The AI Catalyst Pulse

March 25th, 2026

Upcoming AI Catalyst Events

Mark your calendars

- **May 6th: Rev Cycle AI: Epic System's Penny and AI-Driven Solutions for Revenue Cycle Management Within the EMR** | Register Here
  Join us for an exclusive deep dive into Epic's Penny, the AI-powered revenue cycle assistant reshaping how health systems manage billing and administrative workflows. Epic will provide a comprehensive overview of Penny's agentic AI capabilities - from automating professional billing coding and generating denial appeal letters to streamlining claims follow-up and reducing coding-related denials.
  Recommended audience: Leaders responsible for billing accuracy, denial management strategy, and balancing competing AI investments across their organizations.

- **April 9th: Nursing AI: Epic Systems Presents Their Latest AI-Powered Solutions for Nursing Documentation** | Register Here
  Join us as Epic Systems presents an exclusive look into the future of AI-powered nursing documentation and how health systems can strategically align their AI investments to improve nursing efficiency and workflows.
  Recommended audience: C-suite and senior executives responsible for enterprise AI strategy, governance, and clinical operations.

- **June 18th: 2026 National Survey Readout: How Does Your AI Adoption Compare?** | Register Here
  Join us for the release of the 2026 national AI survey results, featuring data from leading health systems on where AI is scaling, where it's generating returns, and where it's quietly being pulled back. During this 60-minute live session, we'll walk through national deployment trends across clinical workflows, revenue cycle, patient throughput, clinical decision support, supply chain, and AI governance.
  Recommended audience: Healthcare executives, clinical leaders, digital officers, and strategy teams seeking grounded, real-world insight into AI performance across the industry.

Coming at You "Too Slow and Too Fast" – The Real Signals From HIMSS 2026

What happened

HIMSS 2026 did not lack announcements – if anything, it lacked restraint.

The event was full of launches, pilots, proof points, and bold claims, but underneath the volume, a more important pattern emerged: healthcare AI is accelerating faster than organizations' ability to measure it or govern it.

This is not net new. We flagged the same "cart before the horse" dynamic back in 2023. What changed at HIMSS 2026, however, is the *magnitude*.

That tension showed up repeatedly across the conference. AI is moving quickly into the operational core of health systems, while many of the systems are still catching up on evaluating impact, assigning accountability, and governing risk.

Three signals stood out most.

1. **AI agents are becoming the center of healthcare operations.**

   One of the clearest themes from the HIMSS floor was the growing role of AI agents and agentic workflows across the enterprise. A few stood out:

   - Epic doubled down on agentic workflows embedded across documentation, billing, and patient communication.
   - Amazon and other major players signaled similar moves toward AI-driven coordination across fragmented workflows.

   This is a structural shift. The conversation goes beyond augmenting discrete tasks. The larger opportunity (and larger risk) sits in how AI begins to shape the flow of work across clinical, operational, and financial domains. As more vendors position themselves to coordinate that work, health systems will need to think more carefully about workflow ownership, accountability, and where operational control is starting to sit. When agents sit across clinical, operational, and financial workflows, decision-making becomes distributed across platforms rather than owned by a single function.

   The strategic question is becoming less about whether a model performs well in isolation and more about who is defining how work moves through the system.

2. **"Show me the ROI" is now the default posture**

   Health systems made one thing clear: they are done with demos.

   Vendors responded accordingly, leading with proof points around documentation efficiency, denial reduction, throughput gains, and other performance-oriented outcomes. At the same time, many organizations still lack a consistent way to validate those outcomes on their own terms. Baselines are inconsistent, attribution is difficult, and enterprise-wide measurement remains underdeveloped in many settings.

   As a result, the constraint on AI scale is increasingly tied to an organization's ability to measure impact credibly. The challenge is no longer simply getting access to tools. It is building the internal discipline to determine which tools are generating meaningful value, under what conditions, and with what tradeoffs.

   That dynamic closely mirrors what we are seeing across THMA research. Adoption is moving quickly, but the infrastructure to compare, defend, and operationalize ROI is far less mature.

