Responsible AI in Action

Bite-sized guidance series

Procure with purpose

This issue covers:

- things to look out for when engaging AI vendors
- key points to include in contract

Read this alongside other issues in the bite-sized AI guidance series

Implementing AI is not just a technology investment. It affects decision-making, trust, privacy, and accountability **(see the *AI in Regulation* issue)**. Getting procurement right helps ensure AI systems are safe, explainable, and fit for a regulatory setting.

When engaging vendors, be clear about the regulatory problem you are solving, what decisions the tool will support, and what guardrails must remain in place. Vendors should be able to demonstrate how their tool manages data, reduces bias, supports transparency, and can be monitored over time .

Practical steps for engaging AI vendors.

1. Start with the basics

Follow the NZ Government Procurement Rules and procurement principles of integrity, transparency, and accountability. Before procuring new tools, check whether suitable All-of-Government AI capabilities are already available through the Marketplace or other shared public service platforms.

2. Apply public service AI guidance

Apply relevant government guidance to test whether the tool is responsible and suitable for public sector use, including fairness, explainability, and privacy expectations.

3. Do vendor due diligence

With the support of your digital services and procurement teams, ask for evidence of real deployments in similar contexts. Vendors may not be able to share everything, as access to a model's internal workings is rarely granted, but the focus should be on what the system actually produces in practice and whether that is good enough for your context.

As a starting point, confirm:

- data storage, access arrangements, and data sovereignty, including where data is held and who can access it
- intellectual property and information management settings
- how the model is tested, updated, and monitored, including what testing and benchmarking data has been used
- the provenance of the data the model was trained on, where this is relevant to your use case
- what contractual commitments the vendor will make regarding bias levels and ongoing performance.

The level of scrutiny should reflect the risk. For lower-risk applications, standard due diligence may be sufficient. Where AI will

inform decisions that affect people's rights or entitlements, a higher standard of interpretability should be required. This means the vendor should be able to explain, in terms your decision-makers can understand, why the system produces the outputs it does **(see the *Opportunities and risks* issue)**. If a vendor cannot provide adequate assurance for a higher-risk application, that is important information. It may mean the product is not suitable for that use case, even if it works well in other settings.

4. Identify and manage key risks early

Common risks include bias in model outputs, security vulnerabilities, and vendor lock-in. However, the risks associated with any particular AI system will depend heavily on the regulatory context it is being used in. A generic risk checklist is unlikely to be sufficient.

Procurement and digital teams should work closely with regulatory staff to develop a full risk profile for each system under consideration. Regulatory staff understand the decisions the AI will be informing, the people those decisions affect, and what good and poor outcomes look like in practice. That operational knowledge is essential for identifying risks that may not be visible from a purely technical or commercial perspective **(see the *Empowering your people* issue)**. Once risks have been identified, require vendors to explain how they will be mitigated and reviewed over time. Risk management should not end at procurement.

5. Lock in accountability through the contract

Include clear clauses on:

- performance measures and audit rights
- security standards
- ongoing monitoring and reporting
- lifecycle management, including updates and decommissioning
- certainty on where data is stored
- data portability and access to outputs generated using agency data
- knowledge transfer if the vendor relationship ends
- exit and transition plans to avoid lock-in
- cost controls, including protections against significant price increases over the contract term.

That last point warrants specific attention. The pricing of AI models, particularly those at the frontier of capability, can change considerably over time. An AI system that is affordable at the point of procurement may become significantly more expensive as vendor costs change. Building price certainty into contracts, or retaining the ability to exit if costs become unsustainable, is prudent risk management.

These arrangements support the Digital Government Target State and whole-of-government digital capability initiatives led by the Government Chief Digital Officer.

To get more practical steps for how regulatory leaders can lead AI innovation with confidence, check out the full guidance: **Responsible AI in Action**.