Deloitte.

Together makes progress

Tech Trends 2026

A life sciences perspective

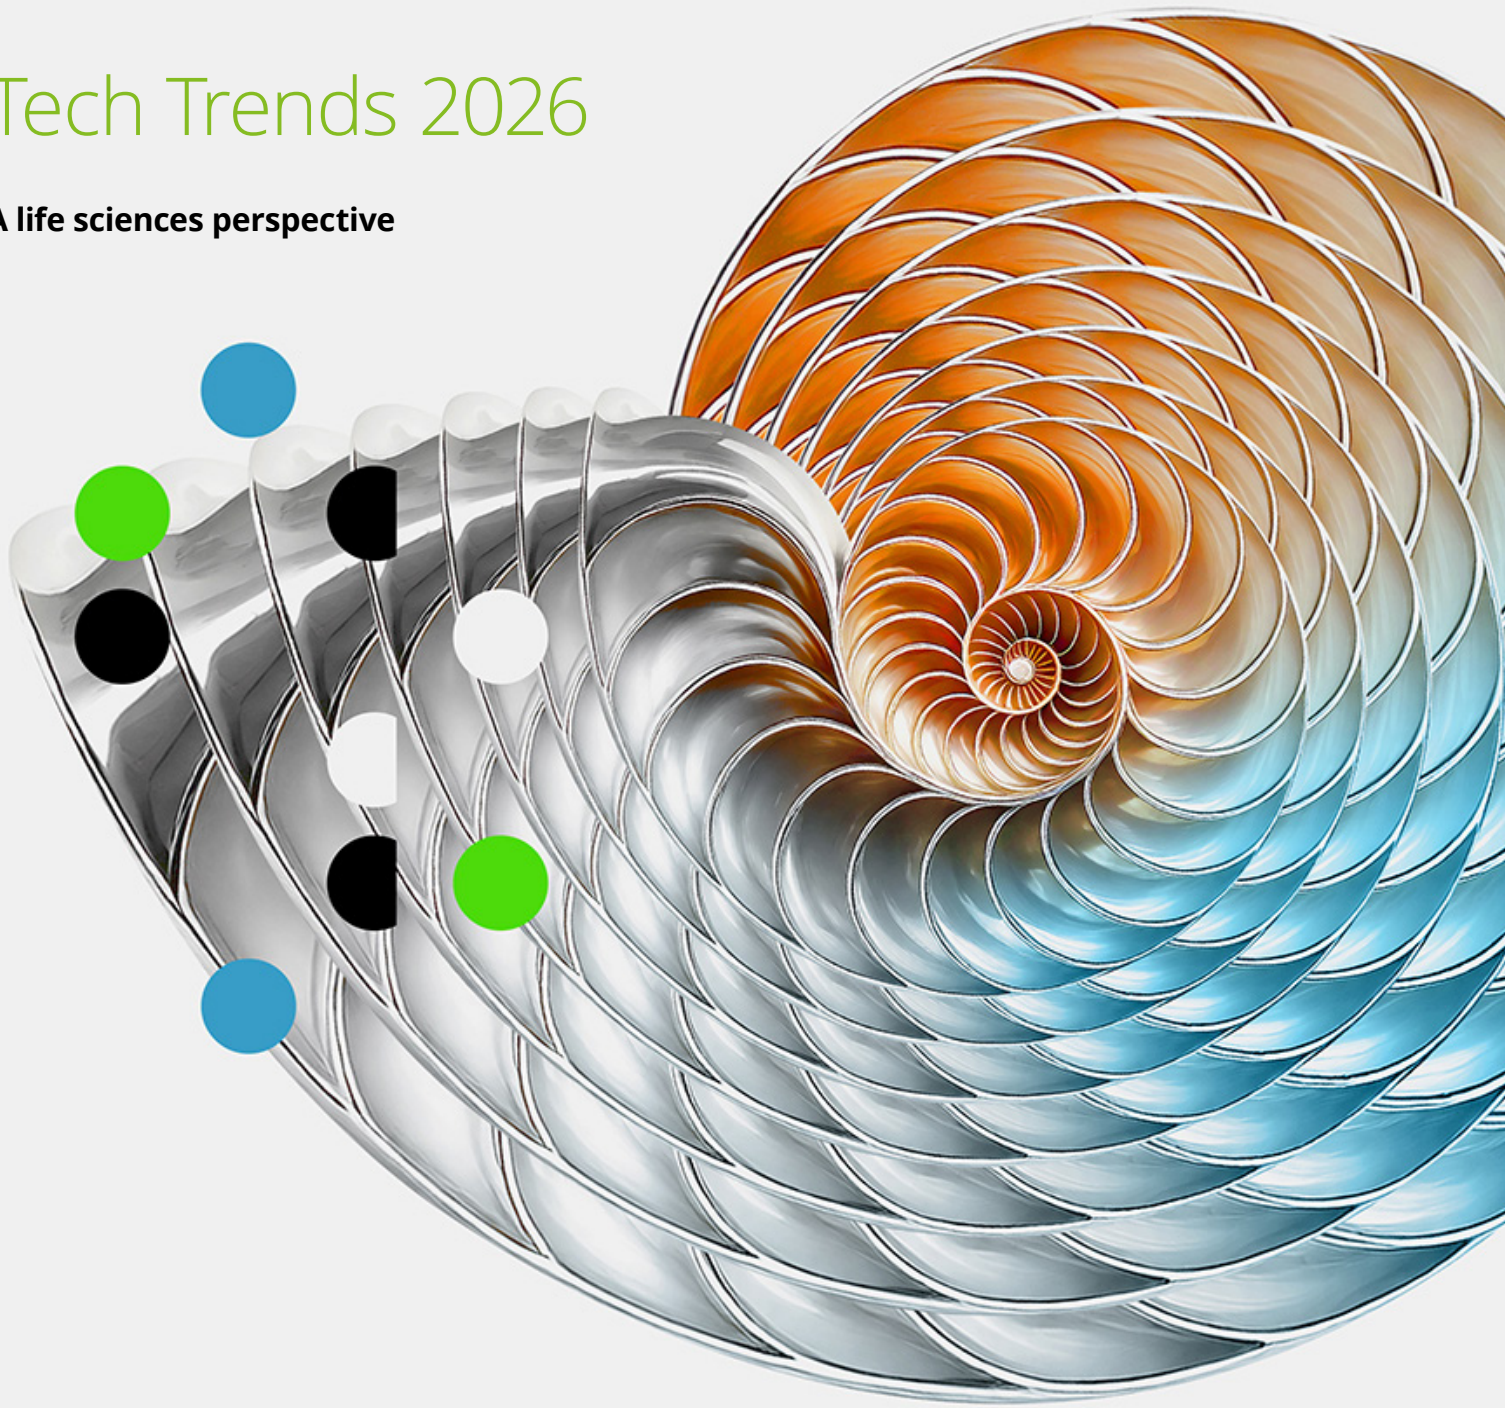

Introduction

Artificial intelligence (AI) is transitioning from proofs of concept to scaled deployment, reshaping how life sciences organizations discover therapies, run trials, manufacture at scale, engage customers, and protect intellectual property. For Biopharma and Medtech leaders, the question is no longer whether to adopt AI, but how to operationalize it responsibly while meeting regulatory expectations, quality standards, and patient-safety commitments.

> "Life sciences stands at the threshold of its most transformative era, where the convergence of biological data, computational power, and AI is rewriting the rules of discovery, development, and delivery. Organizations that move boldly from isolated pilots to enterprise-wide AI adoption will not simply outcompete; they will define the new standard of what's possible in human health."
>
> **Tarun Sharma**, Life Sciences Technology Strategy Leader, Principal, Deloitte Consulting LLP

This life sciences perspective provides an industry-specific view on the five AI-driven technology trends highlighted in Deloitte's Tech Trends 2026, translating cross-industry themes into implications for Biopharma and Medtech.

