

INTELLIGENCE // CONNECTIVITY // INNOVATION

Your AI. Your data. Your premises.

A complete, private AI workforce in a box — installed in your office, yours to command, governed by Australian law. Everyone else rents you intelligence over the wire and keeps your data. We hand you a sovereign appliance — models, agents, every Claude skill you use today, and an open-source software estate — that never lets your data leave the building, and proves it.

On-premises · single-tenant

Local-first · opt-in burst

One-button workflows

Memory that outlives turnover

[Book a sovereignty briefing →](#)



Four compromises no enterprise should have to **swallow**.

Every mid-market business knows it needs AI. Almost none can adopt it without accepting four trade-offs — and the more sensitive your data, the worse they bite.

1

It never stops costing

Per-seat and per-token pricing means the bill grows forever, in lockstep with how useful the AI becomes. Success is punished.

2

Your data leaves the building

Prompts, documents and records are shipped to foreign clouds you don't control, retained on terms you didn't write — and often used to train the next model.

3

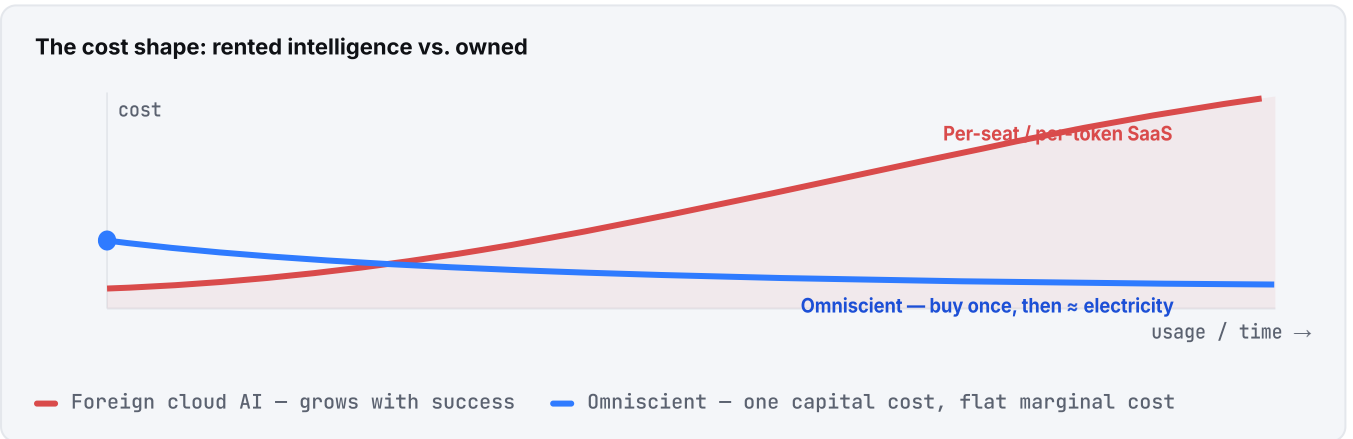
The guardrails aren't yours

Foreign models, tuned by foreign safety teams, refuse legitimate specialist work — medical, legal, mining, defence, agronomy — because a generic policy doesn't know your industry.

