

Accelerate business growth with agentic AI

A guide to realizing agentic AI value from start to scale

August 2025



Contents

Agentic AI: The new frontier in AI	3
5 key considerations for implementing agentic AI	7
Making it work: How to operationalize agentic AI	10
Agentic AI in practice	12
Access agentic AI expertise	13

Agentic AI: The new frontier in AI

The AI market is booming, and for good reason. The technology is already delivering practical, transformative solutions.

By late 2024, more than 60% of organizations globally already believed GenAI would be a game changer within two years, and nearly two-thirds were planning to invest significantly in GenAI during this period.¹

Now, as the demand for autonomy and agility in business operations soars, we find ourselves at the dawn of the agentic AI era. Organizations need smarter, more integrated AI-driven systems that can work seamlessly across functions and domains — and agentic AI is the next big leap in this journey.

As per Gartner®, "Agentic AI is one of 10 current trends identified in 2025 Gartner Top Strategic Technology Trends."²

"By 2028, at least 15% of day-to-day work decisions will be made autonomously through agentic AI, up from 0% in 2024," says the Gartner report.

Unlike GenAI, which primarily boosts productivity and creativity, agentic AI can fundamentally change how tasks are executed by infusing systems with autonomy and intelligence.

Agentic AI systems understand their environment, make decisions based on this understanding and take actions to achieve specific goals — qualities that are especially valuable in dynamic and complex settings where adaptability and real-time decision-making are essential.

This transformative technology is reshaping industries, redefining market dynamics and heralding a future where machines not only assist humans but also autonomously enable innovation and growth.

Agentic AI is a new frontier in AI that creates valuable opportunities for innovation, efficiency and problem-solving, making it a key player in the future of technology.

However, many organizations remain hesitant to fully embrace agentic AI due to a mix of technological, operational and cultural concerns. Overcoming these concerns requires technical readiness, a clear change-management strategy and a thoughtful approach to aligning AI capabilities with business goals.

