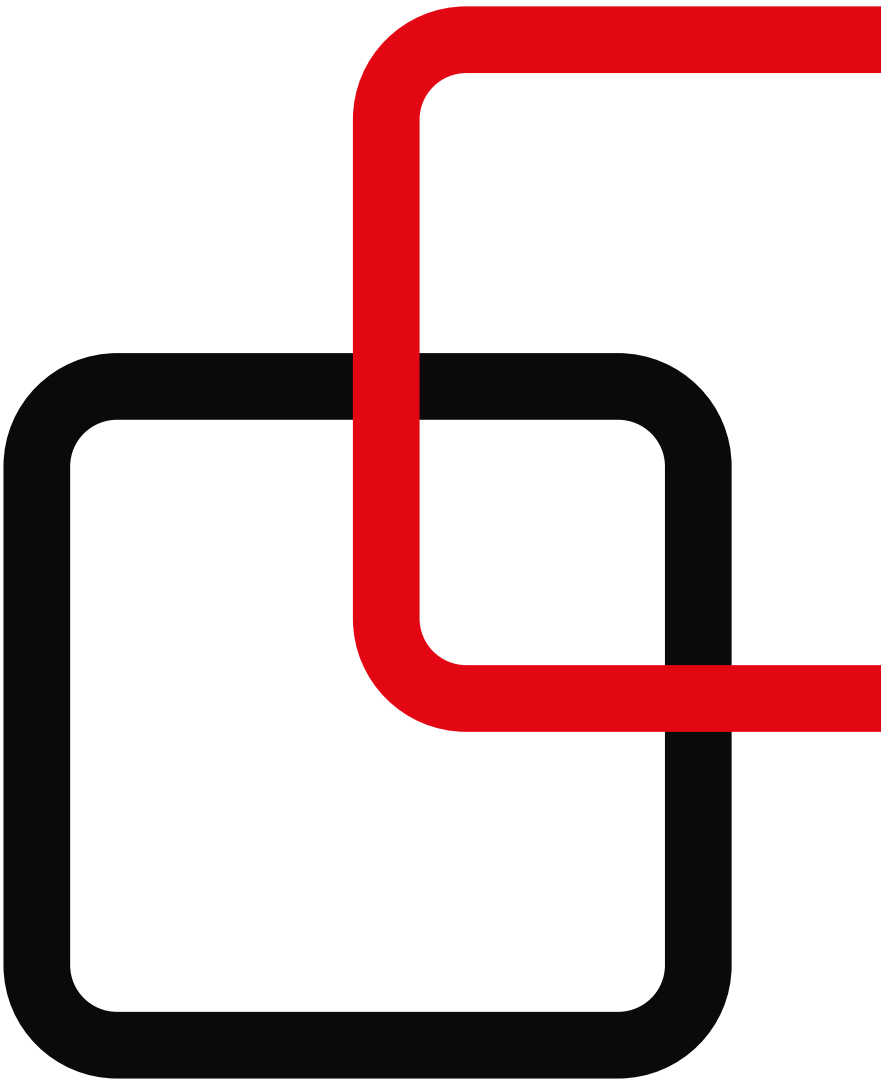
RadMis



TECHNICAL PAPER

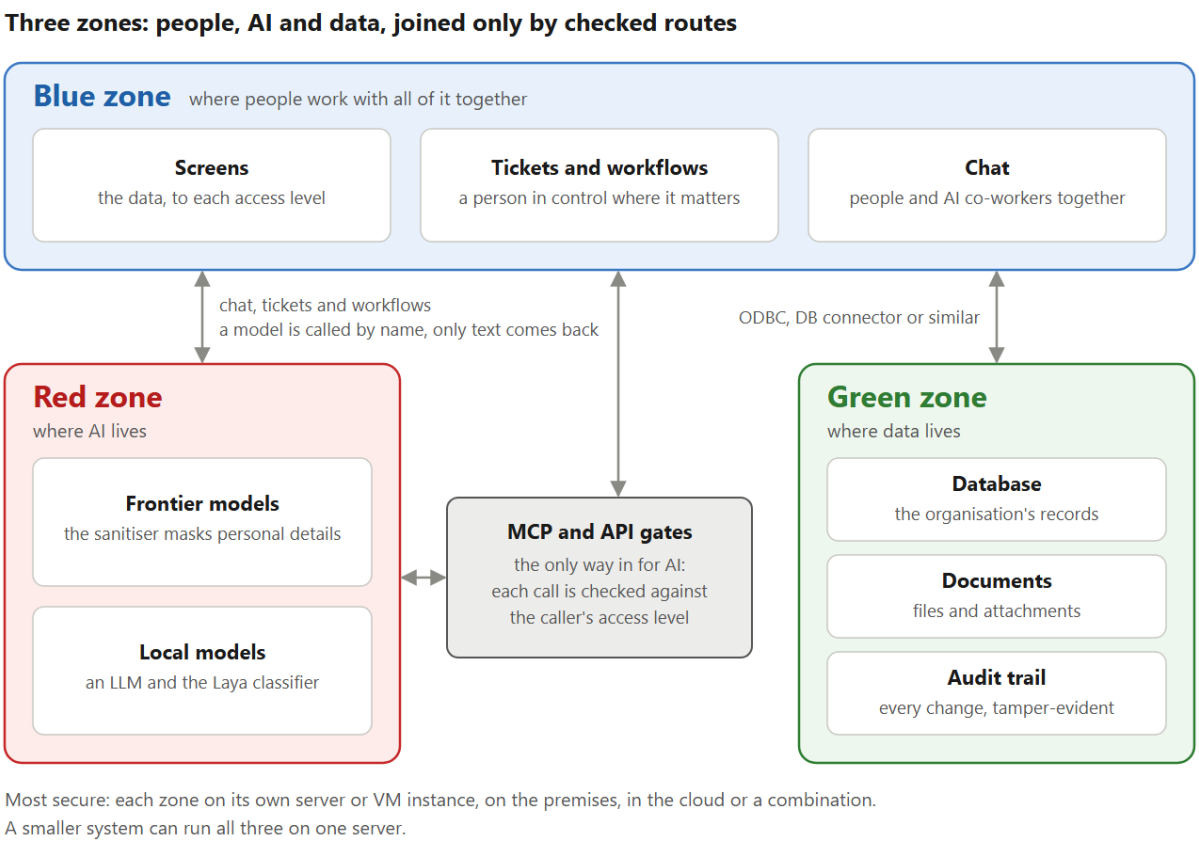
Radmis AI Platform: what is in it, *and how far it goes.*

For architects, security reviewers and auditors. The architecture, the full feature list, the ISO mapping and coverage regulation by regulation.

SECURE INNOVATION NOVEMBER 2026 RADMIS.COM

ARCHITECTURE

Three zones, joined only by checked routes.



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	Frontier and local models run outside the platform itself. When the platform calls a model by name, only text comes back.
AI is a user	A co-worker comes in through the MCP and API gates with an account and an access level. A call above its level does not exist for it. Only registered calls run, and each is logged.
AI never writes code	To build a screen, AI writes a Sketch file. A column's type is one word from a short list. A file cannot carry SQL, and can never rename, retype or drop a column.
The site model	Every function is a row in one tree. The menu and the access check are cut from it, so a link cannot be hidden while its page still answers.

The most secure arrangement puts each zone on its own server or VM. A smaller system can run all three on one, with the same rules between the zones.

One platform for apps, AI co-workers, workflows and tickets.

Foundation

The site model: one tree of every function, which drives the menu and access. One entry point for every request, guarded by access level. Accessibility assessed against AS EN 301 549.

Security and access

