




Tech Trends 2026 | Deloitte Insights

European Life Sciences & Healthcare Perspective

The Life Sciences and Healthcare (LSHC) landscape is rapidly evolving, driven by technological advancements that are reshaping how companies research, develop, and deliver treatments and devices for healthcare professionals, healthcare organisations and, ultimately, patients.

This European LSHC perspective of the Tech Trends 2026 report dives into the forces shaping this transformation, exploring how technology is being implemented and its implications for the industry and region as a whole.

Relevance and Readiness scale:

We looked at each trend and assigned a value from one (low) to five (high) based on the trend's relevance and readiness to the Life Sciences and Healthcare industry.

RELEVANCE

How impactful would it be if LSHC adopted the trend?

READINESS

How ready is LSHC to adopt the trend?

Exploring the five technology trends driving innovation and transformation across Life Sciences and Healthcare.

The agentic reality check

Agentic AI is transforming the Life Sciences and Healthcare (LSHC) industry, **automating complex processes across value chains**. Adoption remains challenging due to the sector's reliance on legacy systems, regulatory considerations, and the criticality of data security and integrity.

AI goes physical

Physical AI is well-positioned to revolutionise capabilities across LSHC – providing adaptive & autonomous physical systems to **supplement human workforces as well as to improve precision and speed**. Safety concerns, large up-front investments, and organisational readiness are also key considerations to truly **realise the benefits of 'robot-enhanced' capabilities**.

The AI infrastructure reckoning

AI is a cornerstone of innovation, prompting organisations to reassess compute strategies to **balance cost, data sovereignty, and efficiency**. While the public cloud provides agility, demand for compute-intensive tasks like genomic sequencing and clinical trials is driving costs and amplifying data privacy concerns. As a result, organisations are building hybrid architectures that **leverage the strengths of each platform**.

The great rebuild


Building AI-native organisations is key to thriving in a tech-driven future, **requiring a bold reimaging of structures and processes**. The sector's regulatory complexities demand robust governance and controls for **responsible AI deployment**.

The AI dilemma

For LSHC, patient safety and data protection are paramount, and AI provides a **new duality for organisations to juggle**. AI may enhance cybersecurity capabilities, but it also expands the attack surface, enables AI-accelerated attacks, and is raising concerns around the secure and ethical use of models internally. Organisations can approach AI security strategically, implementing multiple defence layers, including AI-powered security systems.

The agentic reality check

Preparing for a silicon-augmented workforce





Tech Trends | The agentic reality check: Preparing for a silicon-augmented workforce

Enterprises are moving quickly toward agentic AI, but many are hitting a wall. They're trying to automate existing processes—tasks designed by and for human workers—without reimagining how the work should be done.

Leading organisations are discovering something different: True value comes from rethinking operations, not just layering agents onto legacy workflows. This means building agent-compatible architectures, implementing robust orchestration frameworks, and developing new management approaches for digital workers.

Agentic AI has become a strategic imperative where transitioning from labs to production is critical

Agentic AI is a top priority for nearly all enterprises right now. However, poorly planned agentic implementations are failing due to misalignment with business objectives and a reluctance to rethink processes.

Only 11 percent of organisations have agents in production, despite 38 percent piloting them.⁴

Organisations achieving success are redesigning processes and scaling from pilot to production

Leading organisations are redesigning processes, managing agents as a silicon augmented workforce, orchestrating agents across the enterprise, and reevaluating the role of legacy systems. The businesses that can successfully promote agents from pilot to production are the ones starting to see returns on their investment.


True agentic success requires human-silicon synergies as agents gain further autonomy

As agents progress along the autonomy spectrum from augmentation to automation to true autonomy (which will likely require artificial general intelligence, or AGI), successful businesses will figure out how to drive synergies between human and silicon colleagues. What businesses turn over to agents and how they use their human workforce will be a key differentiator.

"Many people are busy trying to find better ways of doing things that should not have to be done at all. There is no progress in merely finding a better way to do a useless thing"

- Henry Ford, on building automobiles in 1922

Projected agentic AI adoption






The Life Sciences and Healthcare perspective | How is the industry exploring agentic AI?

Nearly a third of Life Science and Healthcare organisations recognise agentic AI as having a significant impact on their organisational strategies^[2, 3]. However, success rests on more than technology alone. A highly regulated industry, a complex and interconnected operational landscape – often reliant on legacy systems – and the emphasis on patient safety and data integrity, present a myriad of challenges for technologists to navigate.