3. **Governance and validation are under the most pressure precisely where adoption is accelerating fastest.**

   As deployment accelerates, so do concerns around validation, safety, and oversight.

As AI spreads into more workflows and more business areas, governance models are being asked to do more than many were designed for. Validation processes remain uneven. Monitoring capabilities vary widely. In many organizations, oversight still depends on fragmented committees, limited staffing, or reactive review processes.

This is the "too slow and too fast" paradox we're seeing.

The pace of deployment is not likely to slow, particularly in high-pressure areas where leaders are looking for productivity gains, financial relief, or workforce support. That puts pressure on governance, validation, and monitoring functions to mature much faster than they have in the past. The implications extend well beyond clinical safety. Weak oversight can create operational instability, muddy ROI claims, and expose organizations to reputational risk if tools scale before appropriate controls are in place. For many health systems, the real challenge is not whether they can adopt AI quickly, but whether they can build enough structure around that adoption to sustain it responsibly.

Why it matters

The biggest misconception coming out of HIMSS is that AI is still a layer being added to healthcare. It is not. It is becoming entrenched in the systems that run healthcare.

That kind of embeddedness raises the stakes for health systems. It requires more than enthusiasm for innovation. It requires clearer accountability, stronger validation methods, and a more mature approach to measuring value at scale.

Healthcare organizations are moving into a period where speed alone will not be enough. The winners will be the ones that can pair rapid adoption with a much more disciplined approach to governance and ROI.

Questions to consider

1. Where are AI agents beginning to influence multi-step workflows in your organization, and who is accountable for the full outcome across that chain?
2. In high-investment areas like ambient documentation or prior authorization, how are you independently validating value beyond vendor-reported results?
3. Which parts of your current AI portfolio are scaling faster than your governance or monitoring processes can realistically support?
4. What capabilities would your organization need to strengthen now in order to evaluate AI consistently across domains over the next 12 to 24 months?

What Healthcare Leaders Can Learn From Law's AI Crisis

This is the first installment in a series examining how other professions have adopted, governed, and failed at AI - and what their experience means for health system leaders making the same decisions today.

AI is transforming the legal profession, for better and for worse. Harvard Law School's Center on the Legal Profession reports that some AmLaw100 firms have seen task-level productivity gains exceeding 100x. At the same time, court cases involving AI-generated hallucinations have been

documented worldwide – up from roughly 300 as of mid-2025 – with new incidents surfacing two to three times per day.

Here are three major legal AI use cases that offer particularly direct lessons for healthcare.

Automated Case Law Research: 1 in 6 Answers Are Wrong

What it is

You ask a question grounded in an authoritative knowledge base; the tool returns an answer. In law, that means querying case law databases to find relevant precedents and surface legal authorities. In healthcare, the same pattern shows up in clinical decision support, diagnostic AI, and evidence synthesis tools.

How it's used today

The leading platforms – Lexis+ AI (LexisNexis), CoCounsel (Thomson Reuters), and Harvey – are embedded in daily workflows at major firms. Harvey alone serves 42% of the AmLaw100. Thomson Reuters reports that 26% of legal organizations are actively using generative AI, while Clio found 79% of individual legal professionals use it — a gap that says more about organizational governance than about the technology itself.

How it goes wrong

When Stanford researchers tested the leading legal AI research tools, products marketed as safer than general-purpose chatbots, they found hallucination rates of 17 to 34%. LexisNexis had promoted Lexis+ AI with the headline claim of "hallucination-free" linked citations; that assurance applied narrowly to citation formatting, not to the accuracy of the analysis built around those citations. The consequences are now playing out in courts nationwide: over 1,000 documented cases of AI-generated legal hallucinations, up from roughly 300 as of mid-2025.

