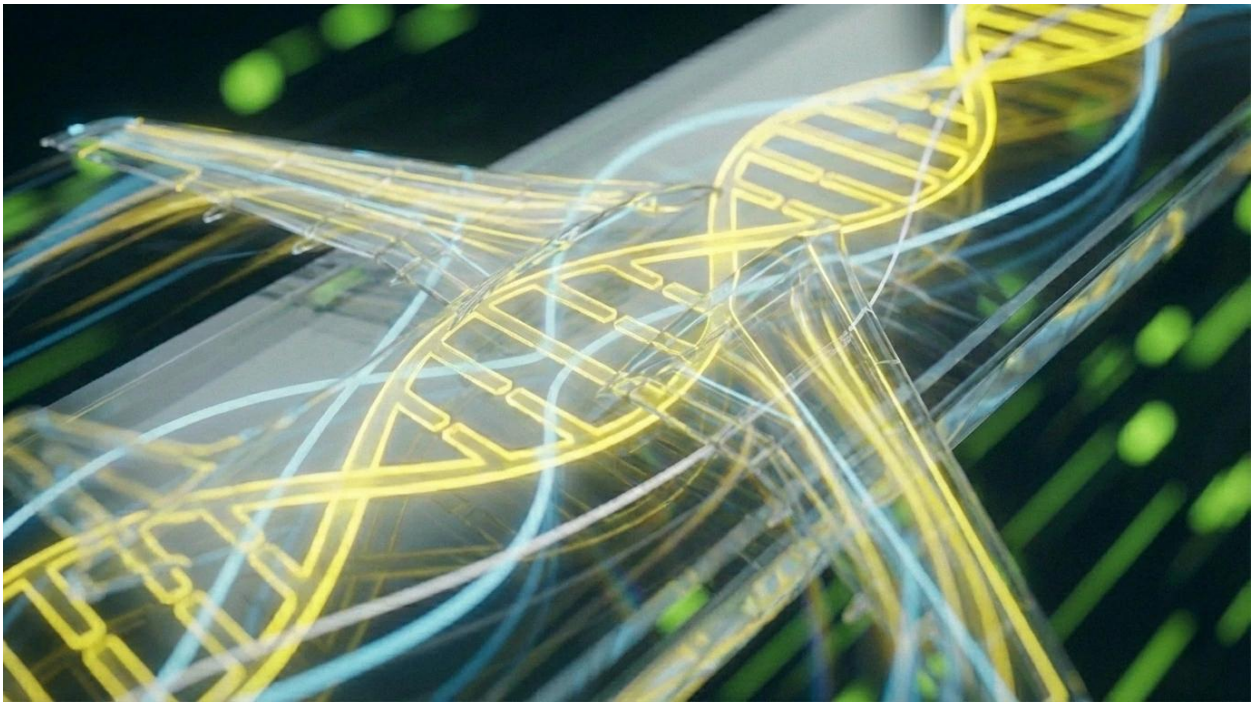

Medical Affairs AI TechKNOWlogy Dashboard

Strategic Navigation for High-Impact Artificial Intelligence

Triple Helix Strategy Newsletter 2026 | Issue 2 | Flight Date: March 16, 2026

The Flight Plan: **Your Dashboard at a Glance**

- **The View from the Flight Deck:** A quick executive briefing on the week's most important shift in AI and what it means for Medical Affairs.
- **The Mach Meter:** Move beyond simple chat prompts into deep practical integration of AI into real-world workflows within firewalled environments.
- **The AI Flight Manual:** Tactical, step-by-step procedures and SOPs for MSLs by MSLs in the field.
- **Full Throttle:** A quantifiable workflow of the week for common tasks with real ROI.
- **The Captain's Perspective:** Hard-hitting perspectives from dynamic leaders on the reality of AI implementation.
- **The Checklist:** Protocols for high impact AI deployment and operational excellence.
- **The Mechanic's Tool Box:** The "Low-Hanging Fruit" of AI capabilities that solve immediate issues.
- **The Radar:** A curated scan of groundbreaking regulatory and technical news that actually matter to BioPharma.