4

Nobody answers to Australian law

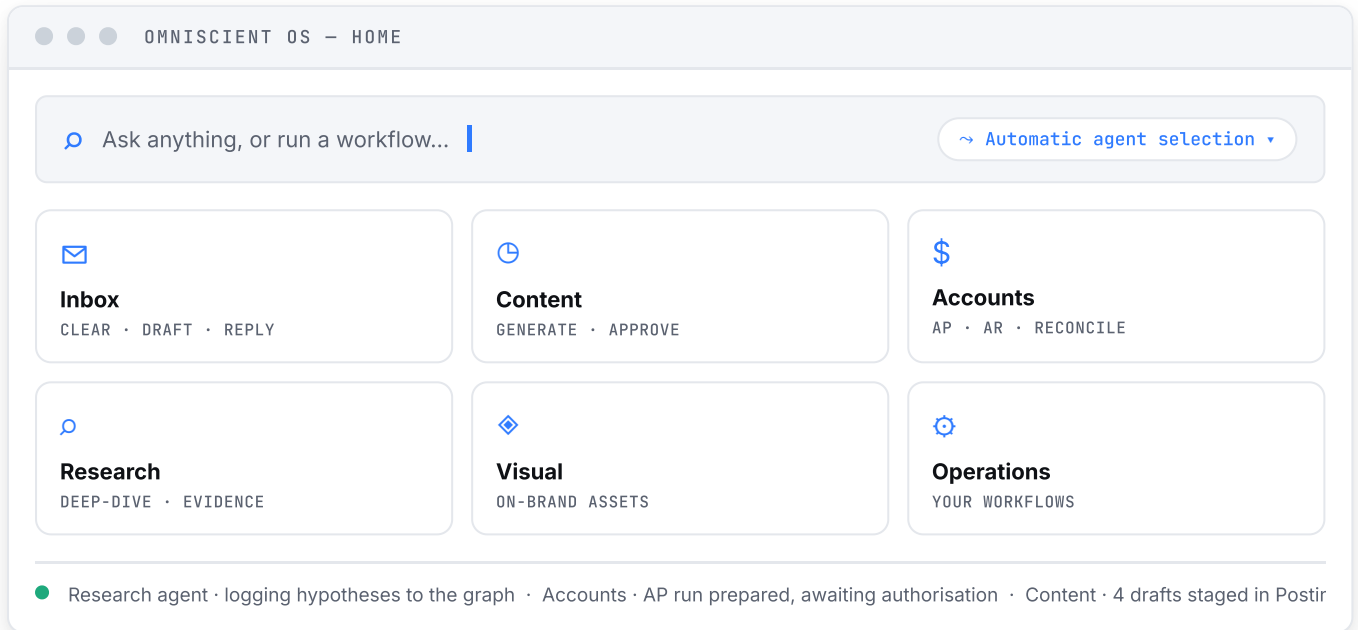
Offshore vendors don't build for the Privacy Act, the emerging AU AI guardrails, or your sector's rules. "A US data centre, we're not sure" is no answer for your regulator.



Mid-market enterprises feel this hardest: **big enough to carry real compliance exposure and sensitive data — not big enough to build a sovereign AI platform in-house.**

Not a chatbot. An **operating system** your whole team runs on.

The hardware is why your CISO says yes. The OS is why your staff love it — productised from a real, working build.



Every Claude skill you use — included

The full library of slash commands, skills and named agents ships pre-installed on **your** hardware. No per-seat AI licence. Ask in plain language; the router picks the right local agent — or bursts, with consent, for the hardest tasks.

- Research
- Finance
- Legal
- Design
- Marketing
- + hundreds

Multi-stage work → one button

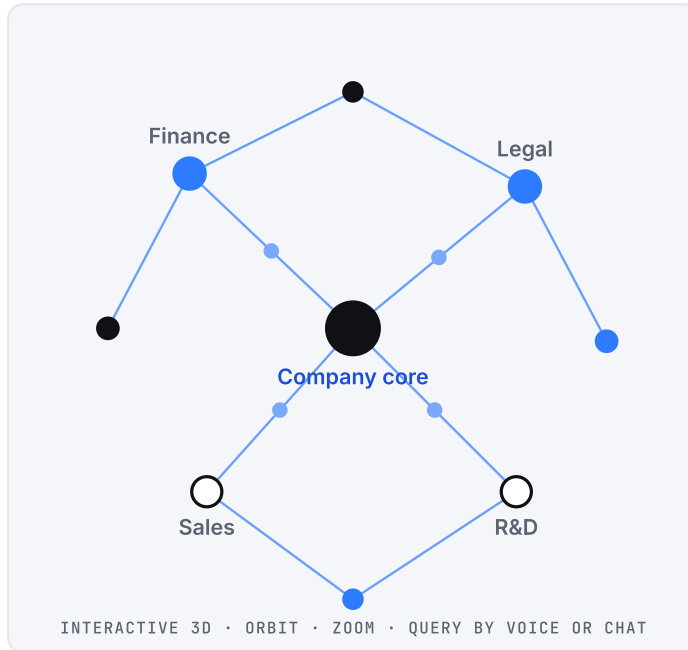
A calm home of work-group tiles, not a wall of dashboards. Each button is a saved orchestration — a dozen manual steps collapsed into one tap. Admins open the same flow as an editable node graph.

DAILY OPS	Email triage · content · AP/AR — scheduled, run overnight on the local GPU, presented for review. The GPU is never idle.
UPDATES	New capabilities arrive with an explainer + one Yes/No; on yes they wire themselves in — no config.


