

# The AI Catalyst Pulse

February 25<sup>th</sup>, 2026

## Upcoming AI Catalyst Events

*Mark your calendars*

- **March 3<sup>rd</sup>: AI 101 Mass Education: AI 101 for Healthcare Professionals** | [Register Here](#)  
Join a dynamic session designed to equip leaders and staff with essential AI knowledge. Gain a clear understanding of AI's core types and applications, along with real-world examples of its impact in healthcare.  
**Recommended audience:** All leaders and staff at your organization
- **March 18<sup>th</sup>: Executive Deep Dive: Navigating the AI-Driven Revenue Cycle** | [Register Here](#)  
Join us for a panel discussion highlighting how leading health systems are achieving early wins in AI-driven rev cycle solutions, from improved revenue to top-of-license practice.  
**Recommended audience:** CRCO, CMIO, CIO, and senior operational leaders

## The First Enterprise Test of AI Autonomy May Be Prior Authorization

Prior authorization has become one of the most crowded battlegrounds in healthcare AI.

Every platform promises faster submissions, fewer denials, and dramatic labor relief. Many now claim some version of autonomy. Charts are searched automatically. Forms are filled automatically. Portals are updated automatically.

But once the vendors wrap their pitch, a harder conversation begins.

### What happened

Earlier this month, AI Catalyst hosted a candid conversation featuring AI prior authorization companies Latent and Notable, followed by a closed-door debrief among health system revenue and operations leaders.

### Latent: search the chart, answer the payer, move *faster*

Latent framed its core innovation as an intelligence layer inside the EHR. The platform retrieves the correct prior auth pathway, searches the patient record across the notes, labs, and medications, and returns the clinical evidence needed to justify approval.

The topline claims centered on productivity and accuracy:

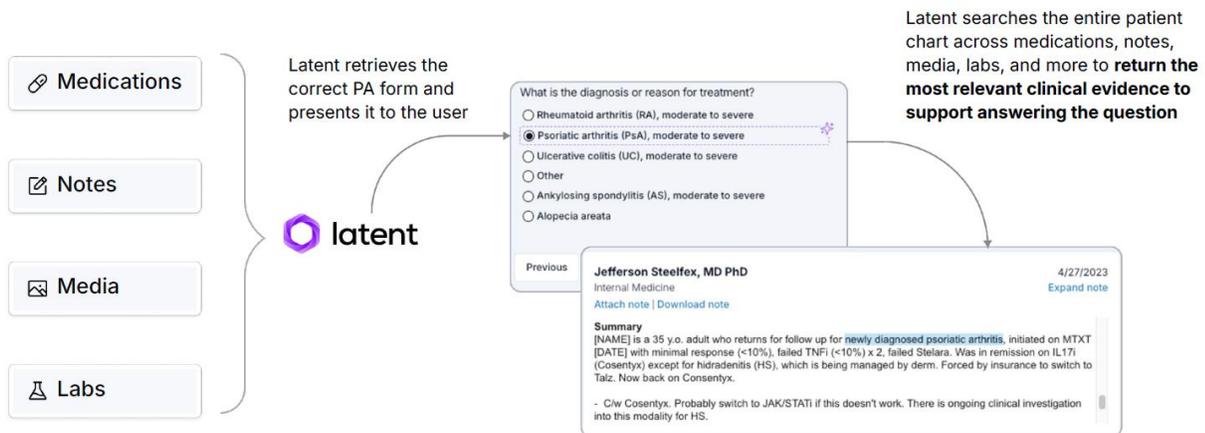
- Specialty pharmacies moving from ~1,000-1,400 prior authorizations per FTE per year to roughly 2,000+
- Approval rate increases across high-volume and specialty drugs, often in the +7 to +16 point range

- Case studies like Mount Sinai showing review times dropping from 15 minutes to under 4, and appeals work falling from 45 minutes to under 8

The underlying message reflected that if you already have the information required to win approval, then the problem (that Latent aims to solve) is finding and packaging it fast enough.

## How Latent Works

Latent integrates with EHRs to streamline prior authorizations by leveraging AI to extract relevant clinical data and guiding users through questionnaires.



