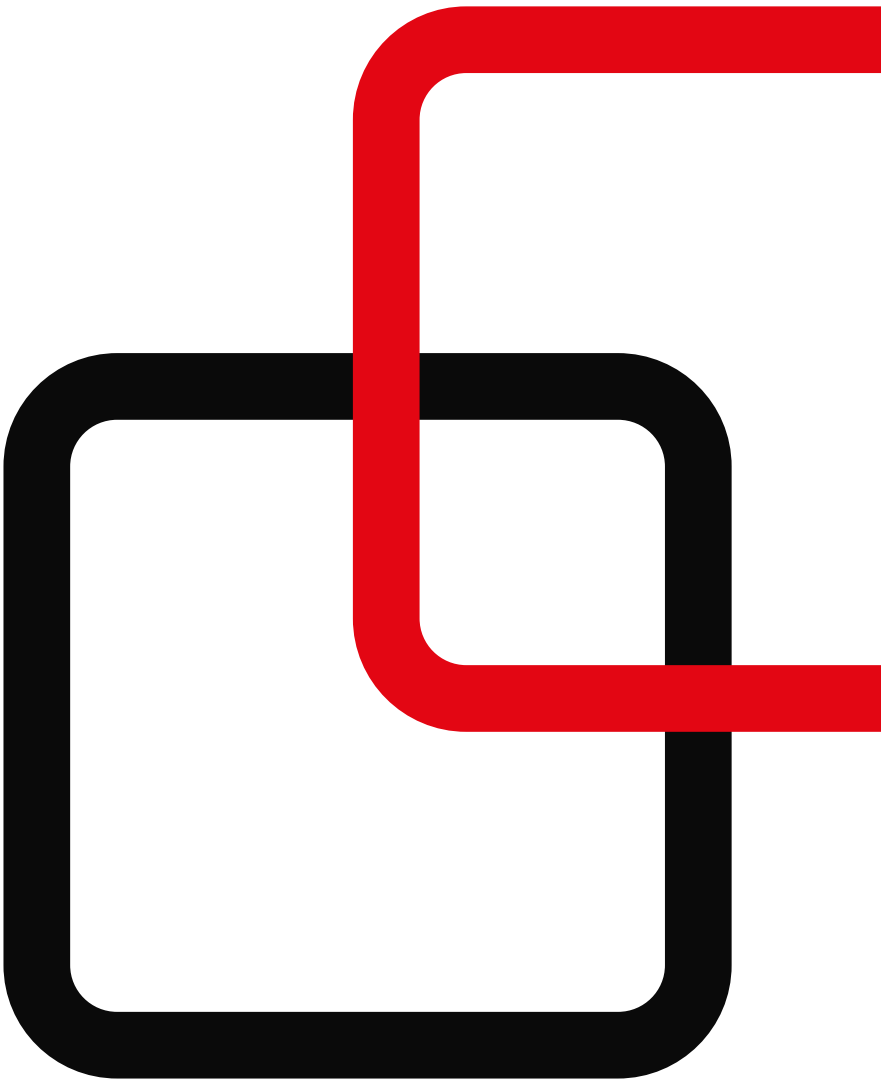
RadMis



A PAPER FOR GOVERNMENT

Governed AI for government: *from policy to practice.*

For accountable officials, and the governance, risk and records teams who make AI policy work day to day.

SECURE INNOVATION NOVEMBER 2026 RADMIS.COM

THE SHORT VERSION

Policy asks five questions. The system doing the work should answer them.

1 Most of what an agency needs from AI is automation of repeatable work, not an autonomous agent.

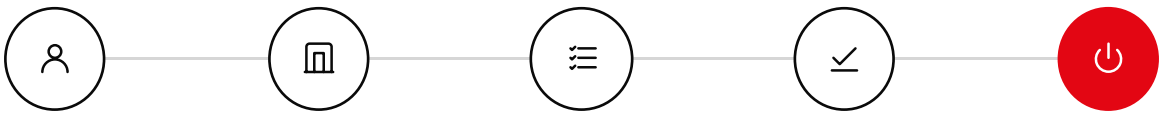
2 Policy asks you to know what AI you run, who owns it, what it did, who checked it and how to stop it.