The following pages reflect conversations with Deloitte life sciences industry leaders and subject matter specialists across Biopharma and Medtech, with a focus on what is changing, where it is showing up first, and what leaders can do next. We explore the five trends expected to influence Biopharma and Medtech over the next 18 to 24 months:

- **AI goes physical:**
  Navigating the convergence of AI and robotics.
- **The agentic reality check:**
  Preparing for a silicon-based workforce.
- **The AI infrastructure reckoning:**
  Optimizing compute strategy in the age of inference economics.
- **The great rebuild:**
  Architecting an AI-native tech organization.
- **The AI advantage dilemma:**
  Securing and leveraging AI for cyber defense.

How to use this report

Each trend chapter includes: (1) what is changing, (2) where the trend is already showing up across the life sciences value chain, and (3) practical actions leaders can take in the next 30 to 90 days. Use the actions as a planning menu: Select a small set of priorities, define measurable outcomes, and establish a learning cadence to refine decisions as adoption expands.

The scoring approach

To help leaders prioritize, we assess each trend using two measures.

Relevance (on a scale of 1 to 5) reflects how strongly the trend can affect business outcomes in the next 18 to 24 months.

Readiness (on a scale of 1 to 5) reflects how prepared organizations are to act today, considering governance, validation, talent, operating model, and technical foundations.

Relevance

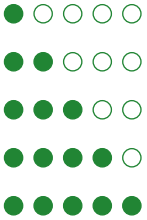

Readiness

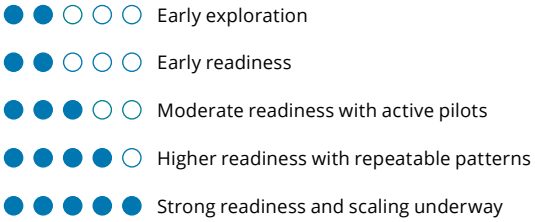

Trend 1: AI goes physical

Navigating the convergence of AI and robotics

Physical AI is evolving robots from pre-programmed machines into adaptive systems that perceive, learn, and operate autonomously in complex environments.

These capabilities show up in industrial robots, autonomous vehicles, drones, and other systems. Current challenges include gaps in training, safety concerns, and cybersecurity risks, but falling costs are extending adoption beyond smart warehousing and supply chain operations into the mainstream. Humanoid robots are the next frontier, with projections of 2 million workplace humanoids by 2035. Future developments may include bio-hybrid robots and quantum robotics.

Key takeaways for life sciences

Biopharma perspective

- The use of AI and robotics is more prevalent in manufacturing and operations than in research and development.
- Sterile injectable manufacturing lines are highly roboticized, with minimal human touch, but these robots have limited advanced AI capabilities.
- In Biopharma, automation and AI have historically focused on optimizing drug targets and lead compounds.
- Current use of automation is focused on repeatable tasks (e.g., pipetting, fill-finish lines) with advanced AI integration still in early stages.
- Industry will continue to experiment with physical AI, but widespread adoption will be gradual due to regulatory and infrastructure challenges.
- Humanoid robots in labs are still emerging and not fully adopted.

Further consideration

Most Biopharma R&D labs are in the early stages of considering how robots could help lab operators or scientists execute their tasks more efficiently. That said, several large Biopharma companies have touted the fact that their injectable drug delivery cartridges are manufactured completely by robots, with no human touch. The process is completely roboticized, from the point where the tank is loaded full of drug substance, to filling and adding sterile water to make the injectable, to the point in time it comes out the other side.

Such mechanisms run on a repeated cycle, with very limited advanced AI tied into the robotic manufacturing process.

The question is, can AI sit on top of those robots so they can actually do and correct things, much like a human would, but with more precision and integrity in terms of data capture and data traceability to the robots.

Medtech perspective

- In the Medtech sector, robots are being used in manufacturing and advanced surgery products.
- Medtech operations use robotics for repeatable tasks, but advanced AI is less prevalent.
- Medtech tends to lag Biopharma in innovative AI adoption by about 12 months.
- Physical safety is a new concern for AI in robotics, alongside traditional AI risks.

Further consideration

In Medtech and other sectors, robots are programmed to avoid contact with humans, but reporting incidents can be a challenge.

Adoption of robots on manufacturing floors is expected to increase over the next 24 to 48 months, as workflows are reimagined. Trust and risk will be key factors in how quickly and comprehensively robots with AI capabilities are accepted.

Actions to consider in the next 30 to 90 days

Identify two to three manufacturing or operational workflows where robotics is already present and AI can add measurable value, such as traceability, precision, or exception handling.

Establish a physical AI risk review that includes occupational safety and incident reporting considerations alongside cyber and model risk.

For R&D environments, **focus** on repeatable lab tasks and vendor-integrated automation while documenting where variability limits near-term standardization.

Trend 'relevance' and industry 'readiness'

Relevance and readiness vary significantly by function. Manufacturing and operations show higher relevance and earlier readiness because validated robotic workflows already exist. R&D readiness is lower because lab variability makes the business case and standardization harder to establish.

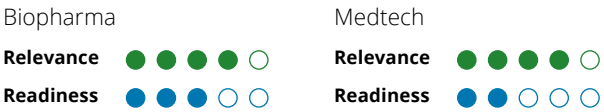

From our leaders

"In Biopharma R&D, the emphasis today is on using AI to help generate and prioritize drug candidates, rather than deploying humanoid robots in the lab to speed up discrete steps. The business case for advanced physical AI robots in R&D is still taking shape, and it is an area we are actively exploring. In the near term, the clearer path is the lab of the future—a digitally enabled, highly automated research environment powered by AI, robotics, and cloud computing—where organizations can streamline scientific workflows, strengthen multimodal data management, and accelerate internal innovation. Over time, this kind of environment can reduce trial-and-error experimentation and help teams identify novel candidates faster."

***Raveen Sharma**, Life Sciences R&D Leader, Managing Director, Deloitte Consulting LLP*

"On the quality control (QC) side, we are still in the seed stage of exploring how to bring robotics into day-to-day actions on the QC lab floor, in part because so much work involves ad hoc problem-solving that has been difficult to standardize. In parallel, our research points to growing investments in robotic automation, instrument connectivity, AI, and generative AI to automate data capture and sample management, improve asset and resource utilization, and strengthen end-to-end traceability. Looking ahead, agentic AI that can sense, reason, and act could connect these threads by enabling intelligent QC workflows with real-time monitoring, early anomaly detection, and dynamic resource allocation, while still supporting the flexibility labs rely on."