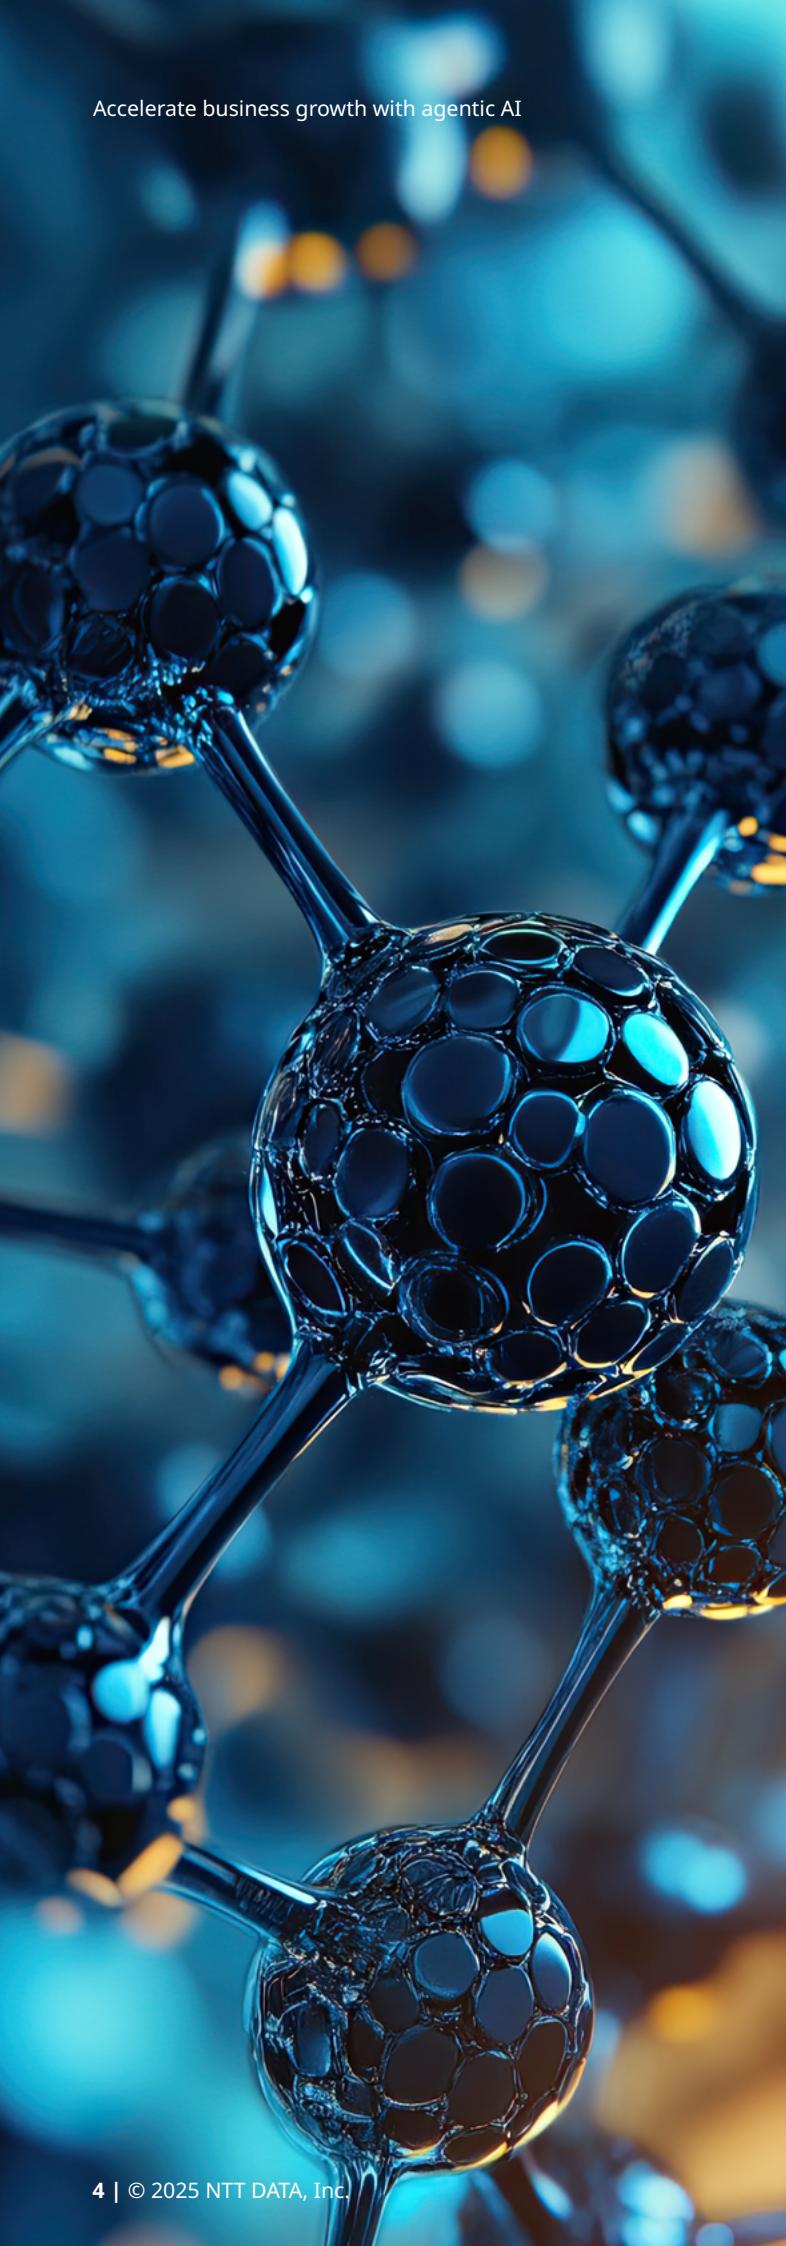
Whether you're exploring agentic AI for the first time or starting to scale it enterprise-wide, this guide offers a roadmap to turn vision into value — with speed, scale and efficiency.

Waiting could mean falling behind in a rapidly evolving, AI-driven economy. By adopting agentic AI now, organizations position themselves to lead rather than follow.



