aws FOR FINANCIAL SERVICES

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

April 2026

# 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.

# Contents

# Abstract

This document provides information regarding governance, risk, and compliance (GRC) considerations for the adoption of responsible 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 can be addressed by AWS services, mechanisms and guidance.

# Updates from previous version (May 2025)

New information:

- OECD G7 Hiroshima AI Process voluntary reporting framework
- Updated AWS ISO 42001 certificate
- Mapping of AWS Well Architected Framework: Responsible AI Lens against AWS responsible AI dimensions
- The role of evaluation within AI governance
- AWS frontier agents
- Sustainability considerations

New prescriptive guidance:

- AWS Well-Architected Framework: Responsible AI Lens
- AWS Security Blog: AI lifecycle risk management: ISO/IEC 42001:2023 for AI governance
- AWS Security Blog: Enabling AI adoption at scale through enterprise risk management framework – Part 1
- AWS Security Blog: Enabling AI adoption at scale through enterprise risk management framework – Part 2

New AWS services and features:

- AWS Well Architected Tool – Generative AI Lens
- AWS Unified Operations
- Guidance for Multi-Provider Generative AI Gateway on AWS (reference architecture)
- AWS Audit Manager – Generative AI Best Practices Framework v2
- AWS Clean Rooms – Synthetic dataset generation
- Amazon Bedrock Guardrails – Automated Reasoning checks
- Amazon Bedrock AgentCore
- Amazon Bedrock AgentCore Evaluations
- AWS Customer Carbon Footprint Tool

# Introduction

The rapid growth of AI brings promising new innovations and raises new challenges. At Amazon Web Services (AWS), we're committed to developing AI responsibly, taking a people-centric approach that prioritizes education, science, and our customers to integrate responsible AI across the entire AI lifecycle.

Of particular significance for Financial Services Industry (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 ISO 42001 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 (see **Appendix 1** for more information about the differences between AI, machine learning, 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 because the board and senior management of financial institutions are ultimately accountable for their activities, including AI use cases[1]. Effective AI governance aims to establish standards and guidelines that promote ethical AI practices, protect individual rights, and mitigate risks associated with AI[2]

[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 and International Electrotechnical Commission (ISO/IEC) 42001 Artificial Intelligence Management System standard states that AI raises the following specific considerations[3]:

- 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 behavior during use. They require special consideration to ensure their responsible use continues with changing behavior.

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

In addition, FSI customers face additional adoption challenges because of 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 (see Appendix 1 for an overview of emerging regulation and guidance and Appendix 4 for AWS customer case studies).

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.

# Responsible AI, the AWS Frontier Model Safety Framework and the OECD G7 Hiroshima AI Process voluntary reporting framework

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

We define the core dimensions of responsible AI as:

[3] ISO/IEC 42001: Information technology – Artificial intelligence – Management system, 2023

| | |
|---|---|
| **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've 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've taken further steps to foster the safe development of future advanced AI models—referred to as frontier models. The Korea Frontier AI Safety Commitments[4] define frontier models as highly capable general-purpose AI

[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 May 21, 2024. Twenty-seven countries and the European Union signed the Seoul Ministerial Statement on May 22, 2024. For more information, see https://www.industry.gov.au/publications/seoul-declaration-countries-attending-ai-seoul-summit-21-22-may-2024

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 frontier models and democratize access to the benefits of AI, we also take responsibility for mitigating the risks of our technology. Consistent with the endorsement of the Korea Frontier AI Safety Commitments by Amazon, this framework outlines the protocols we follow to help ensure that frontier models developed by Amazon don't 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. See Amazon's Frontier Model Safety Framework for more information about 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.

**The OECD G7 Hiroshima AI Process (HAIP) Voluntary Reporting Framework**

As part of the G7 Hiroshima AI Process, the G7 launched a voluntary reporting framework to encourage transparency and accountability among organizations developing advanced AI systems.

Organizations will be able to provide comparable information on their AI risk management actions and practices—such as risk assessment, incident reporting, and information sharing mechanisms—fostering trust and accountability in the development of advanced AI systems.

Amazon published its G7 Hiroshima AI Process (HAIP) Transparency Report in November 2025. This report demonstrates the commitment Amazon has made to responsible AI development and transparency.

# 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, in November 2024, AWS was 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 ISO 42001 AI Management System standard outlines best practices and controls for the responsible development, deployment, and operation of AI systems. The basis of this certification is the establishment of an AI management system that aims to ensure 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 compliance with privacy laws and AI security standards.

As of November 2025, AWS successfully completed its first annual surveillance audit against ISO 42001 with no major or minor non-conformities. The updated ISO 42001 certificate has been published and is available for customers to download on AWS Artifact.

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:** A code of practice that focuses on protecting personal data in the cloud. It's based on the ISO information security standard 27002 and provides implementation guidance on ISO 27002 controls that are applicable to personally identifiable information (PII) in the cloud. 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 with this standard provides assurance of the 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, in addition to

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, & Compliance webpage 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 ISO 42001 certification, SOC reports, 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. See **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, and Governance and transparency. The following table outlines the key considerations for each dimension.

