Guidance for AI adoption: implementation guidance

Strengthen governance for complex and higher-risk artificial intelligence use

On this page

How to use this guidance
Things to keep in mind
1. Decide who is accountable
2. Understand impacts and plan accordingly
3. Measure and manage risks: implement AI-specific risk management
4. Share essential information
5. Test and monitor
6. Maintain human control
Download

This version of the essential practices helps your organisation strengthen how it governs AI.

This guidance is for teams that:

- build or customise AI systems
- use AI in more complex ways
- manage higher-risk use cases
- need stronger controls and oversight
- want more detailed or specialist guidance.

If your organisation is earlier in its AI use, start with the foundations guidance.

How to use this guidance

Work through all 6 essential practices to strengthen governance, manage risk and maintain human oversight.

Start with the areas that matter most to your current systems and risks. You don't need to do everything at once.

Use this guidance to improve testing, monitoring, accountability and supply chain controls as your AI use grows.

Things to keep in mind

Some practices apply across your whole organisation. These includes:

- governance frameworks
- clear roles
- AI registers
- supply chain accountability.

Other actions apply to individual AI systems and specific uses.

The same tool can create very different risks depending on how you use it. For example, an internal chatbot has different risks from AI used in hiring or customer decisions.

Review each use on its own and consider how systems could be reused, combined or misused over time.

Keep clear records of governance decisions, testing, incidents and monitoring. Good records support audits, reviews and ongoing improvement.

For help with key words, refer to the terms and definitions.

1. Decide who is accountable

AI systems can make automated decisions that significantly impact people, communities and businesses. Overall, your organisation is ultimately accountable for how and where AI is used, AI complexity can create gaps where no one takes clear responsibility for outcomes.

Accountability is the first step to using AI responsibly.

1.1 Accountable people

Understanding your role in the supply chain and identifying clear roles for how AI is governed, developed and deployed in the organisation supports accountability and effective oversight.

1.1.1 To ensure AI systems perform as required and obligations are met, assign, document and clearly communicate who is accountable across the organisation (including contractors and third-party providers/systems) for the operation of the AI management system, including:

- safe and responsible policies, practices and procedures
- the development and deployment of every AI system, including ongoing human control and oversight
- oversight of the development and use of AI systems by third parties
- testing of AI systems across the organisation
- oversight of concerns, challenges and requests for redress
- the performance and continual improvement of the AI management system.

1.1.2 For each accountable person, define and communicate the required competencies and their authority. Ensure they are staffed with appropriately skilled people and have the necessary resources.

1.2 Supply chain accountabilities

Understanding your role in the AI supply chain and identifying which parties are responsible for maintaining the performance, safety and integrity of AI systems throughout their lifecycle is key to effective accountability.

1.2.1 Identify, document and communicate accountability for shared responsibility across the AI supply chain (model developers, system developers and system deployers) for:

- monitoring and evaluation of model and system performance, quality and safety
- human oversight and intervention
- processes to raise issues, faults, failures incidents, contested outcomes, issue resolution and system updates.

1.2.2 Clearly document and communicate the accountability and obligations that developers have towards downstream organisations when integrating, customising, enhancing developer provided AI models or systems. This includes transparency of AI model and system risks, expected behaviours, outcomes under expected use cases and changes to the model or system, paying particular attention to any specific contractual obligations, which could vary by customer (DEV).

1.3 AI literacy and training

Delivering effective training in AI across the organisation can build confidence, support AI adoption and ensure accountable people have the right capabilities to perform their roles.

1.3.1 Evaluate and document the training needed to build broad AI understanding and a culture of accountability across the organisation. Source or deliver training to bridge any identified gaps. Regularly check skills are up-to-date as AI development and deployment evolves.

1.3.2. Evaluate the training needs of accountable people and provide appropriate up-to-date training to address gaps, such as those responsible for:

- meeting legal and regulatory obligations
- handling personally identifiable information
- operation, control, intervention or termination of each AI system
- oversight and monitoring of each AI system
- procurement, development or deployment of third-party AI systems
- safe and responsible development of AI systems (DEV).

1.4 AI governance framework

Implementing policies, processes and an overall management system for the development and deployment of AI across the organisation is fundamental to effective and responsible governance of AI.