¹ NTT DATA, [Global GenAI Report](#). November 2024

² Gartner, [2025 Top Strategic Technology Trends](#). 21 October 2024. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.



From task-based to agent-based

Traditional AI initially focused on task-specific applications, making it ideal for automating specific processes.

However, organizations often manage dynamic, interconnected projects that span enterprise functions and departments. This complexity has paved the way for autonomous agents — and a wealth of new business opportunities.

In this context, “agents” refers to autonomous software entities that perceive their environment, make decisions, learn from these decisions and take actions to achieve specific goals. They can be simple, like a program that intelligently responds to user queries, or complex, like communities of agents that orchestrate and administer roles.

An AI agent does what you ask. Agentic AI figures out what needs to be done — and does it.

These agents rely on a combination of large language models, tools and workflows to perform tasks, and can be highly specialized and interconnected through multiagent models.

Agentic AI is the framework that enables these agents to collaborate with other agents and with humans on an enterprise-wide scale. Within this framework, AI agents can specialize in certain tasks and have access to controlled knowledge bases for increased accuracy and more autonomy. This makes agentic AI an incredibly powerful tool for solving complex business challenges.

The strategic case for agentic AI

Advances in AI, coupled with underlying technologies like cloud platforms, GPUs and machine-learning frameworks, have made agentic AI more feasible and accessible than ever before. Agentic AI projects can begin with simple use cases that offer significant value, and using hyperscalers can make it easier to get security-team approval for these projects.

Organizations are under increasing pressure to innovate, streamline their operations and deliver exceptional customer experiences. Agentic AI can help achieve all of these goals by:

- **Boosting productivity:** Automating mundane tasks and enabling employees to focus on higher-value work
- **Supporting innovation:** Creating business models and services that were previously impossible
- **Scaling operations:** Handling large volumes of data and processes with ease
- **Improving customer experiences:** Providing personalized, real-time interactions that delight customers and keep them loyal

A key benefit of agentic AI across all industries is its ability to streamline communication between systems and platforms.

For instance, organizations can create integrated workflows that span multiple applications and systems while maintaining robust security. This enables AI agents to communicate securely without sharing metadata that could pose a risk.

Agentic AI can autonomously plan and execute complex tasks across tools and departments, boosting productivity without scaling costs. As the technology advances and more organizations start using agent-based systems, early adopters will benefit from faster decision-making, better customer and employee experiences and future-ready operations.

Adopting agentic AI now can give organizations a decisive edge in efficiency, innovation and adaptability, and it's not just about traditional automation.



Where will agentic AI make a difference?

To appreciate and understand the power of agentic AI, organizations must focus on high-impact use cases.

Data analysis in healthcare

By analyzing a patient's data, medical history and current symptoms, AI agents in healthcare can recommend a course of action to the patient's physician, who still has the final say. This helps improve patient outcomes and lightens health-provider workloads.

Predictive maintenance in manufacturing

In manufacturing, agentic AI can predict when machines are likely to fail and schedule maintenance proactively. This reduces downtime, extends the life of the equipment and saves money. AI agents can also optimize production schedules and handle quality control.

Streamlining complex processes

In industries such as insurance and automotive manufacturing, agentic AI is streamlining complex processes like claims processing and order management. By integrating AI-powered dispatch and inquiry agents with omnichannel communication gateways, organizations can become more responsive and raise the bar in customer engagement.

Customer service

One of the most significant applications is in customer service, where agentic AI provides 24x7 support, improving both employee and customer experiences. This technology is particularly valuable in areas such as HR and sales, where human interaction is crucial but resources are limited.

IT service management

Another important use case is IT service management, where agentic AI automates routine tasks and lightens the workload for IT teams. AI agents can triage incoming requests and eventually handle many of the tasks independently, enabling human-in-the-loop escalations to be more specialized and precise.

Stock management in retail

In retail, AI agents can predict demand, optimize stock levels and even handle supplier negotiations. This means products are always available when customers need them, which reduces waste and increases profitability.

Agentic AI systems thrive in complexity. They are transforming business processes and customer experiences by offering unparalleled levels of autonomy and adaptability, while presenting numerous opportunities for innovation and continuous improvement.

