


## DELL EMC DSS8440 IPU SERVER

The Dell EMC DSS8440 Graphcore IPU Server lets innovators create new breakthroughs in AI with today's state of the art models and next-generation machine learning architectures.

It is a 4U rack-mounted chassis, containing eight Graphcore C2 IPU-Processor PCIe Cards, each of which has two Graphcore Colossus GC2 IPU (Intelligence Processing Units), supported by Graphcore's Poplar software stack. Designed to support both training and inference on the same hardware for full utilization of your machine intelligence compute, the DSS8440 IPU Server is a new tool in your machine intelligence toolkit.



*Dell DSS8440 containing 8 Graphcore C2 Dual-IPU-Processor PCIe Cards*

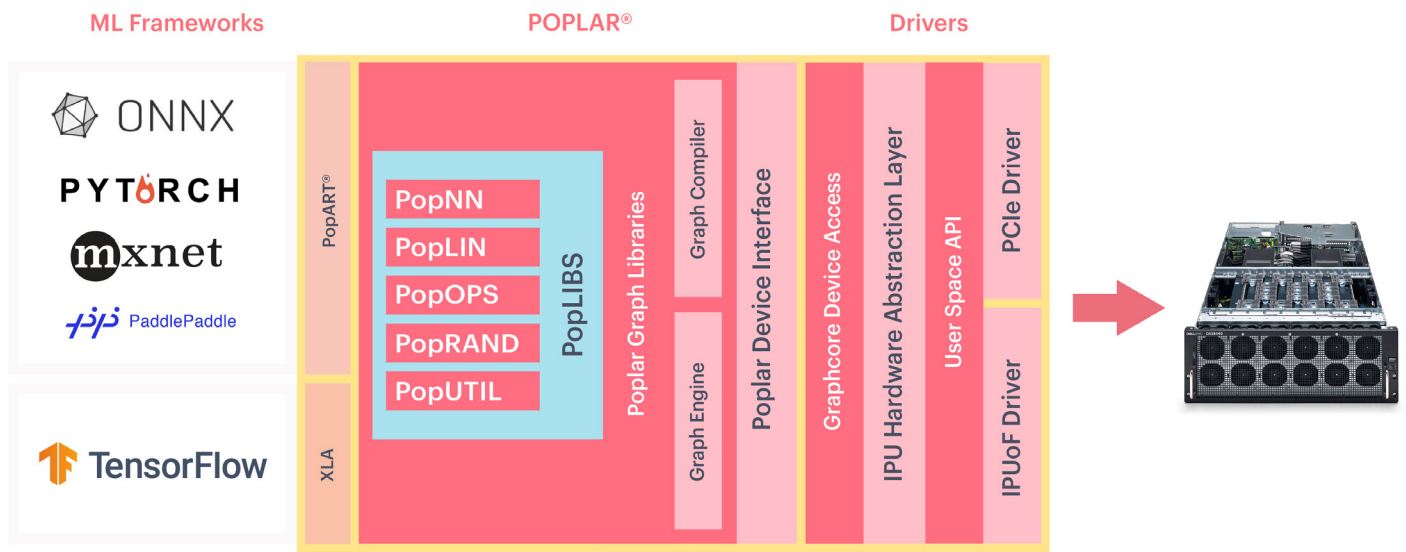
## TECHNICAL SPECIFICATIONS

<b>IPU Cards</b>	8 Graphcore C2 Dual-IPU-Processor PCIe Cards
<b>IPU Processors</b>	16x IPU Colossus GC2 IPU
<b>IPU-Link™ Technology</b>	28 IPU-Link™ cables, enabling high-bandwidth shared pool of compute across the chassis
<b>IPU Compute</b>	1.6 petaFLOPS mixed-precision compute
<b>IPU Memory</b>	4.8GB In-Processor Memory™, with 45TB/s bandwidth
<b>System CPU</b>	2x Xeon Platinum 8168
<b>System Memory</b>	24x 32GB 2.4GHz DDR4 ECC Registered DIMM Modules
<b>OS Storage</b>	2x 480GB SATA 2.5in Enterprise SSD
<b>System Storage</b>	88x 2TB NVMe Enterprise SSD
<b>Network</b>	10Gbps/25Gbps Ethernet
<b>System Software</b>	Ubuntu® Linux; Graphcore Poplar
<b>Gross weight</b>	39kg (86lb)
<b>Dimensions</b>	4U rack. W: 440mm (17.32"); H: 172mm (6.77"); D: 845mm (33.26")
<b>Max power</b>	4800W (at 220–240V)
<b>Operating temperature</b>	10°C to 35°C (50°F to 95°F)
<b>Storage Temperature</b>	–40°C to 70°C (–40°F to 158°F)

# POPLAR<sup>®</sup>

## The Software stack for IPU Platforms

The Dell DSS8440 IPU Server is fully supported by Graphcore's Poplar<sup>®</sup> software stack, providing a complete scalable platform for accelerated development.



- **Use existing ML platforms** such as TensorFlow and PyTorch: no need to recode your projects as Poplar transforms your existing models for optimal IPU performance.
- **Program the IPU directly:** Poplar contains C++ and Python interfaces that allow you to program 'close-to-the-metal'. Open-sourced libraries are provided for all common machine learning data types and manipulations.

### SEAMLESS

- Easy integration into existing software
- Easy integration into existing development workflows

### FLEXIBLE

- New algorithms can be developed quickly
- Libraries can be shared, extended and optimised

### POWERFUL

- Get the highest efficiency from hardware
- Take advantage of hardware nonlinearity and PRNG functions

For more information, visit [www.graphcore.ai](http://www.graphcore.ai)

To contact our sales team, email [sales@graphcore.ai](mailto:sales@graphcore.ai)