***Laks Pernenkil**, Life Sciences Supply Chain Leader, Principal, Deloitte Consulting LLP*

Trend 2: The agentic reality check

Preparing for a silicon-based workforce

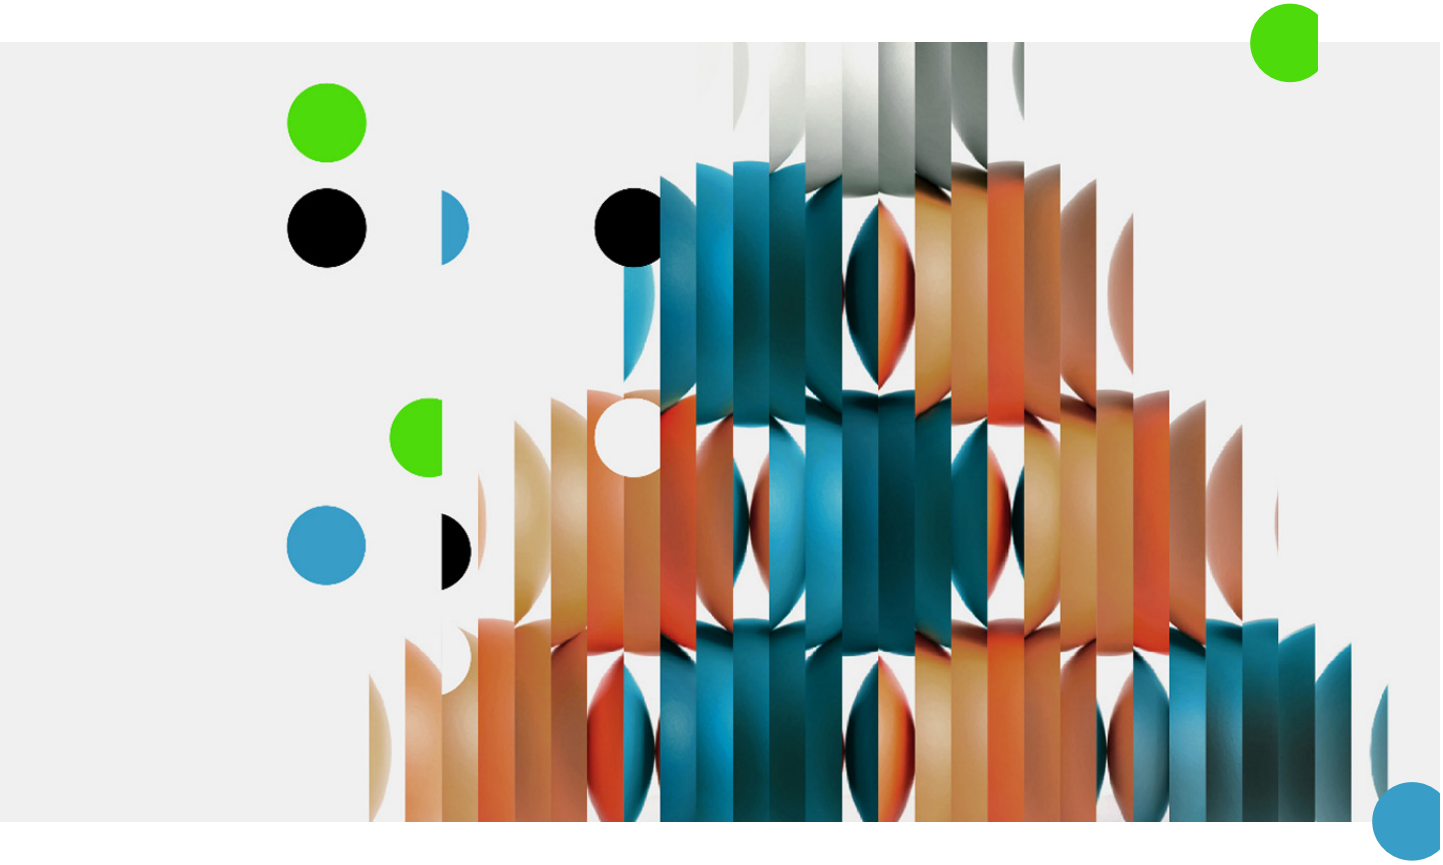

Despite early enthusiasm, many businesses have yet to see significant transformation from agentic AI implementations because most simply automate existing processes rather than fundamentally redesigning operations.

Only 11% of surveyed organizations have deployed agentic systems in production, with challenges including legacy system integration, data architecture constraints, and inadequate governance frameworks. Leading organizations are adopting agent-first process redesign, implementing multiagent orchestration using emerging protocols, and treating agents as a silicon-based workforce requiring specialized management frameworks. This includes agent onboarding, performance tracking, and FinOps (financial operations) cost management. The future points toward graduated autonomy levels, hybrid human-digital workforces, and leveraging agent-generated data for continuous learning, transforming how enterprises operate and compete.

Key takeaways for life sciences

Biopharma perspective

- Biopharma clients are starting to invest significantly in agentic AI, while also paying close attention to the ROI.
- Adoption of agentic AI will require a shift in mindset and a redesign of operations, especially in clinical operations and data management.
- Retaining a human in the loop is critical for GxP-compliant processes, as regulators expect job aids to ultimately receive a human sign-off.
- The workforce should shift from task-based roles to outcome-based roles, where humans are teaching agents to achieve desired outcomes.

Further consideration

Biopharma companies have started investing in both generative AI and agentic AI. Organization leaders have emphasized a pragmatic approach to measuring value, focusing on both the ROI from use cases and the broader shift in organizational mindset toward AI adoption.

Some of the largest Biopharma companies are pursuing agentic AI or AI mandates across their organizations, but none have yet to realize optimal utilization. They are learning and iterating, aiming to reach the right approach. About 30% to 40% of initiatives saw positive ROI, but the main gain was in changing ways of working and thinking about AI as a natural evolution for organizational processes.

Medtech perspective

- Medtech companies typically lag behind Biopharma in AI adoption, but they are now prioritizing cross-functional programs. Examples include cross-channel funnel management, proactive product quality using Internet of Things data, and inventory optimization via agentic control towers.
- Regulatory challenges persist, especially with software validation and clinical decision support. Human oversight is essential for critical decisions.

Further consideration

Medtech has used AI and machine learning in its products for a long time, but regulatory validation remains a challenge. For the foreseeable future, "humans in the loop" will be a necessity for critical decisions. Humans will remain involved in workflows handling auditing and documentation to meet regulatory requirements.

Initial enthusiasm for integrating agentic systems led to pilot projects in the 2022 to 2024 time frame, but true transformation was limited due to a focus on automating existing processes rather than redesigning them.

In one example of AI adoption, a device company created a "cross-channel funnel management," where agentic AI is used to unify customer interactions across channels (calls, emails, website, e-commerce, reps) and provide personalized guidance to both customers and employees. In another example, a cardiovascular company used agentic AI for inventory optimization, starting with a sprint to build a control tower that guides planners in making better inventory decisions.

Barriers to adopting agentic AI include:

- Lack of clear vision and alignment on what organizations are trying to achieve with agentic AI.
- Need for fundamental change in the operating model. Without this, AI projects only deliver incremental value.
- Workforce concerns about job disruption can slow adoption and make outcomes harder to sustain.
- Talent shortages and the challenge of keeping up with the fast-moving tech and business landscape.

Actions to consider in the next 30 to 90 days

Select one end-to-end workflow that crosses functions and redesign it for agent participation, rather than deploying agents into existing steps.

Define human oversight points for regulated decisions, including what triggers require human sign-off and how auditability is documented.

Build a phased roadmap for agent rollout with measurable milestones every few months to support confidence, governance checkpoints, and value tracking.

Trend 'relevance' and industry 'readiness'

Both sectors see agentic AI as highly relevant. Readiness differs. Biopharma is advancing through investment and operating model experimentation, while Medtech progress is often constrained by funding envelopes and validation expectations, especially for decision-support scenarios.

