



2025

AI in Healthcare Decision Brief

Part 1

Insights on navigating the AI platform
shift in healthcare and life sciences



Table of contents

Introduction

The future of work in healthcare and life sciences is here

Section 1

The generative AI platform shift

Leadership perspectives
GenAI's impact in healthcare
by Joe Petro

Section 2

Addressing challenges to generative AI adoption in healthcare

Leadership perspectives
Staying ahead of emerging challenges and threats with AI and security for AI
by Ann Johnson

Section 3

Navigating the AI platform shift

Leadership perspectives
Navigating the future of healthcare together
by Kees Hertogh

Section 4

The winning formula

Leadership perspectives
The role of partners and startups to advance innovation
by Sally Frank

Section 5

Trustworthy AI

Leadership perspectives
Building trust to operationalize responsible AI in healthcare
by Dr. David Rhew

Conclusion

We're all in this together

Resources





Introduction

The future of work in healthcare and life sciences is here

In the last two years, generative AI has experienced explosive growth, creating exciting new ways for healthcare and life sciences organizations to increase productivity, streamline processes, drive innovation, and more. From accelerating the discovery and development of novel therapeutics to improving diagnostics at the bedside, generative AI is impacting every aspect of healthcare—and its prevalence will only grow.

Microsoft has been at the forefront of the AI revolution, building the infrastructure that makes generative AI possible, launching powerful new applications, and investing in cutting-edge research.

While supporting leading innovators and working with healthcare and life sciences organizations to make their generative AI initiatives more impactful, we've studied the strategies that lead to real and sustainable results.

In this Decision Brief, we explore how generative AI is transforming healthcare today and what's possible in the future. We discuss how the industry is embracing AI responsibly, and we share strategies for overcoming challenges and maximizing your organization's success.

Welcome to the future. It's going to inspire us all to achieve more.






Section 1 | The generative AI platform shift

AI is more than just an exciting innovation. It's a rare breakthrough technology driving a platform shift that will fundamentally reshape healthcare. While every person and organization can benefit from this shift, those who envision new possibilities and understand AI's full potential have an opportunity to achieve unprecedented outcomes.

Generative AI in healthcare and life sciences

Generative AI has been adopted with unprecedented speed. While it took seven years for the internet to reach 100 million users, ChatGPT achieved that benchmark in two months.¹ Although generative AI is relatively new to the market, adoption is rapidly expanding. Usage among enterprises jumped to 75% in 2024 from 55% in 2023.²

Today, 57% of life sciences organizations and 45% of healthcare organizations identify generative AI as the most important technology to adopt, and 79% are currently using AI in some form.³ As healthcare and life sciences organizations embrace generative AI, they tend to first prioritize use cases that enhance productivity, then expand to applications that streamline broader business functions. As adoption gains steam, we expect a growing number of use cases to emerge that impact the whole industry—accelerating scientific progress, increasing payor efficiency, improving clinical services, and enhancing the patient experience.



Generative AI usage jumped from 55% in 2023 to

75% in 2024.

The historic impact of generative AI

Just as the pace of its adoption has been historic, the impact of generative AI in healthcare and life sciences is likely to be unlike anything we've ever seen.

Technology platform shifts since the 1960s have been defined by Moore's Law, under which computing power doubled roughly every two years with the number of transistors in an integrated circuit. In contrast, AI ushers in a new dynamic of progress governed by scaling laws. Now, a model can be improved by scaling up any of its key resources, such as the model size, the dataset being used, and the computational power available.

Under scaling laws, AI capabilities are doubling roughly every six months. As research labs, scientific institutions, insurance companies, and providers adopt generative AI, the impact of this accelerated pace will be felt everywhere.

Takeaway

Generative AI is transforming the laborious processes of research, operations, and employee and patient experiences, opening up extraordinary possibilities.

Realizing value from generative AI

As healthcare and life sciences organizations adopt generative AI, they're realizing value in numerous ways, such as improving time to market for novel therapeutics, reducing missed findings, and creating better experiences for clinicians and patients.

By automating time-consuming processes and focusing staff time on more meaningful work, organizations are also generating significant returns on their investments. As a result, they can increase funding to areas that serve their core mission: improving individuals' health and quality of life.

<8 months
Time for AI deployments.

13 months
Time to ROI for generative AI.

\$3.50
AI

\$3.70
Generative AI

ROI for every \$1 a company invested in 2024.

GO DEEPER To learn more about IDC's research on generative AI adoption, business value, and ROI, read the full study, [The Business Value of AI](#).

Takeaway


With deployments taking less time and organizations realizing value faster, generative AI is predicted to add \$10 trillion to the global GDP over the next ten years.⁴

1. Edwards, Benj, "ChatGPT sets record for fastest-growing user base in history, report says: Intense demand for AI chatbot breaks records and inspires new \$20/mo subscription plan," (Ars Technica, Feb. 1, 2023), 2. IDC InfoBrief: sponsored by Microsoft, 2024 Business Opportunity of AI, IDC# US52699124, November 2024. 3. McKinsey, Market perspective: AI and GenAI in Life Sciences., 4. Ibid.

Leadership perspectives

GenAI's impact in healthcare

By Joe Petro, Corporate Vice President, Healthcare and Life Sciences Solutions and Platforms



More than a decade ago, the cover of the *Journal of the American Medical Association* showed a picture that a little girl drew for her doctor. It showed her, her mother, and baby sister in the treatment room. The doctor's back was turned away from them as he sat at his desk typing on his computer. Everyone in the picture is smiling. Even the doctor clicking away at his laptop. It was titled "The Cost of Technology."

