Orchestrating the agentic revolution across life sciences

Building the foundation for scalable, high-impact agentic AI in Pharma and MedTech

Make it real.

Capgemini

Executive summary

Orchestrating the agentic revolution across life sciences

Building the foundation for scalable, high-impact agentic AI in pharma and MedTech.

Our recent executive survey of biopharma R&D leaders reveals a troubling pattern in many AI strategies: Spending is rising steadily and expectations are growing, but value is stalling[1].

Life sciences leaders can perhaps take comfort in the fact that they are not alone. A 2025 MIT report reveals that despite investing $35 to $40 billion in internal AI initiatives, 95% of organizations see near-zero returns[2]. Among the small minority of programs that generate impact, most can be traced to a singular pain point, as opposed to widespread benefit across the value chain.

In a recent interview with Technology Magazine, Capgemini Group CIO Pascal Brier pointed out that this focus on near-term pain points over long-term vision is the AI decision executives will come to regret the most in five years' time[3]. While it can be tempting to focus on quick wins, he says, sustainable, scalable value is built on deeper foundations: strong data, effective governance, a flexible operating model, and a deep understanding of how work is done.

This is certainly the case for life sciences organizations. Navigating this landscape lies in establishing clear, overarching principles for agentic AI, grounded in industry realities and tailored to the specific demands of pharma and MedTech. Hence the reason for this series: ***Rearchitecting life sciences for agentic AI.***

In this point of view, the first in our series, we examine those principles and how to apply them in practice. In subsequent papers, our domain experts will delve deeper, exploring how functional leaders can leverage foundational investments to not only drive success within their area of the business, but extend impact across the value chain.

As we embark on this series, leaders must remember that higher AI investment won't unlock higher value, but rearchitecting around AI will. Join us as we show you why, how, and where to start.

86% expect AI to fundamentally reshape R&D, but 63% lack the data maturity to realize it.

76% anticipate new ways of working, yet fewer than half feel ready to operationalize them.

71% are defining enterprise AI strategies, though just 8% are actively implementing them.

[1]Capgemini Research Institute, Smart bet, only option, or both?: Biopharma R&D turns to AI, 2026. https://www.capgemini.com/fi-en/insights/research-library/how-ai-will-transform-life-sciences-r-and-d/

[2]MIT, The GenAI Divide: State of AI in Business 2025, 2025. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

[3]Technology Magazine, The technology interview: Pascal Brier, 2026. https://technologymagazine.com/magazine/technology-magazine-april-2026?page=24

The uncomfortable truth: AI's promise-reality gap

Pharma and MedTech companies are facing their toughest operating conditions in decades.

In pharma, R&D costs are rising, pricing is tightening, and portfolios are becoming more complex, just as expiring exclusivities threaten $200 to $350 billion in annual revenue[4]. With cost-per-molecule surpassing $2 billion, the industry's blockbuster era is over[5].

MedTech is facing similar pressures: long development cycles of three to seven years[6], coupled with slowing growth in key markets like China are underscoring the need for both acceleration and efficiency[7].

Across both subsectors, the situation is complicated by an evolving regulatory environment, which not only drives complexity, but also increases the need for auditability, transparency, and validation.

In this landscape, agentic AI has emerged as a bright spot, expected to not only accelerate and streamline processes, but reinvent them altogether. Yet despite rising investment, our research shows a persistent gap between promise and reality. To close it, leaders must challenge three long-held assumptions about how value is created, scaled, and sustained:

From isolated use cases to end-to-end workflows

Value is not created in isolated applications, but across end-to-end workflows.

From algorithms to data advantage

Success will not be defined by the best algorithm, but by the ability to build and scale proprietary data flywheels.

From digital evolution to organizational reinvention

Agentic AI is not just the next phase of digitalization, it requires rethinking organizational structures, talent models, partnerships, and investments.

In this paper, we explore how organizations can embrace these shifts to effectively bridge the gap between AI promise and reality.

