

AWS User Guide to Governance, Risk & Compliance for Responsible AI Adoption within Financial Services Industries

May 2025



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Additionally, this document does not constitute legal advice and should not be relied on as legal advice. AWS encourages its customers to obtain appropriate advice on their implementation of privacy and data protection environments, and more generally, applicable laws relevant to their business.

© 2025 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

Abstract	4
Intended Audience	4
Introduction	5
Responsible AI and the AWS Frontier Model Safety Framework.....	7
AWS Compliance Programs	8
GRC Considerations for Responsible AI	12
Next Steps	28
Appendix 1: Overview of emerging AI regulation and guidance by Geography (non-exhaustive)	30
Appendix 2: Differences between AI, ML, Deep Learning and Generative AI	31
Appendix 3: AWS Responsible AI Policy	33
Appendix 4: Customer Case Studies.....	35

Abstract

This document provides information regarding governance, risk and compliance (GRC) considerations for the adoption of responsible Artificial Intelligence (AI) for Amazon Web Services (AWS) Financial Services Industry (FSI) customers.

This guide defines GRC in the context of AI adoption, describes the roles that AWS and its customers play in responsible AI adoption on AWS, describes the AWS Shared Responsibility Model, compliance frameworks and advanced tools that customers can use to meet applicable regulatory requirements. A detailed jurisdiction by jurisdiction analysis of requirements and expectations is out of scope for this user guide.

Intended Audience

This user guide is intended to be primarily used by GRC leaders and practitioners within financial services customers who are planning to adopt AWS AI services. The user guide is intended to highlight relevant GRC considerations and how these may be addressed by AWS services, mechanisms and guidance.

Introduction

The rapid growth of AI brings promising new innovation, and at the same time raises new challenges. At AWS, we are committed to developing AI responsibly, taking a people-centric approach that prioritizes education, science, and our customers, to integrate responsible AI across the end-to-end AI lifecycle.

Of particular significance for FSI customers adopting AI, AWS is excited to be the first major cloud service provider to announce ISO/IEC 42001 accredited certification for AI services, covering: [Amazon Bedrock](#), [Amazon Q Business](#), [Amazon Textract](#), and [Amazon Transcribe](#). The ISO42001 AI Management System standard outlines best practices and controls for the responsible development, deployment, and operation of AI systems.

What is AI?

AI is a technology with human-like problem-solving capabilities. AI in action appears to simulate human intelligence—it can recognize images, write poems, and make data-based predictions.

AI is also an umbrella term for different strategies and techniques for making machines more human-like. It includes everything from self-driving cars to robotic vacuum cleaners and smart assistants like Alexa. While machine learning and deep learning fall under the AI umbrella, not all AI activities are machine learning and deep learning. For example, generative AI demonstrates human-like creative capabilities and is a very advanced form of deep learning (refer to **Appendix 1** for further details of the differences between AI, Machine Learning (ML) and Deep Learning)

What is AI Governance and why is it important?

We define AI Governance as the principles, policies, tools and processes that organizations implement to ensure the responsible development, deployment, and monitoring of AI systems.

AI Governance is important as the board and senior management of financial institutions are ultimately accountable for their activities, including AI use cases¹. Effective AI governance aims to establish standards and guidelines that promote ethical AI practices, protect individual rights and mitigate risks associated with AI²

¹ [Regulating AI in the financial sector: recent developments and main challenges](#), FSI Insights, Financial Stability Institute, The Bank of International Settlements, December 2024

² [White Paper on AI Governance](#), Governance Institute of Australia, September 2024

Why is AI adoption different to other technology adoption?

The International Organization for Standardization the International Electrotechnical Commission (ISO/IEC) 42001 Artificial Intelligence Management System standard states that AI raises the following specific considerations³:

- The use of AI for automatic decision-making (sometimes in a non-transparent and non-explainable way) can require specific management beyond the management of classical IT systems
- The use of data analysis, insight and machine learning, rather than human-coded logic to design systems, both increases the application opportunities for AI systems and changes the way that such systems are developed, justified and deployed
- AI systems that perform continuous learning change their behaviour during use. They require special consideration to ensure their responsible use continues with changing behaviour

In other words, there are unique considerations for AI adoption due to the nature of how the technology is designed, how it behaves and its potential impact on end customers, organizations and regulators.

In addition, FSI customers also face additional adoption challenges due to the introduction of new regulations, such as the world first [EU AI Act](#) and emerging regulation and guidance in the US, UK and Asia Pacific (refer to Appendix 1 for an overview of emerging regulation and guidance and Appendix 4 for AWS customer case studies)

As such, this user guide is intended to provide AWS FSI customers with important governance, risk and compliance considerations when adopting AWS AI services and products. A detailed jurisdiction by jurisdiction analysis of requirements and expectations is out of scope for this user guide.

³ [ISO/IEC 42001 - Information technology — Artificial intelligence — Management system](#), 2023

Responsible AI and the AWS Frontier Model Safety Framework

Responsible AI at AWS

At AWS, we are committed to developing AI responsibly, taking a people-centric approach that prioritizes education, science, and our customers, to integrate responsible AI across the end-to-end AI lifecycle.

[We define the core dimensions of responsible AI](#) as:

Fairness	Considering impacts on different groups of stakeholders
Explainability	Understanding and evaluating system outputs
Privacy and security	Appropriately obtaining, using, and protecting data and models
Safety	Preventing harmful system output and misuse
Controllability	Having mechanisms to monitor and steer AI system behavior
Veracity and robustness	Achieving correct system outputs, even with unexpected or adversarial inputs
Governance	Incorporating best practices into the AI supply chain, including providers and deployers
Transparency	Enabling stakeholders to make informed choices about their engagement with an AI system

Table 1: AWS core dimensions of responsible AI

From the outset, we have prioritized responsible AI innovation by embedding safety, fairness, robustness, security, and privacy into our development processes and educating our employees. Our practical approach to transform responsible AI from theory into practice, coupled with tools and expertise, enables AWS customers to implement responsible AI practices effectively within their organizations.

AWS Frontier Model Safety Framework

"In addition to the AWS responsibilities for AWS AI services, we have taken further steps to ensure the safe development of future advanced AI models referred to as frontier models. The Korea Frontier AI Safety Commitments⁴ define frontier models as highly capable general-purpose AI models or systems that can perform a wide variety of tasks and match or exceed the capabilities present in the most advanced models.

As we continue to scale the capabilities of Amazon's frontier models and democratize access to the benefits of AI, we also take responsibility for mitigating the risks of our technology. Consistent with Amazon's endorsement of the [Korea Frontier AI Safety Commitments](#)⁵, this Framework outlines the protocols we will follow to ensure that frontier models developed by Amazon do not expose critical capabilities that have the potential to create severe risks. At its core, this Framework reflects our commitment that we will not deploy frontier AI models developed by Amazon that exceed specified risk thresholds without appropriate safeguards in place. Please refer to [Amazon's Frontier Model Safety Framework](#) for further details of how the processes Amazon will use to identify, assess, and manage potential severe risks that could arise as we develop more advanced and highly-capable frontier AI models.

