



Australian Government



National
AI Centre

Guidance for AI adoption: foundations

National Artificial Intelligence Centre

October 2025

Contents

Six essential practices for responsible AI governance	2
Introduction	3
Helping all organisations embrace AI	3
Setting up for success: getting AI governance right	3
How to use the guidance	4
Important considerations	4
Essential practices	5
1. Decide who is accountable	5
2. Understand impacts and plan accordingly	6
3. Measure and manage risks	7
4. Share essential information	8
5. Test and monitor	9
6. Maintain human control	10

Six essential practices for responsible AI governance

The *Guidance for AI Adoption* sets out 6 essential practices for responsible AI governance and adoption.

We based this guidance on national and international ethics principles. It is the first update to Australia's Voluntary AI Safety Standard (VAISS), launched in 2024. In this update, published in October 2025, we have:

- condensed 10 guardrails into 6 essential practices
- removed redundant language
- expanded our audience to developers as well as deployers.

There are 2 versions of the guidance. This version, Foundations, outlines the essential practices every organisation needs for good AI governance. Foundations is for businesses and organisations in the early stages of adopting AI. When you're ready for detailed step-by-step instruction on how to implement these practices, refer to the other version of the guidance, [Implementation guidance](#).

Introduction

Artificial intelligence (AI) brings enormous opportunity and potential to Australian business. It's already reshaping the way organisations operate – and if it hasn't yet reached yours, it's likely touching others in your sector.

In Australia, many organisations are moving ahead with their AI plans. Software engineers are using AI to help them write code. Team leaders are boosting employee engagement and speeding up customer interactions. Businesses are automating data entry and, in some cases, using AI tools in their boardrooms.

Helping all organisations embrace AI

For those who get AI right early, advantages abound. AI tools, systems and agents can deliver a competitive edge across everything from marketing and procurement to stakeholder engagement and customer service. Larger operators may be moving faster on AI adoption, but smaller organisations can share in its benefits too.

But there are some practical considerations. Organisations are keen to use systems they and their end users know and trust, and with good reason. Once trust is broken, whether with customers, stakeholders or the public, it is difficult to regain.

That's why AI systems need to be implemented with responsible practices.

If you are using AI in your organisation, you need to put some essential governance practices in place. This will help you safely harness the movement that is reshaping industries and organisations.

By following this guidance, your organisation could improve customer engagement, strengthen employee trust, reimagine products, and increase productivity.

Setting up for success: getting AI governance right

Good AI governance is a way for organisations to adopt AI with confidence. It puts in place the systems, processes and tools to identify and manage risks, while taking advantage of AI's upsides.

Responsible AI practices build in strategic oversight, robust frameworks and transparent, explainable systems.

If you're starting small, responsible AI practices will be simple to implement. They can – and should – grow as your use of AI systems matures.

How to use the guidance

Start where you are and build confidence as your AI use matures.

You don't need to implement everything at once. Get started across all six essential practices to establish basic responsible AI governance. Add more actions as your organisation's AI use grows or your governance capabilities mature.

The six practices are designed for flexibility. Adopt them in ways that fit your organisation's specific context, constraints and use cases.

Important considerations

As you read and adopt the six practices, remember that AI governance happens at 2 levels.

- Some practices apply to your whole organisation (like having an AI policy).
- Some practices apply to each specific AI system in the context that you use it (like performing tests).

The same AI system can present different risks depending on how you use it. For example, using ChatGPT to draft marketing emails is different from using it to assess job applications. Each use case needs its own governance checks.

When an AI system could be used in multiple ways, consider all possible use cases, including ways it might be misused.

Remember that documentation is key. You should document every activity in these essential practices that you carry out. This will let you audit and review your governance when you need to. Good documentation also helps your organisation learn and improve its AI governance over time.

The practices in this document are written for organisational leaders and governance people – they identify the essential practices you need to ensure happen across your organisation and supply chain.

This guidance does not replace your other essential governance frameworks such as those for data, privacy, and cybersecurity. These frameworks should already be in place and reviewed and updated for your use of AI systems. The 6 practices focus on governance matters specific to AI and its unique characteristics.

For clarification on specific terms used in this document, you can refer to the [terms and definitions](#).

Essential practices

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.

Getting started

- 1.1 Assign a senior leader as the overall AI governance owner. They should have enough authority and understanding of AI capabilities and risks to oversee all AI use in your organisation.
- 1.2 Create an AI policy that sets out how your organisation will use AI responsibly. Ensure that the policy provides advice to your staff on AI risks and how to manage them. Take a look at our AI policy template to get started.

Next steps

- 1.3 Make a specific person accountable for every AI system your organisation uses. Make sure the accountable people are familiar with the technology and understand its business implications.
- 1.4 Train your accountable people so they can make informed decisions about AI's risks and behaviours.
- 1.5 Clarify supply chain accountabilities. When AI systems involve multiple parties (vendors, developers, integrators etc), make it clear who is responsible for each part of the AI supply chain. This will let you know who to talk to if something in the system goes wrong.
- 1.6 Turn your AI policy into a governance framework. A comprehensive framework with clear policies and procedures can help your organisation address AI-specific challenges.

