

IPU-POD₁₆ DIRECT ATTACH

Datasheet





Table of contents

Overview	3
Product description	4
IPU-POD₁₆ Direct Attach	4
Software	5
Technical specifications	7
Environmental characteristics	8
Standards compliance for IPU-M2000s	8
Document details	9
Revision history	9




Overview

Graphcore's IPU-M2000 supports scale-up and scale-out for exascale machine intelligence compute. The IPU-POD scale-out options, based on the IPU-M2000, deliver scalable building blocks for the massive levels of compute in next generation machine intelligence workloads.

The IPU-POD is currently available in an IPU-POD₁₆TM configuration with 4 x IPU-M2000s, as well as an IPU-POD₆₄TM configuration with 16 x IPU-M2000s. IPU-POD₆₄ racks can be scaled for systems ranging from 64 to 64K IPU processors in switched or direct 3D torus IPU-FabricTM configurations.

Other configurations such as IPU-POD₃₂ and larger scale-out systems (IPU-POD₁₂₈ and IPU-POD₂₅₆) will be available in 2021 – please contact Graphcore sales for more information.




Product description

IPU-POD₁₆ Direct Attach

Graphcore's IPU-POD₁₆ Direct Attach system combines four IPU-M2000s delivering 4 petaFLOPS of AI compute directly attached to a pre-approved host server from a choice of technology providers including Dell and Supermicro. The IPU-POD₁₆ Direct Attach is a compact and powerful platform ideal for exploration, experimentation and concept and pilot development and is designed to scale with additional IPU-M2000s, host servers and a top of rack switch to a full IPU-POD₆₄. The IPU-POD₆₄ in turn is a building block for larger systems with up to 64,000 IPU processors delivering 16 exaFLOPS of AI compute.

Virtualization and provisioning software allow the AI compute resources to be elastically allocated to users and be grouped for both model-parallel and data-parallel AI compute.

A high-level view of the IPU-POD₁₆ Direct Attach cabling is shown in the figure below.



- Convenient cost effective evaluation platform
- Available through Graphcore channel for on-premise or Graphcloud
- Wide range of benchmarks and examples for IPU-POD16 performance evaluation
- Scale-out with IPU-POD64 and beyond

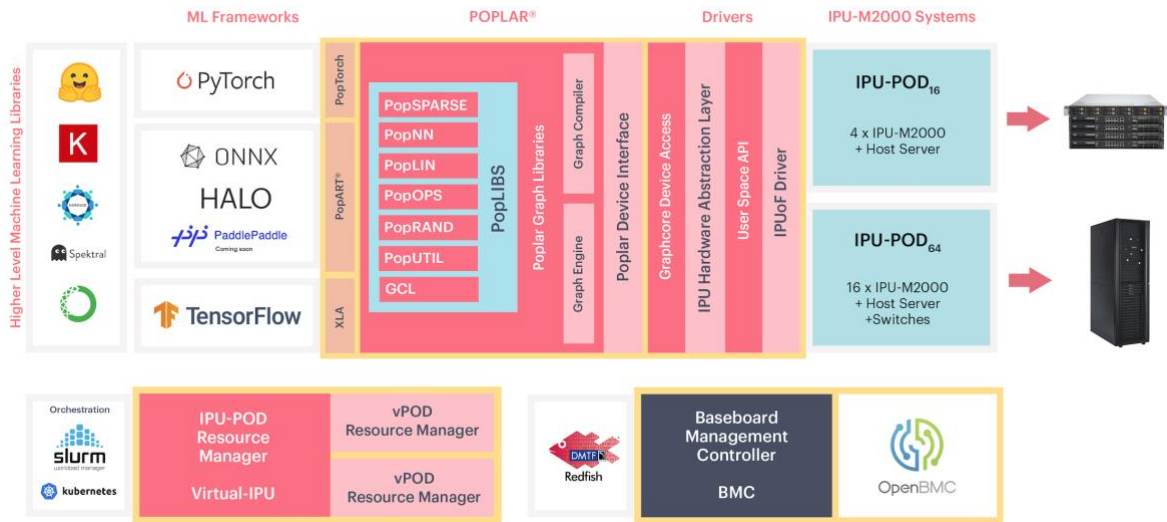
The IPU-POD₁₆ Direct Attach is available through Graphcore's network of reseller and OEM partners.



Software

IPU-POD₁₆ is fully supported by Graphcore’s Poplar® software development environment, providing a complete and mature platform for ML development and deployment. Standard ML frameworks including TensorFlow, ONNX, and PyTorch are fully supported along with access to PopLibs through our Poplar C++ API. Note that PopLibs, PopART and TensorFlow are available as open source in the Graphcore GitHub repo <https://github.com/graphcore>. The Poplar SDK also includes the PopVision™ visualisation and analysis tools which provide performance monitoring for IPU’s - the graphical analysis enables detailed inspection of all processing activities.

