# Our Approach to Adopting AI

Commonwealth Bank

# Contents



**About this report**

This report outlines how Commonwealth Bank of Australia (CBA) is ideating, developing, deploying and managing Artificial Intelligence (AI) at an organisational level with a focus on responsible practices. It highlights our current approach, supported by real-world examples.[1,2]

This report is accurate as at 31 December 2025 unless otherwise stated.

**Acknowledgment of Country**

Commonwealth Bank of Australia respectfully acknowledges the Traditional Owners of the Lands across Australia as the continuing custodians of Country and Culture. We pay our respects to First Nations peoples and their Elders, past and present.

Our registered office is located on the Lands of the Gadigal Peoples.

**Accessibility of this document**

This is an interactive PDF and is best viewed with Adobe Reader. Click on the content tabs on the pages or use the footer to navigate through the report.


Commonwealth Bank of Australia ABN 48 123 123 124
AFSL and Australian credit licence 234945

# CEO message

## I am pleased to present the *Our Approach to Adopting AI* report.

Through conversations with our people, customers, and the wider community, we've heard that stakeholders want to understand how Artificial Intelligence (AI) is being used across the bank and how we're managing the risks associated with its adoption. In response, through this report we are sharing information regarding our AI adoption progress and the actions we're taking to manage potential risks.

Technology is a critical enabler of our strategic priorities, and we are consistently investing in leading technology and data capabilities. Among our technology priorities, investing in AI has been a strategic focus for a number of years. AI offers new ways to enhance customer experiences, strengthen risk management, and unlock business value. It helps us deliver tailored experiences for our customers and makes it easier for our people to get things done. In recent years, we've used AI to help us detect fraud, strengthen cyber security and help identify signs of abuse in transaction descriptions.

We recognise that as AI reshapes industries, it presents both opportunities and risks. We operate in a highly regulated industry where maintaining stakeholder trust is critical.

Our existing risk management frameworks, guided by our AI principles, are central to how we manage AI risk, and we review these frameworks and policies in response to evolving technologies and guidance from regulators and industry bodies.

As AI becomes more accessible, customers may choose to interact with our services through external agents or systems outside our control. While we may not be responsible for these external agents or systems, where we identify material risk, we may seek to implement limitations or controls to help protect our customers, channels and platforms.

Alongside these risk management foundations, we provide guidance and tools to support our teams to implement AI responsibly throughout the AI lifecycle. As we navigate this rapidly changing landscape, we recognise that the adoption of AI is an ongoing journey of learning and improvement. We are committed to implementing practices that help us to adopt the technology responsibly and we will continue to evolve our approach over time.

Our people are central to how we deliver for our customers, so we're expanding skilling programs and helping them to build confidence to thrive in a changing workplace. We are also helping Australia harness AI-related opportunities, by working with global organisations to develop leading capabilities in Australia and create innovation opportunities for our teams.

In 2025, we launched a Tech Hub in Seattle, Washington to collaborate directly with our technology partners, accelerate learning and bring advanced skills back to Australia in CBA roles. We're also seeking to share insights with government, regulators and industry bodies to support the responsible adoption of AI in the wider community.

We look forward to receiving feedback, and further discussion on our approach to ideating, developing, deploying, and managing AI responsibly as we strive to deliver better experiences for our customers and contribute to a brighter future for all.

**Matt Comyn**
Chief Executive Officer
Commonwealth Bank

## About Commonwealth Bank

At Commonwealth Bank of Australia (CBA), we are guided by our purpose – building a brighter future for all.

Our strategy, to build tomorrow's bank today for our customers, reflects our focus on helping build Australia's future economy, reimagining customer experiences, leading in technology and AI and delivering simpler, safer and better banking.

As Australia's leading transaction bank for retail and business customers we serve approximately 16 million customers and employ more than 37,000 full-time equivalent (FTE) employees in Australia as at 30 June 2025.[3,4]

The CBA Group structure includes subsidiaries and branches operating in Australia, New Zealand and other countries. In this report, we are describing CBA's approach, including that of our overseas branches. Where our subsidiaries ideate, develop, deploy and manage AI, they may take different approaches in response to local context and regulation. The content set out in this report applies to CBA only.

# How responsible AI can help us build a brighter future for all

**Banking technology has evolved significantly in the past 30 years, and AI marks another step forward, offering new opportunities for our customers, our people and the broader community. Like any innovation, its impact will depend on how thoughtfully and responsibly it is applied.**

AI has the potential to transform customer experiences and contribute to Australia's long-term prosperity, if its benefits are widely shared. By supporting faster, simpler and more tailored banking, AI helps us better understand and respond to individual customer needs. To help shape the future, we're actively collaborating with experts across Australia and globally to help position the country for success in an AI-driven economy.

As AI has evolved, we have built upon our foundations, including frameworks, policies and approaches to risk management practices, to support our people in the responsible ideation, development, deployment and management of AI Systems.[5] We review these frameworks and policies in light of evolving technologies and guidance from regulators and industry bodies, as well as internal learning.

Our Group AI Policy outlines the principles that help guide how we apply our risk management frameworks. Our AI principles also help inform the tools and resources we provide to our developers to assist them in applying requirements throughout the AI lifecycle.

We recognise that as we continue to adopt AI, it has the potential to transform how we work over time. To support this transition, we're investing in skilling for our people, including technical skilling in areas such as AI-powered engineering, advanced analytics and prompt engineering, to help them develop future ready skills and enable career progression.

**Purpose:** Building a brighter future for all

**Strategy:** Building tomorrow's bank today for our customers

## The AI opportunity

### Enhancing customer experiences

We are using AI to help reduce scams and fraud, protect against phishing, and deliver experiences that feel more tailored and relevant for our customers.

**For more information see pages 7-8.**

### Building Australia's future economy

We are working across the business sector, and with government, regulators, academic partners and community organisations, to support AI exploration and adoption in Australia.

**For more information see pages 9-10.**

## Our approach to adopting AI

### Leveraging our technology and risk management foundations

We are building on existing risk management frameworks to manage risk in the ideation, development, deployment and management of AI Systems, supported by AI-specific requirements that help our people make responsible decisions.

**For more information see page 12.**

### Maintaining governance and accountability

The Board, Executive Leadership Team and management committees play an important role in governing AI risk. In addition to existing forums, an AI-specific forum has been established to support policy development and the deployment of AI at CBA.

**For more information see page 13.**

### Being guided by our AI principles

Our principles help guide the application and development of existing policies, frameworks and resources to support responsible ideation, development, deployment and management of AI Systems.

**For more information see page 14.**

### Applying responsible considerations throughout the AI lifecycle

We have requirements, processes and guidance to support the responsible ideation, development, deployment and management of AI Systems.