However, it is crucial to address the associated compliance and security risks. One of these is shadow AI — the use of AI agents in the business without the approval or oversight of IT. Such agents may have unknown goals and questionable security practices, creating the risk of vulnerabilities, data breaches and operational disruptions. Enforcing governance and oversight ensures that all AI agents are developed and deployed in a secure, controlled environment.

By working closely with stakeholders and gathering feedback, you can identify new and creative ways to leverage agentic AI safely and securely.

But who bears the responsibility for implementing agentic AI correctly so it fulfills its potential? Several members of the C-suite must assume this responsibility. Let's explore how to navigate these uncharted waters with clarity and care.



5 key considerations for implementing agentic AI

It's easy to dive into agentic AI without a game plan, but there are real and pressing issues to take care of first.

Organizations may find agentic AI to be daunting and struggle with misconceptions about its capabilities and costs. Many confuse agentic AI with automation, underestimating its potential and overestimating its expenses. In reality, a single AI agent can handle a vast number of tasks, far exceeding human or legacy infrastructure capacity. However, the journey to implementation is often hindered by internal politics, short-sighted or siloed strategies, unclear ownership, lack of change management, and data-security concerns. By addressing these five considerations, you can start to implement agentic AI with confidence and realize its full potential.

1 Technological readiness

To support agentic AI at scale, you need to prepare your infrastructure to handle the unique demands of these advanced systems. As with any new technology, there are many aspects to consider that will affect IT policies and procedures.

Scalable computing resources

Agentic AI requires significant processing power, memory and storage to handle complex decision-making and autonomous operations. CIOs need to invest in scalable infrastructure, such as cloud services or high-performance computing clusters, to support these needs.

Working with hyperscalers helps make agentic AI scalable. Expert service providers can build private agentic AI "factories" that scale with demand.

Data management and integration

Agentic AI depends on diverse, high-quality data. Data-management systems must be able to handle large volumes of data from multiple sources, provide real-time data integration, and maintain data consistency and integrity.

Interestingly, agentic AI itself can help with this — you can implement specialized AI agents to manage your data. This approach is particularly useful for projects that require very structured data sets.

Interoperability and integration

Agentic AI systems often need to interact with legacy technology, other AI systems, applications and human employees. Your organization's infrastructure must support seamless integration and interoperability between these systems, using standards and application programming interfaces to facilitate communication and data exchange. This is also likely to prompt the reengineering of workflows throughout the organization.

An expert service provider can help enable the connectivity that allows AI agents to interact with the public in measurable and scalable ways while also integrating compliance technology.

Monitoring and maintenance

Tracking system performance, detecting and addressing anomalies, and updating models and algorithms when needed are all essential to ensuring the ongoing and reliable performance of agentic AI systems.

Observability and governance

Implementing a robust observability and governance process is crucial to managing the risks of agentic AI and ensuring its ethical use. Such a process should involve clear policies and procedures for monitoring AI behavior, transparency in decision-making processes, and compliance with regulatory standards.

While addressing these requirements is a significant undertaking, it is the only route to realizing a tangible ROI from agentic AI.

2

Workforce buy-in

You're likely to face internal resistance if employees fear job displacement and a loss of control, and are uncertain about how their roles will evolve alongside agentic AI.

The successful integration of agentic AI demands a distinct set of skills, which can be challenging to find. These skills range from data science, machine learning and systems integration to change management, user experience design, and project management to oversee the implementation and deployment of agentic AI systems.

Even with in-house technical expertise, there's often a gap in understanding exactly how to integrate agentic AI into existing business processes. This can lead to resistance and skepticism from employees and make it difficult to get their buy-in.

Collaborating with a partner to demonstrate the capabilities of agentic AI to key stakeholders will help them identify where it can have the most impact. A virtual roadshow is an effective way of achieving this.

Effective change management and education are crucial to gaining employee buy-in. By investing in employee upskilling and reskilling, and nurturing a culture that embraces AI-driven transformation, CIOs can streamline their organizations' transition to agentic AI.



3

Risk, governance and control

Security, compliance and ethical concerns are major barriers to the adoption of agentic AI.

