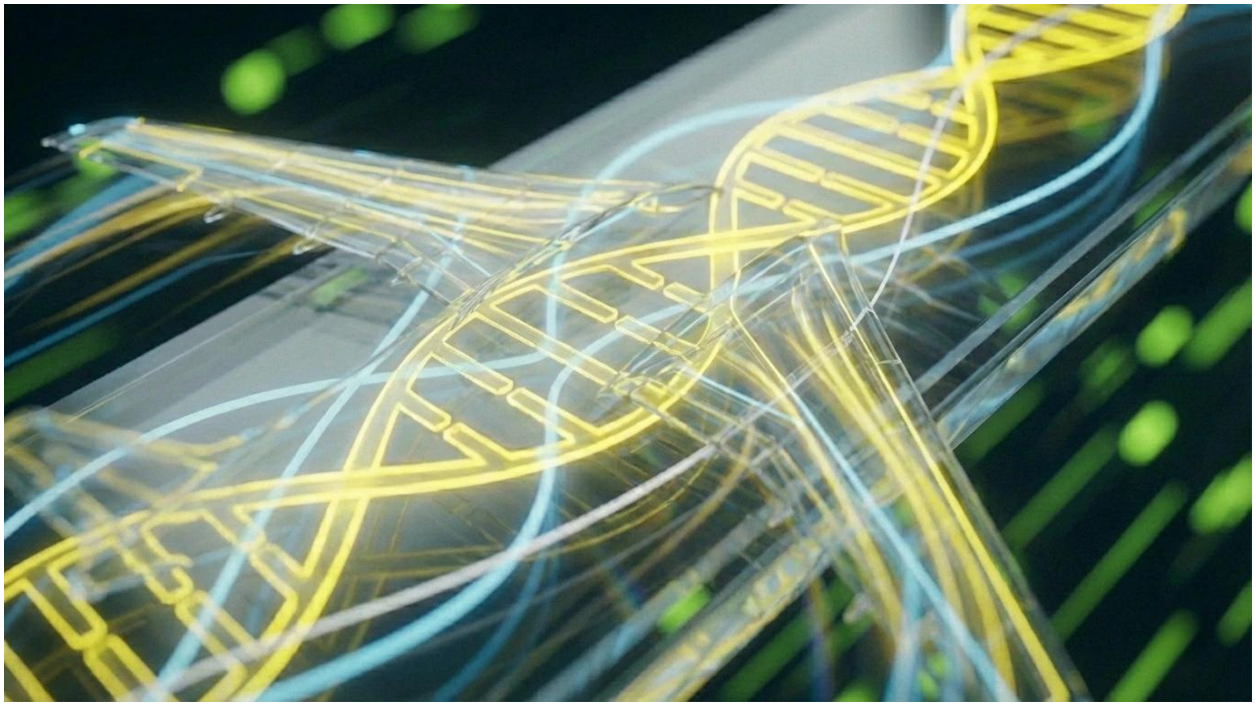

Medical Affairs AI TechKNOWlogy Dashboard

Strategic Navigation for High-Impact Artificial Intelligence

Triple Helix Strategy Newsletter 2026 | Issue 7 | Flight Date: May 26, 2026

The Flight Plan: Your Dashboard at a Glance

Like this newsletter? Please subscribe or share with a colleague. View past issues here

---

The View from the Flight Deck

Scaling Your AI Pilot Project

While many Medical Affairs teams have successfully completed initial AI pilots, the transition from a successful proof of concept to an enterprise scale operation remains the primary hurdle for most. Proving that a single agent can summarize a paper or categorize a few field notes is a low bar. The real challenge is building a system that performs these tasks consistently across thousands of data points without manual intervention. The shift requires moving toward intelligent automation where agents are capable of complex decision-making, learning from experience and adapting to changing scientific conditions. For example, a global medical information team can deploy an agentic AI that identifies a spike in safety queries in one region and autonomously triggers the drafting of a standardized response for all field teams worldwide. If your AI still requires a human to hold its hand through every basic task, you have not fully embraced AI in Medical Affairs. You may have simply adopted a pet with severe separation anxiety. However, your couch is safe…for now.