For further guidance, read [Implementation practice 1](#).

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.

Getting started

- 2.1 Carry out a stakeholder impact assessment. Identify the groups and types of people your AI systems may affect. Pay particular attention to vulnerable or marginalised cohorts. Assess all potential impacts, such as unfair decisions, providing inaccurate or harmful information, encouraging overreliance, etc.
- 2.2 Create contestability channels. Set up channels for people to report problems, challenge or question AI decisions. The strength of those channels should match how serious the impact is. Make sure you can take action to set things right if your AI system has negatively affected someone.

Next steps

- 2.3 Engage your stakeholders early and continue engaging them throughout the AI lifecycle, especially your staff and customers. Understand their needs, concerns and how they might be affected. Use this knowledge to inform how you design, test and deploy your systems.
- 2.4 Identify systemic unwanted impacts by monitoring patterns of feedback. Use this information to proactively fix systemic issues in your systems.

For further guidance, read [Implementation practice 2](#).

3. Measure and manage risks

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.

Getting started

- 3.1 Create a risk screening process to identify and flag AI systems and use cases that pose unacceptable risk or require additional governance attention. See our AI screening tool for help with this process.

Next steps

- 3.2 Set up risk management processes that account for the differences between traditional IT, narrow AI, general purpose AI and agentic AI systems.
- 3.3 Conduct risk assessments and create mitigation plans for each specific use case and identified impacts in that context.
- 3.4 Apply risk controls based on the level of risk for each of your specific uses of AI.
- 3.5 Create processes to investigate, document and analyse AI-related incidents. Make sure you use lessons learned to prevent incidents happening again and to improve AI systems and risk management processes.

For further guidance, read [Implementation practice 3](#).

4. Share essential information

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

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 information from a chatbot, they should know this is not a human specialist.

Getting started

- 4.1 Create and maintain an AI register. This should document all your AI systems, how you use them, and any other important information about them. Your documentation should include any AI system, whether you developed it internally or procured it from elsewhere. It should also cover AI that might be embedded in other systems, like human resources and customer engagement tools. See our AI register template for help with this process.
- 4.2 Disclose your use of AI. Make it standard practice in your organisation to clearly communicate your use of AI to your stakeholders. This is especially important in situations where AI makes or influences decisions, generates content that can meaningfully impact people, or can be mistaken for a human being.

Next steps

- 4.3 Identify, document and communicate AI system capabilities and limitations. Explaining these capabilities and limitations to relevant stakeholders can prevent their overreliance on or misuse of AI.
- 4.4 Be transparent across your AI supply chain. This will let anyone who interacts with the system understand what components they are using, their capabilities and limitations, and key characteristics like training data sources.
- 4.5 Set up ways to explain AI outcomes, especially when they affect people. The detail in your explanation should match how serious the outcome is, and it should be easy for people who are affected to understand.

For further guidance, read [Implementation practice 4](#).

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.

Getting started

- 5.1 Ask for proof. If you're buying an AI system, ask the developer or supplier to show proof that it's been properly tested.
- 5.2 Test before you deploy a system. Consider how you want to use your AI system. Decide how you might test the system to ensure it is performing how you want it to, and conduct them.
- 5.3 Monitor your system after you deploy it. Set up a monitoring process that helps you detect changes in performance and behaviour. Match your monitoring approach to the risk levels identified in your risk assessment. Watch for new risks that were not present when you first deployed the system.
- 5.4 Extend your data governance and cybersecurity practices to your AI systems. Strengthen these practices to address AI-specific needs, like protecting AI models that learned from sensitive data.

Next steps

- 5.5 Stress-test your AI system to spot issues or vulnerabilities before others do. Make sure safety, security and policy controls are sound, even when someone deliberately tries to break or bypass them.
- 5.6 Consider getting independent testing. If your AI use case was identified needing additional governance attention, consider getting a third party to test your AI systems. This testing should happen before you deploy the system, after significant changes and on a regular basis.

For further guidance, read [Implementation practice 5](#).

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.

Getting started

- 6.1 Ensure meaningful human oversight. Make sure a person oversees your AI system in a way that matches how much autonomy the system has, and how high the stakes are. This could mean automated monitoring for low-stakes applications, and mandatory human review for high-stakes decisions.
- 6.2 Build in human override points. Make sure you have clear intervention points where humans can pause, override, roll back or shut down AI systems if needed.

Next steps

- 6.3 Provide training to people overseeing AI systems. Make sure anyone using or overseeing your AI can understand each system's capabilities, limitations and failure points, and when to intervene.
- 6.4 Maintain alternative pathways. Make sure that all of your organisation's critical functions can continue even if your AI systems malfunction or are being retired.

For further guidance, read [Implementation practice 6](#).