Asia Cloud Computing Association

AI Handbook

September 2024

ACKNOWLEDGEMENTS: Research support for this publication was provided by Flint Global. The ACCA is committed to independent, quality research and the conclusions of this report are not determined or influenced by sponsorship.

The ACCA is the apex industry association representing the stakeholders of the cloud computing ecosystem in Asia Pacific. Its mission is to accelerate adoption of cloud computing in Asia by creating a trusted and compelling market environment and a safe and consistent regulatory environment for cloud computing products and services. The association works to ensure that the interests of the cloud computing community are effectively represented in the public policy debate. Drawing on subject-matter expertise from member companies, expert working groups, and special interest groups, it develops best practice recommendations and other thought leadership materials.

To find out more on how to join us, email secretariat@asiacloudcomputing.org, or visit our website at www.asiacloudcomputing.org

Cover image from https://www.istockphoto.com/photo/blue-digital-artificial-intelligence-icon-and-data-matrix-simulation-with-futuristic-gm1495890690-518743456

Table of Contents

Executive Summary

Artificial intelligence (AI) has the potential to significantly spur innovation, increase productivity, and accelerate economic development across Asia-Pacific. McKinsey research has found that generative AI alone could add the equivalent of US$4.4 trillion to global GDP, roughly equivalent to the third-largest economy in the world.[1]

AI is already having a transformative impact across the region. Manufacturers are incorporating AI into their production processes, driving efficiency and economic competitiveness and accelerating the green transition.[2] Farmers are deploying AI to harvest crops more efficiently, improving yields and addressing challenges related to labor shortages.[3] AI solutions are being deployed across healthcare, education and transport in ways that are immeasurably improving ordinary citizens' lives.[4]

While the potential benefits are enormous, AI also presents risks if it is not developed and deployed safely, securely, and responsibly. Members of the Asia Cloud Computing Association (ACCA) are committed to working with governments and other relevant stakeholders to develop and promote robust governance frameworks that enable the safe, secure, responsible and bold adoption of AI across Asia-Pacific.

Effective AI governance must begin with a clear understanding of how AI works and how different actors contribute to that process. Only with this understanding is it possible to distribute responsibilities among those actors best positioned to identify and mitigate the potential harms.

This handbook aims to support policymakers' efforts by:

- Mapping out the **AI lifecycle**, the different stages, and the key actors involved.
- Outlining a **shared responsibility framework** for AI that draws on our experience as cloud service providers of the shared responsibility model for cloud.

As part of this, the handbook sets out a fundamental distinction between AI **developers** and **deployers**:

- **Developers** design, code, or produce AI models.
- **Deployers** implement AI models into their operations or into user-facing applications.

This means that **developers** are better placed to document information on the intended uses, performance expectations, and technical limitations of an AI model. **Deployers** are better placed to

[1] https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#key-insights

[2] https://media.toyota.co.uk/toyota-research-institute-develops-new-ai-technique-with-potential-to-help-speed-up-vehicle-design/

[3] https://enlighten.griffith.edu.au/a-i-and-the-future-of-farming/

[4] https://vinbigdata.com/en/medical-imaging/vindr-ai-to-diagnose-serious-illness-at-early-stage.html; https://www.iasgyan.in/daily-current-affairs/diksha-e-education-platform; https://sites.research.google/greenlight/

undertake risk assessments for a specific use case because only they know in depth the context of that use.

This distinction should form the basis of a shared responsibility framework that provides clarity to different organizations across the AI lifecycle on their responsibilities. A clear framework will ensure that businesses put in place appropriate safeguards to **effectively manage AI risks**, while also providing them with the **clarity and confidence to invest in AI**. This is essential for allowing Asia-Pacific markets to **maintain global competitiveness** in their deployment of AI, while also **providing strong protections** for their citizens and societies against the potential harms.

Asia-Pacific markets have the potential to lead the world in the responsible development and adoption of AI. A shared responsibility approach will enable them to fulfill that potential.

Part 1: The AI lifecycle

Part 1 outlines key aspects relating to how AI is developed and deployed, drawing on ACCA members' experience in providing important functions across the supply chain. Part 1 is divided into:

- Recommending policymakers to use an **internationally aligned and appropriately specific definition of an AI system**, specifically the OECD's globally recognized definition.
- Outlining the key actors at each level of the **AI supply chain** and the function they play.
- Explaining each stage in the development and deployment of AI across the **"AI lifecycle"**.

Definition of AI

Any discussion of AI governance must first begin with establishing a clear and agreed definition of AI. This ensures that policymakers, businesses, and citizens all start from the same place in developing a framework for AI safety.

When defining AI, it is important that:

1) **The definition is internationally aligned**. The development and deployment of AI systems is often a process that takes place across many markets. Aligning the definition provides businesses with greater certainty to develop and deploy AI systems across borders, enabling more markets to share in the benefits of AI.
2) **The definition is appropriately specific**. It should avoid capturing a much wider set of software applications. A targeted definition ensures that policymakers and businesses can focus their governance efforts on those aspects that are unique to AI.

ACCA members recommend aligning with the OECD's globally recognized definition of an AI system. This will allow businesses across Asia-Pacific to collaborate with each other and with other businesses around the world in their development and deployment of AI. When this report discusses AI, it is based on the OECD's current definition of an AI system:

An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment.[5]

The AI supply chain

In addition to establishing an agreed definition of AI, it is also necessary to fully understand the AI supply chain and the different actors involved. This is a vital precondition to effectively allocating responsibility for safety among the different actors.

[5] https://oecd.ai/en/wonk/ai-system-definition-update

At a high level, there is an important distinction between:

1) **Developers**, who design, code, or produce an AI model.
2) **Deployers**, who implement AI models into their operations or into user-facing applications.