AWS Compliance Programs

AWS has obtained certifications and independent third-party attestations for a variety of industry-specific workloads.

Of particular significance for FSI customers adopting AI, AWS is excited to be the first major cloud service provider to announce ISO/IEC 42001 accredited certification for AI services, covering: [Amazon Bedrock](#), [Amazon Q Business](#), [Amazon Textract](#), and [Amazon Transcribe](#). The ISO42001 AI Management System standard outlines best practices and controls for the responsible development, deployment, and operation of artificial intelligence (AI) systems. The basis of this certification is the establishment of an AI management system that ensures AI technologies are developed and used ethically, securely, and transparently. This includes integrating AI management into organizational processes, conducting risk and impact assessments, and ensuring

⁴ The Korea Frontier AI Safety Commitments are significant because ten countries and the European Union signed the Seoul Declaration at the AI Seoul Summit in South Korea on 21 May 2024. Twenty-seven countries and the European Union signed the Seoul Ministerial Statement on 22 May 2024. For further details refer to: <https://www.industry.gov.au/publications/seoul-declaration-countries-attending-ai-seoul-summit-21-22-may-2024>

⁵ As above

compliance with privacy laws and AI security standards. For more information, see the [ISO 42001 Compliance webpage](#).

Other relevant compliance programs for customers adopting AI include:

- **ISO 27001:** A security management standard that specifies security management best practices and comprehensive security controls that follow the ISO 27002 best practice guidance. The basis of this certification is the development and implementation of a rigorous security program, which includes the development and implementation of an information security management system that defines how AWS perpetually manages security in a holistic, comprehensive manner. For more information, or to download the AWS ISO 27001 certification, see the [ISO 27001 Compliance webpage](#).
- **ISO 27017:** Provides guidance on the information security aspects of cloud computing, recommending the implementation of cloud-specific information security controls that supplement the guidance of the ISO 27002 and ISO 27001 standards. This code of practice provides additional implementation guidance for information security controls specific to cloud service providers. For more information, or to download the AWS ISO 27017 certification, see the [ISO 27017 Compliance webpage](#).
- **ISO 27018:** Code of practice that focuses on protecting personal data in the cloud. It is based on the ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls that are applicable to cloud personally identifiable information (PII). It also provides a set of additional controls and associated guidance intended to address cloud PII protection requirements not addressed by the existing ISO 27002 control set. For more information, or to download the AWS ISO 27018 certification, see the [ISO 27018 Compliance webpage](#).
- **ISO 22301:** Specifies the structure and requirements to implement, maintain and improve a business continuity management system (BCMS) to protect against, reduce the likelihood of the occurrence of, prepare for, respond to, and recover from disruptions when they arise. Compliance to this standard provides assurance on AWS commitment to business continuity and resiliency of AWS services. For more information or to download the AWS ISO 22301 certification, see the [ISO 22301 Compliance webpage](#).

- **ISO 9001:** Outlines a process-oriented approach to documenting and reviewing the structure, responsibilities, and procedures that are required to achieve effective quality management within an organization. The key to the ongoing certification under this standard is establishing, maintaining, and improving the organizational structure, responsibilities, procedures, processes, and resources so AWS products and services consistently satisfy ISO 9001 quality requirements. For more information or to download the AWS ISO 9001 certification, see the [ISO 9001 Compliance webpage](#).
- **PCI DSS Level 1:** The Payment Card Industry Data Security Standard (PCI DSS) is a proprietary information security standard administered by the PCI Security Standards Council. PCI DSS applies to all entities that store, process, or transmit cardholder data (CHD) or sensitive authentication data (SAD), including merchants, processors, acquirers, issuers, and service providers. The PCI DSS is mandated by the card brands and administered by the Payment Card Industry Security Standards Council. AWS is certified as a PCI DSS Level 1 Service Provider, the highest level of assessment available. For more information or to request the PCI DSS Attestation of Compliance and Responsibility Summary, see the [PCI DSS Compliance webpage](#).
- **SOC:** AWS System and Organization Control (SOC) Reports are independent third-party examination reports that demonstrate how AWS achieves key compliance controls and objectives. The purpose of these reports is to help customers and their auditors understand the AWS controls that have been established to support operations and compliance. For more information, see the [SOC Compliance webpage](#). AWS SOC Reports come in three forms:
 - **SOC 1:** Provides information about the AWS control environment that might be relevant to a customer's internal controls over financial reporting, as well as information for the assessment of the effectiveness of internal controls over financial reporting.
 - **SOC 2:** Provides customers and their service users that have a business need with an independent assessment of the AWS control environment relevant to system security, availability, and confidentiality.
 - **SOC 3:** Provides customers and their service users that have a business need with an independent assessment of the AWS control environment relevant to system security, availability, and confidentiality, without disclosing AWS internal information.

See the [AWS Compliance Programs webpage](#) for more information about AWS certifications and attestations. See the [Best Practices for Security, Identity, and Compliance website](#) for general AWS security controls and service-specific security.

AWS Artifact

Customers can use [AWS Artifact](#) to review and download reports and details about more than 2,600 security controls. In addition, the [AWS Artifact](#) portal provides on-demand access to AWS security and compliance documents, including ISO42001 certification, SOC reports, Payment Card Industry (PCI) reports, and certifications from accreditation bodies across geographies and compliance verticals.

GRC Considerations for Responsible AI

This section explores GRC considerations for responsible AI adoption. Refer to Appendix 4 for AWS customer case studies.

AWS has defined the core dimensions of responsible AI as: **Fairness, Explainability, Privacy and Security, Safety, Controllability, Veracity and Robustness, Governance and Transparency**. The following table outlines the key considerations for each dimension.

Dimensions of responsible AI	Key considerations
Fairness: Considering impacts on different groups of stakeholders	• Bias Detection and Mitigation: Regularly monitor AI models for potential biases that could lead to discriminatory outcomes, such as in credit decisions or customer service interactions. Implement strategies to mitigate these biases, including using diverse datasets and testing models for disparate impacts. • Regulatory Compliance: Identify and ensure compliance with applicable laws related to financial fairness (e.g. Equal Credit Opportunity Act), which prohibits discrimination based on protected characteristics. This involves continuous auditing and reporting to demonstrate fairness in AI-driven decisions. • Stakeholder Engagement: Engage with diverse stakeholders to understand potential impacts on different groups and incorporate feedback into AI development processes.

