

AI IN ACTION: Creating better experiences with Azure AI apps and agents



CONTENTS

03

Introduction

04

Driving business value with AI

- 05 AI apps and the evolution of agents
- 06 Agentic AI: Expanding what AI can do for your business
- 07 Example use cases for AI apps and agents
 - 07 Accelerate content delivery
 - 07 Transform customer service
 - 08 Optimize employee workflows
 - 08 Prevent fraud and detect anomalies
 - 09 Unlock organizational knowledge
 - 09 Automate document processing

10

Building a foundation for your AI strategy

- 12 Create reliable and scalable AI solutions on Azure
- 13 Create AI apps and agents at scale
- 14 Use a modern, managed app platform
- 16 Ensure that you have robust data capabilities to feed your AI
- 18 Take advantage of the wider capabilities of Azure

19

Start building the AI-powered future with Azure

Introduction

AI in business has evolved from cutting-edge, experimental technology to a core strategic asset, delivering practical, real-world benefits. Sophisticated AI-powered solutions, ranging from productivity-boosting apps to customer-facing tools and powerful agents, are no longer conceptual or experimental. They're in production and creating measurable business impact.

Enterprise-grade AI applications and agents require a comprehensive set of services, models, tools, and data that enable high performance, scalability, and security. As you consider designing, building, deploying, and operating these solutions, you must choose the best cloud provider and services to support AI, apps, and data according to your organization's needs:

→ **AI.** An AI platform helps you build, optimize, and govern AI solutions with the best models, agents, and tools for your use case. It should provide continuous observability, governance, and content safety capabilities.

→ **Apps.** Managed services apps need to provide the runtime environment, APIs, and user interfaces that allow AI apps to operate and scale and must integrate them with other pieces of your infrastructure.

→ **Data.** High-quality, well-managed data is a prerequisite for your AI apps to provide accurate, context-aware, and relevant functionality, while also helping to ensure compliance and governance.

Integrated, AI-powered developer tools accelerate time to market and put developers in charge of innovation while minimizing time spent on unproductive and repetitive tasks. Together, these components support AI solution development that is not only powerful but also reliable, secure, and aligned with business goals.

In this e-book, we explain how AI apps and agents can drive your business growth and how you can use the Azure platform's trusted tools and infrastructure to create these kinds of innovative AI-powered solutions.

Driving business value with AI

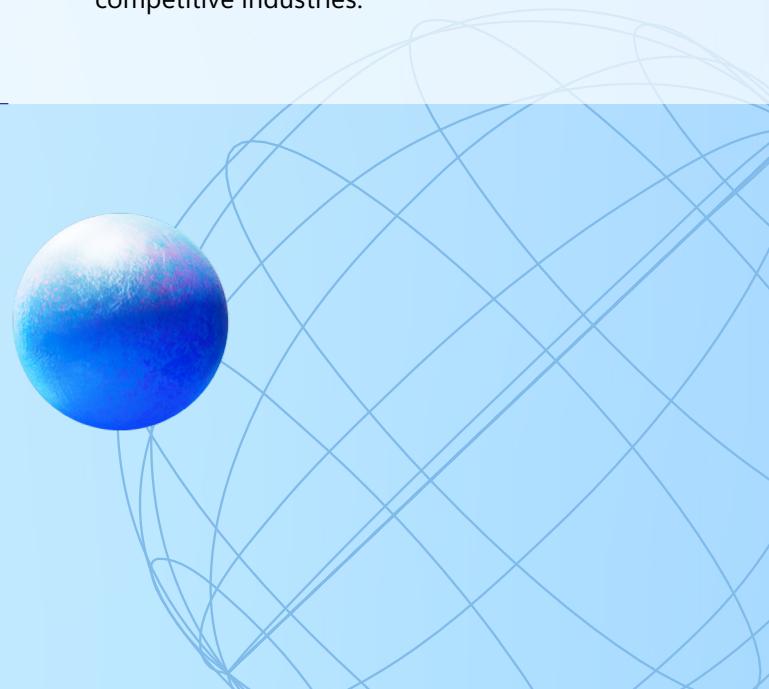
With AI-powered apps and agents, your business can deliver experiences that are more streamlined, personalized, and consistent for your customers. They can give instant support, offer tailored product recommendations, and even help users complete tasks more efficiently. Around the clock, AI solutions help reduce wait times and often anticipate customer needs based on past behavior, creating smoother and more satisfying interactions.

AI-powered solutions also offer advantages for your business operations. Internally, they can reduce costs and increase staff productivity by streamlining tasks like customer support, document handling, and data analysis. For technical and development staff, coding assistance, automated testing, and other AI-powered tools can speed up innovation and time to market for new products. For decision-makers, AI can improve business intelligence capabilities, surface insights from across the organization, and help leaders to make faster, better-informed decisions.

