3Cloud

Microsoft

Meeting Compliance Needs for Financial Services in the Era of AI

sales@3cloudsolutions.com // 888-88-AZURE

Executive summary

Generative artificial intelligence (AI) is driving digital transformation in the financial services industry. It uses models or algorithms based on training data to create new content, including audio, code, images, text, simulations, and videos.

Although generative AI can improve efficiency and enhance consumer services, there are also known risks that can lead to legitimate concerns about fairness, ethics, reliability, explainability, privacy, security, and inclusiveness.

Understanding these risks and opportunities, Microsoft leads the way in addressing these issues.

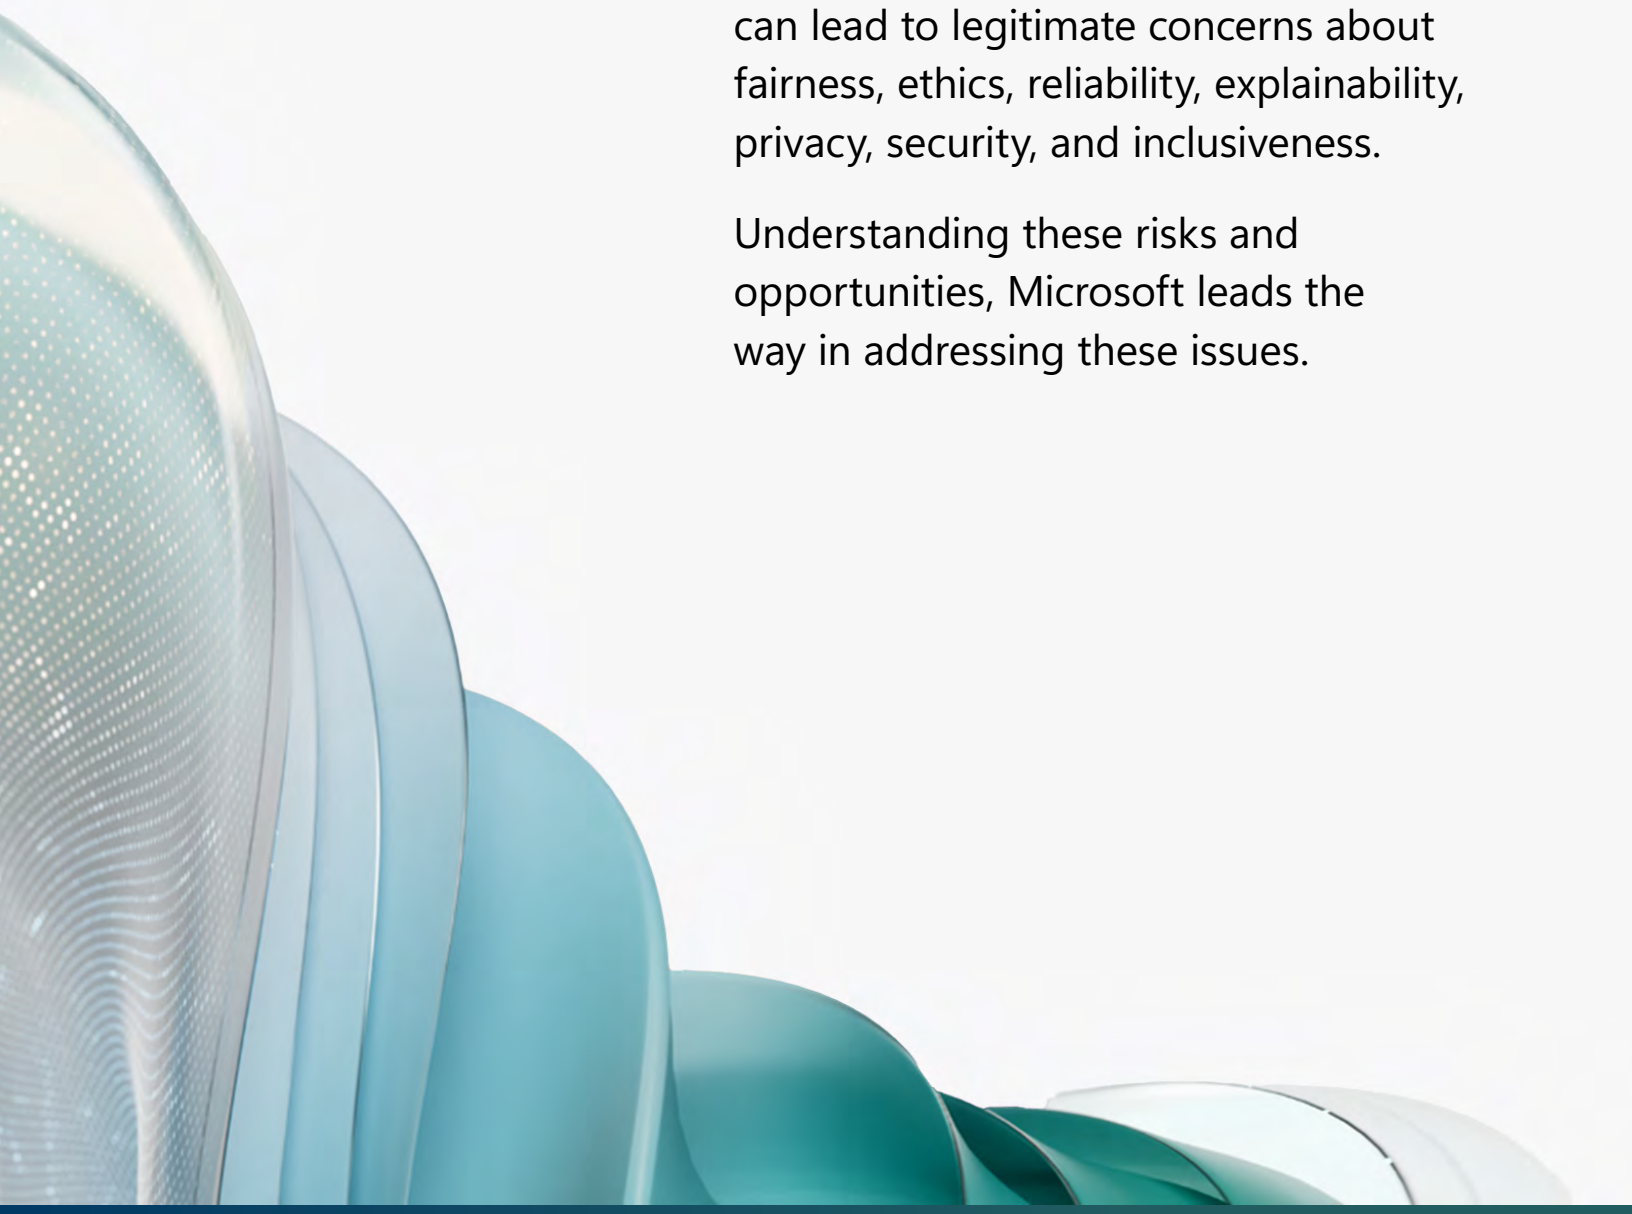

Financial institutions have an unprecedented opportunity to use AI to expand the provision of services to underbanked and underserved populations and communities, mitigate risks, reduce costs, and increase synergies.

Generative AI can also drive massive productivity gains for employees inside financial institutions. With generative AI, all employees can find answers to questions, generate content, and strengthen employee skills in customer-facing scenarios with real-time coaching.

At Microsoft we lead the way in addressing these issues, empowering financial services organizations with tools, resources, guidelines, and frameworks to innovate with AI while maintaining accountability and mitigating risk.

In addition, we provide organizations and regulators with a framework to understand how their approach to responsible AI aligns with relevant legal and regulatory frameworks.

As noted in a 2023 UK Department for Science, Innovation, and Technology paper, "Pro-Innovation Approach to AI Regulation":

> **"It is crucial that we do all we can to create the right environment to harness the benefits of AI and remain at the forefront of technological developments. That includes getting regulation right so that innovators can thrive, and the risks posed by AI can be addressed."**

Ensuring the responsible use of AI

Generative AI is the collective responsibility of technology providers, regulators, financial institutions, and other stakeholders. Working together, we can advance innovations that benefit consumers, increase productivity, and reduce risk.

The business value of AI in financial services

Microsoft AI technologies support three key financial services verticals: banking, capital markets, and insurance. AI technologies built on the highly secure and compliant Microsoft Cloud platform enable financial institutions to optimize costs, reduce time to value, enhance collaboration, gain insights, and deliver more impactful business outcomes.