1.4.1 Document and communicate:

- the organisation's strategic intent to develop and deploy AI systems in line with organisational strategy and values
- the regulations relevant to the development and deployment of AI systems and how the organisation will comply
- appropriately detailed policies, processes and goals for the safe and responsible development and deployment of AI systems which align to the strategy, including:
  - an end-to-end process for AI system design and development (DEV)
  - goals for AI systems to meet organisational policies for the safe and responsible use of AI
  - the consequences for people who act outside of the organisation's policies and defined risk appetite.

1.4.2 Ensure effective operation of the AI management system by:

- documenting and implementing a process to proactively identify deficiencies in the AI management system. This includes instances of non-compliance in AI systems or in their development or deployment, documenting root causes, corrective action and revisions to the AI management system.
- appropriately planning changes to the AI management system
- identifying and documenting the internal and external factors (such as infrastructure or the deployment context) that may affect the organisation's ability to meet its responsibilities through the overarching AI management system
- providing sufficient resources such as human effort and compute to deploy AI systems safely and responsibly over the lifecycle.

1.4.3 Monitor compliance with organisational policies to identify and address any gaps between leadership expectations and staff understanding of how to develop and deploy AI safely and responsibly.

2. Understand impacts and plan accordingly

Because AI systems can operate at speed and scale, their potential impacts are often magnified. Without careful planning, a single AI system can lead to widespread negative outcomes, such as unfair decisions or the provision of inaccurate information.

For example, AI systems can learn from and amplify existing issues such as unwanted bias in data. This can lead to unfair decisions or inappropriate generated content that could affect many people. If an AI system used for shortlisting in hiring has a bias problem, it could unfairly reject hundreds of qualified candidates before anyone notices.

To use AI responsibly, organisations need to understand, plan for and monitor potential impacts of AI systems. Those affected should be able to raise complaints and get help.

2.1 Identify and engage stakeholders

Engaging potentially impacted stakeholders is an important way to identify and understand the impacts of AI systems.

2.1.1 Identify and document key types of stakeholders (such as employees and end users) that may be impacted by the organisation's development and deployment of AI, and their needs.

2.1.2 Prioritise, select and document which stakeholder needs will be addressed in organisational policies and procedures.

2.1.3 Document and communicate the organisation's commitment to preventing harms to people from AI models and systems and upholding diversity, inclusion and fairness.

2.1.4 Document the scope for each AI system, including intended use cases, foreseeable misuse, capabilities, limitations and expected context.

2.1.5 For each AI system, engage stakeholders to identify and document the potential benefits and harms to different types of stakeholders, including:

- impacts to vulnerable groups
- risks of unwanted bias or discriminatory outputs
- use of an individual's personal information
- where the system makes or influences a decision about a person or group of people.

2.1.6 For every documented risk of harm to affected stakeholders, conduct appropriate stakeholder impact analysis.

2.1.7 Monitor for potential harms by engaging affected stakeholders for each AI system on an ongoing basis to identify new stakeholders, including end users throughout the AI lifecycle.

2.1.8 Create processes to support ongoing engagement with stakeholders about their experience of AI systems. Identify vulnerable groups and support appropriately. Equip stakeholders with the skills and tools necessary to give meaningful feedback.

2.2 Establish feedback and redress processes

Establishing processes for people affected by AI systems to give feedback, ask questions, and challenge decisions easily and safely can ensure issues are identified and resolved.

2.2.1 Create, document and communicate a process for:

- Potentially affected stakeholders to raise concerns, challenges, or requests for remediation and receive responses (for example, a human rights grievance and remediation mechanism). This includes when and how frequently to communicate, the level of detail to provide and the communication needs of stakeholders, considering the level of AI knowledge and any regulatory requirements.
- Evaluation of contestability requirements of both internal and external stakeholders and interested parties including accessibility needs.

2.2.2 Implement and document system-level mechanisms to enable contestability of AI use and decisions, enabling stakeholders to understand, challenge and appeal AI use and decisions. These mechanisms must be accessible, understandable and available to users at the appropriate time during interaction with an AI system. Consider mechanisms to share information regarding end user contests and any redress with deployers of AI systems and models (DEV).

2.2.3 Implement and document mechanisms to enable deployers to escalate feedback, report unexpected behaviours, performance concerns or realised harms and support improvements to models or systems (DEV).

2.2.4 Ensure people who monitor and review affected stakeholder feedback can trigger recourse and redress processes where there is an obligation to do so.

2.3 Monitor for systemic issues

Ongoing monitoring of stakeholder feedback and redress processes can ensure that systemic issues are identified and addressed.