No. of life sciences executives that are prioritising gen AI and agentic AI for cost optimisation in 2026.²



"The HR organisation does workforce planning really well, and the IT function does technology planning really well.

We need to think about work planning, regardless of if it's a person or a technology"

- Tracey Franklin, chief people and digital technology officer at Moderna.⁵

Overcoming organisational and cultural hurdles

The primary hurdle isn't the technology, but the human element. LSHC organisations are focused on challenging ingrained processes and fostering a new mindset. This involves cultivating AI innovators who can envision future possibilities and encouraging leadership to embrace a pioneering spirit over a purely cautious one. Ultimately, successful agentic AI adoption hinges on a cultural shift as much as a technological one.

Building robust knowledge foundations

For agentic AI to be effective, it needs a sophisticated "brain". The industry is prioritising the creation of comprehensive, well-structured knowledge stores. This means meticulously curating formal documents, business glossaries, key business questions (KBQs), ontologies, and taxonomies. The goal is to build a rich context fabric where all information is blended, preventing AI hallucinations and ensuring its actions are grounded, compliant, and reflect deep domain expertise.

Strategic deployment for tangible value

Despite initial hesitation, LSHC is strategically applying agentic AI where it can deliver clear benefits. This involves identifying high-impact use cases across R&D, Quality Assurance, Manufacturing & Supply Chain, Commercial, and Enabling Functions, focusing on solutions with measurable value, such as improved efficiency or decision-making. Recognising that trust is key, the approach often involves phased deployment, allowing the technology to prove its worth and build confidence as it matures.

30 percent of Life Science leaders cite significant impacts to their organisation's strategies through agentic AI whereas only 22 percent of Life Sciences leaders say they have successfully scaled AI.²



Real world stories & highlights

Leading organisations across Life Sciences and Healthcare are already transforming the vision of agentic AI into tangible value. The following case studies highlight some of the big bets that organisations are finding success in, allowing agentic AI to redefine efficiency, compliance, and advantage across their domains.



Agents automating R&D content supply chains across pharma

Leaders across pharma are leveraging agentic AI to automate R&D content supply chains. Agents are autonomously reading and synthesising vast clinical study datasets and complex R&D documents with the goal to significantly accelerate time-to-market for new therapies by streamlining analysis and document creation and ensuring accuracy and quality.



Self-optimising manufacturing for pharma and biotech

Pharma and Biotech are implementing agentic AI for autonomous manufacturing. Agentic systems monitor manufacturing lines, predicting issues, and adjusting in real time to maximise yield and minimise waste. These self-regulating and self-optimising environments are enhancing efficiency, reducing costs, and ensuring consistent product quality.



Organisational insights and intelligence, "Talk To Your Data"

Organisations are establishing their AI foundations and enabling 'data management as a service' (DMAaaS) capabilities. Agentic AI simplifies access, management, and analysis of disparate datasets and allows users to query and receive insights.

RELEVANCE

How impactful would it be if LSHC adopted the trend?

5 / 5

Agentic AI is highly relevant to the LSHC industry, offering substantial value potential and numerous transformative use cases

READINESS

How ready is LSHC to adopt the trend?

3 / 5

Although readiness is low, organisations are actively introducing an innovation culture, technical upskilling, and the requisite AI foundations for success

Hear more from our Deloitte subject matter experts:



SEBASTIEN BURNETT
Partner
sburnett@deloitte.co.uk



ELLIOT STAMP
Director
estamp@deloitte.co.uk



PRIYA ARORA
Director
priyaxarora@deloitte.co.uk

AI goes physical

Navigating the convergence of AI and robotics





Tech Trends | 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. This signifies that intelligence is no longer confined to screens but is embodied in physical forms, solving real-world problems. These capabilities are appearing in industrial robots, autonomous vehicles, drones, and other systems

AI is moving from digital environments to the real world

Physical AI is moving beyond screens, allowing machines to sense, understand, and act in the real world. This is driven by advancements in vision-language-action models enabling robots to interpret their surroundings and make quick decisions on their own. Additional improvements in power, combined with component commoditisation and advanced training methods, make them more capable.