Making AI systems secure, compliant and transparent is an organizational responsibility — one that's especially challenging in highly regulated industries such as healthcare and finance.

Because agentic AI systems operate autonomously, they introduce new security risks. These risks require tougher security measures, such as access control, encryption and monitoring, to prevent potential breaches or misuse.

Organizations must also establish clear governance policies to guide the development and deployment of agentic AI. Responsible innovation must take precedence over risky decisions.

Partnering with an accountable provider can ensure that all agentic AI work is observable and results are clearly measured.

4

Making the case for agentic AI

The pressure to demonstrate rapid ROI from agentic AI initiatives is intense.

Enterprise leaders are often asked to justify the cost of AI projects with clear, measurable outcomes. Without proven use cases or industry benchmarks, it can be tough to make a compelling case for bold investments.

To overcome this hurdle, develop robust metrics and evaluation frameworks that can capture the value generated by agentic AI projects — whether through cost savings, revenue growth or improved operational efficiency.

By doing so, you build a stronger narrative to support future AI investments and obtain organizational buy-in. You can also explore emerging industry benchmarks, collaborate with peers and build on early agentic AI successes, ultimately alleviating the pressure to deliver quick returns.



5

The gap between innovation and deployment

The gap between innovation and deployment is where many agentic AI projects stall.

It's one thing to have a great idea but another to turn that idea into a working, scalable solution. To bridge this gap, you need a clear, actionable roadmap for agentic AI adoption that addresses the technological, organizational and cultural challenges inherent in deploying AI at scale (a process that we explore in the next chapter).

Transforming promising AI innovations into real business outcomes that have a meaningful impact — and a clear ROI — requires a multifaceted approach that encompasses not only technical considerations, such as infrastructure and data management, but also change management, talent development and strategic alignment with business objectives.

A structured approach with clear governance and expert guidance helps leaders navigate these challenges.

By focusing on specific, achievable goals and even setting aside more complex issues like unstructured data, you can make tangible progress with agentic AI. This pragmatic approach enables you to capture the real value of agentic AI without getting bogged down by excessive analysis and complexity.

Making it work: How to operationalize agentic AI

Agentic AI requires a sophisticated understanding of autonomous decision-making and multiagent systems. This involves deep technical expertise, knowledge of proven frameworks and access to advanced infrastructure.

As you start using the technology in your organization, set out on a detailed roadmap from the start, with clear overall goals and measurable milestones along the way.



Exploration and roadmap design

Explore the potential of agentic AI by identifying areas of your organization where agentic AI is likely to have the most impact.

These include:

- **Customer service:** Automating support queries and improving response times
- **Supply-chain management:** Streamlining inventory management and logistics
- **Data analysis:** Implementing predictive analytics and improving decision-making

Create a detailed plan that outlines the goals, scope and expected outcomes — such as cost savings or efficiency improvements — of each use case.

Based on this plan, construct a comprehensive roadmap, from proof of concept to full implementation, with measurable milestones. Take into consideration the need for infrastructure upgrades, staff upskilling, stronger security and governance and more, as outlined in the previous chapter.

During this stage, be aware that failing to engage employees poses a major risk. Lack of stakeholder engagement at the outset can lead to resistance and less buy-in once it's time to implement agentic AI.

Adhering to ethical and legal frameworks will limit the risks associated with compliance issues and help to prevent reputational damage.

Stage 2

Pilot with scalable solutions

Once you have a clear plan, it's time to test the waters and gather feedback.

One way of doing this is to "start small and scale fast." Be sure to align your IT and business goals — an approach that is key to successful adoption.

By starting with a small, manageable project, you can demonstrate the value of the technology to stakeholders and build momentum for larger initiatives.

For instance, you could begin by deploying an AI agent to handle customer inquiries related to order status.

Initially, the agent is trained on a limited dataset and integrated into a single customer-service channel, such as a website chatbot. As the agent proves effective in resolving customer queries, it gains understanding and intelligence. You could then introduce additional agents in other customer-service areas, such as technical support or product recommendations, along with agentic AI manager agents to coordinate these agents in working together to triage and resolve customer issues.

