


# Six Strategies for AI Implementation

How to apply responsible AI  
practices on Azure



# Six Strategies for AI Implementation

3 /

Foundation for responsible AI strategy

11 /

Strategy 3: Establish end-to-end governance and security tools

17 /

Strategy 6: Enhance customer experience

6 /

Strategy 1: Rapid experimentation to overcome uncertainty

13 /

Strategy 4: Bring it all together with a responsible AI framework

20 /

Conclusion

9 /

Strategy 2: Integrate AI with unified data to simplify implementation

15 /

Strategy 5: Enable responsible AI innovation in large enterprises

20 /

Next steps

# Foundation for responsible AI strategy

In today's business landscape, companies are compelled to develop and introduce new products, optimize their operational processes, and adhere to regulatory and compliance requirements. This scenario plays out on the dynamic global stage, where innovation plays a pivotal role in shaping the overall business outcome.

Given this level of exposure, AI systems require governance to support more fair, accurate, and responsible decision making. The widespread adoption of generative AI has only intensified this necessity. Today, a vast majority of enterprises face the dilemma of balancing disruption and innovations brought about by modern AI tools with the responsibility and control mandated by business practices and regulations. Developing a modern AI strategy, supported by the corresponding data and infrastructure, is now a priority for most businesses.

Using the full potential of AI tools requires specialized skillsets, well-defined policies, streamlined processes, and the integration of relevant tools and technologies across the vast data landscape of modern organizations. Microsoft Azure provides a wide range of data and AI services that work together to help customers implement AI responsibly and effectively, including [Azure AI Studio](#), [Azure OpenAI Service](#), and [Azure Machine Learning](#).

By integrating responsible AI principles and practices into their AI strategy from the beginning, organizations can build a stable foundation to scale AI successfully. Microsoft recommends six principles for AI development and implementation—fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. These principles serve as the cornerstone in building a corporate AI standard.

Monitoring and auditing, reporting and compliance processes, practices, and tools should be built with responsible AI principles in mind and consistently measured against this set of principles.

