

C600 IPU-PROCESSOR PCIe CARD



The Graphcore® C600 IPU-Processor PCIe Card is a high-performance accelerator card targeted for machine learning inference applications. Powered by the Graphcore Mk2 IPU with FP8 support, the C600 is a dual-slot, full height PCI Express Gen4 card designed for mounting in industry standard server chassis to accelerate machine intelligence workloads.

Up to eight C600 IPU-Processor PCIe Cards can be networked together using IPU-Link™ high-bandwidth interconnect cables, delivering enhanced IPU compute capability.

TECHNICAL SPECIFICATIONS

IPU-Processor	Graphcore Mk2 IPU with FP8 support
IPU-Cores™	1,472 IPU-Cores, each one a high-performance processor capable of multi-thread, independent code execution
In-Processor Memory™	Each IPU-Core is paired with fast, local, tightly-coupled In-Processor Memory. The C600 accelerator includes 900 MB of In-Processor Memory
Compute	Up to 560 teraFLOPS of FP8 compute Up to 280 teraFLOPS of FP16 compute
System interface	Single 8-lane PCIe Gen4 interface
Thermal solution	Passive
Form factor	PCIe full-height/length; double-slot
Dimensions	Length: 267 mm (10.50"); Height: 111 mm (4.37"); Width: 27.6 mm (1.09"); Mass: 1.27 kg (2.8 lbs)
IPU-Link™ support	32 lanes, 128 GB/s bandwidth (64 GB/s in each direction) IPU-Links
TDP	185 W
Auxiliary power supply	8-pin
Quality level	Server grade