



Data is the Rocket Fuel for
Artificial Intelligence









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 digitally-enhanced 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



Data-Driven Digital Transformation Maturity

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




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

Active IQ® Insights via Machine Learning and Community Wisdom



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

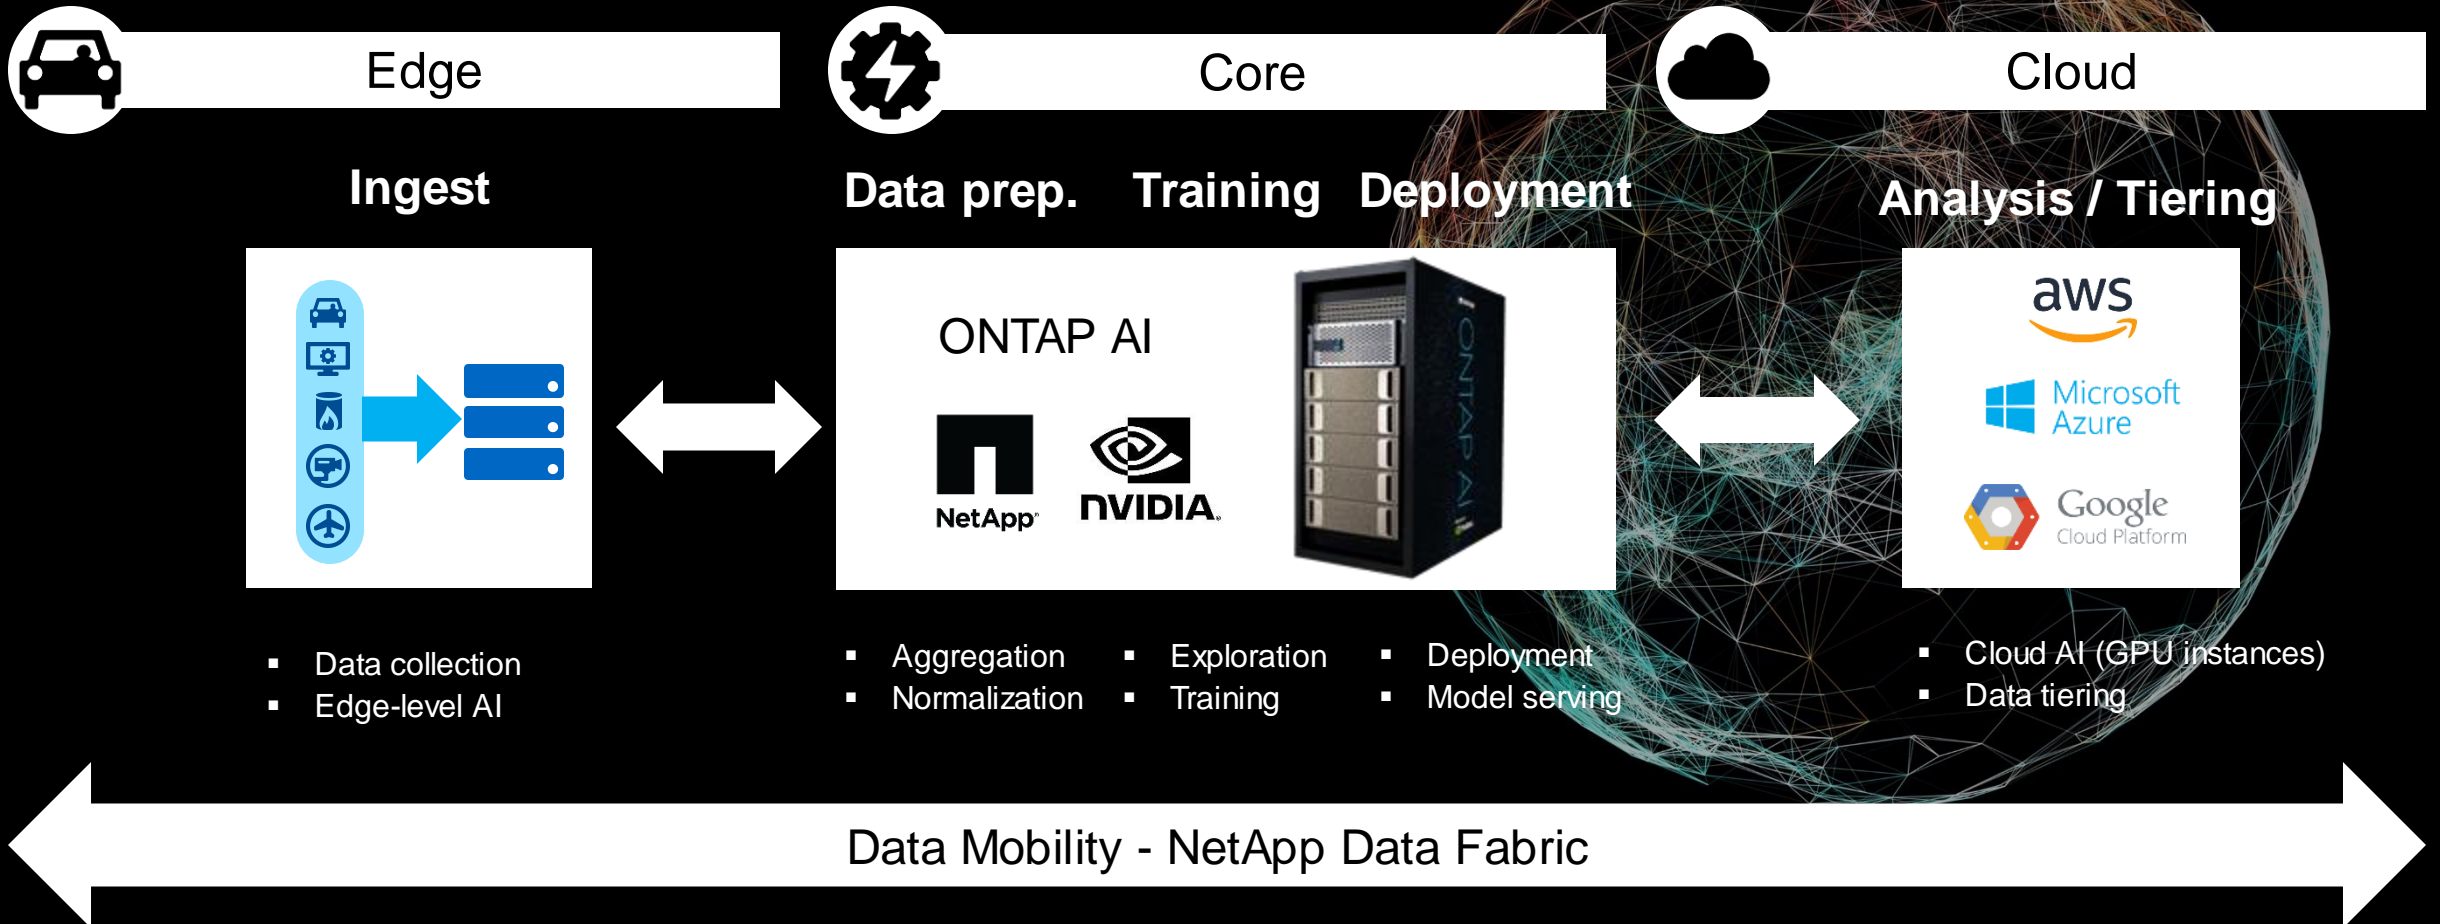
Analytics and insights at your fingertips



