플랜트조선 업계을 위한 GPU 컴퓨팅

NVIDIA 서완석이사 wseo@nvidia.com

Design &
Manufacturing

HPC and Big Data

VDI

NVIDIA

THE VISUAL COMPUTING COMPANY

NVIDIA QUADRO K6000

12GB

World’s most powerful graphics

World’s largest graphics memory

Dassault Systems SolidWorks 2013 SpecAPC

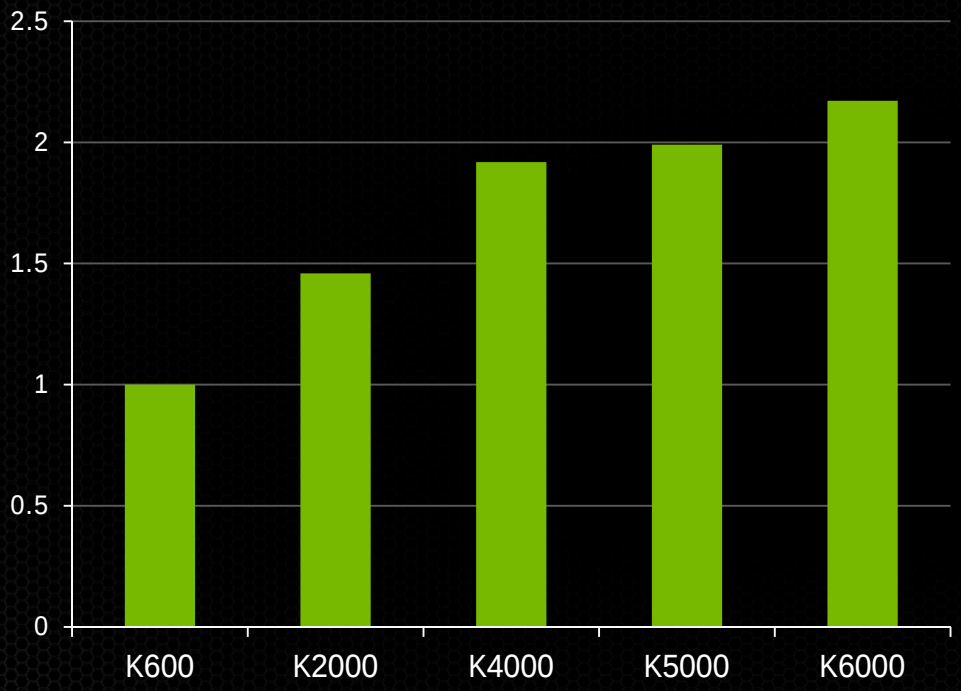

SolidWorks SpecAPC Graphics Composite score, X79(Supermicro) Motherboard, SINGLE XEON E5 2687W (3.10GHz), 32GB RAM, Win7-64(SP1), driver 320.61

PTC Creo 2.0

Target Customer

Large enterprise and SMB in all manufacturing industries, including automotive OEMs and tier 1 and 2 suppliers, Mechanical design engineers