The View from the Flight Deck

FDA TEMPO Pilot and the Life Cycle Mandate

Following the January joint principles, the FDA is moving fast. As of March 2, 2026, the agency has begun sending follow up requests for its TEMPO pilot program. TEMPO (Targeted Engagement for Medical Product Optimization) is the flagship initiative designed to move beyond static snapshots of data toward a continuous exchange between industry and regulators. It essentially turns the post market phase into a live laboratory where the performance of AI enabled tools is monitored in the wild rather than just in controlled settings.

For Medical Affairs leaders, this indicates that the era of one and done validation is over. Regulators are no longer satisfied with a tool that was accurate at launch; they are now requiring proof of ongoing reliability through the use of Pre determined Change Control Plans (PCCPs). This means your strategic focus must shift from initial deployment to active life cycle management. If your team cannot document how it identifies and corrects for Data Drift, the gradual loss of accuracy as scientific literature evolves, you are not just looking at a technical error. You are looking at a fundamental compliance failure. Mastery of this new regulatory landscape requires an infrastructure that treats AI as a living system that must be revalidated as frequently as the science it interprets.

Your AI is like a new hire. It starts off eager and accurate, but if you do not keep it updated on the latest data, it will eventually start giving clinical advice based on a 2024 abstract it found in a digital dumpster.

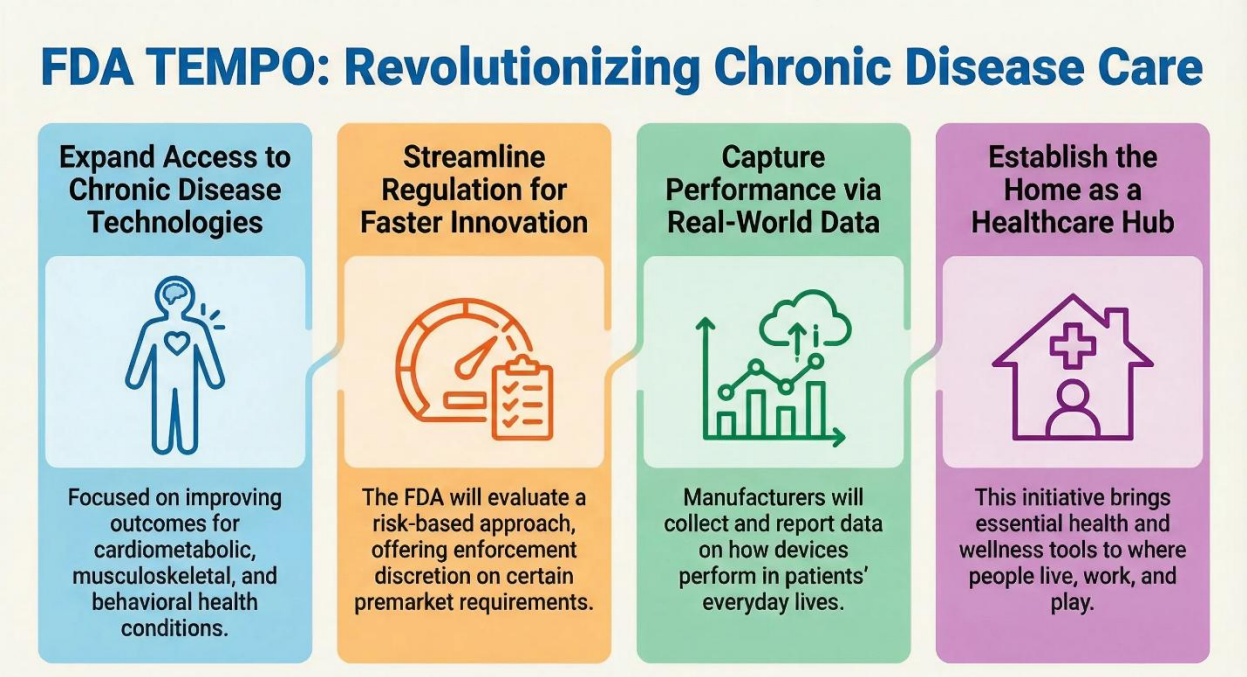

What you need to know about TEMPO: The TEMPO pilot is specifically designed to foster early and frequent interaction between the FDA and product developers. It moves away from the traditional

submit and wait model toward a more collaborative engagement model. Organizations that master PCCPs will be able to iterate their AI tools faster than competitors who are stuck in traditional "submit and wait" cycles.

