

RTX: BRINGING AI & ADVANCED GRAPHICS TO VISUAL COMPUTING

RAJ MIRPURI - VP PROFESSIONAL VISUALIZATION




NVIDIA ACCELERATES ALL WORKLOADS

SCIENTIFIC
RESEARCH





DEEP
LEARNING



RENDERING
& VIZ



MACHINE
LEARNING



CUDA

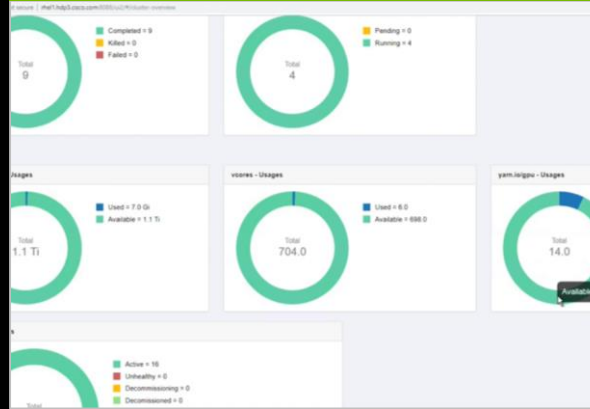
GPU

VISUAL COMPUTING CHALLENGES TODAY

Increased Complexity & Quality




Increased Demand for Data Analysis



Massive Data, Real-time Insight




Need for Mixed Visualization & Compute



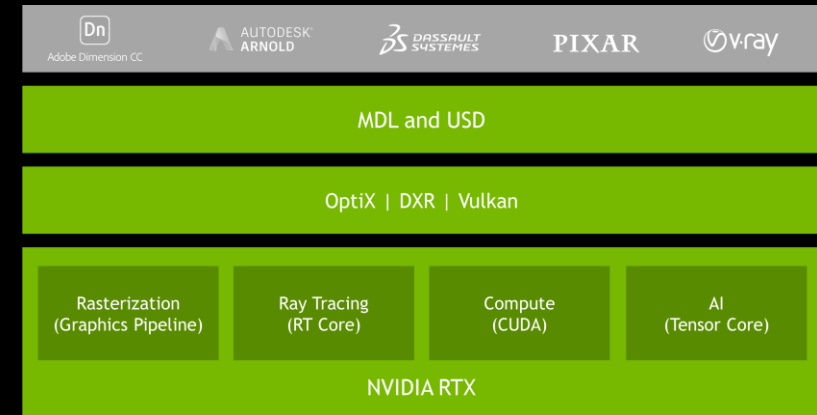
NVIDIA POWERS THE NEXT ERA OF COMPUTING

Turing RTX Architecture



Accelerated offline rendering
Data Science platform
AI at the Edge
Powerful virtual workstations

RTX Platform



RTX SERVER ENABLES AI FOR VISUAL COMPUTING

Rendering



Accelerate offline batch rendering with the power of multi-GPU acceleration

Data Science



Enable multiple GPU configurations to develop on larger data sizes to find business impacting and changing results

AI at the Edge



Unlock benefits as bandwidth, latency and availability to resources are not ubiquitous and plentiful

Virtualization and Mixed Workloads




Empower designers and artists anywhere in the world with the high end RTX GPUs

The background is a dark, almost black, space filled with a network of thin, light green lines. These lines connect various points, some of which are highlighted with small, glowing green circles. The overall effect is a complex, web-like structure that suggests a digital or technological theme. The text 'RTX AI FOR RENDERING' is centered in the middle of the image, with a thin green line extending from the left side of the text towards the left edge of the frame.