Access levels 0 to 13 for people and AI alike. Two-factor sign-in with recovery codes, sign-in blocks and an attack log. OAuth 2.1 and connected apps. The Kill Switch.

Auditing

An audit trail with a SHA-256 hash chain and Verify. Settings Changes records every settings save: who, when, the old value and the new.

Builders

Grid, Record and Report Builders, and Builder Chat, where screens are made by describing them in plain English.

Co-workers

Agent Setup for persona, memory and model. Chat screens. Guardrails and the sanitiser. Test Sets. Local and external models, each behind a label.

Knowledge and RAG

RAG Stores ingest documents by drag and drop or in bulk. Vector databases: Chroma, Postgres with pgvector, or SQL Server. Wikis, each with its own chat.

Classifier

Laya, a local classifier that answers choice, score or yes-or-no questions, wired into named units.

Workflows

Workflows with human approval: a workflow waits at a set step until a person decides.

Governance and privacy are functions, not documents.

Ticketing

Tickets with statuses, priorities, sources and attachments. The ticket is where a workflow's approval is given, by a comment or a reply email.

Mail

The platform's own mail server, on a TLS certificate. Incoming mail becomes tickets or comments. Each co-worker has a mailbox.

MCP and API

The MCP gate and the API gate, with registered calls module by module and the AI Call Log.

Operations

A scheduler for tasks, jobs and watched folders. Platform Health runs automatic checks and notifies when anything looks suspicious.

