Technical Review Summary

Accelerating the Artificial Intelligence Journey with Dell EMC Ready Solutions for AI

With Ready Solutions for AI, Dell EMC has created standardized infrastructure stacks for machine learning (ML) and deep learning (DL) to accelerate the time to business value. These validated hardware and software stacks combine Dell EMC PowerEdge servers, Dell EMC Isilon storage, NVIDIA GPUs, high-speed networking, data science software, and AI libraries and frameworks into preconfigured, scalable, tuned systems. The Deep Learning with NVIDIA design featuring Isilon significantly outperformed the competition in training time, delivering 2.9 times the performance of one competitor for an AlexNet deep learning neural network in a GPU accelerated environment and 2.3 times another competitor for a ResNet50 deep learning neural network.

Read the Full Report at **www.emc.com.**

Two Dell EMC Ready Solutions for AI: ***Machine Learning with Hadoop*** and ***Deep Learning with NVIDIA***, are available today, and provide benefits including:

- **Fast deployment**—Rather than forcing the organization to select, configure, integrate, and tune components into an AI stack, Dell EMC Ready Solutions for AI are validated systems deployed by Dell EMC services, shrinking the time to deploy an AI environment from months to weeks while reducing skillset requirements and operational risk. Dell EMC provides a single point of contact to support the entire hardware and software stack.
- **Self-service Configuration**—Ready Solutions for AI increase data scientist productivity with self-service access to resources for machine learning and deep learning including frameworks and libraries such as BigDL, TensorFlow, Caffe, Neon, cuDNN, and cuBLAS. Deep Learning with NVIDIA includes Dell EMC Data Science Provisioning Portal, which reduces data scientist workspace configuration efforts; Machine Learning with Hadoop includes Cloudera's Data Science Workbench and Dell EMC Data Science Engines—containers that work with Data Science Workbench to configure the BigDL framework.
- **Simplified IT operations**—Deep Learning with NVIDIA includes Bright Cluster Manager and Machine Learning with Hadoop includes Cloudera Manager, providing one console for monitoring and management of the entire cluster.
- **Faster, Deeper Insights**—Dell EMC Ready Solutions for AI are designed for high performance with a scale-out architecture. The Deep Learning with NVIDIA design features the scale-out Dell EMC PowerEdge servers with GPU acceleration and Dell EMC Isilon All-Flash NAS to serve up massive amounts of data. This allows organizations to accelerate AI innovation with faster model training and real-time inferencing on bigger data sets. These solutions start small and then can non-disruptively grow by adding the compute and/or storage to meet the organization's AI needs.

Why This Matters

For AI, more complex models trained with larger data sets provide better results. With data sets in the tens of TBs to tens of PBs range and models with millions of parameters, high performance, high concurrency, and scale-out compute and storage become critical factors for organizations looking to obtain timely results from their AI efforts.

ESG validated that the 16GPU Dell EMC Ready Solutions for AI, Deep Learning with NVIDIA system was able to train the AlexNet model at 27,375 images/sec and the ResNet50 model at 11,126 images/sec. The scale-out Dell EMC solution featuring Isilon proved to be 2.2-2.9 times faster than systems from two other vendors. ESG also validated that Deep Learning with NVIDIA maintained processing speed as the system was scaled, achieving 89-95% of baseline performance as the number of GPUs was doubled and quadrupled. This ensures that organizations can maximize their return on investment as they scale out compute and storage to accelerate AI model development.

ESG Technical Review Highlights

ESG Lab performed hands-on evaluation and testing of Ready Solutions for AI: Deep Learning with NVIDIA, at Dell EMC's facilities. The following is a summary of the results:

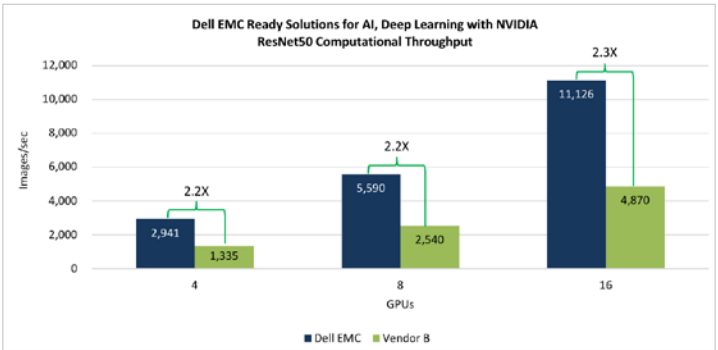

- With four GPUs, ResNet50 on Dell EMC Ready Solutions for Deep Learning processed 2,941 images/sec, scaling to more than 11,000 images/sec with 16 GPUs.
- ResNet50 on the solution ran 2.2-2.3 times faster than a competitor's solution.
- With four GPUs, AlexNet on Dell EMC Ready Solutions for Deep Learning processed almost 7,700 images/sec, scaling to more than 27,000 images/sec with 16 GPUs.
- AlexNet on the solution ran 1.9-2.9 times faster than competitor solutions.
- The Deep Learning solution featuring Isilon was free from I/O bottlenecks, keeping the GPUs at nearly 95% average utilization by saturating them with data delivering more than 1.25 GB/sec (ResNet50) and 3GB/s (AlexNet) with the 16GPU solution.
- The scale-out Dell EMC Ready Solutions for Deep Learning demonstrated linear scaling efficiency, achieving 89-95% of baseline performance as the systems were scaled from four to 16 GPUs, maximizing the return on investment.

The Bigger Truth

Organizations perceive that AI is the next technology that will enable the faster delivery of better business outcomes. According to an ESG research survey on the subject of machine learning and artificial intelligence trends conducted in June 2017, 69% of respondents expect that ML and AI will deliver significant measurable outcomes in the near term, with 17% of respondents indicating that AI and ML were critical to their organization's strategy.

Lacking a standardized AI infrastructure stack, organizations can invest the time, effort, and money to select, acquire, integrate, configure, test, and validate their own custom stack. This complex process can take months, and the organization must juggle purchasing and support across many vendors. Public cloud solutions suffer from huge cost variability and the time and cost necessary to transfer and store terabytes to petabytes of data.

Dell EMC created the Ready Solutions for AI as standardized infrastructure stacks for machine learning and deep learning. These are validated and integrated hardware and software stack solutions, tuned and optimized to accelerate AI initiatives, shortening deployment time from months to weeks. Ready Solutions for AI simplify and accelerate data scientists' efforts, providing self-service workspaces where each data scientist can configure her own environment from a library of AI models and frameworks in just five clicks.

ESG validated that these solutions can accelerate AI model development. With PowerEdge C4140 servers accelerated with (16) NVIDIA GPUs and a chassis of Isilon F800 All-flash Scale-out NAS, Dell EMC Ready Solutions for AI trained the AlexNet model at 27,735 images/sec, and the more computationally complex ResNet50 model at 11,126 images per second. These results were from 2.2 to 2.9 times faster than results published by other vendors.

These integrated solutions for AI demonstrated scaling efficiency, keeping the GPUs pegged at 95% utilization while also achieving 89-95% of baseline performance as the systems were scaled from four to eight to 16 GPUs, maximizing the return on investment as more GPUs were added. These solutions are ideally architected to scale-out and eliminate I/O bottlenecks as they are applied to solve more complex problems with larger and larger data sets.

ESG recommends that organizations investigate how Dell EMC Ready Solutions for AI can simplify and accelerate their AI journey.

Read the Full Report at **www.emc.com**.