| Dimensions of responsible AI | Key considerations | Additional resources |
|---|---|---|
| **Fairness:**<br><br>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.<br>• **Regulatory compliance**: Identify and ensure compliance with applicable laws related to financial fairness (for example, the Equal Credit Opportunity Act), which prohibits discrimination based on protected characteristics. This involves continuous auditing and reporting to demonstrate fairness in AI-driven decisions.<br>• **Stakeholder engagement**: Engage with diverse stakeholders to understand potential impacts on different groups and incorporate feedback into AI development processes. | AWS Well-Architected Responsible AI Lens: RAIBR02-BP01 Identify potential harmful events impacting fairness<br><br>AWS Well-Architected Machine Learning Lens: MLOPS02-BP06 Review fairness and explainability |

| Dimensions of responsible AI | Key considerations | Additional resources |
|---|---|---|
| **Explainability:**<br><br>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.<br>• **Documentation and auditing**: Maintain detailed documentation of AI decision-making processes and conduct regular audits to ensure transparency and accountability.<br>• **Model interpretability**: Use interpretable models and provide clear explanations to stakeholders, enabling them to understand AI-driven outcomes and assess potential risks. | AWS Well-Architected Responsible AI Lens: RAIBR02-BP07 Identify potential harmful events impacting explainability<br><br>AWS Well-Architected Machine Learning Lens:<br><br>MLOPS02-BP06 Review fairness and explainability<br><br>MLPERF06-BP02 Evaluate model explainability |
| **Privacy and security:**<br><br>Appropriately obtaining, using, and protecting data and models | • **Data protection**: Ensure that AI systems handle personal data securely and in compliance with privacy laws such as GDPR. Implement robust data encryption and access controls to protect sensitive information.<br>• **Data governance:** Establish clear policies for data collection, storage, and use. Ensure that data is accurate, relevant, and used only for intended purposes.<br>• **Risk assessment**: Conduct thorough risk assessments to identify potential privacy and security vulnerabilities in AI systems and implement measures to mitigate these risks. | AWS Well-Architected Responsible AI Lens:<br><br>RAIBR02-BP04 Identify potential harmful events impacting privacy<br><br>RAIBR02-BP06 Identify potential harmful events impacting system and data security |

| Dimensions of responsible AI | Key considerations | Additional resources |
|---|---|---|
| **Safety:**<br>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.<br>• **Risk management**: Continuously monitor AI systems for safety risks and update risk management strategies as needed to ensure the integrity of financial operations. | AWS Well-Architected Responsible AI Lens:<br>RAIBR02-BP05 Identify potential harmful events impacting safety |
| **Controllability:**<br>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.<br>• **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.<br>• **Feedback loops**: Establish feedback loops to improve AI performance and adapt to changing conditions, ensuring that AI systems remain aligned with organizational goals. | AWS Well-Architected Responsible AI Lens:<br>RAIUC04-BP01 Map the user journey to identify AI interaction requirements<br>RAIUC04-BP02 Identify human oversight opportunities<br>RAIRC03-BP07 Measure user controllability of system behavior<br>RAIMON02-BP01 Create feedback loops to apply monitoring results to system improvement |

| Dimensions of responsible AI | Key considerations | Additional resources |
|---|---|---|
| **Veracity and robustness:**<br>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.<br>• **Data quality**: Ensure that datasets used for training AI models are of high quality, diverse, ad representative to prevent biases and inaccuracies.<br>• **Continuous improvement**: Regularly update AI models and documentation to maintain their effectiveness and compliance with evolving regulatory standards. | AWS Well-Architected Responsible AI Lens:<br>RAIBR02-BP02 Identify potential harmful events impacting veracity<br>RAIBR02-BP03 Identify potential harmful events impacting robustness |
| **Governance:**<br>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.<br>• **Supply chain management**: Manage the AI supply chain effectively, ensuring that all providers and deployers adhere to best practices and regulatory requirements. | AWS Well-Architected Responsible AI Lens:<br>RAIUC05-BP01 Engage your organization in approving your use case<br>Assess risks<br>Mitigate risks<br>Monitoring |
| **Transparency:**<br>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.<br>• **Documentation and disclosure**: Maintain detailed records of AI decision-making processes and disclose this information to regulators and stakeholders as required.<br>• **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 Well-Architected Responsible AI Lens:<br>RAIBR02-BP08 Identify potential harmful events impacting transparency |

# AWS prescriptive guidance for responsible AI

AWS recommends that customers 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 Responsible AI Lens
- AWS Well-Architected Generative AI Lens
- AWS Well-Architected Machine Learning Lens
- AWS Well-Architected Financial Services Industry Lens (updated for generative and agentic AI)
- AWS Machine Learning Blog: Build safe and responsible generative AI applications with guardrails
- AWS Security Blog: AI lifecycle risk management: ISO/IEC 42001:2023 for AI governance
- AWS Security Blog: Enabling AI adoption at scale through enterprise risk management framework – Part 1
- AWS Security Blog: Enabling AI adoption at scale through enterprise risk management framework – Part 2
- AWS Security Blog: An introduction to the Generative AI Security Scoping Matrix
- AWS Agentic AI Security Scoping Matrix

AWS recommends customers assess their workloads against best practices as recommended within the AWS Well-Architected Responsible AI Lens, Well-Architected Generative AI Lens, Well-Architected Machine Learning Lens, and AWS Well-Architected Financial Services Industry Lens (updated for generative and agentic AI).

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.

# AWS customer considerations

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

**Expectations** reference 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** related to 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 can 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.**

See **Appendix 4** for AWS customer case studies.