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

Edge to Core to Cloud

Seamless data management





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

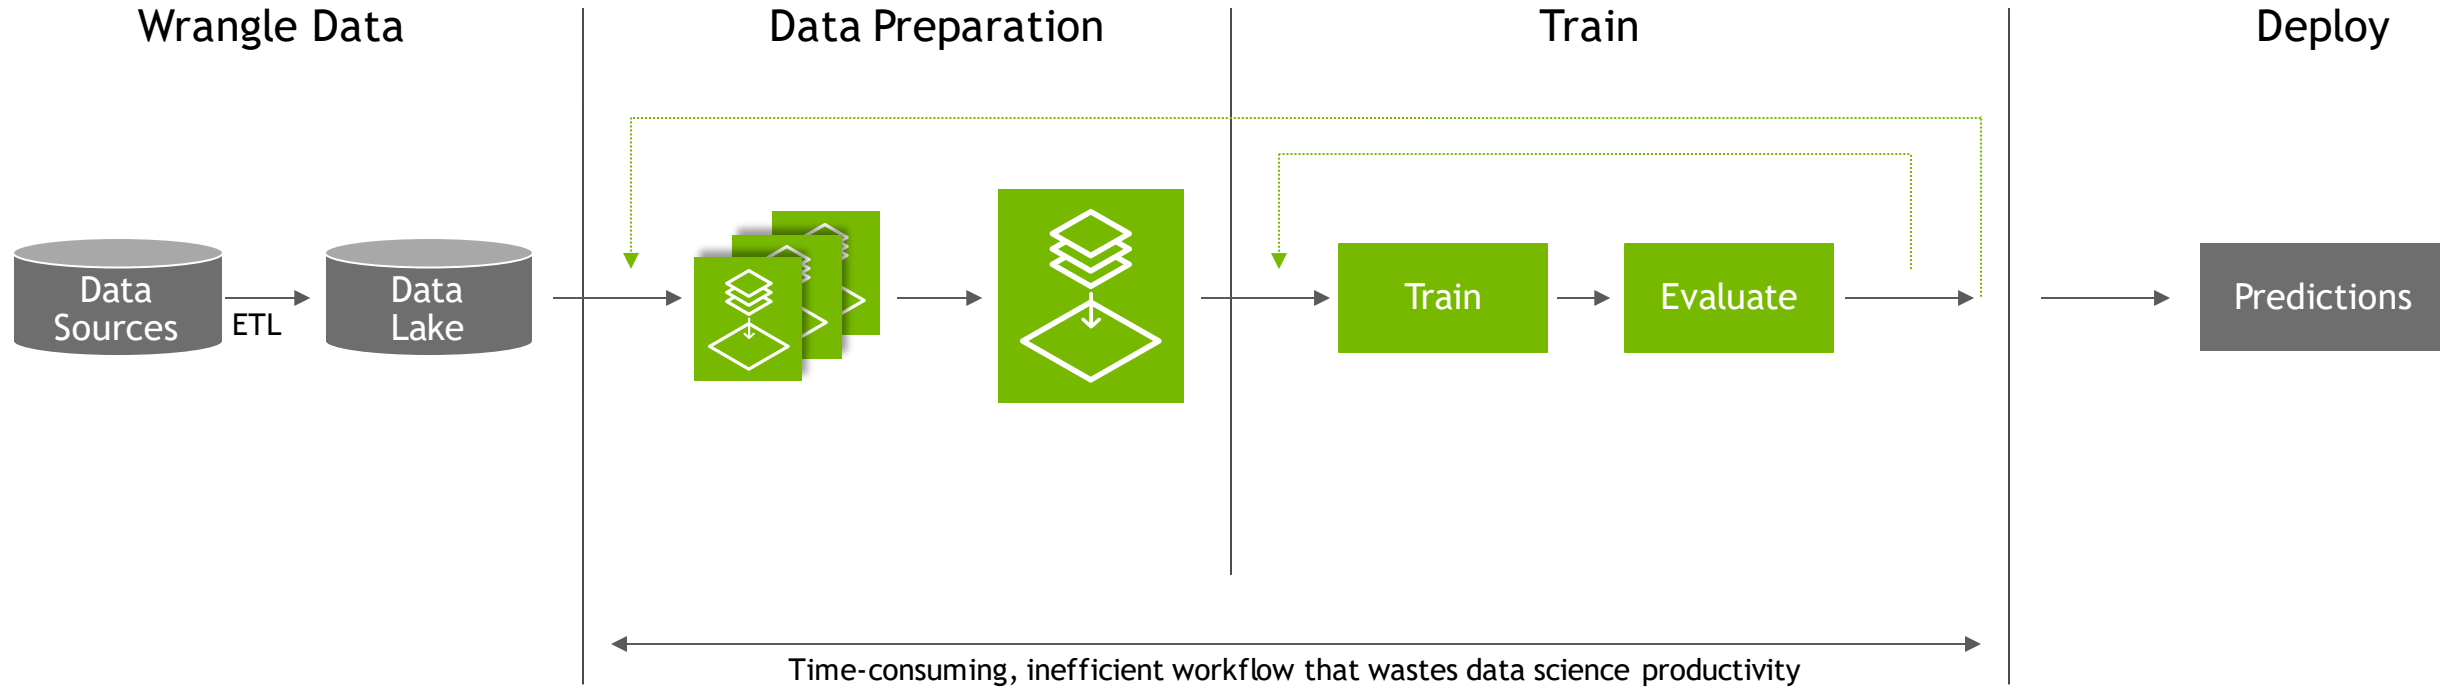
1 DGX-2 | 10 kW

1/8 the Cost | 1/15 the Space



1/18 the Power



ML WORKFLOW STIFLES INNOVATION




DAY IN THE LIFE OF A DATA SCIENTIST



Dataset Collection Analysis Data Prep Train Inference

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

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.