For many of us who worked in healthcare technology at the time, that image will be forever seared into our minds. Research and studies had already begun documenting the side effects of the digitalization of healthcare, and the honesty of that drawing was a wake-up call. It was a stark reminder that in our efforts to do the right thing—in this case, bring about a much-needed digital transformation in healthcare—there were very real and very important unintended consequences.

But the introduction of generative AI has revolutionized the industry and fundamentally changed the way care is delivered, from enhancing patient experiences and supporting better outcomes across the healthcare and life science ecosystem to restoring the joy of practicing medicine to clinicians around the world.

Our mission is to provide AI that empowers people to do more, and as technologists, we're with you on that journey.

Setting people up for success. The adoption of AI is meant to enhance the way people work, but it shouldn't be disruptive to their day. Clinical teams are rightly focused on saving lives, not learning a new application, and they're already in overdrive as most departments are running well beyond capacity. Setting up regular, clear communication about the anticipated benefits and implementation and integration plans is critical. Additionally, it's important to foster ongoing user involvement and create feedback loops with change champions to help ensure timely support and successful adoption.

Process integration is key to adoption. Hospitals produce roughly 50 petabytes of data every year, comprising clinical notes, lab tests, medical images, sensor readings, genomics, and operational and financial data. Yet 97% of it goes unused. Unlocking this requires coordination among multiple stakeholders, including healthcare providers, IT departments, data scientists, and regulatory bodies. It also requires a strong infrastructure that can handle the demands of the algorithms, including high-performance processing resources and scalable storage solutions.

However, it's not just about data integration; this extends to workflow integrations, too. In healthcare, every professional has a different job to do, workflows to follow, and systems to use.

Technology needs to support each one in the right way, at the right time, and in the right place. This means embedding it into existing systems and workflows to provide real-time insights, generate draft responses, and automate administrative tasks.

Technology, for technology's sake, doesn't cut it. There's a simple concept of "human agreeance" in technology: if the application does not have parity, fidelity of experience, or deliver on its promise, then it won't be adopted by users. Nowhere is this more evident than in healthcare. While generative AI holds a massive amount of promise, it comes down to scalable, secure solutions that support high-value use cases with exceptional user experience and that deliver real and tangible outcomes.

Underpinning everything is trust. AI must be deployed and utilized properly to ensure its effectiveness and trustworthiness, adhering to consistent, responsible AI standards that provide confidence to caregivers, administrators, and patients. And, for those of us who are building AI-powered tools, it's important that we share a common, validated framework for security, privacy, and safe use.

We're at an exciting inflection point, and it will take all of us working together—across the healthcare and life sciences industry—to create a better, healthier future.

Highlights

How AI leaders perceive value

Leaders who have made significant progress with generative AI within their organizations tend to see a common goal: transformation. Beyond simply making existing systems more efficient, they view generative AI as a way to reinvent processes and business models, creating a future that's frictionless, personalized, and engineered to accelerate scientific progress and improve patient outcomes.



[Watch](#) Karin Conde-Knape, Novo Nordisk SVP of Global Drug Discovery, shares how AI is helping deliver the next generation of treatments faster.

[Watch](#) A panel of business and research leaders discusses how AI is shaping the future of healthcare and biotechnology.

[Watch](#) Peter Lee, President of Research at Microsoft Research, and Vijay Mital, CVP and Chief Advisor for AI Transformations at Microsoft Research, discuss how AI is transforming healthcare.



Section 2 | Addressing challenges to generative AI adoption in healthcare

As healthcare and life sciences organizations make progress in their AI journeys, their leaders are enthusiastic about the transformative potential, from advancing medical research to improving clinical productivity and patient experiences. Transformation isn't without its roadblocks, but as healthcare and life sciences organizations adopt generative AI, they're overcoming challenges using best practices, innovation, and, in many cases, generative AI itself.

Common challenges organizations face today

Overcoming skill shortages and building expertise. Today, 45% of organizations say a lack of specialized skills is the top challenge when implementing AI technology.⁵ While this challenge is likely to persist, especially in a highly specialized industry like healthcare and life sciences, organizations can take advantage of training opportunities, such as the [Microsoft skills initiative](#), to empower employees with generative AI skills.

As familiar tools increasingly feature generative AI capabilities, deploying these solutions will become less complex. For example, Microsoft and Epic have partnered to deliver robust generative AI tools in clinical settings.⁶

Mitigating security concerns. AI adoption presents unique security challenges, especially for organizations that maintain protected health information (PHI). In healthcare environments, cybersecurity can be further complicated by vulnerable legacy applications, an expanding attack surface, and other factors.⁷

Threat actors are using AI to generate malicious code, increase the effectiveness of social engineering, and multiply the attack volume. As a result, 27% of Microsoft customers say concerns over data or IP loss are holding their AI initiatives back, while 26% cite a lack of governance and risk management.⁸

Although these concerns can slow adoption, security is also a common early generative AI use case. Security copilots and other tools that use generative AI can improve threat detection, accelerate response times, and provide more in-depth incident analysis.⁹ Meanwhile, security built for and embedded in generative AI systems is enabling organizations to mitigate security threats more easily.

Ensuring appropriate data governance. Safeguarding patient privacy and ensuring data security are paramount to maintaining trust and compliance with regulations. Adopting generative AI creates specific data governance challenges.

For example, the health-related data used in laboratory, financial, and clinical settings is often fragmented, inconsistent, or incomplete, which can make generative AI outputs less reliable. Harmonizing data from multiple sources, ensuring accurate labeling and annotation, and structuring data correctly can be resource intensive.