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| **Governance**<br><br>Organizations implement principles, policies, tools and processes to help ensure the responsible development, deployment, and monitoring of AI systems in line with strategy and risk appetite<br><br>**Aligned to:**<br><br>**ISO 42001 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<br><br>**EU AI Act:** Articles 9.8, 14, 26, 22.1, 22.2, 26.3<br><br>**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<br><br>**Hong Kong Monetary Authority (HKMA) Consumer Protection in respect of Use of GenAI:** Principle 1 – Governance and Accountability<br><br>**Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models:** Core Principle 1 – Senior Management Responsibilities | **Considerations**<br><br>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 helps ensure 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, in addition to managing the different cost patterns of AI workloads.<br><br>AWS recommends customers consult the AWS CAF for AI: Governance perspective: Managing an AI-driven organization and AWS Security Blog: AI lifecycle risk management: ISO/IEC 42001 for AI governance to inform the design and operation of their AI governance and control frameworks.<br><br>In practice, mature customers have:<br><br>• Business and technology strategies that have been updated to incorporate AI adoption, or a separate documented AI strategy that maps to business and technology strategy outcomes.<br>• Strategies that have been approved by their Board.<br>• Policies and procedures that have been developed and implemented to provide pathways for AI adoption to progress from idea to proof of concept to production environments.<br>• A clear risk appetite for AI adoption that has been set, approved, and communicated by the Board or Board delegated committee.<br>• A Board or Board-delegated committees charged with overseeing progress against strategy in line with risk appetite and approving high risk AI adoption.<br>• AI oversight governance forums that are cross-functional and include data and analytics, risk management, legal, compliance, technology and audit functions.<br>• AI governance practices that include evaluations such as pre-deployment validation, ongoing monitoring against thresholds, model benchmarks, and reporting to AI governance authorities.<br>• Individuals responsible for AI governance are appropriately trained, skilled, and certified. | AWS CAF for AI<br><br>ISO 42001: AI Management System<br><br>AWS Security Blog: AI Lifecycle risk management: ISO/IEC 42001 for AI Governance<br><br>AWS Well-Architected Generative AI Lens<br><br>Amazon SageMaker AI User Guide<br><br>Amazon SageMaker Model Monitor<br><br>Amazon Bedrock evaluations<br><br>Amazon SageMaker Pipelines<br><br>AWS Well Architected Tool – Responsible AI Lens<br><br>AWS Well Architected Tool – |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | • Crisis management plans and Incident response plans that have been updated to include responses for scenarios involving high risk AI.<br><br>**AWS services**<br><br>Customers should consider using the following AWS services to support AI governance:<br><br>• AWS Well Architected Tool – Responsible AI Lens sets out thoughtful questions and corresponding best practices to help builders address responsible AI concerns throughout development and operation. Based on our experience helping customers run hundreds of thousands of AI workloads and on the experience of responsible AI scientists, this lens provides clear, actionable guidance throughout the AI lifecycle.<br><br>• AWS Well Architected Tool – Generative AI Lens provides a consistent approach for customers to evaluate architectures that use large language models (LLMs) to achieve their business goals. This lens addresses common considerations relevant to model selection, prompt engineering, model customization, workload integration, and continuous improvement.<br><br>• Multi-Provider Generative AI Gateway on AWS (reference architecture) aims to simplify integration while providing access to tools that track LLM usage, manage costs, and implement crucial governance features. This allows straightforward switching between models, efficient management of multiple LLM services within applications, and robust control over security and expenses. This framework helps customers explain AI costs by team and use case, identify which models process sensitive data, have audit trails for all AI requests, and route intelligently between providers based on cost or performance.<br><br>• AWS Trusted Advisor draws on 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.<br><br>• AWS Unified Operations is recommended for customers planning to operate mission critical, high risk AI services on | Generative AI Lens<br><br>AWS Unified Operations<br><br>AWS AI training |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | AWS. Unified Operations assists customers by combining proven expertise with AI-powered insights to help reduce operational and security risks, resolve issues faster, and help architect more resilient cloud solutions from the start. Unified Operations provides customers with planning, evaluations and testing, deployment and operations, response and recovery, and continuous improvement. Unified Operations also provides customers with a 5-minute response for critical incidents raised with incident detection and response.<br><br>• AWS AI training is recommended for customers. Contact your AWS account team to understand options for relevant training and certification. | |
| **Risk management**<br><br>Organizations implement principles, policies, tools, and processes to minimize the potential negative impacts in executing their AI strategy<br><br>**Aligned to:**<br><br>**ISO 42001 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<br><br>**EU AI Act:** Articles 8.1, 8.2, 9.3-9.6, 9.9, 17.1, 27.1<br><br>**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<br><br>**Hong Kong Monetary Authority (HKMA)** | **Considerations**<br><br>AWS recommends customers consult the following sources to inform the design and operation of their AI risk management activities:<br><br>• AWS CAF for AI: Risk Management<br>• ISO 42001: AI Management Systems (section 6.1.2),<br>• ISO 23894: Artificial intelligence – Guidance on risk management<br>• NIST AI Risk Management Framework (NIST AI RMF)<br>• NIST Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile<br>• AWS Well Architected Framework: Responsible AI Lens – Benefits and risks<br>• AWS Security Blog: Enabling AI adoption at scale through enterprise risk management framework – Part 1 and Part 2<br><br>AWS recommends customers address AI specific risks and controls in their risk management activities as follows:<br><br>• Consult the AWS Responsible AI core dimensions (such as explainability and safety), AWS Well-Architected Framework: Responsible AI Lens – Benefits and Risks, and NIST AI RMF Generative AI Profile (section 2) to identify AI specific risks that are functional in nature.<br>• Consult the AWS Security Blog: Securing Generative AI: An introduction to the Generative AI Security Scoping Matrix guidance to identify non-functional risks and controls (such as security and availability). | AWS CAF for AI: Risk Management<br><br>ISO 42001: AI Management Systems (section 6.1.2)<br><br>ISO 23894 - Artificial intelligence - Guidance on risk management<br><br>NIST AI Risk Management Framework<br><br>NIST AI RMF: Generative Artificial Intelligence Profile<br><br>AWS Well Architected Framework: Responsible AI Lens – |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| **Consumer Protection in respect of Use of GenAI:** Principle 2 – Fairness<br><br>**Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models:** Core Principle 2 – AI Model Risk Management | • Consult the AWS Security Blog: Threat modelling your generative AI workload to evaluate security risk.<br>• Consult the AWS Agentic AI Security Scoping Matrix for a structured framework to help understand, classify, and secure autonomous AI implementations.<br>• Consult ISO 42001 for reference controls and implementation guidance in Annex A and B respectively.<br>• Use the AWS Generative AI Best Practices Framework v2, a prebuilt standard framework available within AWS Audit Manager to help you gain visibility into how your generative AI implementation on Amazon Bedrock and Amazon SageMaker AI is working against AWS recommended best practices.<br><br>In practice, mature customers have:<br><br>• A clearly defined risk appetite for AI adoption that has been set, approved and communicated by the Board or a Board delegated committee<br>• Established standards and processes for ongoing revalidations of AI in production, with intensity and frequency based on the materiality of the AI use case<br>• All three lines of defense (for example, business or technology owner – Line 1, group risk and compliance teams – Line 2, and internal audit – Line 3) aligned on context, definitions and incoming demand<br>• Risk management activities that follow a tiered approach where high-risk use cases undergo more extensive scrutiny<br>• Either created or augmented AI-specific risk assessment activities within existing frameworks that interact with their corporate governance structures<br>• Individuals responsible for risk management activities for AI use cases who are appropriately trained, skilled, and certified<br><br>**AWS services**<br><br>Customers should consider using the following AWS services to support AI risk management:<br><br>• AWS AI Service Cards support the entire 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, | Benefits and risks<br><br>AWS Security Blog: Securing Generative AI Security Scoping Matrix<br><br>AWS AI Service Cards<br><br>Amazon Bedrock Guardrails User Guide<br><br>Amazon Bedrock evaluations<br><br>Amazon SageMaker Clarify<br><br>AWS Security Blog: Enabling AI adoption at scale through enterprise risk management framework – Part 1 and Part 2<br><br>AWS Audit Manager: Generative AI Best Practices Framework v2<br><br>AWS AI training |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | training details, evaluation metrics, results, observations, and recommendations. 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.<br><br>• AWS Generative AI Best Practices Framework v2 is a prebuilt standard framework available within AWS Audit Manager to help you gain visibility into how your generative AI implementation on Amazon Bedrock and SageMaker AI is working against AWS recommended best practices.<br><br>• AWS AI training is recommended for customers. Contact your AWS account team to understand options for relevant training and certification. | |