The Mach Meter

Agentic Intelligence: Moving from Chatbots to Strategic Partners

In 2026, we are leaving Chatbot Fatigue in the rearview mirror. We are now integrating Agentic AI autonomous systems that do not wait for your Monday morning caffeine to kick in; they are already working, interrogating the weekend clinical trial updates while you are still clearing your inbox. Unlike the basic chat interfaces of 2024, these agents are designed to execute complex, multi-step workflows with minimal oversight.

To maintain high performance, your agents must be Verticalized. Verticalized agents aren't just smarter, they are traceable, ensuring that every insight comes with a regulatory-grade paper trail that survives an FDA audit. A general LLM is a generalist with a tendency to wander and a verticalized agent is trained on validated medical datasets and ICH GCP standards (International Council for Harmonisation - Good Clinical Practice). These standards represent the unified, international ethical and scientific quality requirements for the design and conduct of clinical trials. When these agents are deployed within your fire walled environment, they transform from simple search tools into persistent scientific infrastructure. They move the needle from reactive data retrieval to proactive strategic navigation by identifying shifts in clinical sentiment or emerging competitor safety signals before they reach your dashboard.

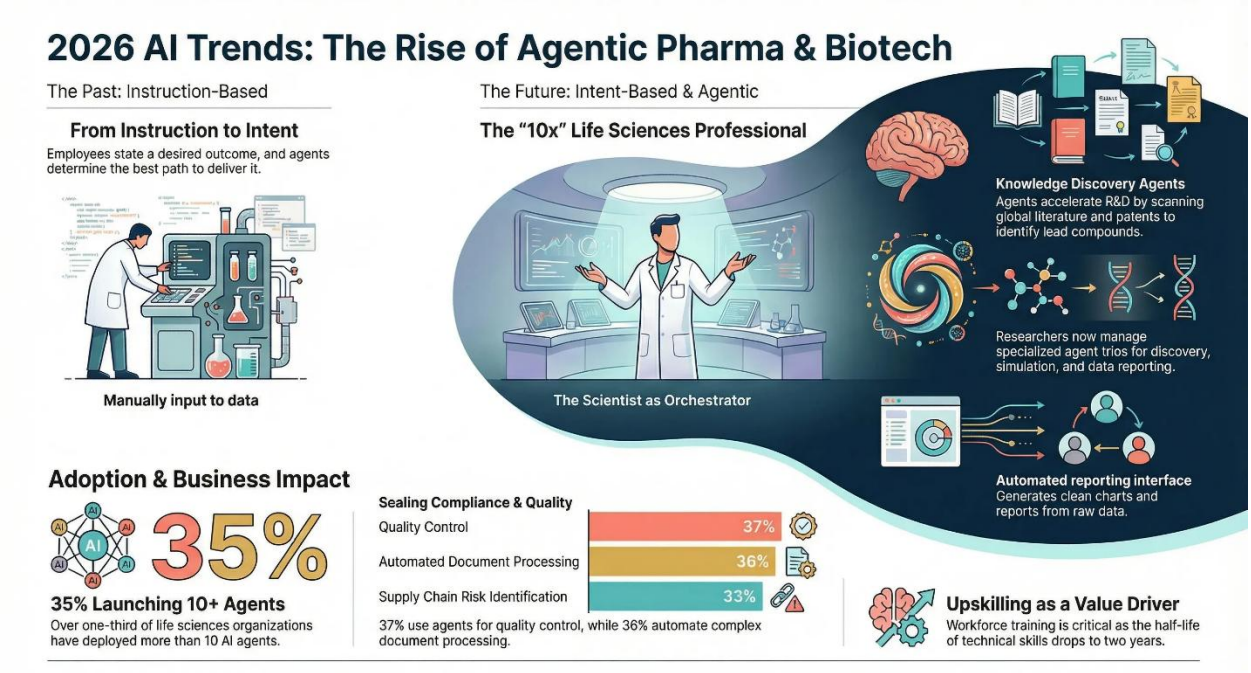

A primary example is the Real World Evidence Navigator. This system architecture is built on the OMOP Common Data Model, an open community standard that allows for the systematic analysis of disparate

