

IBM Spectrum Deep Learning Impact for Power Systems



Quickly put deep learning to work with an easy to install, end-to-end, enterprise solution

Highlights

- Accelerate training times with software optimized for IBM Power Systems
 - Reduce time spent importing, transforming and preparing data
 - Improve ROI by sharing resources among many data scientists running different models
 - Greater model accuracy with hyper-parameter search and optimization
 - Faster time to results with distributed training on multiple GPUs and servers
 - More accurate models with training visualization and tuning
 - Less risk with runtime resiliency in case of server or GPU failures
 - Simplify administration with a consolidated framework for deep learning, monitoring and reporting
-

Deep learning, from data ingestion and preparation to training and tuning doesn't have to be complex or time consuming. Neither does scaling applications and moving them into production. IBM Spectrum Deep Learning Impact for Power Systems software enables you to build a deep learning environment providing an end-to-end workflow that allows data scientists to focus on training, tuning and deploying models into production.

With Deep Learning Impact, organizations can get started working with their data for deep learning while avoiding highly manual and repetitive steps and bypassing the need for specialized domain knowledge. The solution deploys with simple software downloads that gives data scientists everything needed to build a distributed deep learning environment in hours rather than days or weeks—and manage it easily as the environment grows.

Deep Learning Impact is designed to address the deep learning lifecycle with a focus on the steps that are the most time consuming or require highly specialized knowledge—whether the iterative and time-consuming nature of the workflow, the lack of skills to train and tune models, the need to implement open source frameworks, the high demands for computing capacity or the challenges of scale.

In addressing these issues, Deep Learning Impact not only enables powerful deep learning capabilities, it makes achieving them a greater reality for more organizations.



End-to-end support for the deep learning workflow

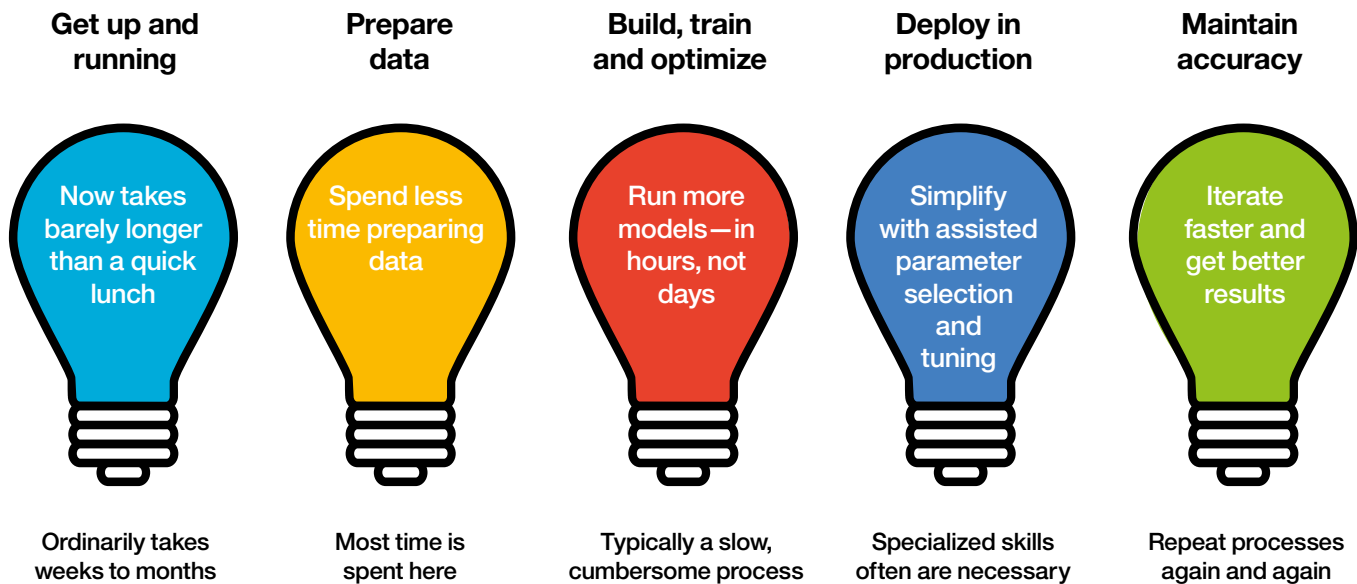
Financial services firms use deep learning for fraud detection and market prediction. Healthcare organizations use it to assist in disease detection and diagnostics. Transportation companies use it for autonomous assisted driving. In fact, for virtually any industry, deep learning presents radical new possibilities.

An electrical power company in Asia, for example, deployed IBM technologies to create a deep learning analysis system for checking 40,000 high-voltage transmission towers with drones. The system is trained to take images of components and analyze them to identify damage. The result has increased the number of inspections possible in a day by ten times¹—while reducing the electrical dangers faced by inspection and repair personnel.