| **Compliance**<br><br>Organizations understand and ensure compliance with applicable AI regulations and standards<br><br>**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<br><br>**NIST AI RMF:** Govern 1.1, Map 4.1 | **Considerations**<br><br>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.<br><br>AWS recommends that customers 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.<br><br>In practice, mature customers have:<br><br>• All three lines of defense (for example, business or technology owner – Line 1, group risk and compliance teams – Line 2, and internal audit – Line 3) aligned on context, definitions, and incoming demand<br>• Policies and procedures that have been updated or created to outline pathways for AI adoption to progress from idea to proof of concept to production environments<br>• Compliance activities that follow a tiered approach where high-risk use cases undergo more extensive scrutiny | AWS CAF for AI: Security perspective – Compliance and assurance<br><br>NIST AI RMF<br><br>Getting Started with AWS Artifact<br><br>AWS AI training |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | • Processes for reporting incidents, safety issues, and non-compliance to relevant authorities and have defined affected stakeholders<br>• Processes to ensure that appropriate disclosures are made for both external and internal facing AI user cases<br><br>**AWS services**<br><br>Customers should consider using the following AWS services to support AI compliance activities:<br><br>• For assurance over AWS global infrastructure, customers can rely on the independent assurance made available at no cost through AWS Artifact. AWS ISO 42001 certification and SOC2 assurance of Amazon AI services are particularly relevant.<br><br>• AWS AI training is recommended for customers. Contact your AWS account team to understand options for relevant training and certification. | |
| **Data management**<br><br>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.<br><br>**Aligned to:**<br><br>**ISO 42001 AI Management System:** B.7 – Data for AI Systems, B.4.3 Data Resources<br><br>**NIST AI Risk Management Framework: GenAI Profile** MP-2.1-001, MP-2.1-002<br><br>**EU AI Act:** Article 10 – Data and Data Governance | **Considerations**<br><br>AWS recommends customers consult ISO 42001 AI Management, Section B7 for an overview of expectations, controls, and recommended implementation with respect to data management for AI systems.<br><br>AWS recommends customers assess their AI deployments against AWS best practices as recommended within:<br><br>• AWS Well-Architected Framework: Machine Learning Lens – Data processing<br>• AWS Well-Architected Framework: Responsible AI Lens – Dataset planning<br>• AWS Well-Architected Framework: Generative AI Lens – Data architecture<br><br>For assurance about AWS global infrastructure, customers should rely on the independent assurance made available through AWS Artifact, particularly the ISO 42001 certification and SOC2 report, which includes AWS AI services such as SageMaker AI and Amazon Bedrock .<br><br>In practice, mature customers have:<br><br>• In-scope data including training, validation, and testing data | AWS CAF for AI: Data Curation<br><br>ISO 42001: AI Management Systems<br><br>AWS Well Architected Framework: Machine Learning Lens – Data processing<br><br>AWS Well-Architected Framework: Generative AI Lens – Data architecture<br><br>AWS Well-Architected |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| **Monetary Authority of Singapore (MAS):** AI model risk management paper<br><br>**Hong Kong Monetary Authority (HKMA) Consumer Protection in respect of Use of GenAI:** Principle 4 – Data Privacy and Protection<br><br>**Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models:** Core Principle 3 – Cybersecurity and Data Risk Management | • Assessment and documentation for:<br>  ○ Data collection processes and data origin<br>  ○ Data-preparation and transformation processes (such as labelling, cleaning, enrichment)<br>  ○ The availability, quantity, and suitability of the data sets required<br>  ○ Assumptions, potential biases, and mitigating measures<br>  ○ Representativeness and relevance of data<br>  ○ Use of personal data<br>• Data protection controls and data validation checks to confirm the completeness, accuracy, relevance, reliability and quality of data<br><br>**AWS services**<br><br>Customers should consider using the following AWS services to support data management for AI:<br><br>• 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.<br><br>• Amazon SageMaker Unified Studio provides a single data and AI development environment that brings together AWS data, analytics, AI, and machine learning services. It provides a place to build, deploy, execute, and monitor workflows from a single interface, so customers can securely build and share analytics and AI artifacts, including data, models, and generative AI applications.<br><br>• 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 helps ensure trustworthy and accurate insights from unstructured, multi-modal data sources.<br><br>• Amazon SageMaker Canvas facilitates data preparation by providing a visual interface for data selection, cleaning, transformation, and feature engineering. | Framework: Responsible AI Lens – Dataset planning<br><br>Amazon SageMaker AI User Guide<br><br>Amazon Bedrock User Guide<br><br>Amazon SageMaker Model Monitor<br><br>Amazon Bedrock Data Automation<br><br>Amazon SageMaker Data preparation<br><br>AWS Clean Rooms – Synthetic dataset generation |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | AWS Clean Rooms – Synthetic dataset generation trains a model that learns the essential statistical patterns of the original dataset, then generates synthetic records by sampling values from the original dataset and using the model to predict the predicted value column. Rather than merely copying or perturbing the original data, the system uses a model capacity reduction technique to mitigate the risk that the model will memorize information about individuals in the training data. The resulting synthetic dataset has the same schema and statistical characteristics as the original data, making it suitable for training classification and regression models. This approach quantifiably reduces the risk of re-identification.<br>• 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 such as Amazon Bedrock, Amazon Simple Storage Service (Amazon S3), Amazon Elastic Block Store (Amazon EBS) , and Amazon Relational Database Service (Amazon RDS) , AWS Key Management Service (AWS KMS) for key management, automatic encryption and decryption, and integration with AWS Identity and Access Management (IAM) for access control. | |
| **Model management**<br><br>Organizations identify and document objectives and implement processes for the responsible design and development of AI systems.<br><br>**Aligned to:**<br><br>**ISO 42001 AI Management System:** A6 – AI System life cycle, A9 – Use of AI systems<br><br>**NIST AI Risk Management Framework:** GenAI Profile | **Considerations**<br><br>AWS recommends customers consult ISO 42001 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.<br><br>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:<br>• Model selection and cost optimization<br>• Model performance evaluations<br>• Energy efficient models<br>• Model customizations<br>• Maintaining model performance<br><br>In practice, mature customers have: | AWS CAF for AI<br><br>ISO 42001: AI Management Systems<br><br>AWS Well-Architected Generative AI Lens<br><br>Amazon SageMaker AI User Guide<br><br>Amazon Bedrock Documentation<br><br>Amazon SageMaker AI Autopilot<br><br>Amazon SageMaker Model Monitor<br><br>Amazon SageMaker MLOps |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| MAP 3.3, Measure 2, Manage 4.1<br><br>**EU AI Act:** Article 9 – Risk Management Systems, Article 11- Technical Documentation<br><br>**Monetary Authority of Singapore (MAS):** AI model risk management paper, section 6<br><br>**Hong Kong Securities and Futures Commission (SFC) Use of GenAI Language Models:** Core Principle 2 – AI Model Risk Management | • Recorded model design decisions that meet objectives and use cases<br>• Performance evaluation approaches and thresholds to assess the model's ability to perform under a range of conditions in accordance to its objectives and use cases<br>• A model validation framework that's used to determine the priority and frequency of validation (and revalidation on deployed models) and the depth of review expected of validators.<br>• Implemented AI-specific system assessment controls, such as NIST AI Risk Management Framework, Measure 2.1-2.13 for model evaluation<br><br>**AWS services**<br><br>• Amazon Bedrock provides access to a wide range of foundation models, so users can select the most appropriate model for their specific use case (including Deepseek, Anthropic, AI21 Labs, and so on). Amazon Bedrock also facilitates customization through techniques like fine-tuning and Retrieval Augmented Generation (RAG).