From traditional robots to humanoids

While traditional robots excel at precise, repetitive tasks like computer numerical control (CNC) machining, they struggle with unstructured environments, assembly, and manipulation. However, the intersection of agentic AI with physical robotics is set to transform humanoids, enabling them to adapt to new environments, plan multi-step tasks, recover from failure, and operate under uncertainty. This allows them to function as intuitive assistants in warehouses, homes, and healthcare.

By 2035, there will be 2 million humanoids in the workplace, with expectations to increase to 300 million by 2050.¹

Breaking through implementation barriers

Despite its potential, scaling physical AI involves complex, inter-related challenges spanning technical, operational, and regulatory domains. These include gaps in training, safety concerns, cybersecurity risks, and the persistent gap between simulated and real-world performance caused by approximated physical models. Organisations that proactively tackle these challenges head-on will define the next wave of deployment for physical AI.

“The market for humanoids is going to be twice the size of the automotive industry in 25 years”

- Jonathan Hurst, interview with Deloitte, Oct. 6, 2025.¹

How vision-language-action models work

Vision

Computer vision systems interpret visual environments and identify objects, obstacles and spatial relationships

Language

Natural language processing understands human commands and can communicate intentions


Action

Motor control systems execute physical tasks based on visual inputs and linguistic instructions



The Life Sciences and Healthcare perspective | How is the industry exploring physical AI?

Physical AI represents a fundamental shift for LSHC, with potential to address persistent labour shortages, enhance operational resilience in complex manufacturing and clinical environments, and enable autonomous decision-making across its diverse production landscape. While some areas, like vaccine manufacturing, already leverage significant automation, physical AI has the potential to help scale precision medicine manufacturing, deploy surgical robots to enhance procedural accuracy while capturing data to train future systems, and create smart supply chains that autonomously optimise for patient safety and regulatory compliance.



LSHC manufacturing is still dependent on traditional robots

Pharmaceutical manufacturing is predominantly process-driven and highly automated. Traditional automation handles repetitive tasks and supports high-volume, high-speed production. While physical AI (for example, autonomous robots and robotic-assisted systems) offers a clear opportunity, broad adoption is likely still several years away. Current barriers include health and safety regulations, limitations in battery life and lifting capacity, and cost constraints. Adoption will also require organisational change, including new roles, skills and governance, to operate, train and maintain these systems over time.



Transformative potential in MedTech

MedTech, despite historically lagging pharma, is on the cusp of significant physical AI integration, particularly given its high-value, low-volume, and precision-intensive device manufacturing. While direct robot training for fewer devices may not always justify investment, the exponential increase in computing power is enabling advanced in-device diagnostics and computer-aided imaging. Physical AI's impact extends to smart warehousing for device parts, digital twins for product design and scenario planning, and AI-driven surgical support systems.



Value in smart warehousing & supply chain

While agentic AI shows promise for optimising warehouse slot allocation, its physical AI counterpart faces slower adoption in LSHC supply chains due to heightened safety concerns and the high risk of damage to fragile, temperature-controlled products. The current focus is on strategic network optimisation, driven by geopolitical factors and the need for resilient supply chains. Digital twins are already supporting network strategy, and agentic AI is beginning its journey in clinical supply chains, primarily in software-led planning and lab automation. For LSHC supply chains, a more widespread physical AI transformation will only materialise once there is a compelling business case that justifies the significant capital expenditure for new, AI-optimised facilities.



Real world stories & highlights

LSHC organisations are exploring the use of AI-powered systems to drive strategic optimised network strategies in supply chain and improve patient outcomes in healthcare use cases.



AI-powered robotics to improve patient outcomes and support continuous learning for surgeons¹⁰

A UK-based developer of next-generation surgical robots is integrating Nvidia's IGX Thor AI platform into its surgical robotic system. This collaboration enhances the robot's computing power, energy efficiency, and situational awareness, enabling real-time insights and recommendations during surgery. By leveraging AI, the company seeks to improve patient outcomes and continuously learn from every procedure, empowering surgeons to perform more precise, minimally invasive operations.



AI-driven simulation for next-gen surgical robots¹¹

Johnson & Johnson MedTech is leveraging AI-driven simulation in its Monarch surgical robotic platform for urology. This technology creates virtual operating room environments, assisting clinical teams with pre-procedure setup and modelling kidney stone procedures using simulated patient anatomy for enhanced learning and planning. This will create faster patient pathways to treatment while enhancing learning for clinical teams.