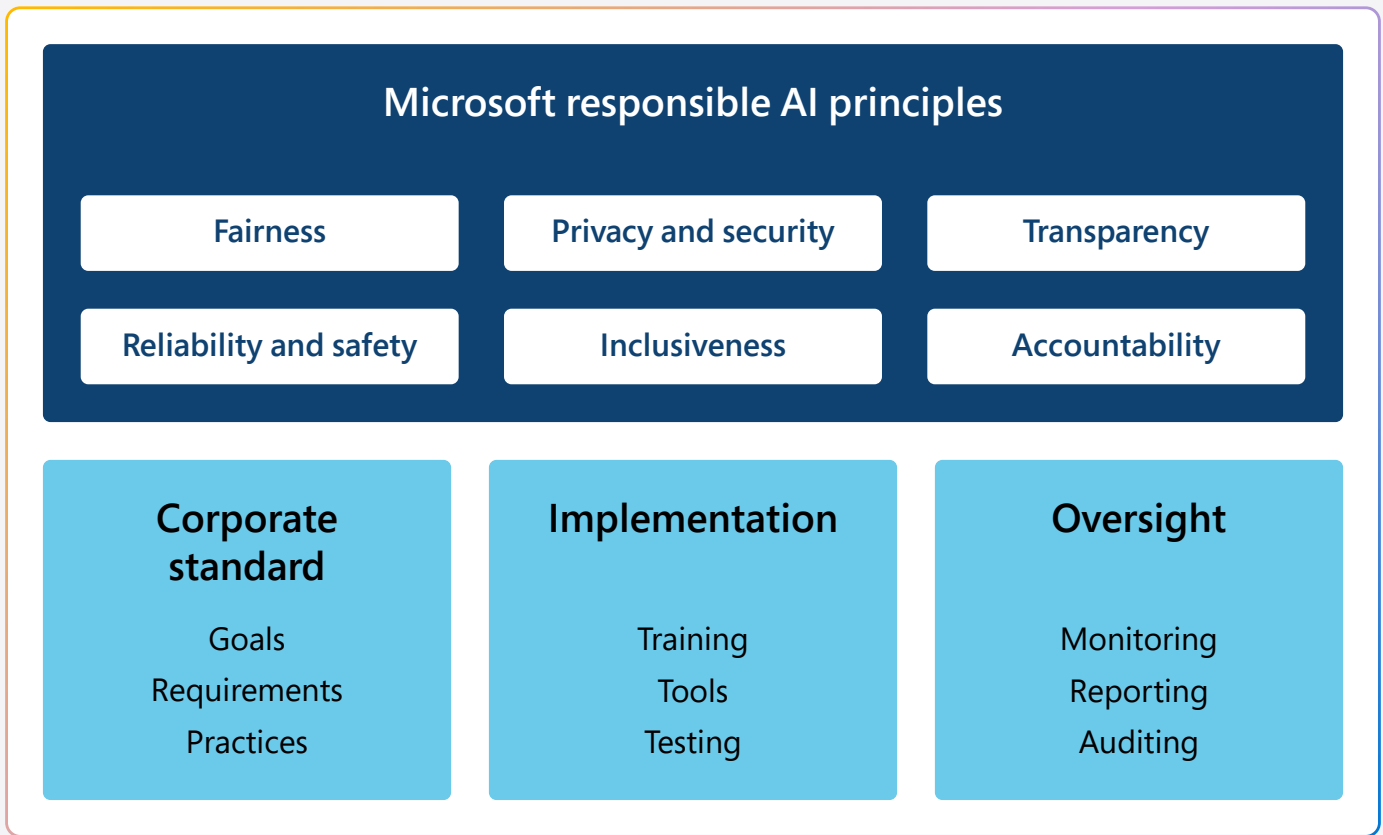


Figure 1: Microsoft's approach to AI governance at scale

# Elements of a successful AI strategy

There are six strategies for AI development and implementation. These strategies complement each other and relate to the six principles recommended by Microsoft. Organizations can scale up their AI adoption efforts and maximize the value offered by AI solutions using Microsoft Azure as the most advanced home platform for responsible AI and supporting workloads.



Figure 2: Six strategies for AI implementation

## Strategy 1:

# Rapid experimentation to overcome uncertainty

AI adoption at scale is a relatively new endeavor for most businesses. Many organizations find themselves in the early stages of acquiring the necessary skills and establishing a foundation for their AI practice. In the absence of well-defined best practices and proven use cases, it's crucial for organizations to experiment with use cases and technologies quickly. To fast track experimentation, organizations should:

1. **Build multi-disciplinary teams to foster innovation:** AI applications, unlike traditional methods, prioritize domain expertise and diverse experiences over predefined deterministic processes. Shortcomings or biases in AI applications can be identified by individuals with expertise loosely related to the subject of AI applications. However, it's helpful for such teams to consist of individuals with extensive expertise in application development, data science, and data

engineering. By cultivating a culture that encourages curiosity, teams can effectively translate cutting-edge technologies into innovative practices and products. These teams necessitate an environment, data, and processes that are conducive to experimentation in a controlled and secure manner.

2. **Establish a controlled experimentation environment:** The process of selecting appropriate data and establishing processes to control the experimentation environment requires collaboration between business users and information security specialists. The best way to strike a balance among all parties is by agreeing on a set of initial use cases to serve as a foundation for developing innovative applications. In most cases, experimentation environments are implemented in the same way as development and testing environments.

Experimentation environments for AI use cases are data dependant, and often use actual company data even during experimentation. Such environments include data services, such as Azure Data Lake Storage Gen2, Azure Cosmos DB, or SQL Database, as well as services from the Azure AI toolset, such as Azure Machine Learning, Azure OpenAI Service, or [Azure AI Search](#). It may also include a variation of application deployment environments, such as Azure Kubernetes. All the preceding services should be governed and protected according to the security baseline of each service.

- 3. Streamline application development and data provisioning using feedback loops:** Establishing internal processes for application development and productization, selecting and provisioning data within the experimentation platform, monitoring data flow, and logging user behavior empower application owners to improve the performance

of applications, and the relevance of data and model outputs, so they can fine-tune use cases based on feedback on the applications and models' performance as well as user experience. This feedback enables organizations to promptly address issues and gather valuable insights into user satisfaction.

- 4. Secure storage and compliance:** An environment with secure and compliant storage and tools for comprehensive governance is essential for organizations to monitor, track, and manage the entire lifecycle of AI applications. This safeguards sensitive data and AI systems, protecting them against unauthorized access and potential data breaches.

By implementing these measures, organizations can produce meaningful results that align with their strategy and are relevant to their unique business context. When results are designed with scalability and reusability in mind across the organization, these outcomes serve as a foundation for establishing company-wide best practices.

Azure offers these capabilities on a global scale and facilitates streamlined interactions with various other cloud and technology providers. Azure AI Landing Zones is a good starting point for organizations as it provides predefined tools and scripts to help organizations jumpstart this process.

Azure AI Landing Zones serve as a robust foundation for deploying advanced AI technologies, including models such as OpenAI's GPT-4. This offering is essentially a reference architecture supported by a set of artifacts, designed to help AI enthusiasts swiftly establish solid and secure groundwork for AI deployments.