Healthcare and life sciences organizations that achieve strong data governance begin with AI-specific governance policies, develop robust systems to ensure results are traceable and auditable, and establish cross-functional oversight that includes IT, compliance, legal, clinical or scientific, and data science teams.

Addressing evolving ethical questions. The integration of generative AI in healthcare brings numerous benefits, such as improved diagnostics, personalized treatment plans, and enhanced efficiency in healthcare delivery. It also raises several ethical considerations organizations need to address to ensure responsible use. These include:

- **Bias and fairness:** Oversight of AI systems is needed to avoid perpetuating bias.
- **Informed consent:** Patients need to be informed about the role generative AI plays in their care.
- **Equity and access:** Efforts need to be made to ensure AI benefits all populations, including underserved and marginalized communities.
- **Professional roles and responsibilities:** Integrating AI requires clearly defined boundaries between human and machine roles.

While each organization will address these and other ethical concerns in its own way, participating in industry associations that study these issues can provide valuable insight and guidance.

For a discussion of approaches to responsible generative AI adoption, see our section on [Trustworthy AI](#).

Managing change across the whole organization. AI adoption can drive significant changes in how organizations operate, giving employees the opportunity to reimagine existing workflows. In healthcare and life sciences organizations, realizing the full benefits of generative AI also means integrating multiple specialized systems.

Effective transformation requires leadership to drive cultural change from the bottom up and from the top down, encouraging and training employees to use their generative AI tools, and discouraging the use of unsanctioned tools.

Variability in regulatory and compliance standards. As new capabilities emerge and AI-powered tools open up new areas of possibility, industry groups and governments are developing standards and regulations to guide the use of generative AI in clinical and research environments. Although some standards exist, they're still evolving as the use and capability of AI expands, and they may not address every use case. In addition, many countries do not yet have regulations in place. As a result, the regulatory landscape for generative AI remains complex, fragmented, and evolving.

While maintaining compliance amid regulatory uncertainty will present an ongoing challenge, AI tools can also make it easier to develop applications that support compliance because of their ability to automate change, with human oversight, to align with evolving guidance.¹⁰

Building and optimizing AI infrastructure. Democratizing access to AI-powered tools can make it difficult to manage infrastructure in a way that both facilitates generative AI projects at scale and maintains appropriate controls. Organizations' AI services are often fragmented, making it difficult to manage costs and ensure that infrastructure can scale and remain resilient.¹¹

Organizations that mitigate this challenge look at AI adoption holistically, creating a centralized management system to monitor AI usage and ensure infrastructure is available, optimized, and administered responsibly. As AI-optimized infrastructure proliferates, driven by both private investment and public initiatives, management will likely become less burdensome.

Recognizing the importance of standardized implementation. In 2009, the US government introduced the Meaningful Use program, which promoted the use of electronic health records (EHRs) for healthcare providers and hospitals. A primary goal was to improve the quality and safety of healthcare. Many EHR vendors entered the market, and there was a widespread belief that certification of EHRs would ensure high quality and safety for all hospitals using the certified EHR.

However, when researchers studied this, they found that outcomes varied considerably amongst hospitals using certified EHRs and that the greatest variability in performance occurred with hospitals using the same EHR system.¹² In other words, implementation of the technology had a greater impact on outcomes than the technology itself.

This has direct relevance to AI. To demonstrate that AI improves outcomes, it's important to understand and document how the AI is being implemented, including who receives the output, when they receive it, how they receive it, and what happens next. Outcomes associated with AI need to be evaluated in the context of AI plus implementation, and healthcare organizations should standardize on implementation approaches that have been associated with the best outcomes.



Mitigating the limitations of generative AI. While generative AI tools are transformative, users must be aware of the technology's limitations. For example, if a model misses important information, there will be omissions in the outputs it generates. Explaining an application's decision making can also be challenging, impacting the transparency essential in healthcare settings.

Most prominently, when the wrong models are used or data isn't correctly structured, applications can fabricate unreliable outputs, or "hallucinations." When those outputs are used to inform research, support diagnoses, or communicate with patients about their health status, inaccurate responses can have serious consequences. Today, 73% of leaders cite hallucinations as a leading concern that deters them from adopting generative AI more broadly.¹³

Organizations that mitigate the risks associated with these limitations adopt frameworks, like Microsoft's Responsible AI Mitigation Layers, that emphasize safety, fairness, and transparency.¹⁴ They're also thoughtful about how they integrate generative AI into workflows, using API tools to ensure that appropriate clinical safeguards are in place.¹⁵

Leadership perspectives

Staying ahead of emerging challenges and threats with AI and security for AI

By Ann Johnson, Microsoft Corporate Vice President and Deputy CISO



What it means to be secure is becoming more nuanced. New technologies and designs introduce new vulnerabilities and risks, and as the amount of sensitive personal and financial data grows, so does its appeal to attackers. Securing AI and its data from the start is a necessary practice to get ahead of cyberthreats before they attack.


The challenge many hospitals and healthcare organizations face is securing a technology that, despite rapid growth, is still in its infancy. Maintaining strong cybersecurity hygiene is fundamental, even in the era of AI.

Establishing a cyber foundation. Fundamental security measures such as multifactor authentication (MFA) and regular software updates are essential. Educating employees on cybersecurity principles is also critical, as many attacks exploit minor vulnerabilities rather than sophisticated AI-powered methods.

Measuring potential risk. Understanding where, what, and how much data is stored is vital for identifying vulnerabilities and restoring operations during incidents. A strong access management strategy is crucial, as human error often leads to breaches.