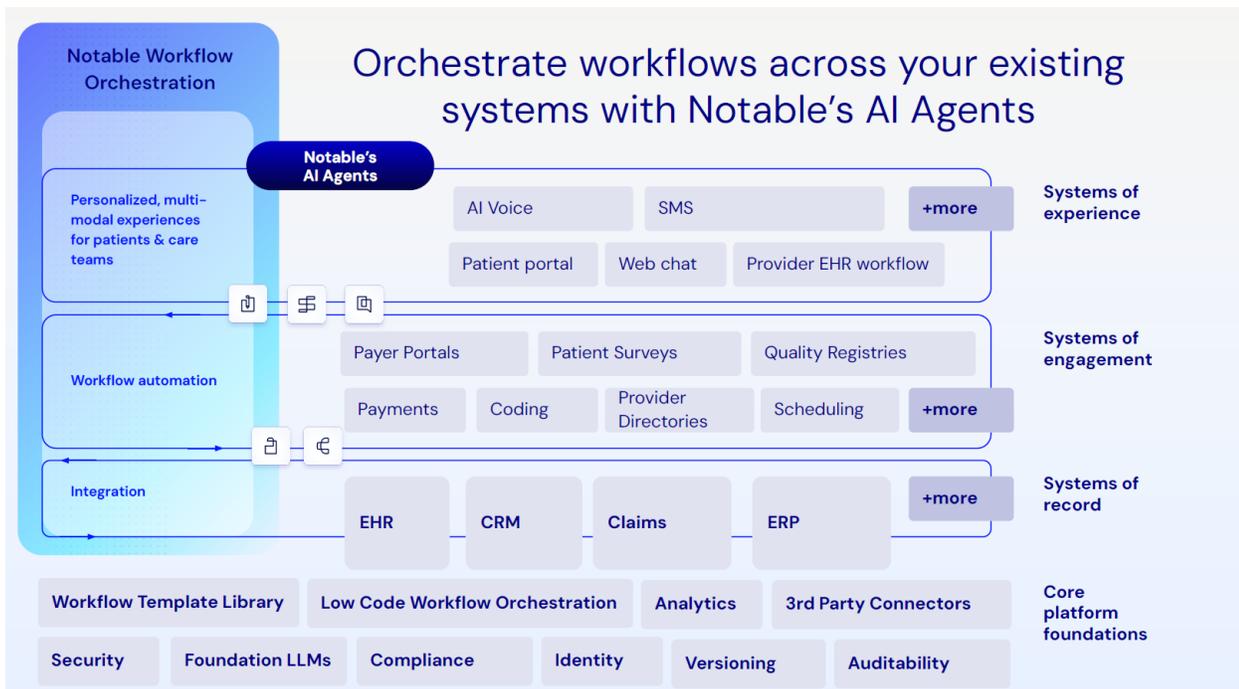
### Notable: orchestrate the workflow end-to-end

Notable tool a broader view. Rather than focusing only on extracting evidence, they emphasized AI agents that monitor work queues, log into payer portals, submit requests, and write results back into the EHR. The overarching frameworks supports capacity creation at scale.

Early implementation in imaging cited:

- Roughly 2 FTEs repurposed in a single department
- Median handling times around 7-10 minutes
- A path toward touchless authorizations for routine volume, with humans reserved for edge cases

Additional partner metrics pushed the scale narrative further: thousands of hours saved, weeks-to-go-live measured in single digits, and meaningful reductions in write-offs.



## When the vendors leave, the truth comes out

With no industry members in the room, the tone changed quickly. It wasn't simply a matter of health system leaders dismissing the claims, but rather because they started asking what has to be true operationally and financially *first* for those claims to actually matter.

A few key realizations surfaced from the conversations:

### 1. We still don't mean the same thing by "automation."

Some organization in the room use RPA, other use NLP assistance, still more are piloting agents that transact in portals., and just a few are contemplating auto-submission without human review. All of these get labeled "AI."

One executive put it plainly, "Until we standardize definitions, we will keep buying tools whose risk profile doesn't match the governance model we put around them."