**Azure AI Landing Zones** encompass a comprehensive guide on how to manage critical aspects of infrastructure, including security, networking, and monitoring. These guidelines ensure the success and security of AI projects.



## Strategy 2:

# Integrate AI with unified data to simplify implementation

The ability of AI to summarize and aggregate information and extract insights from vast datasets is invaluable. The integration of AI with data empowers customers to interact directly with datasets, thereby enhancing their overall experience. Within organizations, the unification of data leads to the development of new, innovative products for customers and adds value to existing business operations.

Unified integration of AI and applications with a comprehensive data estate is pivotal for the successful adoption of AI capabilities at scale. To accomplish this, organizations must implement measures to overcome common challenges, such as:

- **Data quality assurance:** The quality of an AI model directly correlates with the quality of the underlying data. Low-quality, inaccurate, or ambiguous data can compromise the results generated by AI models.

- **Data integration across multiple environments:** Integrating data across multiple environments can be difficult for organizations, especially when these data sources are dispersed across different platforms, such as public clouds, SaaS applications, and on-premises infrastructure. Fortunately, Microsoft Azure provides a comprehensive set of capabilities designed to enable organizations to use data from these diverse environments securely and consistently, regardless of its location.

Data services from Azure, including Azure Data Factory, Azure Arc, and Microsoft Fabric OneLake, are equipped with powerful features tailored for hybrid environments. These services facilitate the integration and management of data across different platforms, streamlining data workflows

and enhancing data accessibility. Azure plays a pivotal role in helping customers utilize their data in a compliant and secure manner, along with providing controls to prevent unintended data sharing or exposure. Azure provides the necessary tools and safeguards the data by maintaining data privacy and security across the data estate.

- **Proficiency with data engineering:** An AI ecosystem employs various tools and techniques for data storage, retrieval, and indexing. This necessitates engineering capabilities to transform data into an optimal format for further indexing, using appropriate tools such as Microsoft Fabric, Azure Cosmos DB, and Azure AI Search.

From an organizational standpoint, an ideal approach to link AI with unified data is through data products. Data products are secured, governed, and controlled entities within the data domain that can be

shared within an organization and adhere to specific data product contracts. For many organizations, adopting a data mesh concept and linking these efforts with AI applications is a logical progression.

Once data products are developed and put into production, they serve as a trusted data source for new AI use cases. This empowers organizations to maintain control and strategically plan the integration of AI with existing and new data products. This approach optimizes an organization's budgets, and allows for the efficient reuse of efforts and resources while modernizing the data estate to facilitate AI adoption.

Azure offers a comprehensive suite of capabilities to address these challenges effectively. With Microsoft Fabric and Microsoft Purview, organizations can simplify the design and implementation of such systems.

## Strategy 3:

# Establish end-to-end governance and security tools

AI signifies a paradigm shift in how individuals interact with machines and applications, setting itself apart from the traditional deterministic applications that rely on hard-coded logic and rules. AI introduces significant variability in outputs, emphasizing the importance of responsible application development, testing, and operations.

In practice, this involves transforming traditional security and data governance practices through the adoption of new tools and processes. These measures must support adherence to organizational principles such as accountability, transparency, inclusiveness, and fairness. This multifaceted effort encompasses people, processes, technologies, and culture across the entire organization.

To effectively navigate this varied endeavor, organizations typically adopt an incremental approach, both vertically and horizontally:

- **Vertical progression:** This involves establishing robust security and compliance governance, followed by data governance, with AI governance as the final piece of the puzzle.
- **Horizontal progression:** Organizations progress through use cases or data products, ensuring that the governance framework not only exists but is effectively applied in application development.

From a security and compliance perspective, organizations can rely on trusted Microsoft tools such as Microsoft Sentinel, Microsoft Purview, and Azure Policy. These tools support the foundational principles of security and governance, allowing organizations to safeguard their data sources, development environments, data engineering pipelines, and applications. This process is notably streamlined when using Azure portfolio of AI services, which operate within Azure subscriptions and are tightly integrated with other Microsoft services.

Responsible AI practices rely on frameworks and alignment across an organization. These practices require practical tools to support the people and processes for AI innovation. The Microsoft responsible AI framework and toolset are constantly evolving, keeping pace with rapid changes in technology. With its comprehensive vision for the future of responsible AI governance, Azure not only continually introduces tools to address specific AI governance needs but also integrates them within the broader Azure stack, including data, applications, and security services.