[4]DrugPatentWatch, A Framework for Multi-Year Pharmaceutical Patent Cliff Impact Modeling and Strategic Response, 2025. https://www.drugpatentwatch.com/blog/a-framework-for-multi-year-pharmaceutical-patent-cliff-impact-modeling-and-strategic-response/

[5]PubMed Central, R&D Costs of New Medicines: A Landscape Analysis, 2021. https://pmc.ncbi.nlm.nih.gov/articles/PMC8576181/#:~:text=The%20evidence%20also%20suggested%20that%20R&D%20costs,50%2D58%25%20of%20R&D%20costs%20per%20new%20medicine.

[6]The Welding Institute, What is the Medical Device Development Process?. https://www.twi-global.com/technical-knowledge/faqs/what-is-the-medical-device-development-process

[7]MedTechDive, Hospital, lab equipment makers brace for China demand slowdown, 2023. https://www.medtechdive.com/news/hospital-lab-equipment-makers-brace-for-china-demand-slowdown/693639/

The comforting truth: The path forward is clear.

In our analysis of 270 processes across the pharmaceutical and MedTech value chains, our team found that the vast majority of workflows could be automated or augmented by agentic AI.

Agentifiable (~84%)

Regulatory intelligence | Adverse event processing | Cold chain qualification | Biologic manufacturing

Not agentifiable (~16%)

First-in-human trial design | QP sign-off | On-site GxP audits | FDA advisory presentations

Source: Capgemini Life Sciences Process Analysis

Industry research confirms this point, with one study estimating that 75-85% of life sciences value chain workflows contain tasks that AI agents could automate or augment, which could free up 25-40% of employee workloads[8].

Next, we take a closer look at these workflows, exploring four key areas where coordinated, multi-agentic systems, coupled with GxP-ready foundations, can create a concrete impact along the value chain for pharma and MedTech companies.

[8]McKinsey & Co., Reimagining life science enterprises with agentic AI, 2025. https://www.mckinsey.com/industries/life-sciences/our-insights/reimagining-life-science-enterprises-with-agentic-ai

01 Building the R&D to clinical bridge

For many life sciences organizations, data fragmentation across labs, CROs, regulatory files, and clinical sites - as well as across engineering and quality in MedTech - continues to stall AI at the pilot stage. Agentic platforms address this by connecting workflows end to end, enabling AI to support candidate identification, study design, and regulatory submission.

Pharma: Accelerating the molecule

When executed on a unified data foundation, agentic AI accelerates every stage of drug discovery and development.

For example, Insilico Medicine has developed PandaClaw[9], an autonomous AI agentic solution that mirrors the reasoning and expertise of experienced biologists and bioinformaticians. Building on its strong quantitative strengths in target ranking and prioritization, Insilico's AI agents deliver real-time, qualitative multi-omics analysis with deep, contextual interpretation, acting as intelligent scientific collaborators throughout the drug discovery process. Launched in 2026, PandaClaw, together with other AI solutions, has the potential to significantly accelerate the drug discovery process.

MedTech: Redesigning the device development cycle

Agentic AI transforms MedTech development by breaking linear, document-heavy processes and reducing reliance on manual coordination across engineering, clinical, and regulatory teams.

For example, design verification agents can be used to cross-reference new device specifications against historical failure mode libraries and applicable standards in real time, flagging non-conformances earlier. At the same time, clinical evaluation agents can continuously monitor published literature and clinical databases, maintaining the clinical evaluation report as a living document rather than a periodic deliverable.

[9]Insilico Medicine, Insilico Medicine Launches PandaClaw: Empowering Biologists with Agentic AI for Therapeutic Discovery, 2026. https://insilico.com/news/sp-jz8fzmb1-insilico-medicine-launches-pandaclaw-emp

02 Creating medical-commercial synergy

Medical, commercial, and compliance functions often operate in silos, creating fragmented messaging and inconsistent engagement. Agentic AI helps create shared intelligence across functions, aligning insights, coordinating actions, and enabling consistent, data-driven engagement.