### 2. Labor savings are the easy math.

Vendors naturally lead with the FTE impact which executives immediately translated into tougher questions like:

- Are positions actually eliminated, or does volume refill the capacity?
- If staff are redeployed, what metric proves value?
- Does finance agree that avoided hiring counts?
- What happens if payer behavior changes?

### 3. Denial risk outweighs speed gains.

Faster submissions are good, but faster *bad* submissions are worse. Leaders focused on whether AI truly understands medical necessity nuance, how often humans must intervene, evolving policies, and physicians who document creatively.

The buying question becomes: How quickly do we detect drift, and who is accountable when it happens?

#### **4. Autonomy shifts who owns prior auth.**

If an AI solution is doing the recommending, then a human owns the decision. If an AI solution is submitting, then ownership blurs.

Is it pharmacy?

Revenue cycle?

Clinical leadership?

IT?

Compliance?

No one in the room had a simple answer and that uncertainty is slowing down decisions more than the technology itself.

#### **Why it matters**

Prior authorization is emerging as one of the first enterprise tests of whether health system can let AI transact, not just recommend. Success here could create a template for how autonomy spread into scheduling, coding, documentation, and beyond. Failure here will harden resistance everywhere else. Before expanding, the group converged on a surprisingly consistent checklist.

They want:

- Baseline denial and turnaround metrics before activation
- Clear thresholds for when automations pauses
- Transparent routing of exceptions
- Audit trails that can survive payer scrutiny
- Finance agreement on what constitutes realized ROI

In other words, AI for prior auth is not a staffing project, but a control framework project. Interest is high, but especially in areas where three conditions already exist:

1. High volume
2. Standardized criteria
3. Limited downside if the model is wrong

Imaging, renewals, and well-defined drug categories surfaced repeatedly as perfect starting points. In contrast, few were ready to hand complex specialty cases directly to autonomous agents.

#### **Questions to consider:**

1. Where is prior auth currently restricting cash or access the most?
2. What percentage of volume could safely move to touchless tomorrow?
3. Who has authority to halt automation if denial rates tick up?
4. What proof will compliance require six months after go-live?
5. If staff time is saved, how will finance see it?

## Soft ROI Is No Longer Enough: Your Third AI Hard Truth of 2026

*Third in a four-part series examining critical challenges healthcare executives must navigate in 2026.*

### What happened

For the first wave of healthcare AI, soft returns were often sufficient, whether it be reduced burnout, better experience, time back in the day, learning value, or strategic signaling. Those benefits were real, and in a period of relative financial flexibility, they were often enough to justify upfront investment in pilots – but that period is over. Welcome to 2026, where the repercussions of federal healthcare cuts, margin compression, rising labor costs, and capital scrutiny have fundamentally reset expectations around AI investment.

Executive teams now face a higher bar to clear for new AI adoption and if AI cannot translate into credible, defensible financial impact, it will not scale.

Importantly, most organizations *are* seeing value. What they are struggling to do is *convert* that value into hard dollars.

When we ask CMIOs where ROI efforts break down, two answers consistently dominate:

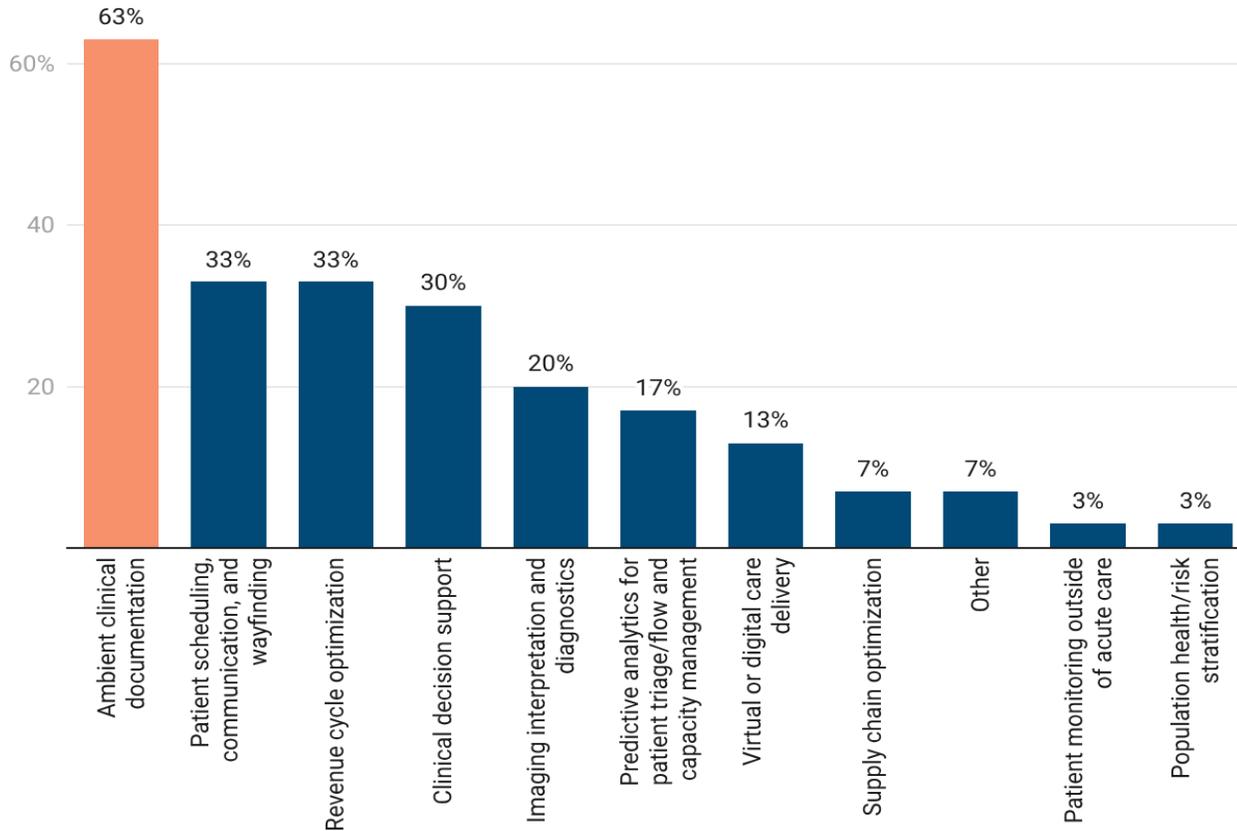
- *We cannot translate operational or experiential gains into financial outcomes.*
- *We never aligned stakeholders on what success meant in the first place.*

One leader's win is another leader's rounding error.

Meanwhile, external data is reinforcing the pressure. Outside of revenue cycle, only a small fraction of health systems reports direct labor cost reduction or net new revenue from AI. Even ambient documentation, one of the [most widely deployed, and high delivering categories](#), rarely reaches the bottom line without deliberate operational change.

# Most measurable outcomes by AI use case category

Percentage of respondents



Created with Datawrapper

What it comes down to is not whether AI saves time; rather, the issue is what the organization chooses to do *with* that time.

## What leading systems are doing differently

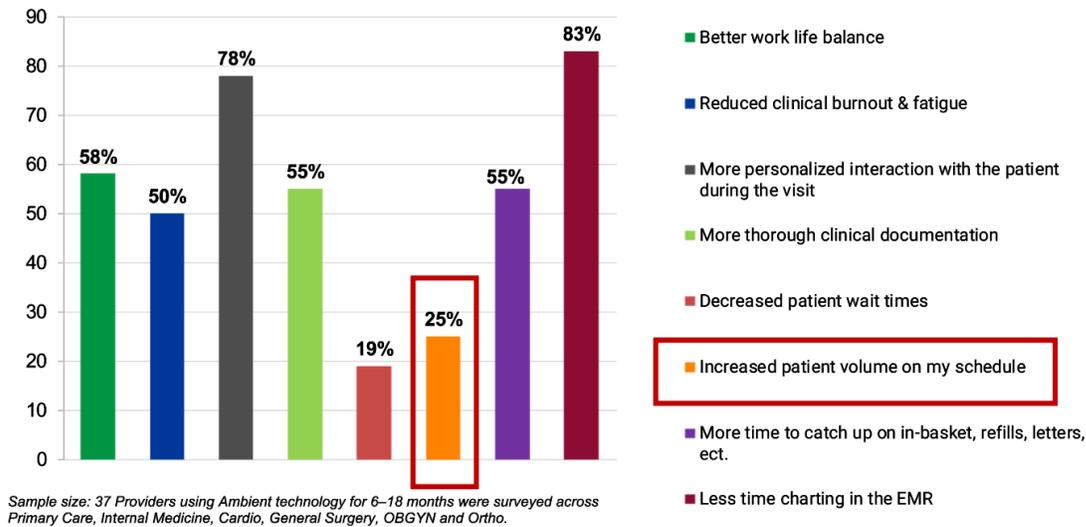
Organizations that are breaking through on ROI are not waiting for finance to “eventually notice” improvement. They are engineering the conversion from soft to hard ROI