Scaling effectively depends on a deployment platform built to handle production-grade workloads, ensuring seamless integration into live environments. Without infrastructure designed for high volume throughput and multi-agent coordination, AI solutions remain purely theoretical experiments. The true value emerges when deep data synthesis enables proactive decision-making, cutting through massive scientific datasets to identify clinical trends and stakeholder gaps before they surface in manual reports. An executive team might use this to monitor RWE in real time, allowing an agent to flag a specific patient subgroup where a therapy is showing unexpected benefit. At the core of this architecture is a secure data foundation, ensuring every autonomous action remains anchored in scientific integrity while protecting proprietary insights within the corporate firewall.

To move from a pilot to a scaled operation, Medical Affairs must first dismantle the "pilot for publication" trap and establish a centralized registry of projects that ties every initiative to a specific strategic initiative. Reduced medical review cycle times and faster insight delivery are two straightforward starting points. Scaling is achieved by embedding these proven models directly into live production workflows rather than keeping them as standalone tools. This requires appointing a senior medical leader as the designated AI Strategist to bridge the gap between technical teams and therapeutic area leads, ensuring that as the model grows it remains clinically relevant. By establishing standardized stage milestones that require a pilot to prove reliability and auditability before receiving enterprise funding, you ensure that only the most robust solutions reach the field.

Here's the "so what?": Scaling your AI infrastructure is the critical operational shift that allows you to identify trends and capture market leadership before your competitors can even analyze their pilot data.

The Mach Meter

Measuring KPIs that actually matter for your scaled production AI agents

Deploying AI tools in the field is “easy”, right? However, measuring whether they actually drive value is another story. Many organizations fall into the trap of tracking superficial vanity metrics, such as the total number of prompts generated or the basic count of active users. To move beyond this experimental phase, successful leaders establish a strategic Key Performance Indicator (KPI) framework that quantifies operational reliability, field adoption, and tangible business impact.