<br><br>• Amazon Bedrock Guardrails provides configurable safeguards to help mitigate model limitations, and consistently applies them across all supported foundation models:<br>  ○ Uses Automated Reasoning to help prevent factual errors from hallucinations<br>  ○ industry-leading safety protections that block up to 88% of harmful content Filters over 75% of hallucinated responses from models for RAG and summarization use cases<br>  ○ Redact sensitive information such as PII to protect privacy<br>  Includes Automated Reasoning to help mathematically validate the accuracy of responses generated by LLMs and prevent factual errors from hallucinations<br>• Amazon SageMaker AI contains various machine learning frameworks (such as TensorFlow and PyTorch), pre-trained models that can be selected | Amazon SageMaker Unified Studio Developer Guide<br><br>Amazon Bedrock Evaluations<br><br>Amazon SageMaker Experiments<br><br>Amazon SageMaker Model Registry<br><br>Amazon SageMaker Clarify<br><br>Amazon Bedrock Guardrails – Automated Reasoning checks |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | from the SageMaker catalog, and you can bring your own model.<br><br>• 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.<br><br>• Evaluating foundation model performance involves considering factors such as 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 and 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.<br><br>• 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.<br><br>• 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.<br><br>• 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— | |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | including evaluations that use LLMs—produce computed scores and metrics that help you assess the effectiveness of a model and knowledge base. Additionally, you can use the LLM-as-a-judge capability of Amazon Bedrock Model Evaluation to select quality metrics such as correctness, completeness, and professional style and tone, in addition to responsible AI metrics such as harmfulness and answer refusal. You can evaluate all available models on Amazon Bedrock, including serverless models, Amazon 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 that you have already fetched into your input prompt dataset for the evaluation job (*bring your own inference responses*). Results can be compared across evaluation jobs.<br><br>• Automated Reasoning checks in Amazon Bedrock Guardrails mathematically verify natural language content against your defined policies, helping to ensure strict compliance with your guardrails. These checks can help to systematically block harmful or non-compliant content before it reaches your users. Unlike pattern-matching approaches, Automated Reasoning aims to deliver higher accuracy with fewer false positives, particularly for complex policy requirements. For customers prioritizing precision, policy rules can be customized to enhance guardrail effectiveness through clear logic statements.<br><br>• 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.<br><br>• Amazon SageMaker model registry provides versioning capabilities for models including capture of metadata of each version. | |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | • Amazon SageMaker clarify aids in fairness and model explainability by providing capabilities for detecting and mitigating bias in the model. | |
| **AI agent management**<br><br>Organizations understand and manage the risks posed by adoption and use of AI agents | **Considerations**<br><br>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 more information, see What are AI Agents?<br><br>AI agents can orchestrate interactions between foundation models, data sources, software applications, and user conversations. In addition, agents can automatically call APIs to take actions and invoke knowledge bases to supplement information for these actions.<br><br>A newer type of agent, frontier agents, are autonomous systems that work independently to achieve goals, scale massively to tackle concurrent tasks, and run persistently for hours or days without intervention. Unlike traditional AI assistants that help with individual tasks, frontier agents aim to act as extensions of your team, delivering complete outcomes across diverse use cases.<br><br>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:<br>• Preventing excessive agency<br>• Cost-informed agents<br>• Distributed availability<br>• Enabling tracing for agents<br><br>In addition, customers should consult the AWS Agentic AI Security Scoping Matrix for a structured framework to help understand, classify, and secure autonomous AI implementations. | What are AI Agents?<br><br>AWS frontier agents<br><br>AWS Well-Architected Generative AI Lens<br><br>Amazon Bedrock Agents User Guide<br><br>Best practices for building robust generative AI applications with Amazon Bedrock Agents – Part 1<br><br>AWS sample customer observability solution<br><br>Amazon Bedrock Agent Trace<br><br>AWS Labs Agent Evaluation<br><br>Design secure generative AI application workflows with Amazon Verified Permissions and Amazon Bedrock Agents<br><br>Amazon Bedrock: Associate guardrails with your agent |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | In practice, mature customers have:<br>• Enabled comprehensive logging and observability and set up a monitoring workflow to continuously analyze logs by enabling Amazon Bedrock model invocation logging and setting up a custom observability solution.<br>• Tracked agents' step-by-step reasoning processes using Amazon Bedrock Agent traces.<br>• Optimized model selection for energy efficiency, cost, and performance. They assessed available foundation models to select the best one for their applications based on energy efficiency, cost, latency, and accuracy requirements.<br>• Implemented robust testing frameworks using tools such as the AWS Labs Agent Evaluation or Amazon Bedrock AgentCore Evaluations to assess agent behavior against predefined criteria.<br>• Implemented robust confirmation mechanisms for critical actions in agents' workflow. For example, they have embedded clearly stated instructions that the agents should ask for user confirmation before running certain functions, especially those that modify data or perform sensitive operations.<br>• Implemented least privilege access and encryption. For example, they have used customer managed keys within AWS KMS to encrypt agents' resources and confirmed that Amazon Identity and Access Management (IAM) privileges limit agents to only have access to required resources and actions. Additionally, fine grained access controls have been implemented using Amazon Verified Permissions.<br>• 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.<br><br>**AWS services**<br>• Amazon Bedrock AgentCore is an agentic platform for building, deploying, and operating highly effective agents securely at scale using any framework and foundation model. With | Use multi-agent collaboration with Amazon Bedrock Agents<br><br>Amazon Bedrock AgentCore Developer Guide<br><br>Amazon Bedrock AgentCore Policy<br><br>Amazon Bedrock AgentCore Evaluations<br><br>AWS Agentic AI Security Scoping Matrix |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | AgentCore, you can enable agents to take actions across tools and data with the right permissions and governance, run agents securely at scale, and monitor agent performance and quality in production.<br><br>• Amazon Bedrock AgentCore Policy enables developers to define and enforce security controls for AI agent interactions with tools by creating a protective boundary around agent operations.<br><br>• Amazon Bedrock AgentCore Evaluations provides automated assessment tools to measure how well your agent or tools perform specific tasks, handle edge cases, and maintain consistency across different inputs and contexts. The service enables data-driven optimization and helps your agents meet quality standards before and after deployment.<br><br>• 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 an AWS Region.<br><br>• 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 to generate the response it gives at that point in the conversation.<br><br>• 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, in addition to bringing your own agent.<br><br>• AWS KMS Customer managed keys are KMS keys in your AWS account that you create, own, and | |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| | 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.<br><br>• AWS Samples: Creating Agent with Guardrails for Amazon Bedrock integration.<br><br>• Amazon Bedrock Guardrails provides configurable safeguards to help mitigate model limitations, and consistently applies this across all supported foundation models<br><br>• 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 use each other's strengths, which leads to more efficient problem-solving. Multi-agent collaboration 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. | |
| **Sustainability**<br><br>Organizations understand and consider the sustainability impacts of AI adoption<br><br>**Aligns to:**<br><br>OECD AI Principle 1.1 | **Considerations**<br><br>As you scale your use of AI, it's important to also minimize its environmental footprint. An effective way to do this is by moving IT workloads from on-premises infrastructure to AWS. AWS infrastructure is up to 4.1 times more efficient than on-premises, and when workloads are optimized on AWS, the associated carbon footprint can be reduced by up to 99%. For more information, see the AWS Sustainability portal. | AWS Sustainability<br><br>AWS Well-Architected Framework: Generative AI Lens – Sustainability Pillar<br><br>AWS Customer Carbon Footprint Tool |

