ASIC
Australian Securities & Investments Commission

Introducing mandatory guardrails for AI in high-risk settings

Submission by the Australian Securities and Investments Commission

October 2024

Contents

Overview

1 The Australian Securities and Investments Commission (ASIC) welcomes the opportunity to make a submission to the Department of Industry, Science and Resources (department) in response to *Safe and responsible AI in Australia—Proposals paper for introducing mandatory guardrails for AI in high-risk settings* (proposals paper).

2 The use of artificial intelligence (AI) is already prevalent in the financial services sector and is continuing to develop in scale and complexity. The use of AI offers the possibility of significant efficiencies in the provision of financial services and in financial markets.

3 However, the use of AI in financial services can amplify existing risks and create new distinct risks to consumers and investors. AI may also pose significant risks to the operation and stability of markets.

4 We acknowledge likely gaps and uncertainties identified in the proposals paper in the application of existing Australian laws to preventing or mitigating AI-related risks, particularly in the development stage of the AI lifecycle.

5 In light of this, ASIC supports the introduction of ex ante regulatory measures to mandate guardrails for the use AI in high-risk settings.

6 We recognise that AI is evolving at a rapid pace and developing regulation in this environment is challenging. ASIC will continue to engage and work with the department on its actions to support safe and responsible AI.

7 Table 1 summarises our responses to specific questions raised in the proposals paper.

Table 1: Overview of ASIC's submission

Submission reference	Summary of ASIC's response
Defining high-risk AI (responses to questions 1–3 and 5–6): Section B	A principles-based approach has the benefit of flexibility and adaptability as new technologies and practices emerge. However, we consider that regulators, organisations and consumers would benefit from further detail and clarity about the principles proposed to guide the use of AI in high-risk settings. To illustrate, we see benefit in: • clarifying the application of the principles to financial services and markets; and • assessing the systemic impact of AI systems. We support the proposal to apply the mandatory guardrails to all general-purpose AI (GPAI) models due to the lack of foreseeability in how these models could be used and their capacity to cause harm on a wide-scale and at speed. The approach taken by ASIC to establish our Indigenous Financial Services Framework could be considered when developing AI principles that better capture potential harms to, and the diverse needs of, First Nations peoples.
Guardrails ensuring testing, transparency and accountability of AI (responses to questions 8 and 10): Section C	ASIC is generally supportive of the proposed mandatory guardrails for high-risk AI. However, we see a number of challenges that will need to be considered. These include: • reconciling the guardrails with existing overlapping laws; • considering the effectiveness of disclosing the use of AI; and • ensuring that responsibility is appropriately allocated across a complex AI supply chain.
Regulatory options to mandate guardrails (responses to questions 13 and 15): Section D	ASIC considers that Option 1 and Option 2 of the options for introducing mandatory guardrails will better support a framework that clearly attributes responsibilities to the relevant regulators with minimal overlap. We recognise that Option 3 may enable a more efficient approach to regulating AI supply chain guardrails, but consider there are some complexities with this approach for existing regulators, organisations and consumers.

A ASIC's position on AI

8 ASIC is Australia's integrated corporate, financial services, consumer credit and markets regulator. Our role includes facilitating and improving the performance of the financial system and the entities within it. Our role also includes promoting confident and informed participation by investors and consumers in the financial system.

9 ASIC supports innovation in the financial system that is balanced with appropriate consumer protections and market integrity safeguards. As such, we consider that the safe and responsible use of AI may be realised through strong governance, transparency and accountability, including human oversight, as well as robust information security.

10 Understanding and responding to the use of AI across the entities we regulate is a key priority for ASIC. We have a range of work underway relating to AI. We are actively monitoring the deployment of AI across our regulated population to better understand how it is used, as well as its benefits and risks. We are also engaging with industry, other regulators (both domestically and internationally) and the public on AI developments.

11 ASIC has issued guidance setting out our expectations around the laws applicable to the use of technologies. We also have projects underway that more specifically target the use of AI by regulated entities:

(a) ASIC recently reviewed the use of AI by 23 licensees in the retail banking, credit, general and life insurance and financial advice industries where AI interacted with or impacted consumers. We analysed information on over 600 AI models (including advanced data analytics) that were in use or being developed as at December 2023, and asked licensees about their risk management and governance arrangements for AI and their future plans. The findings of our review will be released in the fourth quarter of 2024.

