

AI compute — why nations want their own

How AI turned computer chips and data centres into strategic resources

The key idea: capability now has a street address

AI systems are built and run on specialised computing (**compute**) concentrated in data centres full of advanced chips. Unlike software, compute is physical: it sits somewhere, uses that place's energy and networks and belongs to someone. That's why nations increasingly treat it like ports or power stations and why the data-centre proposals communities are seeing carry a national-capability dimension.

Training and inference, in one paragraph

AI work splits in two. **Training** builds a model: months of enormous computation on thousands of chips, the workloads that drive gigawatt-scale “AI factory” proposals. **Inference** is using the trained model to answer questions: smaller per use, but constant and growing with every AI-enabled service. Training can happen almost anywhere with power and fibre; inference increasingly wants to be closer to users. Both are arriving in Australia's development pipeline.

Why the chips became strategic

The advanced chips that train AI are made by very few companies, in very few places and demand exceeds supply. Access to them and to the facilities that house them, has become a matter of national planning in every developed economy. For Australia the realistic question is not self-sufficiency (no nation but the largest has that) but **enough for what?** Enough to run government and critical services? To support local research and industry? Those are policy choices now being made, partly through what gets built here.

What it means for communities

Compute-sovereignty arguments will increasingly appear in planning processes as justification for facilities and their energy use. The argument can be genuine and it still doesn't answer local questions. A nationally significant facility deserves assessment that is *faster and better coordinated*, not weaker; national benefit and local fairness are both legitimate and good process delivers both.

What good practice looks like

- National-capability claims specific: what workloads, for whom, guaranteed how.
- Strategic framing accompanied by, never substituted for, local impact answers.
- Public investments in compute state who gets access: researchers, industry, government.
- Energy and water claims for AI facilities held to the same standards as any other project.

Questions to ask

Is this facility for training, inference or both and what does that mean for its energy profile? Who will use its capacity and under what arrangements? If it's nationally significant, what has been done to coordinate, not shortcut, its assessment?

Want to know more? DII's data-centre series explains the facilities; Renewables sheet 08 covers how AI facilities can support the grid. Insight #13 covers the coming generation of AI factories. This sheet is general information, not legal advice; check the named sources for current detail.