For example, when a bank uses AI to help decide whether to approve a customer's loan application, the AI model may have been trained by another company, the AI **developer**. The developer would have trained the model to have the capability to be deployed in many different sectors and use-cases. The bank, as the **deployer**, procures the model from the AI developer and then determines how to deploy it for the specific use-case of customers' loan applications, first assessing the appropriateness of the model for that use-case. This distinction is widely recognized around the world, including within the Asia-Pacific region.[6]

How an AI model is deployed will often depend on the use-case. For example, the bank may decide to have a "human in the loop" reviewing the AI system's recommended decisions on customers' loan applications, given the implications of that decision to the customer. In other contexts, for instance the use of an AI system to recommend products to customers on an e-commerce site, the deployer may decide it is not necessary to have a "human in the loop".

To expand further beyond this fundamental distinction between developers and deployers, the AI supply chain consists of four main layers:

1) **Infrastructure providers**: they provide the integrated hardware and software environment necessary to train and run AI models. This includes, for example, providers of specialist AI chips that are highly efficient at training AI models, such as Nvidia.[7] It also includes cloud service providers, including Google Cloud, AWS, Microsoft, Salesforce, and Oracle, which provide the computing resources needed to train and deploy AI models at scale, and which now often provide their customers with access to different pre-trained models.[8]
2) **Model developers**: they design, code, or produce an AI model. Google DeepMind, for instance, has trained the Gemini set of models for businesses to integrate into their applications as well as for individuals to use.[9] Amazon's Titan family of models provide customers with a breadth of high-performing image, multimodal, and text model choices, via a fully managed API.[10] Anthropic has developed its Claude 3 family of AI models, making these available for individual and commercial use.[11] Model developers may utilize infrastructure providers' infrastructure to develop models. Anthropic is using AWS's cloud computing infrastructure for mission critical workloads and its Trainium and Inferentia chips to build and train models. Anthropic's models

[6] https://www.adalovelaceinstitute.org/resource/ai-supply-chains/; https://aiverifyfoundation.sg/wp-content/uploads/2024/05/Model-AI-Governance-Framework-for-Generative-AI-May-2024-1-1.pdf, p7; https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20240419_9.pdf, pp 4 - 5

[7] https://www.nvidia.com/en-us/ai-data-science/

[8] Amazon SageMaker JumpStart allows customers to evaluate, compare, and select foundation models from model developers such as Meta, Mistral, and Stability AI. Google Cloud's Vertex AI Model Garden provides customers with the ability to select and customize over 150 high-performing foundation models. https://aws.amazon.com/sagemaker/jumpstart/; https://cloud.google.com/model-garden?hl=en

[9] https://deepmind.google/technologies/gemini/

[10] https://aws.amazon.com/bedrock/titan/.

[11] https://www.anthropic.com/claude

are also available on Google Cloud's Vertex AI Model Garden - Google's curated set of leading AI models.[12]

3) **Model deployers**: they take the AI model produced by the model developer and embed it into a wider system, for example in improving the efficiency of their back-office operations or in serving their own users. HSBC, for example, uses the AI Markets solution, an AI chatbot for HSBC's investor clients, which improves price discovery, client service and distribution using natural language processing.[13] This may also involve finetuning the model with domain-specific proprietary data held by HSBC into a customized AI chatbot which is fluent in HSBC's operations and standards.
4) **End users**: the service-user who ultimately engages and interacts with the AI system. This could be, for example, HSBC's investor client using the AI Markets service to access bespoke financial markets analytics.

A single organization may be involved at multiple levels of the supply chain. Several ACCA members provide cloud computing and chip infrastructure to enable customers to train their own foundation models. AWS's Tranium machine learning chip is purpose built for deep learning training of AI models.[14] ACCA members have also developed their own foundation models for business customers to use, while also providing their customers with access to a range of foundation models developed by third-party developers to choose from. Google Cloud's Vertex AI Model Garden allows customers to choose both from Google's own Gemini models and from third-party models such as Meta's Llama 3 and Mistral's Mistral-7B model.[15] Amazon's Bedrock service allows businesses to access multiple foundation models, including those developed by Amazon, through a unified API.[16]

Organizations and individuals at different levels of the AI supply chain each have their own role to play in the safe development and deployment of AI. The next section will outline the AI lifecycle and the processes involved in more detail.

[12] https://www.aboutamazon.com/news/company-news/amazon-anthropic-ai-investment; https://cloud.google.com/blog/products/ai-machine-learning/announcing-anthropics-claude-3-models-in-google-cloud-vertex-ai

[13] https://www.gbm.hsbc.com/en-gb/products/hsbc-ai-markets

[14] https://aws.amazon.com/machine-learning/trainium/

[15] https://cloud.google.com/vertex-ai/generative-ai/docs/model-garden/explore-models