**For more information see pages 15-17.**

### Strengthening our people's skills

Our people are central to how we deliver for our customers. As we continue to use AI to enhance customer experiences and transform how we work over time, we are creating pathways to help develop our people's skills for the future.

**For more information see page 18.**

# The AI opportunity

# At CBA, we've been working in AI for over 10 years

## AI has supported us to continue to improve our customer experience.

### 2015–2020

- Launched the Customer Engagement Engine (CEE), a machine learning system, for better and more timely customer interactions.[6]
- Joined the Federal Government's pilot to support the development of Australia's AI Ethics Principles.

### 2021–2022

- Announced investment and partnership with H2O.ai, a global leader in AI.
- Established senior accountability for AI risk framework.
- Released model risk framework outlining the model risk controls for AI.
- Participated in the National AI Centre's Responsible AI Think Tank to help Australian businesses with education and support to manage AI safely.
- Established a centre of excellence within CBA to support overall AI governance, connection with existing policies and development of supporting frameworks.
- Began engagement with Gradient Institute to provide guidance on the development of our AI practices.

### 2023

- Commenced implementation of Group AI Policy and a set of principles to supplement our existing policies and risk management frameworks.
- Developed a responsible AI toolkit to help our developers when building AI. The toolkit provides guidance, tools, and code notebooks across key areas like data preparation, model explainability and fairness assessment.
- Made our Abuse In Transaction Descriptions AI model available for use by other banks globally.
- Ranked as #6 globally in AI maturity in 2023 Evident AI Index of the banking sector.[8]

### 2024

- Established a dedicated governance forum to oversee the development of AI risk framework.
- Developed a generative AI (Gen AI) responsible AI toolkit to help developers evaluate performance of their Gen AI models in a more streamlined and standardised way, including various aspects related to quality, robustness and fairness.[9]
- Developed a Gen AI playbook to guide teams through best practices for progressing AI Systems from ideation to deployment, drawing on collaborative InnerSource contributions from across CBA.[10]
- Established our Human-Centred AI unit which supports a research-backed approach to AI at CBA.
- Deepened our Gen AI partnership with Microsoft to improve customer experiences and support ongoing cyber security initiatives.
- Activated an AI factory in collaboration with Amazon Web Services (AWS) to accelerate the adoption of Gen AI across the bank.
- Launched a five-year partnership with the University of Adelaide's Australian Institute for Machine Learning (AIML) to boost foundational AI research in Australia.
- Launched an AI learning series focused on the skilling and education of our people.
- Ranked as #5 globally in AI maturity in 2024 Evident AI Index of the banking sector.[11]

### 2025

- Announced strategic partnership with OpenAI to explore advanced Gen AI solutions, including the rollout of ChatGPT Enterprise to our people, as well as launching a national technology skills initiative to help small businesses grow with AI.
- Expanded strategic partnership with Anthropic, leveraging their AI capabilities and expertise in AI safety and research.
- Launched a Tech Hub in Seattle, Washington USA to provide our technologists with collaborative learning opportunities with global technology leaders.
- Collaborated with human-centred AI leaders and researchers to share their insights and learnings with our people – including MIT Sloan and the University of Technology Sydney.
- Became a founding sponsor of the Tech Policy Design Institute, an independent Australian think tank dedicated to shaping technology policy for the benefit of humanity.
- Joined Future Skills Organisation Skills Accelerator-AI, a national initiative designed to rapidly uplift AI capabilities across Australia's vocational education and training sector.
- Launched the next phase of the AI learning series for our people, focused on practical activities.
- Released a video series designed to help Australians understand and build confidence with AI.
- Ranked as #4 globally in AI maturity in 2025 Evident AI Index of the banking sector.[12]

**80 million events**

**per day screened on average using AI, including payments, password resets and address changes, to help detect fraud and scam activities.[7]**

**27,600 people**

**at CBA have engaged with the AI learning series since it launched in 2024.**

# Enhancing customer experiences

**We are using AI to help reduce scams and fraud, protect against phishing, and deliver experiences that feel more tailored and relevant for our customers.**

AI plays a critical role in helping us to further support our customers' account security – for example in near real time detection and prevention of cybercrime. From helping us to identify and resolve pain points to streamlining interactions across channels, AI supports us in creating more intuitive experiences.

## Using AI to help reduce scams and fraud

Scams and fraud are recognised as a key threat to Australian consumers, with scammers becoming increasingly sophisticated in their tactics including through their use of AI. Helping to keep customers' accounts safe and secure remains a priority. We invested more than $900 million in the 2025 financial year to help protect our customers from fraud, scams, cyber threats and financial crime.[13]

We use AI to help enhance scams and fraud detection strategies by quickly identifying unusual events in highly complex patterns of activity. It also strengthens traditionally manual steps in defensive processes to provide more protection for our customers. As part of our 'prevention, detection and response' approach, our detection platforms use a combination of machine learning and AI technologies to help identify unusual behaviour and alert customers to potential scam events.

CBA now applies these methods to process more than 20 million payments a day on average and send 40,355 proactive warning alerts per day on average to customers via the CommBank app. This has played a crucial role in helping to reduce customer fraud losses by over 20% in the first half of 2026 financial year compared to the first half of 2025 financial year.

We've also partnered with Apate.ai, a cyber-intelligence company, to help fight scams using AI-powered bots that engage scammers in real-time conversations. When a scammer calls or texts, sophisticated bots engage them in extended conversations and gather intelligence. These insights are fed into CBA's fraud detection systems and are shared across the broader anti-scam ecosystem—helping to identify and potentially block threats before they reach customers. This initiative reflects our focus on investing in innovative technologies that help protect the Australian community and support national scam prevention efforts.

Another way we're helping to support anti-scam efforts across Australia is by continuing to share our insights and learning. In 2024 we helped launch the Anti-Scam Intelligence Loop and have already shared over 30,000 pieces of intelligence with industry partners.[14]

## Helping combat phishing using AI

Phishing remains one of the most persistent threats when it comes to cybercrime, and one of the hardest to combat. Phishing involves attackers impersonating people and brands to deceive people into revealing sensitive information, such as passwords, credit card numbers, or other confidential information. This is often done through fake emails, websites, or messages that appear to be from legitimate sources.

We are using AI in new ways to seek to combat phishing attacks. Leveraging AI, we can generate possible domain names that could appear to be CBA related. We then verify if those sites are live and if they are, they are investigated and taken down so they can't be used to deceive customers into thinking they are interacting with a genuine CBA website.

Through this early detection approach, we seek to proactively combat phishing efforts, helping to reduce the ability of fraudsters and scammers to collect information from victims. Ultimately we hope this can help to reduce the number of fraud and scam attacks on Australians overall.