(b) ASIC plans to consult in the 2024–25 financial year on amending the Market Integrity Rules for both securities markets and futures markets to address current and emerging risks, including AI.

12 Internally, we are investing in our use of data and digital technology to become a leading digitally enabled and data-informed regulator. Part of this vision is the safe and responsible adoption of AI to support our regulatory work.

13 We have already undertaken a number of 'proof of concept' trials to test the capability of AI to support ASIC's work, particularly high-volume or manual work. We are continuing to explore and refine how AI could be adopted into our work, as well as developing appropriate governance structures and building capabilities within teams.

B Defining high-risk AI

Key points

This section outlines our feedback on questions 1–3 and 5–6 of the proposals paper.

A principles-based approach has the benefit of flexibility and adaptability as new technologies and practices emerge. However, we consider that regulators, organisations and consumers would benefit from further detail and clarity about the principles proposed to guide the use of AI in high-risk settings.

To illustrate, we see benefit in:

- clarifying the application of the principles to financial services and markets; and
- assessing the systemic impact of AI systems.

We support the proposal to apply the mandatory guardrails to all GPAI models due to the lack of foreseeability in how those models could be used and their capacity to cause harm on a wide-scale and at speed.

The approach taken by ASIC to establish our Indigenous Financial Services Framework could be considered when developing AI principles that better capture potential harms to, and the diverse needs of, First Nations peoples.

14 The proposals paper proposes taking a risk-based approach to regulating AI in Australia, only applying mandatory guardrails to AI systems that are developed and used in settings considered high risk. The paper also proposes effectively treating all GPAI models as high risk (regardless of setting), based on the lack of foreseeability in how these models might be used and their capacity to cause harm to people, community groups and society on a wide-scale and at speed.

15 ASIC's understanding is that the onus will be on organisations to determine whether an AI system is high risk and that for the first category of AI—systems developed and used in high-risk settings—organisations will be required to consider two related questions:

(a) Is the setting or context in which an AI system is intended to be used high risk?

(b) Is the known or foreseeable use of the AI system in that setting or context also high risk?

16 The proposals paper aims to ensure that the assessment of high risk is consistent across organisations by proposing a set of common principles for organisations to contemplate when considering the above questions. These principles focus on the use of an AI system (in its setting or context) and generally relate to the risk of adverse impacts on fundamental rights and safety.

Questions 1 and 3—Clarity, certainty and scope of the proposed principles

17 The 'high risk' determination by organisations will effectively operate as a threshold prerequisite to the application of mandatory guardrails. As a result, it will be important for any framework setting out how to determine high-risk settings is clear and objective. This will provide regulatory certainty to support a consistent application of the principles and ensure that they are able to be enforced.

18 While we acknowledge that the proposals paper is a preliminary step in the law design process, our view is that it is currently unclear how the principles proposed to guide consistent assessment of high-risk settings by organisations would be applied in practice.

19 We consider that there are advantages to a principles-based approach to determining high-risk settings. For example, a principles-based approach may facilitate regulatory flexibility and ensure that the guardrails framework is adaptable to a range of diverse products and industries, and more easily extend to developments in technology.

20 A principles-based approach to determining high-risk settings may also introduce uncertainty and ambiguity for organisations and regulators around whether the guardrail obligations should apply, leading to inconsistent application of a high-risk setting across industries and the economy.

21 For example, without objective criteria or thresholds it will be difficult to ascertain whether—and at what point—a risk that impacts fundamental rights and safety becomes high risk. This question might arise in a situation where an AI system is involved in a single but significant aspect of a decision impacting legal rights (such as a decision on creditworthiness or on an element of an insurance claim), but does not solely determine an individual's legal rights.

22 We consider that the current framing of high-risk settings will require further prescription or guidance (e.g. by regulators) to ensure that it is applied clearly and consistently.

Application of high-risk settings to financial services and markets

23 The application of high-risk AI systems within financial services, consumer credit and financial markets is a key concern for ASIC. We see a number of distinct risks to consumers that could vary in severity depending on how the systems are deployed and the controls put in place. AI systems may also pose significant risks to the operation and stability of markets. However, the

extent to which the proposed principles for high risk would apply to the financial sector is currently unclear.

24 The risks and harms that arise in the financial sector—including banking, credit, insurance, superannuation, financial advice and managed investments—can have real and significant impacts on consumers and small business. The complex nature of financial products and services often amplifies the extent of these risks and harms—the products and services are intangible, may require trade-offs between present and future benefits and often involve risk or uncertainty.