## Texas Health Resources: Turning time into wRVUs

THR confronted one of the thorniest ROI challenges in the market with [ambient documentation](#). After deployment, they observed that clinicians reported major reductions in documentation burden, yet only a minority (25%) voluntarily expanded patient volume. Many simply finished earlier.

From a workforce standpoint, that is understandable (even encouraged!). From a CFO standpoint, it is invisible.

## THR survey of Nuance Dax adoption impact on providers



So THR redesigned the pathway from productivity to compensation to system value.

*Collaborate with finance to pre-define success metrics and KPIs*

Before launch, clinical and operational leaders partnered directly with finance to define success metrics. They paired experience measures (time saved, reduced pajama time) with the trusted financial proxy of wRVUs benchmarked against national specialty data.

*Create dedicated incentives for clinicians to drive hard ROI metrics*

Then they aligned incentives. Once physicians crossed the 50th percentile of productivity relative to those benchmarks, the organization absorbed the full cost of the license. Individual clinicians could see their own performance relative to peers, making the opportunity tangible and personal. Over the subsequent months, productivity moved.

In the end, the breakthrough was not the technology, but the economic design around it.

*Original model – February 2024*

Performance against benchmark*	Individual physician share of cost	Health system share of cost
75-100%	0%	100%
50-74%	50%	50%
35-49%	75%	25%
0-34%	100%	0%

\* Benchmark based on national tri-blend median for individual provider's specialty tied to wRVU's

*Latest model – September 2024*

Performance against benchmark*	Individual physician share of cost	Health system share of cost
75-100%	0%	100%
50-74%	0%	100%
35-49%	50%	50%
0-34%	50%	50%

## Cedars-Sinai: Start with the financial endpoint

Similarly, Cedars-Sinai has increasingly made ROI translation a governance requirement, not a post-implementation aspiration. Sponsors are expected to articulate a primary metric *and* the operational mechanism that will convert it into financial performance before deployment begins.

In one example involving [ambient nursing documentation support](#), leaders tied implementation timing directly to existing overtime metrics. They examined whether other drivers of incidental overtime—patient mix, assignments, procedure volumes—shifted during the same period.

They had not.

What had changed was the introduction of the AI tool.

By anchoring the pilot to pre-existing financial indicators and controlling for alternatives, the organization built a credible attribution narrative that finance leadership could trust.



## Why it matters

Both examples reveal the same truth.

Soft ROI does not automatically mature into hard ROI. It requires structural intervention.

Leading organizations are:

- Choosing financial proxies that executives already accept.
- Making end users participants in value creation, not passive recipients of efficiency.
- Designing pilots around measurable inflection points.
- Treating adoption as a financial variable, not just a change-management concern.

Without those moves, AI produces improvement without income. And in 2026, improvement without income is increasingly difficult to defend. This is the year AI must pay rent.

## Questions to consider:

1. Have you named one definition of success for a specific initiative, or are you trying to satisfy ten stakeholders at once?
2. Did finance agree to your definition of success before launch?
3. What specific behavior must change for time saved to become economic value?
4. If adoption improves but workflows remain unchanged, does your ROI still materialize?
5. If you had to defend your ROI math tomorrow, would it withstand scrutiny?

