

## NVIDIA Launches Magnum IO Software Suite to Help Eliminate Data Bottlenecks for Data Scientists and AI, HPC Researchers

### New Offering Moves Vast Amounts of Data up to 20x Faster Than Previously Possible

SC19 -- NVIDIA today introduced [NVIDIA Magnum IO](#), a suite of software to help data scientists and AI and high performance computing researchers process massive amounts of data in minutes, rather than hours.

Optimized to eliminate storage and input/output bottlenecks, Magnum IO delivers up to 20x faster data processing for multi-server, multi-GPU computing nodes when working with massive datasets to carry out complex financial analysis, climate modeling and other HPC workloads.

NVIDIA has developed Magnum IO in close collaboration with industry leaders in networking and storage, including DataDirect Networks, Exceclero, IBM, Mellanox and WekaIO.

"Processing large amounts of collected or simulated data is at the heart of data-driven sciences like AI," said Jensen Huang, founder and CEO of NVIDIA. "As the scale and velocity of data grow exponentially, processing it has become one of data centers' great challenges and costs.

"Extreme compute needs extreme I/O. Magnum IO delivers this by bringing NVIDIA GPU acceleration, which has revolutionized computing, to I/O and storage. Now, AI researchers and data scientists can stop waiting on data and focus on doing their life's work," he said.

At the heart of Magnum IO is GPUDirect, which provides a path for data to bypass CPUs and travel on "open highways" offered by GPUs, storage and networking devices. Compatible with a wide range of communications interconnects and APIs -- including NVIDIA NVLink™ and NCCL, as well as OpenMPI and UCX -- GPUDirect is composed of peer-to-peer and RDMA elements.

Its newest element is GPUDirect Storage, which enables researchers to bypass CPUs when accessing storage and quickly access data files for simulation, analysis or visualization.

NVIDIA Magnum IO software is available now, with the exception of GPUDirect Storage, which is currently available to select early-access customers. Broader release of GPUDirect Storage is planned for the first half of 2020.

#### Ecosystem Support

"Modern HPC and AI research relies upon an incredible amount of data, often more than a petabyte in scale, which requires a new level of technology leadership to best handle the challenge. DDN, by taking advantage of NVIDIA's Magnum IO suite of software along with our parallel EXA5-enabled storage architecture, is paving the way to a new direct data path which makes petabyte-scale data stores directly accessible to the GPU at high bandwidth, an approach that was not previously possible."

-- Sven Oehme, chief research officer, DDN

"The amount of data that leading HPC and AI researchers now need to access continues to grow by leaps and bounds, making I/O a complex challenge for many to manage. IBM Spectrum Scale is designed to address the needs of any organization looking to accelerate AI and run data-intensive workloads. The use of IBM Spectrum Scale and NVIDIA GPU acceleration can help customers alleviate I/O bottlenecks and get the insights needed from their data faster."

-- Sam Werner, vice president of storage offering management, IBM

"Leading HPC and AI researchers choose Mellanox to provide them with the most advanced technology to move and process immense amounts of data as efficiently and quickly as possible. We have been working together with NVIDIA to ensure Magnum IO works seamlessly with Mellanox's state-of-the-art InfiniBand and Ethernet interconnect solutions and to enable our mutual customers to overcome data bottlenecks altogether and advance their science, research and product development activities."

-- Dror Goldenberg, senior vice president of software architecture, Mellanox Technologies

#### About NVIDIA

[NVIDIA's](#) (NASDAQ: NVDA) invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics and revolutionized parallel computing. More recently, GPU deep learning ignited modern AI — the next era of computing — with the GPU acting as the brain of computers, robots and self-driving cars that can perceive and understand the world. More information at <http://nvidianews.nvidia.com/>.

Certain statements in this press release including, but not limited to, statements as to: the performance, benefits, impact and availability of NVIDIA Magnum IO and NVIDIA GPUDirect Storage; the scale of data growing exponentially and moving it for processing as one of the greatest challenges and costs of data centers; the amount of data that HPC and AI researchers need to access continuing to grow; and the impact of the use of IBM Spectrum Scale and NVIDIA GPU acceleration 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 most recent reports NVIDIA files with the Securities and Exchange Commission, or SEC, including, but not limited to, its annual report on Form 10-K and quarterly reports on Form 10-Q. 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.

© 2019 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, Magnum IO and NVLink 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 Uchiyama

+1-408-486-2248

[kuchiyama@nvidia.com](mailto:kuchiyama@nvidia.com)