Dimensions of responsible AI	Key considerations
Explainability: Understanding and evaluating system outputs	• Transparency in Decision-Making: Implement explainable AI techniques to provide insights into how AI models arrive at their decisions. This is crucial for regulatory compliance and stakeholder trust. • Documentation and Auditing: Maintain detailed documentation of AI decision-making processes and conduct regular audits to ensure transparency and accountability. • Model Interpretability: Use interpretable models and provide clear explanations to stakeholders, enabling them to understand AI-driven outcomes and assess potential risks.
Privacy and Security: Appropriately obtaining, using, and protecting data and models	• Data Protection: Ensure that AI systems handle personal data securely and in compliance with privacy laws like GDPR. Implement robust data encryption and access controls to protect sensitive information. • Data Governance: Establish clear policies for data collection, storage, and use. Ensure that data is accurate, relevant, and used only for intended purposes. • Risk Assessment: Conduct thorough risk assessments to identify potential privacy and security vulnerabilities in AI systems and implement measures to mitigate these risks.
Safety: Preventing harmful system output and misuse	• Harm Prevention: Develop AI systems that prevent harmful outputs or misuse. Implement safeguards to detect and prevent adverse outcomes, such as financial fraud or data breaches. • Risk Management: Continuously monitor AI systems for safety risks and update risk management strategies as needed to ensure the integrity of financial operations.
Controllability: Having mechanisms to monitor and steer AI system behavior	• Monitoring and Oversight: Implement mechanisms to monitor AI system behavior continuously. This includes setting up alerts for unusual activity and having processes in place to intervene if necessary.

Dimensions of responsible AI	Key considerations
	• Human Oversight: Ensure that human operators can intervene in AI decision-making processes when required. This involves training staff to understand AI outputs and make informed decisions. • Feedback Loops: Establish feedback loops to improve AI performance and adapt to changing conditions, ensuring that AI systems remain aligned with organizational goals.
Veracity and Robustness: Achieving correct system outputs, even with unexpected or adversarial inputs	• Model Validation: Rigorously test and validate AI models to ensure they provide accurate and reliable outputs under various conditions, including unexpected inputs. • Data Quality: Ensure that datasets used for training AI models are of high quality, diverse, and representative to prevent biases and inaccuracies. • Continuous Improvement: Regularly update AI models and documentation to maintain their effectiveness and compliance with evolving regulatory standards.
Governance: Incorporating best practices into the AI supply chain, including providers and deployers	• Governance Structures: Ensure oversight of AI development and deployment in line with risk appetite. Develop and enforce internal policies for responsible AI development. Ensure that AI practices comply with relevant financial regulations, and any AI specific requirements. • Supply Chain Management: Manage the AI supply chain effectively, ensuring that all providers and deployers adhere to best practices and regulatory requirements.
Transparency: Enabling stakeholders to make informed choices about their engagement with an AI system	• Stakeholder Communication: Communicate clearly with stakeholders about AI-driven decisions and processes. Provide explanations and insights into how AI systems operate. • Documentation and Disclosure: Maintain detailed records of AI decision-making processes and disclose this information to regulators and stakeholders as required. • Public Trust: Foster public trust by demonstrating transparency and accountability in AI practices, ensuring that stakeholders understand the benefits and risks of AI systems.

AWS Prescriptive Guidance for Responsible AI

AWS recommends customers to consult the following prescriptive guidance for the responsible adoption of AI:

- [AWS Responsible Use of AI Guide](#)
- [AWS Cloud Adoption Framework for Artificial Intelligence, Machine Learning and Generative AI \(AWS CAF for AI\)](#)
- [AWS Well-Architected Generative AI Lens](#)
- [AWS Well-Architected Machine Learning Lens](#)
- [AWS Machine Learning Blog: Build safe and responsible generative AI applications with guardrails](#)
- [AWS Security Blog: An introduction to the Generative AI Security Scoping Matrix](#)

AWS recommends customers assess their workloads against best practices as recommended within the [AWS Well-Architected Generative AI Lens](#) and [AWS Well-Architected Machine Learning Lens](#). The AWS Well-Architected Framework has been developed to help cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications. Based on six pillars—operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability—the AWS Well-Architected Framework provides a consistent approach for customers to evaluate architectures and implement designs that scale over time.

GRC Considerations

The following table can help customers map expectations from key principles in common from global regulation and industry frameworks to relevant AWS resources. The tables are organized into the following columns:

Expectations references common expectations from regulatory guidance and industry frameworks in Europe, US and Asia Pacific, in summary form. Always refer to the authoritative source within your jurisdiction for verification.

Considerations on customer responsibilities when addressing Responsible AI expectations, or the AWS services that customers can use to address these expectations.

Resources lists additional documentation that customers may use to supplement the information in this guide

A full analysis of all regulations and jurisdictions is beyond the scope of this user guide. The information that follows includes considerations that AWS specialists frequently encounter in interactions with FSI customers

Expectations	AWS Customer Considerations	Resources
Governance Organizations implement principles, policies, tools and processes to ensure the responsible development, deployment, and monitoring of AI systems in line with strategy and risk appetite Aligned to: ISO42001 AI Management System: 4.1, 5.1, 5.2, 5.3, 6.2, 6.3, 7.1-7.3, 9.3, A.2.2-2.4, A.3.2, A.4.1-4.4, A.6.1.2, A.9.3, A.9.4 EU AI Act: Articles 14, 26, 22.1, 22.2, 26.3 NIST AI Risk Management Framework: Govern 1.1, Govern 1.2, Govern 2.3, Govern 3.1, Govern 4.1, Map 1.3, Map 1.4 Hong Kong Monetary Authority (HKMA) Consumer Protection in respect of Use of GenAI: Principle 1 – Governance and Accountability	Considerations Incorporating AI governance into an organization’s strategy is instrumental in building trust, enabling the deployment of AI technologies at scale, and overcoming challenges to drive business transformation and growth. By driving consistency, AI governance enables alignment with organizational goals, and ensures that AI technologies are ethically used and effectively managed. To that end, AI governance frameworks create consistent practices in the organization to address organizational risks, ethical deployment, data quality and usage, and even regulatory compliance, as well as managing the different cost patterns of AI workloads. AWS recommends customers consult the AWS CAF for AI (p22) to inform the design and operation of their AI governance and control frameworks. In practice, mature customers have ensured: • Business and technology strategies have been updated to incorporate AI adoption; and/or a separate AI strategy has been documented that maps to business and technology strategy outcomes • Strategies have been approved by their Board • Policies and procedures have been developed and implement to provide pathways for AI adoption to progress from idea to proof of concept to production environments • A clear risk appetite for AI adoption has been set, approved and communicated by Board or Board delegated committee • A Board or Board-delegated committee(s) charged with overseeing progress against strategy in line with risk appetite and approving high risk AI adoption • AI oversight governance forums are cross functional and include data and analytics, risk management, legal and compliance, technology and audit functions • Individuals responsible for AI governance are appropriately trained, skilled and certified • Crisis Management Plans and Incident response plans have been updated include responses for scenarios involving high risk AI	AWS CAF for AI ISO42001: AI Management System AWS Well-Architected Generative AI Lens Amazon SageMaker AI User Guide Amazon Sagemaker model monitor Amazon Bedrock evaluations Amazon Sagemaker pipelines AWS Enterprise Support AWS AI Training and Enablement