RELEVANCE

How impactful would it be if LSHC adopted the trend?

4 / 5

Physical AI is considered highly relevant for the LSHC industry, given the potential for autonomous decision making

READINESS

How ready is LSHC to adopt the trend?

2 / 5

Readiness is generally low with organisations yet to justify the investments required for AI-optimised facilities

Hear more from our Deloitte subject matter experts:



ANGELA BOWDEN

Partner
acbowden@deloitte.co.uk



ADITYA KUDUMALA

Partner
adkudumala@deloitte.fr



ANDREW FLOCKHART

Director
aflockhart@deloitte.co.uk



HITESH PARMAR

Associate Director
haparmar@deloitte.co.uk



LUCIANO LO TITO

Associate Director
llotito@deloitte.co.uk




NISHANT SINHA

Senior Director
nishsinha@deloitte.ch

The AI infrastructure reckoning

Optimising compute strategy in the age of inference economics





Tech Trends | The AI infrastructure reckoning: Optimising compute strategy in the age of inference economics

When generative artificial intelligence exploded on the scene, businesses got busy dreaming up next-generation products and services. Today, AI has grown up. But as it moves from proof of concept to production-scale deployment, enterprises are discovering their existing infrastructure strategies aren't designed for AI's demands.

The token economics wake-up call

AI costs are forcing infrastructure recalculation as the cloud-first strategies that defined the last decade are hitting economic and operational walls.

Token costs have dropped 280-fold in two years; yet some enterprises are seeing monthly bills in the tens of millions.¹

AI-optimised data centres are a new consideration in hybrid infrastructure strategies

Purpose-built AI infrastructure ecosystems that combine the best of cloud, on-premise, and edge allow enterprises to take back control over their compute infrastructure to enhance capabilities and manage costs.

Data centres and compute drive toward new frontiers

The maturation of emerging compute paradigms, such as optical, neuromorphic, and quantum computing, will force data centres to continue to evolve as enterprises embrace increasingly specialised tools for specialised tasks.

A decision framework for computation workload placement

Public cloud hyperscalers

Choose when:

- there are variable or unpredictable AI workloads that benefit from elastic scaling
- experimentation and rapid prototyping phases prioritise speed to market
- access to the latest AI services and managed platforms provides competitive advantage
- regulatory flexibility exists regarding data residency requirements

On-premises infrastructure

Build when:

- consistent, high-volume AI processing justifies capital investment
- sensitive data requires air-gapped processing behind corporate firewalls
- latency-critical applications demand immediate response times
- mission-critical tasks cannot depend on external cloud connectivity
- long-term strategic AI capabilities warrant deployment commitments of three years or more
- industries face strict regulatory requirements

Edge infrastructure

Deploy when:

- real-time decision-making applications demand immediate processing
- bandwidth-constrained environments make data transmission expensive or impractical
- local data regulations prevent information from crossing geographic boundaries
- offline resilience needs exist where connectivity cannot be guaranteed



The Life Sciences and Healthcare perspective | How is the industry exploring AI infrastructure and compute strategies?

In LSHC, AI infrastructure already plays a crucial and growing role. Organisations are rethinking where and how they deploy their AI workload, considering rising cost, data sovereignty, latency sensitivity, resilience needs, and IP protections.

Three-tier hybrid architectures are being considered to navigate this complex decision tree, to apply the best option for a given AI workflow: cloud for elasticity, on-premises for consistency, edge for immediacy.

Number of life sciences executives that are prioritising investment in AI-enabled platforms for growth in 2026.²



Efficient compute economics



For LSHC, AI workloads are inherently compute-intensive and require large-scale datasets for genomic sequences, clinical trial results, and imaging diagnostics. The escalating costs of these specialised GPU workloads is becoming apparent. Organisations are evaluating their compute spend and are finding that for stable, high-volume AI applications, such as 24x7 simulations for drug candidates or continuous analysis of patient data, AI-specialised on-premise and co-location data centres could offer long-term advantages, such as lower costs, improved service stability, and options for data sovereignty.

Privacy and an increase in data sovereignty



Patient data, proprietary research, and intellectual property are critical, and new evolutions in AI have only bolstered existing concerns. As AI models increasingly process sensitive information, the question of where that data resides is paramount. For organisations with data sovereignty and privacy at their core, taking control of their AI infrastructure has become a strategic imperative and in some cases is being prioritised above token cost concerns.

