

Enterprise AI, under your control.

Most enterprise AI stalls between pilot and production. Security reviews, data rules and unclear ownership stop it. ResonAi builds the control layer first: a gateway, a company brain, guardrails and governance. Then we put agents to work on top of it.

WHAT WE BUILD

ResonAi Gateway. One door for every model. A key and budget per team, a Canada-only lock where needed, an audit log without content, and cost per workflow. Swap models without touching your apps.

Company Brain. Your documents, project history and lessons learned, answered with sources. Permission walls are enforced on the server for every question.

Guardrails. Prompt-injection defense, least-privilege tool access, and one firm rule: agents draft, people send.

Agent workflows. PMO status packs, failure analysis, proposals and intake triage. Each agent has one job, a named owner and before-and-after measures.

Governance. AI intake, scoring and gates through your EPMO. Acceptable use, data rules, human-approval thresholds and a one-page CIO view.

Enablement. A five-level training path with a CPMAI®-aligned specialist track, so teams use what is built.

SECURITY & PRIVACY POSTURE

- ✓ Content an agent reads is data, never instructions
- ✓ Roles checked on the server and tested with forged-role attacks
- ✓ No prompt or reply text in gateway logs (tested)
- ✓ Hard budgets and rate limits on every key
- ✓ Commercial model terms: providers don't train on your data
- ✓ Retention agreed in writing before work starts
- ✓ Model changes only after a quality test on the task

DEPLOYMENT OPTIONS

ResonAi Cloud	Multi-tenant, each client isolated, Canadian-region storage
Dedicated	Single-tenant stack in a Canadian region
Inside your tenant	Azure, AWS or Microsoft 365. Data never passes through us
On-premise	Open-weight models on your hardware

Residency, plainly: storing data in Canada is easy. Canada-only model processing narrows the model choice today. We say which workflows need it.

HOW AN ENGAGEMENT RUNS

01

Assess

AI inventory incl. shadow tools, risk and data map, ranked workflows, written security position.

02

Pilot

One governed workflow in production, behind the gateway, measured before and after.

03

Scale

Same controls across more teams, the company brain, and an agent portfolio run by your EPMO.

04

Run

Retained AI operations: monitoring, cost reviews, tested model changes, quarterly review.

PROOF

30-50 h → 4-8 h

Cross-project failure analysis per cycle, Canadian insurer's PMO

CA\$0.21

Measured model cost per run of a 5-agent estimating workflow, tracked by the gateway

In production

PMO status and escalation agent that logs a written reason for every call