But while deep learning can produce exciting results, implementing the technology can be a challenge. Each step in the deep learning workflow can create obstacles that organizations find hard to overcome.

To meet these challenges Deep Learning Impact provides simplifications and optimizations in an end-to-end workflow. This is a process that stretches from installing and configuring the environment to ingestion of data; data preparation and transformation to meet the requirements of deep learning frameworks; building, training and optimizing the neural models that make deep learning possible; deploying the model in production; and improving the model by retraining using new data as needs evolve.

Improving the steps to deep learning



Taking advantage of a distributed server architecture, Deep Learning Impact enables data scientists to quickly ingest, transform, train and iterate by running the processes in parallel. Deep Learning Impact is built to take advantage of IBM Spectrum Conductor, a highly available multitenant application designed to build a shared, enterprise-class environment for deploying and managing modern computing frameworks and services, such as Spark, Anaconda, TensorFlow, Caffe, MongoDB and Cassandra. Spectrum Conductor also provides centralized management and monitoring, along with end-to-end security. IBM Spectrum Deep Learning Impact runs on IBM Power System servers.

“The IBM POWER platform is a great cognitive platform, if not the best out there. The IBM Power team identified the need for and implemented acceleration before anyone else in the industry and are already on their third generation with the highest speed accelerator interconnects (i.e., NVLink) and a coherent architecture (i.e., CAPI) that can share main memory with the accelerator.”²

Faster training on accelerated clusters

IBM Spectrum Deep Learning Impact for Power Systems provides an enterprise-grade solution designed to meet the needs of high performance deep learning applications, including support for:

- Multitenancy, driving higher utilization and ROI by dynamically sharing server resources among many data scientists running multiple models
- Elastic resource allocation, which adds resources to a model in runtime without interrupting the training and provides resiliency in case of server or GPU failures
- Distributed data ingest, transformation and training, so jobs are processed in parallel over a cluster of servers—helping reduce the time spent manipulating data
- A distributed training fabric designed to allow most applications to run in parallel without the need for code changes
- Training visualization and tuning for monitoring the accuracy of the model while training is in progress and for making adjustments or stopping if not converging or low accuracy
- Hyper-parameter search and optimization to improve accuracy with suggestion-based logic while training is running
- Lifecycle support for the deep learning frameworks and Spark infrastructure
- End-to-end security for confident enterprise deployments

With conventional deep learning solutions, most of a data scientist’s time is spent importing, transforming and preparing data for training. By contrast, Deep Learning Impact is designed to reduce that time with a rich set of tools, automation and workflows, enabling the data scientist to spend more time training and optimizing models. Importantly, the distributed implementation that Deep Learning Impact makes it possible to reduce the amount of time needed to import and execute transformations by running the tasks simultaneously.

With software and frameworks optimized to take full advantage of IBM Power servers with NVLink CPUs and NVIDIA GPUs, IBM benchmarks have seen 50x improvements, cutting training times down from days to hours.³

IBM Spectrum Deep Learning Impact for Power Systems at a glance

Hardware requirements	IBM Power System S822LC for HPC (8335-GTB) servers
Software requirements	IBM PowerAI V1.5 base package Red Hat Enterprise Linux 7.4 operating system
Scalability	Up to 64 nodes with up to 256 GPUs
Distribution	Electronic download in multiple eAssemblies Physical media not available.

Why IBM?

For the rapidly growing and quickly evolving artificial intelligence category of deep learning, IBM Spectrum Deep Learning Impact for Power Systems enables organizations to achieve a faster time to results with simplified management. IBM services and support provide distributed deep learning based on parallel processing and elastic training that more easily and effectively deliver the performance advantages of Spark application management with optimized performance and improved time to result.

For more information

To learn more about IBM Spectrum Deep Learning Impact for Power Systems, please contact your IBM representative or IBM Business Partner, or visit:
ibm.com/us-en/marketplace/spectrum-deep-learning-impact

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: ibm.com/financing



© Copyright IBM Corporation 2018

IBM Systems
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
March 2018

IBM, the IBM logo, ibm.com, Power, and POWER are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

This document is current as of the initial date of publication and may be changed by IBM at any time.

All client examples cited or described are presented as illustrations of the manner in which some clients have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions. Contact IBM to see what we can do for you.

It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

¹ Based on IBM customer experience

² Patrick Moorhead, “IBM CEO Ginni Rometty Makes Her Case At InterConnect Why IBM Is The Right Business Platform,” *Forbes*, March 26, 2017.

³ Hillery Hunter, “IBM Research achieves record deep learning performance with new software technology,” *IBM Research*, August 8, 2017.



Please Recycle