## Providing more tailored support

As digital banking continues to evolve, AI is becoming a key enabler in helping to create more intuitive experiences for customers. With over 9.4 million customers using the CommBank app, we're using AI to tailor interactions.

CommBank Yello, our loyalty program, rewards eligible customers with tailored benefits and offers. The program leverages AI to help match customers with the most relevant partner offers.

**$29 million**

**reduction in potential financial losses through Card Not Present Detection project[15]**

## Detecting Card Not Present scams and fraud

The AI-driven Card Not Present (CNP) Detection project was designed to better protect customers from scams and fraud involving online or phone purchases, where the physical card isn't required. These transactions tend to be more susceptible to scams and fraud because the cardholder may not be asked to verify the transaction in real time.

To help combat this, CBA built an AI model using a range of transactional, customer and confirmed fraud indicators. The model selects highly indicative variables and matches these to situational data points to assign a risk score to each CNP transaction. When the risk score is incorporated into our detection platform, alert accuracy has increased, on average, by over 14%, helping us to identify at risk customers sooner and reduce potential financial losses by approximately $29 million.[15]

## Addressing customer pain points

AI has long been the foundation of CommBank app features like Bill Sense, which helps predict upcoming bills so customers can stay ahead.

Released in 2020, Bill Sense uses advanced machine learning techniques to understand the patterns of customers' regular payments, predict when future bills will arrive and how much they will be. While the predictions are informed by an AI algorithm through a feedback loop, throughout the process customers are prompted to take key actions, such as confirmation of a bill and whether to include a specific biller, setting the payment frequencies and adjusting the predicted payment amount.

Our chatbot, Ceba, leverages AI to help customers complete banking tasks themselves. There are currently over 200 available tasks, including activating a card, checking account balances, making payments and accessing cardless cash. Ceba also allows our customers to connect to a customer service representative if they need to. With Ceba housed on our banking systems, it offers a secure channel that is available 24/7.

AI also helps us to solve complex problems for our customers in more efficient ways. AI-driven applications have the capability to quickly analyse large volumes of documents, which would otherwise be a time-intensive task for a human. An example of this is Compass AI, a Gen AI solution used by CBA to support our client-facing people in our business bank. Prior to Compass AI, business bankers spent considerable time searching, navigating and interpreting information to respond to a customer question. Since the roll-out of Compass AI in July 2024, over 500,000 business banker questions have been answered by the Gen AI assistant, delivering inquiries on our knowledge base over three times faster than traditional methods. This has allowed our frontline teams to be able to spend more time focused on the needs of our customers.

## Taking a research-backed approach

In November 2024, we established the Human-Centred AI unit to support an approach to AI that reflects the needs of our people and customers. This unit tests AI experiences, and partners with global research groups to help keep our work grounded in the latest evidence. The interdisciplinary team comprises technologists, psychologists, economists, designers and ethnographers who work with our technology teams to bring a research-backed approach to the ways we conceive, design, and communicate about these technologies.

**Over 3x faster**

**Compass AI has allowed us to deliver inquiries on our business bank knowledge base *over 3x faster* than traditional methods.**

## Utilising CBA's first AI tool

Since 2015, the Customer Engagement Engine (CEE) has helped CBA to better understand its customers and, in turn, help deliver improved and more tailored experiences.

The CEE enabled the development of Benefits finder, a tool available in the CommBank app and NetBank that helps personal and business customers find and apply for benefits, rebates and concessions most relevant to them. It is a free-to-use tool that has supported more than 400 different rebates since the program launched in 2019 – ranging from energy rebates for families, to the wage subsidies rebate for businesses.

Benefits finder was developed as part of the ongoing collaboration between CBA and Harvard University's STAR (Sustainability, Transparency and Accountability Research) Lab. The Harvard STAR Lab comprises a group of scholars at Harvard University who work with organisations to research, develop and test interventions to support better customer, community and societal outcomes.

## Launching an AI factory

CBA has collaborated with AWS to activate an AI factory powered by Amazon EC2 P5 Instances. This initiative provides the high-performance compute capacity needed to support our teams in testing and developing AI solutions, including fine-tuning and training LLMs.

# Building Australia's future economy

**We are working across the business sector, and with government, regulators, academic partners and community organisations, to support AI exploration and adoption in Australia, with the objective that the benefits can be shared widely across our economy.**

AI has the potential to transform the global economy. With fast-paced adoption, estimates indicate AI could add $115 billion annually to the Australian economy by 2030, with the majority (up to $80 billion) coming from productivity improvements.[16]

Capturing this opportunity requires an ongoing and ambitious national approach to AI exploration and adoption, investment in innovation and a considered approach to building Australia's AI capability.

### Positioning Australia to take up the AI opportunity

We anticipate that AI will profoundly shape Australia's economic future. As a national institution, part of our role is to work through this transition with customers, government and regulators, helping to spread the benefits of AI across the economy, support communities and build national prosperity.

We support a national approach that guides Australia's strategic adoption of responsible AI, bringing together all levels of government, the business community, academia, and community organisations.

To share our insights and learn from others, we've participated in key government and industry forums. This included participating in the National AI Centre's (NAIC) Responsible AI Think Tank, which helped Australian businesses strengthen their understanding of responsible and effective AI practices. We also collaborated with the Chair of Australia's AI standards committee, who represents Australia at the International Organization for Standardization (ISO) and helps shape global AI standards.

In 2019, we collaborated with the Australian Government to support the development of Australia's AI Ethics Principles. As part of the Government's pilot, we shared insights on how we applied our risk management frameworks as we sought to develop the CommBank app feature Bill Sense, responsibly. Our insights from developing Bill Sense helped inform key considerations for Australia's AI Ethics Principles. We remain an active participant in government inquiries to inform the safe and responsible use of AI in Australia.

We're taking a collaborative, ecosystem-based approach to addressing the challenges and opportunities of AI security. In May 2025, we hosted the first in a series of cross-industry events, bringing together over 70 peers from 15 organisations to explore key issues at the intersection of AI and cyber security.

In July 2025, we became a founding member of the National Security Tech Alliance (NSTA), a new sub-group of the Tech Council of Australia. NSTA will work to position Australia's tech industry to lead in areas of strategic advantage, including AI, cyber security, and the security of critical infrastructure.

### Driving innovation through AI adoption

Increased research and development investment can help accelerate innovation and help strengthen Australia's global position. This includes supporting start-ups and small businesses to have the tools and skills needed to innovate using AI, and fostering inclusive access and diverse participation from businesses across regions and communities.

In September 2024, CBA and the University of Adelaide's Australian Institute for Machine Learning (AIML) launched a five-year strategic partnership to boost foundational AI research in Australia through the creation of the CommBank Centre for Foundational AI Research.

The Centre aims to keep Australia at the forefront of foundational AI innovation, creating a competitive edge for the nation as developers and adopters of AI.