health databases. By using this standardized framework, the agent can interrogate vast, de identified datasets to identify patient outcome trends or clinical gaps using natural language. This effectively removes the technical barrier between complex data science and strategic medical insights. When these agents are deployed within your fire walled environment, they transform from simple search tools into persistent scientific infrastructure. They move the needle from reactive data retrieval to proactive strategic navigation by identifying shifts in clinical sentiment or emerging competitor safety signals before they reach your dashboard.

For Medical Affairs, using a general purpose AI for medical strategy is like trying to cross the Atlantic in a glider. Maybe technically possible but many times more challenging than using a Boeing 787.

The AI Flight Manual

Are you an MSL using AI in the field and want to share your best practice? Give us your idea to be featured in an upcoming issue

Matt Schmolesky, PhD, a Senior MSL, has woven AI into his daily workflow in two deliberate and thoughtful ways. One of his go-to approaches is what he calls a *"Time Generator of Resources."* During pre-call planning, or when fielding a request for a literature review, he turns to his organization's firewalled AI large language model to accelerate the search process. Unlike traditional databases such as PubMed, the AI doesn't just surface citations; it delivers concise, plain-language summaries that help him quickly grasp the relevance of each paper. Still, Matt approaches these results with healthy skepticism. He knows the AI can miss key publications or return data that's slightly off, so he never takes the output at face value. Instead, he methodically clicks through every source, verifies the citations, and crosschecks the findings against his original intent. If something feels incomplete, he pushes back by asking the AI why an article was excluded or requesting clarification on the data. He describes this iterative process similar to an ongoing scientific dialogue with a KOL. By continuously probing, questioning, and validating, Matt uses AI not as a shortcut, but as a collaborative partner that helps him dig deeper and sharpen his understanding.

Pro Tip: AI is a great time saver, but always verify and seek to clarify. Never settle for AI's first output.

Full Throttle

The ROI of Reclaimed Time

In the last issue, we showed you how to reclaim 50 hours a week for a team of 10 MSLs. This week we measure the Quality of Flight. In the 2026 landscape, executive leadership is moving toward impact metrics that prove scientific strategy is actually landing in the field.

The Metric: First Contact Resolution.

The Workflow: Using a verticalized agent during pre-call planning to analyze an HCP's recent publications, clinical trial involvement, predict specific scientific tension points, and clinical gaps before the MSL even walks through the door.

The Result: **A whopping 93.75 % reduction in time** to task for deep dive research! Reclaiming this time should allow MSLs to identify scientific gaps earlier, which directly impacts how quickly a product's value proposition is understood by key payers and KOLs.

This efficiency gain is grounded in the 2025 METR benchmarks for Frontier AI, which tracked the transition of complex research from a 16 hour human work block down to a single hour of agentic synthesis. By automating the heavy lifting of data integration, your team moves from being reactive information providers to proactive strategic partners. This workflow ensures that every minute of face time with a KOL is spent on high level scientific exchange rather than administrative catch up.

That is enough reclaimed time to finally stop using your lunch break to catch up on email and well…actually…eat!

Strategic AI Navigation with Triple Helix Strategy

TRIPLE HELIX STRATEGY

Is your AI implementation hovering in a holding pattern due to compliance concerns or a lack of clear ROI? Triple Helix Strategy provides specialized consulting services designed to modernize Medical Affairs departments through the integration of Artificial Intelligence. Move beyond generic technology by focusing on clinical validity, data security, and specialized agentic workflows that deliver measurable returns on investment while maintaining scientific integrity and compliance.