2.3.1 Monitor and evaluate contestability and redress processes to identify and address systemic risks as well as improve effectiveness of these processes.

2.3.2 Create, document and communicate a process to review and evaluate stakeholder contests of AI system use across the organisation including any concerns raised by affected stakeholders and requests for information.

3. Measure and manage risks: implement AI-specific risk management

AI risks fundamentally change depending on the type and complexity of your AI systems. Risks often emerge from how the AI system behaves in different situations and use-cases, rather than only from software updates. They can rapidly amplify smaller issues into significant problems.

For example, an AI chatbot that answers simple questions during business hours, when it can be monitored by a staff member, is a low-risk use of AI. The risks expand, however, if that chatbot operates 24/7, without human oversight, and answers more complex questions.

To use AI responsibly, organisations need to be able to identify and manage its risks.

3.1 Establish a fit-for-purpose risk management framework

An effective risk management framework supports organisations to identify and manage the risks of using AI, set clear rules about what risks are acceptable, and regularly check how AI systems are working over the lifecycle.

3.1.1 Create and document:

- a risk management framework that addresses the specific characteristics and risks of AI systems
- organisational-level risk tolerance and criteria to determine acceptable / unacceptable risks for the development and deployment of AI systems. This should include the significance and likelihood of potential harms to affected stakeholders in line with the AI policy and objectives.
- AI impact assessment, risk assessment and risk treatment processes, including criteria for reassessment over the lifecycle of an AI system. Identify and document any specific use cases or qualities of AI systems that represent an unacceptable risk to stakeholders or the organisation, in line with the organisation's risk tolerance.

3.1.2 Ensure that risk management processes include steps to identify, assess and treat risks arising from other parties in the AI supply chain, such as third-party developers and third

party deployers. Specific risks relating to open-source AI models, systems and components should be considered by both providers and consumers of these technologies. For more on cybersecurity supply chain risks, refer to Artificial Intelligence and machine learning: Supply chain risks and mitigations.

3.1.3 Adopt or develop clear and consistent reporting formats, such as data sheets, model cards, or system cards, to communicate appropriate risk management outcomes, including residual risks, to relevant stakeholders (DEV).

3.2 Assess AI system risks

Using proportionate, robust methods to assess AI system risks is a key part of the operation of the risk management framework.

3.2.1 Establish a triage system to determine which AI systems may pose an enhanced or unacceptable risk, aligned to the organisation's context and risk tolerance.

3.2.2 Perform and document a risk assessment and evaluation for the specific requirements, characteristics and documented use cases of each AI system, including systems developed or procured from third party suppliers.

3.2.3 In undertaking AI system risk assessments, take the following steps to evaluate the likelihood and consequence of each risk as well as the consequence of not deploying the AI system:

- Identify potential severity and likelihood of harms to stakeholders, drawing on the Stakeholder Impact Assessment (see 2.1.1 – 2.1.6).
- Identify legal, commercial and reputational risks such as failing to meet legal obligations, organisational commitments to ESG, diversity, inclusion and accessibility or programs supporting diversity, equity and fairness.
- Consider the potential amplified and emerging data governance risks across each phase of the AI system lifecycle including before and after model training.
- Analyse risks systemically using risk models to identify the sources and pathways through which AI systems could produce the identified risks.
- Compare the estimated value or level of identified risks to pre-determined organisational risk criteria (see 3.1.1) or those defined by regulatory bodies or stakeholders.
- Document any specific use cases or qualities that represent an unacceptable level of risk to stakeholders or the organisation.

- Communicate risk assessments in clear reporting formats to relevant stakeholders.

3.3 Implement controls for AI system risks

Where risks are identified, risk treatment plans make it clear how risks will be mitigated.

3.3.1 Create, document and implement a risk treatment plan to prioritise, select and implement treatment options (e.g. risk avoidance, transfer, acceptance, reduction) and controls to mitigate identified risks. Reassess risks after controls are implemented to verify their effectiveness.

3.3.2 Communicate risk treatment plans in clear reporting formats to relevant stakeholders.

3.3.3 Create and document a deployment plan which includes the response, recovery and communications for the realization of residual risks.

3.3.4 Research, document and implement leading practices in safety measures as safeguards, as appropriate for identified risks (DEV).

3.4 Monitor and report incidents

Reporting incidents when they happen and communicating the steps you've taken is essential to build trust with stakeholders and meet regulatory obligations.