Key AI outcomes

Businesses generally adopt AI-powered solutions to achieve any or all of four key outcomes:

- ✓ **Enhance employee experiences.** By automating repetitive, mundane tasks, AI empowers employees to get more done and to focus on more complex, creative, and higher-value work.
- ✓ **Transform customer engagement.** AI enables more personalized and tailored customer interactions, while reducing workloads for employees.
- ✓ **Reimagine business processes.** AI is revolutionizing operations across marketing, supply chain, and finance, optimizing processes and uncovering growth opportunities.
- ✓ **Accelerate innovation.** AI can speed up creative workflows and product development, streamlining time to market and helping companies stand out in competitive industries.



AI apps and the evolution of agents

AI apps have fundamentally changed how we interact with technology. At their most basic, these apps, driven by powerful language models, can respond to questions posed in plain (or *natural*) language with useful answers based on their training data. Organizations are increasingly adopting generative AI-enabled apps for internal support, sales enablement, data analysis, and more.

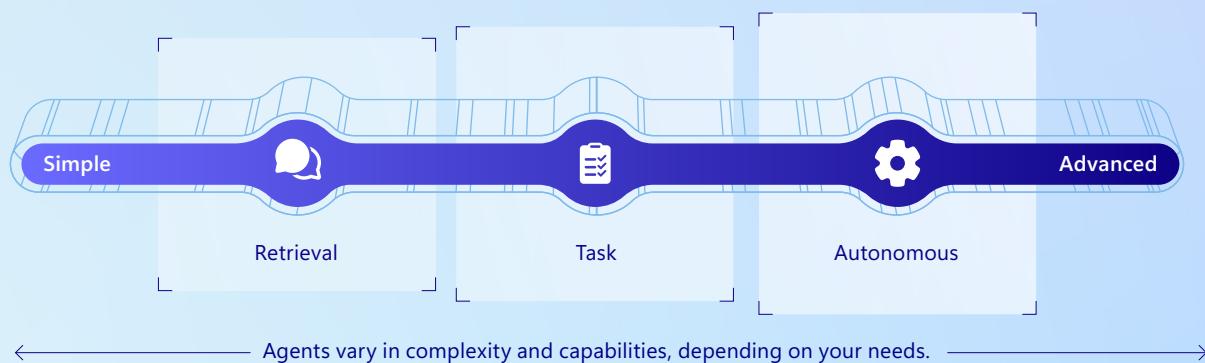
As the adoption of AI apps has grown, so have their capabilities. They can now routinely do things like make decisions based on pictures, automatically translate languages, and ingest and summarize large collections of documents. These AI-powered features are increasingly integrated directly into office productivity suites, human resources (HR) management tools, customer relationship management (CRM) systems, and

development environments. Modern AI apps are also increasingly able to perform more complex tasks, like scheduling meetings, generating reports from analytics dashboards, and assisting in software development, without explicit instructions.

This progression has laid the foundation for the emergence of AI-powered agents. Simply put, agents are systems that use AI to automate and execute business processes. They can provide a spectrum of capabilities. Augmented retrieval agents are on the simpler side of this spectrum, and this enables interaction with external data sources to improve responses. Task-based agents, connected to specific workflows or processes like document generation or facial recognition, are more sophisticated. On the more advanced end of the spectrum are autonomous agents (agentic AI systems), which can operate independently, dynamically plan their actions, and interact with other AI tools to perform specific actions.

What are agents?

Agents use AI to automate and execute business processes, working alongside, or on behalf of, a person, team, or organization.



Agentic AI: Expanding what AI can do for your business

Autonomous or semiautonomous agentic AI systems are more independent than traditional AI apps. They combine language models with goal-oriented logic to carry out complex workflows across systems and data sources, and they're capable of handling complicated behind-the-scenes logic and tasks.

Agentic AI systems consist of several components:

- **Intent recognition and planning.** The agent monitors its environment to identify and interpret user input or system triggers, using AI reasoning models to help determine appropriate tasks or goals and to break them into actionable steps.
- **Tool use and action execution.** Agents are equipped with predefined sets of tools, APIs, or plugins (like calling a REST endpoint, running a query, or accessing CRM data) for performing actions and interacting with external systems.
- **Memory and state management.** AI models are generally not state-aware. Agents, by contrast, can track conversation history, task state, and results, enabling context awareness across interactions and long-running processes.

→ **Agentic identity.** By assigning agents an identity, organizations can extend the same infrastructure, apps, and protections that they already use, leveraging familiar capabilities and policies adapted to agent needs.

→ **Reasoning and feedback loops.** Using the capabilities of language models, agents make decisions at each step, adjust based on real-time feedback, and retry or escalate if needed.

To handle more complex workflows, these types of agents can be combined into larger multiagent systems. They can assign specialized roles to different agents that can operate concurrently, collaborate effectively, and scale independently. By mimicking the dynamics of human teams and delegating tasks, sharing insights, and coordinating actions, multiagent systems can drive smarter automation across workflows like customer service, logistics, and data analysis, delivering more robust and scalable solutions.

