to purposeful, modular orchestration. Health leaders don’t need to have all the answers upfront, but delivering impactful AI that calls for a clear sense of direction and a willingness to build the ‘brain’ of revenue cycle management one carefully connected piece at a time.

Invisible | invisibletech.ai

Your care is integrated. Your systems aren't. Revenue leaks at every step of the cycle, from coding and clearance to collections and denials, while your teams carry the manual work.

If this describes you, let's talk.

We build custom AI for healthcare operations, automating coding, financial clearance, AR prioritization, and denial defense across every facility. The result is faster cash collection, lower denial rates, and fewer AR days.

REQUEST A DEMO ↗

ABOUT THE AUTHOR

Cesar Yazbeck, MD

VP of Healthcare & Life Sciences, Invisible Technologies

Cesar Yazbeck, MD works with enterprise health systems, payers, and life sciences organizations on AI-driven transformation.

A physician by training, he has spent his career operating at the intersection of clinical knowledge and enterprise AI. At CVS Health, he served as Executive Director of Care Delivery Growth & Applied AI, building and scaling a portfolio of AI-driven initiatives across the enterprise. Before that, he was a Senior Engagement Manager at McKinsey & Company, advising payers, providers, and life sciences organizations across North America, EMEA, and APAC, and led AI partnerships for McKinsey Digital.

Having sat on the buyer's side of the table, run transformation programs, and answered for the outcomes, he approaches AI in healthcare with a clear conviction: the constraint is rarely the model. It's the data foundations, the workflows, and the organization around them.

At Invisible, he works with health system, payer, and life sciences leaders to move beyond scattered pilots and into orchestrated, AI-native operations built to compound.