3.4.1 Track, document and report relevant information about serious incidents and possible corrective measures to relevant regulators and/or the public in a reasonable timeframe. Reporting near-misses and corrective measures is good practice. Communication of corrective measures should consider privacy and cybersecurity risks.

3.4.2 Create and document a process to evaluate and fulfil reporting and disclosure obligations such as those under the Online Safety Act relevant to AI systems usage, including documentation of safety measures implemented such as notices and incident reporting.

3.4.3 Conform to and document data breach reporting requirements and liabilities from related standards. For example, under the Notifiable Data Breaches scheme of the Office of the Australian Information Commissioner.

3.4.4 Maintain two-way communication between developers and deployers for incident reporting, sharing performance insights and coordinating responses to identified issues.

3.4.5 Monitor and evaluate risk assessments and treatment plans on a regular, periodic basis or when a significant change to the use case or the system occurs, or new risks are identified. This includes responding to impact assessments or insufficient risk treatment plans.

3.4.6 Monitor and evaluate the overall effectiveness of risk management processes and continually improve them.

4. Share essential information

People should know when they're interacting with AI and understand when AI decisions affect them. For example, when a customer is receiving services and guidance from a chatbot, they should know this is not a human specialist.

To use AI responsibly, organisations need to tell users and stakeholders when and how they're interacting with AI.

4.1 Maintain an AI register

An AI register is a central place that records important details about all of the AI systems across the organisation.

4.1.1 Create and maintain an up-to-date, organisation-wide inventory of each AI model and system, with sufficient detail to inform key stakeholders and support future conformance assessments, including:

- accountable people
- purpose and business goals
- capabilities and limitations of the AI model and/or system

- origin, fine-tuning and updates where applicable
- technical requirements and components
- datasets and their provenance used for training and testing
- acceptance criteria and test results
- any impact and risk assessments and outcomes
- identified risks, potential impacts and the risk treatment plan
- any system audit requirements and outcomes
- dates of review.

See the AI register template for additional guidance.

4.2 AI system transparency and explainability

Clearly communicating when and how AI is being developed or deployed by the organisation and explaining the impacts to users and stakeholders is important to build accountability and trust.

4.2.1 Create, document, implement and communicate a policy and process for how to transparently communicate:

- to AI users, and affected stakeholders that engage directly with AI systems, AI-enabled decisions or AI-generated content about when and how they will be informed about AI development, deployment and use in the organisation.
- the capabilities, limitations and potential risks of the AI systems that users and affected stakeholders may engage with. This should include when and how frequently to communicate, the level of detail and the level of AI knowledge of AI users and affected stakeholders. It should also address communication obligations and the accessibility needs of AI users and affected stakeholders and incorporate feedback mechanisms where appropriate.

4.2.2 For each AI system, evaluate and document:

- the transparency requirements for each user and affected stakeholder group (see 2.1.1 and 2.2.1)
- the transparency and explainability system requirements and measures – including for third-party -provided systems – dependent on use case, stakeholder requirements and risks arising from disclosure.
- how accessibility obligations and commitments are met by implementing human-centred design.

4.2.3 Create, document and implement organisational processes and transparency mechanisms proportionate to the risks arising from the diverse, evolving, and alternative uses of GPAI beyond predefined applications, including their potential for unexpected and hard-to-explain behaviours. (GPAI DEV/DEP).

4.2.4 Wherever possible choose more interpretable and explainable systems.

4.2.5 Wherever possible, provide reasonably interpretable and explainable AI systems and models to accountable people within the organisation or downstream deployers to enable them to meet their own regulatory obligations (DEV).

4.2.6 Conduct internal testing of the AI model and/or system's capabilities and limitations. Clearly communicate results to deployers prior to deployment. (DEV).

4.3 Supply chain transparency

Developers and deployers need to work together to share information and build mechanisms that can clearly communicate information about AI systems to all parties.

4.3.1 Document or request from upstream providers the technical details of the system or model that may be required to meet the needs of users within the organisation or stakeholders.