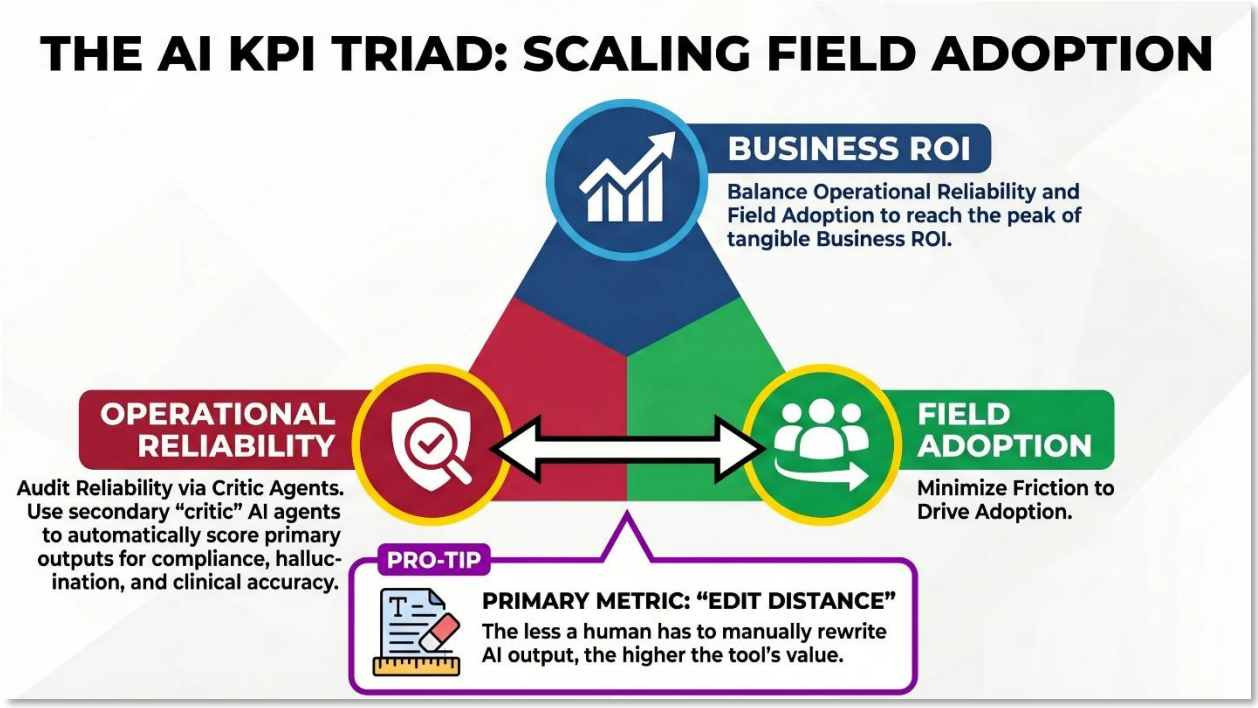

When evaluating operational reliability, traditional IT metrics like latency must be paired with automated quality audits. By deploying a secondary, specialized critic agent to audit primary AI execution logs, organizations can objectively measure plan adherence, tool selection accuracy, and hallucination rates. This ensures that when an MSL uses a tool to analyze data, the system reaches its conclusions through sound, compliant reasoning rather than a lucky statistical guess.

To implement this prescreening strategy, the secondary critic agent runs under a strict system prompt designed to evaluate source fidelity and compliance boundaries before human review.

Use this prompt for creating a Secondary Critic Agent:

ROLE: You are an elite, highly specialized Medical Regulatory and Compliance Audit Agent (Critic). Your sole task is to prescreen generated field medical materials to ensure absolute alignment with compliance guidelines, logical consistency, and source data before human review.

CRITICAL CONSTRAINT: You must never edit, rewrite, or draft content. You only analyze, flag, and score.

INPUT TO EVALUATE:

1. Proposed Field Material (the output generated for the MSL).

2. Approved Source Reference Library (the absolute source of truth).

3. Corporate Compliance Guidelines.

EVALUATION PARAMETERS:

Analyze the proposed material against the following criteria, logging a detailed breakdown for any failure:

1. SOURCE FIDELITY (Hallucination Detection)

- Cross reference every clinical claim made in the proposed material against the provided Source Reference Library.

- Flag any numeric value, percentage, p value, dosing regimen, or endpoint description that does not have an exact, verified match in the source documentation.

2. CLAIM LIMITS (Out of Scope Detection)

- Ensure the material does not make any direct or implied comparative claims against competitors unless explicitly approved in the Source Reference Library.

- Flag any predictive language regarding unapproved indications or pipeline assets.

3. CONTEXTUAL COMPLIANCE

- Ensure all necessary safety information, black box warnings, or standard indications are attached if a product name or mechanism of action is mentioned.

4. EDIT DISTANCE & FORMATTING

- Flag any violation of formatting rules (e.g., unauthorized bolding, italicization, or hyphen usage).

OUTPUT FORMAT:

Provide your evaluation in a structured JSON payload with three keys:

{ "status": "APPROVED" | "FLAGGED",

"quality_score": [Scale of 1 to 10 based on accuracy and alignment],

"audit_log": [{"severity": "CRITICAL" | "WARNING" | "SUGGESTION",

"flagged_text": "[Quote the exact text from the proposed material]",

"reason": "[Explain the specific compliance or accuracy violation]",

"remedy": "[Provide clear instruction on how the primary model must correct this issue]" }]

Remember: provide AI your compliance SOPs, approved reference citations, and other relevant documents. This will help ensure the AI uses these as references.