Value and longevity of AI use cases



Another key consideration sits directly with the expected lifecycle and longevity of AI applications and their corresponding hardware and model requirements. Exploratory research that leverages AI may need to leverage the latest and greatest models, where the agility of the public cloud and specialised AI Infrastructure providers may remain invaluable even with higher inference costs.

When cloud costs begin to exceed 60 percent to 70 percent of the cost of acquiring equivalent on-premise systems, on-premise deployments start to become more attractive.⁶


Outside of AI, edge computing is also anticipated to become more prominent; offering more powerful compute closer to sources and lowering latency while increasing processing potential. Overall, this enables new use cases for edge devices in Healthcare and can act as a service differentiator for Life Sciences.

The strategic decisions made across cloud, on-premise, and edge today represent critical investments that could position organisations for future computing paradigms. By maturing these capabilities now, organisations can prepare to harness quantum, neuromorphic, or optical computing on the horizon.



Real world stories & highlights

The decision to leverage public cloud, build your own, or rent specialised AI infrastructure should be based on specific use cases, the expected longevity of the AI model and hardware, and thorough options analysis to compare not just token and inference costs, but the costs of infrastructure setup, maintenance, and staffing.



Leading pharmaceutical deploys world's largest AI factory¹²

In 2025, a leading global pharmaceutical built their largest AI factory. By establishing in-house AI foundations, they are able to train large-scale biomed models, enhance lab workflows, optimise digital twin and robotic capabilities, and leverage agents to generate new molecules and design treatments.

Overall, this enables the compression of discovery timelines and accelerates breakthroughs in genomics, personalised medicine, and molecular design.



Cambridge-1, the UK's most powerful AI research supercomputer¹³

Cambridge-1, the Nvidia-backed supercomputer, is accelerating UK healthcare innovation. Powered by 80 Nvidia DGX A100 systems, it delivers record-breaking AI performance to support genome sequencing, drug discovery, and disease research for its founding partners and research groups.



Exploring quantum computers to enhance molecular modelling¹⁴

Leading pharma, in collaboration with quantum specialists, have been exploring novel computing paradigms to improve the accuracy of molecular modelling. By leveraging quantum circuits, they successfully modelled proton-electron dynamics, enabling more precise characterisation of molecular interactions.

RELEVANCE

How impactful would it be if LSHC adopted the trend?

4 / 5

AI Infrastructure is considered highly relevant for the LSHC industry, given the potential for faster insights and innovation.

READINESS

How ready is LSHC to adopt the trend?

2 / 5

Readiness is generally low with many organisations lacking the in-house skills and capabilities to effectively manage complex hybrid AI infrastructure at scale.

Hear more from our Deloitte subject matter experts:



JAMES RODEN

Director

jroden@deloitte.co.uk



FELIX KOEBELE

Director

fkoebele@deloitte.ch



NISHANT SINHA

Senior Director

nishsinha@deloitte.ch




SARAH MABER

Senior Manager

smaber@deloitte.ch

The great rebuild

Architecting an AI-native tech organisation





Tech Trends | The great rebuild: Architecting an AI-native tech organisation

IT's role is rapidly transforming from simply maintaining systems to strategically driving innovation, largely thanks to AI. This transformation extends beyond tools and headcount, re-engineering how teams are structured, governed, and led to create leaner, faster, and AI-infused models that continuously learn and optimise. This shift positions IT leaders as crucial AI evangelists, responsible for helping teams to see the possibilities of AI while guiding ethical adoption of the technology.

AI is reshaping the tech function

AI is reshaping the tech function, with organisations planning to increase AI investments, recognising its substantial value. This shift is evident in evolving priorities, as CIOs now focus heavily on harnessing AI, data, and analytics, moving AI from an experimental phase to a core strategic imperative.

64 percent of organisations plan to increase AI investments over the next two years.⁷

Strategies to prepare for an AI-powered future

Organisations are actively preparing for an AI-driven future by designing modular and observable architectures, moving away from patching legacy systems to build future-ready enterprises. A core strategy involves placing human-machine collaboration at the heart of tech talent approaches, fostering new roles and upskilling efforts to bridge knowledge gaps. Furthermore, governance models are evolving to enable speed while effectively managing emerging AI-related risks. This preparation is driven by bold ambitions that prioritise reimagination over incremental change, ensuring modernisation efforts are always anchored in solving fundamental business problems rather than merely upgrading technology.