Agentic AI systems are well-suited to handle repetitive processes and can enhance your decision-making by providing real-time data access and logic. With the ability to integrate securely into enterprise workflows and adapt to changing contexts, these agents can take on your organization's more tedious tasks, so your staff can focus on work that only humans can handle.

Example use cases for AI apps and agents

AI apps and agents can help your business automate complex tasks, enhance decision-making, increase efficiency across operations, and improve customer experiences.

Following are a few examples of how AI solutions are helping organizations meet their real-world business needs.

Assembly

Accelerate content delivery

Read the story: [Assembly Software saves law firms up to 25 hours per case with Azure](#)

Scenario: A legal tech provider accelerates legal workflows by automating manual document drafting, data entry, and case review using embedded generative AI.

AI in action: Using Microsoft Foundry, Assembly Software created an embedded generative AI solution that helps their clients automate many aspects of legal case management. Integrating OCR, data extraction, retrieval-augmented generation (RAG) via Azure AI Search, and Azure Document Intelligence in Foundry Tools to streamline legal workflows, their solution reduced legal document drafting from 40 hours to mere minutes and saves their customers up to 25 hours per case.

CARVANA

Transform customer service

Read the story: [Carvana builds efficient, seamless customer experience agent on Azure](#)

Scenario: An online automotive retailer enhances customer support by analyzing every customer interaction with AI-powered review and agent tools.

AI in action: Using Microsoft Foundry, Carvana built an AI tool that analyzes 100% of customer calls and chats, enabling faster feedback loops and higher service quality. It also created an AI assistant that helps customers throughout the buying and selling journey, reducing inbound support calls by 45% per sale. The approach has improved efficiency, lowered operational costs, and delivered a more seamless, consistent customer experience.



Optimize employee workflows

Read the story: [The ODP Corporation transforms HR, sales, and retail workflows with Microsoft Foundry](#)

Scenario: A large retail services organization streamlines internal operations by deploying AI-powered assistants across HR, sales, and in-store retail workflows.

AI in action: The ODP Corporation equipped 16,000 associates with AI assistants that improved HR response times, simplified employee training and benefits inquiries, and supported retail staff in nearly 900 stores. AI-powered tools have also reduced quote creation time from hours to minutes, saving representatives over 25 hours per quote and contributing to \$800,000 in new annual revenue.



Prevent fraud and detect anomalies

Read the story: [Kinectify fights financial crime in the gaming world with Microsoft Azure](#)

Scenario: A risk-reduction company that specializes in combating money laundering enhances its fraud detection and decision-making speed by taking advantage of AI to identify suspicious activities and patterns in real time.

AI in action: To detect potentially fraudulent activity, an intelligent application enables real-time AI-enabled monitoring of vast numbers of transactions. By making the most of predictive models in machine learning and AI algorithms, this application identifies patterns and anomalies associated with fraudulent behavior, automates risk scoring, and accelerates decision-making. This helps individual analysts to be more efficient, speeds up fraud detection and investigations, and ultimately helps minimize financial losses.



Unlock organizational knowledge

Read the story: [NFL Combine gains real-time insights for informed draft day decisions with Microsoft Foundry](#)

Scenario: At the NFL Combine, scouts and coaches need faster access to prospect data to make accurate decisions during live drills.

AI in action: The NFL was able to equip its Combine App with a real-time AI assistant that delivers natural-language insights in seconds to over 300 elite football prospects. Instead of labor-heavy manual reports, scouts now receive instant answers during drills. These tools help turn hours of post-event analysis into immediate, sideline-enabled insight of up to 570 hours. This time savings allows employees to focus on advanced analysis and creative tasks, enhancing productivity and operational efficiency.

LayerX

Automate document processing

Read the story: [AI Workforce built on Azure helps LayerX client save 570 hours of processing docs](#)

Scenario: A software as a service provider for fintech firms was able to automate error-prone manual processes and reduce inefficiencies so their clients can process documents more efficiently.

AI in action: LayerX developed an AI-powered service to help its clients automate complicated document processing and see annual savings of up to 570 hours. This time savings allows employees to focus on advanced analysis and creative tasks, enhancing productivity and operational efficiency.

Building a foundation for your AI strategy

AI-powered apps and agents offer many ways to add powerful intelligence capabilities for your business. You also have many choices for how you build AI solutions. However, stitching together AI models, agents, and tools with different application-hosting services and data infrastructure can be tremendously complex, often resulting in steep learning curves and costly, time-consuming development efforts.

