

NetApp CONDQ







Those at the forefront of digital transformation use technology to radically improve the performance and reach of their enterprise

By 2020 it is expected that **50% of the G2000** will see the majority of their business depend on their ability to create digitallyenhanced products, services and experiences

IDC Directions 02/17

67% of Global 2000 CEOs will put digital transformation at the center of their growth and profitability strategies.

Forbes 12/15

When successful in their data-driven digital transformation, organizations:

Enable new customer touchpoints Create innovative business opportunities

Optimize operations

Only 11% of institutions are using data aggressively to disrupt their industries



IDC research

Active IQ: AI-Powered Insights

A data-driven service using AI/ML and community wisdom to provide intelligent insights

AutoSupport telemetry built into NetApp[®] Products and Solutions Active IQ[®] Insights via Machine Learning and Community Wisdom

Analytics and insights at your fingertips







300,000 assets create over 200 billion data points each day

4 petabyte data lake processes over 100TB of data each month

Active IQ portal and mobile app for proactive, AI-assisted support



Edge to Core to Cloud

Seamless data management



Data Mobility - NetApp Data Fabric





TRADITIONAL DATA ANALYTICS CLUSTER

Workload Profile:

Fannie Mae Mortgage Data:

- 192GB data set
- 16 years, 68 quarters
- 34.7 Million single family mortgage loans
- 1.85 Billion performance records
- XGBoost training set: 50 features

300 Servers | \$3M | 180 kW



GPU-ACCELERATED DATA ANALYTICS CLUSTER

NVIDIA Accelerated Data Science Software Platform with NVIDIA Servers

1 DGX-2 | 10 kW

1/8 the Cost | 1/15 the Space

1/18 the Power





ML WORKFLOW STIFLES INNOVATION



DAY IN THE LIFE OF A DATA SCIENTIST



DATA SCIENCE WORKFLOW WITH NVIDIA

Open Source, End-to-end GPU-accelerated Workflow Built On CUDA



DATA PREPARATION

GPUs accelerated compute for in-memory data preparation Simplified implementation using familiar data science tools Python drop-in Pandas replacement built on CUDA C++. GPU-accelerated Spark (in development)

DATA SCIENCE WORKFLOW WITH NVIDIA

Open Source, End-to-end GPU-accelerated Workflow Built On CUDA



MODEL TRAINING

GPU-acceleration of today's most popular ML algorithms XGBoost, PCA, K-means, k-NN, DBScan, tSVD ...

DATA SCIENCE WORKFLOW WITH NVIDIA

Open Source, End-to-end GPU-accelerated Workflow Built On CUDA



VISUALIZATION

Effortless exploration of datasets, billions of records in milliseconds Dynamic interaction with data = faster ML model development Data visualization ecosystem (Graphistry & OmniSci), integrated with RAPIDS

TRANSFORMING RETAIL WITH RAPIDS

Inventory Forecast



"My previous bottleneck was I/O. ...15 seconds to pull in data for 10 stores (about 1 Million rows). With RAPIDS, we can pull in data for about 600 stores (60 Million rows) in less than 5 seconds. ... plain awesome."

- A mid-market specialty retailer with 6000 stores

IMPROVING DEMAND FORECASTS

Meat & Poultry

Rollback

Accurate demand forecasting is a critical but challenging science for retailers requiring massive amounts of data and compute cycles. Walmart is optimizing machine learning with NVIDIA RAPIDS open-source software on GPUs.

GPUs deliver 50x faster processing speed allowing Walmart to benefit from more sophisticated algorithms, reduce forecasting errors, and increase the efficiency of its supply chain.

Walmart

ANNOUNCING VOLVO CARS SELECTS NVIDIA DRIVE PLATFORM

DRIVE AGX Xavier to Pilot Next-generation Production Cars



AUTOMOTIVE WORKFLOW LARGE-SCALE DEEP LEARNING MODEL DEVELOPMENT

- Workflow, Tools, Supercomputing Infrastructure
- Data Ingest, Labeling, Training, Validation, Adaptation Automation, Best Model Discovery, Traceability, Reproducibility
- Purpose-built for Safety Standards of Automotive

"Data is the new source code"



DESIGNING INFRASTRUCTURE THAT SCALES

Insights gained from deep learning data centers

	Rack Design	Networking	Storage	Facilities	Software
 Example: Autonomous vehicle = 1TB / hr Training sets up to 500 PB 	 DL drives close to operational limits Similarities to HPC best practices High- bandwidth, ultra-low latency 	 Datasets range from 10k's to millions objects 	 assume higher watts per-rack Higher 	 Scale requires "cluster- aware" software 	
 RN50: 113 days to train Objective: 7 days 6 simultaneous developers = 97 node cluster 		inter- connect • High- bandwidth, ultra-low latency	 terabyte levels of storage and up High IOPS, low latency 	FLOPS/watt = DC less floorspace required	

DELIVERING DATA SCIENCE VALUE





NETAPP ONTAP AI

Simplify, Accelerate, and Scale the Data Pipeline for Deep Learning

HARDWARE

- NVIDIA DGX-1 | 5x DGX-1 Systems | 5 PFLOPS
- NETAPP AFF A800 | HA Pair | 364TB | 1M IOPS
- CISCO | 2x 100Gb Ethernet Switches with RDMA

SOFTWARE

- NVIDIA GPU CLOUD DEEP LEARNING STACK | NVIDIA Optimized Frameworks
- **NETAPP ONTAP 9** | Simplified Data Management
- **TRIDENT** | Provision Persistent Storage for DL

IMPLEMENTATION AND SUPPORT BY CONOA

- Single point of contact support
- Proven support model backed by NetApp and Nvidia



AI and Deep Learning Infrastructure Specialists

Fundamentals for AI

Models and libraries

Your own data is golden





Compute power



Fastest way to find the right use case



ACCELERATE YOUR JOURNEY TO AI

Visit booth E28

- Watch demo
- Learn how to get started
- Next steps
- Training offer
- Have fun and play
- Win Nvidia Shield TV 4K



NetApp CONDƏ