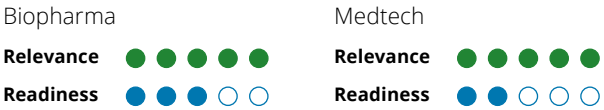

From our leaders

"On the agentic horizon, I see agents taking on more of the execution, while people shift toward oversight—confirming the outcome is correct, addressing exceptions, and moving the work forward. That shift calls for a mindset change: We have spent years talking about skills, but skills are ultimately in service of tasks. What starts to matter more is the outcome—and how we redesign work so agents can contribute to it, rather than layering agents onto processes built for humans. In that model, the question becomes, 'How do I teach agents to work with me toward the outcome?' Your focus shifts to: Am I providing the right context and intent? Am I setting the right boundaries for autonomy? Am I supervising performance in a way that is measurable and observable? And am I applying the ethical, risk, and governance guardrails that help build trust—so orchestration is disciplined, decisions remain accountable, and human judgment stays in the loop where it counts most."

Puneet Sharma, *Medtech Commercial Leader, Principal, Deloitte Consulting LLP*

"A lot of our clients are rethinking how they have organized clinical trial operations, especially across clinical data management and site orchestration. Historically, these capabilities often sat in separate teams because the workflow depended on handoffs and specialized roles. With agents in the picture, teams can automate more of the repetitive coordination work and respond to signals in near real time, which is prompting a different question: Do our operating models still make sense for how trials run today, or do we redesign them around outcomes and end-to-end orchestration across people, platforms, and partners?"

Nirav Mehta*, Life Sciences and AI & Engineering Leader, Principal, Deloitte Consulting LLP*

Trend 3: The AI infrastructure reckoning

Optimizing compute strategy in the age of inference economics

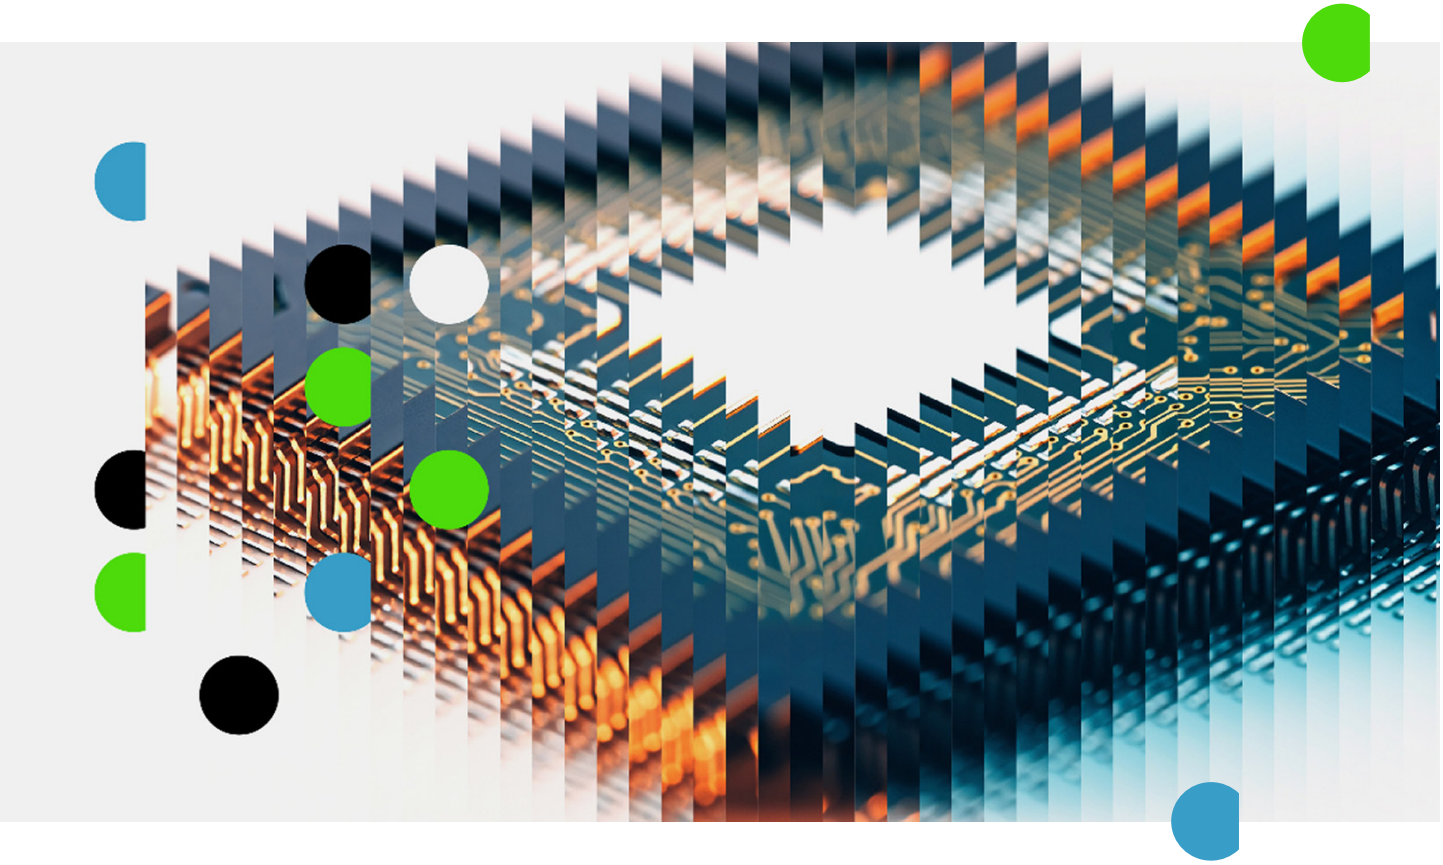

As AI moves from experimentation to production, enterprises face an infrastructure dilemma. While token costs have dropped substantially, overall AI spending is exploding due to massive usage growth.

Organizations are hitting a tipping point where cloud services become cost-prohibitive for high-volume workloads, with monthly bills reaching tens of millions of dollars. Leading enterprises are adopting strategic hybrid architectures: cloud for variable workloads, on-premises for consistent production inference, and edge for latency-critical applications. This requires purpose-built AI data centers featuring graphics processing unit (GPU)-optimized hardware, advanced networking, and specialized cooling. Future challenges include workforce reskilling, AI agents managing infrastructure, and sustainable computing innovations like renewable-powered and potentially orbital data centers.

Key takeaways for life sciences

Biopharma perspective

- Most Biopharma organizations do not yet have a mechanism to analyze energy costs on their business model, but the issue is emerging.
- Current discussions focus more on token-level costs and optimizing computing power on cloud providers rather than deeper supply chain or energy sourcing.
- Large Biopharma companies are investing in enterprise-scale GPU clusters but have not yet moved into deeper supply chain investments, like chip or data center ownership.
- Biopharma has not yet expressed interest in mini-nuclear reactors for powering data centers.
- Biopharma organizations are leading consumers of hyperscale cloud capacity and generally rely on external providers, shifting select workloads in-house only when specialized, high-intensity compute is required.
- Computing latency is generally not a problem in Biopharma, except possibly in manufacturing lines where immediate action is needed to prevent production issues. In R&D and other areas, thoughtful, accurate answers are prioritized over speed.

Further consideration

Biopharma companies are not as concerned about latency issues, unless it's to stop the manufacturing line if there is a problem. Drug development takes decades sometimes, so a few seconds or a minute or two likely won't make a difference. It's better to get the answer right versus getting it fast. That said, latency does matter at times, and companies may need split-second decisions in manufacturing so that they're creating optimal golden batches each time.

Medtech perspective

- Optimization of infrastructure costs, such as computing, storage, and tokens, is often neglected, with manufacturers relying heavily on hyper-scaler guidance and deal-centric decisions.
- AI models are typically built for accuracy rather than infrastructure efficiency, with optimization taking a back seat.
- Edge computing is prioritized in Medtech due to device proximity to hospital and patient settings, a factor that influences infrastructure strategy for device monitoring.
- Infrastructure concerns are most prominent in commercialization, especially in contact center operations (handling large volumes of patient/doctor calls and outreach).

