TRANSFORMING AI DEVELOPMENT FOR DATA SCIENTISTS

WHAT A DATA SCIENTIST LOOKS LIKE

Combining Mathematics, Coding, and Research

Curious and creative, enjoys problem solving, data manipulation, and data exploration

Strong interest in math and statistics

Proficient in procedural languages such as C++

Skilled with Python (and, to a lesser extent, R)

Needs: high performance GPU and software for fast data processing and continuous experimentation

~3M DATA SCIENTISTS AT WORK

CONSUMER INTERNET

Ad Personalization
Click Through Rate Optimization
Churn Reduction

FINANCIAL SERVICES

Claim Fraud
Customer Service Chatbots/Routing
Risk Evaluation

HEALTHCARE

Improve Clinical Care
Drive Operational Efficiency
Speed Up Drug Discovery

RETAIL

Supply Chain & Inventory Management
Price Management / Markdown Optimization
Promotion Prioritization And Ad Targeting

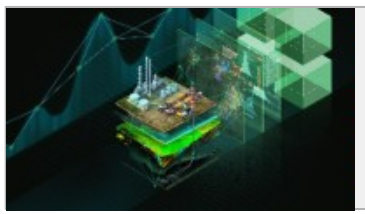

OIL & GAS

Sensor Data Tag Mapping
Anomaly Detection
Robust Fault Prediction

MANUFACTURING

Remaining Useful Life Estimation
Failure Prediction
Demand Forecasting

TELECOM

Detect Network/Security Anomalies
Forecasting Network Performance
Network Resource Optimization (SON)

AUTOMOTIVE

Personalization & Intelligent Customer Interactions
Connected Vehicle Predictive Maintenance
Forecasting, Demand, & Capacity Planning