25 Other factors that can amplify the impact on consumers and small business include the long-tail exposure of any loss or harm, the amount of money involved, and the disproportionate impact on consumers when things go wrong.

> Note: For examples of the risks and harms exposed by ASIC, see Report 751 *Disputes and deficiencies: A review of complaints handling by superannuation trustees* (REP 751), Report 778 *Review of online trading providers* (REP 778), Report 788 *Credit card lending in Australia: Staying in control* (REP 788), and Report 790 *Anti-scam practices of banks outside the four major banks* (REP 790). For further discussion about complexity in financial services and product design, processes and choice architecture, see Report 632 *Disclosure: Why it shouldn't be the default* (REP 632).

26 Given the factors above, we see a number of settings in our remit that are high risk for the use of AI systems—for example:

(a) In the superannuation industry, poor or incorrect advice in relation to an individual's superannuation could lead to a deterioration or depletion of retirement savings.

(b) In the credit industry, AI-powered credit scoring to evaluate creditworthiness may unfairly discriminate and risk financial exclusion.

(c) Financial markets may experience market misconduct where machine learning-based trading algorithms can learn to manipulate market prices when designed to seek out and generate profitable trades.

27 Although these high-risk examples may be intended to be captured by the principles—because, for example, they may have adverse legal effects on an individual or have systemic impacts to the broader economy—the current framing of the principles does not provide sufficient certainty to understand whether the guardrails will apply. We think there is opportunity to provide further clarity and detail around high-risk settings and how they apply to financial services and markets.

> Note: For example, Annex III of the *Artificial Intelligence Act* (EU) specifies a list of high-risk areas for the purposes of determining high-risk AI. One section relates to 'access to and enjoyment of essential private services and essential public services and benefit'. It provides further clarity around what constitutes 'essential private services' by identifying certain AI systems used in the credit and insurance industries.

Assessing systemic impacts of AI systems

28 One of the principles to guide high-risk use of AI proposed in the proposals paper relates to the systemic impact of AI systems—that is, the risk of adverse impacts to the broader Australian economy, society, environment and rule of law.

29 In relation to the assessment of whether the use of AI systems will meet these criteria, we query whether organisations will always be best placed to assess whole-of-ecosystem impacts.

30 There may be circumstances where an AI system or model is considered low risk by an organisation, but, when used collectively or at scale by multiple organisations within an industry through the use of similar models or similar underlying data (i.e. 'herding behaviour'), there are significant risks if the system or model produces adverse outcomes. This can result in impacts at a systemic level, affecting market stability, increasing volatility and exacerbating extreme market events.

> Note: This kind of risk was the subject of commentary by the chair of the Securities and Exchange Commission in July 2023:
>
> > AI may heighten financial fragility as it could promote herding with individual actors making similar decisions because they are getting the same signal from a base model or data aggregator. This could encourage monocultures. It also could exacerbate the inherent network interconnectedness of the global financial system.

31 While the proposals paper does discuss emergent systemic risk, it is unclear how the proposed principles would address situations such as those outlined above. We consider that regulators will need to play an important role in understanding the industries in which this kind of 'systemic high risk' might occur.

Questions 5 and 6—Applying the guardrails to GPAI models

32 In our view, the principles proposed to guide whether AI is being used in high-risk settings are unlikely to be flexible enough to capture GPAI models, given their focus on the foreseeable or intended use of AI systems in a defined context or setting.

33 In light of this, ASIC supports the proposal to effectively treat all GPAI models as high risk and to apply the mandatory guardrails to these models on the basis of a lack of foreseeability in their use or misuse and their capacity to harm people, community groups and society on a wide-scale and at speed.

Question 2—Capturing harms to First Nations communities

34 The proposals paper notes that AI systems have the potential to create highly unequal or damaging outcomes for specific communities and perpetuate existing inequalities. We have observed the risks and harms that can arise more broadly through the unscrupulous conduct of some financial service providers. This includes practices that target particular groups, with disproportionate impacts not only on individuals but also groups of people—for example, recipients of Centrelink payments in remote First Nations communities.

35 ASIC is also aware of potential algorithmic bias that may lead to unfair outcomes for individuals and groups and could create harmful practices that impact First Nations consumers, communities and Country.