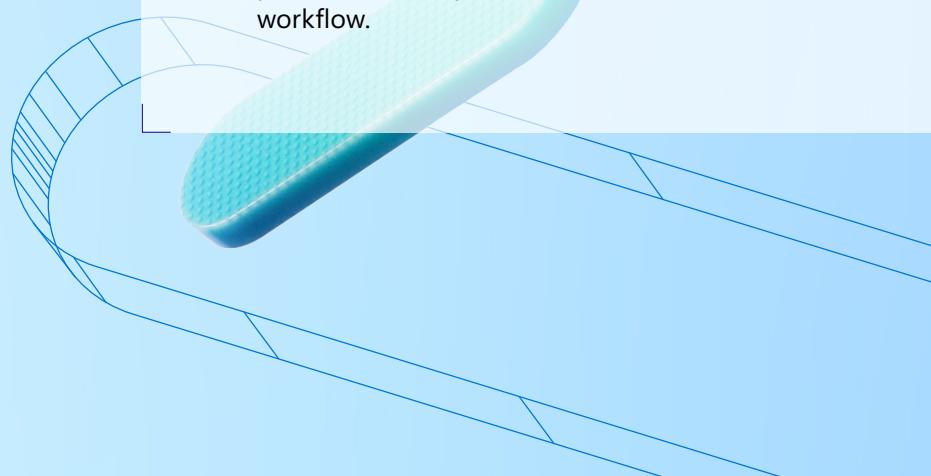
Enterprise AI applications do have some common requirements that impact your cloud technology choices, including:

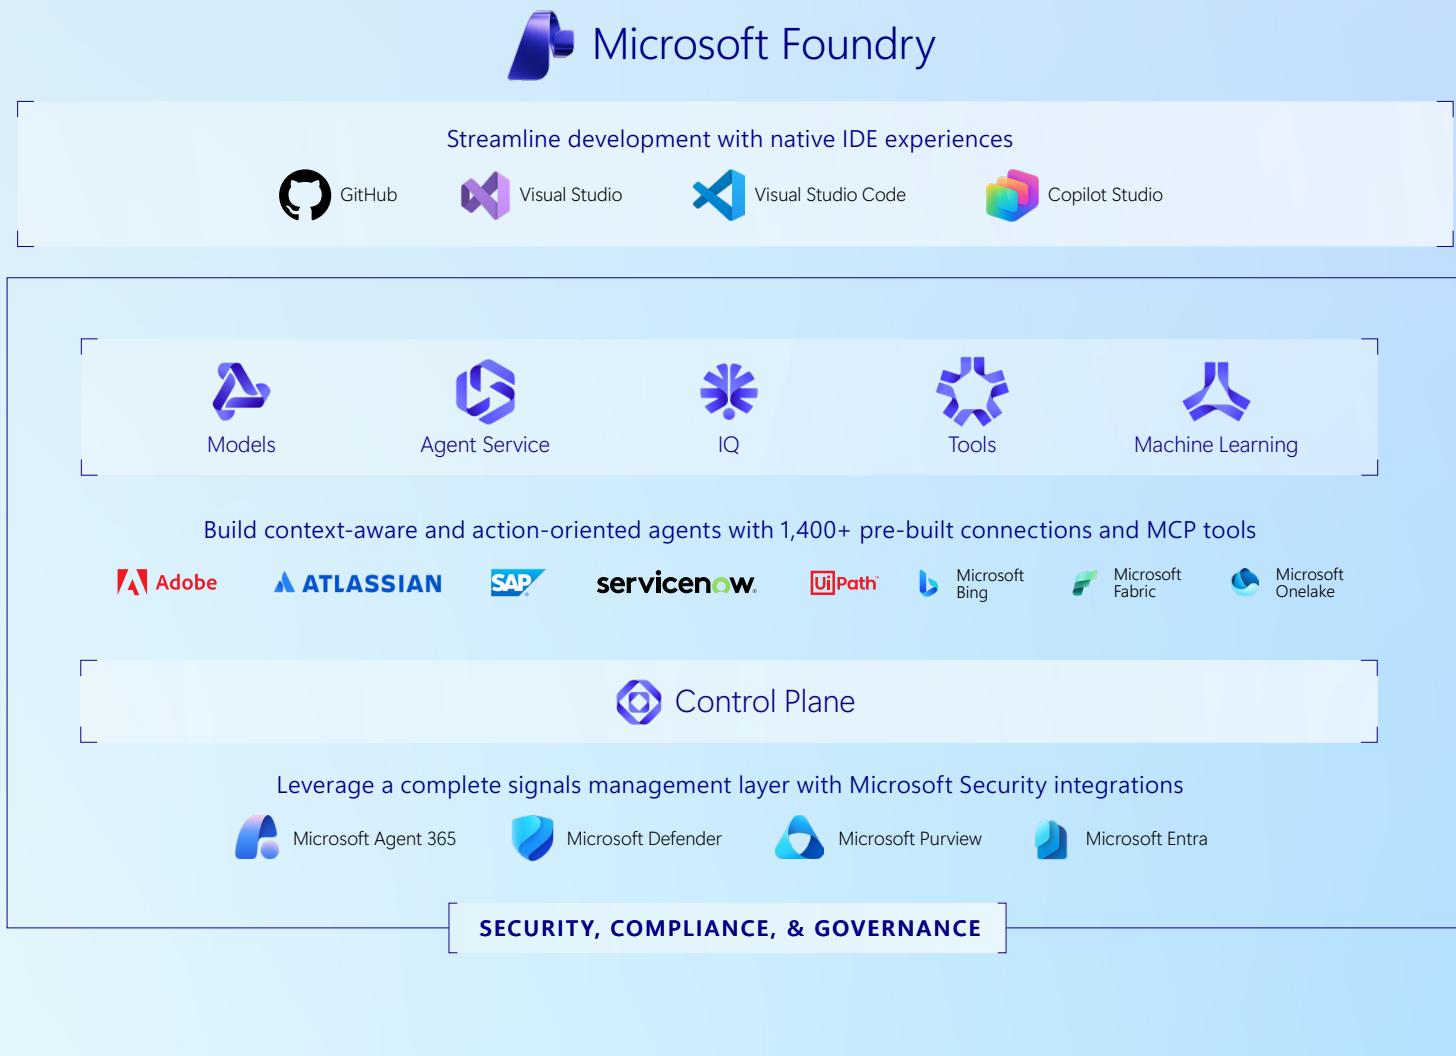
- Robust AI services that support the full model lifecycle, including training, deployment, monitoring, and retraining, often at scale.
- Highly scalable, secure, and resilient cloud-native infrastructure to support continuous availability and global reach.
- Seamless access to high-quality, governed, and real-time data sources.