3 Those answers should come out of the system itself, not a spreadsheet kept beside it.

4 A gifts and hospitality register shows the whole approach in four minutes.

WHAT POLICY NOW ASKS OF YOU

Government AI policy has moved from principles to obligations: register your use of AI, assess its impact, name an accountable official and publish a transparency statement. The difficulty is rarely the policy. It is that the obligations arrive with no operational tooling behind them.



01
Who owns it?
The AI Register: an owner, purpose, risk band and review date for every AI system.

02
What may it do?
AI is an account with an access level, on the same access tree as your people.

03
What did it do?
A hash-chained audit trail, and every AI turn kept with its sources.

04
Who checked?
Approval steps that wait for a named person. The decision is recorded.

05
How do we stop it?
One AI Kill Switch. Each use is recorded and mailed.

AGENT OR AUTOMATION?

Not everything needs an agent.

There is a common misconception about when an agency really needs a large language model. Registers, forms, approvals, reports and routing are most of what agencies ask for, and none of them needs one at run time.



Illustrative. The split varies by organisation.

MOST OF THE WORK

If it repeats, automate it.

A language model builds the automation quickly. Once it is built, it runs without one.

THE EDGE CASES

If it needs judgement, reason.

Keep the model for the complex cases your current systems cannot cover.

ALWAYS

Guardrails, so it is compliant.

Access levels, human approval and an audit trail on both. The compliance report comes out of the same record.

WHY THIS MATTERS IN GOVERNMENT

Smaller assessments

Automation that runs without a model has far less to assess.

Lower risk

The model is kept to the cases that need it, inside guardrails.

One record

The compliance report comes from the same trail as the work.

Automate what repeats. Reason where it is needed. *Prove all of it.*

MONITORING AI POLICY

What you have to do, and how it is done.

Know what AI is in use	The AI Register holds every AI system. A check finds any that is not registered, registers it and raises a ticket.
Assess risk before go-live	Risk ratings with approval roles, and a go-live gate that waits for the approval and a test run.
Review on a schedule	A review ticket by risk band, with its checklist and sample. Test sets are run and recorded.
Report and be transparent	Monthly reports on AI activity, refusals and incidents, a dated transparency statement, and the evidence for the yearly attestation.
Disclose AI-written content	Mail and comments an AI wrote are marked. A workflow cannot send AI-written mail without a person's approval.
Handle incidents and appeals	Thirteen incident types, each with its clock. “Report this answer” on chat screens. Appeals decided by another officer.
Meet privacy obligations	A Data Inventory, personal data requests on a 30-day clock, retention with Records Holds, and a Destruction Log.

The evidence is produced as the work happens. Nobody assembles it afterwards.

A SIMPLE THING, DONE PROPERLY

The gifts register. From an email to working screens in four minutes.

The story is made up. The files in it are real: on 5 November 2026 they passed the platform's own check.

THU 4.10 PM	An auditor wants the gifts register by Monday. The head of governance emails Eva, an AI co-worker, from a phone.
4.11 PM	Eva asks two questions: who should see what, and is the decision one of three?
4.13 PM	Eva writes three short text files: a form, a list and a report. They are not code. The platform builds the screens.
4.14 PM	**Three working, access-controlled screens. Opened to staff by changing one number.**
FRI 9 AM	The auditor asks whether any giver was a government official. One more line adds the tick box.
MONDAY	The auditor gets the export. The screens are locked against change.

NOBODY ASKED FOR WHAT CAME WITH IT

Access and logging

Sign-in, the access check and the request log, from the first request.

Privacy

The table joins the Data Inventory. Personal columns are marked.

Retention

Old rows removed on schedule, with a Destruction Log entry.