Assessing emerging technology risks. While AI enhances the sophistication of cyberattacks, traditional techniques remain effective. The reliance on digital communication channels increases the risk of AI-powered phishing attacks, especially in hybrid work environments. Educating staff about these threats can enhance vigilance and protect business infrastructure.

Security is, fundamentally, a data problem. Building cyber resilience starts with mastering the basics, from educating staff to ensuring an up-to-date access management strategy. While AI poses new threats, these foundational principles remain the most effective defense against many attacks today.



Highlights

Identifying the first use case—and growing it into a “super use case”

Implementing high-impact generative AI use cases early leads to better outcomes and accelerated progress toward AI goals.¹⁶ Collaboration between healthcare and life sciences leaders and AI experts in value-creation exercises helps to pinpoint areas where generative AI can deliver the greatest value. These exercises involve:

- Aligning stakeholders so administrative and technical leaders agree on priorities
- Identifying core processes that can benefit from generative AI integration
- Establishing metrics to assess the effectiveness of implementation
- Assessing data availability to ensure it's ready for generative AI applications



Section 3 | Navigating the AI platform shift

In Section 2, we discussed specific challenges for successful generative AI adoption. There are also global concerns requiring a holistic strategy that encompasses the entire organization. In studying our customers' AI journeys, Microsoft has defined two frameworks any organization can use to guide AI adoption.

AI Readiness outlines the maturity stages from exploring AI to delivering value. **Drivers of AI Value** outlines best practices for engaging leaders and empowering employees during AI implementation.

The 5 stages of generative AI readiness

In Microsoft's research, organizations progress through five stages of maturity in their AI journeys: **Exploring, Planning, Implementing, Scaling, and Realizing**. By identifying their organization's maturity stage, leaders can implement the strategies that will best help them drive generative AI transformation while maintaining compliance and navigating their unique challenges.

- 1 Exploring**
Organizations are educating themselves and experimenting with generative AI in select areas.
- 2 Planning**
Organizations are defining a generative AI strategy. Building on their strategic foundation, they run proofs of concept and plan generative AI deployments.
- 3 Implementing**
Organizations are moving from proofs of concept and pilots into production.
- 4 Scaling**
Organizations have deployed generative AI applications and systems for a few business functions and are scaling into new areas.
- 5 Realizing**
Organizations are achieving repeatable, measurable value. Not only have they piloted, implemented, and scaled applications, they've also built strong processes and guardrails for further development and have a robust development pipeline.

Where organizations are today

As more healthcare and life sciences organizations feel confident and comfortable with generative AI, AI usage and sentiment are trending upwards. This increases AI maturity as more organizations move from Exploring to Planning strategy and progress from Scaling AI to consistently Realizing value.



To help you understand your organization's level of AI maturity, we've created a simple [10-question assessment](#).


Progress in AI readiness

Customers moved forward on their AI journeys in 2024, with an increasing number scaling and realizing value.

Key

- Exploring
- Planning
- Implementing
- Scaling
- Realizing

Adjusted scores based on AI Readiness overstatement adjustment. Base: Commercial (n=400). Weighted. BIZ_AIMATUR: On a scale of 1-5, how would you describe your organization's overall adoption with regards to AI? Please answer where 1 is beginning stage of exploring the role of AI in the organization, and 5 is a relatively advanced stage where the organization can realize and measure the value driven by AI. Select only one. Data visualization on this page is not to mathematical scale.



The 5 drivers of AI value

As Microsoft customers move through the stages of generative AI maturity, we've found those that establish cross-functional teams to manage adoption across the organization see greater impacts from their investments and advance more quickly. Through our research, we've identified five key drivers of AI value that organizations can use to accelerate and harmonize AI adoption: **Business strategy, Technology and data strategy, AI strategy and experience, Organization and culture, and AI governance.**¹⁷

1

Business strategy

An effective business strategy focuses generative AI adoption on achieving the organization's most important goals, aligning investment with desired, measurable outcomes. Organizations that accelerate their generative AI journey use this strategy to:

- Identify the outcomes they want to accomplish
- Understand how generative AI can help them achieve those outcomes
- Define the most meaningful objectives and use cases
- Create a clear investment plan for achieving strategic goals

2

Technology and data strategy

To address technology and data challenges effectively, leaders should focus on ensuring access to the data and infrastructure needed to run generative AI applications at scale. Organizations that successfully implement generative AI initiatives successfully can:

- Provide access to quality data that's accurate and properly formatted
- Understand where applications and data are stored
- Ensure cloud infrastructure supports scaling
- Decide strategically whether to buy or build generative AI applications

3

AI strategy and experience

Organizations employ generative AI with greater impact as they accumulate experience and expertise. Through a combination of hiring and skilling programs and organizational change, successful organizations can:

- Assess whether the workforce has the diverse skills and experiences needed
- Create cross-functional teams that work across business units to build experience
- Apply the right models for each generative AI initiative

4

Organization and culture

A focus on organization and culture is essential to encourage the use of generative AI tools and establish a mindset that invites change and seeks continuous growth. Organizations that manage this most effectively can:

- Provide top-down leadership support
- Define an operating model to industrialize generative AI adoption and use
- Implement effective change management and continuous learning
- Build strong partnerships to support and accelerate generative AI initiatives

5

AI governance

Successful organizations meet their security and data governance obligations by developing policies using the Trustworthy AI framework which allows them to:

- Implement effective and transparent processes, controls, and accountability structures
- Govern data privacy, security, and responsible use
- Adhere to regulatory mandates and maintain awareness of changing regulations
- Assess the impact of generative AI investments on people, organizations, and society

For a discussion of the framework, see our section on [Trustworthy AI](#).