Pharma: From mass marketing to intelligence-led engagement

In-field teams, including sales reps and medical science liaisons (MSLs), have traditionally spent hours preparing for physician calls by manually aggregating data from disparate sources. Pre-call briefing agents can now pull from clinical databases, HCP engagement histories and the latest literature to generate personalized, compliance-cleared summaries in seconds, shifting the MSL role toward curating AI-driven insights rather than producing them from scratch.

Across commercial and regulatory functions, agentic systems enable a new engagement model. Conversational agents allow commercial teams to ask natural language questions and receive data-grounded answers in real time. In regulatory affairs, intelligent compliance agents can embed regulatory requirements directly into content creation workflows, enabling compliance officers to shift from reactive box-checking to strategic ethics engineering.

MedTech: Selling outcomes, not equipment

For MedTech companies, the shift from product- to outcome-centric models makes the commercial challenge both more complex and more urgent. Agentic AI provides the enabling infrastructure.

Health economics and outcomes agents can continuously synthesize real-world evidence from device performance, clinical studies, and health system data to update value dossiers in near real time. Clinical specialists can bring agent-generated insights into surgical discussions, shifting the conversation from product claims to patient-specific evidence. Post-sales, predictive maintenance agents can continuously monitor connected device performance, alerting clinical engineering teams to potential failures before they occur.

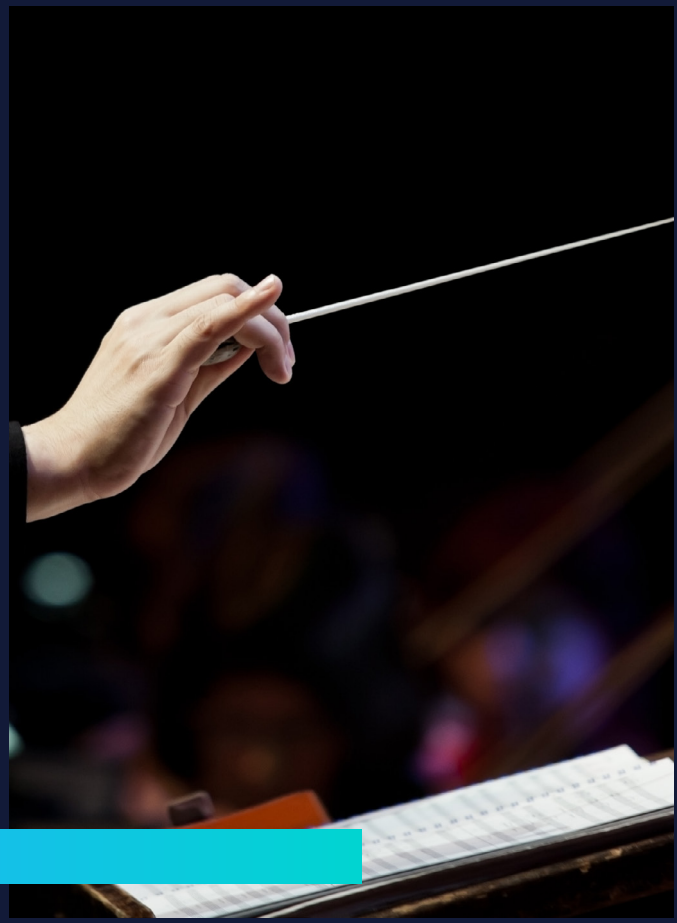

03 Closing the loop with post-market data

For decades, life sciences organizations have sought to use real-world evidence as a continuous, connected input across discovery, development, and strategy. Agentic AI finally makes this possible.

Pharma: Indication finding and lifecycle intelligence

AI-driven indication discovery shows the power of this model at scale. By analyzing millions of patient records, it can rapidly identify new uses for existing drugs, accelerating repurposing and label expansion at far lower costs.

This capability also creates a bidirectional and continuous cycle: real-world data informs trials, AI optimizes design and recruitment, and post-launch signals refine commercial strategy in real time. What was once a multi-year process is now ongoing.

MedTech: From post-market surveillance to continuous improvement

In MedTech, post-market data has traditionally served a compliance role, confirming devices are safe and perform as intended. Agentic AI transforms it into a strategic asset, surfacing insights beyond periodic human review, from subtle performance degradation to real-world outcomes that validate or challenge pre-market evidence.