The benefits of an integrated AI application platform

An integrated AI application platform is a comprehensive environment that combines tools and services for designing, customizing, and managing AI-powered applications from end to end. There are many benefits to creating your AI apps on these types of platforms, including the ability to:

- ✓ **Accelerate time to market.** Go from idea to production faster with a streamlined platform that can reduce development timelines with Azure's built-in resiliency features.
- ✓ **Iterate faster with feedback.** Test, deploy, and scale AI applications rapidly, refining them based on user insights to help ensure continual improvement.
- ✓ **Securely integrate AI with enterprise data.** Connect AI models to enterprise data sources for advanced personalization, search, and analytics, while maintaining robust security.
- ✓ **Empower developers with AI tools.** Enable developers to build AI apps faster using their preferred development tools for a seamless workflow.





Operationally, you need to ensure that your AI apps can reliably access the data and services they need while deploying and scaling quickly to meet demand. From a business perspective, you also need to manage costs and make sure AI solutions meet your business goals. Additionally, you may be subject to

regulations, such as data sovereignty, privacy laws, bias mitigation, and AI safety standards.

To address these issues, you need an integrated application platform that provides the hosting, data, and services you need to quickly and efficiently build trusted, scalable, and performant AI solutions.

Create reliable and scalable AI solutions on Azure

Azure offers an integrated AI application platform that brings together a comprehensive suite of cloud services to support the full lifecycle of designing, customizing, and managing AI-powered applications. It streamlines app development, data management, and AI models, agents, and tools with a unified environment, so you can create applications more efficiently.

High-performance AI applications depend on the app-hosting, data, and AI capabilities that your underlying infrastructure provides. Managed cloud AI services are critical because they support overall scalability, performance, development speed, and long-term maintainability. The right services provide a choice of models for every use case, including pretrained and open-source models.

Microsoft Foundry is the unified, interoperable platform that empowers enterprises, startups, and software development companies to build, optimize, and govern AI innovation at scale. It provides easy access to powerful tools, such as Foundry Models, Foundry Agent Service, Foundry Tools, Foundry IQ, and Foundry Control Plane, all from a single platform.

A commitment to building trustworthy AI

Microsoft is committed to ensuring that its AI services are secure, safe, private, and aligned with ethical principles. These efforts are organized around three main pillars:

- ✓ **Security.** The [Microsoft Secure Future Initiative \(SFI\)](#) reflects our company-wide commitment to security, with secure-by-design principles implemented throughout the Microsoft internal development lifecycle. It also includes internal governance improvements, accountability, and a dedicated security workforce of over 34,000 engineers.
- ✓ **Safety.** The [principles of responsible AI at Microsoft](#), established in 2018, guide the company in safely building and deploying AI, including properly developing, testing, and monitoring systems to avoid undesirable behaviors, such as harmful content, bias, misuse, and other unintended risks.
- ✓ **Privacy.** Another critical Microsoft priority is to help ensure that customer data is protected and compliant with Microsoft privacy principles, as itemized in [Microsoft Trust Center: Security and privacy in the age of AI](#). These include user control, transparency, and legal and regulatory protections.