RTX AI FOR RENDERING

AN EXPLOSION OF RICH CONTENT



1993: Jurassic Park had a total of 63 VFX shots taking a year to create

2018: Marvel's Avengers: End Game had more than 3,000 shots

fxguide

articles podcasts fxinsider forums contact


RenderMan's Visuals for Coco

By Mike Seymour
March 3, 2018

The shot had over 8 million instanced lights, which initially took 1,000 hours to render each frame

By the end of production, it still took 50 hours to render each fully optimized frame

CREATE MORE, WAIT LESS WITH QUADRO



Note: CPU = i9-7900X; GPU = NVIDIA RTX GPU; video playback at 2x speed

EXPONENTIAL POWER AT 1/4 THE COST



240 Dual 12-core Skylake CPU Servers
144 kW
\$2M Render Farm



4 RTX 8-GPU Servers
13 kW
\$500,000
1/4 the Cost 1/10 the Space 1/11 the Power

Interactive Content Creation

Batch Rendering

RTX-Accelerated Rendering

25X ACCELERATED RENDERING FOR NETFLIX



NETFLIX Lost In Space scene: renders in a fraction of the time using RTX Server

- 6x faster for a single frame
- 25x faster for the entire shot

	CPU Node (Dual Skylake)	RTX Server (4 x RTX 8000)	Improvement
Render time (1 frame)	38 min	6 min	6x
Total render time (120 frames)	76 hours	3 hours	25x
# of nodes	25	1	25x
Power (kW)	13.2	1.9	7x
Acquisition cost	\$188k	\$28k	7x
Cost of power (5 yrs)	\$68k	\$10k	7x
Total cost	\$256k	\$38k	7x


RTX FOR DESIGN AND RENDERING

Accelerated Workflows with RTX

DESIGN



RENDERING



OpenGL

DirectX

Vulkan

OptiX

IndeX

VXGI

VRWorks

PhysX

MDL

vMaterials

GVDB

DESIGNWORKS

NvPro

Video Codec SDK

NVIDIA
QUADRO

RTX SERVER

NVIDIA
GRID

NVIDIA
VRREADY

DESKTOP | MOBILE WORKSTATION



DATACENTER




CLOUD



The background is a dark, almost black, field filled with a complex network of thin, light green lines. These lines connect various points, some of which are highlighted as small, glowing green dots. The overall effect is reminiscent of a data network or a neural network diagram. In the upper left quadrant, there is a small, thin, light green L-shaped line that frames the top-left corner of the text area.