04 Enabling predictive resilience in manufacturing and supply chain

Manufacturing and supply chain processes remain highly document-centric and reliant on paper-based workflows, even as regulatory expectations for real-time data availability, traceability, and end-to-end transparency increase.

Agentic AI has the potential to reshape operations by reducing cycle times, accelerating capacity ramp-up, and shifting organizations from reactive compliance to continuous, data-driven assurance.

Pharma: Agentic AI in manufacturing and supply

In pharmaceutical manufacturing, agentic AI enables a shift from siloed execution to a connected digital thread spanning chemistry, manufacturing and controls (CMC), quality, manufacturing, and supply chain. Agents, underpinned by a governed, contextualized data foundation, can orchestrate critical workflows such as technology transfer, deviation classification, inspection readiness, and demand-supply synchronization, enabling a resilient manufacturing network that is able to scale faster, operate right-first-time, and maintain regulatory confidence as complexity grows.

MedTech: Modernizing operations and quality processes

In MedTech, where product lifecycles are shorter, change velocity is high, and post-market obligations are especially demanding, agentic AI enables a shift from episodic, human-driven reviews to continuous, system-level assurance across quality management, field performance, and supply operations; shortening response times and reducing operational friction, while strengthening patient safety and regulatory readiness.

For example, one high-impact use case is corrective and preventive actions (CAPA) assessment and resolution wherein agents can continuously ingest complaints, vigilance data, and device performance signals, classify issues, propose root-cause hypotheses, and guide remediation workflows, while maintaining full traceability and compliance.

What it takes to win: 5 interconnected capabilities to enable change

In life sciences, GxP requirements fundamentally elevate the role of human–AI interaction, placing stricter demands on transparency, oversight, and accountability than in typical agentic systems. At the same time, the complexity of building a full semantic layer (across diverse data types and embedded scientific logic) makes a "fix the data first" approach impractical. Instead, organizations must take a pragmatic, workflow-led path: prioritizing a small number of high-impact use cases and incrementally building the semantic and governance foundations needed to support them.

Crucially, these capabilities cannot be developed in isolation, each reinforces the others, and value only emerges when they are orchestrated as a connected system. Here we explore these interconnected capabilities and how organizations can operationalize them to effectively leverage AI and scale its value.

Workflow redesign

To understand where agentic AI can deliver the greatest value, we analyzed 270 life sciences processes and assessed each across business value, cross-functional collaboration, technical complexity, and level of effort. The results highlight where organizations can move fastest, where to invest strategically, and where human-led decision-making remains essential.

High

Business impact

Low

Implementation complexity

Low complexity · High impact

~28%

of all workflows

Sweet spot

Regulatory Intelligence

Agents continuously monitor global regulatory authority updates, flag guideline changes relevant to specific products, and deliver structured intelligence briefings to regulatory and clinical teams replacing weeks of manual website scanning per product cycle.

Start here: fast ROI, low risk

High complexity · High impact

~26%

of all workflows

Adverse Event Case Processing

Agents handle the end-to-end lifecycle of individual safety case reports from intake and MedDRA coding through causality assessment, narrative writing, and quality review to regulatory submission operating across safety databases, medical ontologies, and authority portals simultaneously.

Strategic: build after wave 1 foundations

Low complexity · Lower impact

~12%

of all workflows

Cold Chain Qualification

Agents generate the validation protocol for a new temperature-controlled shipping lane and process sensor data from qualification runs to produce the summary report. The workflow is straightforward and data-driven, but it runs only when a new lane or packaging configuration is introduced so the cumulative time saving is modest.

Opportunistic: pick up alongside other work

High complexity · Lower impact

~17%

of all workflows

Biologic Manufacturing

Agents coordinate cell culture, purification, sterile fill-finish, in-process QC, batch release and validation management across deeply integrated manufacturing systems. The GxP validation overhead and novel biology make this a long-cycle transformation the value is real but the path to ROI is complex and organization-dependent.