Microsoft Foundry directly supports AI safety by integrating comprehensive tools and practices that align with principles of responsible AI. Products like [Content Safety in Foundry Control Plane](#) and features like [Risks & Safety monitoring](#) can help to identify issues with AI output and to prevent problematic output or hallucinations from reaching end users. You can also configure custom guardrails and filters to meet your organization's specific needs.

Create AI apps and agents at scale

A unified AI platform makes building AI apps and agents easier by bringing data ingestion, model selection, AI app and agent orchestration, deployment, monitoring, and governance into one integrated platform. By integrating these services into developer workbenches, your development teams can work in a consistent environment with built-in access to AI, compute, and data services—without needing to context switch and stitch together disparate tools.

Microsoft Foundry brings the full capabilities of Azure to your AI projects. It includes powerful models, such as Azure OpenAI in Foundry Models, support for open frameworks and protocols, and native integrations with Azure and the broader Microsoft Cloud to help developers accelerate innovation.

Microsoft Foundry can also take advantage of GitHub and Azure DevOps for streamlined development and deployment, inherits global scalability and reliability from the worldwide Azure infrastructure, and helps to ensure enterprise-grade security, governance, and compliance through services like Microsoft Purview, Microsoft Defender, and Microsoft Entra ID. These integrations make Microsoft Foundry a powerful hub for building, optimizing, and governing trustworthy AI applications.

You can use Microsoft Foundry to help you manage data integration, model training and customization, deployment orchestration, and experimentation. It also provides enterprise-grade governance and management, and it integrates AI safety tools, helping to ensure that your AI apps and agents are secure and compliant. By unifying data, models, and operations into a single platform, Microsoft Foundry radically simplifies the process of creating intelligent agents and integrated AI solutions that deliver real-world business impact.

The recently released [Foundry Agent Service](#) offers powerful new capabilities so developers can more easily build and operate intelligent agents. This enterprise-grade service enables the customization, deployment, and scaling of context-aware, multturn agents and multiagent apps capable of orchestrating complex workflows across diverse systems. With built-in security, interoperability, and orchestration capabilities, the service empowers organizations to accelerate innovation through AI-driven automation. Implementations by companies such as [Fujitsu](#), [NTT DATA](#), and [YoungWilliams](#) demonstrate the platform's potential to deliver measurable impact across industries.



Empowering your developers with AI

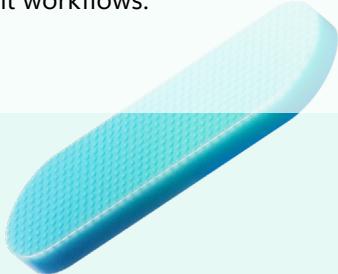
AI technology is not just a feature that's included in end-user apps. It also can help your developers create useful applications that solve complex business problems faster and more efficiently.

Within the Microsoft AI development ecosystem, developers can choose tools and environments to suit their skills and preferences and to reduce integration complexities. AI assistance, including the powerful GitHub Copilot, interoperates smoothly with core development tools. Microsoft Foundry can also simplify your developers' work by offering broad support for an array of programming languages, models, and AI frameworks.

With Microsoft Foundry, developers can integrate available AI models into their apps by using SDKs for Python, C#, Java, and JavaScript. Developers can also take advantage of many popular AI frameworks, including LangChain, LlamaIndex, Semantic Kernel, and AutoGen.

Observability in Foundry Control Plane features integrate with your development cycle in multiple ways, further improving your dev processes. This can include automated testing and debugging, smart documentation, continuous integration/continuous delivery (CI/CD) optimization, and project management and planning.

By embedding intelligence throughout your development lifecycle, Microsoft Foundry can help reduce friction, shorten feedback loops, and create more seamless and productive development workflows.



Use a modern, managed app platform

Depending on your use case, Azure offers an array of managed app platforms to choose from that can support scalable, high-demand applications. These include Azure Functions for serverless applications, Azure App Service for traditional web apps, and Azure API Management for API-based services that support your users' other applications. Azure Essentials reinforces API governance with policy patterns and identity controls so integrations remain resilient and secure across teams.

However, many modern AI apps and agents are built around containerized microservices architectures. This allows these solutions to be flexible, scalable, modular, and better suited for creating AI-powered solutions.

On Azure, Azure Kubernetes Service (AKS) serves as the cloud-native-app-hosting backbone that you need to deploy, scale, and manage AI-powered, containerized applications and agents. AKS offers key features that AI-powered applications require, including:

→ **Scalable infrastructure.** AKS can automatically scale compute resources up or down, based on need, so it's ideal for handling demanding enterprise AI workloads. AKS Automatic further simplifies Kubernetes for all teams, integrating with Azure services and tools. It enables reliable, secure, and performant deployment, even for first-time Kubernetes users.