| Expectations | AWS customer considerations | Resources |
|---|---|---|
| EU AI Act:<br>• Article 59.1(a)(iii)<br>• Article 95.2(b)<br>• Article 112.7<br><br>Australian Sustainability Reporting Standard AASB S2 | With respect to AI adoption, the primary consideration should always be to ask if generative AI is right for your workload. There is no need to use computationally intensive AI when a simpler, more sustainable approach might achieve the same outcome. For instance, when searching for information, search engines can return results using fewer resources than generative AI.<br><br>If generative AI is determined to be the right solution, AWS recommends customers assess their AI deployments against AWS best practices within the AWS Well-Architected Generative AI Lens, particularly the Sustainability pillar, including practices regarding:<br>• Energy efficient infrastructure and services<br>• Sustainable data processing and storage services<br>• Energy efficient models<br><br>In practice, mature customers have:<br>• Embedded consideration of generative AI against alternative existing solutions for proposed workloads within their AI governance mechanisms.<br>• Optimized model selection for energy efficiency, cost and performance. They assessed available foundation models to select the best one for their applications based on energy efficiency, latency, and accuracy requirements.<br>• Established mechanisms for tracking, measuring, and optimizing the carbon emissions of their AWS usage by Region and AWS service (particularly Amazon Bedrock, Amazon Q Business, Amazon Transcribe, Amazon Textract, and Amazon Quick Suite) using the AWS Customer Carbon Footprint Tool (CCFT).<br><br>**AWS services:**<br>• AWS Customer Carbon Footprint Tool (CCFT). The CCFT now includes Scope 3 emissions data alongside existing Scope 1 and 2 emissions, providing the carbon impact of hardware manufacturing, buildings, equipment, and fuel | |