Deprioritise: revisit once foundations are set

~16% **· Not agentifiable · human-led**

First-in-human trial design

Designing a first-in-human dose escalation study requires a qualified investigator to make judgements about participant safety margins, ethical risk thresholds, and regulatory accountability that cannot be pre-specified or delegated. Agents can support by surfacing precedents and assembling protocol templates but the decision itself must be owned by a human. Includes: QP sign-off · on-site GxP audits · FDA advisory presentations.

2 AI-ready data foundation and semantic layer

Building a full semantic layer in life sciences is inherently complex, given the diversity of data types and the challenge of embedding domain-specific business logic. Rather than attempting to solve everything upfront, organizations need to take a pragmatic, workflow-driven approach, building only the parts of the semantic layer required to enable high-value use cases.

In practice, this means focusing on targeted semantic layers, such as knowledge graphs, which create a shared, precise understanding of research, clinical, and commercial data, enabling automation and AI reasoning at scale with governance built in. Leveraging standards such as CDISC and Pistoia further establishes FAIR data, bridging internal silos and enabling cross-company collaboration in data spaces like EHDS and Sphin-X.

For example, in pharma R&D, the biggest barrier to unlocking the full value of AI is not a lack of data, but fragmentation and the loss of scientific context across systems such as ELNs, LIMS, instruments, and external databases. A semantic layer addresses this by unifying meaning across silos and creating a single, queryable view of research data, transforming disconnected facts into a living network of context. Trained on this foundation, AI agents can detect patterns, surface hidden insights, and predict outcomes with far greater precision, enabling intelligent, scalable action that accelerates R&D and improves time to market.

3 Agentic orchestration

As organizations move from isolated AI use cases to agentic AI embedded across end-to-end workflows, effective coordination becomes a limiting factor. Without a dedicated orchestration layer, agents are difficult to govern, integrate, and scale.

An enterprise agentic orchestration layer provides the shared control and coordination needed to turn individual agents into a coherent capability rather than a collection of disconnected tools. This enables governed interactions, seamless integration of internal and third-party agents, and the flexibility to evolve capabilities without disrupting core workflows.

4 Governance control panel

Governance is a prerequisite for scaling agentic AI. Without a unified control plane and central coordination, AI agents behave inconsistently and deviate from controlled workflows. This can increase GxP compliance risk, while weakening auditability and accountability.

A governance control panel provides the central oversight required to safely and sustainably deploy autonomous agents in regulated environments. It also ensures consistent controls across data access, model behavior, decision authority, and auditability, while aligning agent actions with scientific, clinical, and regulatory intent.

Guardrail catalog

Operational:

Ensure that generated content aligns with user expectations, does not drift away from its main purpose, or contain factually wrong or misleading information.

Privacy:

Use secure private environments, encryption, and pseudonymization to protect sensitive data before and during AI processing.

Regulation:

Validate that the generated content meets regulatory requirements and is compliant with frameworks like FDA (US), EMA (EU), ICH, GDPR, HIPAA, etc.

Security:

Extend security across the AI platform to protect sensitive data (e.g., PHI, clinical and R&D IP) with robust risk controls, monitoring, and mitigation mechanisms.

Ethical:

Verify if the generated content is not toxic, harmful, biased, or based on stereotypes, and filter out any such inappropriate content before it reaches customers.

5

Human-AI chemistry

In GxP contexts, human-AI interaction is not just a usability concern but a regulatory requirement. Compared to standard agentic systems, it must be designed for heightened oversight, traceability, and accountability.

Effective human-AI chemistry requires three elements integrated from the start: clarity on what the agent is responsible for and where human judgement takes over; transparency in how the agent reasons and what data it uses; and feedback loops that enable continuous learning and evolution based on scientific knowledge, regulatory expectations, and operational realities.

Staying ahead: What's next for life sciences

Proprietary data

As AI commoditizes access to models, tools, and even syndicated data, competitive advantage is shifting decisively toward organizations that build and scale proprietary data flywheels. The differentiator is no longer who has access to data, but who continuously generates, enriches, and learns from data that competitors cannot replicate, such as patient registries, real-world outcomes, longitudinal device performance, and novel biomarker insights.