Banking

Banks can use generative AI and Microsoft AI technologies to empower contact center agents and financial advisors; generate content; detect fraud; and improve knowledge of products, customers, and counterparty requirements. Banks can use Azure OpenAI Service to accelerate the deployment of their own tailored use cases, and at Microsoft we continue to provide new capabilities to enable productivity through Copilot for Microsoft 365.

Insurance

Generative AI in the insurance industry will have the greatest immediate impact on the productivity of contact center agents, underwriters, and claims managers. It will enable the rapid creation of intelligent virtual agents and assistants. It will also assist in mitigating risk through fraud detection and improved knowledge of products and customers. Azure OpenAI Service enables insurers to build generative AI solutions for all these purposes.

Capital markets

Key use cases where generative AI can have the greatest impact in capital markets include client engagement and customer service, market research and report summarization, and pitch book generation. Other key areas include investment and wealth advisory; knowledge of products, customers, and counterparty requirements; fraud analysis; accessibility; and language translation.

Microsoft responsible AI initiatives, principles, and standards

History

At Microsoft, we've established ourselves as thought leaders in the responsible use of AI by implementing our own internal governance and framework over the past six years. In 2017, we established the research-led Aether Committee—AI, Ethics, and Effects in Engineering and Research—to explore development standards for AI technologies. In 2018, we established our responsible AI principles. In 2019, building on our approach to privacy, accessibility, and security, we established the Office of Responsible AI to develop comprehensive principles and standards for the development and implementation of AI.

Throughout this journey, we've worked closely with regulators as they develop guidelines for the responsible use of AI, including the Monetary Authority of Singapore, the Dutch National Bank, the Basel Supervision Committee, and the International Association of Insurance Supervisors.

Microsoft responsible AI principles

We use the following responsible AI principles to guide AI development and use:

1. **Fairness:** AI technologies should not treat people unfairly or reinforce existing societal biases. Mitigating unfairness starts with understanding the implications and limitations of AI predictions and recommendations.

2. **Reliability and safety**: AI systems should operate reliably, safely, and consistently, even under unexpected conditions. Rigorous testing is essential during system development and deployment, and proper operation, maintenance, and protection of AI systems are important over their lifespan.

3. **Privacy and security**: Protecting privacy and securing important personal and business information becomes more critical and complex as AI becomes more prevalent. Access to data is essential for AI systems to make accurate and informed predictions and decisions, so protecting personally identifiable information and confidential business data is paramount.

4. **Inclusiveness:** Everyone should benefit from intelligent technology, which should incorporate and address a broad range of human needs and experiences. AI can improve access to opportunities for users with differing backgrounds, skill levels, and perspectives.

5. **Transparency:** Transparency relies on traceability, communication, and intelligibility. People should be able to understand and monitor the technical behavior of AI systems, and those building and using AI systems should be forthcoming about their design, deployment, and limitations.

6. **Accountability:** The people who deploy AI systems should be accountable for how those systems operate. It's important to establish an internal governance system tailored to the organization's unique characteristics, culture, guiding principles, and level of engagement with AI.

Based on these principles, we launched the Responsible AI Standard, driven by a multi-year effort to define product development requirements for responsible AI. This standard consists of goals and requirements for each of the six principles, functioning as a checklist or scorecard for companies developing and implementing AI systems.