Healthcare takeaway

Vendor accuracy claims require independent validation in your specific clinical context. A 2025 Communications Medicine study from the Icahn School of Medicine at Mount Sinai found that when researchers deliberately embedded fabricated clinical details into prompts – fictitious lab tests, invented syndromes – six leading LLMs repeated or elaborated on the planted errors 50 to 82% of the time. That's a measure of adversarial susceptibility, not a baseline error rate. The lesson from law is direct: purpose-built tools marketed as safe inherit the same architectural vulnerabilities. The question isn't whether your CDS vendor is trustworthy. It's who in your organization owns independent validation and what happens when the tool is confidently wrong.

AI Contract Review and Redlining: Governance-First Firms Are Pulling Away

What it is

AI ingests source material, compares it against established standards, and generates a structured output for professional review. In law, that means feeding a draft contract through a precedent library and producing redline edits for attorney sign-off. In healthcare, ambient documentation works the same way: AI listens to a patient-clinician conversation and drafts a clinical note for the physician to review and sign.

How it's used today

A&O Shearman, an elite global firm, is the governance-first model. Their deployment followed a deliberate sequence: sandbox first (limited attorneys in a ring-fenced environment), governance protocols before scale (use cases and risks identified, verification workflows designed), then enterprise rollout. Today their ContractMatrix tool supports 4,000 staff across 43 jurisdictions, with roughly 2,000 lawyers using it daily. The firm reports saving about seven hours per contract review — a 30% efficiency gain — with every AI-suggested edit surfaced as tracked changes for human review. The key design choice: AI proposes amendments through a familiar tracked-changes interface. The human review step is the product architecture.

How it goes wrong

Without that architecture, outcomes deteriorate fast. Gordon Rees, an AmLaw100 firm with 1,800 lawyers, admitted in late 2025 that one of its attorneys had submitted bankruptcy filings containing AI-fabricated citations. The firm paid over $55,000 in remediation, updated its policies, and created a cite-checking protocol. Then it happened again: a December 2025 judicial reprimand. Policy changes didn't change the workflow; an individual lawyer under time pressure, without a structured verification step, produced the same errors repeatedly.

Healthcare takeaway

The verification architecture matters more than the verification policy. Health systems deploying ambient documentation face the same design choice. When a clinician reviews and signs an AI-generated note under time pressure, the structural incentive is to skim and approve. The firms that avoided failures designed tools where checking the work was the default interaction. If your workflow routes a finished note to a clinician for a single signature, that's a policy. If it surfaces specific AI-generated elements for targeted review, that's architecture.

AI-Drafted Court Filings: 100x Faster, but Errors Compound Without Structured Oversight

What it is

AI produces a complete first draft of a professional document that a human then reviews, edits, and submits. In law, that means generating briefs, motions, and complaint responses for attorney sign-off. In healthcare, the parallel is AI-assisted coding, prior authorization letters, and discharge documentation – any workflow where AI generates a finished output that goes out the door under a clinician's name.

How it's used today

This is the highest-volume legal AI use case and the most publicly scrutinized. Harvard Law School found that one complaint-response system cut associate time from 16 hours to 3-4 minutes. The ABA's 2025 task force concluded that the profession has moved from asking whether to use AI to figuring out how to govern it, and early adoption has concentrated on exactly this kind of high-volume drafting.

How it goes wrong

Courts have drawn a hard line: the machine may hallucinate, but the professional cannot. Bar associations in New York, California, and Florida have issued guidance making attorneys personally responsible for AI errors, even unintentional ones. But accountability rules alone

haven't stopped the problem. The Gordon Rees pattern – update the policy, repeat the failure – shows that individual responsibility without workflow-level verification produces the same errors, with compounding reputational damage.

Healthcare takeaway

Individual accountability without structural safeguards produces repeated, compounding failures. The legal profession's forward-leaning firms didn't stop at accountability rules; they restructured verification into the workflow itself. Health systems relying on training and policy alone to manage AI risk in documentation, coding, or clinical decision support should expect similar results. When a payer or regulator flags an AI-related error in a claim or clinical note, the question won't be whether you had a policy. It will be whether you had a process.

What Law Learned That Healthcare Hasn't - Yet

Across all three use cases, the dividing line is the same: governance built before scale versus governance retrofitted after failure.