[16] https://docs.aws.amazon.com/bedrock/latest/userguide/what-is-bedrock.html

Figure 1: The AI supply chain

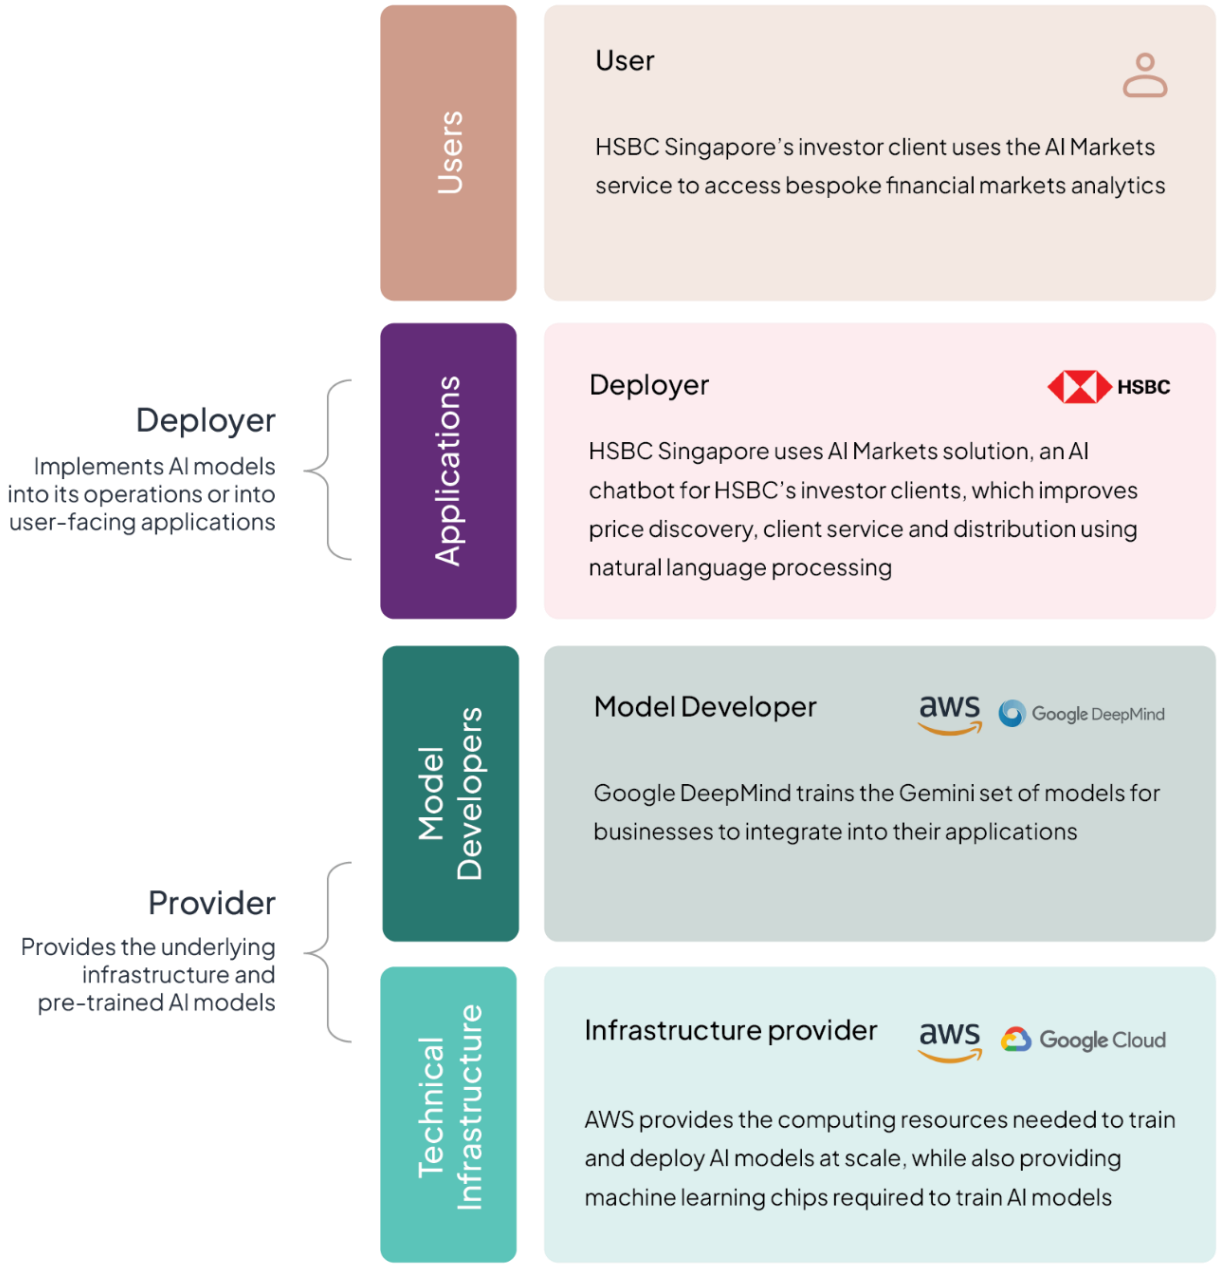

The AI lifecycle

As the previous section made clear, the AI supply chain is complex and involves many different types of organizations. This section will explore the AI lifecycle in more detail. The aim is to provide clarity on all the steps that take place to deploy an AI model in a real-world use-case.

There are five main phases in the AI lifecycle:

Phase 1 - Establishing AI infrastructure

This phase includes establishing the underlying hardware, software, networking, and system processes required to train AI models and develop AI applications. The development and deployment of AI requires consideration of data storage and management, compute resources, robust networking and high bandwidth to support the transfer and processing of large datasets. This infrastructure must meet the vast computational and data processing demands of AI workloads.[17]

Cloud computing enables organizations to access computing resources on-demand and pay only for what they use, rather than having to invest in and maintain their own IT infrastructure. Cloud computing provides the computing resources and infrastructure needed to train and deploy AI models at scale. Cloud service providers - including ACCA members - now also offer a wide range of tools aimed at helping their business customers to develop and deploy AI safely. Amazon Bedrock, for example, allows customers to experiment with and evaluate foundation models from leading AI companies through a single API.[18] Google Cloud's Vertex AI provides customers with APIs for leading foundation models, allowing them to experiment and test different models and to select the model most appropriate to their use case.[19]

Phase 2 - Model development

This phase includes building models, which are designed to optimize for generality and versatility of output. AI models are a set of instructions or rules that enable machines to learn, analyze data and make decisions based on that knowledge.[20]

There are different types of machine learning models. Experts increasingly discuss a distinct category of "foundation models". These may be defined as: *"A machine learning model that is trained on broad data at scale, is designed for generality of output, and can be adapted to a wide range of downstream distinctive tasks or applications, including simple task completion, natural language understanding, translation, or content generation."* These encompass, for example, OpenAI's ChatGPT, Meta's Llama models, Google DeepMind's Gemini models, and Amazon's Titan models.

[17] https://www.computerweekly.com/feature/Top-AI-infrastructure-considerations
[18] https://aws.amazon.com/bedrock/
[19] https://cloud.google.com/generative-ai-studio?hl=en
[20] https://www.techtarget.com/searchenterpriseai/tip/Types-of-AI-algorithms-and-how-they-work

As a subset of foundation models, large language models (LLM), are models that can help computers analyze, understand and respond to human inputs using speech and written text.[21] A regression model estimates the relationship between different variables, for example using different variables to predict whether a company's stock price will increase in the future.[22]

Phase 3 - Model pre-training

At this stage, foundation model developers pre-train their models using a large and diverse dataset with the aim of enhancing their adaptability and applicability to a wide range of real-world scenarios. This allows the AI model to perform effectively and safely when, say, a car manufacturer uses it to accelerate the design of electric vehicles, or when a farmer uses it to automate the harvesting of crops.