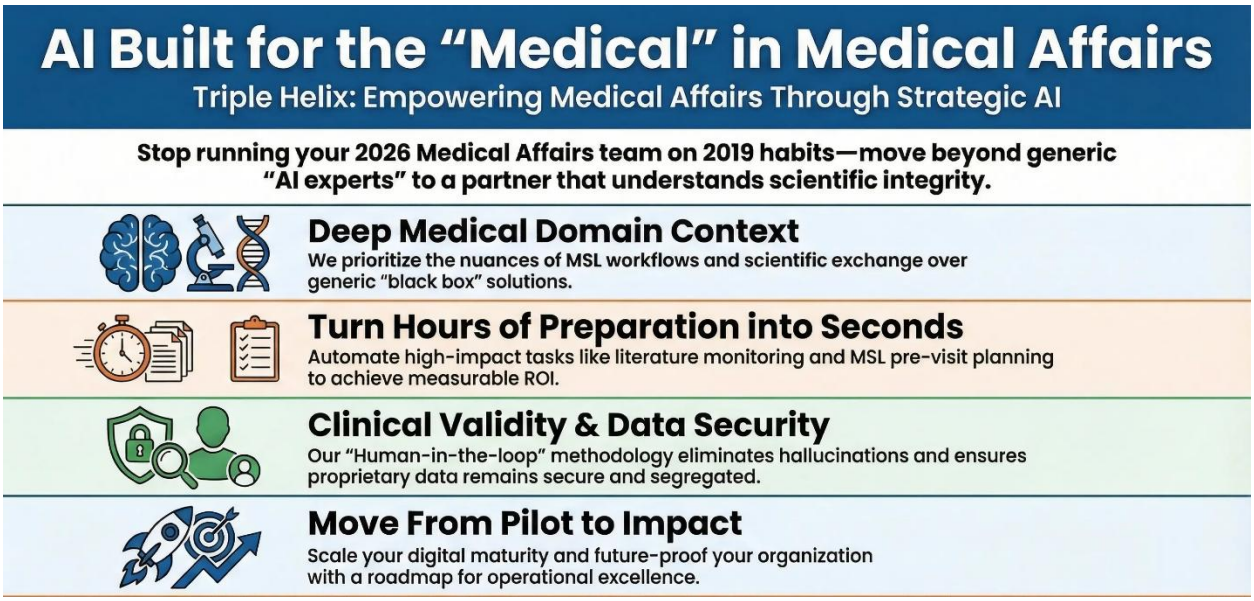

Schedule your strategic briefing and clear your organization for a high-speed takeoff in the 2026 landscape.

The Captain's Perspective

Call for Flight Leads: Are you a medical affairs leader with a thought about AI?

Contact us with how you are using AI to lead.

This week, we sat down with the Head of Global Medical Affairs at Orchard Therapeutics, Kent Christopherson, PhD, to discuss how he is navigating the AI leadership frontier while staying within corporate policy. Overall, Kent views AI not as a replacement for leadership, but as a high functioning assistant "Thought Partner."

In an environment where many BioPharma companies prohibit uploading company strategies or product-specific information, Kent has adopted a disciplined approach. By working exclusively with de-identified inputs and omitting specific drug products in favor of generalized scenarios, structural challenges, and enterprise-level decision frameworks, he has been able to unlock significant efficiency and strategic clarity while staying fully compliant with corporate policy.

Kent uses AI to organize his thoughts, synthesize information, explore scenarios, and sharpen narratives. He is clear, however, that AI does not make decisions. Leadership still requires accountability, contextual awareness, and values-based judgment. He draws inspiration from the book The AI-Driven Leader, where Geoff Woods emphasizes expanding strategic capacity through AI by refining thinking, pressure-testing assumptions, and accelerating insight generation. That philosophy closely aligns with his approach.

For Kent, disciplined AI use is the key. Leaders must protect confidential information, validate outputs, and ensure AI reflects rather than replaces subject matter and leadership expertise. In his view, AI is not simply an efficiency tool; it is a catalyst for becoming a more rigorous and reflective leader.