36 ASIC published our Indigenous Financial Services Framework in February 2023, outlining four Key Learnings relating to First Nations peoples' access to and engagement with the Australian financial system.

37 The long-term outcomes set out in the Framework that ASIC and other stakeholders aim to achieve (e.g. reduced impact of harms and misconduct on individuals and communities) were developed in consultation and collaboration with First Nations peoples. A similar approach in the development of the AI principles could better capture potential harms to, and the diverse needs of, First Nations peoples.

C Guardrails ensuring testing, transparency and accountability of AI

Key points

This section outlines our feedback on questions 8 and 10 of the proposals paper.

ASIC is generally supportive of the proposed mandatory guardrails for high-risk AI. However, we see a number of challenges that will need to be considered. These include:

- reconciling the guardrails with existing overlapping laws;
- considering the effectiveness of disclosure of the use of AI; and
- ensuring that responsibility is appropriately allocated across a complex AI supply chain.

38 The proposals paper proposes addressing the risk of AI in high-risk settings by introducing 10 mandatory requirements for organisations seeking to develop or deploy high-risk AI systems—that is, the 'mandatory guardrails'. These mandatory guardrails will impose testing, transparency and accountability obligations on organisations. They will allocate different levels of responsibility across the AI supply chain (between developers and deployers of high-risk AI systems) applicable throughout the AI lifecycle. The proposed guardrails do not extend to end-users of these systems.

Question 8—Appropriateness of the guardrails

39 ASIC is generally supportive of the proposed mandatory guardrails for high-risk AI systems, as we consider that the safe and responsible use of AI may be realised through strong governance, transparency and accountability.

40 To ensure the effectiveness of the guardrails, we consider they should be reconciled with existing laws where they overlap to prevent unnecessary complexity in how an organisation's obligations are interpreted and to assist regulators in enforcing the obligations.

Interaction with existing ASIC-administered laws

41 Australia's financial system is heavily regulated and includes a range of laws administered by ASIC relating to corporate governance, transparency and accountability. As a consequence, the guardrails proposed in the proposals paper overlap with existing requirements under corporations and financial services legislation.

42 ASIC agrees with the observation in the proposals paper that obligations under the guardrails should not replace or exempt organisations from existing statutory obligations.

43 However, it will be important to reconcile and align the guardrails with overlapping laws. In particular, this will be the case where the guardrails are intended to impose different standards of conduct on regulated entities or where the allocation of the guardrails along the AI supply chain conflicts with general obligations that may apply to an entity. We would not support implementation of the guardrails in a way that weakens the general standards of conduct that apply to entities under other legislation.

> Note: For example, under Guardrail 6 it is proposed that participants meet a 'best efforts' standard for informing end-users on how AI is being used. In contrast, other laws relating to product disclosure or representations made in trade or commerce impose a higher standard. See, for example, the prohibition against false or misleading representations under s12DB of the *Australian Securities and Investments Act 2001*.

44 Following are examples of where the proposed guardrails will overlap with ASIC-administered legislation and may require more detailed consideration:

(a) **Preventative action**—The proposed guardrails establish a framework of preventative action, imposing testing, transparency and accountability requirements on developers and deployers to ensure the safety of AI in high-risk settings. Australian financial services (AFS) and credit licensees are similarly required to take preventative action to comply with the general obligation requiring licensees to do all things necessary to ensure that services under their licence are provided efficiently, honestly and fairly.

> Note: See s912A(1)(a) of the *Corporations Act 2001* (Corporations Act) and s47(1)(a) of the *National Consumer Credit Protection Act 2009* (National Credit Act).

(b) **Accountability and risk management**—Guardrails 1 and 2 propose that organisations establish accountability and risk management processes to address risks arising from high-risk AI systems and outline governance policies and clear roles to ensure compliance with the guardrails. Similar requirements are imposed on:

(i) AFS licensees under the Corporations Act and credit licensees under the National Credit Act;

(ii) the directors and officers of companies under the Corporations Act;

(iii) market operators and participants as licensed market operators and AFS licensees under the Corporations Act and under the *ASIC Market Integrity Rules (Securities Markets) 2017* and *ASIC Market Integrity Rules (Futures Markets) 2017*; and

(iv) accountable entities and persons under the *Financial Accountability Regime Act 2023*.