### Accelerating AI skills for the next generation

The demand for skilled professionals in AI and related fields is growing rapidly. Research published in 2025 indicated global AI hiring had increased by more than 300% over the previous eight years and in 2024 alone, over 1,500 organisations sought workers with AI-related skills in Australia.[17, 18] To remain competitive, Australia must continue investing in AI education, so that the next generation graduates with foundational AI skills.

Research has shown that having and using AI skills can close the gap between low and high-skilled workers.[19] However, without equitable investment, the adoption of AI has the potential to accelerate social disparities. To support this transition, education initiatives need to consider the needs of diverse communities, so all Australians have the opportunity to thrive in an AI-driven economy. CBA supports Day of AI Australia's Classroom Champion program, which works with lower-Index of Community Socio-Educational Advantage (ICSEA) public schools across the country to bring hands-on AI literacy experiences to primary and high school students. Between April and September 2025, together with other Day of AI Australia partners, we helped to support activities with 39 schools, reaching more than 1,820 students.

As more businesses adopt AI, it has the potential to transform how people work. Harnessing the full opportunity that AI will bring to Australia requires fostering partnerships between educational institutions and industry. CBA is a supporting partner of the FSO Skills Accelerator-AI, led by Future Skills Organisation in partnership with Microsoft. Referenced in the Australian Government's National AI Plan, this initiative brings together leaders from vocational education and training (VET) and industry to connect, collaborate, and share best practices, expanding access to AI skills for the VET workforce.

## Supporting Australia's future workforce

We believe investing in Australia's future workforce is key to building a more inclusive and innovative economy. That's why we contribute to forums shaping the future of work, including our collaborations with Future Skills Organisation, NSW Digital Skills and Workforce Compact, and the Tech Council of Australia's Partners Forum. We support programs that promote equal access to opportunities in the fields of science, technology, engineering and mathematics (STEM), helping make AI skills more accessible.

### *Developing future AI talent*

We established the Data Science and Analytics Graduate Program to support the development of data scientists of the future. The 12-18 month program enables new Data Science and Analytics graduates to rotate across a range of business units at CBA; providing them with hands-on experience in quantitative and predictive modelling, AI, machine learning and more.

### *Expanding access to tech careers*

In February 2024, in partnership with the Tech Council of Australia and Year13, we sponsored a Data Scientist virtual work experience to offer young Australians, including those experiencing disadvantage or increased barriers to employment, the opportunity to experience careers in tech. The data science module has garnered positive interest, achieving more than 300,000 social video views, with 45% of the engaged audience identifying as female.

### *Hosting an Immersion Day for UNSW Business School's Australian Graduate School of Management (AGSM)*

CBA collaborated with AGSM to host an AI Immersion Day in October 2024. This event brought 40 MBA students from around the world together with our AI leaders to gain insights on our AI strategy and review case studies aimed at delivering better customer outcomes. Participants also engaged in a hands-on predictive AI model-building competition.

### *Inspiring young people to pursue careers in STEM*

We are working with the Australian Business and Community Network (ABCN) and Year13 to deliver programs to young people to build their AI awareness and skills. Since 2024, we delivered AI taster workshops to over 140 students as at 1 September 2025. These workshops were developed and delivered by volunteers from our Data Science and AI teams.

We also deliver the Innovate program, which seeks to inspire young people to pursue jobs in STEM. The Full STEAM Ahead program supports critical thinking development and coding skills. The Girls in Tech program is aimed at high school girls from across Sydney and nearby areas. As part of this program, CBA invites partners and other industry representatives to present interactive tech demonstrations showcasing careers in STEM, focusing on fields like Engineering, AI, Cyber Security and Data Science to help encourage more women to go into these fields.

# Our approach to adopting AI

# Leveraging our technology and risk management foundations

**We are building on existing risk management frameworks to manage risk in the ideation, development, deployment and management of AI Systems, supported by AI-specific requirements that help our people make responsible decisions. Our approach continues to mature as we integrate AI into our business and adapt to changes in the technology.**

### Being guided by our Code of Conduct

The use of technology at CBA, including AI, is guided by our Group Conduct Policy and Code of Conduct. Our Code of Conduct guides our decision making by describing the standards of conduct that we expect of our people. The 'Must We?', 'Can We?' and 'Should We?' questions help our people make decisions that comply with our legal and regulatory obligations, as well as being consistent with our Group Values of Care, Courage and Commitment, Policies and the outcomes we expect (including fair customer outcomes).

### Safeguarding personal information

We are conscious of the trust our customers place in us to collect, handle and protect their personal information and meet our regulatory obligations. The integrity and quality of data is fundamental to addressing customer needs and making sound business decisions. We have policies and standards in place to manage customers' personal information, according to our privacy obligations and we continue to invest in data platforms to manage the quality of our data across its lifecycle.

### Responding to customer complaints

Customers can make complaints about our products, services, and outcomes. We have processes in place to acknowledge, record and resolve customer complaints in line with our standards and regulatory obligations. Throughout the process, we require that customers be kept informed and seek to use what we learn to improve our products, services and outcomes.

### Managing AI incidents

Our risk framework requires that operational risk and compliance incidents - including those involving an AI System - are identified, recorded, assessed and resolved, and that a root cause analysis of each incident takes place. Accountable individuals are assigned to manage the incident and determine whether customer remediation is required.

### Protecting against cyber risk

AI and cyber risk are closely connected. Bad actors can exploit AI to scale their operations and increase the effectiveness of attacks, such as scams or phishing, while AI itself can introduce new cyber vulnerabilities. At the same time, AI offers powerful opportunities to help strengthen our defences. Understanding and managing these interdependencies is an important part of using AI safely and responsibly, including by:

- assessing cyber risk during the development and operation of AI Systems;
- leveraging AI to help protect our systems against AI-powered attacks; and
- using AI to help protect our customers and the Bank, by processing security data to identify signals that may indicate anomalous behaviour that could be an attack on our systems.

### Managing our technology environment

We leverage our existing foundational capabilities when building AI, which aim to keep our systems


safe, stable, and resilient. We continue to focus on maintaining robust engineering practices, resilient architecture, automated deployment pathways, and continuous monitoring. We also use AI to enhance and evolve these technology foundations over time.

### Working with third-party service providers

Third-party service providers who help develop, deliver or support an AI System must go through engagement and onboarding steps to help identify and manage any risks they might introduce.

We have processes for reassessing providers in the event of a material change to the service or during a routine review, in accordance with the Group's regulatory obligations including those relating to the use and management of CBA data.[21]

### Understanding and managing environmental impacts

We recognise that some AI technologies may have adverse environmental impacts. For example, the training and operation of AI models require significant computing power. The data centres that provide this computing power can consume water to maintain stable operations and require electricity, which can generate greenhouse gas emissions.