4.3.2 Share as much of the following information as possible about AI models and systems with downstream deployers (while protecting commercially sensitive information and meeting legal compliance) (DEV):

- Technical details such as model architecture, description of data, components and their characteristics
- Test methods, use cases and testing resulting
- Known limitations, risks and mitigations (such as potential bias and corrective actions) and external audit findings
- Data management processes for training and testing data including data quality, meta data, and provenance
- Privacy and cybersecurity practices including conformance to standards and best practice
- Transparency mechanisms implemented for AI-generated content, interactions and decisions

- Document and share the following key information for GPAI systems with downstream organisations, stakeholders, researchers and regulators (GPAI DEV):
  - Training data sources and compliance details with relevant privacy, intellectual property and copyright laws
  - Model cards and system cards, including risk assessment results particularly evaluation of dangerous and emerging capabilities in the deployment and scaffolding context of tool access and agent design.
- Restricted and managed access to model weights and other associated artefacts.

4.3.3 Share as much of the following information as possible about AI systems with upstream developers (while protecting commercially sensitive information and meeting privacy obligations) (DEP) (See also incident reporting 3.4.1 and 3.4.4).

- Issues, faults, failures, incidents and any other observed risks that can be addressed by developers
- Any unexpected and unwanted bias resulting from use of the system.
- Ensure you've included the required information in contracts with suppliers of systems, including when to update information.

4.4 AI-generated content transparency

Being clear about when and how content is AI-generated or modified is important to build trust in digital content with stakeholders.

4.4.1 Implement fit-for-purpose and proportionate transparency mechanisms for AI generated content as set out in the *Be clear about AI-generated content* guidance.

5. Test and monitor

AI systems can change their behaviour over time or act in ways that are less predictable than conventional software. For example, an AI system that worked well last month might start giving different answers today if it is trained on additional data.

To use AI safely, organisations should test and monitor their AI systems.

5.1 Pre-deployment testing

Conducting testing before an AI system is deployed and documenting the outcomes supports ongoing risk mitigation.

5.1.1 Establish oversight mechanisms to review and approve testing methodologies and results, and to monitor system performance, user feedback and operational impacts post-deployment.

5.1.2 Define and communicate clear acceptance criteria and test methodologies that reflect the intended use, context and potential risks (as identified in Essential Practice 3). Conduct pre-deployment testing.

5.1.3 Clearly document tests and outcomes to support external audits and oversight

5.1.4 When testing is conducted by upstream providers during the development of the AI system and model, request test methodologies and results from the provider and ensure its alignment with your acceptance criteria.

5.1.5 Obtain documented deployment authorisation and rationale from the accountable person for the AI system based on test results.

5.2 Monitor system performance

Setting performance metrics, closely monitoring and reviewing the performance of AI systems ensures that they operate as intended.

5.2.1 Establish monitoring systems for each AI system to track key performance metrics and indicators relevant to the identified risks.

5.2.2 Implement a deployment process for AI systems that maps business targets to system performance metrics for both internal and third-party developed systems.

5.2.3 Establish and document response processes for addressing all foreseeable issues and harms during system operation.

5.2.4 Establish regular system performance review cycles with stakeholders and subject matter experts to evaluate testing criteria, effectiveness and outcomes.

5.2.5 For each AI system, create and document monitoring requirements including human oversight prior to deployment and evaluate as part of continuous improvement cycle.

5.3 Conduct additional testing proportionate to risk

Determine whether AI systems require further safety evaluations, independent testing or auditing which are proportionate to their risks.

5.3.1 Conduct safety evaluations that scale with model capabilities (GPAI DEV), for example: assessment for cyber-offensive capabilities and vulnerabilities; testing for potential chemical, biological, radiological and nuclear information risks; evaluation of model behaviours beyond intended use cases; testing for jailbreaking or prompt manipulation; data privacy risks; or comprehensive red teaming to identify vulnerabilities.

5.3.2 For use cases requiring enhanced practices and GPAI systems, conduct independent (internal or external) and thorough review of testing and evaluation methodologies and results. Document and report any issues to the accountable person for the system.

5.3.3 Create a process for and determine whether an AI system requires regular auditing, appropriate to the level of risk identified by its risk assessment. Conduct audits when required.

5.4 Implement robust data and cybersecurity measures

Effective data governance, privacy and cybersecurity practices are fundamental to supporting the responsible operation of AI systems.

5.4.1 Implement and evaluate the effectiveness of policies and procedures in addressing AI-specific risks and adapt as necessary:

- Data governance processes covering the use of data with AI models and systems. This includes the management of data usage rights for AI including intellectual property (including copyright), Indigenous Data Sovereignty, privacy, confidentiality and contractual rights.
- Privacy policies covering the collection, use and disclosure of personal or sensitive information by AI models and systems, including for model training purposes. This needs to support teams to comply with the Australian Privacy Principles for all AI systems.
- Cybersecurity processes to cover the emerging and amplified risks of AI systems interaction with existing systems and data, such as AI systems unintentionally exposing sensitive information or bypassing security controls. This includes application of the Essential Eight Maturity Model for cybersecurity risks to AI systems. This includes the application of AI cybersecurity advice available on **cyber.gov.au** to address unique risks associated with the development and deployment of AI systems.