In addition to these Poplar development tools, the IPU-POD₁₆ is enabled with software support for industry standard converged infrastructure management tools including OpenBMC, Redfish, Docker containers, and orchestration with Slurm and Kubernetes.





Poplar SDK	Complete end-to-end software stack for developing, deploying and monitoring AI model training jobs as well as inference applications on the Graphcore IPU
ML frameworks	ONNX, TensorFlow, and PyTorch
Deployment options	Bare metal (Linux), VM (HyperV), containers (Docker) Supported operating systems are Ubuntu 18.04 and CentOS 7.6 For other OS options please contact sales
Host-Links	RDMA based disaggregation between a host and IPU over 100Gbps RoCEv2 NIC, using the IPU over Fabric (IPUoF) protocol Host-to-IPU ratios supported: 1:1 up to 1:64
Graphcore Communication Library (GCL)	IPU-optimized communication and collective library integrated with the Poplar SDK stack. Support all-reduce (sum,max), all-gather, reduce, broadcast Scale at near linear performance to 64k IPUs
PopVision	Visualization and analysis tools
Graphcore Virtual IPU SW	IPU-M2000 and IPU-POD ₁₆ resource manager IPU-Fabric topology discovery and validation
Provisioning	REST API and SSH/CLI for IPU allocation / de-allocation into isolated domains (vPODs) Plug-ins for SLURM and Kubernetes (K8)
Resource monitoring	REST API and SSH/CLI for accessing the IPU-M2000 monitoring service Prometheus node exporter and Grafana (visualization) support
Lights out management	Baseboard Management Controller (OpenBMC) Dual-image firmware with local rollback support Console support, CLI/SSH based Serial-over-Lan and Redfish REST API



Technical specifications

IPU processors	16 Colossus GC200 IPU processors
IPU-Cores™	23,552
Worker Threads	141,312
In-Processor-Memory	14.4GB
AI compute	4 petaFLOPS AI (FP16.16) compute 1 petaFLOPS FP32 compute
Exchange Memory	Up to 526.4GB (includes 14.4GB In-Processor Memory (4x 3.6GB per IPU-M2000) and 512GB Streaming Memory (4x 64GB DIMM x2 per IPU-M2000))
Default	
Internal SSDs	16 x 1TB M.2 SSD for program and data store
IPU-POD₁₆ host server(s)	1 x pre-qualified host server (see the list of approved servers here) Please contact Graphcore sales or your channel partner for options
IPU-POD₁₆ switches	The IPU-POD ₁₆ can also be implemented in a switched configuration. Please contact Graphcore sales
Optional	
IPU-POD₁₆ thermal	Air cooled with built-in N+1 hot-plug fan cooling system in each of the individual components (IPU-M2000s, servers and switches)
Rack airflow	All IPU-POD ₁₆ components (IPU-M2000s, and server) are mounted for airflow direction front of rack (single door, cold aisle side) to back of rack (split door, hot aisle side)
Airflow rate	1750 CFM

For information on IPU-POD₁₆ integration with datacentre infrastructure, please contact Graphcore sales.



Environmental characteristics

Operating temperature and humidity (inlet air)	10-32°C (50 to 90°F) at 20%-80% RH (*)
--	--

Operating altitude 0 to 3,048m (0-10,000ft) (**)

(*) Altitude less than 900m/3000ft and non-condensing environment

(**) Max. ambient temperature is de-rated by 1°C per 300m above 900m

Standards compliance for IPU-M2000s

EMC standards	Emissions: FCC CFR 47, ICES-003, EN55032, EN61000-3-2, EN61000-3-3, VCCI 32-1 Immunity: EN55024, EN61000-4-2, EN61000-4-3, EN61000-4-4, EN61000-4-5, EN61000-4-6, EN61000-4-8, EN61000-4-11
---------------	--

Safety standards	IEC62368, IEC60950
------------------	--------------------

Certifications	North America (FCC), Europe (CE), UK (UKCA), Australia (RCM), Taiwan (BSMI), Japan (VCCI), South Korea (KC), China (CQC) CB-62368, CB-60950
----------------	---

Environmental standards	EU 2011/65/EU RoHS Directive, XVII REACH 1907/2006, 2012/19/EU WEEE Directive
-------------------------	---



Document details

Revision history

This document's revision history is as follows:

Version	Date	Notes
0.1	7 th of March 2021	First draft
0.2	24 th of March 2021	Second draft
1.0	26 th of March 2021	First release

Trademarks & copyright

Graphcore® and Poplar® are Registered Trademarks of Graphcore Ltd.

Colossus™, IPU-Core™, In-Processor-Memory™, Exchange Memory™, Streaming Memory™, IPU-Tile™, IPU-Exchange™, IPU-Machine™, IPU-M2000™, IPU-POD™, IPU-Link™, Virtual-IPU™, AI-Float™, IPU-Fabric™, PopART™, PopLibs™, PopTorch™ and PopVision™ are Trademarks of Graphcore Ltd.

All other trademarks are the property of their respective owners.

© Copyright 2020, Graphcore Ltd.