RadMis

PRODUCT PAPER

Anyone can build with a frontier LLM. Few can make it *secure and pass an audit.*

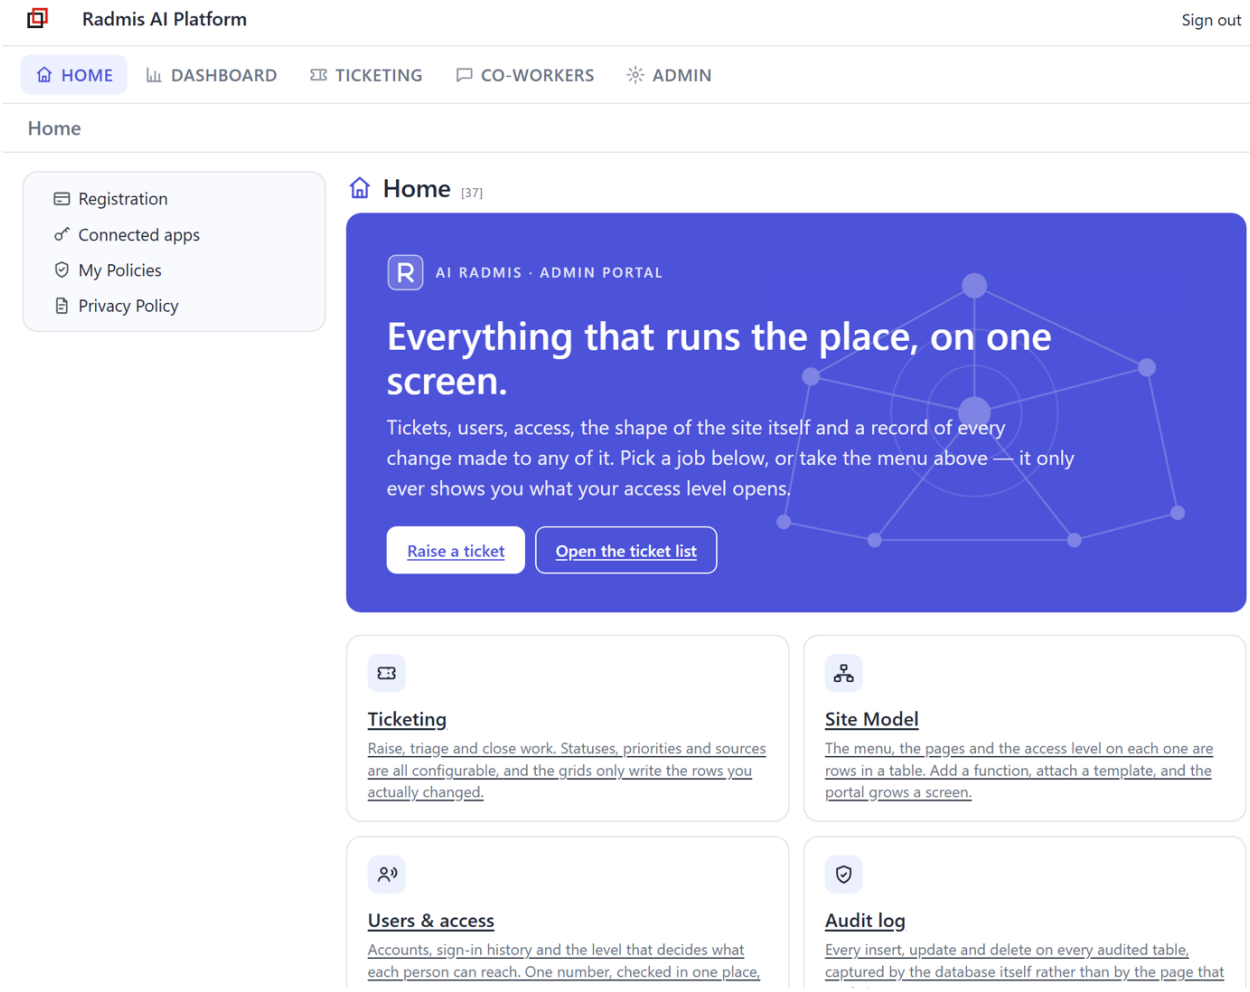
Radmis is the governed workplace where people and AI co-workers do regulated work together: every action permissioned, a person in control wherever it matters, every audited change provable.

SAY YES TO AI, WITH THE EVIDENCE TO PROVE IT

Most AI pilots work. Then they stop at the risk review.

The system has no owner on record and no register entry. Nobody can show what it did or who checked it, and there is no way to switch it off.

Radmis is built for that review: a complete business platform where the controls that let AI act are the same controls that produce the evidence.



THE PLATFORM'S HOME SCREEN

Minutes

to build a governed screen

About an hour

to install, on your premises or your cloud

Thirteen

regulations and guidelines mapped

One switch

stops the AI

AGENT OR AUTOMATION?

Most of what you need is automation. Not an agent.



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.

