

Simplifying Deep Learning Infrastructure with Dell EMC

Boni Bruno, CISSP, CISM, CGEIT
Chief Solutions Architect



TRADITIONAL CLUSTER

Workload Profile:

Financial Services Company

192TB data set

16 years, 68 quarters

34.7 Million Customers

1.85 Billion performance records

XGBoost training set: 50 features

300 Servers | **\$3M** | **180 kW**



Dell EMC Solution

1/6 the Cost

1/14 the Space


1/9 the Power



Why Dell EMC Isilon for AI?



**High
Performance**



Flexibility



**Extreme
Scale**



**Enterprise
Data Governance**

**ALL THIS WITH JUST STANDARD ACCESS PROTOCOLS
NO OS MODIFICATION or SPECIAL DRIVERS NEEDED**

Dell EMC Isilon Deep Learning Solution Portfolio

Only vendor to offer flexibility & informed choice with NVIDIA, the leader in AI



Augment existing Isilon with F800 All-Flash



Dell EMC PowerEdge servers
accelerated with NVIDIA GPUs

↑↓ Massively Parallel IO ↑↓



Ready Solution for AI: Deep Learning with NVIDIA




Extreme Deep Learning Isilon with NVIDIA DGX-1



Isilon with NVIDIA GPUs Benchmark

Image detection with Tensorflow and 22 TB dataset

ResNet-50 (Training)



Highlights


- **Record performance:** 96% or more of theoretical max
- **Linear Scaling:** From 8 to 72 GPUs
- Only vendor to publish above 32 GPUs




Not STAC Benchmarks

Ready Solution for Deep Learning: Benchmark results

Training: Image Classification with TensorFlow and ImageNet Data Set



ResNet-50 Results




*Based on an Enterprise Strategy Group Technical Review "Accelerating the Artificial Intelligence Journey with Dell EMC Ready Solutions for AI", August 2018 AD# G18000181


Not STAC Benchmarks

Ready Solution for Deep Learning: Benchmark results

Training: Image Classification with TensorFlow and ImageNet Data Set



AlexNet Results




*Based on an Enterprise Strategy Group Technical Review "Accelerating the Artificial Intelligence Journey with Dell EMC Ready Solutions for AI", August 2018 AD# G18000181

Not STAC Benchmarks


Isilon with DGX-1: Benchmark Results

Training: Image Classification with TensorFlow and ImageNet Data Set



Training

ResNet-50



- 97% GPU utilization or higher
- 96% of local memory throughput with Isilon
- Linear Scaling from 8 to 32 to 72 GPUs


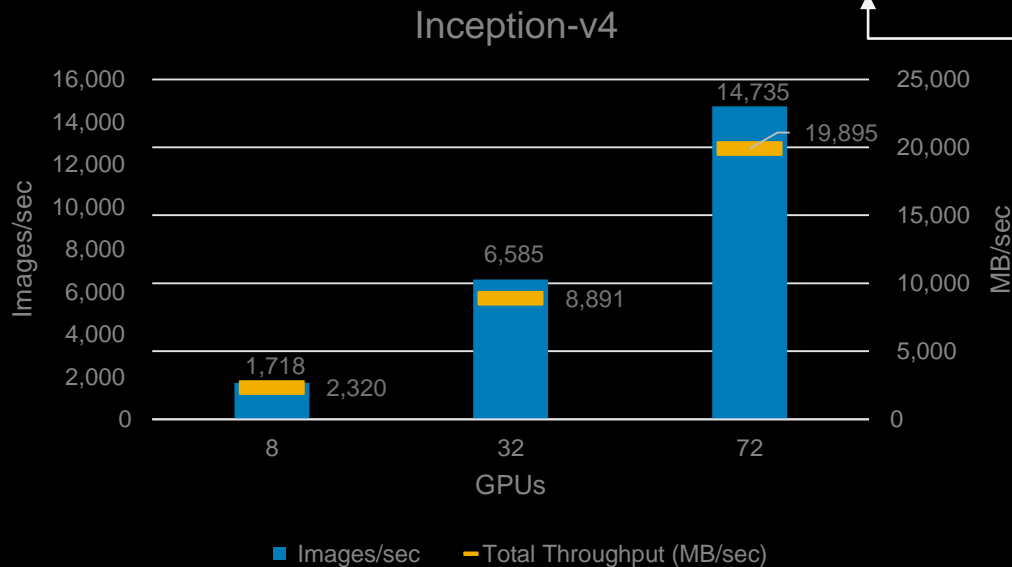
Not STAC Benchmarks

Isilon with DGX-1: Benchmark Results

Training: Large Image Classification with TensorFlow and ImageNet



Training

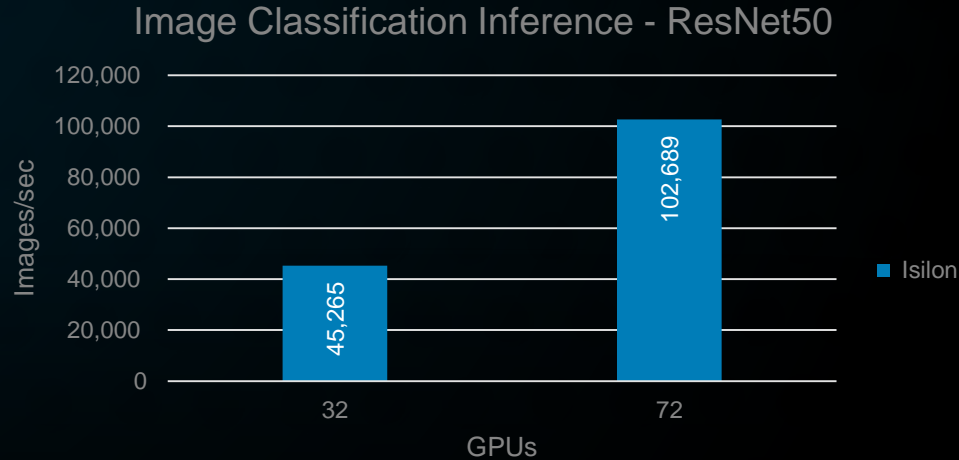



- 98% GPU utilization
- Linear Scaling from 8 to 32 to 72 GPUs
- 19.9 GB/s throughput with 72 GPUs

Not STAC Benchmarks

Isilon with DGX-1: Benchmark results

Inferencing: Image Classification with TensorFlow and ImageNet Data Set



Inferencing

- 100% of local memory throughput with Isilon
- Linear Scaling from 32 to 72 GPUs

Not STAC Benchmarks

Isilon eliminates the I/O bottleneck for AI at any scale



Ultra Dense

72TBs up to 924TBs in 4U

Seamless Scale to Over 58PBs

In a Single File System

No Special Storage Drivers Needed

Simply use NFS for data access




Faster training and validation of AI models



Higher model accuracy



Improve data science productivity



Maximize ROI of compute investments

Additional Reading Material...

- [Tick Data Analytics - Scaling Concurrency and I/O Performance](#)
- [Deep Learning with Dell EMC Isilon](#)
- [Digital Banking](#)

Thank You