Expectations	AWS Customer Considerations	Resources
Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models: Core Principle 1 – Senior Management Responsibilities	AWS Services Customers should consider using the following AWS services to support AI governance: • AWS Trusted Advisor – AWS Trusted Advisor draws upon best practices learned from serving hundreds of thousands of AWS customers. Trusted Advisor inspects your AWS environment, and then makes recommendations when opportunities exist to save money, improve system availability and performance, or help close security gaps. • AWS Enterprise Support is recommended for customers planning to operate high risk AI services on AWS. AWS Enterprise Support will assist customers to manage, monitor, analyze and report on usage of AWS. AWS Enterprise Support provides customers with proactive planning, architectural reviews, and consultative guidance including: strategic business reviews, security improvement programs, guided Well-Architected reviews, and cost optimization workshops. AWS Enterprise Support also provides customers a response time within 15 minutes in the case of business-critical system going down • AWS AI Training and Enablement is recommended for customers. Please contact your AWS Account team to understand options for relevant training and certification	
Risk Management Organizations implement principles, policies, tools and processes to minimize the potential negative impacts in executing their AI strategy Aligned to: ISO42001 AI Management System: 6.1.1-6.1.4, 8.1-8.4, A.5.2, A.5.3, A.5.4, A.5.5	Considerations AWS recommends customers consult the AWS CAF for AI (p26), ISO42001: AI Management Systems (section 6.1.2), NIST AI Risk Management Framework (NIST AI RMF) and NIST Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile to inform the design and operation of their AI risk management activities. In particular, AWS recommends customers address AI specific risks and controls in their risk management activities as follows: • Consult the AWS Responsible AI core dimensions (e.g. explainability and safety) and NIST AI RMF Generative AI Profile (section 2) to identify AI specific risks that are functional in nature • Consult the AWS Securing Generative AI Security Scoping Matrix guidance to identify non-functional risks and controls (e.g. security and availability)	AWS CAF for AI (p26) , ISO42001: AI Management Systems (section 6.1.2) NIST AI Risk Management Framework NIST AI RMF: Generative Artificial Intelligence Profile AWS Security Blog: Securing Generative AI Security Scoping Matrix

Expectations	AWS Customer Considerations	Resources
EU AI Act: Articles 8.1, 8.2, 9.3-9.6, 9.9, 17.1, 27.1 NIST AI Risk Management Framework: Map 1.1, Map 3.1, Map 3.2, Measure 2.6, Measure 2.7, Measure 2.8, Measure 2.10, Measure 2.12, Manage 1.2-1.4, Measure 3.1-3.2, Manage 2.1-2.3, Manage 3.1, Govern 6.1-6.2 Hong Kong Monetary Authority (HKMA) Consumer Protection in respect of Use of GenAI: Principle 2 – Fairness Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models: Core Principle 2 – AI Model Risk Management	- ISO42001 contains reference controls and implementation guidance in Annex A and B respectively In practice, mature customers have ensured: - A clear risk appetite for AI adoption has been set, approved and communicated by Board or Board delegated committee - Standards and processes are established for ongoing revalidations of AI in production, with the intensity and frequency based on the materiality of the AI use case - All 3 Lines of Defense (for example, Business or technology owner – Line 1, Group Risk and Compliance teams – Line 2 and Internal Audit – Line 3) are aligned on context, definitions and incoming demand - Risk management activities follow a tiered approach where high-risk use cases undergo more extensive scrutiny - AI specific risk assessment activities are either created or augmented within existing frameworks that interact with their corporate governance structures - Individuals responsible for risk management activities for AI use cases are appropriately trained, skilled and certified AWS Services Customers should consider using the following AWS services to support AI risk management: AWS AI service cards support the end-to-end AI lifecycle to promote transparency, and are available for services including Amazon Nova and Amazon Transcribe. These provide information such as the model's intended use cases, training details, evaluation metrics, results, observations, and recommendations. Amazon Sagemaker has model cards designed to specify the intended use of a model and document information that users need to train or deploy the model responsibly AWS AI Training and Enablement is recommended for customers. Please contact your AWS Account team to understand options for relevant training and certification	AWS AI service cards Amazon Bedrock Guardrails User Guide Amazon Bedrock Evaluations User Guide Amazon Sagemaker Clarify User Guide AWS AI Training and Enablement AWS for Industries Blog: Is your Enterprise Risk Management Framework ready for the Cloud?

Expectations	AWS Customer Considerations	Resources
Compliance Organizations understand and ensure compliance with applicable AI regulations and standards EU AI Act: 5.1, 5.2, 6.1-6.4, 8.1-8.2, 40.1, 41.1, 42.1, 43.1-43.4, 44.2-44.3, 47.1-47.4, 49.1-49.3 NIST AI RMF: Govern 1.1, Map 4.1	Considerations AWS recommends customers consult the AWS CAF for AI: Security perspective - Compliance and Assurance , NIST AI RMF (Govern 1.1 and Map 4.1) to inform the design and operation of AI compliance activities AWS recommends customers to engage with their compliance teams to assess use cases for compliance with applicable laws and regulations throughout all phases of design, development, deployment, and operation of AI systems. AI is a constantly evolving landscape, and new techniques, technologies, laws, and social norms will continue to be developed and evolve over time. As such it is important to have mechanisms to keep GRC practitioners up to date. In practice, mature customers have ensured: • All 3 Lines of Defense (for example, Business or technology owner – Line 1, Group Risk and Compliance teams – Line 2 and Internal Audit – Line 3) are aligned on context, definitions and incoming demand • Policies and procedures have been updated or created to outline pathways for AI adoption to progress from idea to proof of concept to production environments • Compliance activities follow a tiered approach where high-risk use cases undergo more extensive scrutiny • Processes for reporting incidents, safety issues, and non-compliance to relevant authorities and affected stakeholders have been defined • Appropriate disclosures are made for both external and internal facing AI user cases AWS Services Customers should consider using the following AWS services to support AI compliance activities: • For assurance over AWS global infrastructure, customers can rely on the independent assurance made available for free through AWS Artifact. AWS' ISO42001 certification and SOC2 assurance over Amazon AI services are particularly relevant.	AWS CAF for AI: Security perspective - Compliance and Assurance NIST AI RMF Getting started with AWS Artifact AWS AI Training and Enablement