Pre-training can take place through **supervised** or **unsupervised** learning:

- **Supervised** learning uses labeled training data. With supervised learning, an algorithm uses a sample dataset to train itself to make predictions, iteratively adjusting itself to minimize error. These datasets are labeled for context, providing the desired output values to enable a model to give a "correct" answer. Supervised learning models therefore have a baseline understanding of what the correct output values should be. Supervised models are often more focused on learning the relationships between input and output data, for example in predicting flight times based on parameters such as weather conditions and airport traffic.

- **Unsupervised** learning does not use labeled training data. Unsupervised learning algorithms work independently to learn the data's inherent structure without any specific guidance or instruction. The developer simply provides unlabeled input data and lets the algorithm identify any naturally occurring patterns in the dataset. This is more helpful for discovering new patterns in raw, unlabeled data, for example in identifying buyer groups that purchase related products together.[23]

Phase 4 - Model fine-tuning

This involves customizing the parameters of a pre-trained model to suit the characteristics of a new dataset tailored to a specific domain or task. This is aimed at improving a model's effectiveness in preparation for deploying it for a specific use-case.

During fine-tuning, the model starts with pre-trained parameters from the initial training. This allows the organization that is fine-tuning the model to save the time and resources needed to train a model from scratch. Fine-tuning involves using domain-specific data to adapt the model. For instance, a bank wanting to improve the effectiveness of a foundation model for detecting fraud may wish to use its own proprietary data of financial transactions to enhance the model.

[21] https://www.elastic.co/what-is/large-language-models

[22] https://h2o.ai/wiki/regression/

[23] https://cloud.google.com/discover/supervised-vs-unsupervised-learning

Several cloud service providers provide tools to help their customers fine-tune AI models. Google Cloud provides customers with supervised tuning for Google DeepMind's Gemini models, helping them to improve the performance of the models for specific tasks.[24] Amazon SageMaker provides business customers with tools for fine-tuning pre-trained foundation models, including by using their own data.[25]

Phase 5 - Model deployment

At this stage, the model is deployed in a real-world application. For example, in the Philippines, SM Supermalls deployed an AI-powered portal to automate the process of business tenants submitting tax forms to the authorities. This significantly accelerated the process of submitting the forms, enabling the businesses to spend more time on other, more business-critical tasks.[26] At this stage, the deployer must decide, firstly, whether the model is suitable for its intended purpose, secondly, what risks there may be for deploying the model for that purpose, and thirdly, what safeguards it should implement to mitigate that risk. SM Supermalls, for example, used Google Cloud's Secret Manager to secure storage of sensitive data such as API keys, passwords and certificates to protect against intrusion.

The deployment of the model is not always linear or a one-off event, but a continuous process of monitoring deployment and how the model is used and adjusting the model accordingly. For example, an insurance company deploying an AI chatbot that provides inappropriate or inaccurate responses to consumers may respond by implementing additional safeguards such as content filters.

Varying approaches to model deployment

The roles that developers and deployers play in each phase can vary. Deployers have the option of procuring an AI model "out of the box", deploying an off-the-shelf AI model such as OpenAI's GPT, Anthropic's Claude, Amazon's Titan, or Google DeepMind's Gemini for a specific application without further fine-tuning. This is often a preferred option for organizations that lack the resources required to fine-tune the model. However, some deployers choose to fine-tune the model for their specific use-case, as off-the-shelf models may need further tailoring before being ready for domain-specific deployment. As Part 2 will show, the distinction between different types of deployment is important for allocating responsibility around the safe development and deployment of AI.

[24] https://cloud.google.com/vertex-ai/generative-ai/docs/models/tune-models#gemini