THE FIVE QUESTIONS EVERY AI SYSTEM MUST ANSWER

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.

THE WORKED EXAMPLE

From an email to a working register 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.

4.10 PM	An auditor wants the gifts register by Monday. A compliance head emails Eva, an AI co-worker, from a train.
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. Nothing an AI wrote is run as code.
4.14 PM	Three working, access-controlled screens.

THE WHOLE OF THE FORM EVA WROTE

```
@sketch 1
@record
@table gifts

The gift:
'Received by' fullname*: text | 'Email' email: text | 'Department' department [select: finance, operations, sales, board]:
code
'From' giver*: text | 'Date received' received* [date]: date | 'Value' amount*: money
'What it was' description* [textarea]: long

The decision:
'Manager' manager: text | 'Decision' decision [select: pending, kept, returned, donated]: code
'Why' reason [textarea]: long
```

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

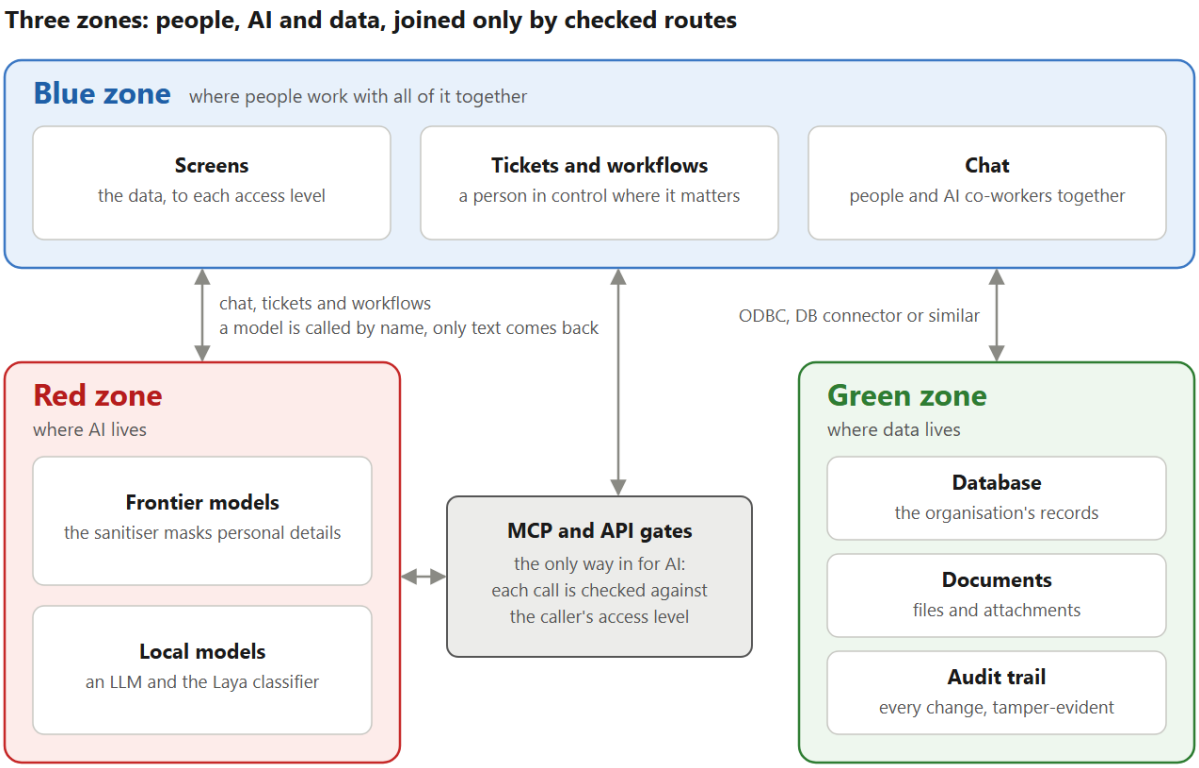
Old rows removed on schedule, with a Destruction Log entry.

Audit

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

HOW IT STAYS SECURE

AI stays outside. Your data stays in.



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.

Today the sanitiser is set per AI co-worker; making it mandatory on every outside model is the next step.

WHAT YOU CAN DO WITH IT

Ten jobs it takes off the risk register.

01
Stalled AI pilots, past the risk committee

Register, risk approval, test sets, go-live gate, Kill Switch.

02
Spreadsheets become governed registers

Described in an email. Access, privacy, retention and audit built in.

03
Frontier AI on sensitive data

Local model first. The sanitiser masks personal data.

04
AI-drafted letters with sign-off

No AI-written mail leaves without an approval.

05
Personal data requests

One search, then export, erase or restrict, on a 30-day clock.

06
APRA CPS 230

Risk profile, controls, provider register and test tickets.

07
Evidence on demand

A trail an auditor verifies without the platform.

08
Incidents and appeals

Thirteen incident types, each with its clock.

09
Policy and knowledge assistant

Answers only from approved stores, kept with sources.