Leadership perspectives

Navigating the future of healthcare together

By Kees Hertogh, Vice President, Healthcare and Life Sciences Marketing



Generative AI has a transformative potential for healthcare, but navigating the shift will require careful preparation, significant change management, and collaboration across the healthcare and life sciences ecosystem.

A secure data and cloud foundation is essential. Healthcare organizations sit on vast, rapidly growing volumes of clinical information and operational data. Generative AI can turn this data into valuable insights for improving clinician and patient experiences, enhancing care quality, accelerating therapeutic breakthroughs, and delivering better patient outcomes. But first, the challenge is to integrate data from numerous siloed systems—and keep it protected from constantly shifting cybersecurity threats. Microsoft is committed to providing a complete cloud foundation that enables healthcare organizations to fuel AI tools with high-quality, integrated data while maintaining security and privacy.

Organizational transformation requires change management. Like any technological transformation, the AI platform shift is only partially about the technology itself. Without buy-in from every level of the organization, it's difficult to implement meaningful change or sustain it in the long term. By fostering a culture of collaboration and continuous learning, organizations can successfully navigate the AI platform shift and achieve lasting transformation.

Success will depend on an efficient and connected healthcare ecosystem. No organization can transform healthcare alone, but the AI platform shift offers an opportunity to bring the healthcare and life sciences ecosystem together to overcome the sector's most persistent challenges. Often, technology companies and their ecosystems of partners are uniquely positioned to connect organizations and empower them to collaborate and innovate together to accomplish their shared goal: improving the health of individuals, communities, and society.



Highlights

Effective and sustainable AI in healthcare

In 2022, Microsoft collaborated with The Health Management Academy and leading health systems across the United States to develop the AI Maturity Roadmap. Similar to AI Readiness and Drivers of AI Value, the AI Maturity Roadmap is a flexible framework specifically tailored for healthcare providers. It can be used as:

- A step-by-step process that helps health systems think about adopting generative AI sequentially
- A wayfinding tool with a loose structure that allows for each health system's unique priorities
- A communication tool that can drive alignment between leadership, implementation teams, and users

GO DEEPER Read the [AI Maturity Roadmap](#) for healthcare providers.

Section 4 | The winning formula

Perhaps the most striking aspect of the current state of generative AI is the difference between healthcare and life sciences organizations in the Exploring stage and those that have fully embraced the AI platform shift. Today, early-adopting high performers realize almost four times the value from their AI investments as those taking a more measured approach. They're also deriving value from more use cases and implementing generative AI projects in a fraction of the time.¹⁸

Generative AI creates impact throughout the healthcare and life sciences ecosystem. Examining what success looks like in **Patient care and education**, **Drug discovery and development**, **Medical research**, and **Health insurance** offers insight into what differentiates those who have found a winning formula.

Patient care and education

Identifying and following the right course of treatment delivers the ultimate benefit: helping patients live longer, healthier lives. Healthcare leaders are using generative AI to enrich the provider-patient relationship in important ways to improve patient outcomes.

→ What can organizations serving patient populations with advanced AI for patient care and education teach us about success?

They're making care more personal. The time patients spend with their clinicians is precious. Leaders maximize the value of this time and enable clinicians to focus more on their patients by using generative AI to summarize histories and transcribe interactions to create more accurate records.

They're improving equity and access. Underserved populations typically have worse patient outcomes. To change this, leaders are using generative AI to mine data, surfacing risk factors in underserved patients' histories, and identifying high-risk patients to drive more equitable resource allocation.¹⁹ They're also using generative AI translation capabilities to overcome language and cultural barriers in clinician-patient interactions.²⁰

They're empowering patients. Patients have a key role to play in improving their health outcomes, and leaders are using generative AI to help them become better stewards of their own health. Personalized educational materials and generative AI-powered chatbots that act as 24/7 virtual health assistants give patients tools to become more informed, proactive, and healthy.²¹



Customer story

Providence accelerates AI innovation at scale to improve care delivery and outcomes

Investing in generative AI has helped Providence increase support for scientific discovery while improving operational efficiency and clinician and patient experiences.

GO DEEPER [Read the full story.](#)

Drug discovery and development

Generative AI can dramatically accelerate the creation of new therapeutics by replacing the long, complex, and costly process of screening known compounds with the ability to generate entirely new ones with desired properties.

→ What can organizations on the cutting edge of AI-driven drug development teach us about success?

They're going big. Using generative AI to accelerate drug discovery and development requires significant data, infrastructure, and model resources. Leading organizations make investments and form partnerships to assemble and optimize these resources at scale, building a strong foundation for molecular generation and optimization and predictive modeling.

They're addressing a wider range of conditions. Leaders use generative AI to enhance the speed and precision of delivering high-revenue therapeutics to market. They're also using faster, less costly discovery and development to address underfunded areas, such as alternatives for drug-resistant bacteria.²²

They're balancing machine and human intelligence. Leaders take a deliberate approach to balancing the benefits of automation with the importance of human researchers' expertise. This allows them to accelerate experimental timelines while maintaining a rigorous scientific approach to critical areas of research.²³

→ As one of the most data- and compute-intensive applications for generative AI in healthcare and life sciences, drug discovery stands to benefit from investments and advancements in AI-optimized infrastructure. Look for significant gains in system efficiency for this use case in the near future.



Customer story

Novo Nordisk accelerates life-changing research

By creating a robust center for AI innovation within the company, Novo Nordisk is accelerating scientific progress, expanding the scope of its research, and exploring the potential of precision medicine.