Further consideration

Medtech companies rarely strategize infrastructure costs up front, with only one or two out of 10 clients considering infrastructure early.

Of the various infrastructure models, edge-first computing is highly relevant for Medtech due to device proximity to patients/hospitals. Regulatory approvals for AI-embedded devices are increasing, and edge computing is rising for both business and internal operations. Infrastructure cost optimization is often deferred until after use cases are implemented.

On the other hand, only a handful of clients pursue sustainability goals. Green computing maturity is low in Medtech, and concepts like orbital computing with data centers located in space are not widely considered.

In general, some Medtech clients are beginning to consider computing and energy costs together to standardize AI workloads and manage costs. Computing costs are a near-term concern as resources diminish.

Actions to consider in the next 30 to 90 days

Create an infrastructure view for priority AI workloads that include compute, storage, token economics, and where edge architectures are required.

For Medtech, **assess** commercialization and post-market surveillance workloads first, since they often surface cost and architecture constraints early.

Establish a basic measurement approach for AI workload cost and performance so optimization does not start after scale pressure builds.

Trend 'relevance' and industry 'readiness'

Relevance is rising as AI moves from pilots to production, but readiness remains limited in many organizations because infrastructure planning often starts after use cases are selected. This is especially visible in high-volume commercial workflows and device monitoring programs, where costs and architecture decisions surface quickly.

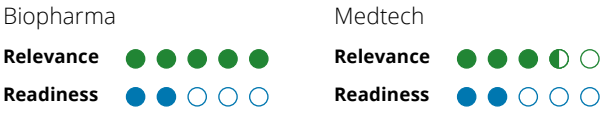

From our leaders

"One way to think about AI compute cost is to ask, 'How far down the stack do we want to go?' Many leaders start at the surface with token spend: What is driving consumption, and how do we optimize prompts, models, and agent workflows so inference does not run continuously by default? As AI moves from pilots into production, that question gets sharper because usage can grow faster than unit costs fall, and organizations can see bills rise quickly when agentic systems are running near-constant inference. From there, the conversation often shifts to where the work should run. Instead of treating it as a cloud-versus-on-premises debate, teams are beginning to take a workload-driven approach: using elastic capacity for variable demand; shifting steady, high-volume inference to private infrastructure; and pushing latency-sensitive workloads closer to where decisions are needed. The deeper question is what sits behind those choices: data center design, power and cooling, resilience, and the operating model to monitor and orchestrate workloads over time. Most organizations are actively working the token and platform layers today, and only a smaller set are pushing further down into the underlying infrastructure and supply chain decisions."

***Subadhra Parthasarathy**, Life Sciences AI & Data Leader, Managing Director, Deloitte Consulting LLP*

"Infrastructure tends to come into focus in two places. First is commercialization, especially service and engagement at scale—contact centers and omnichannel outreach where high volumes of inbound and outbound interactions create constant demand for orchestration, personalization, and faster resolution. As agentic capabilities mature, the conversation shifts from isolated automation to end-to-end service models that coordinate work across channels and teams, with people staying involved for judgment and exception handling. The second place is post-market surveillance and connected care, where devices in real-world settings generate continuous data streams. When you have connected devices—including wearables and smart implantables—the value is not only the device data itself, but how well it integrates with clinical systems and workflows to produce usable signals. That puts the spotlight on architecture choices: interoperability and workflow integration, secure and privacy-aware data exchange, clarity on who can use the data and for what purpose, and the ability to manage complex data feeds over time. The question becomes, 'How do we design the platform so we can capture the right data, make it actionable, and scale responsibly across care settings?'"

***Mukund Lal**, Medtech Senior Manager, Deloitte Consulting LLP*

Trend 4: The great rebuild

Architecting an AI-native tech organization

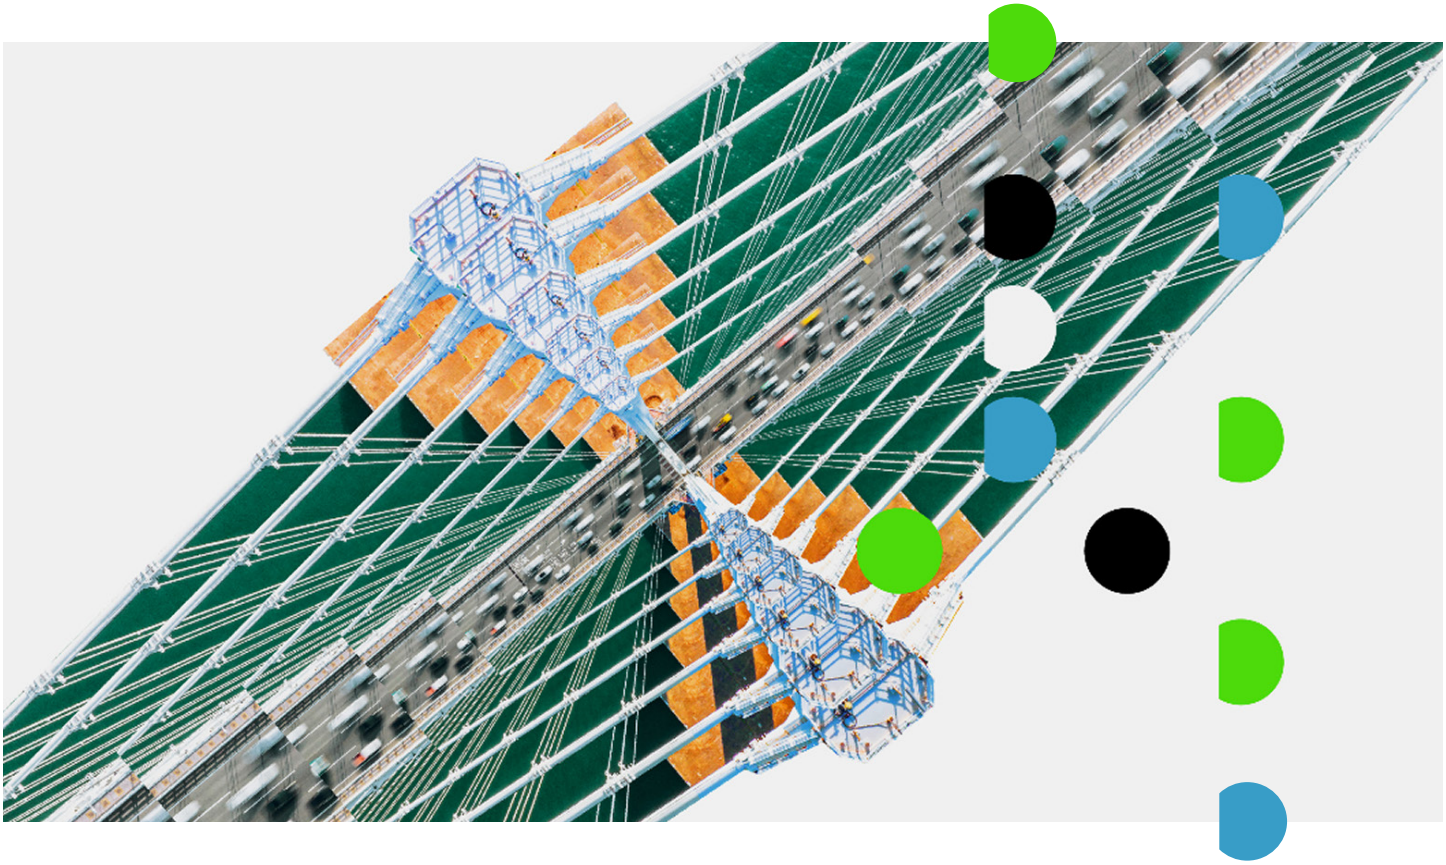

AI is fundamentally restructuring technology organizations beyond simple automation.