In line with the transition to a net zero emissions economy, we believe electricity generation, including for AI, should gradually shift to lower-emissions technologies. We are taking steps to monitor and manage key environmental risks associated with our use of AI. At CBA, we purchase renewable energy to match 100% of the electricity used across our Australian operations, including in the Australian data centres we operate. We also actively seek to improve water efficiency across our Australian-based operations.

As much of our AI computing power is provided by third-party service providers we work with our largest AI providers to understand their strategies for sourcing renewable or low emissions electricity, and where available, their water-related commitments.

As Australia's AI infrastructure grows, we are helping the electricity grid to decarbonise by financing renewable electricity. We are committed to supporting our customers and welcome policies that can help Australia deliver more affordable and reliable energy as the country transitions to net zero emissions.

# Maintaining governance and accountability

**The Board, Executive Leadership Team (ELT) and management committees play an important role in governing AI risk. In addition to existing forums, an AI-specific forum has been established to support policy development and the deployment of AI at CBA.**

## Board and Management-level committees

The Board has ultimate responsibility for the Bank's risk governance, including the Risk Management Framework (RMF) and oversight of its operation by management. The Board is supported by the Risk & Compliance Committee, Audit Committee, Nominations Committee and People & Remuneration Committee, as well as specific management-level committees that oversee material risks.

AI is a material risk type under the RMF. The Board sets risk appetite annually for all material risks. The policies for managing AI risk are subject to periodic review under the Group Policy Framework.

The Board holds the CEO and ELT accountable for the management of AI related risks and opportunities. The ELT's management of these risks and opportunities is supported by the application of a range of internal policies, standards and procedures that govern the way we deliver our products and services.

## Business-level governance forums

Business Unit Financial Risk Committees and Business Unit and Support Unit Non-Financial Risk Committees help manage risk in line with our RMF. As an example, these forums may review and evaluate AI models.

The Model Risk Governance Committee derives its authority from the ELT Financial Risk Committee. The purpose of the Model Risk Governance Committee is to oversee the design and operation of the Group's Model Risk Framework and support approval of the Group's models, including AI models.

The purpose of the AI Risk Committee is to oversee the design and operation of the Group's AI risk framework and provide risk management challenge and advice for higher-risk AI use cases. This includes advice on changes to the Group AI Policy and AI risk-related content in Group policy documents.

### Group Governance

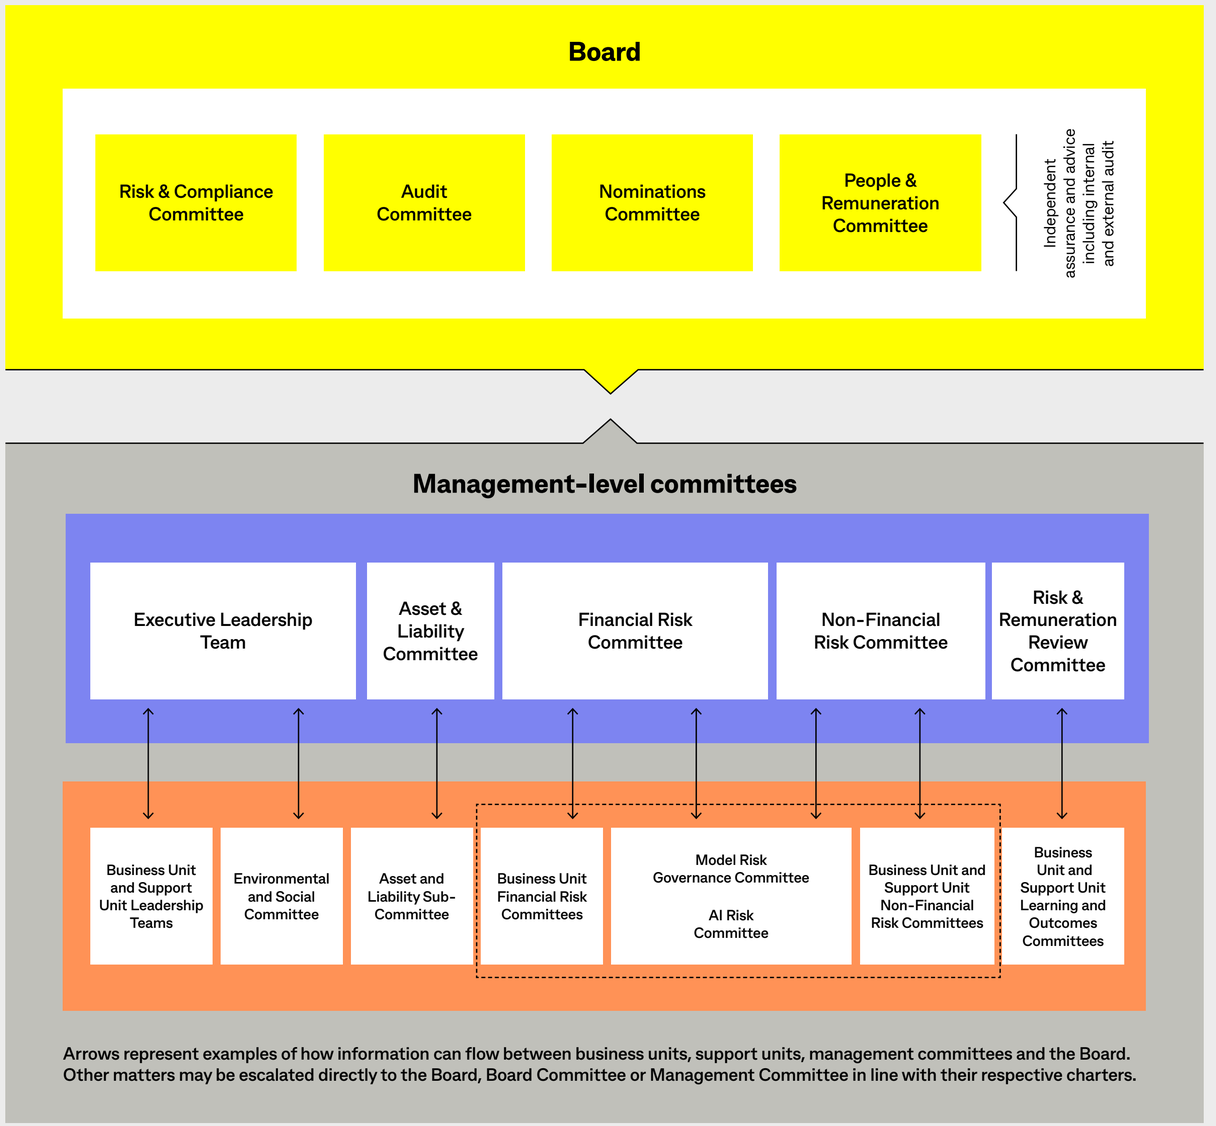


Arrows represent examples of how information can flow between business units, support units, management committees and the Board. Other matters may be escalated directly to the Board, Board Committee or Management Committee in line with their respective charters.