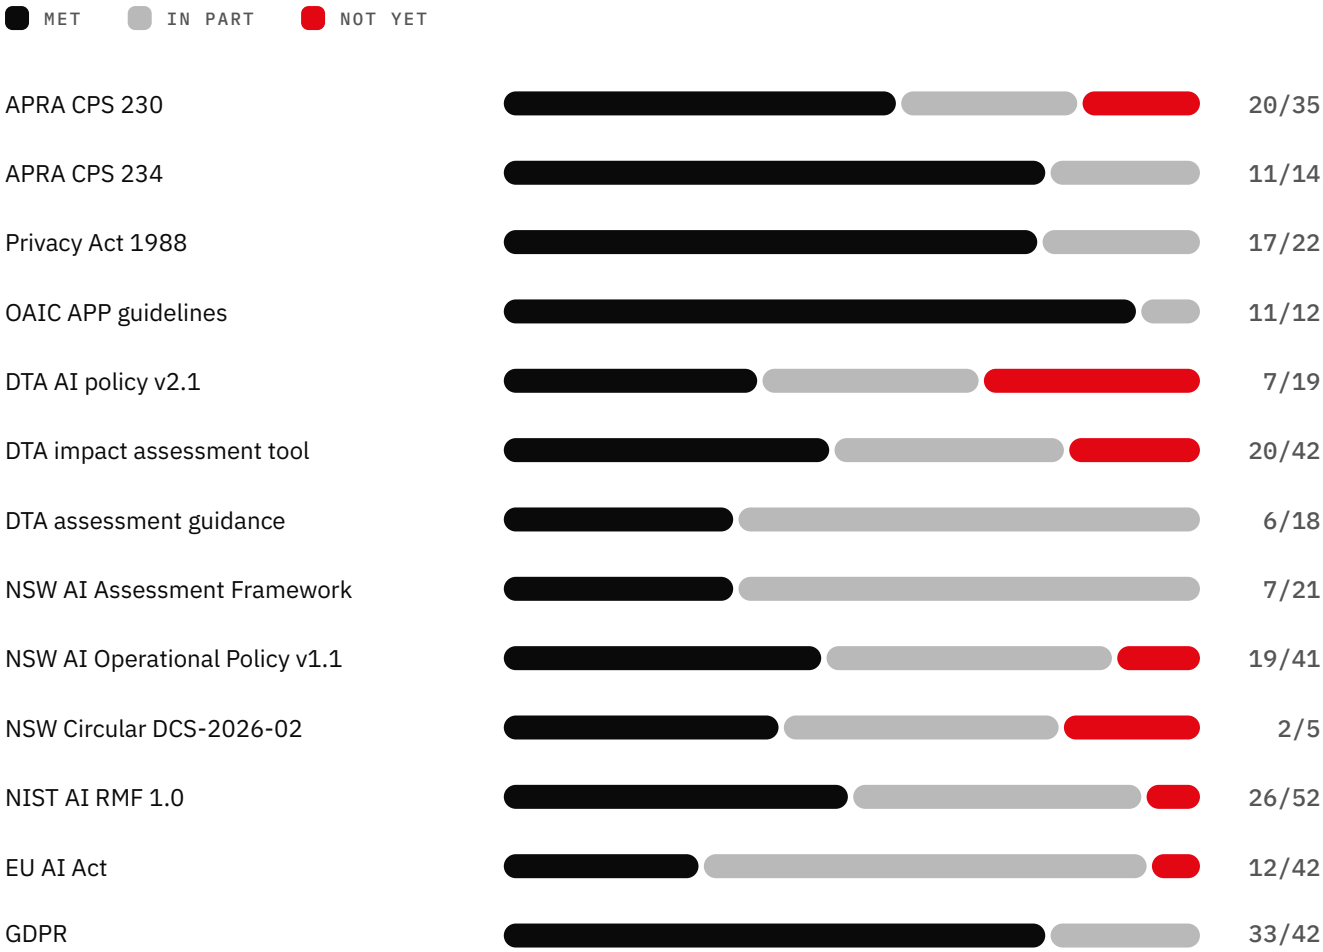
10
Mail triage on your servers

A local classifier sorts, scores and flags.

COVERAGE

Mapped, not certified.

Thirteen regulations and guidelines, point by point, against the code on 28 October 2026. Points only your organisation can meet are counted apart.



WHAT IS STILL OPEN

- The sanitiser required on every outside model
- One stop per AI system
- Ask before a write
- Explanation on request
- Co-worker checks in Platform Health
- An AI mark on an account
- A six-month retention floor

26 years
of research and development

100+
deployments

5 sectors
health, education, telco, payments, government

Sydney
wholly Australian-owned

RadMis

BEFORE YOU BUY FROM ANYONE

Five questions to ask any AI vendor. Including us.

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. The full compliance map and the technical appendices are available on request.

EMAIL
info@radmis.com

WEB
www.radmis.com

OFFICE
Sydney, Australia