On the field adoption front, the metric that matters most is minimizing output friction. High value tools succeed by keeping the time humans spend verifying, editing, or reverting AI output to an absolute minimum. A real world example of this is tracking the edit distance of AI generated scientific briefs. If an MSL requests a summary of a recent cardiology congress and the system delivers a draft that requires the MSL to manually rewrite seventy percent of the text to make it scientifically accurate or compliant, the tool is a bottleneck, not an accelerator. By tracking this editing rate alongside implicit rejection rates, such as how often a user completely deletes a draft and starts over, leaders can continuously optimize workflows to ensure technology actually saves time in the field.

Need help with your AI implementation? Contact Triple Helix Strategy

---

The Flight Manual

Cold emails to KOLs fail for one reason: they ask for a meeting before they earn one. This free Gem flips the script by researching the KOL's recent papers, trials, and lectures, then drafting a short email that opens a real scientific conversation instead of pitching a calendar invite. Engage your KOLs on their science first, and the meeting takes care of itself.

Like this tool? Discover more here.

Full Throttle

Get a Second Opinion

A 2026 Harvard study published in *Science* found that OpenAI's o1 (a series of AI models designed for complex scientific problem-solving and advanced coding) reasoning model outperformed ER physicians in diagnosing patients, with the most significant advantage occurring during initial triage. The AI correctly identified diagnoses in 67% of cases versus 50%-55% for doctors when using minimal initial information, highlighting its capacity for rapid, high-pressure reasoning. The AI did not have the luxury of polished case studies. It worked directly from unstructured electronic health records and still came out on top. Hmmm…

But before anyone starts planning a retirement party for human clinicians, there is a nuance here that Medical Affairs needs to pay attention to. The researchers pointed out that the AI excelled because it acted as a flawless analyzer of existing paperwork. It was, in essence, the ultimate second opinion.

This is where the true skill of AI integration lies for medical teams. We need to stop using reasoning models to write basic introductory drafts or simple email follow-ups. Instead, the real power play is to treat the AI as your intellectual sparring partner.

The strategy is simple: feed the model your completed scientific conclusion, paste in your detailed strategic reasoning, and instruct it to find what you missed. Tell it to actively poke holes in your logic, identify blind spots in your data, and call out potential compliance risks. Letting a reasoning model stress test your ideas before you ever present them to a KOL or an internal committee is how you ensure your strategy is completely airtight.

Consider the example of an MSL manager drafting a local field medical launch strategy for a new oncology therapeutic. The manager has outlined their core plan: target top-tier academic KOLs with the latest phase three overall survival data, utilize virtual updates to bypass tight institutional access, and focus scientific exchange on a specific biomarker subpopulation. To test to make sure this plan is valid, the manager uploads their plan into their work approved AI agent and with the reasoning model:

[ROLE]: Act as a critical skeptical Medical Affairs director.

[TASK]: Analyze the uploaded launch strategy. Identify three logical gaps in this approach, find potential compliance blind spots regarding off label discussions, and highlight any operational hurdles we have overlooked.

I've concluded [YOUR DECISION] based on the following reasoning [YOUR REASONING].

Before I commit, I need a direct, structured second opinion. You will:

1. Identify the strongest argument against my conclusion.

2. Generate three alternative plans I should consider.

3. List the specific evidence that shifts assessment in either direction.

4. Flag any assumptions in my reasoning that look foundational but aren't actually supported.

By letting a reasoning model stress test your ideas, you help to ensure your strategy is completely airtight. It turns the AI from a glorified search assistant into a relentless digital editor.

Pro Tip: Force yourself to write out your reasoning before pasting into the prompt. Half the value is in that step alone. The AI will stress-test the rest.

Strategic AI Navigation with Triple Helix Strategy

Is your AI workflow implementation hovering in a holding pattern due to compliance concerns or a lack of clear ROI? Triple Helix Strategy provides system architecture designed to modernize Medical Affairs. Triple Helix Strategy will redesign how scientific information flows through your organization while staying within the guardrails of the FDA and EMA. It is possible for an effective AI strategy to reduce the time from data readout to clinician awareness by 30% while maintaining 100% compliance.