This step-by-step approach will allow you to incrementally demonstrate the value of agentic AI and its broader business impact to stakeholders.

And, once you have proof of value and business buy-in, you can scale the agentic AI system across the organization — for example, by linking it to customer relationship management and financial systems — to deliver even greater business impact.

Stage 3

Stage 4

Manage your agentic AI ratios

As you progress along your agentic AI roadmap, change management is critical in keeping your agentic AI systems not just technically sound but also culturally aligned within your organization.

This includes helping employees to understand how their roles will evolve alongside those of AI agents.

The relationship between humans and AI agents is a delicate one. It's important to strike a balance where AI complements human work rather than replaces it. This means designing systems that allow for human oversight and intervention, and creating a culture where employees feel empowered and supported.

Offer training and education programs to help your team understand AI and how to work with it effectively. This includes executive education to align leaders in the business with your organization's vision for agentic AI so they can make informed decisions.



Mature with multiagent orchestration

The final stage is to mature your agentic AI system by implementing multiagent orchestration and change management. This involves coordinating multiple agents to work together, rethinking workflows and making sure your organization is ready for this new operational paradigm.

As you start to deploy more AI agents, the need for orchestration becomes critical.

Just like a symphony needs a conductor, your agents need direction to work together seamlessly. This is where agent-to-agent management comes in. It allows you to coordinate multiple agents, prioritize tasks and track outcomes so that your agentic AI systems are always operating in harmony.

This delivers new levels of accuracy and automation, and it enables agents to work together as part of a system or workflow.

When you're managing agents, you need to define clear roles and responsibilities for each one, along with robust communication protocols to facilitate their smooth interaction. You also need monitoring and feedback mechanisms to track agent performance and adjust your strategy as needed.

Additionally, prioritize transparency and explainability in agent decision-making to build trust and reliability in the overall agentic AI system.

There is plenty of complexity involved in agent-to-agent management, making the use of proven frameworks and platforms essential. These structures provide the tools and processes you need to orchestrate complex agentic workflows, manage dependencies and ensure that each agent is contributing to your overall business objectives through a coordinated approach.

Agentic AI in practice

Insurance assistance

An insurance company faced significant challenges with high churn rates among life-insurance policyholders, mainly because these policyholders did not understand their policy benefits and struggled to obtain answers to policy-related queries.

The goal was to reduce churn rates by providing a tool that improved the user experience and answered customers' questions.

NTT DATA developed a conversational multiagent system accessible through the insurance company's WhatsApp channel. The system manages and orchestrates the following agents:

- **Supervisor agent:** Interprets customers' questions and directs them to the appropriate agent
- **Policy-details agent:** Answers specific questions about a customer's life-insurance policy, such as coverage and beneficiaries
- **Retrieval agent:** Answers general procedural questions, like how to change beneficiaries
- **Health and wellness agent:** Provides recommendations about the company's health and wellness program and tracks a customer's progress in using it
- **Dumpster agent:** Handles unrelated queries by providing generic responses

Now, customers receive immediate and accurate responses, and these personalized and informative interactions foster greater loyalty and customer retention.

User profiling for targeted advertising

A television streaming service wanted to improve ad relevance for their 8 million users by creating psychographic profiles based on content preferences, user behaviors and viewing patterns. The company also wanted to identify detailed user preferences across demographics. This involved handling large volumes of user data to develop meaningful insights efficiently.

NTT DATA used a multiagent autonomous system integrated within Amazon Web Services to automatically analyze and categorize user profiles.

Now, AI-driven agents (designed to assume roles such as "profiler," "chief director" and "advertising specialist") work together to build psychographic profiles by processing demographic, behavioral and transactional data.

The system also continually learns and adapts, allowing for real-time refinement in ad-personalization strategies.

Access agentic AI expertise

NTT DATA is a global leader in agentic AI products, tools and managed services. We have the expertise to build, deploy, secure, optimize and manage AI agents.

Our [Agentic AI Services for Hyperscaler AI Technologies](#) offer a comprehensive suite of cloud-managed services to help organizations implement agentic AI by using hyperscaler AI technologies.

