

## Product Support Engineer Job Description:

### OVERVIEW

The use of machine learning to reimagine software application and service development is exploding. Companies from every corner of the industry -- the biggest cloud service providers to corporate industrials to financial services to healthcare to retailers -- are exploring new ways of building products and services using data-centric learning models in place of traditional explicit programming. The drive to deliver more timely and more accurate results is compelling an ever greater need for specialized computing power. GPUs have been hailed as the solution to those computing needs but the industry is actively searching for a better, more efficient solution. Graphcore has that solution.

Graphcore's Intelligence Processing Units, or IPU's, are specifically designed for artificial intelligence and compute-dense graph applications. These are not GPUs, Graphics Processing Units, but rather graph processors especially adept at the kinds of computations used in understanding relationships within a sea of data. Architecturally, IPU's look nothing like GPUs. They offer performance, latency, and power efficiency advantages a GPU will be unable to match. Graphcore offers support for popular industry ML frameworks and a full tool suite for developers to innovate both within and outside those frameworks.

We are looking for Product Support Engineers who excel in the resolution of highly complex, technical, and varied support issues with customers on the leading edge of machine learning development. Our Product Support team engages with customers directly via web, email, chat, phone, and outreach to web-based developer communities. The successful candidate will have the drive and intellectual horsepower to resolve the most difficult technical customer issues, often through collaboration with other technical specializations. He or she will be part of Graphcore's face to the customer and able to work effectively and efficiently with customers and other Graphcore employees across the company.

In your work you will support some of the world's top machine learning innovators at deep learning research groups, at academic institutions, at innovative machine learning start-ups, at leading automotive companies, and at some of the world's largest cloud and internet companies. You will need to develop a deep understanding of the IPU architecture and the associated Poplar™ tools and become familiar with leading machine learning frameworks.

## REQUIREMENTS

- 2+ years of related experience in a high tech electronics industry in a technical support, programming or technical customer-facing role
- Bachelors in Engineering, Computer Science, Mathematics, Physics or similar field
- Proficient with C++ and Python, performance optimization, and software debug
- Possess deep technical expertise and can solve complex technical problems through sound, creative troubleshooting
- Demonstrate critical thinking, strong communication skills and ability to develop strategic on-going customer relationships
- Enjoy team work, and actively contribute to their peer group
- Keen interest to learn about the exciting new field of AI

## DIFFERENTIATORS

- Masters in related computationally intensive science or engineering field
- Experience with high performance parallel programming in C++, CUDA, OpenCL
- Experience implementing high performance numerical computation, such as Matrix Multiply and/or CNN convolution
- Experience working with modern deep learning software architecture and frameworks including: Tensorflow, MxNet, Caffe, Caffe2, Torch, and/or PyTorch
- Experience working and debugging issues with Linux, including system administration and drivers

This position will be based at our office in our Bristol, UK and could require occasional domestic and international travel.

Graphcore is an equal opportunity workplace and we would like to hear from all qualified applicants. If you have a disability or any special needs that we might need to accommodate, please just let us know.