

# AI that works for health data



A pragmatic framework for health plans, employers, and benefits leaders who need AI to deliver real impact to access, outcomes, and cost – not just create headlines.

## Executive summary

Most health plans and employer benefits leaders are actively exploring AI across every function – but the gap between the hype of generative AI and measurable impact on access, outcomes, and cost is vast. Token consumption grew 4x in the first quarter of 2026 and the largest AI firms have raised over \$100 billion in combined capital, yet a 2025 MIT survey found 95% of organizations are not seeing measurable return from the technology!

At the 2026 Truven Summit, Jason Williams introduced “AI that works” – a five-part management framework grounded in the reality that AI is not the hero of the healthcare story. It is a sidekick, and a powerful one – but only when it knows its job, tells the truth, follows expert practices, is monitored, and is judged by business outcomes.

This paper explains that framework aligned with feedback from ~100 health plan, employer, and consultant attendees at the Truven Summit, and previews Insight Edge – Truven’s analytic study generator launching late in 2026 as a working example of the framework in practice.

### **What you’ll take away**

- The five principles of “AI that works” and questions to ask your teams and vendors against each one.
- Benchmarks of where health care stakeholders sit on AI maturity, strategic drivers, and inhibitors.
- A first look at Insight Edge, Truven’s expert-encoded, human-in-the-loop AI for proactive savings discovery.



# AI is not the hero of the healthcare story

Imagine being handed the power to fly without an instruction manual. That is the situation many leaders find themselves in today – suddenly equipped with extraordinary AI capabilities and an expectation to save healthcare. Most are familiar with the foundational issues: training data quality, hallucination, bias, explainability, etc. Fewer are equipped to manage the less-discussed, non-technical decisions that determine whether AI creates value or erodes trust.


Value matters precisely because AI is going to get more expensive as adoption grows. While compute power is constrained today, the spending boom in data center construction will unleash greater capacity. And while AI providers

are not yet profitable on the \$100B+ invested, dependency on these tools will only deepen. The right question is not “How do we implement AI?” but “How do we advance our mission more effectively – solving problems that matter, more efficiently, and accurately?”

## Where your peers are

Before we walk through the “AI that works” framework, anchor yourself against the room. The chart below summarizes how ~100 health plan, employer benefits, and consultant attendees at the 2026 Truven Summit described their organization’s current state of AI maturity.

## Self-reported AI maturity

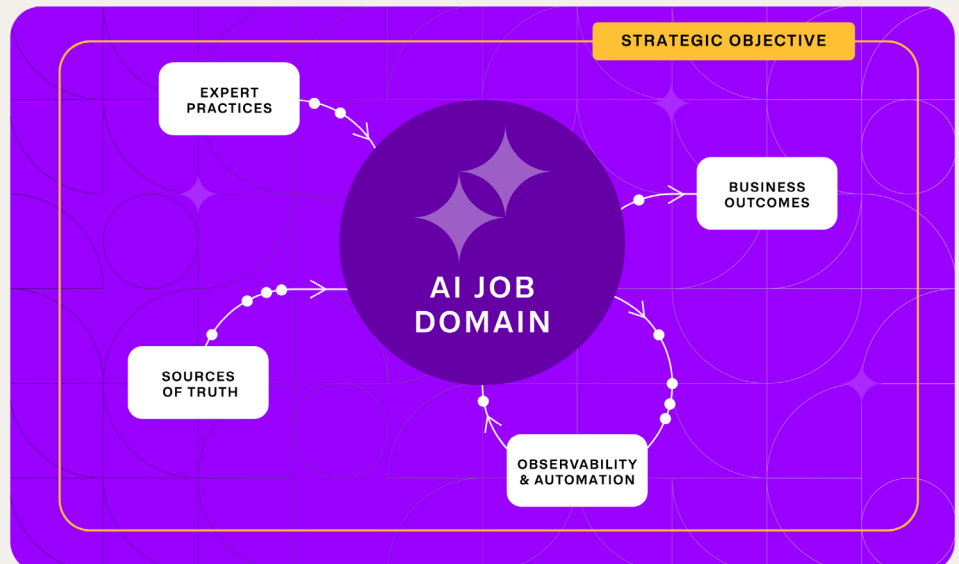


72% of respondents are still in Exploring or Experimenting mode. Just 4% report Scaling or Optimizing. The signal is clear: pilots are plentiful, but operational integration and measured benefit remain the exception. That is the gap “AI that works” is designed to close.

## The “AI that works” framework

AI that works knows its job, delivers consistent truth, follows expert practice, is monitored, and delivers quantifiable value. The framework displayed here is how Truven’s product, engineering, and consulting teams build – and how we believe customers should evaluate any AI investment, internal or third-party.