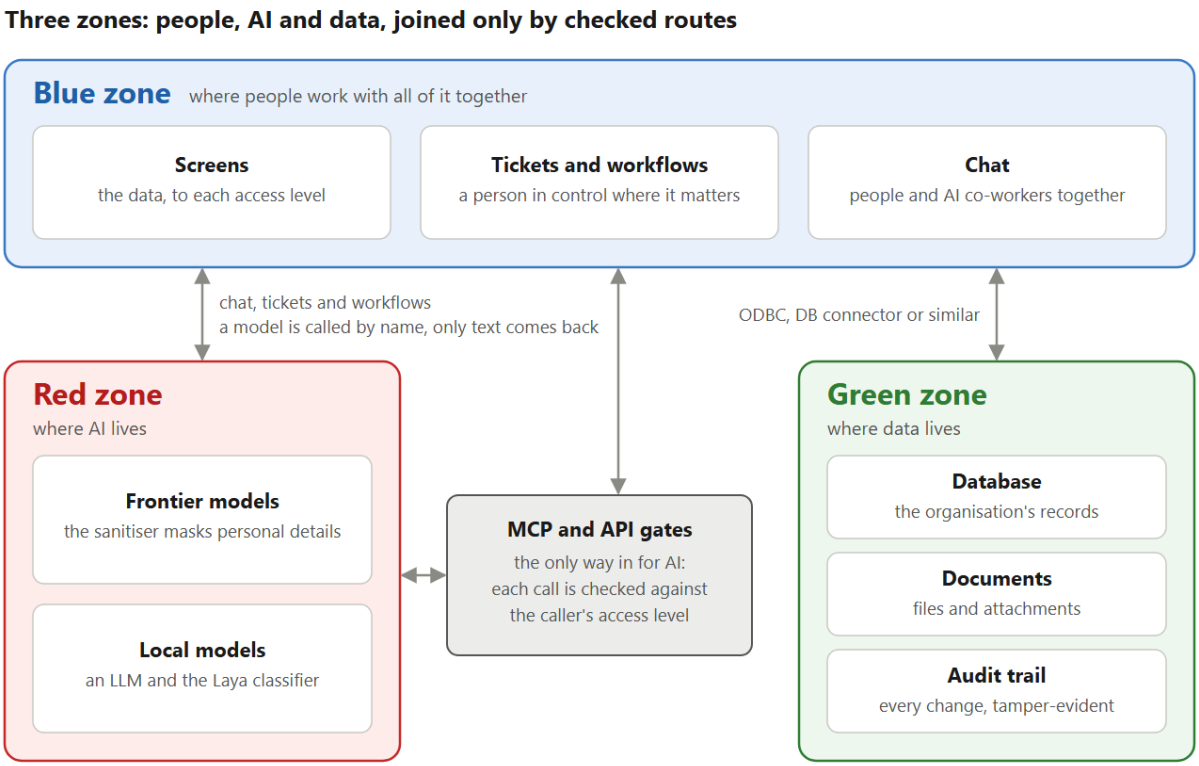
Audit

One tick, and every change is in the hash chain.

The same register can be built with no AI at all, on the platform's builder screens. Conflicts of interest, breaches and complaints registers follow the same pattern.

SECURITY

Why it is safe to run in an agency.



Most secure: each zone on its own server or VM instance, on the premises, in the cloud or a combination.
A smaller system can run all three on one server.

AI stays outside

Models run in their own zone and never touch the data directly.

AI is a user

Two gates, an access level, registered calls only, each one logged.

AI never writes code

It describes a screen in plain text. The platform builds it.

Personal data is masked

Local model first. The sanitiser is set per AI co-worker today.

Sign-in is defended

Two-factor sign-in, attack blocking, tokens stored only as hashes.