YOUR PROPRIETARY KNOWLEDGE, ALIVE

A second brain your company owns — and can watch grow.

Every file, email, contract, design and record is graphed into a living network of entities and relationships the AI can traverse — institutional memory, not a folder of files. It grows per employee and per department, and it **endures even when people leave**.



- Split by department**
Finance sees finance, Legal sees legal — under the same access boundaries as the rest of the box. Each team gets its own slice.
- It saves you money**
Complete pre-organised context means no long documents re-pasted on every request — the AI queries the graph for exactly what's relevant. Fewer tokens → faster, cheaper inference and smaller burst payloads — **and better answers**.
- You can watch it grow**
A live graph. As work happens, new nodes and connections appear — teams see their knowledge base grow alongside the business.



1 Work happens → 2 captured to the graph → 3 richer context → 4 cheaper, better automation ↻

...AND IT RESOLVES TO EVERY INDIVIDUAL – THE EMPLOYEE BRAIN

AI-adoption lens

See how each person uses, shares and builds with the tools. Real adoption and ROI per seat — not guesswork.

Live role capture

Every client, open task, commitment and decision — captured as a byproduct of work, not a status report.

Zero-loss handover

Leave, transfer, departure: the successor inherits the whole picture — full context, every client, every detail. Institutional knowledge no longer walks out the door.

A continuity & capability tool — not surveillance. Access-controlled, audit-logged, deployed transparently in line with Australian workplace & privacy law.

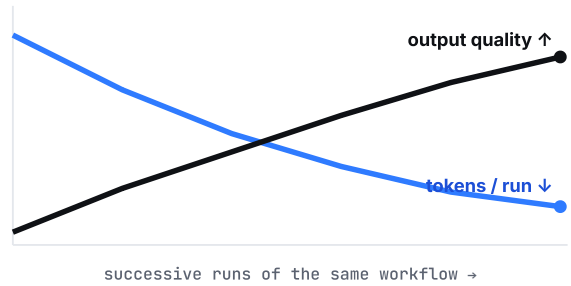
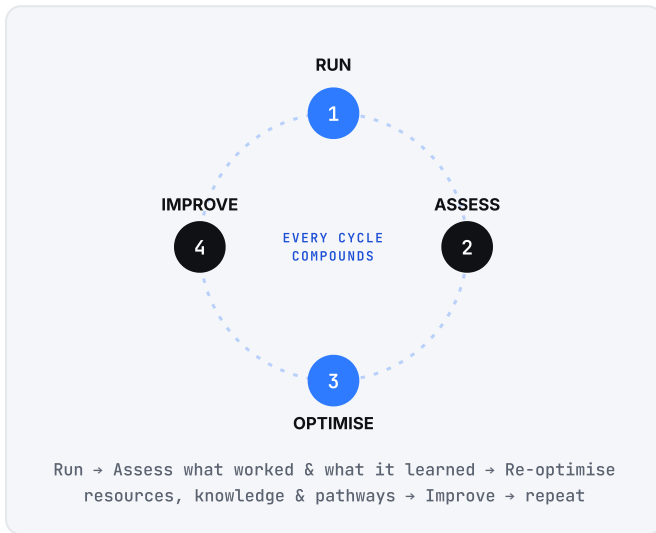
The difference: memory that doesn't leave **when people do**.

Today's AI forgets every session, and what your people know lives in their heads. Omniscient keeps a complete, enduring graph — **per employee and per department** — of every task, client and project. So when someone leaves, nothing is lost: the next person is handed full context and every detail. **The painful handover, and the erosion of institutional knowledge through staff turnover, is a thing of the past.**

It doesn't just do the work. It gets better at it — every run.



Conventional AI does the same task the same way every time, and waits to be told. Omniscient runs your repetitive work and workflows as self-improving loops: it reviews each run, learns, and re-optimises how it uses resources, knowledge and pathways — so cost falls and quality rises the longer it runs.



Across successive runs of the same workflow, **tokens per run fall** while **output quality rises**. Per-seat SaaS runs the opposite way — every token is full price, every time, forever.

↓

Fewer tokens each run

It finds the shortest correct path and stops re-deriving context it already holds — so the same task costs less every time it runs.

⚡

Faster each task

Proven workflows are saved, refined and re-used. A dozen manual steps collapse into a tuned, one-tap routine.