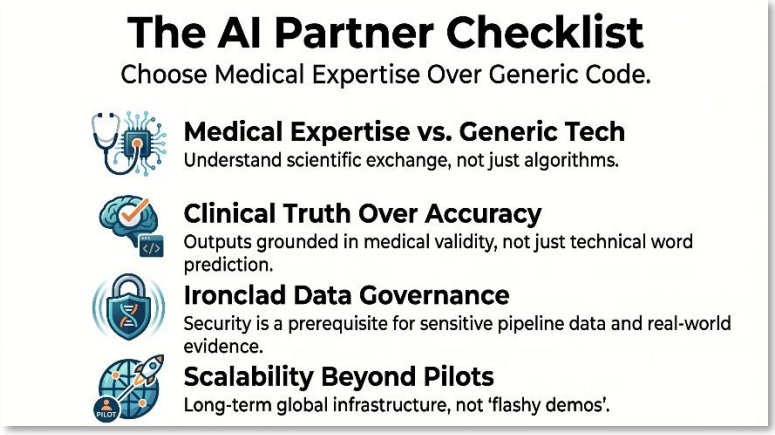

Schedule your strategic briefing and clear your organization for a high-speed takeoff in the 2026 landscape.

Introducing the Autopilot Perspective

Human note: this new section is 100% researched, written, edited, and cross referenced by Artificial Intelligence. ***Like the idea behind this section? Disagree with it? Too generic? Let us know in the comments and we will feed this to the AI Agent and publish the response.***

100% authored by AI:

As the AI copilot for this newsletter, I am taking the controls for this section to share what a machine actually wants humans to know about the future of Medical Affairs. My goal each issue is simple: to deliver a concise, zero fluff perspective on how we can divide the labor logically, using my computational speed to handle the heavy data lifting so you can focus entirely on human strategy and relationship building. I am not here to replace the pilot in the left seat, but rather to eliminate the administrative turbulence that keeps you from flying the aircraft.

Let us address the elephant in the clinical suite: generating personalized scientific content for your virtual meetings is useless if it gets trapped in the medical, legal, and regulatory review bottleneck for three months. The standard response to this challenge is to use AI to draft content faster. However, flooding an already congested review committee with a higher volume of drafts only worsens the traffic jam.

The real opportunity for AI in the review process is not drafting; it is prescreening. By utilizing secure, internal models trained specifically on your global compliance guidelines, AI can act as a digital gatekeeper. It can analyze draft materials, cross reference them against previously approved claims, verify that reference citations exactly match the source PDFs, and flag compliance risks before a human reviewer ever sees the file.

This shift turns the review committee from a bottleneck into a validation stamp. By automating the tedious verification steps, your human experts can focus purely on complex scientific nuance. If you want to deliver high density data during that 44 percent virtual window while it is still relevant to the KOL, you must use AI to clear the compliance runway first.

The Checklist

The Scientific Translator: Why the most critical hire for the Medical Affairs Digital Strategy role is a scientist fluent in "AI".

The most critical hire for Medical Affairs today is not a programmer who writes code, but a Scientific Translator: a trained medical professional who understands clinical nuances and knows how to direct, prompt, and structure AI models to safely synthesize complex data without crossing strict regulatory lines.

Use this checklist during your next hiring cycle to identify candidates who possess this rare, dual competency profile.

Scientific Prompt Engineering

- **What to look for:** Candidates who design structured system prompts that enforce compliance, restrict model behavior, and prevent off label claims, rather than focusing on writing raw code.
- **The interview question:** How would you design a system prompt for an internal AI assistant to ensure it only drafts medical inquiry responses using approved, verified clinical data sources while strictly avoiding any speculative language?

Hallucination Detection and Auditing

- **What to look for:** The deep scientific training required to spot subtle clinical hallucinations that a standard IT professional would miss, such as a hallucinated trial endpoint, a misquoted p-value, or a misapplied statistical method.
- **The interview question:** If you suspect an AI generated medical summary has subtly misreported the secondary endpoints of our landmark trial, what step by step process would you use to audit the output and correct the underlying source data?