DATA SCIENCE WORKFLOW

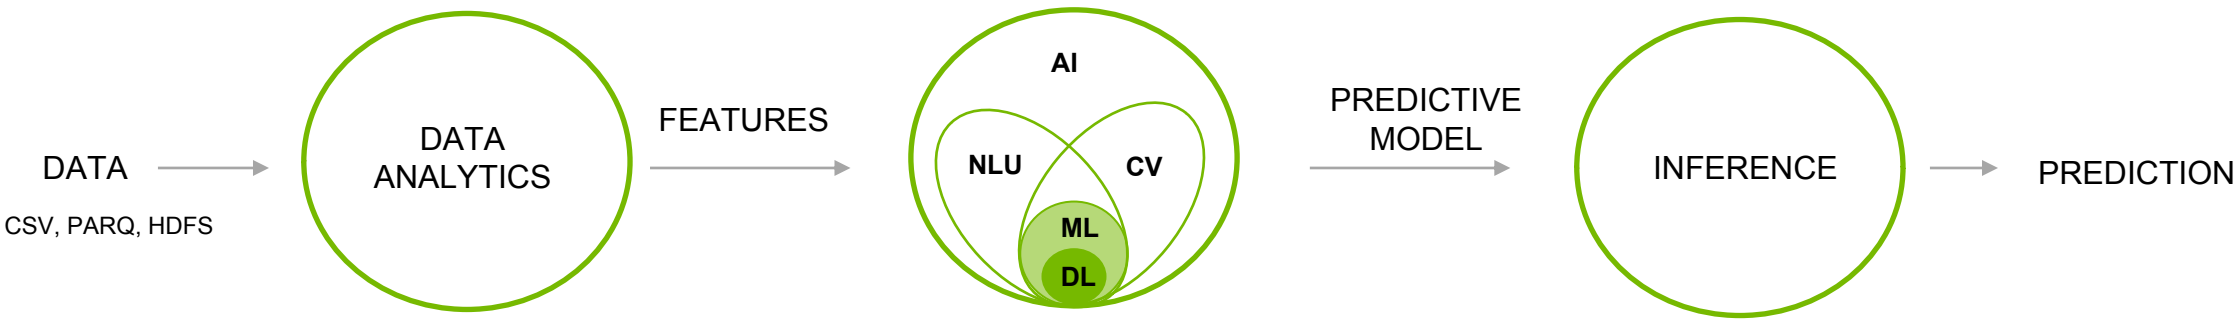

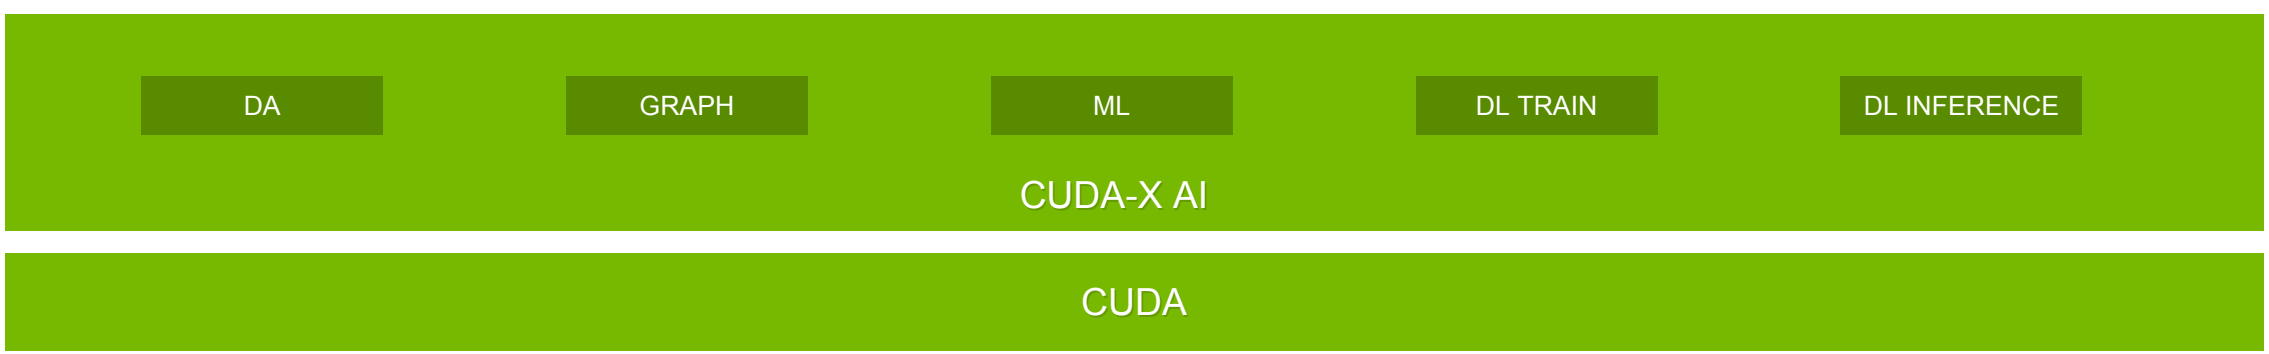

TRADITIONAL SETUP

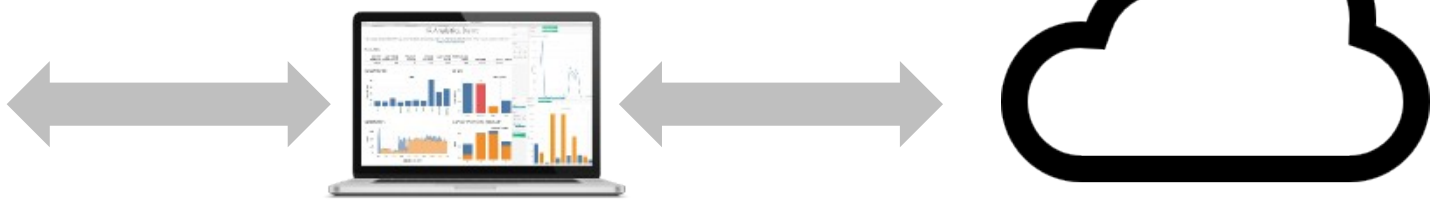

- Resource availability depends on job queues
- Slower model iteration process
- Software stack needs IT management and support

- Too expensive for everyday development
- Slow data migration
- Security and privacy concerns