RTX AI FOR DATA SCIENCE

3M DATA SCIENTISTS AT WORK TODAY



CONSUMER INTERNET

Ad Personalization
Click Through Rate Optimization
Churn Reduction




OIL & GAS

Sensor Data Tag Mapping
Anomaly Detection
Robust Fault Prediction



FINANCIAL SERVICES

Claim Fraud
Customer Service Chatbots/Routing
Risk Evaluation



MANUFACTURING

Remaining Useful Life Estimation
Failure Prediction
Demand Forecasting



HEALTHCARE

Improve Clinical Care
Drive Operational Efficiency
Speed Up Drug Discovery



TELECOM

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



RETAIL



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



AUTOMOTIVE

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

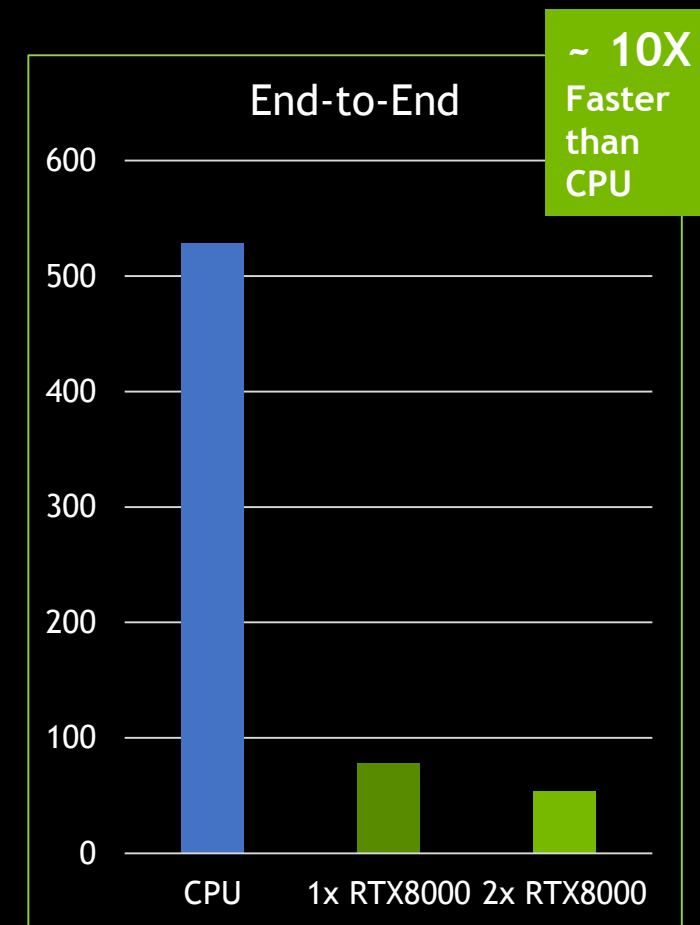
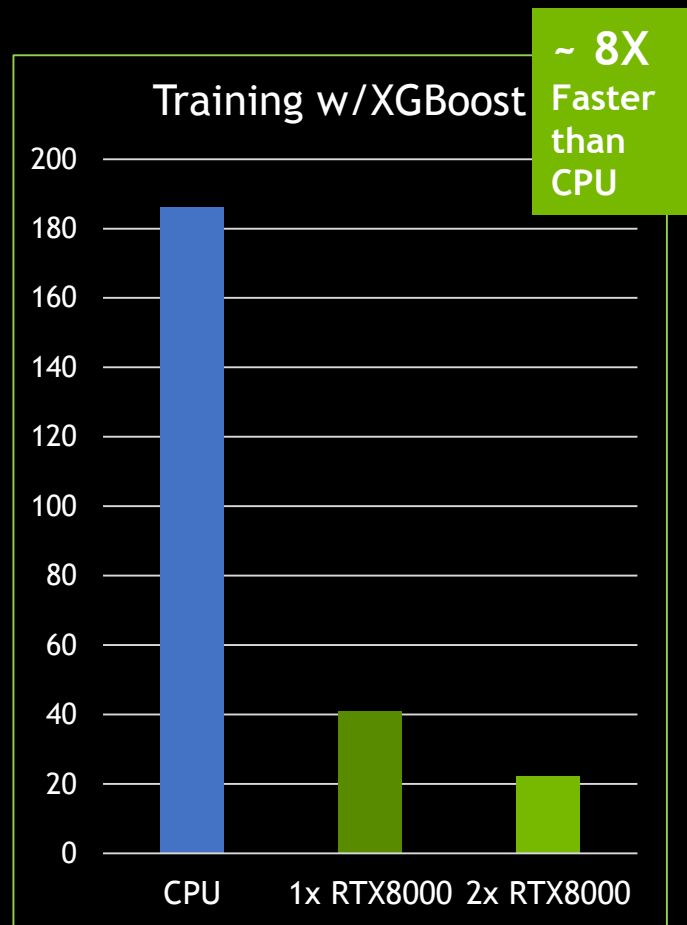

RTX STREAMLINES DATA SCIENCE WORKFLOWS



Dataset Collection Analysis Data Prep Train Inference

10X FASTER PERFORMANCE WITH RTX

Seconds (lower is better)



