

IBM Unveils New AI Software, Reduces Barriers for Data Scientists to Fuel Cognitive Development

New PowerAI solution reduces training of Deep Learning from weeks to hours

ARMONK, N.Y., May 10, 2017 /PRNewswire/ -- IBM (NYSE: [IBM](#)) today announced a significant new release of its [PowerAI](#) deep learning software distribution on Power Systems that attacks the major challenges facing data scientists and developers by simplifying the development experience with tools and data preparation while also dramatically reducing the time required for AI system training from weeks to hours.

Data scientists and developers use deep learning to develop applications ranging from computer vision for self-driving cars to real time fraud detection and credit risk analysis systems. These cognitive applications are much more compute resource hungry than traditional applications and can often overwhelm x86 systems.

"IBM PowerAI on Power servers with GPU accelerators provide at least twice the performance of our x86 platform; everything is faster and easier: adding memory, setting up new servers and so on," said current PowerAI customer Ari Juntunen, CTO [at Elinar Oy Ltd](#). "As a result, we can get new solutions to market quickly, protecting our edge over the competition. We think that the combination of IBM Power and PowerAI is the best platform for AI developers in the market today. For AI, speed is everything —nothing else comes close in our opinion."

The new PowerAI roadmap announced today offers four significant new features that address critical customer needs for AI system performance, effective data preparation, and enterprise-level software:

- **Ease of Use:** A new software tool called "AI Vision" that an application developer can use with limited knowledge about deep learning to train and deploy deep learning models targeted at computer vision for their application needs.
- **Tools for data preparation:** Integration with IBM Spectrum Conductor cluster virtualization software that integrates Apache Spark to ease the process of transforming unstructured as well as structured data sets to prepare them for deep learning training
- **Decreased training time:** A distributed computing version of TensorFlow, a popular open-source machine learning framework first built by Google. This distributed version of TensorFlow takes advantage of a virtualized cluster of GPU-accelerated servers using cost-efficient, high-performance computing methods to bring deep learning training time down from weeks to hours
- **Easier model development:** A new software tool called "DL Insight" that enables data scientists to rapidly get better accuracy from their deep learning models. This tool monitors the deep learning training process and automatically adjusts parameters for peak performance.

"Data scientists and an emerging community of cognitive developers will lead much of the innovation in the cognitive era. Our objective with PowerAI is to make their journey to AI as easy, intuitive and productive as possible," said Bob Picciano, Senior Vice President, IBM Cognitive Systems. "Power AI reduces the frustration of waiting and increases productivity. Power Systems were designed for data and this next era of computing, in great contrast to x86 servers which were designed for the client/server programmable era of the past."

PowerAI Support for New NVIDIA Volta Data Center GPU

PowerAI is optimized for IBM Power Systems S822LC for HPC, which are designed for data-intensive workloads like deep learning, machine learning and AI. The tight integration of IBM POWER processors and NVIDIA GPUs is enabled by the NVIDIA® NVLink™ high-speed interconnect. This "super-highway" between the POWER processor and NVIDIA GPUs helps enables extremely fast data movement between the two types of processors. This exclusive CPU-to-GPU coupling delivers

higher performance in AI training, which is a key metric for developer productivity. It enables innovation at a faster pace, so developers can invent and try new models, parameter settings and data sets.

PowerAI will support the NVIDIA Volta® architecture announced today. Volta features the next generation NVLink, with two key enhancements that benefit PowerAI customers: (a) the data transfer between the Power9 CPUs and Volta GPUs is ten times faster than between Volta GPUs and x86 CPUs, which rely on the old PCI-e 3.0 interface, first introduced 4 years ago, and (b) it is memory coherent, which makes programming GPU accelerators much easier for software developers by automatically moving data between the system memory connected to the Power9 CPU and the GPU memory.

About PowerAI

PowerAI's curated, tested, and pre-packaged distribution of the major deep learning frameworks run on IBM Power System servers built for AI. This combination of best-in-class hardware and software is capable of handling compute-heavy workloads with ease. PowerAI provides data scientists and developers with the fastest time to development, both on-premises and in the cloud.

PowerAI is an enterprise software distribution of popular machine learning and deep learning open-source application frameworks. PowerAI is a single download, easy to install binary distribution that includes TensorFlow, Caffe, Torch, Theano, Chainer, NVIDIA DIGITS, and other machine and deep learning frameworks, along with the associated libraries and packages.

The PowerAI ecosystem includes software like Anaconda from Continuum Analytics, the H2O machine learning libraries from H2O, the Bons.ai AI software developer tools, and many others. IBM also offers enterprise support and services to developers using deep learning for their applications.

Courtney Lowell
IBM Media Relations
courtney.lowell@ibm.com
512-599-1676

SOURCE IBM