DAY IN THE LIFE OF A DATA SCIENTIST

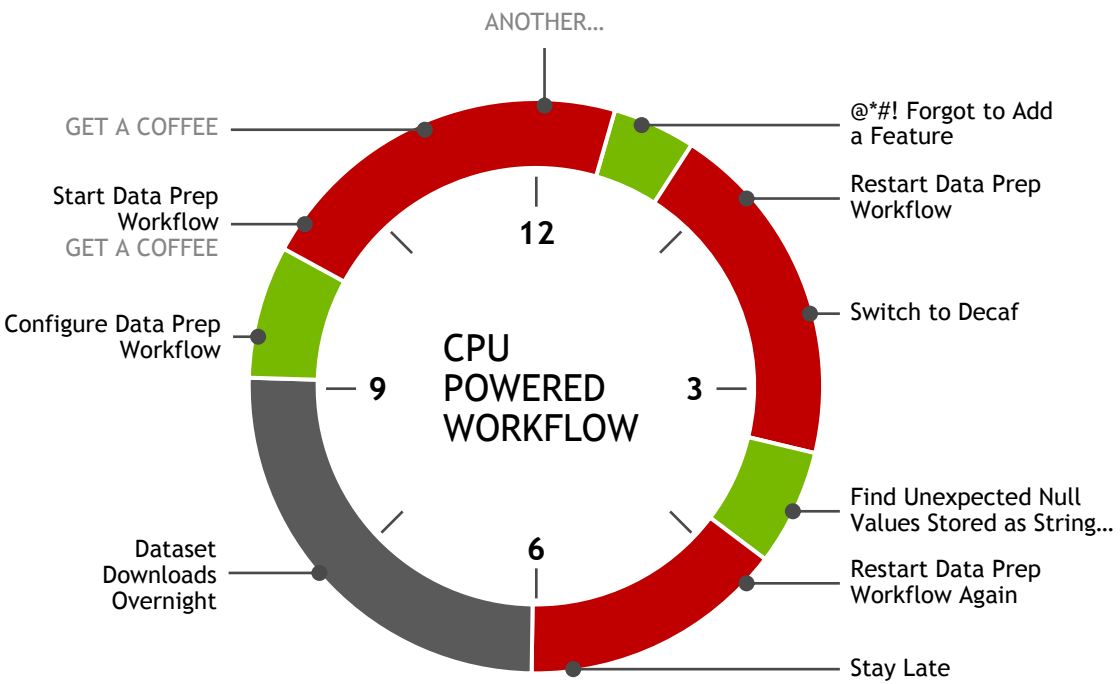

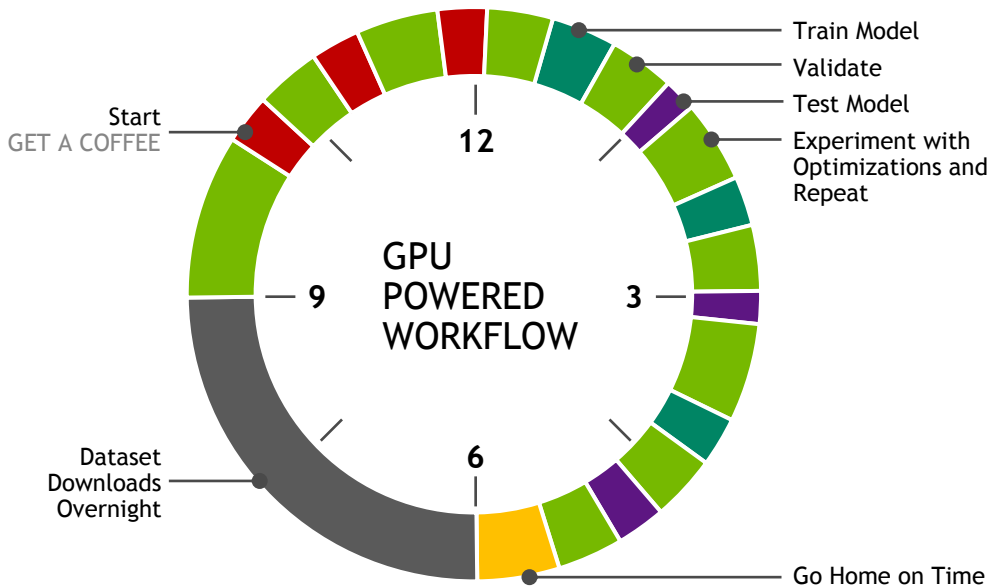

Dataset Collection | Analysis | Data Prep | Train | Inference

NVIDIA POWERED DATA SCIENCE WORKSTATION

Integrated hardware and software solution for Data Science

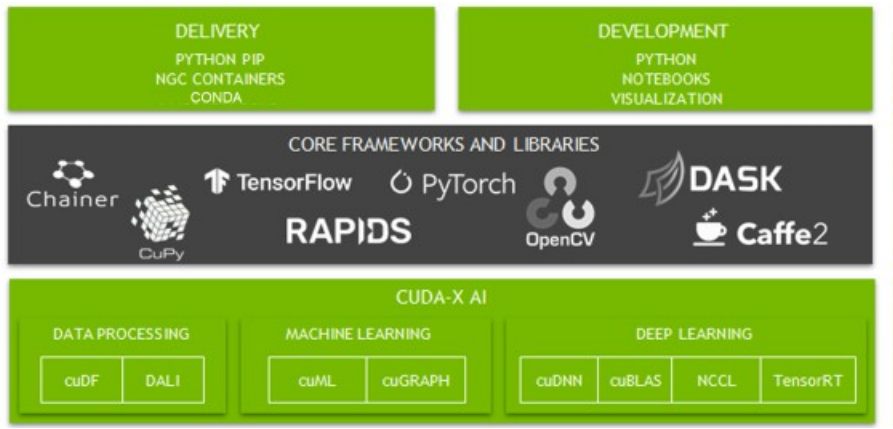

QUADRO RTX FOR DATA SCIENCE

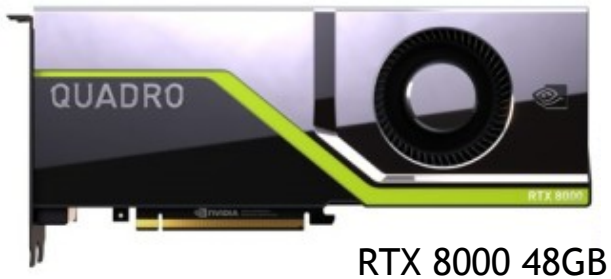

RTX 8000 48GB / 96GB w/NVLink

RTX 6000 24GB / 48GB w/NVLink

GPU Architecture	**Turing**
CUDA Cores	**4608**
RT Cores	**72**
Tensor Cores	**576**
Memory BW	**Up to 672 GB/s**
NVLink	**2-way (2 & 3slot)** **100 GB/s bidirectional**
Display Support	**4x DP + 1x VirtualLink**

GV100 32GB / 64GB
Double Precision (FP64)

CUDA-X AI DISTRIBUTION

Delivering Best User Experience with Validated Interoperability