→ **Microservices and API management.**

Because it's Kubernetes-based, AKS supports containerized microservices architectures. Each component of your AI app can be deployed and managed independently as a container, so the system is modular and resilient.

→ **Integration with AI and machine learning pipelines.**

AKS integrates well with Microsoft Foundry for model training, deployment, and monitoring. You can deploy machine learning models as containerized inference endpoints, retrain them in the background, and manage their lifecycle in production.

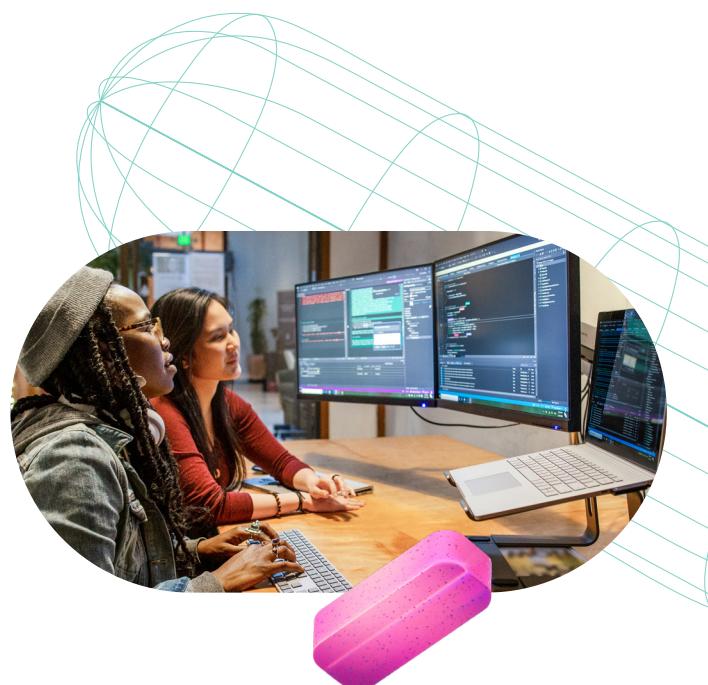
→ **Enterprise-grade security and governance.**

AKS includes built-in capabilities for identity management, role-based access control, network policies, and integration with Azure Policy and Microsoft Defender for Containers, helping you to be sure that your AI workloads meet enterprise security and compliance standards.

→ **DevOps readiness.**

AKS can integrate into your organization's DevOps processes and into features such as CI/CD pipelines, canary deployments, and infrastructure as code tools like Terraform and Bicep.

For those who want to leverage the flexibility and power of Kubernetes without the overhead operations, AKS Automatic offers a powerful solution for hands-off cluster management and intelligent automation, while still maintaining access to the underlying API. In addition to AKS, Azure Container Apps is an alternative service that offers fully managed serverless container hosting. Azure Container Apps has also recently added support for serverless GPUs, which can give you access to GPUs on demand without the need to manage underlying infrastructure.



Ensure that you have robust data capabilities to feed your AI

AI-optimized databases are critical to storing, managing, and accessing the vast amounts of data that power intelligent systems. AI models rely on high-quality, well-organized data, both structured and unstructured, for training, validation, and real-time inference.

Scalable data services help to ensure that data can be ingested from various sources, transformed, and retrieved efficiently, even under heavy workloads. They also support essential regulatory compliance and governance capabilities. Azure offers several powerful data platforms that you can integrate with your AI apps and agents.

Azure SQL Database

Azure SQL Database provides a resilient, fully managed relational database service that can store and manage any structured data your AI applications rely on. SQL Database has integrations with Microsoft Foundry and Foundry IQ. With AI search and SQL, you can index and enrich data in SQL stores with powerful semantic and vector searches, so your AI applications can combine relational data with natural language, keyword, and embedding-based retrieval to deliver more accurate and contextually relevant results.

With the Hyperscale tier of Azure SQL Database, you can accommodate data growth associated with modern applications without compromising performance in one cost-effective solution.

Azure Database for PostgreSQL

Azure Database for PostgreSQL offers a powerful, fully managed database service built on top of the popular PostgreSQL open-source database software. Although PostgreSQL is a relational database technology by default, integrated support for the pgvector PostgreSQL extension and newly released DiskANN search algorithm allows you to use Azure Database for PostgreSQL as both a relational and vector data store.

With this flexibility, you can build, manage, and scale your entire AI application in a single data service, significantly simplifying your architecture, reducing latency, and ensuring data consistency.

Azure Cosmos DB

For your highest-scale and most demanding AI applications, Azure Cosmos DB can provide all the power you need. It offers a flexible, enterprise-grade data platform that meets the demanding requirements of modern AI systems, including real-time intelligence, global reach, and seamless integration across an evolving ecosystem of AI tools and technologies. These capabilities make Azure Cosmos DB a foundational component for building scalable, responsive, and context-aware enterprise AI systems.