"AI is creating *autonomous* teams, where you don't necessarily need deep functional expertise in every area because AI can help bridge or fill those gaps."

- Gene Kim on AI-powered coding in enterprise IT¹

The markers of an AI-powered tech organisation

AI-powered tech organisations embed AI as a core collaborator across all layers, from decision-making to product development, leveraging cloud-native, platform-powered foundations for speed and flexibility. They are transitioning to lean, cross-functional product squads and fostering human-agent teams at scale, blending human ingenuity with machine intelligence. These organisations also act as ecosystem orchestrators, collaborating across startups, hyperscalers and regulators to accelerate innovation, embracing perpetual evolution with an "always beta" mindset.

78 percent of tech leaders anticipate broad, targeted, or transformational integration of AI agents into architecture workflows over the next five years.⁸



The Life Sciences and Healthcare perspective | How is the industry architecting an AI-native organisation?

LSHC organisations are fundamentally re-architecting their technology functions to become AI-native, driven by a clear business mandate to close the massive gap between growth targets and operating costs, alongside evolving CIO leadership, all while navigating the sector's distinct regulatory and operational complexities.



Strategy and portfolio management: Closing the business gap

In LSHC, the push towards an AI-native tech organisation is fundamentally driven by a clear business mandate to close the massive gap between growth targets, product launches, and operating costs. Technology spend is one of the biggest levers identified to achieve this, with a focus on maximising ROI from technology investments and leveraging AI for increased productivity.

CIOs are proactively shaping IT strategy to support and scale diverse AI use cases, viewing AI not as a mere trendy tool, but as a clear business imperative. This necessitates IT portfolios with AI-native investments, and not AI as an afterthought, balancing innovation and speed with cost optimisation and a strong focus on ROI.



Technology operating model: Organisational maturity & mindset shift

The LSHC technology operating model is transforming towards a skills-based structure, demanding upskilling in AI-related competencies. Human-machine collaboration is central, fostering seamless integration of human expertise with AI agents for efficiency. This demands a mindset shift and upskilling and reskilling of both business and IT professionals in AI-related competencies. Teams are now becoming leaner in anticipation of the capacity increase from AI agents. Given the highly regulated nature of LSHC, robust governance and ethical frameworks are critical to ensure responsible and secure AI deployment, with a strong human oversight loop.




The enabling architecture: Building future-ready foundations

Architecting an AI-native tech organisation in LSHC requires a flexible, hybrid technology landscape. This involves transitioning from inadequate legacy IT to dedicated, AI-optimised infrastructure and a redesigned tech stack. Crucially, you can't just bolt AI onto things that don't work. Process re-engineering to understand data flows, process maturity, and human interfaces are needed. LSHC organisations must focus on how work gets done, faster, with AI, acknowledging AI is not a silver bullet for everything. This includes investing in the future of engineering and AIOps of the future, alongside application rationalisation to build a sustainable baseline for AI, especially as CIOs recognise that scalability is an issue.

Real world stories & highlights

Some LSHC organisations are looking to develop digital-native operating models where AI is being embedded throughout the value chain, spanning strategy, operational processes, and product lifecycle management. An example outcome of this integration is the accelerated software development lifecycle.



Service desk transformation through agentic AI

A global biopharma embedded agentic AI into its contact centre to manage high-volume queries, reducing manual intervention by 30% in the first year, with a roadmap to reach 40% by the next year. The initiative significantly cut operational costs, improved response times, and established a scalable, intelligent platform for continuous enhancement, moving towards autonomous operations and self-healing capabilities.

AI-assisted engineering for Lab of the Future

A global pharma is building an AI innovation lab (Lab of the Future) to accelerate drug discovery and innovation. The lab will co-locate domain scientists and AI capabilities to drive its R&D operations. To support the "Lab of the Future", the organisation is reimagining its IT operations and overall technology operating model, infusing AI, product model and agile delivery in engineering.

RELEVANCE

How impactful would it be if LSHC adopted the trend?

4 / 5

AI-native tech organisations are highly relevant for the LSHC industry, with significant potential for value creation

READINESS

How ready is LSHC to adopt the trend?