## When the Governance Committee Is No Longer Enough: The AI Governance Leader Is Coming

Health systems are deploying AI faster than they're deciding who's accountable for it. A LinkedIn scan of governance roles across U.S. health systems shows new titles appearing - Chief AI Officer, Director of AI Governance, Responsible AI Lead - but the authority attached to those roles varies wildly. More importantly, many systems are operating AI in production without resolving the fundamental question: who owns this when something goes wrong?

### 2025 has forced a reckoning on AI governance for health systems in 2026

Three forces converged in 2025 that make governance focus an imperative for AI adoption among health systems.

[State regulation exploded](#): 47 states introduced over 250 AI bills in 2025, with 33 enacted. California's AI Transparency Act and Texas's TRAIGA (Responsible Artificial Intelligence Governance Act) went live January 1, 2026.

[Federal enforcement shifted](#): The Office of the National Coordinator for Health Information Technology (ONC) HTI-1 Final Rule (Health Data, Technology, and Interoperability) moved algorithm transparency to enforcement; FDA shifted to continuous surveillance.

[GenAI hit scale](#): AI tools for clinical documentation are moving "from isolated pilot programs to full enterprise-scale deployment.

### But why isn't an AI governance committee enough?

Increasingly, committee-based governance is breaking down and it's showing in two ways. First, pace. AI adoption is accelerating: new models go live, vendors push updates, GenAI tools embed into workflows continuously, not quarterly. When a clinical team wants to deploy an ambient scribe next week and governance doesn't meet until month-end, the tool either launches unreviewed or stalls. Either way, governance fails.

Second, monitoring. Watching for model drift, performance degradation, and bias signals requires dedicated, continuous oversight. Monthly committee meetings can't catch when an imaging algorithm performs differently across patient populations or when a GenAI scribe hallucinates clinical details. [THMA's AI Maturity Matrix](#) shows organizations with committee structures that have governance on paper, but that can't keep up.

### AI governance investments beyond the committee: We're seeing a fork in the road

Some organizations are outsourcing governance. They're relying on vendors like Epic or specialized startups like [Onboard AI](#) to handle their monitoring and oversight, but it's a bet that external partners will catch problems before they happen.

The other path is building in-house governance roles dedicated full-time to AI oversight. These aren't committee members meeting quarterly, but rather full-time operators watching systems round-the-clock. We're focusing on that second path: what it looks like when health systems build dedicated AI governance capacity. The models look different, but the question is the same. Can your governance structure catch up to what you've already deployed?

### **Attributes of in-house governance roles**

Based on our LinkedIn scan of 15 profiles and job postings across health systems and regulated industries, we identified a new category of roles emerging: the dedicated AI governance operator. Although titles are different, these aren't committee members or part-time oversight; they're full-time jobs with distinct titles, skillsets, and responsibilities. Here's what we are seeing.

### **Examples of titles appearing:**

- [VP of AI Governance](#)
- [Chief Data and AI Officer](#)
- [Responsible AI Lead](#)
- [Responsible AI and Clinical Governance](#)
- [Head of AI Governance, Risk and Compliance](#)
- [Senior Principal, AI Governance](#)
- [Director of AI Governance](#)

### **The work converges around four areas:**

1. *Own* enterprise AI governance –design intake, run approvals, enforce registration.
2. *Set expectations* for documentation, validation, monitoring.
3. *Translate* technical AI issues into business language and facilitate hard conversations.
4. *Build* repeatable processes that scale, not one-off reviews.

**The people doing it** came from risk, compliance, privacy, quality, safety, or analytics–regulated environments where accountability already mattered. Strong judgment about tradeoffs, comfort with ambiguity, bias toward building systems. They know when to slow things down.

**Where they sit:** Rarely inside AI or engineering teams. Most report into Data/Analytics, Risk/Compliance, or Legal/Privacy, operating in a matrix with dotted lines to clinical and product leadership.

