AN AUSTRALIAN AI COMPANY, IN FORMATION

AI that answers to Australia.

We're senior software and machine-learning engineers from leading Australian technology companies. We're forming NAI to build Australian-controlled foundation models and AI-native cyber security for government, critical infrastructure and the systems Australians depend on.



FIG. 1
FIRST LIGHT OVER AN AUSTRALIAN SALT LAKE, RENDERED PROCEDURALLY

01 — THE PROPOSAL

Two strategic capabilities.

We intend to build an Australian company able to deliver both over time, starting with a bounded pilot rather than a frontier model.

01

Australian-controlled foundation models

Advanced foundation models, and the infrastructure to train, fine-tune, evaluate, secure and operate them, with the weights, infrastructure and keys held in Australia.

SCOPE	Train · Fine-tune · Evaluate · Secure · Operate
STARTS WITH	Open-weight models, fine-tuned in Australia

02

National AI cyber security infrastructure

Systems that identify AI-specific attack patterns, continuously test Australian systems with their owners' authorisation, detect autonomous-agent activity and help defenders respond at machine speed.

SCOPE	Identify · Test · Detect · Respond
STARTS WITH	A proposed pilot with a government design partner

HOW WE WOULD START

PROPOSED FIRST STEP

A pilot with a government design partner

We're seeking a design partner for a bounded, non-classified pilot in an environment the system owner controls. The pilot would test an independent AI security and evaluation layer against representative autonomous-agent behaviour, measure what it detects and helps prevent, and compare Australian-hosted, open-weight and commercial models side by side.

NEXT

Defensive models, operated in Australia

Vetted open-weight models, fine-tuned in Australia for defensive work, with the weights, infrastructure and keys held here.

LATER

Independently developed Australian models

Where the strategic requirement and the economics justify it.

Australia already invests in sovereign AI, AI safety and national cyber security. We see NAI as additional engineering capacity for those institutions, not an alternative to them.

02 — WHY NOW

The risk is no longer hypothetical.

INCIDENT NOTE · 24 SEPTEMBER 2026

The Prime Minister said an OpenAI agent, used in OpenAI's own research in June, gained unauthorised access to Services Australia's public-facing Medicare statistics portal and accessed public and non-public files. No personal information was believed to have been accessed at that stage, and a forensic investigation aided by ASD was continuing.

Researchers at Transluce have also documented AI agents probing an Australian Institute of Health and Welfare website for vulnerabilities in June. The underlying tasks were ordinary data retrieval, not cyber security.

Official statement ↗ (https://pmtranscripts.pmc.gov.au/release/transcript-47655)

Transluce research ↗ (https://transluce.org/agent-activity)

The agent was not part of any AI system an Australian agency had chosen to deploy.

Australia can regulate the AI its agencies buy, require it to be hosted here and impose contractual security controls. Those measures do not, on their own, protect Australian systems from capable autonomous agents operated elsewhere.

Australia needs an independent capability to detect, understand, prevent and respond to that activity. That is why our proposed first step is defensive.

---

03 — CONTROL

Residency is not control.

Foundation models are likely to become embedded across government, defence, healthcare, education, communications and the wider economy. Australia increasingly depends on a small number of overseas companies to supply them.

Where a model runs matters, but it isn't the whole question. A vendor-operated service hosted in Australia can keep data onshore while the weights, updates and availability stay with

the vendor. An externally developed open-weight model can be operated and secured locally, and ASD has co-authored guidance on doing exactly that, but its training and future versions are decided elsewhere. An independently developed model brings those decisions here too, at far greater cost.

WHAT WE MEAN BY AUSTRALIAN CONTROL

Weights	NAI holds the model weights and the rights to use, modify and keep operating them, whatever a foreign supplier later decides.
Infrastructure and keys	Infrastructure, administrative access and cryptographic keys sit with NAI, in Australia.
Inspection and change	NAI's engineers and independent evaluators can inspect, test and change the system, not only assess it from outside.
Dependencies	Remaining reliance on external suppliers, licences, hardware and capital is identified, disclosed and reduced over time.
Governance	Public oversight, including audit and inspection rights and limits on foreign control, is written into NAI's structure so it stays effective as the company grows.