Microsoft Azure provides a comprehensive set of tools to support responsible AI. These include open source resources such as the [Responsible AI Toolbox](#), technologies such as the [Azure Machine Learning Responsible AI dashboard](#) and [Azure AI Content Safety](#), and content filters inside Azure OpenAI Service. Microsoft also provides information and guidelines on how to responsibly use [Azure AI services](#), along with [Transparency Notes](#) to help organizations understand how the technology works and the choices system owners can make that influence system performance and behavior.

## Strategy 4:

# Bring it all together with a responsible AI framework

A responsible AI governance framework brings together the responsible AI principles that support policies, similar to the Microsoft Responsible AI Standard. It encompasses resources for training, tools, and guidance that further advance responsible AI practices. Oversight of the process and procedures for the responsible AI governance framework allows organizations to achieve consistency in goals and implementation across the organization.

The responsible AI framework should apply throughout the AI lifecycle starting with the design and development phase. Organizations can begin with analysis of stakeholders, benefits, and potential risks for a use case. In order to help with such an assessment, Microsoft provides the [Responsible AI Impact Assessment Template](#), serving as an example for organizations to develop their own frameworks. Once the impact assessment is underway, it can be used

to inform considerations across design, development, testing, and deployment. It may also be used later in the lifecycle to establish correct management, and governance processes and practices.

It is crucial that responsible AI practices, along with their respective procedures are followed during the development phase. Organizations can utilize tools to assess data and ensure responsible AI practices are respected during model training and testing, while operationalizing and monitoring these practices through MLOps and LLMOps methodologies. To align AI-powered applications and processes with responsible AI practices, Microsoft provides a suite of tools such as the [Responsible AI scorecard \(preview\) in Azure Machine Learning](#), which can be employed to define and establish organization-specific frameworks for responsible AI.

Effective oversight, which includes monitoring, auditing, and reporting is part of the responsible AI governance framework and modeled after an Enterprise Risk Management framework. Consistent monitoring and oversight of AI systems, models, and data is key to managing safety and reliability.

Standing up a responsible AI governance framework (or even tailoring an existing framework to an organization's needs) is a collaborative effort that involves multiple functions within an organization. It requires finding a balance between

the interests and needs of different stakeholders. This framework should align with enterprise policies and the operational model, offering guidance through every stage of the product and application lifecycle. One way to start developing this framework is to get acquainted with the [Microsoft Responsible AI Standard](#). By using Microsoft Responsible AI Standard as a learning point, organizations can get a head start on building their own tailored version of responsible AI standard, policies, and procedures.

## Strategy 5:

# Enable responsible AI innovation in large enterprises

Implementing large-scale changes at the enterprise level requires more than just process and policy adjustments; it requires a significant cultural shift. It is well known that culture transformation can only be achieved via people.


To harness the potential of AI for positive outcomes, organizations must cultivate a culture that aligns with the following approach:

- **Invest in talent:** In the modern world, organizations have to expand their view on AI and its role way beyond technology changes. This requires special talent profiles looking at AI use cases and process automation through the lens of a socio-technical approach where technology, culture, processes, and people are equally important. These talents require significant diversity of experiences and viewpoints to help build responsible AI solutions.
- **Embed accountability of responsible AI in every role:** Responsible AI touches every aspect of the organization and requires careful handling and proper governance. As AI becomes a part of everyone's job, each individual has a role to play in fostering responsible AI practices, leading to more innovative AI use cases.
- **Develop knowledge and skills:** Influenced by continuous advancements in AI, every role within the organization will experience a transformation in skills and responsibilities. Every employee needs to understand the transformation in strategic and tactical terms, including how processes and tools will change and what effect it will have on daily work and outcomes. Initiating skill-up programs, such as on-demand training, in-person workshops, and supporting guidance documents, as well as knowledge of the organizational units and company as a whole, is vital.

It allows employees and organizations to continue to retain ownership over processes, tools, and policies while keeping full control of the outcomes.

- **Encourage open feedback:** Invite feedback from individuals engaged in learning or utilizing tools and processes. Cultivate an environment where feedback is not only sought after but also highly valued.
- **Enable information flow:** Facilitate free exchange of experiences and information throughout the organization. This can be achieved through communities, internal social networks, sharing groups, and other channels.

However, the most intricate challenge lies in ensuring that this feedback translates into actionable insights and drives tangible change. Actionable feedback should serve as a catalyst for refining processes, policies, and tools to align with the evolving needs and objectives of an organization. This positive feedback loop serves as a pillar for successful, adaptable, and ever-evolving organizational growth.