GO DEEPER [Read the full story.](#)

Medical research

Research in healthcare and life sciences is arduous, time-consuming, and costly. Leaders in generative AI adoption are using this technology to minimize complexity and drive better patient outcomes.

→ What can organizations leading in AI-accelerated medical research teach us about success?

They're advancing genomics and proteomics. Generative AI models can predict protein structures and create accurate 3D models of proteins, deepening researchers' understanding of biological functions and disease mechanisms. Leaders are using generative AI to create synthetic genomic and proteomic data to augment data sets, preserve patient privacy, and enhance the training of predictive models.²⁴

They're optimizing clinical trials. Recruiting diverse and representative clinical trial populations can be difficult and imprecise, leading to less effective trials. By using generative AI to analyze patient data and automate candidate selection support, leaders are improving trial quality while eliminating mundane tasks that slow down selection.²⁵

They're unlocking value in healthcare data. Healthcare and life sciences organizations tend to be data-rich but converting that data into real value can be challenging. Leaders are using generative AI to create synthetic health records that retain the statistical properties of real data while maintaining patient privacy. Using these records, they're modeling disease progression more accurately, increasing researchers' understanding of potential outcomes, and informing treatment strategies.²⁶



Customer story

Boehringer Ingelheim speeds up knowledge sharing

A generative AI interface makes precise, factual answers from Boehringer Ingelheim's knowledge management platform instantly available, saving time and accelerating the development of needed therapeutics.

GO DEEPER [Read the full story.](#)





Customer story

Cognizant & Microsoft collaborate to power healthcare with generative AI

By using generative AI to enable process automation, improve regulatory compliance, enhance data security, and provide access to innovation and data insights, Cognizant and Microsoft are reducing hidden costs and creating better outcomes.

GO DEEPER [Read the full story.](#)

Health insurance

Members rely on health insurance payors to process claims efficiently and transparently. Leading insurers are using generative AI to make it easier for members to get the care they need without confusion over claims processing or authorization.

→ What can organizations improving processes with AI-driven applications for health insurance teach us about success?

They're enhancing the member experience. Every member's needs are different, and when insurers and providers are aligned in determining, authorizing, and initiating treatment, health outcomes improve.²⁷ By using generative AI to drive personalized communication, enhance customer service, and help improve access to care, leading payors are increasing member satisfaction and loyalty.

They're achieving operational enhancements. Enhancing benefit design, responding to member needs, and processing claims require significant staff resources and can lead to delays in care. By using generative AI to streamline the claims management process, leaders are seeing dramatic improvements in operational efficiency, such as increasing daily claim processing by as much as 60%.²⁸

They're reducing costs. By driving efficiency in their operations, payors can reduce administrative costs and dedicate more resources to their ultimate purpose: supporting member care. By fully integrating AI into every process, leading payors can reduce administrative costs by \$150 million to \$300 million and generate \$260 million to \$1.2 billion in additional revenue per \$10 billion in current revenue.²⁹



The emergence of AI agents

Organizations have already achieved dramatic impacts with generative AI. Now, AI agents are poised to take those results to the next level. Beyond assistance, AI agents can be tailored for specific expertise and work alongside or on behalf of users, handling tasks from simple queries to complex assignments.

AI agents work by receiving information, processing that information, and taking actions to achieve specific goals. Typically, they operate using a cycle of perception, processing, and action.

- 1 Perception.** The agent gathers data. This could be from a database, sensors, or a variety of other sources.
- 2 Processing.** The agent analyzes the data using models, logic, and algorithms to make decisions about the best course of action.
- 3 Action.** The agent performs an action, such as updating data, controlling a device, or interacting with users.

→ How are AI agents being deployed in healthcare and life sciences organizations?

They're enhancing workflows. Pharmaceutical companies are creating experiences that give clinicians a conversational interface that streamlines access to medical knowledge. By putting clinical documentation and relevant information in clinicians' hands, they're making it easier to make well-informed medical decisions.

They're improving the use of healthcare content. Organizations are creating tools for healthcare professionals that improve access to clinical knowledge, reducing the time spent looking for information and enhancing the efficiency of patient care.

They're supporting clinicians' administrative workload. Using agentic chatbots to answer questions about medical protocols and treatment information, hospitals are reducing clinicians' workloads, ensuring compliance, and enhancing patient safety.

GO DEEPER For insights into how to use AI agents with built-in healthcare safeguards, learn about Microsoft's [healthcare agent service](#).

Leadership perspectives

The role of partners and startups to advance innovation

By Sally Frank, Worldwide Lead - Healthcare and Life Sciences, Microsoft for Startups



Microsoft's partners, and particularly the early-stage companies in our ecosystem, play a pivotal role in driving innovation and AI adoption. They are often trailblazers, solving challenges to improve patient care and accelerate drug discovery and medical research. Our Microsoft for Startups [Pegasus Program](#) is designed to identify those trailblazing product-market fit startups and fast track their growth by building relationships between them and leading providers, payors, pharma, and medical device companies.

Drug discovery and development. Artificial supports drug discovery by leveraging AI and digital twins to accelerate the pace of life science discoveries. They provide a comprehensive software platform called the Artificial Digital Lab Platform that connects and controls all aspects of a lab, including hardware, software, people, and data, enabling scientists to make breakthroughs faster. Furthermore, by [combining their \(Azure-based\) solution with NVIDIA's BioNeMo](#), they're able to offer seamless adaptability and scalability for drug discovery and a broader range of laboratory environments.