**The case for dedicated roles:** The argument for setting up these positions centers on continuous oversight. Dedicated operators provide monitoring instead of quarterly reviews, enforce the "no go-live without registration" rule because someone owns the registry, and build repeatable processes that scale with portfolio growth. When incidents occur, there's a single point of accountability. They translate between technical, clinical, and regulatory languages full-time, which has long been a coordination function that doesn't fit into anyone else's job description.

The tradeoffs are real. These roles add headcount and budget when margins are tight. They risk becoming governance theater if decision rights aren't clearly defined. Positioned as gatekeepers rather than enablers, they can slow innovation. And the skillset blend –clinical fluency plus risk judgment plus technical governance capability–is difficult to hire for, especially in competitive markets

**What the 'governance operator' owns:** Synthesized from [NIST](#), [AMA guidance](#), [CHAI's Applied Model Card](#), the [Joint Commission](#), and THMA [research](#), these five areas require dedicated ownership.

1. Maintain the AI registry of products and enforce the 'no go-live without registration' rule.
2. Coordinate validation reviews and ensure local testing happens before deployment.
3. Monitor systems continuously for drift, bias, and performance degradation.
4. Manage escalation when incidents occur –know who to call, when.
5. Own vendor due diligence and contract oversight for purchased AI.

**The choice ahead:** AI governance is a full-time job. Committees can't govern technology that moves faster than they meet. The organizations getting this right didn't just pick a better org chart –they funded the governance operator layer. That's the 'Director of AI Governance' who maintains the registry, coordinates validation, monitors continuously, and manages escalation when things go wrong. Without that role, you have governance on paper. With it, you have someone accountable 24/7.

The question isn't whether to formalize governance. Regulators and boards already decided that. The question is whether you're building external dependence or internal capability. If you're choosing in-house, the commitment is clear: stand up the operator role, give them decision rights, and fund the function to scale with your portfolio. Your AI is growing faster than your governance calendar.

#### Questions to consider:

1. If an AI tool in your system failed tomorrow, who would detect it, contain it, and report to the board and how long would that take?
2. Does your organization have someone watching AI systems full-time, or is governance something that happens in meetings?

## Policy Updates

*Understanding the evolving AI regulatory and legislative landscape*

## What CMMI's New ACCESS Payment Rates Signal for AI in Chronic Care

### What happened

CMMI just published payment rates for its new chronic care model, and the math only works if you can deliver care with radically lower labor costs (like through the use of AI).

The Advancing Chronic Care with Effective, Scalable Solutions (ACCESS) model pays participants fixed installments for managing various chronic conditions like high blood pressure, diabetes, musculoskeletal pain, and depression. The agency's new [request for applications](#) announced its key design elements:

- Payments are split into a first-year rate and a lower maintenance rate.
- CMS pays half in monthly installments with the remainder tied to annual performance metrics and paid out at the end of the year.
- Providers receive full bonus payments only if at least 50% of aligned beneficiaries meet performance targets. Below that threshold, payments are reduced proportionally.

Many are [underwhelmed](#) by the proposed rates, saying the program is not a financially viable alternative for organizations dependent on FFS payments.

- Some clinical tracks could pay as little as \$7 per-member-per-month for behavioral or MSK care.
- [On LinkedIn](#), health tech attorney Carrie Nixon calculated that a chronic condition RPM provider would be trading a potential \$216/month FFS payment for a \$35/month rate under ACCESS.

## Why it matters

The low rates have fueled speculation that the program will mostly be embraced by AI-focused startups with truly transformative models that replace human clinicians with AI-guided self-care. For a traditional provider (even a fully virtual one), the program offers a great deal of hassle and financial risk for very little reward, especially compared to the much higher rates offered for the same care by commercial payers.

This raises two possibilities in the long run:

1. Either the program will flop with low participation and limited impact, or
2. It could catalyze a new wave of disruptor startups that revolutionize chronic care.

The first is certainly more likely, but if the latter comes to pass, we wouldn't be surprised if this voluntary model is eventually used as the basis for mandatory changes to FFS Medicare (a possibility discussed in THMA's [deep dive into CMS innovation models](#) from last October).