# Being guided by our AI principles

**Our principles help guide the application and development of existing policies, frameworks and resources to support responsible ideation, development, deployment and management of AI Systems.**

Our AI principles were informed by our Code of Conduct, Australia's AI Ethics Principles and the Organisation for Economic Co-operation and Development (OECD) AI Principles. The OECD AI Principles, adopted by the OECD in 2019, were the first intergovernmental standard on AI and outlined a values-based approach to AI principles.

The environmental impacts of AI Systems are being considered at a Group level. AI technology and the conversations around it are evolving rapidly. We'll continue to refine our approach as we learn.

### Environmental and Social

Impacts on environmental and social (including human rights) issues, should be considered in our approach to AI.

### Fairness

AI System outputs should not unfairly discriminate.

### Transparency

What an AI System does should be understandable.

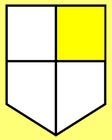

### Privacy and Data Protection

AI Systems should comply with privacy and data protection laws.

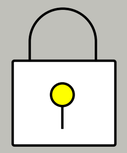

### Reliability and Security

AI Systems should perform consistently with their specifications and in accordance with their intended purpose. AI Systems should consider security measures that are proportionate to the potential risks.

### Accountability

Human oversight of AI Systems is necessary. The Group Employee(s) accountable for an AI System should be identified, and accountabilities for an AI System during its lifecycle should be documented.

# Applying responsible considerations throughout the AI lifecycle

**We have requirements, processes and guidance to support the responsible ideation, development, deployment and management of AI Systems.**

Requirements are incorporated into our Group AI Policy framework, which maps them across each stage of an AI System's lifecycle. They are taken from our existing policies and frameworks, supported by our AI principles.

The AI System lifecycle spans from ideation through to development and deployment and ongoing management. At each stage, there are requirements that may be applied, depending on the AI System and how it is intended to be used.

To support our developers through this process, we have toolkits that help our teams to build AI responsibly. These toolkits include practical guidance to support key modelling steps. The goal is to help our people understand AI, the risks, and help them to be able to identify and manage potential negative outcomes.

We expect our approach will continue to evolve over time.

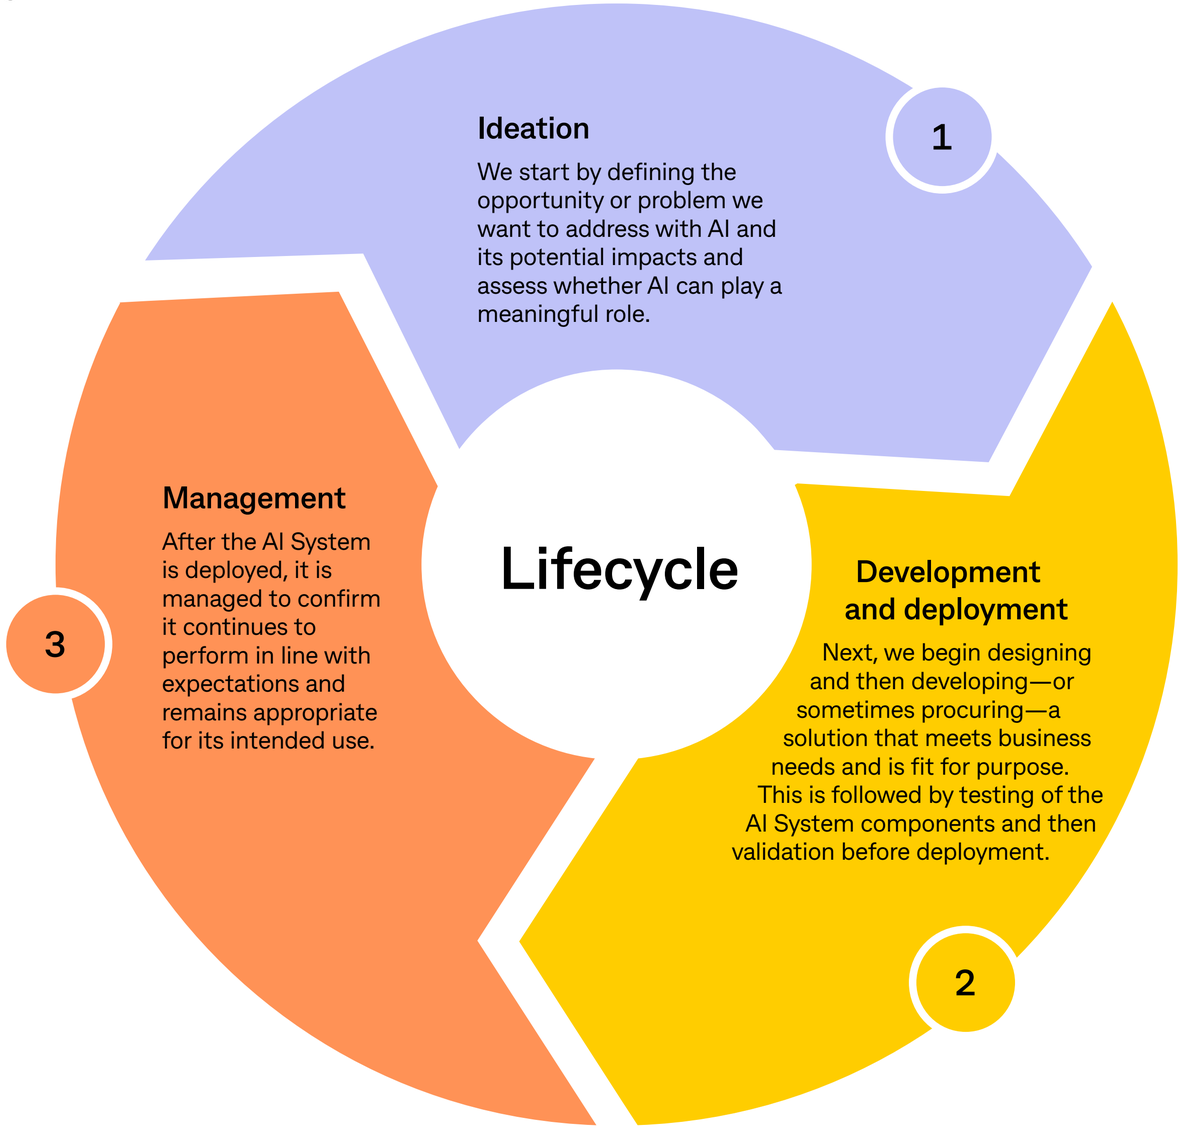


## Applying Guardrails in Ceba

We've developed AI Guardrails-as-a-Service (GaaS) to help manage key risks when using Gen AI. These guardrails are intended to act as safety mechanisms that help prevent the technology from producing harmful, inaccurate or inappropriate outputs.