5.4.2 For each AI use case:

- Define and document the data quality, data/model provenance and data preparation requirements.
- Understand and document the data sources and collection processes each AI model / system relies on to function including personal and sensitive data. Put in place systems to manage the data and document the data used to train and test each AI model or systems and data used for inference.
- Define and document processes for protecting AI models and systems to address emerging cybersecurity and privacy risks (DEV).
- Where appropriate, report to relevant stakeholders on data, model and system provenance.
- Document how the Australian Privacy Principles have been applied including in models and systems developed by third parties. Refer to the OAIC Guidance on privacy and developing and training generative AI models and Guidance on privacy and the use of commercially available AI products.
- Document data usage rights including intellectual property (including copyright), indigenous data sovereignty privacy confidentiality and contractual rights.
- Monitor for and detect any leakage of personal and sensitive information from AI models and systems.

6. Maintain human control

Unlike traditional software that follows explicit instructions, AI systems learn patterns from data and make their own opaque decision logic. This means they need human oversight to

make sure they operate safely. For example, while regular software does exactly what you program it to do, AI might interpret your instructions differently than you intended.

To responsibly use AI, organisations need to make sure a human appropriately oversees any AI systems in use. The person overseeing your AI systems should know how to do so appropriately, and what they need to do to override the system if something goes wrong.

6.1 Maintain human oversight and control

Ensuring that people in the organisation retain oversight of AI systems, with the ability to intervene where necessary is important to the safe ongoing operation of the system.

6.1.1 Maintain operational accountability, capability and human oversight throughout the lifecycle of AI systems.

6.1.2 Implement mechanisms to enable human control and intervention during the operation of the AI system (DEV).

6.1.3 Implement mechanisms to enable human oversight and intervention to address systemic risks and emerging capabilities such as capability evaluation, training, pause, independent oversight, dynamic guardrails and tool/system access controls (DEV).

6.1.4 Ensure appropriate training is provided to anyone overseeing or using AI systems to understand each system's capabilities, limitations and failure modes and when human intervention is needed.

6.2 Decommission when appropriate

Establishing processes to decommission AI systems when they are no longer needed or performing as intended can protect ongoing service delivery and data assets.

6.2.1 Define and determine:

- the criteria or reasons that termination of an AI model or system might need to occur, and at what point intervention should take place
- the most appropriate role or person to oversee the intervention and decommissioning process
- whether the model or system is essential to any critical infrastructure or service delivery
    - Assess the risks and impacts of shutting down the AI model or system, including impacts on end-users and interdependencies with other integrated systems on which the model, data or outputs rely to function.
    - Develop a timeframe and treatment plan to minimise impacts or disruption caused by the decommissioning process.
    - Determine a method to extract data and for the return or deletion of assets. Establish which information should be preserved for record keeping purposes.

6.2.2 Create a process for how your organisation will inform relevant parties (such as employees, customers, and upstream or downstream parties) within a reasonable timefram of the retirement or shutdown of an AI model or system. Establish a channel for people to raise concerns, request support and receive responses.

6.2.3 Determine whether alternative systems or processes will need to be provided to address any issues or gaps.

6.2.4 Maintain alternative pathways for critical functions so operations can continue if AI systems malfunction and/or are taken offline.

Download

Guidance-for-AI-adoption-implementation-guidance_0_0_0.pdf

PDF 4.89 MB · Published 05 May 2026

Next up

Be clear about AI-generated content

Related pages

Guidance for AI adoption: foundations

Start with practical governance for low-risk and early artificial intelligence use

Learn more →

Planning tools and templates

Access resources that help you manage and use artificial intelligence in your work

Learn more →

Activities for teams

Explore practical activities to build confidence in using artificial intelligence

Learn more →

Subscribe to our newsletter

The National AI Centre acknowledges organisations whose research and expertise inform the guidance and resources on this platform, including the Safe AI Adoption Model (SAAM) AI Adopt Centre, CSIRO's Data61 and the Human Technology Institute, University of Technology Sydney.