Medical research. In the field of medical research, innovators like [Triomics](#) leverage generative AI to streamline clinical trial processes. Their platform, [PRISM](#), aids in prescreening patients for trials, enhancing cohort feasibility, and continuously tracking unmatched patients. This accelerates research timelines and improves the accuracy and efficiency of clinical trials, leading to better health outcomes. Additionally, their [Harmony](#) quality improvement solution assists with tumor registry completion, precision oncology workflows, and business intelligence/advanced analytics requests. Triomics' [Symphony](#) care operations solution aids in pre-charting and EHR search, triage protocol implementation, and identification of missing clinical data. They've [collaborated with Mount Sinai Health System](#) on decentralized trial research and have worked closely with the [Medical College of Wisconsin](#) to advance oncology research and care operations.

Patient care and experience. For patient care and experience, [Strolll](#) offers a clinically validated, AI-driven digital therapeutic using off-the-shelf augmented reality (AR) glasses, enabling patients with Parkinson's, those recovering from a stroke, or others engaged in neurological rehabilitation to improve their gait and balance. Additionally, the solution enables clinics to scale services and improve access to neurorehabilitation by offering both in-clinic and at-home options. They're [working closely with Cleveland Clinic](#) and have shared patient results in [impactful videos](#).

Technology partnerships. The foundation of meaningful and lasting innovation is often the technology behind visionary leaders and ground-breaking concepts. [NVIDIA Inception](#) and [Microsoft for Startups](#) are working together to foster the next wave of breakthroughs in healthcare and life sciences. [This collaboration](#) combines Microsoft's AI, cloud, and enterprise expertise with NVIDIA's pioneering advancements in AI hardware and software, enabling startups with the tools, resources, and support they need to fast track their ideas and deliver life-changing outcomes.

Microsoft's partners, particularly early-stage companies, are crucial drivers of innovation in the healthcare and life sciences sector. These trailblazing startups are addressing critical challenges in patient care, drug discovery, and medical research, ultimately improving health outcomes.

The Pegasus Program and partnerships with technology leaders like NVIDIA Inception and Microsoft for Startups provide the necessary support and resources for these startups to thrive. By leveraging AI and other advanced technologies, these collaborations are paving the way for groundbreaking solutions that will shape the future of healthcare.



Section 5 | Trustworthy AI

The AI platform shift opens up extraordinary possibilities. But as organizations progress on their AI journeys, a failure to foster and maintain trust can have serious consequences, from increasing skepticism among users to exposing sensitive data or falling victim to a cyberattack.

With any AI deployment, success depends on the trustworthiness of the tools being used. That's the foundation of [Trustworthy AI](#), a framework for developing, deploying, and using generative AI applications.

What is Trustworthy AI?

Trustworthy AI comprises three pillars: **Security, Privacy, and Safety**.

Security requires protecting AI systems, data, and users from threats, misuse, and vulnerabilities. It addresses the risks posed by threat actors using AI to make their attacks more persistent and sophisticated and by configuration challenges in complex cloud environments. Under this pillar, applications should be secure by design, secure by default, and subject to continuous monitoring and improvement to secure operations.

GO DEEPER Read about our commitments to security in the [Secure Future Initiative](#).

Privacy focuses on protecting the data organizations collect and use throughout the application lifecycle. It addresses the risk of AI applications exposing sensitive or proprietary data, which could violate regulatory requirements and severely undermine trust. Under this pillar, organizations must maintain control of their data, know where it resides, and ensure it remains secure at rest and in transit.

GO DEEPER Read about our commitments to privacy in the [Trust Center](#).

Safety includes security and privacy and ensures that AI systems operate reliably and predictably across various environments and conditions. It addresses the risk of generative AI applications producing inconsistent, unreliable, or culturally biased content. Rooted in the principles of Responsible AI, this pillar also involves upholding the values of fairness, reliability, inclusiveness, transparency, and accountability.

GO DEEPER Read about our commitments to safety in [Responsible AI](#).

➔ For healthcare and life sciences organizations, incorporating clinical safeguards into AI systems is a critical element of building Trustworthy AI. Collaboration with a technology partner can streamline this, as Microsoft customers have found while employing the healthcare agent service in [Microsoft Copilot Studio](#).

Trustworthy AI in healthcare and life sciences

While Trustworthy AI applies to all organizations that adopt AI, each sector has unique needs that must be addressed to maintain trust. In healthcare and life sciences, these include strict regulatory mandates, the obligation to protect PHI, and sector-specific ethical considerations, such as ensuring equitable and accessible care.

In line with our commitment to the responsible development and use of AI, Microsoft is a supporting member of several collaborative bodies dedicated to advancing the deployment of safe, reliable, transparent, and fair AI throughout the healthcare ecosystem.



Coalition for Health AI

Coalition for Health AI (CHAI) is a diverse group of stakeholders, including healthcare and technology leaders, who collaborate to develop guidelines and best practices for how healthcare AI should be designed. Members also help establish how AI tools can be deployed with equity and fairness to ensure the benefits are felt by every clinician and patient.

Through a nationwide network of laboratories, CHAI also tests and validates healthcare AI models and tools, publishing details about their performance in advance of their widespread adoption.

GO DEEPER [Read about the Coalition for Health AI.](#)

Trustworthy and Responsible AI Network

In order to operationalize responsible AI (RAI), healthcare organizations need to invest additional time and resources beyond implementing AI. This may include:

- Taking an inventory of all the AI models running in their system
- Testing the models on local data sets
- Monitoring post-deployment for AI drift
- Evaluating the impact that the AI model has on both expected and unexpected outcomes
- Evaluating for bias in the AI model and in its impact on subpopulations
- Developing streamlined governance processes

Unfortunately, the amount of AI that needs to be assessed, monitored, and governed has increased beyond the current capacity of health systems. In 2024, healthcare systems in the US, Europe, and the Middle East, working in collaboration with Microsoft, formed the Trustworthy and Responsible AI Network (TRAIN). TRAIN’s goal is to help healthcare systems operationalize RAI in the most time-, resource-, and cost-efficient manner possible.