With 64% of organizations increasing AI investments, and tech budgets for AI rising, priorities are shifting from infrastructure maintenance to strategic leadership. Leading organizations are anchoring AI initiatives to measurable business outcomes, designing modular architectures for flexibility, and redefining talent strategies around human-machine collaboration. New roles are emerging, such as AI collaboration designers, edge AI engineers, and prompt engineers, while CIOs evolve from tech strategists to AI evangelists and orchestrators. Future tech organizations will feature agentic architectures, lean product-led teams, blended human-agent workforces, adaptive governance, and ecosystem-oriented innovation. Success requires embracing continuous evolution and boldly reimagining operations rather than incremental change.

Key takeaways for life sciences

Biopharma perspective

- Current business models may not be agile or ready enough to adopt the new AI-driven models.
- Organizations are facing difficulties in upskilling or reskilling existing staff, even as much effort is being put into these efforts.
- Over the next 24 months, organizations will need to decide whether to retain or disrupt their workforce based on future needs.
- Challenges to adoption include deeply embedded work habits and ingrained organizational cultures.
- Transitioning to AI-native operations demands a fundamental change in how employees perceive their roles and add value, which is challenging for many.

Further consideration

Over the next 24 months, organizations will start making choices on whether to keep or disrupt the existing workforce based on future needs and AI capabilities. This moment of reckoning hasn't hit yet because an immense amount of time and effort is still going into training and upskilling employees on the new AI models.

Strategies to help companies accelerate AI adoption include clearly defining the specific roles AI agents will play within business processes, rather than relying on general agentic capabilities. Such clarity helps differentiate and streamline transformation efforts.

Some organizations already are seeing positive outcomes through AI integration. These include shortened cycle times around compound testing and development. But overall adoption is still in the early phases and not yet widespread.

Medtech perspective

- The Medtech sector is far from being truly AI native, with organizations currently embedding AI mainly through legacy platforms and vendor solutions.
- While there is a push to leverage AI, much of the focus is on training the workforce and addressing concerns about job changes.
- Most AI adoption is system-centric, with organizations considering themselves AI native only if their systems offer quick AI features.
- Some companies have seen successes with AI adoption, but not at scale. Improvements have occurred only in small pockets within large enterprises.
- In commercial functions, such as marketing and ad agency work, AI has led to minor cost savings but not significant cost reductions.
- The most notable productivity gains have been in internal software development, where AI tools have shortened development and testing cycles, but this mainly benefits internal operations rather than revenue generation.
- Some organizations are hiring more people who understand advanced AI concepts (like agentic language and retrieval-augmented generation models) and can architect AI solutions. But these are not yet standard job titles or roles across the workforce.

Further consideration

Most Medtech organizations are far from becoming AI native at the enterprise level, even though some are moving in that direction and could see progress over the next three to five years.

When it comes to adopting new technologies such as AI, life sciences organizations historically have been more inclined to buy a ready-made solution rather than build one.

Such a buy-first approach results in quick adoption of the technology but fails to prepare the workforce for holistic change. In AI, this leads to a cycle where people are not fully ready for AI-driven transformation.

On the positive side, individual productivity in life sciences organizations has increased due to widely adopted AI-enabled productivity and knowledge-assistance tools. These gains, however, are still largely personal rather than scaled across enterprise workflows.

A top-down push is needed to accelerate AI adoption across business units.

Actions to consider in the next 30 to 90 days

Define what "AI native" means for your organization beyond platform features, including governance, validation, and outcome measurement.

Identify two workflows where leadership can set clear direction and accountable outcomes, helping teams move beyond experimentation.

Prioritize orchestration, architecture, and process redesign skills to support agent patterns, even if job titles vary across organizations.

Trend 'relevance' and industry 'readiness'

Leaders recognize the importance of AI-native transformation, but enterprise-wide readiness is still developing. Progress often appears in pockets, with stronger evidence in internal technology delivery than in end-to-end business transformation.

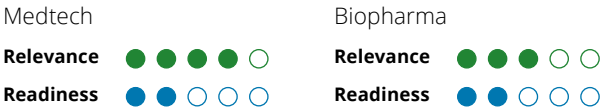

From our leaders

"Organizations that define up front the specific roles they want agents to play inside a specific business process tend to make faster progress and stand out in how they scale. Instead of starting with a broad 'agentic capability,' it helps to truly reimagine work and assign clear, role-based responsibilities—for example, an agent that drafts, an agent that reviews, an agent that validates against policies, and an agent that routes exceptions for human judgment. That role clarity makes orchestration more practical, improves consistency, and creates a cleaner path from pilots to production because teams can measure performance, monitor outcomes, and refine the design over time. The key is not 'an agent that does X, Y, or Z.' The key is a defined set of agents with defined responsibilities in a defined workflow—coordinated through orchestration and paired with the right governance and oversight to keep decisions accountable as autonomy increases."

***Todd Konersmann**, Global Life Sciences Sector Leader, Principal, Deloitte Consulting LLP*

"We have seen some clients where the CEO or CIO sets a clear mandate—for example, embedding hundreds of AI agents within the next two years. That kind of direction tends to activate every business unit and functional leader because it moves the conversation from pilots to operating-model change and measurable impact. In many cases, the momentum comes from redesigning how work gets done, not simply layering agents onto processes built for humans, and from treating agents as a form of digital labor that needs orchestration, governance, and ongoing management. By contrast, organizations that take a wait-and-see approach often stay closer to incremental adoption—looking for standard tools to buy and then figuring out where they fit. The risk is that they optimize tasks in pockets but miss the compounding value that can come from coordinated, multiagent workflows and a workforce model built for human-plus-agent collaboration."