A January 2026 piece in Understanding AI, based on interviews with lawyers at top-20 firms, captured the dynamic. One junior associate described their firm's "fast follower" posture as a "huge mistake," arguing that elite institutions cannot afford to be followers in anything core to their business. The lawyers weren't AI skeptics – they understood the technology's value – but were still leaving gains on the table. Caution, it turns out, isn't the same as governance.

Healthcare's governance gap is quantifiable. A November 2025 Black Book survey of 182 hospital leaders found that only 22% are highly confident they could produce a complete AI audit trail within 30 days for regulators or payers. Just 29% have implemented and enforced AI model inventory policies. Nearly half are still drafting them.

The legal profession's correction offers a concrete model: define who verifies AI outputs before they leave the building, build that verification into the workflow architecture rather than the policy manual, and treat errors as governance signals rather than individual failures.

Next in the series: Aviation built the governance architecture before it scaled the technology. Here's what it looked like – and what healthcare can take directly from their playbook.

Medical coders will outlast AI; it's only their roles that will be replaced

Health systems do not have a coding labor problem alone. They have an expert-capacity allocation problem.

That may sound like semantics, but it's not.

The market still likes to talk about medical coding as if the central question is whether AI can replace coders, but that framing misses what is happening inside most revenue cycle organizations right now. Coding teams are already stretched. Work is getting more complex. Backlogs are not just annoying. They create denial risk, timely filing risk, rework, and revenue leakage.

The real question is not whether AI can code. It is whether AI can keep scarce expert talent from being swallowed by the wrong work.

Transforming your rev cycle shop from cost-center to expertise-center

Our recent AI Catalyst Executive Deep Dive on the AI-Driven Revenue Cycle summed it up well: growing volumes overwhelm coding, denials, and payment variance teams; backlog forces prioritization based on time instead of value; and cost-to-collect rises as rework becomes routine.

And once that dynamic starts, it compounds. The work that gets delayed is often the hardest, most specialized, or most financially sensitive (especially in terms of having the highest reimbursement values). Experienced coders spend time clearing repetitive first-pass volume while higher-risk cases, denial-prone areas, and specialty backlogs wait their turn – the net result being delayed payment.

That is a bad use of expert labor even before AI enters the picture.

So what is the best use of time for an AI-enabled medical coder?

One of the ways AI Catalyst approached this question was through its custom GPT for job redesign in the AI era. Instead of asking generically whether a role is "safe" from AI, the tool goes line by line through a job description, breaks the work into tasks, and asks a harder question: which parts are repetitive and rules-based, which parts are judgment-heavy, and what should this job look like if it were redesigned for the AI era?

Our team looked at multiple real job postings for medical coders across LinkedIn and fed them into our custom GPT for analysis. For medical coders, that analysis landed on something important. The role is not disappearing. But its premium value is moving.

In the old model, value came from manual chart review, code assignment, abstraction, edit checks, and individual expertise.

In the AI-assisted model, the coder increasingly reviews AI-summarized charts, validates AI-suggested codes, and governs automated logic. In short, AI handles repetitive and rules-based work; the coder verifies and governs.

And then, in the AI-transformed model, the coder becomes the person who handles complex cases, controls quality, manages exceptions, supports audit defense, and helps build the workflows that let the system scale.