[25] https://docs.aws.amazon.com/sagemaker/latest/dg/jumpstart-fine-tune.html

[26] https://cloud.google.com/customers/sm-supermalls

Figure 2: The AI lifecycle

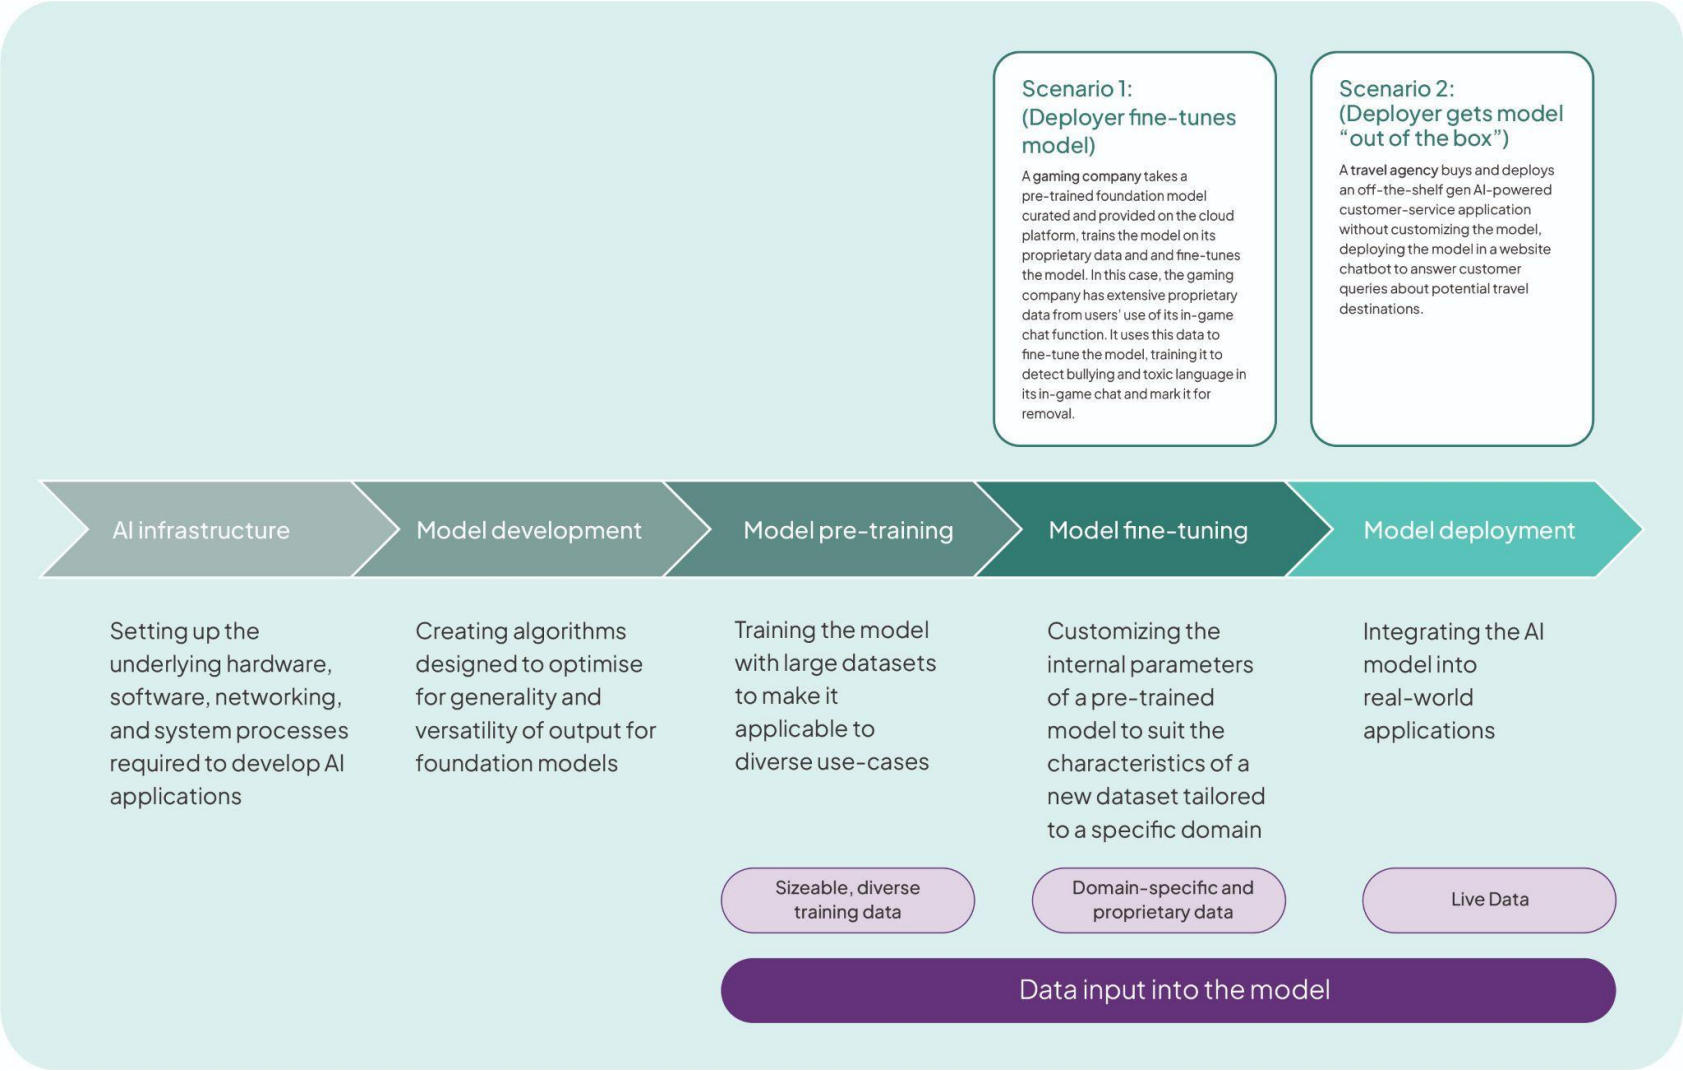

Closing Part 1

Part 1 of this handbook has proposed a globally recognized definition of AI, mapped the different layers of the AI supply chain, and summarized the five stages of the AI lifecycle.

Part 2 of the handbook will draw on this understanding of the AI lifecycle to explore how responsibility for AI safety can be effectively distributed among different organizations involved in the development and deployment of AI.

Part 2: A shared responsibility framework for AI safety

Part 2 of this report provides a framework for how policymakers and industry can work together to effectively allocate responsibility for AI safety. It is divided into:

- Highlighting the importance of clearly **differentiating the responsibilities of AI developers, deployers, and users**.
- Outlining a **comprehensive shared-responsibility model** for AI safety, and providing case studies of how this model works in practice.

To maximize the benefits of AI while effectively managing the risks, it is vital to develop policy approaches that provide **certainty** and **clarity** to companies across the AI lifecycle on their responsibilities. A lack of certainty and clarity creates barriers to investment and innovation, as well as undermining public confidence in the use of AI. Insufficient clarity may also lead to confusion among companies as to how they can meet their responsibilities, potentially leading to gaps in AI safety measures and exposing citizens to greater risk. For example, since developers hold key information on the models themselves, they are best-placed to document information on the intended uses, performance expectations, and technical limitations of the model. By contrast, imposing these responsibilities solely on deployers may lead to unintended consequences. In particular, deployers may have to compel developers to reveal their intellectual property so that the deployers could themselves comply.

This shows how, to be effective, *AI policy approaches must carefully allocate responsibilities on the organization that is best positioned to identify and mitigate the potential harms that could arise from the use of the model*. It is therefore critical to start with an in-depth understanding of the AI lifecycle and the different actors involved, as outlined in Part 1 of this report.

Responsibilities of developers, deployers, and users

Part 1 of this Handbook mapped out the distinct stages of the AI lifecycle. This included the different roles of **developers** (who build the AI technology), **deployers** (who implement AI in a user-facing product or model), and **users** (who engage and interact with the AI model).