***Mukund Lal**, Medtech Senior Manager, Deloitte Consulting LLP*

Trend 5: The AI advantage dilemma

Securing and leveraging AI for cyber defense

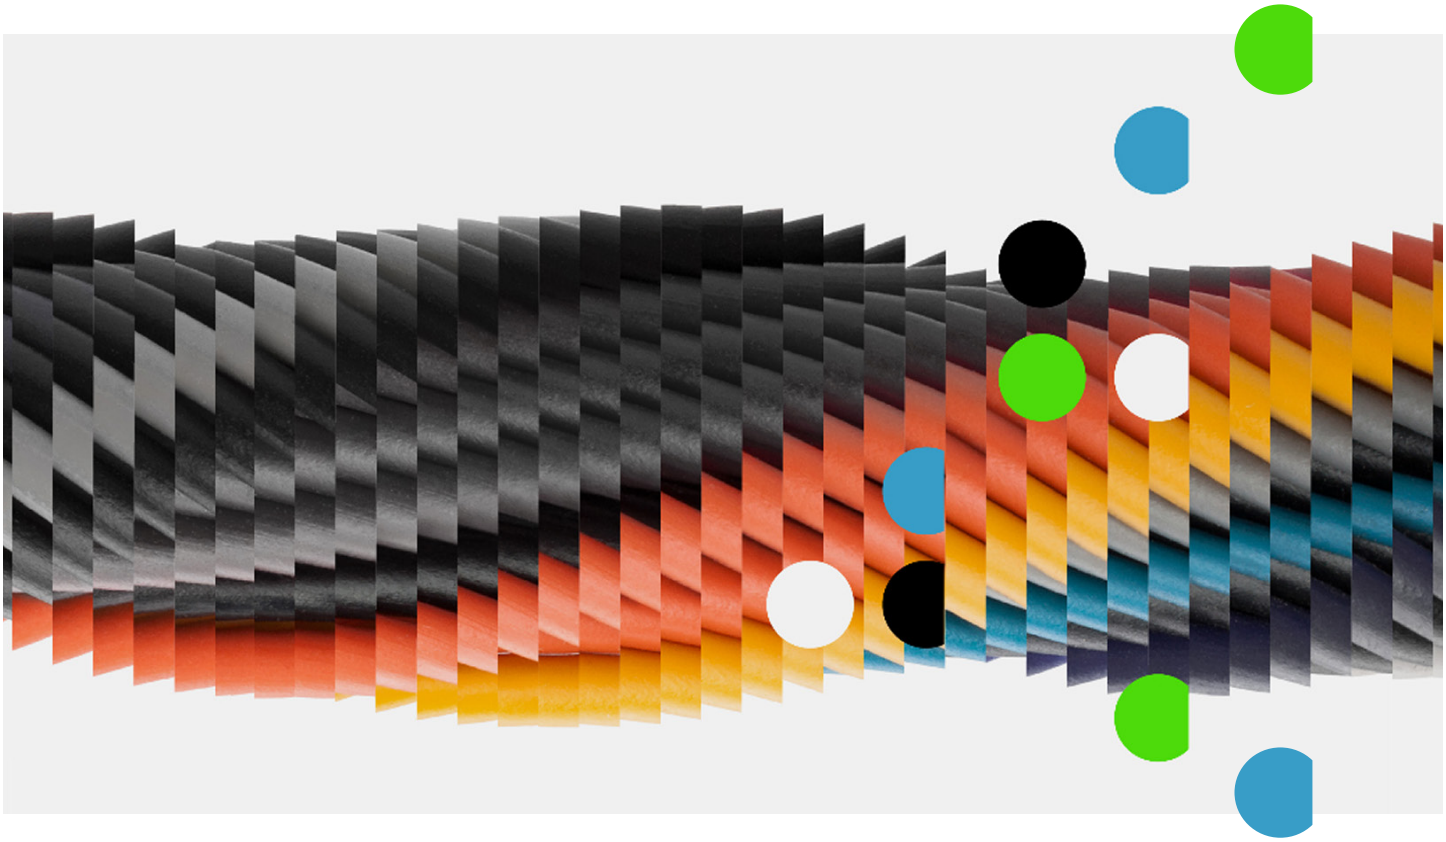

AI creates a cybersecurity paradox: The same capabilities driving business innovation are also introducing new risks.

Organizations face threats from shadow AI deployments, adversarial attacks, and intrinsic AI system weaknesses across four domains: data, models, applications, and infrastructure. Existing security practices can be adapted to address AI-specific risks through robust access controls, model isolation, and secure deployment architectures. And AI offers powerful new capabilities to counter the very vulnerabilities it creates. Leading organizations are leveraging AI defensively through red teaming with AI agents, adversarial training, and automated threat detection at machine speed. Future challenges include AI-physical infrastructure convergence, autonomous cyber warfare, and quantum/space security threats. Success requires embedding security into AI initiatives from inception, treating it as an enabler rather than a constraint on innovation.

Key takeaways for life sciences

Biopharma perspective

- A model extraction is the most critical security risk for Biopharma, as it could result in loss of valuable intellectual property (IP) and years of R&D investment. A model extraction attack happens where someone replicates an AI language model by systematically querying it and using its outputs to train a surrogate model.
- Key concerns also include shadow AI (i.e., the unauthorized use of AI by an organization's employees) and uncontrolled AI usage, which could expose proprietary data, and misconfigurations at various levels (data, model, application, infrastructure).
- Security safeguards include restricting access to unofficial tools, providing employees access only to approved tools, implementing access controls, continuous monitoring, network monitoring, and alerting security teams to anomalies.
- While data in Biopharma is more sensitive, due to the presence of patient health information, the fundamental controls—data governance, data classification, and access management—are similar to those used in other industries.

Further consideration

Biopharma companies face significant risks related to AI exposing their proprietary data. If a company's AI model is extracted by an attacker or nation-sponsored hacker, it could mean losing decades of research and millions of dollars.

Standard platforms for electronic health records, finance, HR, and CRM come with built-in security and privacy measures, which companies can leverage when building AI solutions. These controls provide some relief but do not eliminate the need for additional safeguards.

Additional safeguards include applying data encryption, strict access controls, and role-based identity management to acompany AI models.

Scenario-based security training should include showing research teams what can happen if they copy intellectual property into common AI tools. Such actions could expose proprietary information to the public or competitors.

Medtech perspective

- Intellectual property leakage is a top concern for Medtech clients, especially with the use of generative and agentic AI models in R&D and commercial functions.
- Medtech has used machine learning in its products for years, so the focus is on protecting data and proprietary algorithms as AI adoption grows.
- Model manipulation is especially critical for Medtech, as it could lead to device failures with direct patient-safety implications (e.g., incorrect diagnostics or dosing).
- Medtech companies have limited control over device security once the device is deployed, which increases risk.

Further consideration

Medtech companies tend to be highly risk averse, and they view AI-related security risks as significant to both their business and patients they serve. Data breaches can lead to regulatory intervention or product withdrawal. Even worse, a compromised device used for diagnosis or administering medicines could be lethal if it fails to detect a tumor or administers the wrong dosage.

Actions to consider in the next 30 to 90 days

Clarify approved tools and usage boundaries, and implement monitoring that reduces shadow AI exposure.

Add scenario-based training for high-risk groups such as R&D, commercial, and technology teams, focusing on realistic examples of unintended exposure.

For device and product teams, **reassess** post-deployment threat scenarios, including model manipulation pathways and monitoring expectations.

Trend 'relevance' and industry 'readiness'

Foundational cybersecurity practices are mature in many organizations because privacy, safety, and compliance are long-standing priorities. AI-specific cyber readiness is still evolving, particularly for threats such as model extraction and manipulation, and for scaling approaches where AI helps test and defend AI-enabled applications.

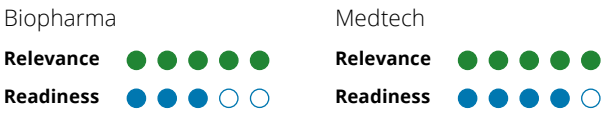

From our leaders

"You can put strong technical controls in place, but outcomes still depend on how people use the tools day to day. That is why role-based, scenario-driven training matters—tailored for researchers, IT teams, and commercial teams, and grounded in real situations such as what happens when sensitive information is entered into an unsanctioned, publicly available AI tool. In practice, the human layer and the governance layer work together. Clear data classification and usage policies help people understand what information can be used with AI systems, what needs added safeguards, and what stays off limits. When that guidance is reinforced through training, communication, and ongoing guardrails, technical controls have a better chance of scaling beyond a single team or platform and supporting responsible adoption across the organization."

***Biju Chacko**, Senior Vice President, Life Sciences, Deloitte Consulting LLP*

"IP leakage is top of mind for many Medtech clients, especially as generative and agentic capabilities become widely accessible. Medtech organizations have long embedded machine learning into products and software, so the bigger shift is not 'AI is new'—it is that the interface to AI is new. These systems can ingest and generate large volumes of content, including code, business plans, and technical formulae, and that changes the exposure profile for confidential information and trade secrets. The practical questions become, 'How do we protect the data and IP, and what differentiates our products while still capturing value?' That often points to a trustworthy AI approach that combines security and privacy controls, clear governance and accountability, and guardrails that reduce the likelihood of sensitive material being used in the wrong context or leaving approved boundaries. In parallel, it calls for clarity on IP risk pathways unique to generative AI—from what data is used in training, to how outputs may unintentionally incorporate protected content, to how organizations set policies for acceptable use."