The TRAIN approach relies on three principles:

- 1 Make RAI technologies accessible.**
Technologies can make things more efficient and streamlined. TRAIN intends to introduce and help evaluate RAI technologies for its members.
- 2 Collaborate with other health systems.**
Collaboration involves identifying the most commonly used and/or high-impact AI models, documenting the implementation approach for each model, monitoring for expected and unexpected outcomes, and sharing results between TRAIN members.
- 3 Collaborate with technology partners.**
TRAIN will work with AI developers and RAI technology enablers to create programs where they can participate in TRAIN.

GO DEEPER [Read about the Trustworthy and Responsible AI Network.](#)



TRAIN founding members

Healthcare AI models

For models to be trustworthy, they need to emulate the dynamics of the practices they support. Healthcare providers consider the whole picture of a patient’s health, including medical imaging, genomics, clinical records, and other data sources. Trustworthy models must be equally multimodal, so Microsoft has collaborated with partners, including Providence and Paige.ai, to develop multimodal models that enable healthcare organizations to create generative AI solutions rapidly while maintaining Trustworthy AI standards.

GO DEEPER [Read about multimodal healthcare AI models.](#)



Takeaway

Trust is vital as healthcare and life sciences organizations adopt generative AI. Emerging resources help organizations reach their goals without jeopardizing trust.



Customer story

Kenya Red Cross expands access to mental health support with Azure AI-powered chatbot

The [Kenya Red Cross Society](#) faced a significant challenge in providing mental health support due to a shortage of mental health professionals and services, making care inaccessible for many individuals. To address this issue, the Kenya Red Cross utilized [Microsoft Azure AI](#) to develop an AI-powered chatbot named Chat Care. This innovative solution aims to bridge the gap in mental health care availability and accessibility, particularly in regions with limited resources.

GO DEEPER [Read the full story.](#)

Leadership perspectives

Building trust to operationalize responsible AI in healthcare

By Dr. David Rhew, Chief Medical Officer and VP for Healthcare



The tremendous potential of AI to revolutionize the healthcare and life sciences sector is undeniable. However, that potential will remain unrealized unless there's trust in the models and AI-powered tools we use to transform the diagnosis, treatment, and prevention of disease. Adhering to responsible AI principles and the Trustworthy AI framework will help foster trust, but there are other important factors that will determine whether the promise of generative AI becomes an operational reality.

Healthcare needs a blueprint for adopting trustworthy and responsible AI. Fortunately, thanks to the efforts of the Coalition for Health AI (CHAI), such a blueprint now exists. CHAI has defined key principles and best-practice guidelines for healthcare AI development and equitable deployment. It also assesses the safety and efficacy of emerging healthcare AI models and tools, providing transparent evaluation and a trusted source of validation prior to widespread adoption.

Standardized processes for responsible implementation are essential. To operationalize the principles and guidelines developed by CHAI, healthcare organizations need to understand how to apply AI safely and effectively. That's why the work being done by the Trustworthy and Responsible AI Network (TRAIN) is so important. As well as helping healthcare organizations inventory AI models and their expected outcomes, TRAIN is also establishing AI governance models and enabling the development of a standardized and federated AI outcomes registry to help ensure that AI tools remain safe throughout their lifecycle.

Ongoing assessment is the key to sustaining trust. Unlike traditional technologies, AI can change over time, so assessment is never a one-off activity. For example, organizations need to continually monitor for model drift that leads to declining AI performance. It's also vital to periodically assess the impact of model performance on subpopulations to ensure AI outcomes are always unbiased and fair.

This combination of responsible AI standards developed by the healthcare community, proven implementation best practices, and continuous performance monitoring will help organizations build the trust that's so essential for the AI-fueled transformation of healthcare.



Highlights

AI for good

From discovering ways to make communities more resilient to supporting the fight against infectious disease, the AI for Good Lab confronts the world's most persistent and complex challenges. "AI for Health," a program within the AI for Good Lab, is investing \$60 million over five years to support projects that use AI to accelerate medical research, build research capabilities, increase global health insights, and address health inequities. The program has partnered with over 200 grantees on projects addressing population health, imaging analytics, genomic and proteomic research, and more.

GO DEEPER Read about the [AI for Good Lab](#).



In some cases, AI isn't just a solution. It's the only solution we have."

Juan Lavista Ferris
CVP and Chief Data Scientist
AI for Good Lab



Conclusion

We're all in this together

As the real-world impacts of the AI platform shift take shape, organizations of all kinds will need to learn, adapt, and refine their approaches.

Collaboration, partnership, and shared learning will be vital for everyone.

Microsoft remains committed to sharing our learning in the age of generative AI to further our mission of empowering every person and every organization on the planet to do more.

In the forthcoming Part 2 of the AI in Healthcare Decision Brief, "Perspectives on the role of AI in R&D and addressing global health challenges," we'll explore the potential of generative AI and a practical approach for accelerating adoption in your organization.

Resources

Throughout this report, we've provided various resources to assist you in your generative AI journey. Here are three additional recommended resources that may prove particularly useful:

- 1 Explore the latest AI trends and news at [Microsoft for Healthcare](#).
- 2 Assess your AI readiness with the [AI Readiness Wizard](#).
- 3 Learn more about [Microsoft AI](#).



© 2025 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet website references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.