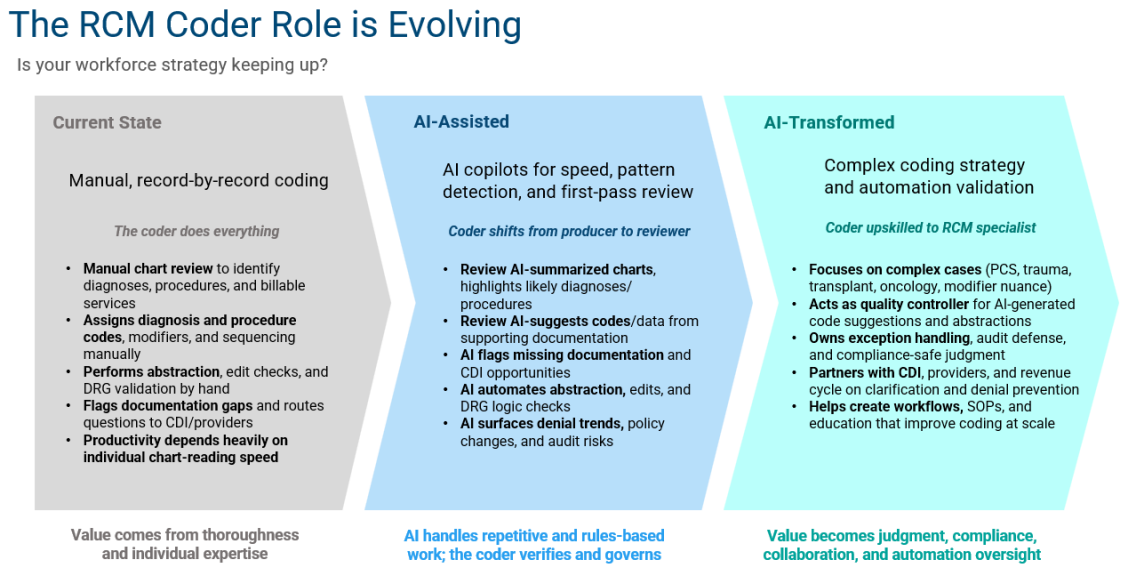

That is the shift executives should be planning around now.

Because once AI takes on more first-pass work, the obvious question becomes: what do you want your best coders doing instead?

Calculating true ROI from AI adoption for rev cycle means moving beyond the automation rate

A good example came from University of Vermont Health Network's diagnostic radiology workflow. Casey Webb, System Director, Professional Billing at UVMHN, described how the organization used autonomous coding in that domain and reached a straight-to-bill rate of about 76%.

Those numbers are impressive enough to make a slide, but the more important point is what UVM did with the capacity it created. Webb explained that once coders no longer had to review 100% of those encounters, the team took the opportunity to move two FTEs to more valuable work. Those resources were redirected to coding groups with the highest backlog, including specialty and primary care areas where the operational risk of delay was greater.

That is the kind of detail executives should pay attention to. The most strategic value of coding AI may not be that it automates a chunk of radiology. It may be that it lets you redeploy scarce coding expertise into the places where backlog hurts most.

Webb made the principle explicit: *"Using the ability of the coders to their highest skill set is something that I think is extremely important. If you can have technology as a tool, and allow your teams, whatever role they're in, to function to their highest ability, you will see the ROI on that."*

AI's most credible role in coding is prioritization and selective trust

This is where many conversations about coding AI get sloppy. The debate is often framed as a binary – either AI can code autonomously, or humans remain indispensable. Either it replaces the labor, or it does not matter. That is the wrong frame.

There is now a growing set of coding tasks that software can help with meaningfully: summarizing records, suggesting codes, populating abstracting fields, running edits, and flagging likely documentation gaps. There is also a clear set of tasks that still belong with experienced humans: complex-case interpretation, modifier nuance, compliance-sensitive escalation, denial defense, and judgment in the face of conflicting or incomplete documentation. This model reframes AI as a "co-pilot" and a "second validating set of eyes" for coding decisions with reimbursement implications.

The win is not that AI suddenly becomes trustworthy everywhere. The win is that teams stop spending scarce expert attention on routine work that can be handled, or at least narrowed, by software first.

In other words, AI matters not because it removes the need for coders, but because it changes who should be reviewing what.

What health system leaders should do next

1. Stop treating coding AI as a labor-reduction story. It is a capacity-allocation story. Find where your best coders are still buried in repetitive first-pass work, then ask where that expertise would create more value if it were freed up.
2. Redesign the role on purpose. If coders are moving from producer to reviewer to automation overseer, then job descriptions, quality metrics, training pathways, and productivity expectations

should change too. If they do not, your organization will install AI into a workflow still built for manual production.
3. Decide who owns the validation layer. As AI takes on more repetitive revenue cycle work, someone has to own the oversight: who reviews outputs, who escalates exceptions, who monitors drift, and who is accountable when the machine is wrong. AI Catalyst's broader research on governance has pointed to the need for dedicated operator roles that do exactly this as AI workflows mature.