### How it works

Ceba, our customer chatbot in the CommBank app, can use Gen AI to answer certain user queries based on information publicly available on the CommBank website, through a method called Retrieval-Augmented Generation (RAG).

In the RAG process, a customer question triggers the retrieval of information on the CommBank website, referred to as 'retrieved context.' This context is combined with the customer question and provided to a large language model (LLM) to generate a response.[20]

LLMs do not always guarantee accurate or factually correct answers. Instead, they generate responses based on probabilities and language patterns, which means they can potentially generate inaccurate statements.

To address this, Ceba uses the groundedness guardrail. When an LLM generates an answer, the customer question, retrieved context, and model response are sent to the groundedness guardrail. The guardrail evaluates the response to check it is fully supported by the retrieved context and flags whether the response may potentially be inaccurate or contains hallucinated information.

The groundedness guardrail is part of a wider collection of guardrails within Ceba that review and validate responses in real-time before they are sent to the customer.

## Providing frameworks that help teams apply requirements.

### 1. Ideation

During this stage we define the opportunity, or problem we aim to address. Key considerations during the ideation stage include:

***Understanding the potential impacts on customers, community and our people***

We've introduced a pre-screening process to help teams identify AI Systems that may require enhanced oversight.

If enhanced oversight is identified, the AI System is subject to additional reviews by subject matter experts and senior leaders before progressing into the development and deployment stage.

***Understanding the appropriate level of governance and accountability***

A preliminary model risk assessment is completed to determine the potential risks associated with the proposed AI model and the governance needed throughout the model's lifecycle. It may also help us to identify and evaluate actual and potential risks before we move into model development.

### 2. Development and deployment

We then begin to develop the solution. Key considerations during this stage include:

***Establishing AI transparency***

Documenting an AI model's development helps relevant teams understand the inputs, decisions and intentions behind it. Documentation is essential for later validation, where the model is assessed to confirm it performs as intended. For example, to support explainability of non-Gen AI models, we provide toolkits that recommend explainability methods. For Gen AI models, recording model inputs and outputs is required to support explainability. Explainability is commonly understood as the ability to understand and interpret how a model generates its output.

Our AI Systems are required to make clear to customers when they're directly interacting with AI. In rare cases, an exception may be made for a specific use case, such as in relation to fraud.

***Evaluating outcome fairness***

High quality data is important to help AI models produce reliable and accurate results. Developers are required to assess the data quality and its representativeness. They do this by selecting a representative sample on which to develop the AI model. Our toolkits provide methodologies and tools to identify and evaluate differences in outcomes, accuracy or error rates among different customer groups. They also help teams consider factors that support the development of AI models that are fit for their intended use.

***Applying safety measures***

An important aspect of maintaining reliable AI is having appropriate safety measures in place to help prevent unintended outcomes. These safety measures may include appropriate model shutdown or fall-back mechanisms that can safely deactivate an AI model in case of unintended behaviours.

***Understanding privacy and data protection***

To help identify and address privacy risks in AI Systems that involve personal information, all AI Systems are required to comply with our privacy obligations and policy requirements. We have processes in place to conduct privacy assessments to support developers in understanding the privacy and data protection laws that are relevant to AI Systems.

***Evaluating, assessing and validating model performance***

Before deploying a model, developers are expected to assess the data and test the model, including for accuracy. An independent review of the model is also required, which is conducted by a separate CBA team member independent of the model developer. The specifics of the review approach differ depending on the model's risk rating. This review includes assessing the relevance, completeness and potential bias in data, as well as evaluating the model performance. Our toolkits provide our developers and independent reviewers with practical guidance to assess model performance, including methods to assess the model's stability over time.

***Evaluating potential security vulnerabilities***

In addition to performance testing, AI-enabled applications may undergo security testing, such as penetration testing, prior to launch. These exercises are designed to test for cyber security vulnerabilities and assess both the AI models and surrounding infrastructure's response to various simulated and real world attack scenarios.

***Final checks before deployment***

Models require approval by an authorised individual before deployment. If the model's risk rating requires it, the approval process may also include technical committee oversight.

### 3. Management

After deployment, AI models are expected to undergo ongoing monitoring to confirm they remain compliant and effective.[21] Key considerations during this stage include:

***Establishing ongoing monitoring requirements***

Live models undergo periodic monitoring including fit for purpose reviews. These reviews assess the model's performance, as well as its operational risk and control environment, using both quantitative and qualitative evaluations of potential risks due to its use.

***Updating models***

Over time, models may be updated to improve performance. Depending on the type and materiality of the change, this may require independent review and approval.

***Decommissioning models***

If it is determined that the model is no longer required, or has been replaced, the model will be decommissioned.

## Detecting harmful messages in financial transactions

***A CBA developed AI model***

Family and domestic violence is a serious societal issue with far-reaching negative impacts on individuals, families, and communities. Since 2020, CBA's specialist Next Chapter team has been supporting customers experiencing domestic and family violence (DFV), financial abuse and problem gambling. The team provides free, confidential support to victim-survivors of DFV and financial abuse, connecting them to specialist services that help rebuild financial independence. These services can include crisis support, employment readiness programs, and services that help build financial confidence and capability.

In 2019, we identified more than 8,000 of our customers had received multiple low value (under $1) transfers with abusive transaction descriptions over a three-month sample period. To help combat this abuse, as a first step, a filter for key abusive words was implemented. However, a better solution was required to help detect the nuances of abuse. Leveraging AI, we built a monitoring system to help identify cases of serious abuse and began developing processes to support victim-survivors and deter perpetrators.

The AI model helps to identify digital payment transactions that include harassing, threatening or offensive messages – referred to as technology-facilitated abuse. Our AI model continues to identify over 3,600 cases each year which are further reviewed by our Next Chapter team.

The system was developed using a combination of machine learning, natural language processing and pre-trained LLMs. It was trained on verified anonymised high-risk abuse cases from CBA transaction data and is designed to detect patterns of sustained abuse based on key criteria. Multiple model configurations were tested to determine the most effective approach for identifying abusive transaction descriptions. To guide our response to the model's outputs, we consulted with individuals with lived-experience of domestic and family violence to help shape an intervention framework tailored to this form of abuse.

The anonymised high-risk cases identified by the model are reviewed by the Next Chapter team. Depending on the severity of abuse identified and with the consent of the customer, the Next Chapter team can issue warning letters and place digital banking restrictions on customers sending abusive messages. Pathways have also been established to provide customers with access to a network of external specialists as appropriate.