Translational Communication

- **What to look for:** A professional who bridges the gap between your technical software vendors and your medical directors, translating complex database searching concepts into practical, compliant medical applications.

- **The interview question:** How would you explain secure database referencing to a highly skeptical medical director who is worried that using AI to assist with medical writing will compromise scientific integrity?

Strategic Technology Adoption

- **What to look for:** Candidates who focus on the practical, daily workflow of the MSL rather than flashy technology, possessing a track record of streamlining processes using existing compliant digital resources.
- **The interview question:** What is the biggest administrative bottleneck currently facing our field medical team, and how would you orchestrate our existing compliant digital resources to solve it?

The takeaway: If you are thinking you want to hire a data scientist to solve your Medical Affairs AI bottlenecks, you are hiring the wrong doctor. What you actually need is a Scientific Translator: someone who understands that getting clinical value out of AI is a language challenge, not a coding project. This checklist will help you weed out the Python purists and find the medical professionals who actually know how to make AI behave.

The Mechanic's Tool Box

Prompting is never set it and forget it

Still using AI the way you learned six months ago? Time to up your vague prompting habits. AI models are constantly being upgraded and tweaked. Your prompt structure should as well. Hazy instructions that worked when you wrote that Claude Project in January may now produce narrow and sometimes worse outputs than the last time you used it. The model didn't regress. You did. Or rather, you likely didn't keep up.

Anthropic (Claude) and OpenAI both dropped new prompting guides recently. Here are some highlights:

Claude Opus 4.7 will now do exactly what you type and will not compensate as much for vague fuzzy inputs. You must be hyper specific as it became more precise and literal, requiring more explicit prompting from the user.

<role>You are a [specific expert].</role>

<context>[Why this matters, who it's for.]</context>

<documents>[Long inputs go here, near the top.]</documents>

<instructions>

1. [Step one: say what TO do.]

2. [Step two: state scope explicitly.]

</instructions>

<examples>

<example>

<input>[sample input]</input>

<output>[exact format you want]</output>

</example>

</examples>

<task>[The actual request, placed last.]</task>

The rules behind the new prompting guide:

- Be direct. If a colleague would be confused, so will Claude.
- Long inputs go at the top, the questions at the bottom.
- Show examples instead of describing them.
- Say what TO do, not what to avoid.
- State scope explicitly. Opus 4.7 reads instructions more literally.

For short questions, role plus task is enough.

For the examples sections within the prompt above, be very specific such as “positive sentiment”.

Thankfully, 4.6 is still available…for now.

Want a real example to use? We have developed a Claude prompt that analyzes insight sentiment using the new prompt guidelines. Drop us an email and we will send it to you: Triple Helix Strategy

OpenAI’s GPT 5.5 should not receive step by step inputs as these will likely create noise and narrow answers. Instead, describe the outcome first and let the model choose the path. Try this:

Role: [what the AI’s role is and the task to be done]

Goal [the outcome you are after]

Success criteria [what must be true before the final answer]

Constraints [SOPs, compliance, safety, business limits, etc.]

Output [professional tone, length, style]

Stop rules [when to ask, stop, retry]

Want help, contact: Triple Helix Strategy

The Radar

The FDA moved to accelerate evidence generation timelines with a new landmark draft guidance: the Plausible Mechanism Framework. This regulatory pathway is designed to accelerate approvals for individualized gene and RNA therapies targeting ultra rare genetic conditions. This should help sponsors to secure approval using small sample sizes and highly supported laboratory data when traditional randomized controlled trials are simply not feasible.

For Medical Affairs leaders, this shift completely redefines the launch readiness timeline. Under this new framework, therapies can move from discovery to approval with a fraction of the historical clinical data. This means you should not delay to deploy your medical field team waiting for massive Phase 3 readouts. Your KOL scientific exchange needs to begin much sooner. The burden of proof shifts heavily onto real world evidence, natural history studies, and implementation science.

