Cirrascale CLOUD SERVICES

Graphcore IPU Cloud Instances and Servers
Cirrascale provides customers with the option to use Graphcore IPU servers in the cloud or on premises

GRAPHCORE

Designed for machine intelligence workloads such as Natural Language Processing and Financial Risk Analysis.

Accelerating Next Generation Artificial Intelligence

Graphcore systems excel at both training and inference. The highly parallel computational resources together with graph software tools and libraries, allows researchers to explore machine intelligence across a much broader front than current solutions. This technology lets recent success in deep learning evolve rapidly towards useful, general artificial intelligence in NLP, financial risk analysis and more.

Flat-Rate Billing Model

Using a Graphcore IPU cloud instance with Cirrascale ensures no hidden fees with our flat-rate billing model. You pay one price without the worry of fluctuating bills like those at other providers.

No Ingress / Egress Fees on Data

Cirrascale doesn't charge any ingress or egress fees on transferring data in and out of our cloud service. We respect the fact that using your server means using big data.

High-Speed Storage and Networking

Cirrascale has partnered with the industry's top storage vendors to supply our customers with the absolute fastest storage and networking options available. Specialized NVMe hot-tier storage offerings such as WekaIO Matrix and IBM Spectrum Scale ECE deliver more than 10GBytes/second per node across an InfiniBand network. Cirrascale can also provide connectivity through 100Gb Ethernet or 100Gb EDR InfiniBand so that your Graphcore IPU Cloud Instances are connected together in a fast, private, protected network.

Graphcore IPUs with Dell EMC DSS 8440 Server

Cirrascale offers Graphcore Intelligence Processing Units (IPUs) technology to customers as a monthly cloud service, as well as an on-premises solution for sale. The Dell EMC DSS 8440 is a machine learning server with Graphcore technology that can deliver more than 1.6 PetaFlops of machine intelligence compute, using 8 Graphcore C2 IPU-Processor PCIe cards, each with 2 IPU processors, all connected with high speed IPU-Link™ technology in a standard 4U chassis.

Full Software Stack and Framework Support

Graphcore has designed the worlds very first graph tool chain specifically designed for machine intelligence - the Poplar software stack. Poplar seamlessly integrates with TensorFlow and Open Neural Network Exchange (ONNX) allowing developers to use their existing machine intelligence development tools and existing machine learning models. Graphcore also delivers a full training runtime for ONNX and is working closely with the ONNX organisation to include this in the ONNX standard environment. Initial PyTorch support is available in Q4 2019 with full advanced feature support becoming available in early 2020.

Graphcore IPU Cloud

Dual Intel Xeon Gold 6248 Processors
768GB RAM / 2x 480GB SSD / 4x 1TB NVMe
8x Graphcore® C2 IPU-Processors (IPU-Link)
Dual 25Gb Bonded, Non-Blocking Connectivity

Cloud Instances Starting at
$4,000 per server per week **$15,999** per server per month

Graphcore IPU Servers

- Dell EMC DSS 8440 Chassis
- Dual Intel Xeon Gold 6248 Processors
- 768GB DDR4-2666MHz RAM
- 2x 480GB SSD and 4x 1TB NVMe
- 8x Graphcore® C2 IPU-Processors connected via IPU-Link technology
- Mellanox ConnectX-5 EDR Adapter
- Quad 2400W Redundant Power (2+2)
- 3 Year Onsite Parts Warranty

Immediately Available for Sale Starting at
$124,499

Cirrascale Graphcore IPU Cloud Instances and Graphcore IPU Servers are AVAILABLE NOW!

http://graphcore.cirrascale.com