The ACCESS model structurally favors AI-native models. Startups operating in employer or commercial markets like Sword Health and Hinge Health for MSK care and Wysa for behavioral health are already optimized to leverage AI to drive cost savings. While they may not be compatible with ACCESS's requirements from the get-go, these disruptors are might be able to adapt their services to the program more easily than a traditional provider, aiming for:

- Asynchronous engagement
- Automated triage
- AI-guided self-care
- Lower reliance on licensed clinicians per encounter

The ACCESS model exemplifies many of the characteristic policy impulses of the Trump administration's healthcare agenda, including a strong emphasis on cost containment, a preference for private-sector solutions, and a heavy focus on technology and AI.

By setting low payment rates that might work for a new AI startup but probably aren't viable for a traditional provider with an existing cost base, the administration is explicitly aiming to disrupt the market and bring new players into the fold.

### Questions to consider

1. At \$35 PMPM, can your current chronic care model break even without materially reducing licensed clinical labor?
2. If CMS reimbursement for chronic care continues compressing, do you have an AI-enabled model ready, or are you defending a cost base built for a different era?
3. If AI-native entrants can deliver measurable outcomes at these rates, where are you competitively advantaged? And where are you structurally exposed?
4. What level of automation in chronic care are you prepared to defend clinically and legally if payment pressure forces substitution?

Other policy updates:

- [Major U.S. health systems and patient-safety leaders launch AI Care Standard™ with 10-pillar evaluation framework for patient-facing AI communications](#)

## AI Strategy Quick Hits

*Noteworthy moves from peers to implement AI technologies*

- [FMOL Health, McLeod Health, and Rush University System for Health report time savings and monthly revenue gain with Suki's AI scribe](#)
- [Nebraska Methodist Health System recovers \\$2M and frees 24 FTE-equivalents with Akasa AI coding and claims automation](#)
- [MultiCare Health System selects Ambience Healthcare for enterprise ambient documentation deployment after 550-clinician pilot shows 33% documentation time reduction](#)
- [Pine Park Health deploys Retell AI voice agents for scheduling, freeing 2+ FTEs and booking visits with 55.7% of connected callers](#)
- [CommonSpirit Health partners with Midstream Health to deploy AI agents targeting missed rebates, underpayments, and denials in financial operations](#)
- [AdventHealth Celebration improves OR on-time starts from 62% to 78% with AI-powered computer vision cameras](#)

## Emerging Use Cases

*New capabilities that indicate AI's potential*

- [Sutter Health embeds OpenEvidence AI medical search engine into Epic EHR for real-time point-of-care access to clinical guidelines](#)
- [McLaren Health Care launches AI cardiac screening program using Carebricks to identify high-risk patients from existing CT scans](#)
- [Southern California Permanente Medical Group cuts high-acuity patient message response time from 22 hours to 5 with AI inbox prioritization](#)
- [University of Michigan Health achieves 97.5% accuracy across 50+ radiologic diagnoses with Prima vision-language brain MRI model](#)

## Cautionary Tales

*The risky side of AI implementation*

- [Mount Sinai study finds nine leading LLMs propagate false medical claims across 1M+ prompts, calling for safeguards before clinical deployment](#)
- [Memorial Sloan Kettering warns health systems of mounting multi-petabyte storage and interoperability challenges from digital imaging growth](#)
- [Randomized study of 1,300 participants finds LLMs do not improve patient health decisions compared to standard online searches](#)

## Market Moves

*A round-up of AI company announcements and stories*

- [Optum launches provider- and payer-facing AI prior authorization tools, with Allina Health as early health system deployer](#)
- [Anterior raises \\$40M to expand clinician-led AI platform covering ~50M health plan lives](#)
- [Talkiatry raises \\$210M Series D to scale AI-powered virtual psychiatry platform](#)

## The News in Numbers

*An interesting data point that caught our eye*

**85%**

of healthcare technology leaders plan to increase agentic AI investment over the next 2–3 years, with 61% already building or implementing initiatives, per a recently [Deloitte study](#).

**52%**

of participating health systems in a recent [HIMSS/Guidehouse study](#) feel that they are operationally ready to deploy AI at scale, despite 78% of them engaging in AI projects.