✦

Better outputs each project

Every run banks what worked. Quality compounds as the box accumulates and applies its own best approaches.

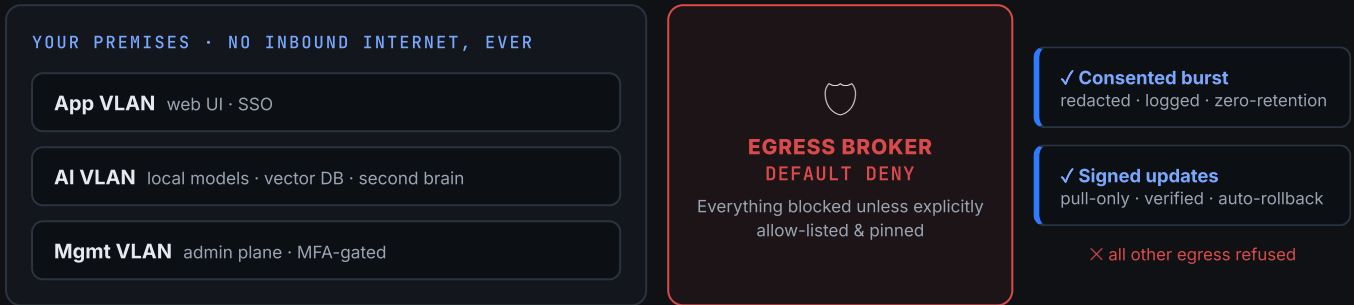
“We don't prompt our AI. We allocate **skills to tasks — and let efficient **workflows and loops** optimise themselves.”**

A different operating model from prompt-by-prompt chatbots: you define the outcome and the skills; the loops own the how — and get better at it on their own, run after run.

Sovereign by architecture — not by promise.

Two controls do most of the work: a single default-deny way out of the box, and a policy broker that gates every action an agent can take.

1 · NETWORK — THE ONLY PATH OUT IS DEFAULT-DENY



2 · ACTION TIERING — AGENTS CAN'T ACT UNILATERALLY



THE APPLIANCE — WHAT'S IN THE BOX

128–256 GB RAM	48 GB+ GPU · SIZED TO TEAM	NVMe + NAS STORAGE · RIGHT-SIZED	AMD·Intel LATEST CPU	Compact→rack FORM FACTOR	On-prem MODELS + DATA, LOCAL
--------------------------	--------------------------------------	--	--------------------------------	------------------------------------	--

Preloaded, expandable open-source estate — CRM · content · accounting · document editors · spreadsheets · BI · comms — matching the tools you use today with no performance penalty. Need a proprietary or outside API (bank feed, gov portal, industry service)? **We ensure interoperability**, through the same consented, logged egress controls.

The security story a CISO says yes to.

Each answer is a technical control that produces evidence — not a promise.

“Can our data leave?”

No inbound internet path, **ever**. A single default-deny egress broker permits only two pinned destinations — consented burst and signed updates. Flip one switch (or the physical air-gap) and the box is **an island**.

“Can the AI go rogue?”

Agents **can't call tools directly**. A policy broker gates every action; anything consequential — money, personal data, sends, deletions — needs **human or dual approval**. Injected instructions are tainted and can't escalate.

“Is there a backdoor for Omniscient?”

No standing access. Support is customer-initiated, time-boxed and recorded to **your** ledger. You hold the recovery keys; our escrow share can't decrypt alone. You can lock us out — you can never be locked out.

“How is it updated without one?”

The box **pulls** signed updates — no push, no open port — verifies signature + hash, and applies them as atomic, **auto-rolling-back** swaps in windows you approve.

Proof, not promises

Every control writes to a hash-chained, tamper-evident audit ledger — AI actions, approvals, admin actions, each burst and its consent record, every update and support session. When the regulator, the board or a customer asks what happened, **you can show them**. On demand, the box generates an auditor-ready evidence pack: data-residency attestation, access reports, model provenance, control status.

Same budget. Multiples of the output.

You don't win on per-seat price — you win on capacity. Per-seat SaaS prices access but caps usage, and seats sit idle outside working hours. An owned appliance turns the same spend into uncapped capacity that runs ~22 hours a day.

3-year cost — 200 seats vs one appliance

Per-seat SaaS (200 × \$150/mo)

\$1.08M

Omniscient appliance (3-yr lease, all-in)

~\$222K

Illustrative · 200 seats @ US\$150/mo · Business-tier 3-year lease + subscription + power · tune to your numbers in the live model.