arrow-cpp 0.12
cython 0.29
pyarrow 0.12
xgboost 0.8
nltk 3.2.5
numba 0.42
caffe-gpu 1.0.0
lime 0.1.1.32
anaconda-navigator 1.9.6
bokeh 1.0.4
cmake 3.12

dask 1.1.1
dask-core 1.1.1
dask-cuda 0.0.1
dask-cudf 0.0.1
dask-xgboost 0.1.5
distributed 1.25.3
faiss-gpu 1.5.0
ipywidgets 7.4.2
jupyterlab 0.35.4
numpy 1.15.4

opencv 3.4.2
pandas 0.23.4
pytorch 1.0.1
scikit-learn 0.20.2
scipy 1.2.0
statsmodels 0.9.0
torchvision 0.2.1
cffi 1.11.5
chainer 5.1.0
ipyvolume 0.5.1
pytest 4.3.0

python-graphviz 0.8.4
setuptools 40.8.0
cuda100 1.0.0
cudatoolkit 10.0.130
cudf 0.5.1
cuml 0.5.1
cupy-cuda100 5.2.0
nvstrings 0.2.0
matplotlib 3.0.2
python 3.7
matplotlib 3.0.2

(Almost) One Click Install: Assembled into one script, installed and tested on proven hardware

Highlighted libraries are GPU accelerated.