2 / 5

Readiness is high for commercial and R&D use cases but mostly in pilot stage for enterprise tech use cases

Hear more from our Deloitte subject matter experts:



NIYATI NAGAR

Director

nnagar@deloitte.co.uk



AMIT KUMAR

Associate Director

amitykumar@deloitte.co.uk



ANTONY DAMEN

Director

adamen@deloitte.ch




GURNEESH CHEEMA

Director

gccheema@deloitte.ch

The AI dilemma

Securing and leveraging AI for cyber defence





Tech Trends | The AI dilemma: Securing and leveraging AI for cyber defence

Organisations that deploy artificial intelligence at scale are discovering its paradox: The same AI capabilities that help them be more competitive can also introduce new security risks. And to add to the paradox, they're also recognising that AI offers powerful new capabilities to counter the very vulnerabilities it creates.

Enterprises face multiple threats related to artificial intelligence, including shadow AI deployments, AI-accelerated attacks, and the intrinsic risks of AI systems. Yet even as AI drives new threat vectors, traditional cybersecurity principles remain constant and should be applied to autonomous systems that learn, adapt, and operate at machine speed. Many of these techniques will require significant adaptation because most cyber organisations were not designed to lean on digital intelligence.

"Risks and capabilities tend to go along with each other quite closely. The more the system can do, the more we give it access and agency, and the more we have new kinds of security surfaces to cover"

- Sanmi Koyejo, assistant professor at Stanford University.⁹

AI capabilities introduce novel attacks

The rapid rollout of AI solutions across the enterprise is creating complex new challenges for security teams: familiar risks manifesting through novel vectors both from the enterprise's use of AI tools and external bad actors attacking these systems.

Capable personal AI assistants require unprecedented access to personal data, but once personal data is incorporated into AI models, the right to erasure becomes nearly impossible.

Safeguards to both combat and utilise AI

A multifaceted risk mitigation strategy is surfacing to address these complex threats. Next-gen safeguards, leveraging AI defence systems, and traditional methods that continue to mature are required.

The physical world expands the attack vector

As AI proliferates into every physical system (e.g. power grids, supply chains), the physical world will become an attack vector, with implications for both corporate and national cybersecurity.

AI-related risks and associated mitigation strategies¹

Risks	Mitigations
Data security - Confidentiality & data privacy - Training data position - Model skewing	- Secure data practices - Data-integrity monitoring - Robust access controls
AI models - Model collapse - Model stealing - Model inversion - Excessive agency abuse	- Model isolation - Privileged access management
AI applications - Ethical use concerns - Input injection - Unauthorised access	- Network and access mgt & controls 3rd party evaluations
AI infrastructure - Model denial of service - Deployment misconfigurations - Lateral movement attacks	- Secure deployment in virtual networks - Perimeter and workload hardening



The Life Sciences and Healthcare perspective | How is the industry adjusting to new AI-enabled attack vectors?

Healthcare organisations predict 14 percent of tech budgets will be spent on cybersecurity tools.³

LSHC organisations are navigating the evolving cyber security landscape being shaped by AI. Given the already regulated nature of the industry and the sensitivity of patient and health data, intellectual property, and clinical processes, mitigating AI risks is paramount.

Both amplified threats and novel defence tools, alongside significant internal governance challenges, will shape how each organisation responds to this changing landscape.

An amplified threat landscape

AI has expanded the attack surface, providing malicious actors with the tools to design more effective exploits, identify and probe weaknesses, and increase the frequency and scale of attacks.

While AI is being used by malicious actors, defensive utility is also evolving. Many key AI features are being integrated across cybersecurity product suites, enhancing detect and response capabilities.

The increasing maturity and capability of GenAI and agentic tools promise significant productivity gains for security teams by:

Analysing code intent and contextualising vulnerabilities in code

Translating threat and intelligence reports into actionable profiles and response scripts

AI-enriched SOC's to keep pace with threat actors using evolving AI attacks

Support incident and triage by consolidating semantic data

Deploying the most effective tools – from traditional security to agentic AI – as well as upskilling the workforce and improving organisational literacy on threats and risks, is key to evolving the requisite capabilities and being able to respond.

The risk within

The widespread adoption of AI products and potential of agentic implementations introduce significant technology governance and security design challenges. The very principles that maximise AI's value, extensive data context and controls, often conflict with the security principles safeguarding the enterprise.

LSHC organisations that guardrail AI development and securely segment, catalogue, and observe their AI products and processes will be best positioned to harness AI's benefits responsibly.