~\$858K

SAVED OVER 3 YEARS
While capacity goes up, not down

~\$286K

SAVED EVERY YEAR
SaaS spend – lease, subscription & power

Unlimited

USAGE CEILING
No rate or token caps · runs ~22h/day

~85B

TOKENS / YEAR CAPACITY
≈ \$1.28M/yr of metered-API value

~19x

CHEAPER PER TOKEN
Local vs metered frontier API

~4x

VOLUME AT EQUAL SPEND
vs what the budget buys on demand

The point isn't "cheaper." It's **uncapped output for the same money.**

Cost holds flat or falls — but the AI's work is no longer limited by a plan, only by hardware that runs almost around the clock. Full, editable model at omniscientai.io.

Adopt it as R&D — claim up to 43.5% back.

Replacing capped subscriptions with a sovereign, always-on appliance isn't just procurement — it's a genuine experiment in new ways of working. For eligible Australian companies, that experimentation, and the appliance used to run it, can qualify for the R&D Tax Incentive: a **43.5% refundable offset** for companies under \$20M turnover.

EXAMPLE R&D HYPOTHESIS

“We hypothesise that an on-site, locally-served AI model — replacing subscription services — lets our team develop new workflows and methods that increase productivity, reduce costs, fully secure our institutional knowledge, eliminate context loss, and improve throughput through always-on AI.”

Systematic, hypothesis-driven work to develop and test new methods — where the outcome can't be known in advance — is exactly what the R&DTI is designed to support.

Worked example — a 10-person team

THE TEAM	10 people using the AI ~80% of the day to design & test new AI-enabled workflows
ELIGIBLE R&D LABOUR	~80% of those salaries, apportioned to the experimental activity
PLUS THE APPLIANCE	the 3-year lease payments , claimable as incurred — you lease, Omniscient owns
OFFSET	x 43.5% refundable (companies under \$20M turnover)

43.5%

REFUNDABLE OFFSET

On eligible R&D expenditure · < \$20M turnover

Self-documenting

CONTEMPORANEOUS EVIDENCE

The box's tamper-evident ledger generates the timestamped records an R&DTI claim needs

The appliance can help fund itself.

The team's salaries and the appliance's R&D-apportioned **lease payments** may both be claimable — leasing (rather than owning) means the hardware is claimable as incurred — and the box's evidence engine produces the records to substantiate it. **Not tax advice:** eligibility must be assessed and activities registered with AusIndustry; companies over \$20M turnover receive a non-refundable offset at a different rate. Omniscient can help structure and evidence the claim.

Compliant by physics. Yours by **ownership**.

- **Privacy Act 1988 & reforms** — onshore data, no offshore disclosure by default, access/retention/deletion controls, automated-decision transparency.
- **Australian AI guardrails** — human oversight on consequential actions, record-keeping, testing, model provenance, accountable roles.
- **Sector overlays** — APRA CPS 234, health records, government classification, defence-adjacent — on hardened tiers.
- **Essential Eight alignment** — patching, app control, MFA, least privilege, backups, logging — out of the box.

AU AI & privacy law is evolving fast. Your compliance pack is built and maintained with an Australian privacy/AI lawyer and shipped as versioned, signed updates — which is why the Sovereignty Subscription is a genuine ongoing service, not rent.

Edge

Branch / pilot · single node

MID-MARKET DEFAULT

Business

Single node · 99.5% target · HA upgrade → 99.9%+

Enterprise

Multi-site · negotiated SLA

Regulated

Defence / health / gov hardening

Commercial model

3-year appliance lease (Omniscient owns the hardware — so the lease + staff salaries are R&D-claimable) + **Sovereignty Subscription** (updates, patching, compliance, support) + optional **professional services**.

FROM ORDER TO PRODUCTIVE — WEEKS, NOT QUARTERS

01

Discovery

Audit workflows, data & obligations

02

Provision

Build, harden & load your agents

03

Install

On-site, into your network & identity

04

Onboard

Train staff; graph your knowledge

05

Manage

Updates & support, never your data

Book a sovereignty briefing.

See the box. Ask the hard questions. Get evidence, not slideware.

omniscientai.io

enquiries@omniscientai.io

Melbourne · vendor-neutral · Australian-sovereign