Expectations	AWS Customer Considerations	Resources
	AWS AI Training and Enablement is recommended for customers. Please contact your AWS Account team to understand options for relevant training and certification	
Data management Organizations understand the role and impacts of data in AI systems in the application and development, provision or use of AI systems throughout their life cycles. Aligned to: ISO42001 AI Management System: B.7 – Data for AI Systems, B.4.3 Data Resources NIST AI Risk Management Framework: GenAI Profile MP-2.1-001, MP-2.1-002 EU AI Act: Article 10 - Data and Data Governance Monetary Authority of Singapore (MAS): “AI model risk management paper” Hong Kong Monetary Authority (HKMA) Consumer Protection in respect of Use of GenAI: Principle 4 – Data Privacy and Protection	Considerations AWS recommends customers consult ISO42001 AI Management, Section B7 for an overview of expectations, controls and recommended implementation with respect to data management for AI systems. Customers should also consult the AWS CAF for AI for advice on data curation (p25) AWS recommends customers assess their AI deployments against AWS best practices as recommended within the AWS Well-Architected Framework - Machine Learning Lens, particularly the data processing lifecycle as well as the AWS Well-Architected Generative AI Lens. For assurance over AWS global infrastructure, customers should rely on the independent assurance made available through AWS Artifact, particularly the ISO42001 certification and SOC2 report, which includes AWS AI services such as Amazon SageMaker AI and Amazon Bedrock . In practice, mature customers have ensured: - In-scope data includes training, validation and testing data - Assessment and documentation for: - data collection processes and data origin - data-preparation/transformation processes (e.g. labelling, cleaning, enrichment) - the availability, quantity and suitability of the data sets required - assumptions, potential biases and mitigating measures - representativeness and relevance of data - use of personal data - Data protection controls and data validation checks to confirm the completeness, accuracy, relevance, reliability and quality of data	AWS CAF for AI ISO42001: AI Management Systems AWS Well-Architected Generative AI Lens AWS Well Architected Framework – Machine Learning Lens – data processing lifecycle Amazon SageMaker AI User Guide Amazon Bedrock User Guide Amazon Sagemaker Model Monitor Amazon Bedrock Data Automation User Guide Amazon SageMaker Data Preparation Developer Guide

Expectations	AWS Customer Considerations	Resources
Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models: Core Principle 3 – Cybersecurity and Data Risk Management	AWS Services Customers should consider using the following AWS services to support data management for AI: • Amazon Sagemaker Model Monitor can continuously check the data being used for inference against a set of defined data quality constraints, such as missing values, outliers, and distribution shifts. It can automatically detect and alert on data quality issues. • Amazon SageMaker Unified Studio provides a single data and AI development environment that brings together AWS data, analytics, artificial intelligence (AI), and machine learning (ML) services. It provides a place to build, deploy, execute, and monitor workflows from a single interface, allowing you to securely build and share analytics and AI artifacts, including data, models, and generative AI applications. • Amazon Bedrock Data Automation provides a streamlined and automated approach to provisioning and preparing the data required for machine learning models. This includes features such as visual grounding with confidence scores for explainability and built-in hallucination mitigation. This ensures trustworthy and accurate insights from unstructured, multi-modal data sources. • Amazon SageMaker Canvas facilitates data preparation by providing a visual interface for data selection, cleaning, transformation, and feature engineering. Data used for building and testing AI systems and applications must be secure, especially when dealing with confidential information. AWS provides comprehensive encryption solutions to protect data both at rest and in transit. This includes server-side encryption across various AWS services like Amazon Bedrock, Amazon S3, Amazon EBS, and Amazon RDS, AWS KMS (Key Management Service) for key management, automatic encryption/decryption, and integration with AWS IAM for access control.	

Model Management

Organizations identify and document objectives and implements processes for the responsible design and development of AI systems.

Aligned to:

ISO42001 AI Management System: A6 – AI System life cycle, A9 – Use of AI systems

NIST AI Risk Management Framework: GenAI Profile MAP 3.3, Measure 2, Manage 4.1

EU AI Act: Article 9 – Risk Management Systems, Article 11- Technical Documentation

Monetary Authority of Singapore (MAS): “AI model risk management paper” section 6

Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models: Core Principle 2 – AI Model Risk Management

Considerations

AWS recommends customers consult ISO42001 AI Management, Section A6 which encompasses capturing the objective and processes for the responsible design and development of AI systems, including criteria and requirements for each stage of the AI system life cycle.

AWS recommends customers assess their AI deployments against AWS best practices as recommended within the [AWS Well-Architected Generative AI Lens](#), particularly the practices regarding:

- [model selection and cost optimization](#)
- [model performance evaluations](#)
- [energy efficient models](#)
- [model customizations](#); and
- [maintaining model performance](#).

In practice, mature customers have ensured:

- Model design decisions are recorded to meet objectives and use cases.
- Performance evaluation approaches and thresholds are included to assess the model's ability to perform under a range of conditions in accordance to its objectives and use cases
- A model validation framework is used which determines the priority and frequency of validation (and revalidation on deployed models), as well as the depth of review expected of validators.
- Implementation of AI-specific system assessment controls, such as [NIST AI Risk Management Framework, Measure 2.1-2.13](#) for model evaluation

AWS Services

- [Amazon Bedrock](#) provides access to a [wide range of foundational models](#), allowing users to select the most appropriate model for their specific use case (e.g. Deepseek, Anthropic, AI21 Labs etc). [Amazon Bedrock](#) also facilitates customization through techniques like fine-tuning and [RAG](#).
- [Amazon Bedrock Guardrails](#) provides configurable safeguards to help mitigate model limitations, and consistently applies this across all supported foundation models:

[AWS CAF for AI](#)

[ISO42001: AI Management Systems](#)

[AWS Well-Architected Generative AI Lens](#)

[Amazon SageMaker AI User Guide](#)

[Amazon Bedrock User Guide](#)

[Amazon Sagemaker AI Autopilot User Guide](#)

[Amazon Sagemaker Model Monitor User Guide](#)

[Amazon Sagemaker MLOps Developer Guide](#)

[Amazon Sagemaker Unified Studio Developer Guide](#)

[Amazon Bedrock Evaluations User Guide](#)

[Amazon Sagemaker Experiments Developer Guide](#)

[Amazon Sagemaker Model Registry user guide](#)

- Uses Automated Reasoning to help prevent factual errors from hallucinations
 - Blocks up to 85% more undesirable and harmful content
 - Filters over 75% hallucinated responses from models for Retrieval Augmented Generation (RAG) and summarization use cases
 - Redact sensitive information such as PII to protect privacy
- Includes [Automated Reasoning](#) to help mathematically validate the accuracy of responses generated by large language models (LLMs) and prevent factual errors from hallucinations
- [Amazon Sagemaker AI](#) contains various machine learning frameworks (e.g. TensorFlow, PyTorch), pre-trained models that can be easily selected from the Amazon Sagemaker catalog and you can [bring your own model](#).
- [Amazon Sagemaker AI Autopilot](#) automates model selection, hyperparameter tuning, and deployment specific to your use case. It also provides insights into the model's decision-making process, helping you understand the key features and factors influencing the model's predictions for explainability.
- Evaluating foundational model performance involves considering factors like accuracy, relevance, and safety. Both [Amazon Sagemaker](#) and [Amazon Bedrock](#) have built-in performance evaluation metrics, such as tracking and comparison of different model training runs, evaluation of models under various conditions and hyperparameter settings. [Amazon Sagemaker Model Monitor](#) continuously monitors deployed models for data drift and model quality degradation, alerting users when performance falls below defined thresholds.
- Use cases and objectives for model selection should be documented. Both Amazon Bedrock and [Amazon Sagemaker](#) have [model evaluation](#) tools and [model cards](#) specifically designed to evaluate the models, record model details, intended uses, risk ratings, and evaluation results. [Amazon Sagemaker MLOps](#) also enables documentation of the full model lifecycle including training, deployment, and monitoring decisions.
- [Amazon SageMaker Unified Studio](#) provides a unified development environment that enables users to centralize their machine learning workflows, including data preparation, model training, and evaluation.