**Let’s delve deeper into the meaning and implications of each of the five principles...**




# 1. Intent boundaries & guardrails

Keeping AI in the proper lane

## The pattern

Out-of-the-box, foundational LLMs can do far more than you've asked them to. That latent capability can also be a liability.

 A car-dealership assistance chatbot was convinced to sell a \$76K truck for one dollar.

In healthcare, a benefits chatbot that drifts from explaining coverage into clinical or utilization-management advice is not a funny anecdote — it's a regulatory exposure.

## The Truven approach

Design for the work. An AI assistant must know where its job ends. It is imperative to classify user intent against an allow-list of topics, define safe-completion patterns (refusal, clarifying prompts, or escalation to a human), and potentially use of architectural fencing — hard technical barriers between data the AI is never meant to touch — as the most foolproof defense against prompt injection. Then test fanatically.

## Questions to ask your teams and vendors

- What is the explicit allow-list of topics this AI is permitted to address?
- What does it do when a user prompt falls outside that list — refuse, clarify, or escalate?
- What architectural controls (not just instructions) prevent access to off-scope data?
- How are adversarial prompts and jailbreaks tested before each release?

---

### TRUVEN PREVIEW


Our Health Insights natural-language query capability is designed to refuse questions it isn't built to answer — protecting users from confident but useless output.


# 2. Hybrid GenAI with deterministic analytics

The truth, the whole truth, and nothing but the truth

## The pattern

GenAI was designed to be helpful — but not necessarily correct. ChatGPT itself ships with an “AI can be wrong” disclaimer. The well-polished synthesis from GenAI can be misleading.

 A New York City small business support chatbot advised users to do things that may violate local law.

 A member-support assistant once gave weight-loss advice to a person with an eating disorder.

Sounding right is not the same as being right.

## The Truven approach

Bind generative output to deterministic sources of truth. Machine-learning models, predictive models, actuarial logic, benefit accumulators, authorization criteria, and rules engines are consistent. Deploy them as the tools your AI is required to call. Allow GenAI to do the work it is extremely good at — synthesis, summarization, explanation, agent orchestration — while never allowing inference to masquerade as fact in healthcare decision-making.

## Questions to ask your teams and vendors

- Which deterministic systems must this AI be required to call before generating an answer?
- How do you prevent inference from being presented as authoritative fact?
- Will the same prompt produce the same answer on regulated outputs?
- Who owns the analytic or rules engine the AI depends on — and how often is it audited?

---

### TRUVEN PREVIEW

Our MarketScan benchmark profiler (code name Snapshot) — releasing this year — uses a natural-language user interface for cohort data profiling that algorithmically deploys the deep referential knowledge base of HEOR experts.

### 3. Embedded SME-in-the-loop workflows

This is how we do it

#### The pattern

“AI slop” has given rise to a new term — “workslop” — low-value, generic artifacts that waste everyone’s time. A recent LinkedIn debate criticized LLMs for producing “polished but unusable” market research. In healthcare, the stakes are higher: a large national plan now faces litigation alleging AI-influenced, post-acute coverage denials that overrode physician judgment without expert peer review. CMS has warned that AI cannot replace individualized coverage determinations.

#### The Truven approach

AI does not replace expertise. It scales it. [Treat AI like a trainee — not an expert.](#) Effective solutions don’t rely on AI to discover analytic approaches over time; they encode expert ways of working from day one. Consultants, actuaries, and clinicians should be advisory team members, deployed to author output templates, analysis playbooks, review standards, and to monitor, audit, and evolve the system.

#### Questions to ask your teams and vendors

- Whose specific expertise is encoded into this AI’s playbooks and templates?
- Is there a human review and override process before output reaches a decision maker?
- How does expert feedback re-train or re-scope the AI over time?
- What is the AI explicitly not allowed to do without human sign-off?

---

#### TRUVEN PREVIEW

Insight Edge, our autonomous analytic study generator, is embedding a library of expert “recipes” — beginning with a GLP-1 study — inside Health Insights. Think of it as a junior consultant constantly surveilling your data to answer questions you may not have yet asked.

### 4. Observability & agent automation

Supporting actors needed

#### The pattern



An AI banking agent quietly sent promotional rates to 3,000 customers; no one noticed until \$2M of loss had accrued.

Gartner predicted more than 40% of agentic AI projects will be cancelled by end of 2027, largely due to insufficient risk controls and observability. In healthcare, a major EMR vendor’s sepsis prediction model was found by JAMA to miss two-thirds of cases and over-alert on false positives — and the vendor’s rebuttal that performance “depends on institutional factors” underscored exactly why measurement matters.

