



COMMUNITY INFORMATION SHEET

GPUs explained

The chips that power AI — and why they shape energy and cooling needs

What is a GPU?

A GPU (Graphics Processing Unit) is a type of computer chip originally designed to draw graphics for games. It turns out the same ability — doing thousands of calculations at the same time (in “parallel”) — is exactly what training and running AI needs. So GPUs (and similar AI accelerators) have become the workhorse of modern AI.

Why they matter for a data centre

AI facilities pack thousands of GPUs into racks. Each high-end AI GPU can draw roughly 0.7–1.2 kilowatts — many times more than a home computer — and a single rack of them can use as much power as several houses. They also run hot, which is why high-density AI sites increasingly use liquid cooling (see the Cooling systems sheet).

What	Roughly
A typical home computer	0.1 - 0.3 kW
One high-end AI GPU	~0.7 - 1.2 kW
A dense AI rack (many GPUs)	tens of kW — like several homes
A large AI campus	tens to hundreds of MW — a small town

The takeaway: the number and power of the chips is what drives a facility's electricity use, heat and cooling — the issues the rest of this series explains.

Want to know more? Your local council, the EPA Tasmania and ARPANSA publish further information. This sheet is general information, not medical, legal or planning advice; figures are indicative and a specific proposal is confirmed by qualified assessment.