Your premises or your cloud

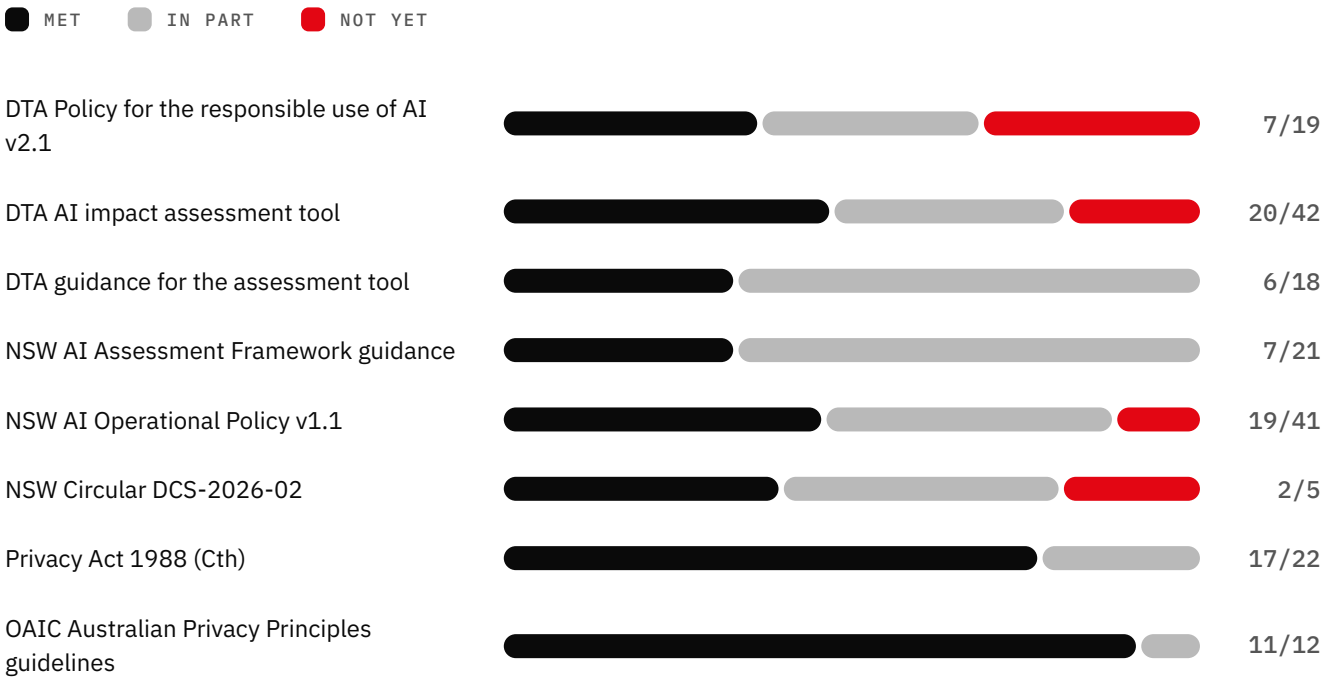
Or both. The same rules hold between the zones.

The sanitiser masks personal details before text goes to an outside model. Today it is set per AI co-worker; making it mandatory on every outside model is next on the roadmap.

COVERAGE

How far, against government frameworks.

The platform keeps its own compliance map. Points only your agency can meet, such as a decision of your executive, are counted apart and not shown.



Mapped, not certified, against the code on 28 October 2026. The full map, point by point, is available on request.

WHAT IS STILL OPEN

- The sanitiser required on every outside model
- One stop per AI system, with a restart path and drill log
- Ask before a write
- Explanation on request

We would rather you heard the gaps from us.

RadMis

BEFORE YOU BUY FROM ANYONE

Five questions to put to any AI vendor.

1. Can our auditor verify the audit log without your help?
2. If an agent is tricked by a prompt, what can it write, delete or read?
3. When AI builds a screen or a workflow, is any code generated?
4. Where does personal data go before it reaches the model?
5. How long before the first governed use case is running?

Our answers are on radmis.com. We would be glad to build your gifts register in front of you.

EMAIL
info@radmis.com

WEB
www.radmis.com

OFFICE
Sydney, Australia