***Puneet Sharma**, US Medtech Commercial Leader, Principal, Deloitte Consulting LLP*

Path forward

From AI momentum to measurable outcomes

Across Biopharma and Medtech, AI is shifting from pilots to production. The upside is clear, and delivery remains uneven. In life sciences, progress depends on more than model selection. It depends on operating model design, governance, validation, and trust across regulated workflows.

As we explored the five trends expected to influence life sciences, one theme stands out: Relevance is high, and readiness varies by function and by organization. That variation is not a drawback. It is a planning advantage when leaders use it to focus on the right starting points.

What it can look like in practice

Start where value and feasibility overlap: Use the relevance and readiness scores to identify one to two workflows where the business case is clear and the path to validation is workable.

Design for regulated reality: For many decisions that touch quality, patient safety, or compliance, human oversight remains part of the operating model. Treat auditability, traceability, and documented sign-off as design inputs, not downstream tasks.

Build toward scale through disciplined sequencing: Many organizations see early gains in pockets. The next step is connecting those pockets through consistent governance, shared data foundations, orchestration patterns, and measurable outcomes.

Treat security as an enabler of adoption: Foundational security and privacy practices are established across life sciences, and AI-specific risks continue to evolve. Clear guardrails, approved-tool boundaries, and scenario-based training can support progress while reducing exposure.

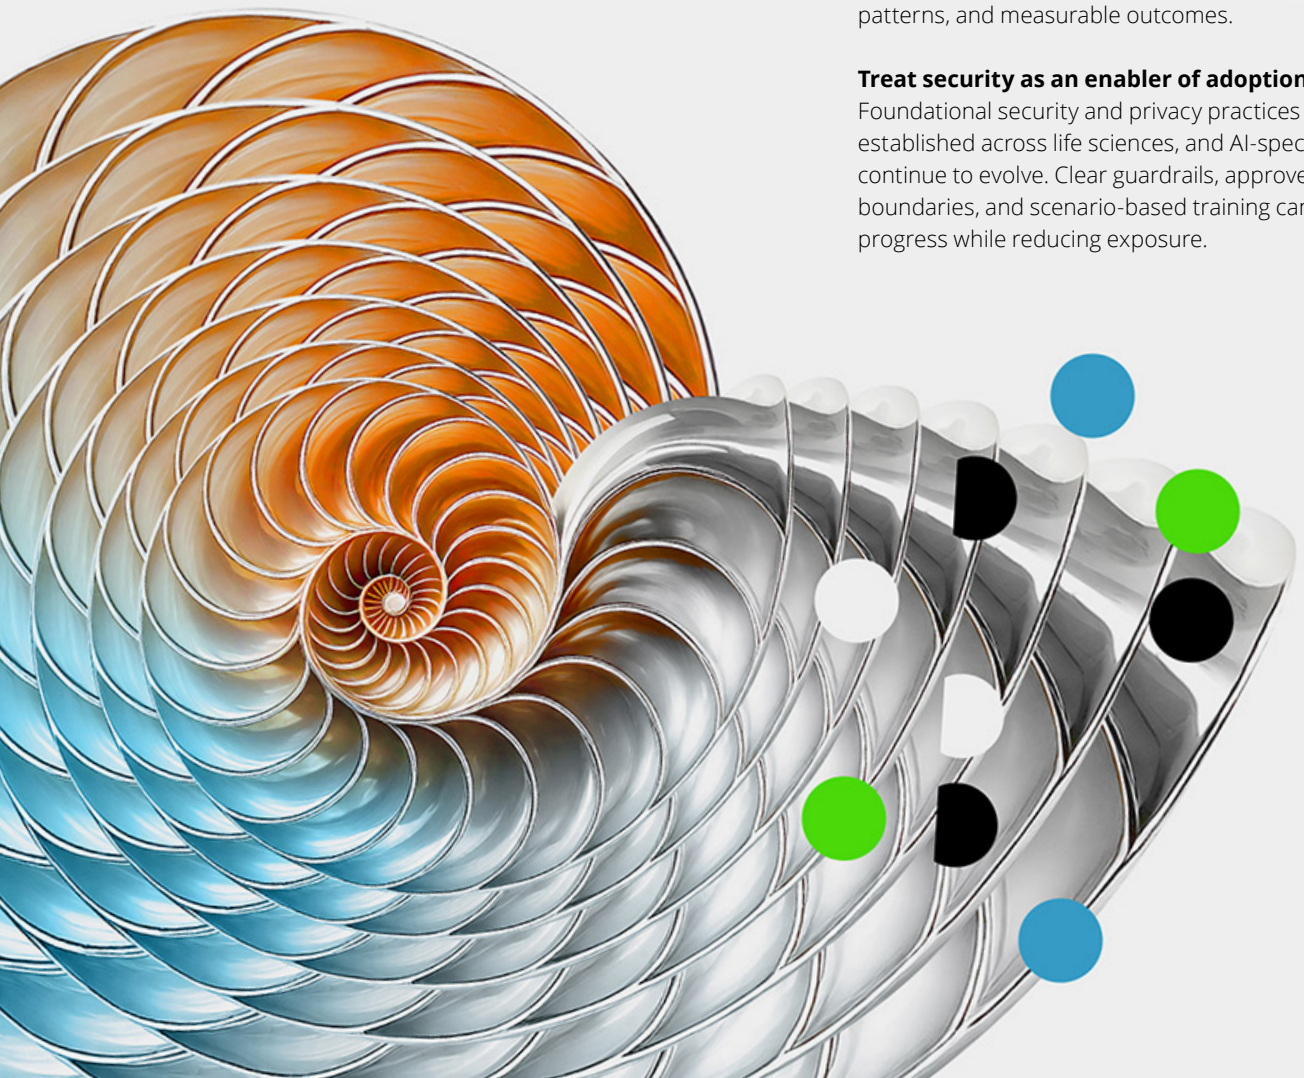

Continue the conversation

AI adoption in life sciences is now a leadership agenda. If you would like to translate these trends into your organization's priorities, we invite you to connect with us for an executive conversation or a working session.

Together, we can explore which workflows to target first, what governance and validation patterns fit your risk posture, and how to move from pilots to repeatable outcomes.

To continue the conversation, reach out to us.

Deborshi Dutt
Principal
Deloitte Consulting LLP
debdutt@deloitte.com

Luis Hakim
Principal
Deloitte Consulting LLP
lhakim@deloitte.com

Todd Konersmann
Principal
Deloitte Consulting LLP
tkonersmann@deloitte.com

Aditya Kudumala
Managing Director
Deloitte Conseil (Global)
adkudumala@deloitte.fr

Mukund Lal
Senior Manager
Deloitte Consulting LLP
muklal@deloitte.com

Nirav Mehta
Principal
Deloitte Consulting LLP
niravmehta@deloitte.com

Ram Paravasthu
Principal
Deloitte Consulting LLP
rparavasthu@deloitte.com

Subadhra Parthasarathy
Managing Director
Deloitte Consulting LLP
suparthasarathy@deloitte.com

Laks Pernenkil
Principal
Deloitte Consulting LLP
lpernenkil@deloitte.com

Puneet Sharma
Principal
Deloitte Consulting LLP
punesharma@deloitte.com

Raveen Sharma
Managing Director
Deloitte Consulting LLP
ravesharma@deloitte.com

Tarun Sharma
Principal
Deloitte Consulting LLP
tasharma@deloitte.com

Deloitte.

1365762