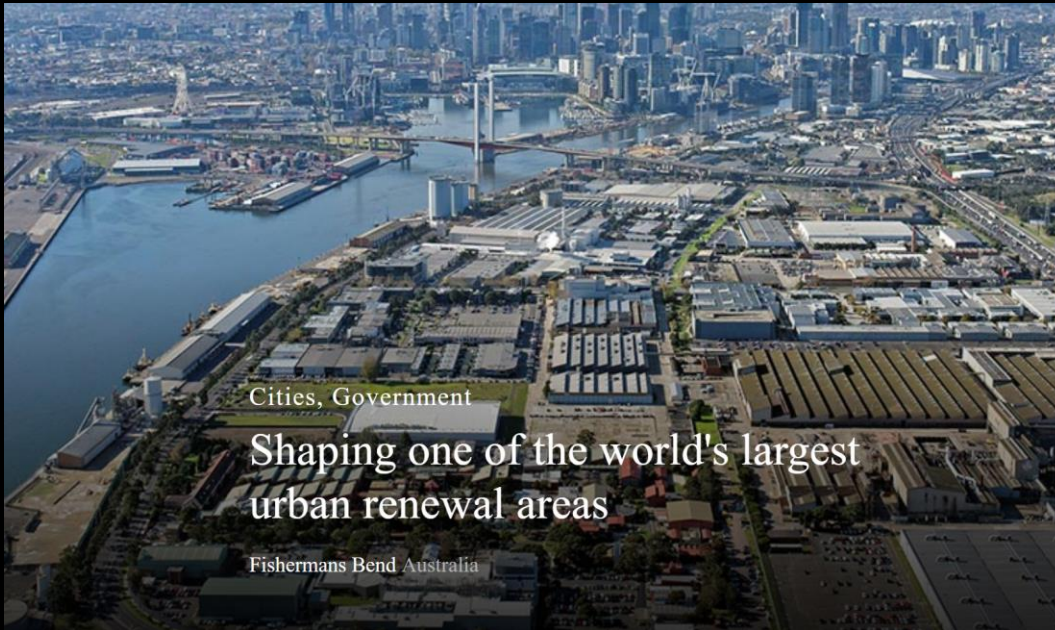
20X SPEEDUP FOR ARUP ANALYTICS

Opportunity: 200 Data Scientist

Use Case: Distributed infrastructure assets, precipitation risk assessment

Data Preparation: 50,000 CSV files to extract data and then aggregate in multiple formats

Model Training: Open source python for analytics, pandas, GeoPandas, Shapely, NumPy



