Digital Infrastructure
INSTITUTE

COMMUNITY INFORMATION SHEET · SOVEREIGN AI

Open and closed AI models

A genuinely unsettled debate that matters for national capability — both sides, fairly

The key idea: how AI is shared shapes who can build

AI models come to the world two ways. **Closed models** are used through a provider's service: powerful, polished, but controlled by the company and priced by use. **Open-weights models** can be downloaded, studied, adapted and run on your own computing. Which approach dominates matters enormously for a middle-sized nation's AI capability and honest people disagree about it.

The case for open

Open models let Australian researchers, businesses and governments build without permission or dependency: run models on local infrastructure (data stays here), adapt them to local needs and languages, inspect them for safety research and avoid lock-in to any provider's pricing. For a country that won't build frontier models from scratch, open weights are the practical path to **model sovereignty**: capability without needing to own the frontier.

The case for caution

Once weights are released they can't be recalled: safety measures can be removed by whoever downloads them and misuse can't be switched off centrally the way a provider can cut off a service. Closed providers can monitor for abuse, patch problems fleet-wide and be held accountable as identifiable operators. Some risks scale with capability, so the debate sharpens as models improve.

Where it lands for Australia

In practice Australia will use both: closed services for some needs, open models on local infrastructure for others. The mix is a live policy question, not a settled one. For the digital-infrastructure story: open models running onshore need **local compute**, which links this debate directly to the data-centre build-out. DII takes no side; we flag the trade-offs so claims of either kind get proper scrutiny.

What good practice looks like

- Capability claims say whether they rest on open models, closed services, or both.
- Public-sector AI choices weigh dependency and data-location alongside capability and cost.
- Safety arguments engage the actual trade-offs on both sides, not caricatures.
- Local-compute implications of “run it onshore” proposals costed honestly.

Questions to ask

Does this project or service depend on open models, a closed provider, or both? What happens to capability, cost and data if the provider changes terms? What infrastructure does the onshore option actually require?

Want to know more? Sheets 01–03 cover the sovereignty layers and compute; the AI Safety Institute publishes technical analysis as this debate evolves. This sheet is general information, not legal advice; check the named sources for current detail.