To increase the model's robustness, it is regularly retrained using the anonymised cases reviewed by the Next Chapter team that are confirmed to be abusive. The model also undergoes an annual review to confirm it continues to operate as intended.

CBA has made its research publicly available to share learnings and receive feedback from the wider research community. In addition, the team collaborated with H2O.ai to make CBA's pre-trained model available to other financial institutions both domestically and globally.

# Strengthening our people's skills

**Our people are central to how we deliver for our customers. As we continue to use AI to enhance customer experiences and transform how we work over time, we are creating pathways to help develop our people's skills for the future.**

### Supporting our people through change

Technology is transforming the way we work. We acknowledge that, as a result, over time, some roles will change or will no longer be required; however, new roles will also emerge. Throughout this change, human qualities such as judgement, empathy, adaptability and creativity remain essential.

Our people provide oversight and accountability, and we recognise that their experiences, integrity and judgement are central to how we deliver on our purpose. Where the adoption of AI results in changes to roles, employees will be engaged through our established consultation processes, with support from Human Resources.

To help our people build sustainable careers as roles and skills they require evolve, we are expanding skilling programs and access to advanced AI tools. Ongoing employee surveys explore views and experiences of AI at CBA, with insights informing our approach to AI adoption, change management and employee engagement over time.

### Building our people's AI capabilities

AI can support our employees in their day-to-day work and we encourage our people to experiment, learn from each other, and apply AI in ways that improve customer outcomes, while remaining firmly in control of decisions. To take full advantage of the benefits AI offers, we are taking a new approach to skilling our people.

In 2024, we launched an AI learning series, which is designed to help build confidence and knowledge, so employees understand when and how to use AI responsibly. Since launch, over 27,600 employees have engaged with this learning series. Building on this momentum, in July 2025 we launched the next phase of our AI learning which has seen us shift to experiential learning, with many employees participating in fortnightly hands-on activities, prompting workshops, and using AI to support their work.

We have also designed a learning program for our leaders to develop the skills to harness the power of AI and lead it responsibly. Through this program our ELT and over 300 senior leaders have been trained on topics such as Gen AI and responsible AI. In addition, as at 1 September 2025 more than 220 product leaders were supported to identify when Gen AI might be an appropriate solution for their customers.

To further support responsible AI ideation, development, deployment and management, we launched an AI risk learning pathway in November 2024. This pathway aims to build risk professionals' understanding of AI-related risks and mitigations, with over 8,700 participants.

### Developing the capabilities of our technical teams

We are investing in the ongoing education and professional development of our technical teams, providing them access to the latest AI tools, technology and knowledge to help them excel in their roles. We partner with leading academic partners and tech organisations to help us achieve this.

Seattle Tech Hub team onsite, March 2025

### Establishing a Tech Hub in Seattle, Washington

In March 2025, we launched our newest technology hub in Seattle, USA. This hub offers our technology teams a three-week exchange in a collaborative learning environment with global leaders like AWS and Microsoft, combining formal training, case studies, and hands-on learning in advanced AI.

Through multiple cohorts, technologists are accelerating AI adoption and exploring new ways to improve how we respond to customer needs through these learning experiences. The Seattle Tech Hub reflects our focus to fast-tracking technologies such as agentic AI while fostering growth and career development for our people. Teams return to Australia with practical learnings, while boosting the knowledge and expertise in Australia's tech ecosystem.

### Collaborating for innovative solutions

As we continue learning from global experts, we're also focused on sparking innovation through collaboration with our people. One way we do this is through hosting hackathons – short, intensive events where teams from across the bank come together to develop solutions to specific challenges or problems. These events foster creativity, problem-solving, and cross-functional teamwork in a fast-paced, collaborative environment. In 2025, we held a hackathon focused on Gen AI and how we might utilise its potential to create innovative solutions for our customers.

# Sources and notes

**For ease of navigation, each source is linked directly to the page where its corresponding end note appears.**

1. Definition informed by the ISO: Artificial intelligence (AI) refers to computer systems capable of performing tasks that typically require human intelligence, such as reasoning, learning, perception, and language understanding.

2. The report focuses on CBA's approach to the ideation, development, deployment and management of AI Systems, and does not address employee use of AI or the creation of AI assistants by employees for internal use.

3. Customers includes the number of customers who have a relationship with the Commonwealth Bank of Australia (CBA) and Bankwest as at 30 June 2025. A CBA customer is defined as anyone who is currently associated with an open account as either the owner, joint owner, trustee or primary cardholder. A Bankwest customer is defined as anyone who holds an open account. Includes retail and non-retail customers and deceased estates.

4. Total full-time equivalent (FTE) employees for Australia as at 30 June 2025. FTE includes full-time, part-time, job share employees, employees on extended leave and contractors. One full-time role is equal to 38 working hours per week.

5. Definition informed by OECD: AI System is a machine-based system(s) 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.

6. Machine learning is a subset of AI that involves models that learn from large amounts of data, detect patterns and make predictions and recommendations.

7. Time frame reviewed for average events per day screened using AI models is 1 July – 31 December 2025.

8. Evident AI Index: Key Findings Report, November 2023.

9. Definition informed by the Department of Industry Science and Resources: Gen AI is a branch of AI that develops generative models with the capability of learning to generate new content such as images, text, and other media with similar properties as their training data.

10. InnerSource is a software development strategy that applies open source practices to proprietary code. InnerSource can help establish an open source culture within an organisation while retaining software for internal use.

11. Evident AI Index: Key Findings Report, October 2024.

12. Evident AI Index: Key Findings Report, October 2025.

13. All currency in this report is in Australian dollars (AUD), unless otherwise stated.

14. Data sourced from April 2024 – December 2025.

15. Data sourced from January – December 2025.

16. Australia's Generative AI Opportunity, Microsoft and Tech Council of Australia, July 2023.

17. Work Change Report: AI is Coming to Work, LinkedIn, 2025.

18. Bratanova A*, Hajkowicz SA*, Evans D, Chen H, Bentley S, Pham H, Hartman S (2025): Australia's artificial intelligence ecosystem: The 2025 update. Australian Government Department of Industry, Science and Resources National Artificial Intelligence Centre, CSIRO. *Joint first authors.

19. Artificial Intelligence Index Report 2025, Stanford University.

20. Large Language Model (LLM) is an AI model trained on extensive amounts of textual data capturing deep knowledge on how humans communicate and use language. LLMs can be used for several language-related AI tasks, such as conversation and generation. Chatbots, for example, are often LLMs.

21. Engagement of a third-party service provider for AI System development, delivery or support follows the process as outlined on page 12.

## Who to contact

For further feedback or questions on the report please contact us at **socialimpact@cba.com.au**.

Commonwealth Bank