Azure Cosmos DB features several key capabilities that help you effectively build powerful AI solutions:

- **Built for AI performance at scale.** AI apps need fast access to data. Azure Cosmos DB offers single-digit-millisecond latency for both reads and writes so your AI apps and agents can respond in real time, whether they're generating recommendations, navigating complex user interactions, or executing multistep tasks. With automatic scaling, Azure Cosmos DB adjusts to unpredictable workloads, handling everything from sudden surges in chat usage to fluctuating inferencing demands.
- **Built-in vector search.** Azure Cosmos DB supports document (NoSQL) and native vector data, enabling seamless storage and retrieval of relevant data in the same database at scale. This flexibility is essential for powering RAG systems, knowledge-grounded agents, and other emerging AI patterns.



→ **Real-time AI pipelines.** AI is only as good as the data feeding it. Azure Cosmos DB enables real-time data ingestion and querying, making it ideally suited for the personalization, session memory, chat history, and feedback functionality that enable continuous machine learning. It also integrates deeply with Microsoft Foundry, Azure OpenAI in Foundry Models, Azure Functions, and Microsoft Fabric, letting you build event-driven, inference-ready pipelines without friction.

→ **Secure and available AI.** Azure Cosmos DB comes with enterprise-ready security, compliance, and SLA-backed 99.999% availability, offering encryption at rest and in transit, multiregion and per-partition automatic failover, that guarantee the availability and consistency of your AI workload.

This combination of performance, flexibility, and integration makes Azure Cosmos DB the ideal-fit data backbone for AI-powered solutions.

Azure Databricks integration

Having your AI apps and agents ingest and work with truly large datasets can be challenging. In addition to traditional database services, the Microsoft Foundry Azure Databricks connector can enable high-performance integrations between your AI apps and data sources managed through Azure Databricks. This integration allows your AI solutions to take advantage of Apache Spark capabilities and integrated tools that provide seamless big data processing, machine learning, and data engineering capabilities.

Take advantage of the wider capabilities of Azure

The services discussed in the previous section are all built on the powerful, globally distributed infrastructure of Azure. Modern AI applications, built around Azure services, can benefit from the many advantages that the Azure platform offers, including:

- **Service availability.** Most Azure services offer redundant systems and automated failover capabilities that help keep your workloads operational. With over 300 datacenters, the multiregional deployment options of Azure help maintain business continuity.
- **Global resiliency.** Azure datacenters are distributed across more than 60 geographical regions around the world. This worldwide Azure infrastructure delivers consistent performance and offers region-specific availability zones to increase resilience.
- **Validate Designs.** The Azure Well-Architected reliability pillar and Azure Essentials checklists use monitoring/chaos testing where appropriate.
- **Performance SLAs.** Microsoft guarantees up to 99.999% uptime across all services underpinned by Azure Essentials guidance.
- **Dynamic scaling.** Azure services emphasize intelligent capacity management that adjusts to workload demands in real time.

→ **Support framework.** Dedicated technical resources handle critical issues with strictly defined SLAs, and regular service health updates and maintenance notifications keep you informed.

Build AI apps faster with GitHub Copilot

[GitHub Copilot](#) is an AI-powered coding assistant that helps your developers write code faster by generating real-time suggestions, functions, and entire workflows directly in their editor. It seamlessly integrates with the Azure ecosystem so developers can quickly move from code ideation to deployment.

Copilot integrates with Microsoft Foundry by streamlining code generation and task automation in Microsoft Foundry projects, helping your developers design, customize, and manage AI solutions more efficiently using Azure-native tools and infrastructure. It also helps with resource configurations, deployment scripts, and creating connections to services such as Azure Cosmos DB, Azure Functions, and AKS.

The new [GitHub Copilot coding agent](#) takes these capabilities further by acting as an autonomous teammate that can own and execute entire development tasks. Instead of just suggesting code snippets, it can be assigned GitHub issues, spin up dev environments, clone repos, make code changes, run tests, and submit pull requests. For your AI apps, this means faster iteration cycles, easier onboarding of new features, and automation of common tasks, like model integration, prompt engineering refactors, and API wrapper development. With the coding agent, teams can spend less time on setup and more time on experimentation and innovation.

Start building the AI powered future with Azure

Start building the AI-powered future with Azure. Whether you're creating powerful customer-facing products, building AI apps to boost team productivity, designing custom agents to automate complex workflows, or empowering your developers with best-in-class tools, Azure gives you the integrated platform, flexibility, and trust to confidently embrace AI technologies.

Explore what's possible with Azure services, like Microsoft Foundry, Azure Cosmos DB, and AKS, and start building the intelligent solutions that can define your future.

- Learn more about [building AI apps and agents](#) and [Microsoft Foundry](#).
- Access Microsoft experts and funding with [Azure Accelerate](#) and explore [Azure Essentials](#) for technical guidance and resources
- Talk to an [Azure expert](#) about your organization's needs