(c) **Dispute resolution**—Guardrail 7 proposes to introduce a requirement for organisations to establish processes for people negatively impacted by high-risk AI systems to contest AI-enabled decisions or make complaints about their experience or treatment. Similar requirements are imposed on entities governed by credit and financial services laws. For example, AFS and credit licensees are required to have a dispute resolution system in place comprising:

(i) internal dispute resolution procedures that comply with standards made or approved by ASIC and cover complaints made by retail clients against the licensee about the provision of all services covered under the relevant licence; and

(ii) membership of the Australian Financial Complaints Authority scheme (i.e. an external dispute resolution body).

Informed end-users and the limits of disclosure

45 The proposals paper proposes to introduce a requirement at Guardrail 6 for organisations to inform end-users about how AI is being used and where it affects them. Under this requirement, organisations must inform people when AI is used to make or inform decisions relevant to them and when they are directly interacting with an AI system. This could include methods such as content labelling, watermarking or unique markers.

46 We consider transparency is important as it allows for greater engagement and more informed decision making and generally contributes to better operating financial markets. Therefore, we support the objective of Guardrail 6 to ensure people are informed when AI is being used for decision making or when they are interacting with an AI system.

47 While disclosure is a necessary part of this, our previous experience and evidence shows that there are limits to the effectiveness of disclosure (e.g. warnings) in protecting consumers, enabling good decision making and driving competition from the demand side. For example, warnings can be ignored, overlooked, misunderstood or misremembered. They can have no impact on peoples' behaviour, and can even have adverse impacts. We consider these factors should be taken into account when further developing Guardrail 6.

Note: See REP 632 for further information.

Question 10—Guardrail responsibility across the supply chain

48 There is likely to be a range of complexities in applying the guardrails across the supply chain, particularly around how 'developers', 'deployers' and 'end users' are defined and how these definitions will capture the complex ecosystem of supply chain participants (such as those outlined in Figure 3 of the proposals paper, as well as agents, outsourced arrangements and offshore participants).

49 It will be important to delineate supply chain roles clearly, especially in circumstances where responsibility for the guardrails is intended to be allocated (and therefore monitored and enforced) differently for developers and deployers. Because of this, we think that there will need to be appropriate safeguards in place to address potential gaps in oversight or accountability-shifting between supply chain participants. One way of achieving this could be through imposing a general overarching obligation on all supply chain participants regarding the safety of an AI system.

> Note: For example, although AFS and credit licensees are able to outsource functions relating to their licence—such as administrative or operational functions—they remain responsible for meeting and complying with their obligations as a licensee, including their responsibility for ensuring that services are provided efficiently, honestly and fairly: see Regulatory Guide 104 *AFS licensing: Meeting the general obligations* (RG 104) and *ASIC v MobiSuper Pty Ltd and Others* [2021] FCA 855, at [2]–[3].

50 There is a high degree of interdependency between the guardrails, with a number of requirements dependent on the actions of other supply chain participants. Given this, the guardrails framework should include appropriate mechanisms to deal with situations where one or more supply chain participants do not meet their obligations, impacting the ability of other supply chain participants to meet their requirements.

51 It will also be important for the framework to anticipate and address situations where participants perform multiple roles or transition between developer, deployer and end user during the AI lifecycle, as well as situations where supply chain participants may no longer be covered by the guardrails—for example, if they no longer develop high-risk AI or undergo changes to their business through mergers, acquisitions or insolvency.

D Regulatory options to mandate guardrails

Key points

This section outlines our feedback on questions 13 and 15 of the proposals paper.

ASIC considers that Option 1 and Option 2 of the options for introducing mandatory guardrails will better support a framework that clearly attributes responsibilities to the relevant regulators with minimal overlap.

We recognise that Option 3 may enable a more efficient approach to regulating AI supply chain guardrails, but consider there are some complexities with this approach for existing regulators, organisations and consumers.

52 The proposals paper outlines three regulatory options to mandate the proposed mandatory guardrails:

(a) Option 1 (domain specific approach) where the guardrails are adopted into existing regulatory frameworks as needed;

(b) Option 2 (framework approach) where framework legislation is introduced and incorporated into regulatory frameworks across the economy; and

(c) Option 3 (whole-of-economy approach) where a new cross-economy AI-specific Act is introduced, overseen and enforced by a single AI regulator (either a new regulator or an existing regulator with expanded powers).

