Securing Australia's AI Future: Sovereign Infrastructure and Open-Source Innovation for Sensitive Data

Amr Hassan, David Powell, Gabriel Noaje, Luc Betbeder-Matibet, Sue Keay, Andrew Rohl, Mark Gray

Question 1: Does Australia need to build sovereign AI infrastructure?

- Does Australia need to build sovereign AI infrastructure for **sensitive-data workloads**, or can commercial solutions (neoclouds like Firmus, ResetData, SHARON AI, or hyperscalers with local regions) adequately serve our needs in healthcare, defence, and critical infrastructure?
- If sovereign capacity is essential, what specific capabilities—**technical, regulatory, or strategic**—would it provide that commercial offerings cannot?

Question 2: Technical Requirements and Priorities

If we commit to building sovereign AI infrastructure, what are the **critical technical requirements**? Specifically:

- What compute architectures and scale are needed?
- Which workload patterns must we prioritise in the next 12-24 months—base model training, fine-tuning, inference at scale, or federated learning approaches?
- How do we balance cutting-edge capabilities with practical deployment timelines?

Question 3: Sustainability and Governance Models

Sovereign AI infrastructure requires sustained investment in power, cooling, specialised staff, and regular hardware refreshes. What governance and funding model is most viable for Australia?

- Single-institution ownership dedicated to that institution's research priorities
- Cross-institutional cost-recovery consortia?
- Public-private partnerships with clear sovereignty protections?
- Designation as a national research infrastructure (NCRIS/NDRI) with ongoing Commonwealth support

Question 4: Open-Source Models in Regulated Contexts

- What role should **open-source and open-weight models play** in Australian regulated sectors?
- How do we balance the **benefits**—transparency, auditability, **adaptability**, avoiding vendor lock-in—against concerns around **security vetting, model provenance, and ongoing maintenance responsibilities**?
- Are there specific regulatory or operational requirements that make open models more (or less) suitable than proprietary alternatives for sensitive Australian workloads?

Discussion Questions

- What's the minimum viable sovereign capability—could we start small and scale, or do we need large-scale infrastructure from day one?
- In 5 years, will we regret building our own infrastructure, or regret not building it?
- For those working with sensitive data—what's your current approach? Are you using commercial cloud, on-premises solutions, or avoiding certain AI applications entirely?