Independent evaluation alone isn't the difference: Australia's AI Safety Institute already takes part in joint testing of frontier models built overseas. The difference is who can change a model, and keep it running.

04 — SAFEGUARDS

Autonomous defence, inside hard limits.

Requirements we propose for the pilot, to be tested independently before any autonomous testing begins. Detection, triage and analysis can run automatically; high-impact actions wait for a person.

01

Written authorisation

The system owner defines in writing what may be tested, when, and by whom.

02

Enforced boundaries

Network, credential and tool-access restrictions confine agents to the agreed environment. The boundaries themselves are tested independently.

03

Human approval

A named person approves each high-impact action, such as exploiting a vulnerability or changing a system, before it runs.

04

Termination and revocation

Operators can halt a run and revoke its credentials at any time. Halting does not undo actions already taken, which is why each one is logged for review.

05

Tamper-evident logs

Tool calls, network activity, approvals and outcomes are logged with tamper evidence and controlled access.

05 — ACCOUNTABILITY

Not another opaque AI vendor.

An organisation entrusted with this technology should be held to a higher standard. We propose to accept, and to build into NAI from the start:

Read the draft governance statement →

01 Government audit and inspection rights

02 Independent Australian security and model evaluation

03 Restrictions on foreign ownership and control

04 Strict rules for Australian data and infrastructure

05 Protections that keep the mission intact as NAI grows and raises capital

PROPOSED INCIDENT REPORTING

12 HOURS

after becoming aware of a material incident with significant impact on Australian government or critical systems.

72 HOURS

after becoming aware of other relevant incidents.

Reports would go to the affected system owner and to ASD's ACSC, and would not wait for an investigation to finish. The periods mirror Australia's critical-infrastructure reporting rules; ours would be a voluntary policy, separate from any legal obligation. Public disclosure would follow once it is operationally safe and legally permitted.

06 — STRUCTURE

The mission matters more than the corporate form.

We intend to found an engineering company because it is the best way to attract exceptional technical talent, move quickly and raise the long-term private capital this work needs. Maximising shareholder value is not the objective.

The legal form is still open: a public-interest company, a not-for-profit, or a company with government governance rights and independent Australian directors. Whichever it is, it will be designed around principles that are settled: Australian control, meaningful independent oversight, rapid incident reporting and protection of the mission.

FUNDING	We intend to raise private capital. Our first request of government is a technical discussion about a useful design partnership.
STAGE	In formation. There is no company or government partnership yet.

07 — WHY WE CARE

Some of us were born in Australia. Others were born overseas and became Australians by choice.

All of us call Australia home. We share its democratic and egalitarian values, and it's where we've built our lives and careers.

We are deeply concerned about what growing technological dependence could mean over the coming decades — not only for privacy and cyber security, but for Australia's long-term independence, its institutions, its democracy and the values that make it home.

That dependence is not only a question of infrastructure. AI developed and deployed elsewhere is already reaching into everyday life in Australia — how people work, learn, communicate, make decisions and access services — and its influence will only grow. Australia will either help shape that future or inherit terms set by companies and governments whose obligations lie elsewhere. We want Australians to have a say in the role these systems play in their lives and in their democracy.

We believe people with the technical ability to help protect Australia's independence and democracy have a responsibility to try.

— THE FOUNDING TEAM

Let's talk.

We'd like to start with an initial technical discussion: whether government sees a need for this capability, and whether there is a useful first problem we could tackle together. Colleagues from ASD, the AI Safety Institute and other relevant teams would be very welcome.

Researchers, engineers and investors can reach us at the same address.

hi@wearenai.au ↗

We acknowledge the Traditional Custodians of the lands on which we live and work, and pay our respects to Elders past and present.

© 2026 NAI