## Strategy 6:

# Enhance customer experience

Most enterprises embark on their AI journey by focusing on internal use cases. From a user satisfaction perspective, the most successful use cases often involve addressing questions outside the domain of expertise, such as providing technical support for marketing specialists or summarizing business logs of applications. Another category of use cases pertains to direct customer interactions, which must be handled in a controlled and secure manner. Therefore, it is ideal for initial use cases to avoid using confidential or personal data, tolerate some margin of errors or inconsistencies, and offer immediate value to a company's employees or customers.

Some examples of use cases for AI and large language models (LLMs) within enterprises include:

- **Automating call centers:** AI can be used to automate responses in call centers, improving the efficiency of agents via intelligent search over the knowledge base and previous communications. Copilots for call center agents are already improving the quality of interactions with customers in multiple organizations. Another valuable use case is the summarization of conversations at scale, performing analysis on high-priority issues, and sentiment analysis in conjunction with the topic and context of the situation.
- **Replacing enterprise chatbots:** Traditional chatbots can be replaced or enhanced by LLMs, allowing more natural and context-aware interactions with both internal and

external users. For internal users, such chatbots can be used to trigger actions and optimized for working with enterprise-specific domains. Internal LLM-powered enterprise chatbots can evolve into a major interface and tool for efficiently handling tedious tasks. However, for external users, chatbot functionality is typically limited to providing basic information on tools, services, and procedures.

- **User-facing applications:** Adding LLM capabilities to user-facing applications can enhance user experience by providing intelligent assistance, answering queries, offering recommendations, and automating processes such as generating visualizations from data.
- **Dataset summarization and analysis:** AI and LLMs can be employed to summarize and analyze large datasets within a company, using reasoning engines to derive valuable insights for decision making.
- **Summarization:** AI and LLMs can be used to summarize interactions, whether in customer support chats, emails, or meetings, enabling more efficient review and analysis.
- **Report and document analysis:** LLMs can assist in analyzing reports, documents, and other text-based content, extracting key information and insights. Some recent examples include simplification of complex domain-specific texts such as legal contracts and communication from governments. Another set of use cases includes aggregating complex texts (such as several scientific articles) into a single piece of knowledge and performing analysis on conclusions, contradictions, and assumptions.
- **New value-added products:** Carefully managed and governed datasets can be transformed into new value-added products for customers, such as data summaries, reports, and insights.
- **Interactive user experience for complex operations:** Handling complex operations such as flight rebooking or financial transactions using AI may require secure connections to APIs and data sources, demanding meticulous dataset preparation, and in some cases, custom LLMs.

It is important to note that AI and LLMs provide a new way to communicate between humans and machines, essentially serving as a unique user interface enriched with reasoning capabilities. AI and LLMs have the potential to revolutionize various levels of human-to-machine interactions, including interactions with hardware. However, to unlock their full potential, organizations must foster a robust application development and data engineering culture, along with the development of specific best practices that continue to evolve with rapid advancements in the field.

For organizations creating user-facing AI products, Microsoft provides the [HAX Toolkit](#), which offers hands-on tools for building effective human-AI experiences. The [HAX Design Library](#) features an interactive collection of guidelines for human-AI interaction. These guidelines ensure AI systems are transparent by setting expectations regarding the [purpose of the AI system, its performance,](#) and providing users access to an [explanation of the AI system's behavior.](#) Teams within organizations can mix and match tools depending on their needs, use cases, and where they are in their product's lifecycle. For instance, disclosing the use of AI in a product increases trust in the organization delivering the product.

# Conclusion

Understanding the potential benefits of scaling AI adoption within an enterprise is paramount, whether for internal operations or customer-facing applications. The six strategies for AI implementation presented in this e-book offer guidance on the key components required for successful organization-wide AI adoption.

Specifically designed and optimized for responsible AI practices, Azure provides native capabilities that are integrated to connect data estates with AI in a governed, secure, and controllable manner. Over 95% of Fortune Global 500 companies trust Microsoft with their business. By using Azure's unique capabilities, and tightly integrated and highly secure ecosystem, some of these companies were able to achieve over 150%<sup>1</sup> increase in their work output using Azure AI to automate and scale processes, and up to 88%<sup>2</sup> improvement in developer productivity using GitHub Copilot.

---

1 <https://tools.totaleconomicimpact.com/go/microsoft/azureai/index.html?lang=en-us>

2 <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/>

## Next steps

- Get in touch with [Azure Sales](#)
- Learn more about [responsible AI on Azure](#)
- Learn more about the [Microsoft Responsible AI Partners Initiative](#)
- Use [Azure AI](#) built-in tools and guidance to ensure safe, secure, and responsible AI