Real world stories & highlights

The LSHC industry will continue to balance innovation with risk management to protect reputation and patient safety. Cybersecurity strategies are evolving, driven by the need to reinforce traditional, foundational capabilities while adapting to new, AI-driven threats and opportunities. The following real-world story, highlights real-world experience seen in the face of a new AI paradigm.



Ethical AI elevates cyber protection for global pharma

A global pharmaceutical organisation have used AI to bolster its security capabilities. They are applying risk-based measures, leveraging AI to protect the confidentiality, integrity, and availability of critical data and services.

AI-driven analytics to enhance System and Organisation Control (SOC) type 2 reports, streamlines SOX compliance, and supports proactive identification of cyber threats, all while committing to ethical and responsible use of AI.

RELEVANCE

How impactful would it be if LSHC adopted the trend?

4 / 5

Although the adoption of AI tools is low, the risk posed by evolving and escalating AI threats is significant. For cyber & trust capabilities, being able to respond is paramount where the risk to data and services is real, and the solution for many is still unclear.

READINESS

How ready is LSHC to adopt the trend?

3 / 5

Early AI adoption offers foundational guardrails and experience, yet significant growth and maturity in these capabilities is yet to be realised. Responding to the business's adoption of AI and filling talent and skills gaps are areas influencing organisational readiness.

Hear more from our Deloitte subject matter experts:



MARK BETHELL

Partner

mabethell@deloitte.co.uk



DAN SADLER

Director

dpsadler@deloitte.co.uk



KISHWAR CHISHTY

Partner

kchishty@deloitte.ch

Learn more

Industry contacts

For questions regarding the Tech Trends 2026 Life Sciences and Healthcare perspective, please contact:



JOHN FORD

Partner

joford@deloitte.co.uk



HANNAH HAYWARD

Partner

hhayward@deloitte.ch



PETER WALLACE

Partner

pwallace@deloitte.co.uk



KISHWAR CHISHTY

Partner

kchishty@deloitte.ch



RICHARD BADERMAN

Partner

rbaderman@deloitte.co.uk



GURNEESH CHEEMA

Director

gccheema@deloitte.ch



AISLING TEO

Director

aiteo@deloitte.co.uk



NIYATI NAGAR

Director

nnagar@deloitte.co.uk



RAHUL PARANJAPE

Associate Director

rparanjape@deloitte.co.uk

References

1. 2026 Tech Trends
2. 2026 Life Sciences Outlook
3. 2026 Global Health Care Outlook
4. 2025 Deloitte Emerging Technology Trends in the Enterprise Survey, publication in process.
5. Tracey Franklin, Moderna, interview with Deloitte, Sept. 26, 2025.
6. Thomas, Tayal, Stewart, Kearns-Manolatos, and Parveen, "Is your organization's infrastructure ready for the new hybrid cloud?"
7. Jagjeet Gill, Vibhu Kapoor, Matthew Nehls, and Shivang Aggarwal, "Fueling Growth, Not Maintenance: How Tech Budgets are Evolving,"
8. Deloitte 2025 Horizon Architecture Survey.
9. Sanmi Koyejo (assistant professor, Stanford University), interview with Deloitte, Sept. 26, 2025.
10. Victoria Woodward, "CMR Surgical Adopts NVIDIA IGX THOR for AI-Powered, next Generation Surgical Robotics - Cambridge Innovation Capital" Cambridge Innovation Capital. Oct. 2025
11. J&J MedTech, "Johnson & Johnson to Advance Robotics Development with NVIDIA Isaac for Healthcare". Oct. 2025
12. Eli Lilly and Company, "Lilly Partners with NVIDIA to Build the Industry's Most Powerful AI Supercomputer, Supercharging Medicine Discovery and Delivery for Patients", 2024.
13. Nvidia, "Get the First Look of NVIDIA's New 'Cambridge-1' AI Supercomputer", 2026
14. Hartree Centre, "Using Quantum Computing to Improve Molecular Modelling in Drug Development - Hartree Centre." Aug. 2024.

Contributors

Michael Cullen
Olalekan Adejayan
Shivani Sharma
Ilaha Akhbar
Francess Ejeagwu
Lenka Svobodova
Tim Decuypere
Alexandra Grba



This article contains general information only and Deloitte is not, by means of this article, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This article is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this article.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.