The different roles played by developers, deployers, and users is an important starting point for allocating responsibility. Effective risk management is integral to ensuring there are robust safeguards in place against AI-related harm, but developers, deployers, and users are not in a position to carry out the same type of risk management as each other. The OECD has recognized that risk management responsibilities should be different for different types of stakeholders:

> *"AI actors, should, based on their roles, the context, and their ability to act, apply a systematic risk management approach to each phase of the AI system lifecycle on an ongoing basis and adopt responsible business conduct to address risks related to AI systems".*[27]

[27] https://oecd.ai/en/dashboards/ai-principles/P9

Firstly, **developers** build algorithms and pre-train their models with the aim of ensuring they can be applied to a wide range of real-world scenarios. They therefore hold important responsibilities in relation to making information on their models available for deployers. This should include guidance on the model's intended purpose, operating boundaries, and known, likely and specific high risks. However, developers are not best-placed to conduct risk assessments, given the enormous diversity of potential use cases (see next point). Requiring developers to comprehensively identify and mitigate risks may lead to a situation where AI models are primarily developed by a small group of large AI model developers. This would reduce the pool of developers who are willing to innovate and deprive economies and societies of the benefits of AI innovation.

Secondly, **deployers** integrate the AI model into real-world applications, in some cases first fine-tuning the model with their own data. They have direct oversight over how they are integrating the AI model into their operations. Since this will vary by sector and even by individual company, deployers should be responsible for conducting a context-specific and company-level risk assessment on their intended use of AI. Deployers even in the same industry do not take an identical approach to deploying AI. One bank may deploy an AI chatbot to answer customers' queries about the general availability of loan products, whereas another bank may choose to task the AI chatbot with using the customer's data to provide personalized quotes. The differing nature of deployment leads to differences in the risk profile, requiring separate risk assessments and mitigation measures.

Thirdly, **users** interact with the AI system, for example as customers of a bank or as citizens accessing information about government services online. While developers and deployers must shoulder their own responsibilities for safe AI, effective risk management also requires users to take responsibility. Developers and deployers put in place terms of service with the end user, based on the application's intended use. The end user is obliged to abide by these terms and deployers and developers cannot be held responsible where this does not happen. Where users abuse AI models in order to perpetuate scams, for example, those users should be held to account. Providing users with information on managing AI-related risks and holding users accountable for deliberate misuse of AI models is an integral part of ensuring AI safety.

It should be noted that these roles will not always be fixed and are context-dependent. For instance, organizations often choose to build and deploy their own models. In such instances, the organization in question would play the roles of both the developer and deployer, and would accordingly become responsible for both developing and deploying AI safely. This handbook will explain this further in subsequent sections.

Figure 3: Responsibilities between AI developers and deployers

A shared responsibility framework for AI safety

This section will outline in further detail the actions that developers and deployers should each take. So that each organization understands its responsibilities, there needs to be a clear framework for how organizations at the different stages cooperate and share responsibility, based on their ability to identify and mitigate potential harm.

ACCA members have extensive experience of how such a shared responsibility framework can function from our experience as cloud service providers. The cloud security shared responsibility model is a long-established framework for sharing responsibility for the security and availability of data and workloads in a cloud service.[28] This framework has provided robust protections for privacy and security, underpinning the role of cloud computing in accelerating economic growth across the Asia-Pacific region.[29]

Similarly, as Part 1 of this handbook outlined, there are variations in the roles played by AI developers and deployers, and in the level of control that deployers choose to have over the process. A shared responsibility framework for AI safety can account for these variations. *Figure 4* maps out an appropriate allocation of responsibility across different stages of the AI lifecycle, according to the nature of deployment. This includes considering responsibilities at each stage:

- **AI infrastructure** - It is necessary to build and safeguard the infrastructure required to train AI models. For example, this includes maintaining the availability and physical protections of data centers. In most cases, the **infrastructure providers** themselves will be responsible for the security of the underlying infrastructure. However, if an AI **developer** decided to build services in its own data centers ('on-premises'), it would be responsible for the security of that infrastructure.

- **Model development** - The organization developing the model should publish key information on capabilities and limitations of the models. This should include general information on the intended uses, performance expectations, and technical limitations of the AI model, as well as on known risks that could occur, and steps taken to mitigate those risks. In many cases, this responsibility will fall to an AI **developer** that develops models for a wide variety of use-cases and makes them available to third parties. However, in some cases an organization may choose to build, develop, and deploy its own model from scratch, or substantially modify an AI model developed by a third-party developer. In this case, it would be considered as playing the role of both a **developer** and a **deployer** and should be responsible for the applicable responsibilities for both developers and deployers. A further consideration is when deployers use open-source AI models. Open-source AI is where the source code, algorithms, and/or training data are made publicly available to third parties. If the deployer takes an open-source AI model and substantially modifies core elements of the model, it would also be considered as playing the

[28] https://www.ncsc.gov.uk/collection/cloud/understanding-cloud-services/cloud-security-shared-responsibility-model

[29] https://www.adb.org/publications/cloud-computing-policies-and-their-economic-impacts-in-asia-and-the-pacific

role of both a **developer** and **deployer** and should be responsible for the applicable responsibilities for both.

- **Model pre-training** - The pre-training of models to ensure their applicability to a wide range of real-world scenarios comes with an important set of responsibilities. This includes responsibility for ensuring that this pre-training complies with relevant obligations relating to data privacy and intellectual property protection. It also includes responsibility for testing the models with the aim of surfacing risks such as the potential for bias or leakages of sensitive or personal data. As above, in most cases these responsibilities will fall to the AI **developer**, but where the **deployer** itself chooses to develop and pre-train the model, it would be considered as playing the role of both a **developer** and **deployer** and should assume the applicable responsibilities for both roles accordingly.

- **Model fine-tuning** - Fine-tuning the model - training the model on a smaller, task-specific dataset - involves further responsibilities in relation to data governance and risk management. Again, these include ensuring that they fine-tune data in line with data privacy and intellectual property protection obligations, as well as the management of risks such as data bias. Deployers performing fine tuning prior to deployment should test before deployment, as they are modifying the model for their own specific uses.