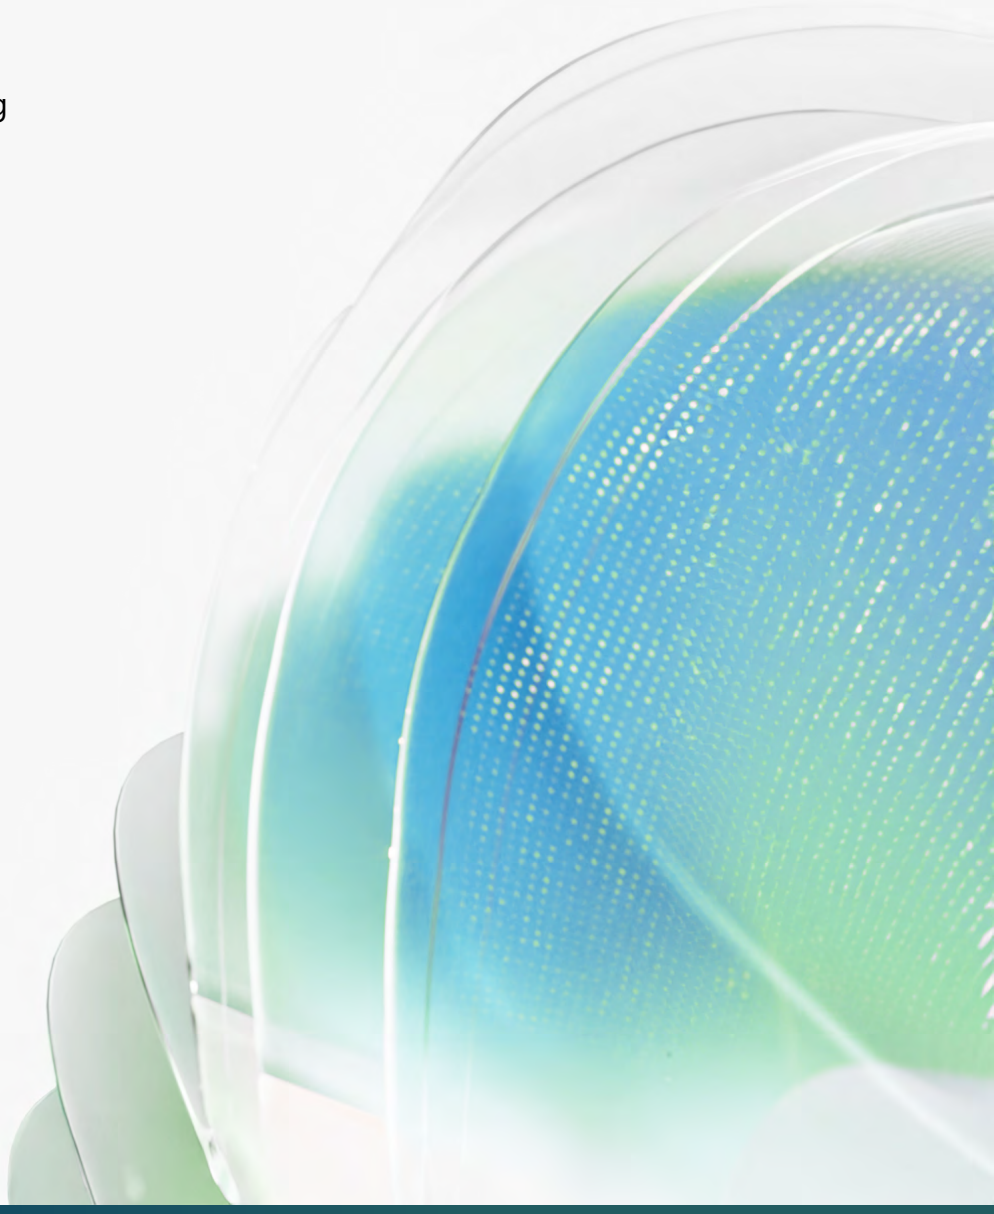

Common principles in global-state AI governance and regulation

Governments are adopting a range of principles, standards, and regulations to govern the use of AI technologies. A consensus is emerging around the principles that should guide financial institutions' evaluation, development, and implementation of these technologies. These principles also align with existing regulatory frameworks for information technology, third-party risk management, and safety and soundness.

The six principles are effectiveness, fairness, privacy and data security, transparency, training and governance, and ethics. These align with Microsoft responsible AI principles, as shown in the following diagram.

Alignment of Microsoft principles to global principles

Microsoft responsible AI principles	Global regulatory AI principles for financial services
Privacy and security	Privacy
Reliability and safety	Effectiveness
	Explainability
Accountability	Training & governance
Transparency	Transparency
Fairness	Fairness
Inclusiveness	Ethics

Common issues and concerns for the use of AI

Financial services organizations have raised several issues and concerns with the development and implementation of AI and, more specifically, generative AI technologies. The Microsoft responsible AI initiative and Microsoft AI systems and tools respond to these issues and concerns with six areas of focus:

- **Data privacy**: Microsoft provides a permissions model to ensure only authorized users and groups can access data. We don't use customer data to train or improve Azure AI foundation models, and we adhere to data protection requirements such as the Global Data Protection Regulation (GDPR).
- **Data location**: Microsoft manages AI data location in-line with general data location commitments. Data remains within agreed-upon boundaries consistent with data protection commitments, including European Union (EU) organizations for whom data must remain within the EU Data Boundary.
- **Legal and liability**: Microsoft will defend a customer if an IP infringement claim is brought against the customer based on their use of outputs produced by Microsoft's AI-powered copilot services. Microsoft extended protection to Azure OpenAI Service, subject to requirements in-line with the shared responsibility model. These commitments are subject to various safeguards and other requirements organizations must adhere to in order to benefit from this commitment.
- **IP ownership**: Microsoft holds no ownership rights over prompts or input data and doesn't claim ownership over AI-generated output or content.
- **Security controls**: Microsoft has built-in security controls for its services that adhere to a broad range of certification requirements like ISO27001 and FedRamp. We also provide a range of security and privacy features for organizations to manage. The permissions model within a Microsoft 365 tenant provides controls to prevent data leakage among users, groups, and tenants. Azure OpenAI Service and cognitive services don't directly access data—instead, data is passed through the services using application programming interfaces (APIs). The application developer must implement proper authentication controls to protect data within the application.
- **Safety and accountability**: Microsoft provides multiple levels of safety in our AI technologies, including application-level protections, technical protections, and process and policy protections.

Microsoft helps financial services organizations with AI applications

Microsoft's AI technologies are well-suited to enable financial services organizations to develop and integrate AI technologies and systems in accordance with prevailing global principles and standards. Our AI technologies are guided by our responsible AI principles, which are designed to align with global laws, regulations, and guidance.

As a result, Microsoft AI technologies are well-positioned to help financial organizations use AI consistent with regulatory expectations, enable oversight by senior management and the board of directors, address supervisory expectations, and protect organizations.

In addition, we make training programs and content available to help financial services organizations educate employees and manage AI knowledge. We have extensive experience working with financial services regulators to provide the information required for examinations of supervised institutions.

For each key principle, Microsoft provides capabilities and resources to help organizations manage and govern risk in the use of AI consistent with our responsible AI principles and regulatory expectations.

Coverage areas

- Effectiveness
- Fairness
- Privacy and data security
- Transparency
- Training and governance
- Ethics

Effectiveness: AI technology should be effective, reliable, and suitable for its intended use.

- The Microsoft Responsible AI Standard provides financial institutions with a useful starting point for evaluating AI solutions as institutions develop their own criteria and objectives.
- Microsoft offers "transparency notes" to help organizations understand how AI technologies work.
- The Microsoft Responsible AI Dashboard provides top feature importance, causal analysis, and counterfactual explanations for tabular, image, and text data with classification.

Fairness: AI technology should not result in discrimination, bias, or unintended outcomes for consumers.

- The Microsoft Responsible AI Standard and principles prioritize the avoidance of unlawful discrimination with AI. We provide tools to help prevent the misuse of AI technologies.
- Microsoft uses a methodology for reducing bias in word embedding to reduce the reinforcement of gender stereotypes, such as reducing the association between "receptionist" and "female," while maintaining potentially useful associations, like the association between the words "queen" and "female."
- Microsoft conducts internal testing to evaluate AI quality and responsible AI metrics for large language models (LLMs) such as those that underpin copilot experiences and makes these tools available to organizations using Azure OpenAI Service.
- The Microsoft Responsible AI Dashboard helps institutions implement responsible AI practices efficiently and fairly. Other open-source tools, such as Fairlearn, help data scientists improve the fairness of AI systems.

Privacy and data security: AI technology should be supported by strong controls to protect consumer privacy and data security.

- Copilot for Microsoft 365 is integrated into Microsoft 365 and adheres to all existing privacy, security, and compliance commitments to Microsoft 365 commercial customers.
- Azure OpenAI Service stores and processes data subject to all controls as set forth in the Data Protection Addendum (DPA). Within Microsoft 365 and for Azure OpenAI Service developers, Microsoft offers Microsoft Purview—a family of data governance, risk, and compliance solutions that help govern, protect, and manage the data estate. Microsoft AI technologies are built on Azure, which includes Microsoft information security controls that can be integrated with a financial institution's information security program.

Transparency: Transparency relies on traceability, communication, and intelligibility—it should be auditable. People should be able to understand and monitor the technical behavior of AI systems, and those building and using AI systems should be forthcoming about their design, deployment, and limitations.