| Expectations | AWS customer considerations | Resources |
| --- | --- | --- |
| | and energy-related activities associated with your AWS usage. | |

# Next steps

Each organization's AI adoption journey is unique, 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 help their staff to thrive in the cloud.

For customers, the recommended next steps include the following:

- Assessment of current state compared to 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 and machine learning (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 – Responsible AI Lens
  - AWS Well-Architected – Generative AI Lens
  - AWS Well-Architected – Machine Learning Lens
  - AWS Well-Architected – Financial Services Industry 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 Security Blog: Agentic 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
  - Amazon Bedrock Data Automation
  - Amazon Bedrock Evaluations
  - Amazon Bedrock Guardrails
  - Amazon Bedrock: Associate guardrails with your agent
  - Amazon Bedrock AgentCore Developer Guide
  - Amazon Bedrock AgentCore Policy
  - Amazon Bedrock AgentCore Evaluations
  - Amazon SageMaker AI User Guide
  - Amazon SageMaker Clarify
  - Amazon SageMaker Data Preparation
  - Amazon SageMaker Experiments
  - Amazon SageMaker – Implement MLOps
  - Amazon SageMaker – Model Monitor
  - Amazon SageMaker Model Registry
  - Amazon SageMaker Pipelines
  - Amazon SageMaker Unified Studio User Guide
  - AWS AI Service Cards
  - AWS AI Training
  - AWS Enterprise Support
  - AWS Labs Agent Evaluation
  - AWS sample customer observability solution
  - Amazon Bedrock multi-agent collaboration

- Review the use of AWS GameDays, security incident response simulations, and other practical testing exercises to validate and optimize the operational resilience of cloud deployments.

- Engage with your AWS Enterprise Support or Unified Operations Team, especially if you're planning to adopt high risk AI use case. AWS Enterprise Support will effectively manage, monitor, analyze, and report on usage of AWS services, and provide proactive planning, architectural reviews, and consultative guidance.

- Contact your AWS representative to discuss how the AWS Partner Network, and AWS Solution Architects, AWS 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 Ensuring A National Policy Framework For Artificial Intelligence, December 2025 (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), 2023 (Federal)
- FINRA Regulatory Notice 24-09
- Colorado Consumer Protections for Artificial Intelligence Act (SB 24-205)
- California Privacy Protection Agency – Automated decision-making technology
- Texas Responsible AI Governance Act (HB 2060)

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 Basic Act on the Development of Artificial Intelligence and the Establishment of Foundation for Trustworthiness ("AI Basic Act") (South Korea), December 2024
- The National Artificial Intelligence Centre's Guidance for AI Adoption (Australia), October 2025
- The National Artificial Intelligence Centre's Voluntary AI Safety Standard (Australia), updated December 2025
- The Monetary Authority of Singapore's Consultation Paper on Proposed Guidelines on Artificial Intelligence Risk Management for Financial Institutions, November 2025
- The Digital Policy Office Hong Kong, Generative Artificial Intelligence Technical and Application Guideline, December 2025
- The Hong Kong Monetary Authority's Generative Artificial Intelligence in the Financial Services Space, 2024
- The Act on Promotion of Research, Development, and Utilization of Artificial Intelligence-related Technologies (Japan), September 2025

- 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 (see What's the Difference Between AI and Machine Learning? for more information).

### Machine learning

While you might see the terms AI and machine learning (ML) being used interchangeably in many places, machine learning is one among many subsets of AI. It's 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 (see What is Deep Learning for more information).