Cities, Government

Shaping one of the world's largest urban renewal areas

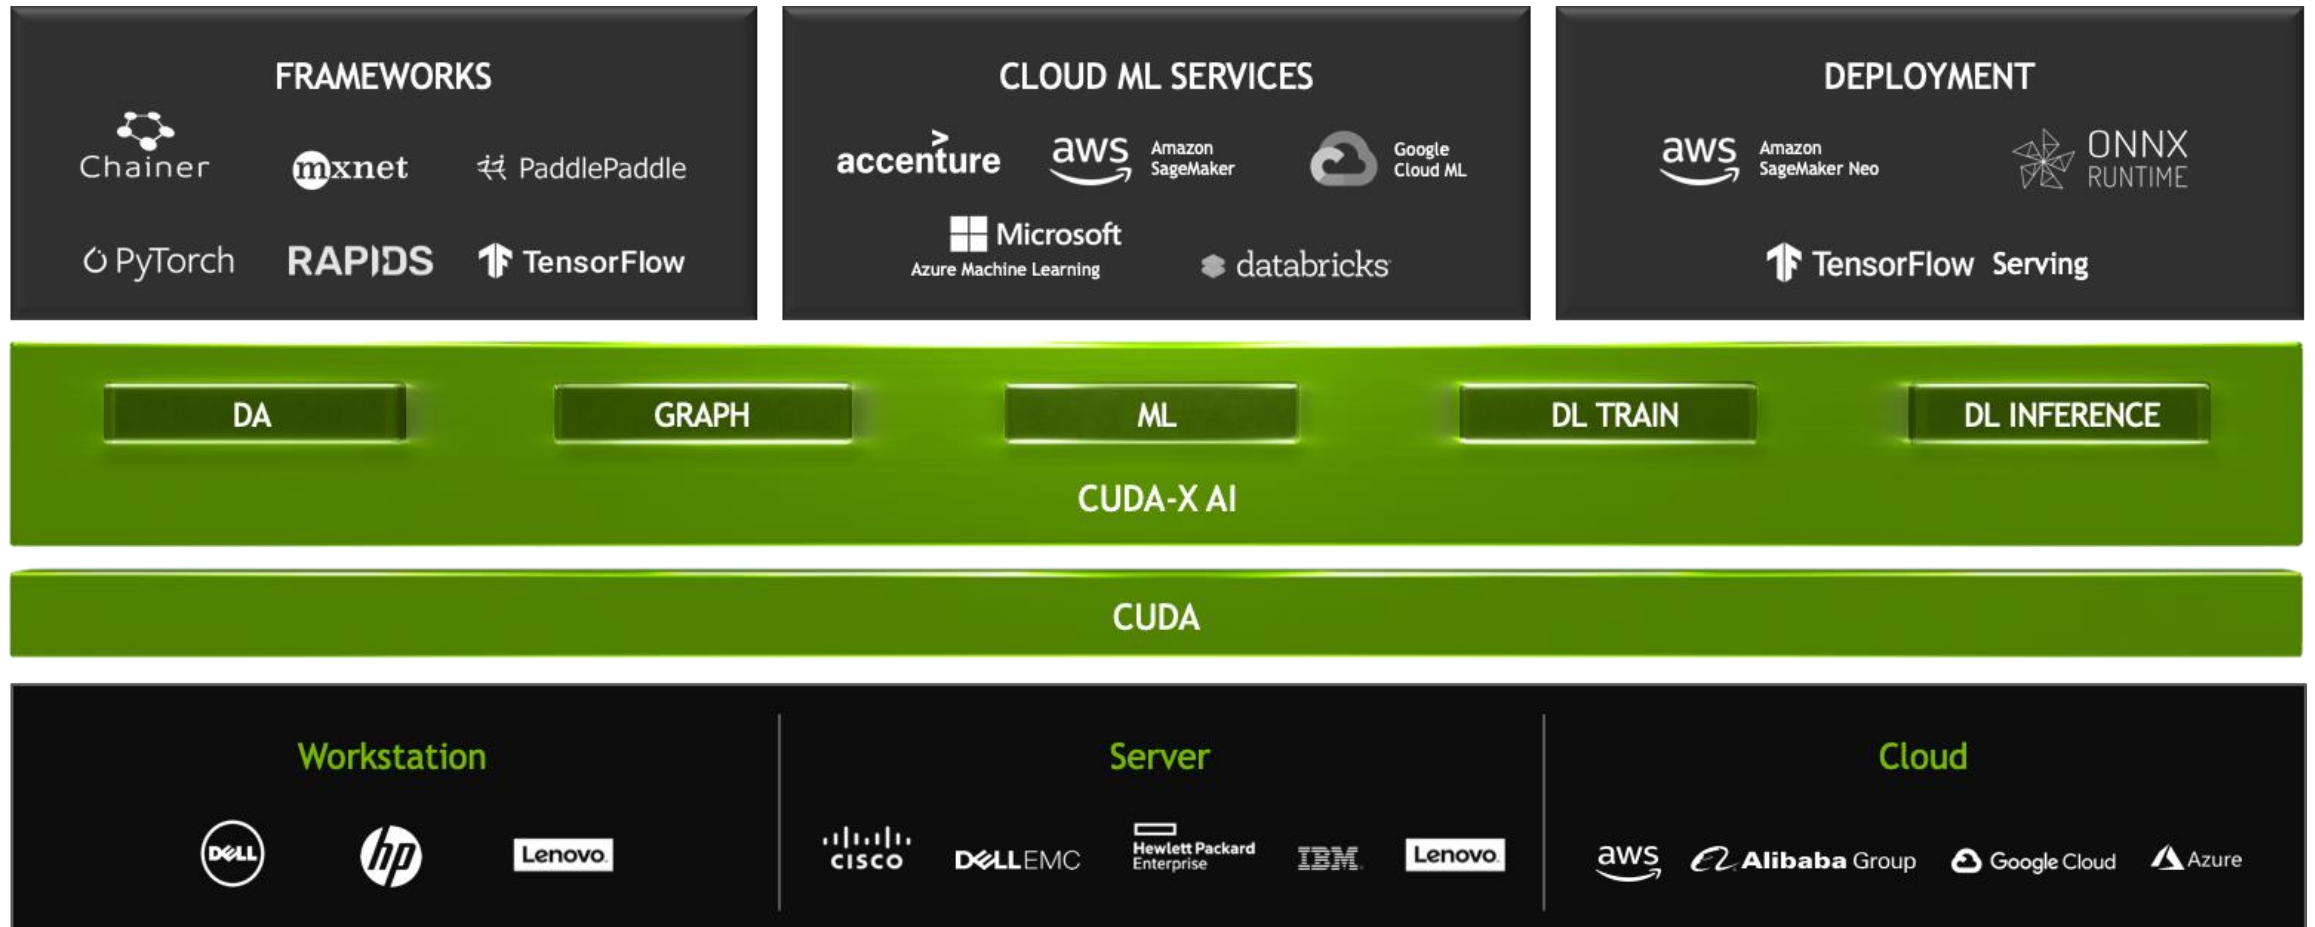
Fishermans Bend Australia

*“The NVIDIA-powered Data Science Workstation[’s]... 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

RTX FOR DATA SCIENCE & AI


From Data Science to NVIDIA Accelerated Data Science with CUDA-X AI



The background is a dark, almost black, field filled with a complex network of thin, light green lines. These lines connect various points, some of which are highlighted as small, glowing green dots. The overall effect is reminiscent of a neural network or a data visualization of interconnected nodes. The text 'RTX AI AT THE EDGE' is positioned on the left side, partially enclosed by a thin green L-shaped line that extends from the top and left edges.

RTX AI AT THE EDGE


SMART AND SAFE CITIES NEED AI



1 billion installed security cameras WW (2020)
30 billion frames per day



Challenging real world conditions
Traditional video analytics not trustworthy



AI achieves super human results
AI driven intelligent video analytics

RTX ENABLES AI AT THE EDGE

EXTRACTING NEW VALUE



VALOSSA

SAFETY COMPLIANCE




SMARTVID.IO

RETAIL STORE



DeepBlue

VIDEO SECURITY & PUBLIC SAFETY




NVIDIA
QUADRO

















REAL-TIME INSIGHTS TO FIND RESOLUTION




Finding lost people
in 30 minutes



Capturing a criminal
before he strikes again



Adding 100 virtual
security guards



Improving traffic
safety




Reducing peak traffic
congestion by 15%

The background is a dark, almost black, space filled with a network of thin, light green lines. These lines connect various points, some of which are highlighted with small, glowing green circles. The overall effect is reminiscent of a neural network or a complex data structure. A thin, light green L-shaped line is positioned to the left of the main text, consisting of a horizontal segment and a vertical segment that meet at a right angle.