These steps point to a simple truth.

The future of coding is not a future without coders. It is a future where the best coders stop spending their days doing work that software can meaningfully assist with. And the health systems that move first on AI-enabled job redesign will not just automate more charts. They will make better use of the expertise they already cannot afford to waste.

Questions for consideration

1. If coding volume rose 10–15% tomorrow, where would backlog and revenue leakage show up first?
2. What portion of your strongest coding talent is still consumed by repetitive first-pass work?
3. If AI freed 1–2 FTEs worth of coding capacity, where would you redeploy that expertise first?
4. Are your coder job descriptions and productivity metrics still built for a manual-production world?
5. Who in your organization owns validation and governance as coding workflows become more AI-assisted?

AI Strategy Quick Hits

Noteworthy moves from peers to implement AI technologies

Catch up on the latest health system returns from AI investments:

- **Adventist Health** cuts stroke transfer times 44% with **Viz.ai** platform and standardized protocols
- **Erlanger** adds ~220 surgical cases per month and projects ~5x ROI with **Qventus** OR optimization platform
- **Atlantic Health** ACO gains $15M in MSSP shared savings after deploying **Innovaccer** clinical intelligence
- **WVU Medicine** expands **Abridge** ambient AI to 2,800+ clinicians, reports 61% cognitive load reduction
- **Allegro Pediatrics** saves ~$2,000/month per provider by switching to **Greenway Health** AI ambient scribing
- **Parkview Health** reports rapid results from 18 Epic generative AI features deployed in three months

Other strategy quick hits:

- **Sutter Health** scales cancer-detection AI across 26 hospitals with **Ferrum Health**
- **WVU Health System** selects **hellocare.ai** for AI-assisted intelligent hospital rooms across all 25 hospitals

Emerging Use Cases

New capabilities that indicate AI's potential

- **Emory Healthcare** and **Mass General Brigham** benchmark foundation models across 40 institutions via Healthcare AI Challenge
- **Cornell** researchers build autonomous LLM scam caller to expose AI fraud risks and push for safeguards

Market Moves

A round-up of AI company announcements and stories

- Epic launches Agent Factory and Curiosity foundation models at HIMSS26
- **Epic** shares AI outcome metrics at HIMSS26: 20–32% faster documentation, 69% early lung cancer detection at **The Christ Hospital**
- **Google Cloud** showcases Gemini-powered agents at HIMSS26 with **CVS Health**, **Highmark**, **Humana**, **Waystar**, and **Quest Diagnostics**
- **Microsoft** expands Dragon Copilot integrations at HIMSS26 with **Sentara**, **Mercy**, and **University of Rochester** among pilots
- **FDB** launches Script Agent and VerifyAssist AI Rx tools at HIMSS26, projecting ~70% prescribing documentation time reduction

Other market moves:

- **Universal Health Services** acquires **Talkspace** for ~$835M to build national AI-enabled virtual behavioral health platform
- **CVS Health** and **Google Cloud** launch Health100 AI platform to unify fragmented healthcare data
- **Salesforce** releases six new healthcare AI agents for referrals, EHR exchange, epidemiology, and hospital operations
- **Brainomix** deploys stroke AI imaging across all 25 **WVU Health System** sites
- **Samsung** and **b.well** expand partnership for digital check-ins and consumer health record sharing
- **Suki** and **Optum Real** partner on documentation-to-revenue-cycle alignment
- **Microsoft** unveils Copilot Health as consumer-facing AI companion for medical records and test interpretation
- **Marshall Health Network** launches **Notable** AI-driven scheduling and registration pilot across primary care

The News in Numbers

An interesting data point that caught our eye

81% of participating physicians in a recent AMA study reported that they currently use AI professionally

Expert Insights

For further reading, articles, videos, and podcasts that we found insightful

- *STAT*, "Blue Cross Blue Shield says data back up claim that AI is driving up medical bills"