This requires a fundamental shift in data strategy: from aggregating what exists to systematically creating and compounding exclusive data assets. These flywheels not only differentiate the organization, but strengthen over time, improving models, sharpening decisions, and reinforcing competitive advantage.

A semantic layer is critical to activating proprietary data. It structures and contextualizes information so AI systems can interpret it consistently, operate within regulatory constraints, and generate the insights that will create advantage.

Workforce shifts

Agentic AI deployed at scale will reduce headcount in specific functions, not through dramatic restructuring, but through gradual redeployment as agents absorb cognitive tasks currently performed by junior-to-mid level professionals. Leaders must plan for this honestly and deliberately rather than discovering it retrospectively.

On the other hand, demand for new roles will grow. For example, organizations may require human orchestrators to manage cross-functional agent networks or clinical AI validators to audit outputs for scientific and regulatory accuracy. Building these human capabilities is a strategic investment in the skills that create and sustain competitive advantage.

The agentic imperative in life sciences

The life sciences industry stands at an inflection point. Patent cliffs, pricing pressure, hardware commoditization, and rising evidence costs have created a burning platform demanding a fundamentally different operating model. Agentic AI, deployed on redesigned workflows across a unified data and semantic foundation, is the architecture of that model.

The goal is not to replace human expertise, but to free it, using AI to take on cognitive tasks so experts can focus on where their judgment matters most.

In this paper, we offered a path forward. The organizations that move with conviction now, will shape the competitive landscape for the next decade. Those who wait for certainty will find that the landscape has been defined without them.

Authors

Neeraj Goel Mittal

Vice President, Head of Life Sciences Industry Europe, Capgemini Insights & Data

neeraj.goelmittal@capgemini.com

Thorsten Rall

Global Head of Life Sciences Industry, Capgemini

thorsten.rall@capgemini.com

Dr. Parvin Moyassari

Global Head of Life Sciences, Capgemini Insights & Data

parvin.moyassari@capgemini.com

Contributors

Gayathri Balappa
Capgemini

Brian Eden
Capgemini

Ivana Knyght
Capgemini Invent

Geoff McCleary
Capgemini

Frédéric Burger
Capgemini Engineering

David Ghesquieres
Capgemini

Ross Landles
Frog

Justin Melnick
Capgemini Invent

Jacques-Antoine Duret
Capgemini Engineering

Amit Kakatkar
Capgemini

Barnabé Lecouteux
Capgemini Invent

Mini Nair
Capgemini

Further reading

Smart bet, only option or both? Biopharma R&D turns to AI

AI is set to redefine how new therapies are discovered and developed, with most executives expecting a fundamental shift in R and D in the coming years. Yet unlocking this potential requires more than technology alone. This report combines data, industry benchmarks and leadership perspectives to highlight where AI is already improving outcomes and what it takes to scale it across the organization.

Download the report here.

Life Sciences engineering and R&D pulse 2026

Life sciences engineering leaders are under pressure. Innovation cycles are lengthening. Compliance demands are intensifying. AI is disrupting existing operating models. Global volatility is testing the resilience of supply chains, talent models, and legacy systems. The Life Sciences Engineering and R&D Pulse 2026 provides a practical, benchmark driven view of how engineering leaders and their C suite colleagues are responding and where execution is still falling short. Based on insights from 200 C level and senior life sciences leaders globally, this report examines how organizations are closing the gap between strategic ambition and operational readiness.

Download the report here.

From pilots to real impact: How enterprises actually scale Agentic AI

Most enterprises have no shortage of AI pilots. What they lack is real impact at scale. This report explores why the vast majority of agentic AI initiatives fail to move beyond the demo stage and what sets the top performers apart. It highlights how leading organizations are embedding AI into core workflows, governance and operations from day one, and outlines the practical foundations needed to turn promising experiments into measurable business performance.

Download the report here.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real.
www.capgemini.com