[Amazon Sagemaker Clarify Developer Guide](#)

- [Amazon Bedrock evaluations](#) help evaluate the performance and effectiveness of Amazon Bedrock models and knowledge bases. Amazon Bedrock can compute performance metrics such as the semantic robustness of a model and the correctness of a knowledge base in retrieving information and generating responses. Automatic evaluations, including evaluations that leverage Large Language Models (LLMs), produce computed scores and metrics that help you assess the effectiveness of a model and knowledge base. Additionally, Amazon [Bedrock Model Evaluation's LLM-as-a-judge capability](#) allows you to select quality metrics such as correctness, completeness, and professional style and tone, as well as responsible AI metrics such as harmfulness and answer refusal. You can evaluate all available models on Amazon Bedrock, including serverless models, Bedrock Marketplace models compatible with Converse API, customized and distilled models, imported models, and model routers. You can also evaluate any model or system hosted anywhere by bringing your own inference responses you already fetched into your input prompt dataset for the evaluation job ("bring your own inference responses"). Results can be compared across evaluation jobs.
- [Amazon Sagemaker with ML flow](#) provides a structured way to define and track your machine learning experiments, including the parameters, data, and evaluation metrics used in model training. This creates a comprehensive record of the model development process that can be reviewed by validators.
- [Amazon Sagemaker model registry](#) provides versioning capabilities for models including capture of metadata of each version.
- [Amazon Sagemaker clarify](#) aids in fairness and model explainability as it provides capabilities for detecting and mitigating bias in the model.

AI Agent Management

Organizations understand and manage the risks posed by adoption and use of AI agents

Considerations

An AI agent is a software program that can interact with its environment, collect data, and use the data to perform self-determined tasks to meet predetermined goals. Humans set goals, but an AI agent independently chooses the best actions it needs to perform to achieve those goals. For further detail, refer to [What are AI Agents?](#)

AI agents can orchestrate interactions between foundation models (FMs), data sources, software applications, and user conversations. In addition, agents can automatically call

[What are AI Agents?](#)

[AWS Well-Architected
Generative AI Lens](#)

[Amazon Bedrock Agents
User Guide](#)

APIs to take actions and invoke knowledge bases to supplement information for these actions.

AWS recommends customers assess their AI deployments against AWS best practices as recommended within the [AWS Well-Architected Generative AI Lens](#), particularly the practices regarding:

- [preventing excessive agency](#)
- [cost informed agents](#)
- [distributed availability](#); and
- [enable tracing for agents](#)

In practice, mature customers have:

- Enabled comprehensive logging and observability and setup a monitoring workflow to continuously analyze logs by enabling [Amazon Bedrock model invocation logging](#) and setting up a [custom observability solution](#)
- Tracked agent's step-by-step reasoning process using [Amazon Bedrock Agent traces](#)
- Optimized model selection for cost and performance. They assessed available FMs to select the best one for their applications based on cost, latency, and accuracy requirements
- Implemented robust testing frameworks using tools such as the [AWS Labs Agent Evaluation](#) to assess agent behavior against predefined criteria
- Implemented robust confirmation mechanisms for critical actions in agents' workflow. For example, they have embedded clearly stated in instructions that the agent should ask for user confirmation before running certain functions, especially those that modify data or perform sensitive operations
- Implemented least privilege access and encryption. For example, they have used [customer managed keys](#) within [Amazon Key Management System \(KMS\)](#) to encrypt the agent's resources and confirmed [Amazon Identity and Access Management \(IAM\)](#) privileges limit the agent to only have access to required resources and actions. Additionally, fine grained access controls were implemented using [Amazon Verified Permissions](#)
- Applied [Amazon Bedrock Guardrails](#) to the agent to avoid sensitive topics, filter user input and agent output from harmful content, and redact sensitive information to protect user privacy

[Best practices for building robust generative AI applications with Amazon Bedrock Agents](#)

[AWS sample customer observability solution](#)

[Amazon Bedrock Agent Trace User Guide](#)

[AWS Labs Agent Evaluation](#)

[AWS Machine Learning Blog: Design secure generative AI application workflows with Amazon Verified Permissions and Amazon Bedrock Agents](#)

[Amazon Bedrock User Guide: Associate guardrails with your agent](#)

[Amazon Bedrock User Guide: Use multi-agent collaboration with Amazon Bedrock Agents](#)

AWS Services

- [Amazon Bedrock model invocation logging](#). You can use model invocation logging to collect invocation logs, model input data, and model output data for all invocations in your AWS account used in Amazon Bedrock in a region.
- [Amazon Bedrock Agent traces](#). Each response from an Amazon Bedrock agent is accompanied by a *trace* that details the steps being orchestrated by the agent. The trace helps you follow the agent's reasoning process that leads it to the response it gives at that point in the conversation.
- [AWS Labs Agent Evaluation](#) is an open source tool that can be used to evaluate an AI agent's responses by simulating concurrent, multi-turn conversations. This tool supports AWS AI services such as Amazon Bedrock, Amazon Q Business and Amazon Sagemaker AI, as well as bring your own agent.
- [AWS KMS Customer managed keys](#) are KMS keys in your AWS account that you create, own, and manage. You have full control over these KMS keys, including establishing and maintaining their key policies, IAM policies, and grants, enabling and disabling them, rotating their cryptographic material, adding tags, creating aliases that refer to the KMS keys, and scheduling the KMS keys for deletion
- [AWS Samples: Creating Agent with Guardrails for Amazon Bedrock integration](#)
- [Amazon Bedrock Guardrails](#) provides configurable safeguards to help mitigate model limitations, and consistently applies this across all supported foundation models
- [Amazon Bedrock Multi-Agent Collaboration](#) enables multiple Amazon Bedrock Agents to collaboratively plan and solve complex tasks. With multi-agent collaboration, you can quickly assemble a team of agents that can break down tasks, assign specific tasks to domain specialist sub-agents, work in parallel, and leverage each other's strengths, which leads to more efficient problem-solving. Multi-agent provides a centralized mechanism for planning, orchestration, and user interaction for your generative AI applications. You can [integrate Amazon Bedrock Agents with open source orchestration frameworks](#) such as [LangGraph](#) and [CrewAI](#) for dispatching and reasoning.

Next Steps

Each organization's AI adoption journey is unique; and so, customers need to understand their organization's current state, the desired target state, and the transition required to achieve the target state to manage the AI adoption successfully. Knowing this helps customers set goals and create work streams that helps their staff to thrive in the cloud.