FASTER PERFORMANCE, REAL WORLD BENEFITS

Dataset: Mortgage Data 2015-2016
Seconds (lower is better)

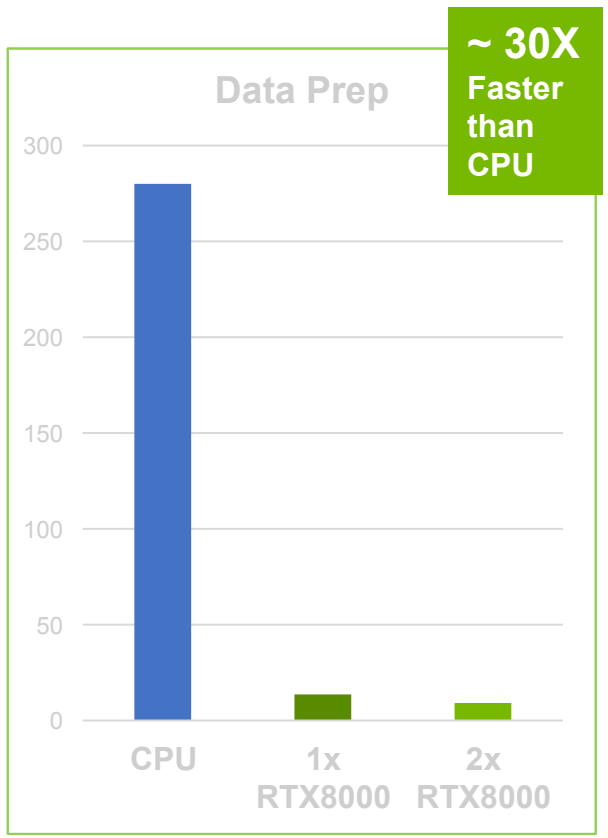

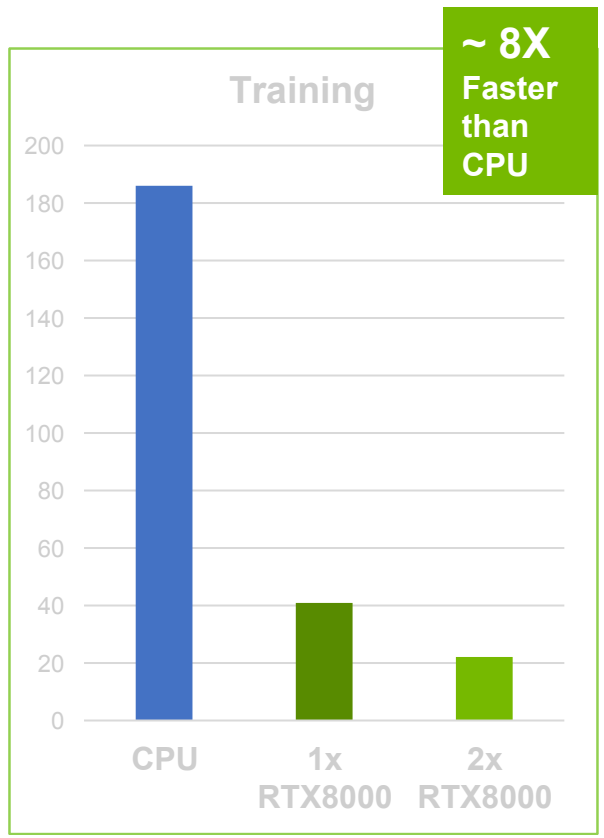

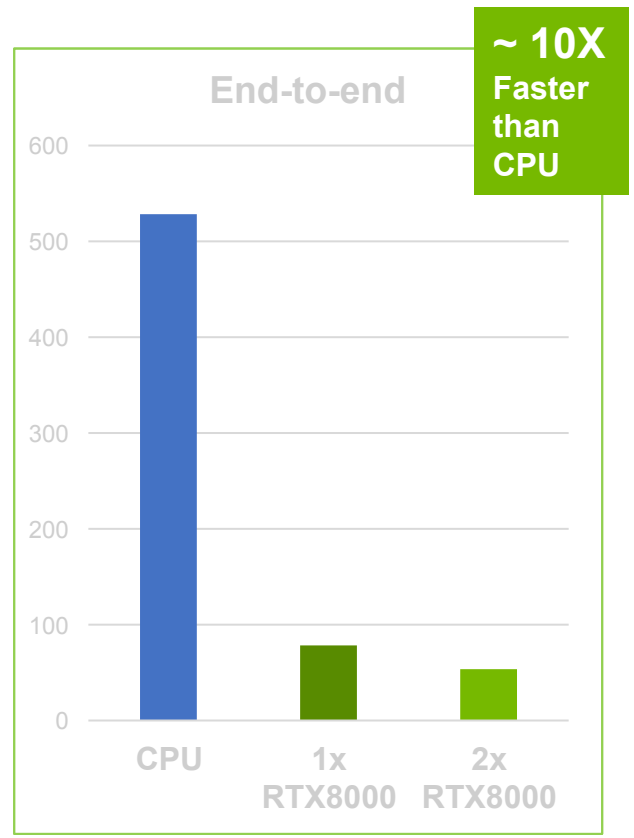

End-to-end time = Data Prep + Conversion + Training + Validation

CPU: dual Gold 6140@2.30GHz 3.7GHz Turbo (Skylake), ETL with Dask + Pandas

NVIDIA DATA SCIENCE WORKSTATION SOFTWARE STACK

How to Get Up and Running

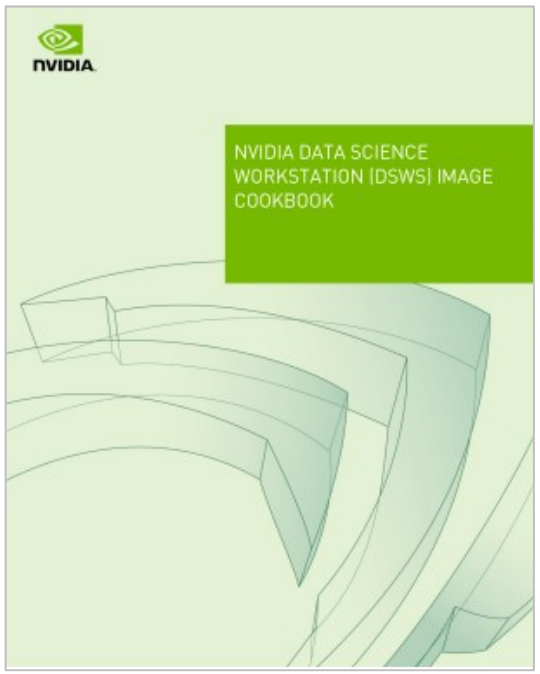

NVIDIA Data Science Workstation Cookbook

NVIDIA Data Science Workstation Cookbook

- Detailed instruction on how to set up your Data Science Workstation