For example, by using Microsoft Azure AI Foundry Agent Service, we can build, manage and orchestrate multiagent workflows across multiple platforms. This approach simplifies complex multiagent deployments.

Additionally, we're making our tooling available across clients' cloud platforms. This means that any agent we build can easily be transferred to your cloud environment.

We work with expert partners to achieve this, including Microsoft, whose cutting-edge platforms like Azure OpenAI Service, Foundry, Copilot and Fabric are powering the next generation of intelligent agents.

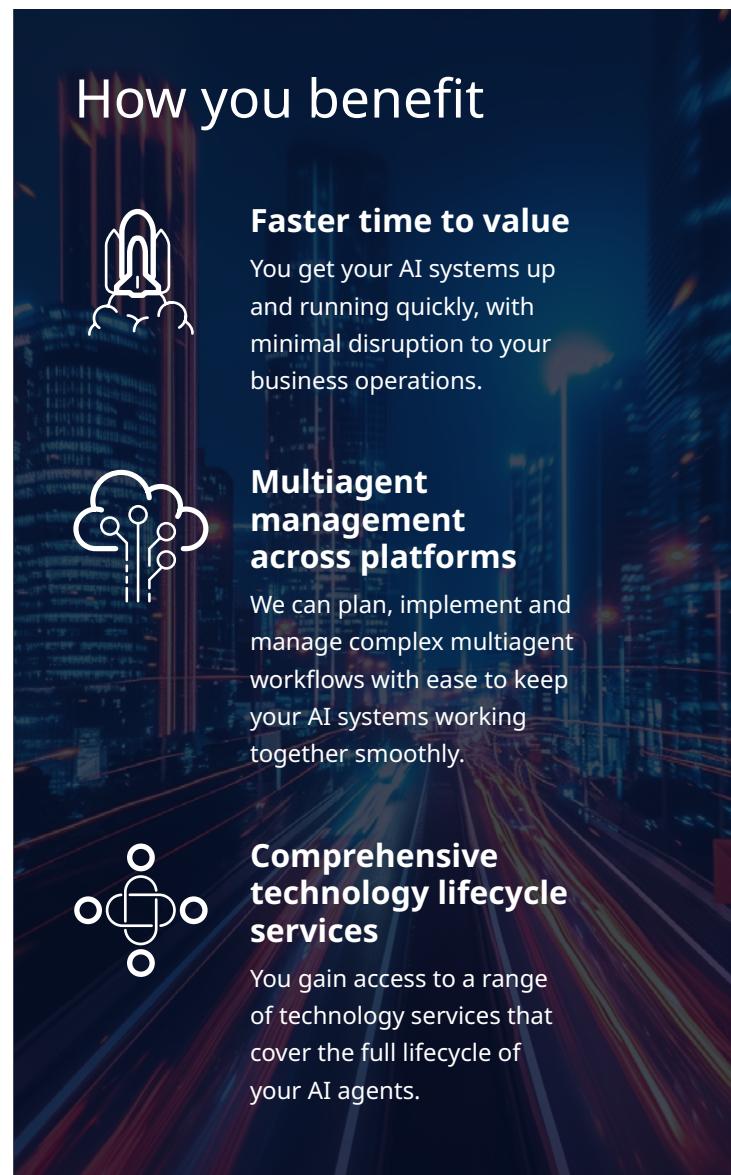
Microsoft is also a part of our [Smart AI Agent™ Ecosystem](#), which speeds up and simplifies agentic AI adoption.

Embrace agentic AI

Don't let the complexity of agentic AI adoption hold you back. **This new technology is already reshaping industries and the way we work. Waiting to adopt it is the real risk.**

Partner with us to transform your agentic AI vision into tangible business results.

[Learn more](#) about NTT DATA's Agentic AI Services for Hyperscaler AI Technologies.



Visit [nttdata.com](#) to learn more.

NTT DATA is a global innovator of digital business and technology services, helping clients innovate, optimize and transform for success. As a Global Top Employer, we have experts in more than 50 countries and a robust partner ecosystem. NTT DATA is part of NTT Group.

 **NTT DATA**