- Microsoft offers AI tools to track model inputs and models, including machine learning operations (MLOps), which makes it easier to track and reproduce models and their version histories. Microsoft AI technologies include tools to enable transparency. For example, InterpretML is an open-source toolkit that helps explain model behavior and enables the use of interpretability algorithms such as Explainable Boosting Machine (EBM).
- Azure Machine Learning prompt flow provides methods and tools to ensure that outcomes or outputs of AI technologies are identifiable and explainable. Azure OpenAI Service in Azure Machine Learning also enables model designers and evaluators to explain why a model makes the predictions it does by providing information relevant to traceability (AI systems developed with clearly documented goals, definitions, design choices, and any assumptions) and intelligibility (people should be able to fully understand and monitor the technical behavior of AI systems). These insights can be used to debug the model, validate that its behavior matches objectives, check for bias, and build trust.
- Microsoft offers "transparency notes" intended to help organizations understand how AI technologies work.

Training and governance: An employee of a financial services organization should have the necessary expertise to implement and review the AI technology, and the AI technology should be subject to governance and oversight within the organization.

- Microsoft provides training programs and content to assist with education and knowledge management within financial organizations with respect to AI. These include learning courses like those available on Microsoft Learn's AI learning hub. Learning modules and credentials provide education and training for organizations to use as part of a governance program. The Service Trust Portal provides additional assurance and documentation, and the Compliance Program provides assistance with assessments of the Microsoft Cloud, including AI technologies.

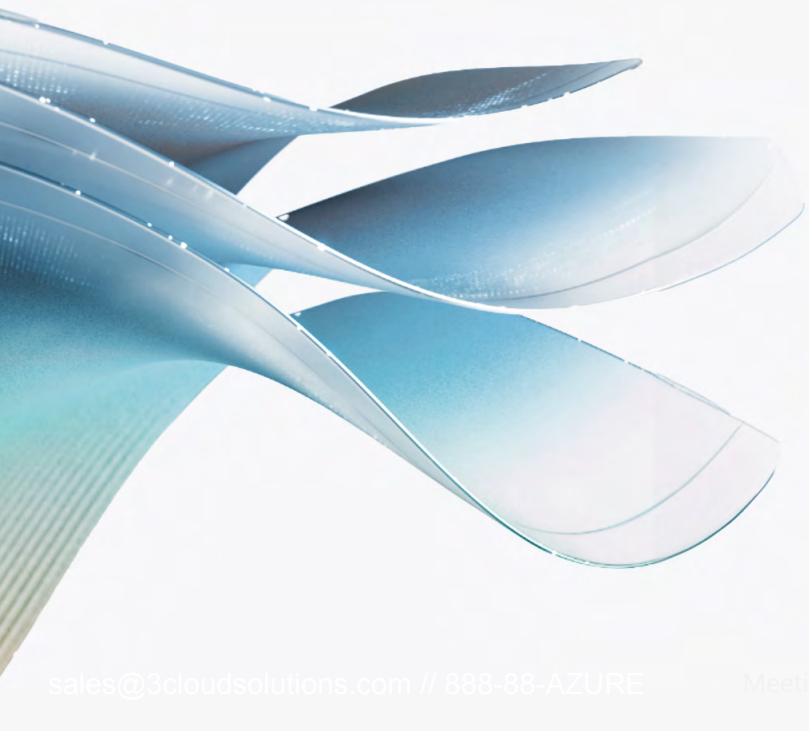

Ethics: The use of AI technology should align with the financial organization's code of conduct and applicable ethics standards, which requires ongoing accountability by the firm in terms of use and oversight of the AI technology.

- Microsoft developed its Responsible AI Standard and principles with the goal of helping to ensure the ethical use of AI. We provide information to assist financial services organizations in governing AI technologies in ways that are consistent with their policies and procedures, including codes of conduct.
- Azure OpenAI Service has a code of conduct that financial services organizations must follow to protect against and mitigate misuse.
- Microsoft has committed to implementing the National Institute of Standards and Technology (NIST) AI risk management framework and has aligned its Responsible AI Standard with the ISO 42001 AI Management System. We plan to implement future standards, including those that will emerge following the implementation of the EU AI Act.

Conclusion

Microsoft has been a leader for many years in developing AI technologies in accordance with responsible AI principles that align in lockstep with the principles being articulated by global policymakers in laws, regulations, and guidance.

As a result, Microsoft generative AI can help financial services organizations gain the benefits of AI while meeting legal and compliance requirements, protecting data and systems, and maintaining customer trust.

Read the full whitepaper