- Information on OS, drivers, CUDA, & containers

NVIDIA GPU CLOUD (NGC)

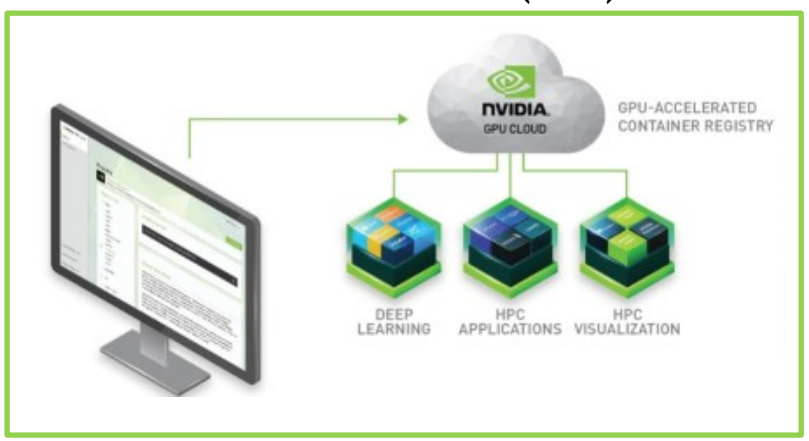

AI, HPC, HPC Visualization container repository

NVIDIA NGC SUPPORT SERVICES

Minimize Downtime And Maximize System Utilization

Support Coverage

- NGC DL & ML containers
- NVIDIA drivers
- NV-docker
- CUDA

Support by NVIDIA's subject matter experts

24x7 portal, phone and email access to create support cases

Live support during local, regional business hours for technical assistance

Availability

- Exclusively for NGC-Ready workstations
- Availability starting in Q2
- Service agreement between NVIDIA & customer
- Purchase from OEM