Meat & Poultry



ANNOUNCING VOLVO CARS SELECTS NVIDIA DRIVE PLATFORM

DRIVE AGX Xavier to Pilot
Next-generation Production Cars

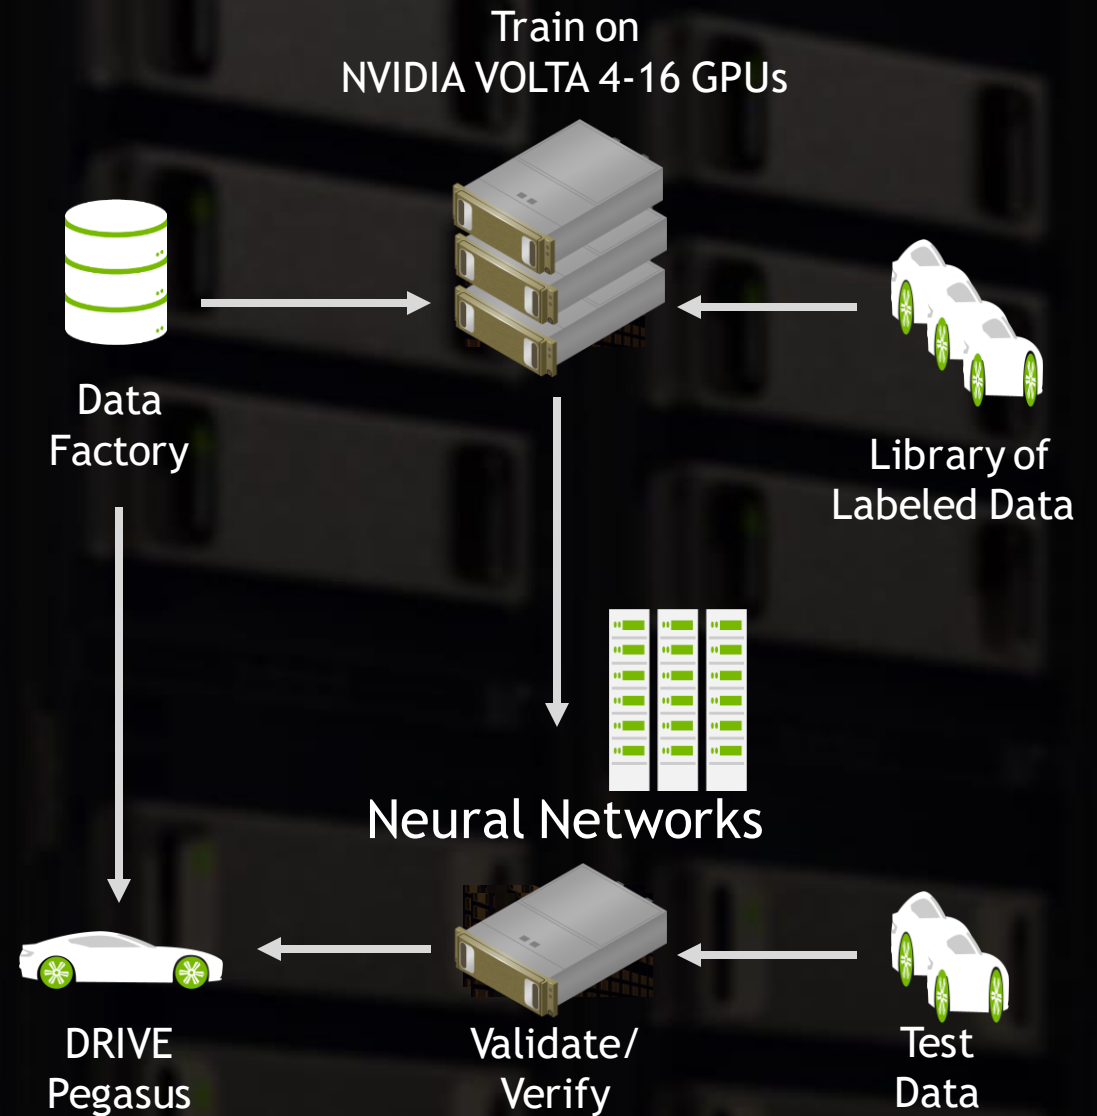


AUTOMOTIVE WORKFLOW

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



Example:

- Autonomous vehicle = 1TB / hr
 - Training sets up to 500 PB
 - RN50: 113 days to train
 - Objective: 7 days
 - 6 simultaneous developers
- = 97 node cluster

Rack Design



- DL drives close to operational limits
- Similarities to HPC best practices

Networking



- Ethernet or IB based fabric
- 100Gbps inter-connect
- High-bandwidth, ultra-low latency

Storage



- Datasets range from 10k's to millions objects
- terabyte levels of storage and up
- High IOPS, low latency

Facilities



- assume higher watts per-rack
- Higher FLOPS/watt = DC less floorspace required

Software



- Scale requires "cluster-aware" software

DELIVERING DATA SCIENCE VALUE



Maximized Productivity



Top Model Accuracy



Lowest TCO

Oak Ridge
National Labs

215x

Speedup Using RAPIDS
with XGBoost

Global
Retail Giant

\$1B

Potential Saving with
4% Error Rate Reduction

Streaming Media
Company

\$1.5M

Infrastructure
Cost Saving



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

CONOA

AI and Deep Learning Infrastructure Specialists

Fundamentals for AI


Your own data
is golden



Models and libraries



Compute power



Fastest way to find the right use case

AI owner



Gather stakeholders
from Business and IT



Do Workshop





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