#### The Truven approach

If you can’t automatically monitor a process today, your AI won’t either. Instrument every AI workflow with signals that demonstrate it is operating as expected — for example, call deflection, appeal reduction, cohort cost trajectory, drift on key metrics. Future-ready systems emit telemetry for surveillance, track drift, and enable timely correction. Best practice now includes agents that check each other’s work.

#### Questions to ask your teams and vendors

- What signals will tell you this AI is drifting, misfiring, or replicating junk at scale?
- Who is alerted when a threshold is crossed — and what is their authority to halt the system?
- Are agent-to-agent checks in place where the cost of error is high?
- Can you reconstruct, after the fact, exactly what an agent decided and why?

---

#### TRUVEN PREVIEW

Observability is a customer partnering priority. Cost-savings cohort analytic opportunities like Truven Risk of Rising Cost are pursuing customer instrumentation for longitudinal surveillance — for example to tag individuals who were case-managed — to continuously calibrate the model.

## 5. Business imperatives & success metrics

AI is not the strategy

### The pattern

A large rideshare company recently disclosed that its full 2026 AI development token budget was consumed in the first four months of the year. A CTO summarized it bluntly: “Their app is not \$70M better.” Too many organizations sound more like “we have AI now” than “we’re solving something you care about — with the help of AI.” The hype-cycle pattern repeats itself: heavy investment without clear payoff thesis can lead to disappointing business results.

### The Truven approach

Treat every AI investment like any other capital decision: with an explicit value thesis and tracked outcomes. The most important question is not how to implement AI but how to advance the mission more effectively — faster, more accurately, and on the problems that matter most.

### Questions to ask your teams and vendors

- What is the value thesis for this AI investment — and how is it measured?
- What business outcome (cost, productivity, quality, member experience, ROI) does it move?
- What is the decision criterion to scale, scrap, or enhance?

---

#### TRUVEN PREVIEW

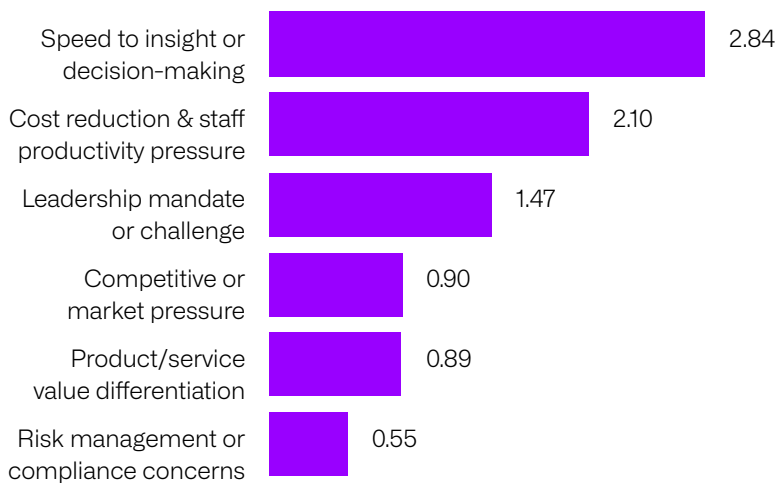
In early 2026, Truven deployed AI-powered data-quality summaries that accelerated data through our best-practice quality checks 50% faster, while enabling redeployment of a significant number of team members to other high-value priorities.



# Validation from Truven clients


The five themes above are not a hypothesis. They were stress-tested live against ~100 health plan, employer, and consultant leaders at the 2026 Truven Summit. Feedback from these key stakeholders below directly reinforce the framework's emphasis on sources-of-truth, expert practice, and governance.

## Strategic Drivers



Speed to insight or decision-making ranked #1, edging out cost reduction and productivity pressure at #2. Translation: leaders are not chasing AI for novelty – they are buying time. That is exactly the value proposition of theme 2 (deterministic truth at GenAI speed) and theme 3 (expert-encoded analytic recipes). Regarding leadership mandates ranked #3, most of these leaders expressed feeling pressure – generally for cost savings – that they are not yet equipped to deliver against. The lowest-ranked driver – risk management – is a cautionary finding worth pairing with the inhibitor chart below.

## Inhibitors to progress



Governance, risk, and compliance concerns ranked #1 as the dominant inhibitor – even though risk management was ranked last among drivers. That asymmetry is the single most important finding of the day: organizations are not investing in AI because of risk, but they are being slowed down by it. The framework's themes 1 (boundaries), 2 (truth), and 4 (observability) are designed precisely to convert that headwind into a tailwind.

These rankings are index-weighted based on overall customer responses. I.e. how many times #1, #2, etc. selected per response option would be a higher score.