- **Model deployment** - The deployment of the model in a real-world use case requires a comprehensive approach to assessing and mitigating risk, as well as monitoring the AI model's use for any harmful uses or impacts after deployment. The **deployer** has direct visibility over exactly how it plans to integrate AI into its operations and is therefore best-placed to conduct this risk assessment and to test the model before deployment. This often requires deep-domain specific insights. As one example, a medtech company deploying AI in a diagnostic device will have deep expertise on the diseases being tested for and how these affect different parts of the population, as well as on how the device is likely to be used in a medical setting. This company will therefore be best-placed to assume responsibility for managing the risks around deployment, such as the risk of biased outcomes and higher rates of false negatives for certain groups of patients, and for testing the model prior to deployment.

Figure 4: A shared responsibility framework for AI

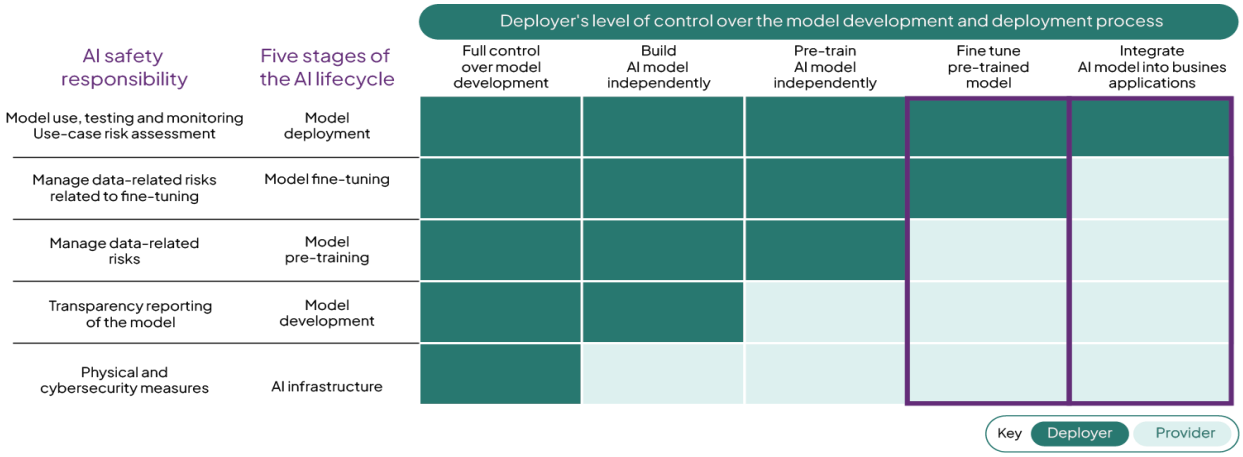

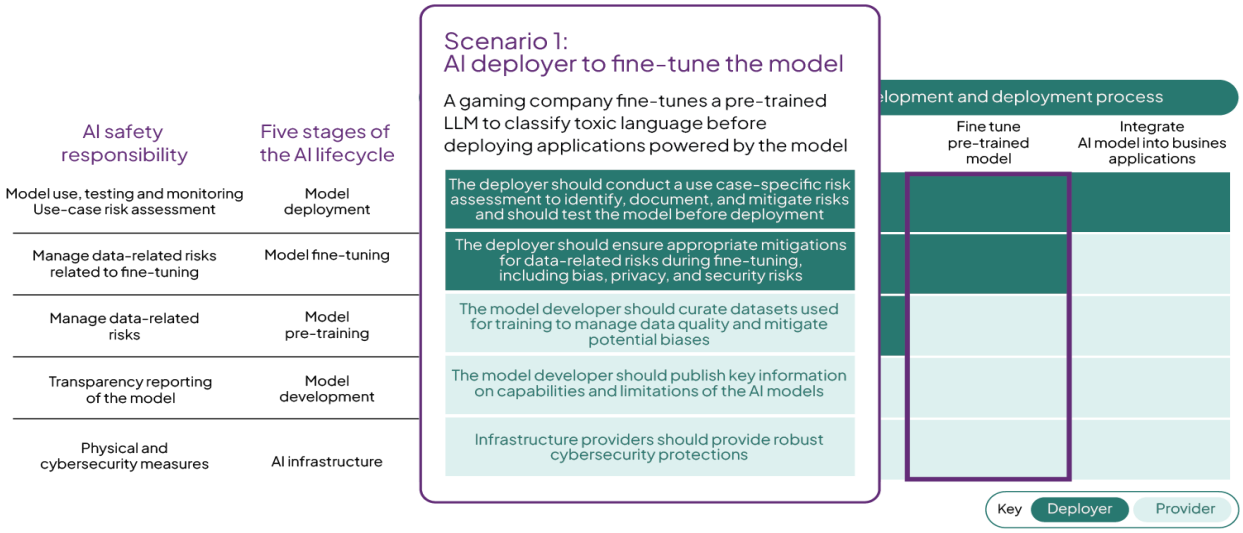

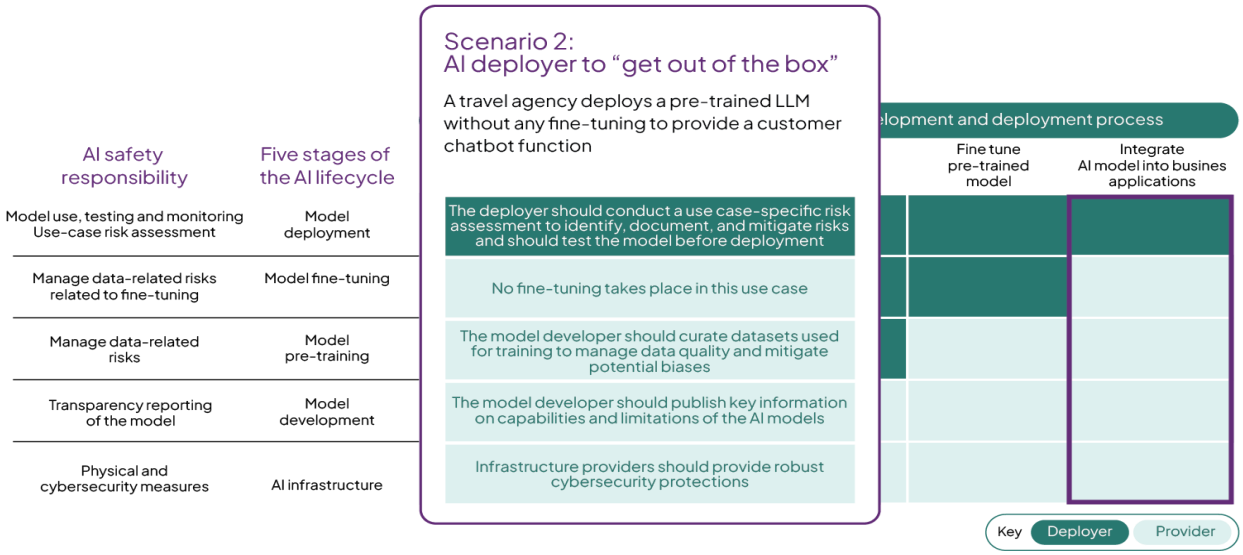

The shared responsibility framework in practice

As policymakers consider the allocation of responsibility across the AI lifecycle, they would be well advised to build on existing systems of cooperation between AI developers and deployers. These systems are already proving to be effective in leading developers and deployers to work together to manage AI risks.

ACCA members recognize the importance that, when we develop AI models, we work cooperatively with deployers to provide them with the tools and information required to understand the models developed and potential risks or unintended consequences that could arise from their use. In working with developers, ACCA members are putting the shared responsibility approach into practice.

ACCA members support deployers to fulfill their responsibilities through some of the following approaches to sharing information:

- **Model cards** - short documents accompanying trained machine learning models that typically include information such as the model's intended use case, the model's performance on different metrics, any known biases or limitations of the model, and any potential risks or unintended consequences that could arise from its use. AWS's AI Service Cards, for example, provide customers with a single place to find information on the intended use cases and limitations, responsible AI design choices, and deployment and performance optimization best practices for our AI services.[30]

- **Data cards** - dataset documentations framework aimed at increasing transparency across dataset lifecycles internally within an organization. They provide structured summaries of ML datasets with explanations of processes and rationale that shape the data and describe how the data may be used to train or evaluate models. Google's Data Cards Playbook is a self-guided toolkit that both AI developers and deployers can use to guide their efforts to manage internal data governance challenges such as the use of sensitive data.[31]

- **Technical reports** - technical documents, white papers or guidance that provide high-level information on model architecture, training infrastructure, and pre-training dataset, sometimes alongside details of model evaluations and benchmarking of model performance on key capabilities. These documents can also include discussion of the broader implications of AI models, including their limitations and potential applications.[32]

- **Tools to support risk assessment and mitigation** - ACCA members provide their customers with a wide range of tools aimed at simplifying the process of assessing and mitigating risks relating to their deployment of AI. On risk mitigation, AWS's Guardrails for Amazon Bedrock