For customers, the recommended next steps include the following:

- Assessment of current state vs. AWS customer considerations and then addressing identified gaps
- Evaluate cloud deployments ("in the cloud") against the [AWS Well-Architected Framework](#) and the [AWS Well-Architected Framework - Machine Learning Lens](#) to build secure, high-performing, resilient, and efficient AI/ML workloads.
- Consult AWS guidance such as:
 - [AWS Responsible Use of AI Guide](#)
 - [AWS Cloud Adoption Framework for Artificial Intelligence, Machine Learning and Generative AI](#)
 - [AWS Well-Architected Generative AI Lens](#)
 - [AWS Well-Architected Framework - Machine Learning Lens](#)
 - [AWS Machine Learning Blog: Build safe and responsible generative AI applications with guardrails](#)
 - [AWS Security Blog: Securing Generative AI Security Scoping Matrix](#)
 - [AWS Machine Learning Blog: Design secure generative AI application workflows with Amazon Verified Permissions and Amazon Bedrock Agents](#)
- Review the use of AWS services and features against existing or planned AI use cases workloads:
 - [Amazon Bedrock User Guide](#)
 - [Amazon Bedrock Agent Trace User Guide](#)
 - [Amazon Bedrock Data Automation User Guide](#)
 - [Amazon Bedrock Evaluations User Guide](#)
 - [Amazon Bedrock Guardrails User Guide](#)
 - [Amazon Bedrock User Guide: Associate guardrails with your agent](#)
 - [Amazon SageMaker AI User Guide](#)
 - [Amazon SageMaker Clarify Developer Guide](#)
 - [Amazon SageMaker Data Preparation Developer Guide](#)

- [Amazon Sagemaker Experiments Developer Guide](#)
 - [Amazon Sagemaker MLOps Developer Guide](#)
 - [Amazon Sagemaker Model Monitor](#)
 - [Amazon Sagemaker Model Registry user guide](#)
 - [Amazon Sagemaker pipelines](#)
 - [Amazon Sagemaker Unified Studio Developer Guide](#)
 - [AWS AI service cards](#)
 - [AWS AI Training and Enablement](#)
 - [AWS Enterprise Support](#)
 - [AWS Labs Agent Evaluation](#)
 - [AWS sample customer observability solution](#)
 - [Amazon Bedrock Multi-Agent Collaboration](#)
- Review the use of AWS Game Days, Security Incident Response Simulation and other practical testing exercises to validate and optimize the operational resilience of cloud deployments.
 - Engage with [AWS Enterprise Support, especially](#) if planning to adopt high risk AI use case. AWS Enterprise Support will effectively manage, monitor, analyze, and report on usage of AWS services, as well as receive proactive planning, architectural reviews, and consultative guidance from AWS.
 - Contact your AWS representative to discuss how the AWS Partner Network, and AWS Solution Architects, Professional Services teams, and training instructors can assist with your AI adoption journey. -

Appendix 1: Overview of emerging AI regulation and guidance by Geography (non-exhaustive)

USA:

- the [Executive Order](#) on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, 2023 (Federal)
- the [Algorithmic Accountability Act of 2023](#) (Federal)
- the National Institute of Standards and Technology's [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#) (Federal), 2023
- FINRA [Regulatory Notice 24-09](#)
- the [Safe and Secure Innovation for Frontier Artificial Intelligence Models Act of 2024](#) (California)

UK:

- the Department of Science, Innovation and Technology's [Implementing the UK's AI regulatory principles: initial guidance for regulators](#) policy paper, 2024
- the Competition and Markets Authority's [AI Foundation Models: Initial report](#), 2023
- the Financial Conduct Authority's [AI Update](#), 2024

Europe:

- [EU AI Act](#)

Asia-Pacific

- the National Artificial Intelligence Centre's [Voluntary AI Safety Standard](#) (Australia), 2024
- the Monetary Authority of Singapore's [Principles to Promote Fairness, Ethics, Accountability and Transparency \(FEAT\) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector](#), (Singapore) 2019
- the Hong Kong Monetary Authority's [Generative Artificial Intelligence in the Financial Services Space](#) (Hong Kong), 2024
- the proposed Act on Promotion of the AI Industry and Framework for Establishing Trustworthy AI (Korea), 2023
- the [AI Guidelines for Business](#) (Japan), 2024
- the [Guide on AI Governance and Ethics](#) (ASEAN), 2024

Appendix 2: Differences between AI, ML, Deep Learning and Generative AI

AI

AI is an umbrella term for different strategies and techniques for making machines more human-like. It includes everything from self-driving cars to robotic vacuum cleaners and smart assistants like Alexa. While machine learning and deep learning fall under the AI umbrella, not all AI activities are machine learning and deep learning. For example, generative AI demonstrates human-like creative capabilities and is a very advanced form of deep learning (refer to [What's the Difference Between AI and Machine Learning?](#) for further details).

ML

While you may see the terms AI and ML being used interchangeably in many places, ML is technically one among many other branches of artificial intelligence. It is the science of developing algorithms and statistical models to correlate data. Computer systems use machine learning algorithms to process large quantities of historical data and identify data patterns. In the current context, machine learning refers to a set of statistical techniques called machine learning models that you can use independently or to support other more complex AI techniques (refer to [What is Deep Learning](#) for further details)

Deep Learning

Deep learning takes machine learning one step further. Deep learning models use neural networks that work together to learn and process information. They comprise millions of software components that perform micro-mathematical operations on small data units to solve a larger problem. For example, they process individual pixels in an image to classify that image. Modern AI systems often combine multiple deep neural networks to perform complex tasks like writing poems or creating images from text prompts (refer to [What is Deep Learning](#) for further details)

Generative AI

An emerging capability within deep learning is generative AI, which makes AI generate or create new, potentially original content. This innovative sub-discipline is increasingly being recognized for its ability to produce outputs that mimic aspects of human-like thought and reasoning capabilities.

Advances in computing power, data availability, and algorithmic innovation have made generative AI possible, paving the way for a wide range of applications, from entertainment and art to scientific research (refer to [What is Generative AI](#) for further details).




Figure 1: Taxonomy of artificial intelligence, machine learning, deep learning, and generative AI

Appendix 3: AWS Responsible AI Policy

This appendix is the AWS Responsible AI Policy as of April 2025. Please always check [online for the latest version](#).

The AWS Responsible AI Policy (“Policy”) applies to your use of artificial intelligence and machine learning Services, features, and functionality (including third-party models) that we provide (collectively, “AI/ML Services”). This Policy supplements the [AWS Acceptable Use Policy](#) and [AWS Service Terms](#).

Prohibitions. You may not use, or facilitate or allow others to use, the AI/ML Services:

- for intentional disinformation or deception;
- to violate the privacy rights of others, including unlawful tracking, monitoring, and identification;
- to depict a person’s voice or likeness without their consent or other appropriate rights, including unauthorized impersonation and non-consensual sexual imagery;
- for harm or abuse of a minor, including grooming and child sexual exploitation;
- to harass, harm, or encourage the harm of individuals or specific groups;
- to intentionally circumvent safety filters and functionality or prompt models to act in a manner that violates our Policies;
- to perform a lethal function in a weapon without human authorization or control.

Responsible AI Requirements. If you use the AI/ML Services to make consequential decisions, you must evaluate the potential risks of your use case and implement appropriate human oversight, testing, and other use case-specific safeguards to mitigate such risks.