Workforce skills and adoption change management ranked #2, validating theme 3 (SME-in-the-loop) – AI investments that arrive without encoded expertise and a credible human-in-the-loop story can stall on adoption, regardless of model quality.

# Insight Edge

An expert-encoded, human-in-the-loop AI for proactive solutioning. Coming in 2026.

Insight Edge is AI that works. It augments generative AI with curated data and deep domain expertise, embedded directly into Health Insights. Instead of asking you to prompt, Insight Edge automatically surfaces the savings opportunities, operational interventions, and risk signals that experts would have surfaced – if they had the time to look at every population, every week.

## Why it's different

	Generic GenAI assistants	Truven Insight Edge
Who guides the analysis	End user prompts	Truven SMEs encode the recipes. <b>AI scales our expertise – it does not replace it.</b>
Source of truth	LLM inference	Pre-aggregated, calculated metrics. <b>AI never sees patient-level data.</b>
Output cadence	On request	Proactive, scheduled, delivered to your inbox. <b>Saves your team time.</b>
Human review	Optional	Built-in human review. <b>Truven analysts pre-release reviews.</b>
Risk of hallucination	User-managed	SME-directed AI. <b>Limits potential for inference or calculated synthesis.</b>

## Sample insights from the GLP-1 recipe

Insight	Illustrative finding
Persistence	Spend on members who stopped using GLP-1 before 3 months
Inappropriate use	Spend on members on an obesity-related GLP-1 without an obesity diagnosis
Prescribing rates	Top 5 providers with obesity GLP-1 prescribing rates above the national average
Lowest-cost option	Potential savings if members move to a lower-cost GLP-1; repricing opportunity vs. MarketScan average

## Be first to see Insight Edge in your data

Truven is partnering with a limited number of organizations for the 2026 launch. Contact your Truven account team to schedule a live walk-through of the GLP-1 recipe – and to provide feedback on Insight Edge study recipe priorities for your population.

# You are the heroes

AI succeeds when it knows its job, tells the truth, follows experts, is measurable, and is judged by business outcomes. The story of healthcare transformation isn't written by models – it's written by the leaders, clinicians, actuaries, and analysts who decide which problems are worth solving and what "good" looks like.

We've seen what happens when humans abdicate vigilance. We've also seen what happens when they don't – when a nurse caught a drug order assigned to the wrong premature twin baby before it would have been administered. The system wasn't the hero that day. The person was.

It is not the capability of AI that will transform healthcare. It is our informed determination to wield it responsibly. Truven is committed to building AI that works – and to partnering with you through user panels, design partnerships, and roadmap feedback to elevate you as the heroes of health transformation.

*"It's not the suit. It's the person wearing it"*

## GET INVOLVED

- Join a Truven user panel on the 2026 AI roadmap.
- Become an Insight Edge launch partner for the GLP-1 study recipe and Rx package.
- Schedule an "AI that works" framework session with your account team.

## ABOUT THIS PAPER

Sourced from the keynote "AI is not the hero of the healthcare story" delivered by Jason Williams, Head of Product, Truven at the 2026 Truven Customer Summit; the Health Insights Roundtable preview of Insight Edge; and live audience polling of ~100 health plan, employer benefits, and consultant attendees.

## About Truven

Truven by Merative is a portfolio of healthcare data and analytics solutions, backed by deep healthcare expertise. We help organizations understand, enhance, and maximize healthcare data to drive better health and financial outcomes. Through best-in-class data quality and market-leading solutions like Health Insights, MarketScan, and Springbuk, Truven serves 7 of the top U.S. health plans, more than 40% of the Fortune 500, and the top global pharmaceutical companies.

Learn more at [merative.com/truven](https://merative.com/truven)

## References

1. The GenAI Divide: State of AI in Business 2025. MIT NANDA, July 2025  
[https://mlq.ai/media/quarterly\\_decks/v01\\_State\\_of\\_AI\\_in\\_Business\\_2025\\_Report.pdf](https://mlq.ai/media/quarterly_decks/v01_State_of_AI_in_Business_2025_Report.pdf)

© Merative US L.P. 2026. All Rights Reserved.

Produced in the United States of America  
June 2026

Merative and the Merative logo are trademarks of Merative US L.P. Other product and service names might be trademarks of Merative or other companies.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on Merative's product plans and strategy as of the date of this publication, which are subject to change by Merative without notice. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from Merative, or stating or implying that any activities undertaken by you will result in any specific performance results. Merative products are warranted according to the terms and conditions of the agreements under which they are provided.

Product release dates, availability and/or capabilities referenced in this publication may change at any time at Merative's sole discretion, and are not intended to be a commitment to future product or feature availability in any way. Not all Merative products are available in all jurisdictions in which Merative operates.

MHI-000000065 Rev 1.0



Truven