AI governance

AI Register and Models register, risk approval with a go-live gate, scheduled reviews, incidents and appeals, transparency reports, Risk Profile and Controls, Providers and Training.

Privacy

Data Inventory, personal data requests, retention with Records Holds and a Destruction Log, policies and user consents.

Every new screen inherits sign-in, access control, logging, privacy classification and retention from its first request.

THE ISO STANDARDS, BY THE FIVE QUESTIONS

Each question, what answers it, and the clauses it maps to.

QUESTION	DONE IN THE PLATFORM	ISO/IEC 42001:2023	ISO/IEC 27001:2022
1. Who owns it, and for what?	The AI Register entry, its owner and its risk approval	5.3, A.2.2, A.3.2, A.9.4	A.5.1, A.5.2, A.5.9
2. What is it allowed to do?	Access levels, the two gates, guardrails and the sanitiser	A.6.2, A.7, A.9.2	A.5.15, A.5.18, A.8.3, A.8.12, A.8.22
3. What did it do?	The hash-chained trail, Settings Changes, the AI Call Log and each conversation's record	9.1, A.6.2.8, A.8	A.8.15, A.8.16, A.8.17
4. Who checked?	Ticket approvals, workflow approval steps and the approval before AI mail	6.1.4, A.6.2.6, A.9.3	A.5.3, A.8.32
5. How do we stop it?	The AI Kill Switch: gates, models and classifier, each flip recorded and mailed	10.2, A.6.2.6	A.5.24 to A.5.30

Every reference was checked against the text of the standard on 3 October 2026. This is a mapping, not a certification.

THE WORDS WE USE

Co-worker	An AI with its own account, access level and owner.
Gate	One of the two controlled entrances, MCP and API, through which a co-worker reaches the platform.
Sketch	The small plain-text language in which a screen is described, by a person or by an AI, and which the platform builds from.
Sanitiser	Masks personal details before text goes to an outside model.
Site model	The one tree of every function on the platform, which decides the menu and who can reach what.

HOW FAR, REGULATION BY REGULATION

Mapped, not certified.

Every point of thirteen regulations and guidelines, set against what is built, on the code of 28 October 2026. Points the platform cannot meet alone, such as a board's decisions, are counted apart as organisational.

REGULATION OR GUIDELINE	PLATFORM POINTS	MET	IN PART	NOT YET	ORGANISATIONAL
APRA CPS 230 Operational Risk Management	35	20	9	6	26
APRA CPS 234 Information Security	14	11	3	0	11
Privacy Act 1988 (Cth)	22	17	5	0	10
OAIC Australian Privacy Principles guidelines	12	11	1	0	9
DTA Policy for the responsible use of AI in government v2.1	19	7	6	6	14
DTA AI impact assessment tool	42	20	14	8	14
DTA guidance for the AI impact assessment tool	18	6	12	0	1
NSW AI Assessment Framework guidance	21	7	14	0	12
NSW AI Operational Policy v1.1	41	19	17	5	26
NSW Circular DCS-2026-02	5	2	2	1	2
NIST AI Risk Management Framework 1.0	52	26	22	4	20
EU AI Act, as amended by the AI Omnibus	42	12	27	3	15
GDPR	42	33	9	0	33

STILL OPEN

- The sanitiser required on every outside model
- Explanation on request, and a rating on each answer
- One stop per AI system, a restart path and a drill log
- Ask before a write
- Co-worker checks in Platform Health
- An AI mark on an account
- A six-month floor on retention for high-risk systems
- Other regulators' fields on the AI Register

The full map, point by point, is available on request.

QUESTIONS TO ASK ANY AI VENDOR

And our answers.

1	**Can our auditor verify the audit log without your help?**	Yes. The trail is a SHA-256 hash chain written by the database itself, so a change made outside the screens is recorded too, and the chain shows if the trail has been altered.
2	**If an agent is tricked by a prompt, what can it write, delete or read?**	Only what its access level opens, through calls you have registered. The builder language has no delete command and returns no data.
3	**When AI builds a screen or a workflow, is any code generated?**	No. The AI describes and the platform builds.
4	**Where does personal data go before it reaches the model?**	To a local model first. Before an outside model, through the sanitiser where it is set.
5	**How long before the first governed use case is running?**	The platform installs in about an hour, and a first register can be on screen the same day.

Minutes

to build a governed screen

About an hour

to install

Thirteen

regulations and guidelines mapped

One switch

stops the AI

BEFORE YOU BUY FROM ANYONE

The same five questions, for your own review.

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?

The full compliance map, point by point, is available on request.

EMAIL
info@radmis.com

WEB
www.radmis.com

OFFICE
Sydney, Australia