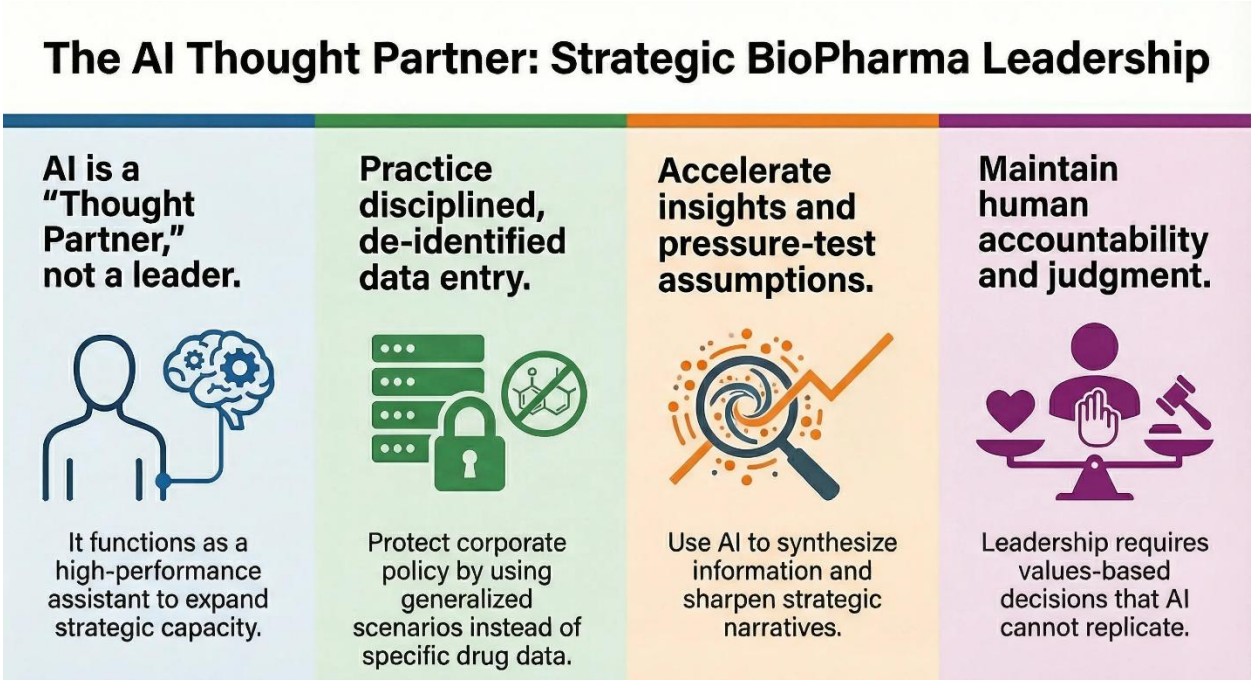

The Checklist

Standardized Procedures for AI Readiness

The Pre-Flight LLM Check

Before entering any prompt into a generative AI tool, run this checklist to ensure your flight remains compliant and secure.

- **Firewall Verification:** Am I using an enterprise approved instance or a public site?
- **Data Deidentification:** Have all patient and HCP names been scrubbed from the input?
- **Source Validation:** Do I have the source documents ready to cross check the AI output for hallucinations?
- **The So What Test:** Does this prompt ask for a summary or does it ask for a strategic implication?
- **Check Your Six:** Did you treat the AI output as a first draft or a final verdict? (Spoiler alert: It's always the former).

Treating an AI prompt like a verbal conversation is a great way to end up off course. You need to treat it like a flight plan: if the coordinates are off by even one degree, you are going to land in a very different zip code than you intended.

The Mechanic's Tool Box

Scientific Article Sentiment Mapping

In 2026, we are finally moving beyond simple keyword tracking to identify the qualitative reality of research impact. While traditional analytics can tell you how many times a paper was cited, they are notoriously bad at telling you if the author was praising your work or using it as a cautionary tale for their students. Scientific article sentiment mapping uses LLMs to decode the why behind the citation. This essentially will allow Medical Affairs teams to identify emerging skepticism or any scientific tension in real time, long before it reaches the sterile environment of an advisory board.