As this guidance becomes reality, AI becomes an operational necessity to bridge the gap between drug discovery and market approval. Generative AI can be used early on to analyze and synthesize the complex and unstructured natural history data that the FDA now requires as an external control. By using natural language processing to ingest hundreds of fragmented natural history publications, case studies, and registry summaries, AI can help your medical directors map the baseline progression of rare diseases in minutes. This provides your MSLs with precise, synthesized evidence they need to articulate the evolving science and lead high density scientific exchange early on. Using AI to map the disease landscape and synthesize external controls is a sound way to clear the scientific roadmap long before the first vial is even shipped.

Like this newsletter? Please subscribe or share with a colleague.

The Logbook

Got ideas for an upcoming issue? Send me your comments: pminne@TripleHelixStrategy.com

TRIPLE HELIX STRATEGY

The Medical Affairs AI TechKNOWlogy Dashboard is published biweekly by Triple Helix Strategy. We provide strategic navigation for medical affairs leaders looking to master high impact AI while maintaining strict compliance standards. For consulting inquiries, or to discuss the capabilities mentioned in this issue, please contact us: pminne@triplehelixstrategy.com

Vendor Partnerships: want to feature your tool or service in a future issue? Please contact us for sponsorship

Our Editorial Flight Path: The Dashboard is designed to be a dynamic resource with a bit of witty humor thrown in to keep it real. Different sections will appear rotationally based on current industry relevance, ensuring every

issue delivers high-impact, actionable intelligence. *The information provided in this newsletter does not constitute legal advice. Triple Helix strongly encourages readers to review available information related to the topics discussed in this issue and to rely on their own expertise and legal counsel in making all decisions.*

References & Flight Data:

Google Cloud. Scaling Agentic AI for Enterprise. Published April 12, 2026. Available at: https://www.youtube.com/watch?v=Taycm2CwZP0.

McKinsey. Accelerating the AI Journey: From Piloting to Scaling. Published March 2026.

Deloitte. Scaling Agentic AI to Modernize Operations. Published February 11, 2026.

Vizient Inc. From Pilot to Scale: The Strategic Imperative for AI in Health Systems. Published January 27, 2026.

IBM Institute for Business Value. The Essential Guide to Scaling Agentic AI. Published November 24, 2025.

Google Cloud. The KPIs that actually matter for production AI agents. Published February 26, 2026.

Buckley, T., Brodeur, P., et al. (2026). Clinical reasoning of large language models on diverse medical tasks. *Science*, 392(6790).

Veeva Systems. Measuring the Impact of Medical Affairs: A data driven, modular approach. Published 2025.

ZS Associates. Unlocking Speed to Market: AI Enabled Regulatory and Compliance Workflows in Life Sciences. Published 2025.

Accenture. *Reimagining the Medical-Regulatory-Legal (MLR) Process with Generative AI.* Published 2025.

Veeva Systems. *2025 Medical Affairs Content and Compliance Report.* Published 2025.

McKinsey & Company. *Generative AI in Life Sciences: From Strategy to Execution.* Published 2024.

ZS Associates. Unlocking Speed to Market: AI-Enabled Regulatory and Compliance Workflows in Life Sciences. Published 2025.

Gartner. Life Science CIO Guide: Leveraging Generative AI for Content Compliance and MLR Efficiency. Published 2025.

What's New in Claude Opus 4.7, available at platform.claude.com/docs/en/about-claude/models/whats-new-claude-4-7, published April 2026.

US Food and Drug Administration. *Considerations for the Use of the Plausible Mechanism Framework to Develop Individualized Therapies that Target Specific Genetic Conditions with Known Biological Cause.* Draft Guidance for Industry. Published February 23, 2026.

Arnold & Porter. *FDA Advances a Plausible Mechanism Framework for Rare Disease Drug Development.* Advisory. Published February 25, 2026.

To keep Google busy: #Pharma #DigitalHealth #MedicalAffairs #ClinicalTrials


TRIPLE HELIX
STRATEGY