[30] https://aws.amazon.com/blogs/machine-learning/introducing-aws-ai-service-cards-a-new-resource-to-enhance-transparency-and-advance-responsible-ai/

[31] https://sites.research.google/datacardsplaybook/

[32] https://storage.googleapis.com/deepmind-media/gemini/gemini_1_report.pdf; https://cdn.openai.com/papers/gpt-4.pdf; https://aws.amazon.com/blogs/machine-learning/use-amazon-titan-models-for-image-generation-editing-and-searching/

supports customers to implement safeguards such as content filters customized to their unique requirements.[33] Amazon's Sagemaker Clarify also detects potential bias from data prep to deployment.[34] Google provides a range of tools and guidance such as a Responsible Generative AI Toolkit to help developers design, build and evaluate AI models responsibly.[35]

By working with our customers in this way, ACCA members are helping to maintain the pace of AI adoption across the Asia-Pacific, while ensuring that AI is rolled out safely. Without this support, some deployers would either **choose not to incorporate AI into their operations** or would alternatively **implement AI but with insufficient safeguards**. *Figure 5* provides three examples of this shared responsibility framework in action.

There are significant risks involved in seeking to replace existing systems of cooperation with an entirely new framework. For one, it risks allocating responsibilities to entities that are not best-placed to fulfill them, meaning safeguards are likely to be less robust. Secondly, an overly rigid and prescriptive policy approach could lead to both developers and deployers moving from their current cooperative mindset to a more legalistic, do-the-bare-minimum compliance-focused approach. In short, this could mean both that **companies are more risk-averse about deploying AI**, depriving Asia-Pacific economies of the potential benefits of AI applications, and that where AI is deployed, **AI governance is less effective**, exposing citizens and societies to greater risk.

[33] https://aws.amazon.com/bedrock/guardrails/

[34] https://aws.amazon.com/sagemaker/clarify/

[35] https://ai.google.dev/responsible

Figure 5: The shared responsibility framework in action

Final thoughts

Developing and deploying AI safely is a whole-of-society effort. It requires government and businesses across the economy to work together in the best interests of society as a whole.

If we can get this right, economies across Asia-Pacific stand to reap the rewards of fast and safe AI adoption. AI has the potential to accelerate economic development, boost competitiveness and productivity, and address some of the region's most pressing societal and environmental challenges.

But it will not be easy. Only through a clear framework for cooperating and sharing responsibility across the AI lifecycle can governments and industry ensure AI safety, security, and innovation go hand-in-hand. ACCA members look forward to working with governments and partners on implementing a shared responsibility framework for the safe implementation of AI.

The ACCA is the apex industry association representing the stakeholders of the cloud computing ecosystem in Asia Pacific. Its mission is to accelerate adoption of cloud computing in Asia by creating a trusted and compelling market environment and a safe and consistent regulatory environment for cloud computing products and services.

The association works to ensure that the interests of the cloud computing community are effectively represented in the public policy debate. Drawing on subject -matter expertise from member companies, expert working groups, and special interest groups, it develops best practice recommendations and other thought leadership materials.

To find out more on how to join us, email secretariat@asiacloudcomputing.org, or visit our website at www.asiacloudcomputing.org

ACCA Member Companies

aws
EQUINIX
Google Cloud
HSBC
Microsoft
ORACLE
PDG
salesforce