### 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 (see What is Deep Learning for more information).

### 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 (see What is Generative AI for more information).

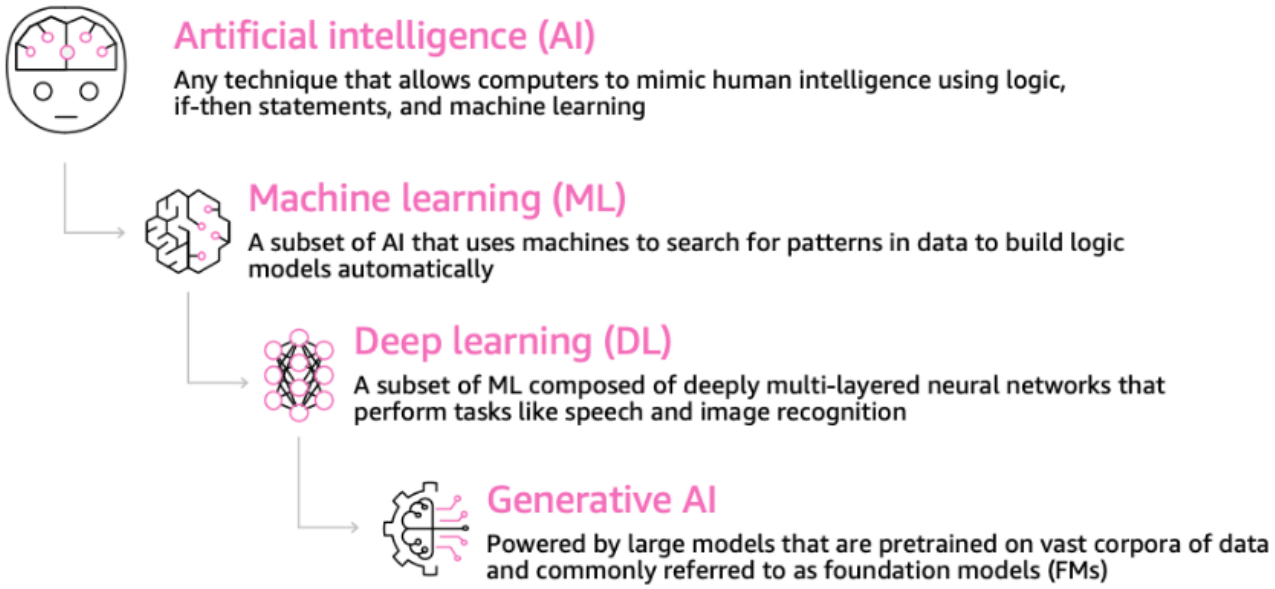


Figure 1: Taxonomy of AI, machine learning, deep learning, and generative AI

# Appendix 3: AWS Responsible AI Policy

The AWS Responsible AI Policy applies to your use of AI and machine learning services, features, and functionality (including third-party models) that we provide (collectively, *AI/ML Services*).

# 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.

## Canada Life transformed its contact center with AI and Amazon Connect


### Overview

Canada Life faced contact center challenges including 5-minute average wait times, complex routing systems, and limited self-service options across their 21 business units.

### Solution

The company implemented Amazon Connect as their cloud contact center solution, integrating it with Salesforce and Calibrio for workforce optimization, while deploying AI capabilities including call summarization, chatbots, and automated authentication across their entire operation in just seven months.

### Impact

Canada Life achieved remarkable results including a 94% reduction in wait times, 92% reduction in average speed to answer, 10% reduction in average handle time, $7.5 million in savings within the first six months of 2025, and improved employee engagement scores that were 5% higher than the national average while maintaining 100% uptime throughout the migration.


*"Using Amazon Bedrock models, we feel comfortable knowing that the data isn't leaving our control. We have a lot of confidence in the security and reliability of the solution."*

***Karthik Kenchaiah*** *Financial Crimes Engineering Lead, Robinhood*

**Overview**

Robinhood needs to keep pace with evolving risks and the rapid growth on its platform, so it's always looking for innovative, scalable tools to streamline operations while maintaining precision. When the company identified an opportunity to augment its financial crimes (FinCrimes) investigations using generative AI, Robinhood chose to work alongside AWS.

**Solution**

Robinhood used Amazon RDS and generative AI on Amazon Bedrock to build FinCrimes Agent—a generative AI solution that analyzes data and produces summaries that surface key insights. The agent helps investigators handle cases quickly and confidently. As a result, Robinhood can create consistent, high-quality investigative evidence and supporting documentation.

**Impact**

The fraud investigation assistant built on Amazon Bedrock improved efficiency of Fraud operations team by 20% and reduced time spent manually collecting and summarizing data, with additional optimizations actively under development.

# Rocket Mortgage transforms client service with AI-powered assistants


*"There really has not been a huge need to build a lot of proprietary stuff here, because Amazon has provided the pieces for us…and showed us how to glue them together securely."*

***Dan Vasquez*** *VP of AI Strategy, Rocket Mortgage*

## Overview

Despite being a household name, Rocket Mortgage captures only 8–9% of the $5 trillion mortgage market. They needed to scale operations efficiently while maintaining personalized service across thousands of team members serving millions of clients, without compromising security or regulatory compliance.

## Solution

Rocket implemented a multi-layered AI system including a client-facing assistant for 24/7 service, their Synopsis solution for call center optimization, and their Navigator platform for internal AI experimentation. Rocket built a secure architecture using services such as Amazon Bedrock and Amazon Transcribe, while maintaining existing security controls.

## Impact

Rocket's AI suite also enabled 15,000 team members to safely experiment with AI applications and saved 40,000 team member hours. Rocket expects a 33% increase in lead conversion through 24/7 AI assistance and enabled 70% of servicing clients to either self-serve or achieve first-call resolution.

# 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 use 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 used code scanning capabilities of Amazon Q 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 used 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 use this generative AI solution in its U.S. equity market surveillance.

# Contributors

Contributors to this document include:

- Krish De, Principal FSI Governance, Risk & Compliance (GRC) specialist
- Brenda Fong, Senior FSI GRC specialist
- Kelvin Leung, Security and Compliance Lead
- Stephen Martin, Head of Financial Services Compliance and Security for EMEA and APAC