DELIVERING DATA SCIENCE VALUE

Maximized Productivity

Highly optimized cross-compatible stack of data science libraries

Faster model design, development and iteration

Greater flexibility using Python and conda package management

Ease of Use

Turnkey system for GPU accelerated data science

End-to-End software stack acceleration from data preparation to visualization

Orchestration compatible software stack to help scale on clusters

Enterprise Support

Built for enterprise level reliability and robustness

Quick and easy deployment using “NGC-Ready” containers and conda

Tested across GPUs and systems for compatibility and performance

CUSTOMER EXPERIENCES

LOCKHEED MARTIN

"The NVIDIA-powered data science workstation enables our data scientists to run end-to-end data processing pipelines on large data sets faster than ever. Leveraging RAPIDS to push more of the data processing pipeline to the GPU reduces model development time which leads to faster deployment and business insights."

-Mike Koelemay, Lockheed Martin Fellow

MIT

"The MIT FAST Lab uses the RTX 8000 GPUs for both graphics and deep learning. The RTX 8000 GPUs provide significant capabilities for training deep neural networks for robot perception, especially through increasing batch size thanks to very large VRAM. By switching to RTX 8000, our team was able to accelerate learning by using large batch sizes."

- Sertac Karman, Associate Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology

Raytheon

"Raytheon technology defends thousands of miles of borders for nations around the world, and much of that work requires monitoring huge expanses of land and open ocean for weapon smuggling. By using the NVIDIA-powered Data Science Workstation, we've been able to reduce the time it takes to develop machine learning models that detect weapons in imagery. We've drastically reduced development time from days to hours for our data scientists. That means less time spent sifting through data, and more time stopping weapons from getting into the wrong hands."

– Dr. Shane Zabel, Chief Artificial Intelligence Officer, Raytheon Intelligence, Information and Services

ARUP

"We have a diverse, multi-disciplinary environment and are looking to couple data science and analytics to a wider range of our technical practices throughout our business. The NVIDIA-powered Data Science Workstation promises to ease the transition and democratize the application of data science. We find it extremely well suited to experimentation, exploration, solution discovery, and early prototyping work. Its combination of well-designed software and highly performant hardware provides a 20x and higher speed-ups in our analytics work and our team found its ease of use liberating."

-Steve Walker, Associate Director, Arup, Advanced Digital Engineering

RESOURCES

NVIDIA-POWERED DATA SCIENCE WORKSTATION WEB PAGE

RAPIDS WEB PAGE

PRODUCT DATASHEETS

DATA SCIENCE WORKSTATION SOLUTION BRIEF

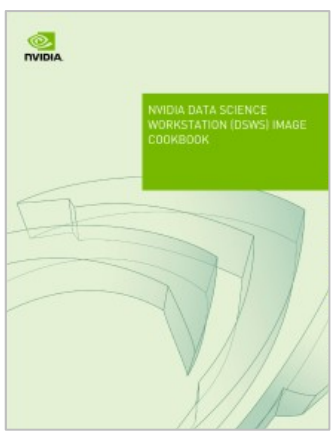

DATA SCIENCE WORKSTATION COOKBOOK

PRODUCT DECK

WORKSTATION FOR DATA SCIENTISTS

POWERED BY NVIDIA GPU AND CUDA-X AI

- Dual Quadro GPUs with up to 96 GB Memory
- Pre-installed for CUDA-X Accelerated Data Science – RAPIDS, TensorFLow, PyTorch, Caffe Anaconda Distribution
- 10X Faster


NVIDIA