Questions 13 and 15—Adaptability, responsiveness and appropriateness of regulatory options

53 ASIC is experienced in administering laws under multiple regulatory frameworks, including where they overlap with or relate to the responsibilities of other regulators. We regularly share information and collaborate with other regulators to find efficient and effective ways to regulate overlapping industries.

54 We are generally supportive of regulatory arrangements that contain clear areas of regulator responsibility with minimal overlap. We are also supportive of regulatory frameworks that include clear pathways for regulatory action—including enforcement—clear avenues for consumer redress, and penalties for non-compliance that go beyond the cost of doing business.

55 In our view, both Option 1 (domain specific approach) and Option 2 (framework approach) will achieve these objectives by leveraging existing regulatory frameworks and arrangements. However, Option 2 is more likely to provide baseline consistency in how AI in high-risk settings is addressed across the economy. This option is also likely to be sufficiently flexible to ensure that the guardrails can be adapted to diverse areas of regulatory responsibility and their unique features.

56 At the same time, we recognise that it will be challenging to monitor and enforce AI supply chain guardrails across the economy where there might be existing gaps in regulatory coverage or jurisdictional limitations. Option 3 (whole-of-economy approach), with its proposed standalone AI act and regulator, would likely cure these limitations.

57 However, we do consider that a new regulator (or an existing regulator with an expanded remit) would add further complexity for:

(a) existing regulators when enforcing laws where there are overlapping remits or unclear boundaries or where the regulated population of a new regulator is not clearly defined;

(b) organisations in understanding how multiple regimes apply; and

(c) consumers when determining how to seek redress where harms have occurred.

58 Existing regulators also hold a deep body of industry-specific knowledge and expertise, which can be applied effectively to new risks and harms and may be utilised less effectively under Option 3.

59 Regardless of the option adopted, we think there will be a range of additional issues that may impact the effectiveness of the guardrails regime. We suggest these should be considered when designing the regulatory framework—for example:

(a) whether the regime will operate extraterritorially given that many developers and deployers are based outside of Australia;

(b) the investigations and enforcement powers available to regulators and the penalties that will apply for breaches of the mandatory guardrails;

(c) sufficient information sharing and delegation powers between regulators to ensure efficient and effective monitoring and enforcement, particularly where there are overlaps between the guardrails and existing laws; and

(d) the intended role of regulators and whether this will extend beyond monitoring and enforcement—for example, conducting conformity assessments under Guardrail 10 and whether it is most efficient for this to be undertaken by regulators given the significant and growing number of AI models being used in Australia.

60 With the proposals likely to impact all regulators in some way, even under Option 3, there will need to be processes established to enable access to relevant expertise to support the regulation of the proposed guardrails. ASIC will continue to engage with the department to support the safe and responsible use of AI. As part of this, we are interested in any initiatives to ensure a cohesive and collaborative framework.

Key terms

Term	Meaning in this document
AFS licensee	A person who holds an Australian financial services licence granted under s913B of the Corporations Act
AI	Artificial intelligence
AI system	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
AI lifecycle	All events and processes that relate to an AI system's lifespan. This spans from inception to decommissioning, including its design, research, model development, training, deployment, integration, operation, maintenance, sale, use, and governance
AI supply chain	The complex network of actors and organisations that enable the use and supply of AI throughout the AI lifecycle from model design, testing and fine tuning to deployment and integration into the local IT system
ASIC	Australian Securities and Investments Commission
Corporations Act	*Corporations Act 2001*, including regulations made for the purposes of that Act
credit licensee	A person who holds an Australian credit licence under s35 of the National Credit Act
department	Department of Industry, Science and Resources
deployer	Any individual or organisation that supplies or uses an AI system to provide a product or service. Deployment can be for internal purposes or used externally impacting others, such as customers or individuals
developer	Organisations or individuals who design, build, train, adapt or combine AI models and applications
end user	Any intended or actual individual or organisation that consumes an AI-based product or service, interacts with it or is impacted by it after it is deployed
GPAI model	General-purpose AI model—An AI model that is capable of being used, or capable of being adapted for use, for a variety of purposes, both for direct use as well as for integration in other systems

Term	Meaning in this document
National Credit Act	*National Consumer Credit Protection Act 2009*
proposals paper	Department of Industry, Science and Resources, *Safe and responsible AI in Australia—Proposals paper for introducing mandatory guardrails for AI in high-risk settings* [proposals paper], September 2024
s912 (for example)	A section of the Corporations Act (in this example numbered 912), unless otherwise specified