Use sentiment mapping to establish an early warning system for your key assets. By analyzing the nuanced language used in digital medical forums comments, specialized models can now classify mentions on a granular 7-point scale. Using Entity Based Sentiment Analysis, these models can differentiate between a "negative" word directed at your data versus a "negative" word used to defend your data against a competitor. For example, when a KOL uses your paper to prove a competitor "wrong," the AI recognizes your asset as the hero of the sentence, not the target. Peer reviewed studies from 2026 show these models now reach over 90 percent accuracy in identifying these complex scientific intents.

This could be the difference between knowing people are talking about your data and knowing they are currently tearing your methodology apart in a private group chat. One is a vanity metric; the other is a call to action for your MSL team to provide the right evidence before the "noise" becomes the narrative.

Key Takeaway: Technology such as this allows Medical Affairs to move from defensive (responding to criticism) to offensive (leveraging KOL praise to drive scientific narrative).

Want to create a tool that can do this? Contact Triple Helix Strategy

---

The Radar

The Colorado AI Act: Turbulence Ahead (SB 24-205)

The only thing more unpredictable than a spring snowstorm in the Rockies is the local legislative weather. Just when the Colorado AI Act (SB 24-205) was cleared for takeoff with a February 2026 effective date, the state legislature pulled a last-minute go around. Following a contentious special session and months of industry lobbying, the implementation has been officially pushed back to June 30, 2026.

Despite the delay, the fundamental requirements remain unchanged. The act is the first of its kind in the United States to target "algorithmic discrimination" in high-risk systems, which specifically includes AI used in healthcare services. If your AI agents are a substantial factor in making consequential decisions about patient access or treatment, you are officially in the crosshairs.

The Compliance Flight Plan:

- **Developers** must provide detailed documentation and risk disclosures to anyone using their systems.
- **Deployers** (may apply to Medical Affairs organizations) must implement a robust risk management program and conduct annual impact assessments.
- **The Attorney General** has exclusive enforcement authority, with fines reaching up to $20,000 per violation.

While the tech industry successfully lobbied for more time, they did not get a change in course. Consider this your final approach: you have exactly four extra months to ensure your AI governance is airworthy before the June deadline.

The Logbook

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

The Medical Affairs AI TechKNOWlogy Dashboard is published 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 or 303-219-0303

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 a witty humor thrown in to keep it real. While the sections below represent our core operational pillars, they 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:

1. FDA. Digital Health Center of Excellence: TEMPO Pilot Program Updates. March 2026.
2. FDA & EMA. Guiding Principles for the use of Pre-determined Change Control Plans (PCCPs) in AI/ML Enabled Medical Devices. August 2025
3. International Council for Harmonisation. ICH E6(R3) Guideline for Good Clinical Practice. Integrated Addendum to ICH E6(R2). Current version dated 2026.
4. Narrativa. AI Trends in Pharma for 2026: Verticalization and Agentic AI. Published January 15, 2026.

5. OHDSI. The OMOP Common Data Model (CDM). Observational Health Data Sciences and Informatics. Current version: v5.4. Available at: https://www.ohdsi.org/data-standardization/
6. METR. Task-Completion Time Horizons of Frontier AI Models. Published March 19, 2025. Updated February 20, 2026. Available at: https://metr.org/time-horizons/
7. Zeng, H. Evaluating Sentiment Analysis Models in Healthcare: Addressing Bias and Enhancing Interpretability. Frontiers in Public Health. Published November 14, 2025. Available at: https://www.frontiersin.org/journals/public-health/articles/10.3389/fpubh.2025.1663871/full
8. Colorado General Assembly. SB24-205 Consumer Protections for Artificial Intelligence. Enacted May 2024; Effective June 30, 2026. Available at: https://leg.colorado.gov/bills/sb24-205