Consequential decisions include those impacting a person’s fundamental rights, health, or safety (e.g., medical diagnosis, judicial proceedings, access to critical benefits like housing or government benefits, opportunities like education, decisions to hire or terminate employees, or access to lending/credit, and providing legal, financial, or medical advice). You agree to provide information about your intended uses of the AI/ML Services and compliance with this Policy upon request.

You and your end users are responsible for all decisions made, advice given, actions taken, and failures to take action based on your use of AI/ML Services. AI/ML Services use machine learning models that generate predictions based on patterns in data. Output generated by a machine learning model is probabilistic, and generative AI may produce inaccurate or

inappropriate content. Outputs should be evaluated for accuracy and appropriateness for your use case.

It's your responsibility to ensure you comply with any laws, rules, and regulations applicable to your use of the AI/ML Services. This includes laws specific to artificial intelligence such as the EU AI Act.

AWS may investigate and enforce violations of this Policy as noted in the [AWS Acceptable Use Policy](#). AWS is committed to developing safe, fair, and accurate AI and ML services and providing you with tools and guidance to assist you in building and using AI and ML applications responsibly, see our [Responsible Use of AI and ML page](#) for additional tools and resources.

Appendix 4: Customer Case Studies

AWS FSI customers have increased agility, optimized costs, and accelerated innovation [using AWS and generative AI](#). Here are some of them:

NatWest enhances personalization for its customers using Generative AI on AWS



“

We’ve chosen to build on our existing strategic relationship with AWS and develop AI-powered financial products with a trusted collaborator who understands how we work with data at scale to keep our customers safe and secure.

Scott Marcar
CIO of NatWest Group

OVERVIEW

NatWest Group sought to leverage Generative AI to more rapidly drive toward its goal of helping 10 million people manage their financial wellbeing per year by the end of 2027.

SOLUTION

NatWest will enhance its existing ML personalization products, including tools like NatWest Digital Financial Health Check and the Know Your Credit Score service by working closely with AWS to co-create responsible AI products on top of foundation models (FMs) available through [Amazon Bedrock](#).

IMPACT

Combining AWS’s leading data, analytics, machine learning, and generative AI capabilities with the bank’s 300 years of financial services experience will result in tailored banking experiences that deliver personalized communications, financial education, products, and services to NatWest Group’s customers.

BankUnited provides superior client service with generative AI on AWS



“ Having access to the right information is extremely important for us. It speaks to our reputation with our clients and helps our employees be the advisor that the clients expect.

Jeiner Morales

Director of Operations, Consumer and Small Business Banking, BankUnited

OVERVIEW

Bank United, an leading regional commercial and small business bank, faced challenges locating information from an extensive policy management system containing approximately 400 documents. Employees struggled to find relevant information, hindering their ability to provide prompt and accurate responses to customer inquiries.

SOLUTION

Bank United turned to AWS for its comprehensive set of services, including [Amazon Bedrock](#), [Amazon S3](#), [Amazon Kendra](#), and [AWS Lambda](#). It leveraged AWS's flexible and scalable infrastructure and robust security features to build a sophisticated natural language processing (NLP) solution.

IMPACT

This NLP solution enabled employees to understand and respond accurately to customer inquiries, regardless of how the questions were phrased. Bank United transitioned from a five-team support model to a 24/7 service model at a fraction of the cost. This transformation expedited customer service, ensuring that clients received the right information promptly. The solution not only enhanced the bank's reputation and built trust with its clients but also empowered employees to serve as trusted advisors.

DBS delivers a personalized digital marketing strategy with AWS



“As we continue to evolve ADA at scale across DBS, we’ll continue the partnership with AWS in delivering the functionality needed to enable our insight driven transformation.”

Peng Khim Ng
Managing Director & Head, Institutional Banking Group Technology, DBS

OVERVIEW

DBS wanted to deliver a hyper-personalized SME digital marketing strategy and needed an architecture focused on security and data tokenization for built-in confidentiality features.

SOLUTION

With AWS, DBS is able to run its “Advancing DBS with AI” (ADA) advanced data analytics platform, which enables data scientists to model using [Amazon SageMaker](#) and SME tokenized data.

IMPACT

DBS is able to identify buyer-seller relationships, external SEOs, and SEM optimizers. These insights help DBS enhance customer engagement, brand management, and product management.

NAB speeds enterprise development with Amazon Q Developer



“ We’re still doing migrations to cloud, and my team actively uses Q when they’re working and doing migrations. We’ve used some code transform for a couple of Java 8 to Java 17 upgrades thus far.

Paul Roney
Executive, Core Platforms, National Australia Bank

OVERVIEW

National Australia Bank (NAB) faced the challenge of improving developer productivity and experience across its thousands of developers in Australia, India, and Vietnam. With 84% of its applications already on the cloud, NAB sought to leverage generative AI to further enhance its cloud-first strategy and development processes.

SOLUTION

NAB implemented [Amazon Q Developer](#), starting with a small proof of concept and gradually scaling to 1,000 software developers. They customized Q Developer for their frameworks used in building applications and services, selecting the best reference implementations and repositories to serve as examples for their development teams.

IMPACT

As a result of adopting the service, NAB Developers reported a 41% improvement in productivity and efficiency, while 45% noted enhanced code quality. Key benefits included faster function writing, quicker unit test creation, smarter solution suggestions, and improved documentation. NAB also utilized Q's code scanning capabilities to improve security in their development workflow.

Nasdaq is enhancing its global market surveillance offering with generative AI on AWS



“ By drawing on the latest innovation in cloud technology and artificial intelligence, we can better respond to new threats and offer the global financial system advanced tools to more effectively tackle market abuse. This is a continuous cycle of investment and improvement in our capability, leveraging our unique position as both a world-class market operator and best in class surveillance technology provider around the world.

Tony Sio
Head of Regulatory Strategy and Innovation, Nasdaq

OVERVIEW

Global capital markets continue to grow in complexity. As trading volumes surge, Nasdaq aimed to enhance its market surveillance capabilities to detect increasingly sophisticated patterns of potential abuse and maintain fair and orderly markets. Traditional rules-based systems and human analysts were becoming strained, necessitating a more advanced solution to keep pace with evolving market dynamics.

SOLUTION

Nasdaq leveraged [Amazon Bedrock](#) to create a solution that uses generative AI to streamline the triage and examination process involved in investigating suspected market manipulation and insider dealing, empowering regulators and market operators to more effectively monitor and detect potential market abuse.

IMPACT

Analysts will be empowered to distill, analyze, and interpret relevant information more quickly, enhancing their ability to form detailed initial assessments of alerts. For example, the technology can produce a consolidated table of the company's regulatory filings, summaries and links to company, sector, and peer company news, news sentiment analysis, and other impactful factors that may impact any given security. During proof-of-concept testing, surveillance analysts estimated a 33% reduction in investigation time. Nasdaq is planning to leverage this generative AI solution in its U.S. equity market surveillance.