

IBM and NVIDIA Collaborate to Expand Open Source Machine Learning Tools for Data Scientists

RAPIDS, new GPU-accelerated libraries to be made available in PowerAI on IBM POWER9, IBM Cloud, and IBM Watson Studio

IBM (NYSE: IBM) today announced that it plans to incorporate the new RAPIDS™ open source software into its enterprise-grade data science platform for on-premises, hybrid, and multicloud environments. With IBM's vast portfolio of deep learning and machine learning solutions, it is best positioned to bring this open-source technology to data scientists regardless of their preferred deployment model.

"IBM has a long collaboration with NVIDIA that has shown demonstrable performance increases leveraging IBM technology, like the IBM POWER9 processor, in combination with NVIDIA GPUs," said Bob Picciano, Senior Vice President of IBM Cognitive Systems. "We look to continue to aggressively push the performance boundaries of AI for our clients as we bring RAPIDS into the IBM portfolio."

RAPIDS will help bring GPU acceleration capabilities to IBM offerings that take advantage of open source machine learning software including Apache Arrow, Pandas and scikit-learn. Immediate, wide ecosystem support for RAPIDS comes from key open-source contributors including Anaconda, BlazingDB, Graphistry, NERSC, PyData, INRIA, and Ursa Labs.

IBM is planning to bring RAPIDS to key areas across on-premises, public, hybrid, and multicloud environments, including:

- PowerAI on IBM POWER9, to leverage RAPIDS to expand the options available to data scientists with new open source machine learning and analytics libraries. Accelerated workloads have been proven to get a direct benefit from the special engineering that NVIDIA and IBM have done around POWER9, including integration of NVIDIA NVLink® and NVIDIA Tesla® Tensor Core GPUs. PowerAI is IBM's software layer that optimizes how today's data science and AI workloads run on heterogeneous computing systems, and our goal is for this improved performance trajectory for GPU accelerated workloads on POWER9 to continue with RAPIDS.
- IBM Watson Studio and Watson Machine Learning, to take advantage of the power of NVIDIA GPUs so that data scientists and AI developers can build, deploy, and run faster models than CPU-only deployments for their AI applications in a multicloud environment with IBM Cloud Private for Data and IBM Cloud.
- IBM Cloud, to users who choose machines equipped with GPUs will be able to apply the accelerated machine learning and analytics libraries in RAPIDS for their cloud applications and tap the benefits of machine learning.

"IBM and NVIDIA's close collaboration over the years has helped leading enterprises and organizations around the world tackle some of the world's largest problems," said Ian Buck, vice president and general manager of Accelerated Computing at NVIDIA. "Now, with IBM taking advantage of RAPIDS open-source libraries announced today by NVIDIA, GPU accelerated machine learning is coming to data scientists, helping them analyze big data for insights faster than ever possible before."

Machine learning is a form of AI that enables a system to learn from data rather than through explicit programming. Enterprises across multiple industries like retail, finance, and telecommunications, are either actively using machine learning or exploring machine learning for the potential value it offers to companies trying to leverage big data to help them better understand the subtle changes in behavior, preferences, or customer satisfaction.

Earlier this year, [IBM set a record in a tera-scale machine learning benchmark](#), beating the previous record holder by 46x. Using an IBM Research-developed machine learning algorithm called IBM Snap Machine Learning (Snap ML) running on IBM Power Systems AC922 servers with NVIDIA Tesla V100 Tensor Core GPUs, IBM researchers trained a logistic regression classifier in 91.5 seconds using [an online advertising dataset released by Criteo Labs](#) with over 4 billion training examples.

ⁱ Statements regarding IBM's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.

Certain statements in this press release including, but not limited to, statements as to: the benefits and impact of IBM incorporating RAPIDS into its data science platform; and enterprises across multiple industries actively using or exploring machine learning are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the reports NVIDIA files with the Securities and Exchange Commission, or SEC, including its Form 10-Q for the fiscal period ended July 29, 2018. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2018 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVLink, RAPIDS and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Media Contacts

Kristin Bryson
+1 203 241 9190

kbryson@nvidia.com

Sam Ponedal

(916) 217-0145

sponeda@us.ibm.com