RTX AI FOR MIXED WORKLOADS

WORKFLOWS DEMAND MIXED WORKLOADS

VISUALIZATION



Windows 10



2D EDA




3D Apps

COMPUTE




Deep Learning



HPC

MIXED WORKLOADS ARE EVERYWHERE


Professional workflows are getting more demanding




Manufacturing



Architecture



Medical Imaging




Energy



Media & Entertainment

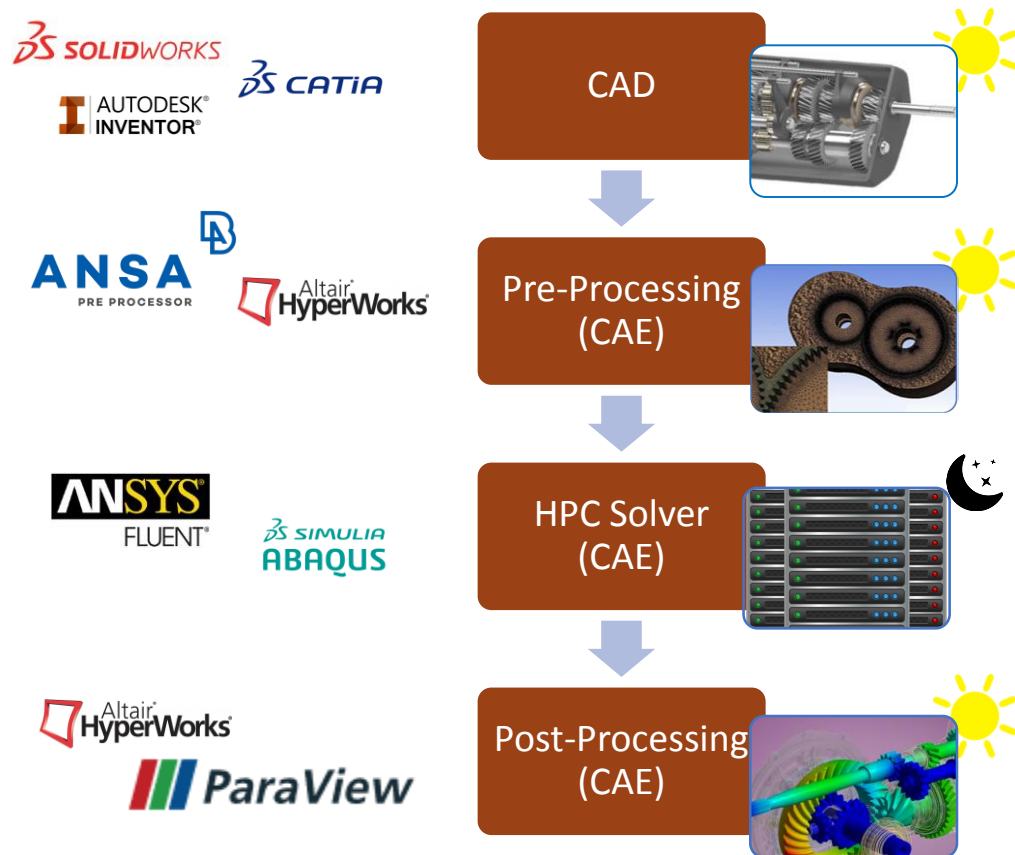


Automotive



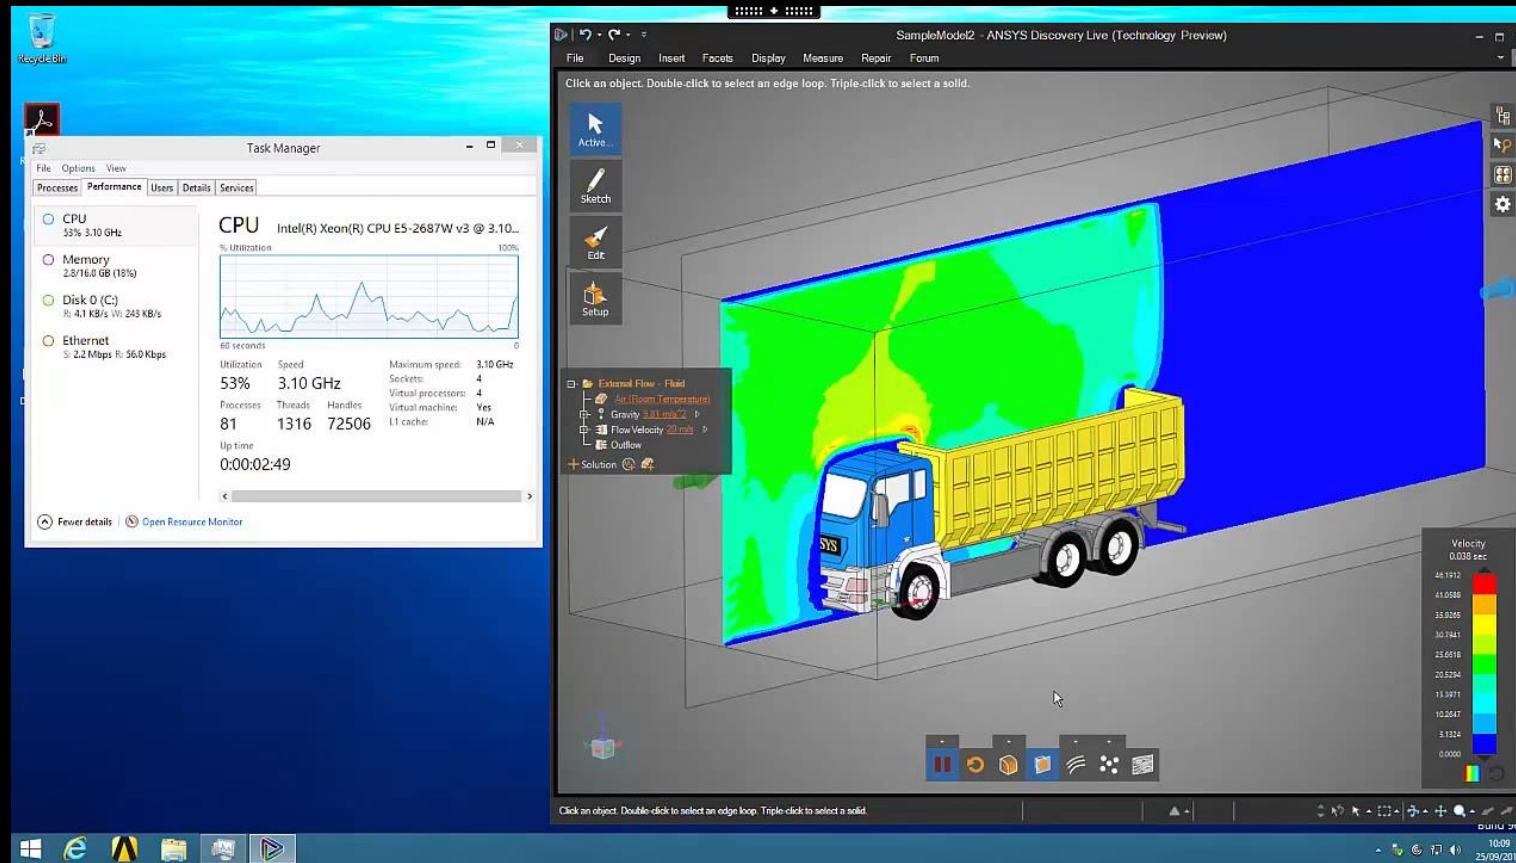
Defense

COMPUTER AIDED ENGINEERING WORKFLOW



Addressing the Bottlenecks of FEA Simulation
Source: Tech-Clarity 2017

SIMULATION WITHIN DESIGN POWERED BY RTX



QUADRO VIRTUAL WORKSTATION STREAMS FROM RTX SERVER

To Any Device, Anywhere

NVIDIA Quadro Virtual Data Center Workstation Software

- Now supports RTX 6000 and RTX 8000
- Brings world's most powerful virtual workstation to RTX Server platform
- Flexibly provision virtual workstations or a combination of virtual workstations and render nodes from a single RTX Server
- Extends power of RTX platform to designers, on any device, anywhere



QUADRO RTX VIRTUAL WORKSTATIONS

Type of User	Light Users	Medium Users	Heavy Users
Recommended Solution	NVIDIA T4 or P6 Quadro Virtual Data Center Workstation	NVIDIA T4 or P6 Quadro vDWS	NVIDIA Quadro RTX 8000, RTX 6000, P40 or V100 with Quadro vDWS
GPU Memory	16 GB	16 GB	48 GB/32 GB/24 GB
Equivalent Performance	Multiple Quadro P1000	Up to Quadro P4000	Up to Quadro RTX 8000
Replaces	K2, M60, P4, M6	K2, M60, P4, M6	N/A
			NEW TO VDI

RTX POWERS AI FOR VISUAL COMPUTING




Accelerate offline batch rendering with the power of multi-GPU acceleration



Enable multiple GPU configurations to develop on larger data sizes to find business impacting and changing results



Unlock benefits as bandwidth, latency and availability to resources are not ubiquitous and plentiful



Empower designers and artists anywhere in the world with the high end RTX GPUs