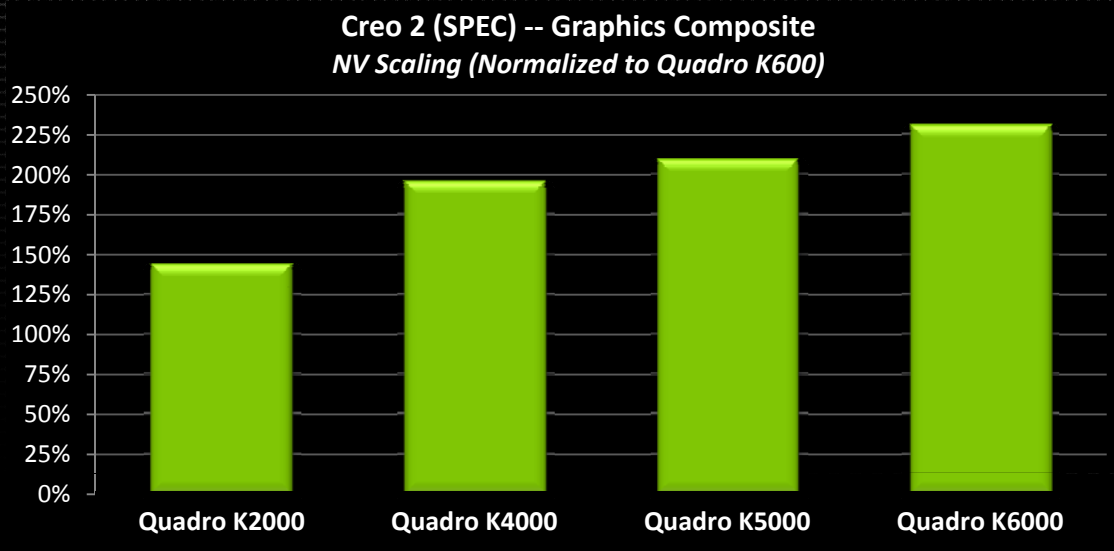

X79(Supermicro) Motherboard. SINGLE XEON E5 2687W (3.10GHz). 32GB RAM. Win7-64(SP1)

Value

- Fast, Interactive Assemblies with Creo Parametric 2.0
- Better clarity while designing with *anti-aliasing* support in Creo Parametric 2.0

What's new

Creo 2.0 graphics sub-system redesigned to take advantage of GPU performance scaling

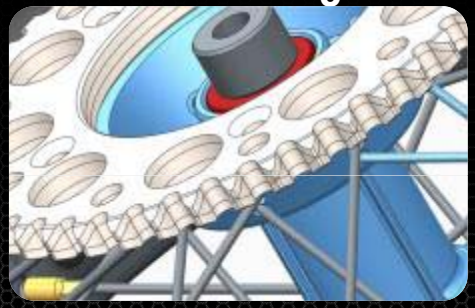

Tesla K40
World's Fastest Accelerator

FASTER
1.4 TF| 2880 Cores | 288 GB/s

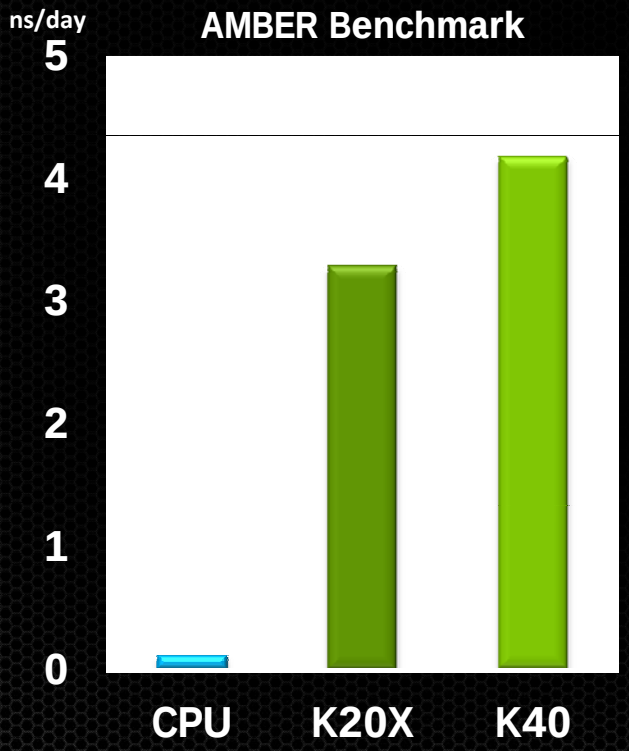

LARGER
2x Memory Enables More Apps

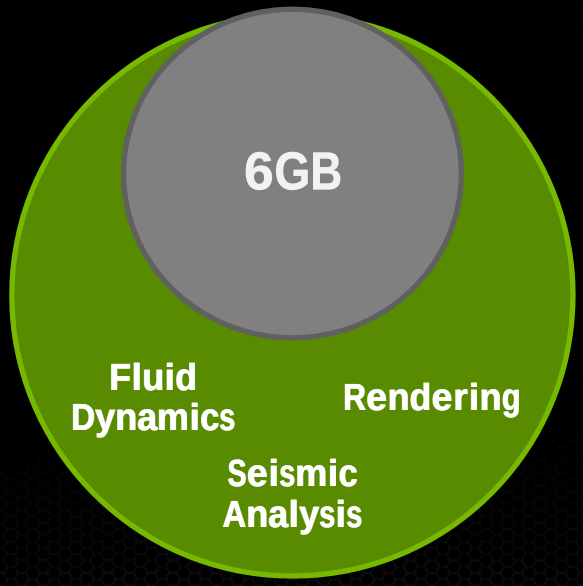

12GB

SMARTER
Unlock Extra Performance
Using Power Headroom

GPU Boost

Abaqus 6.13-DEV Scaling on Tesla GPU Cluster

(Shipbuilding/Energy sector)

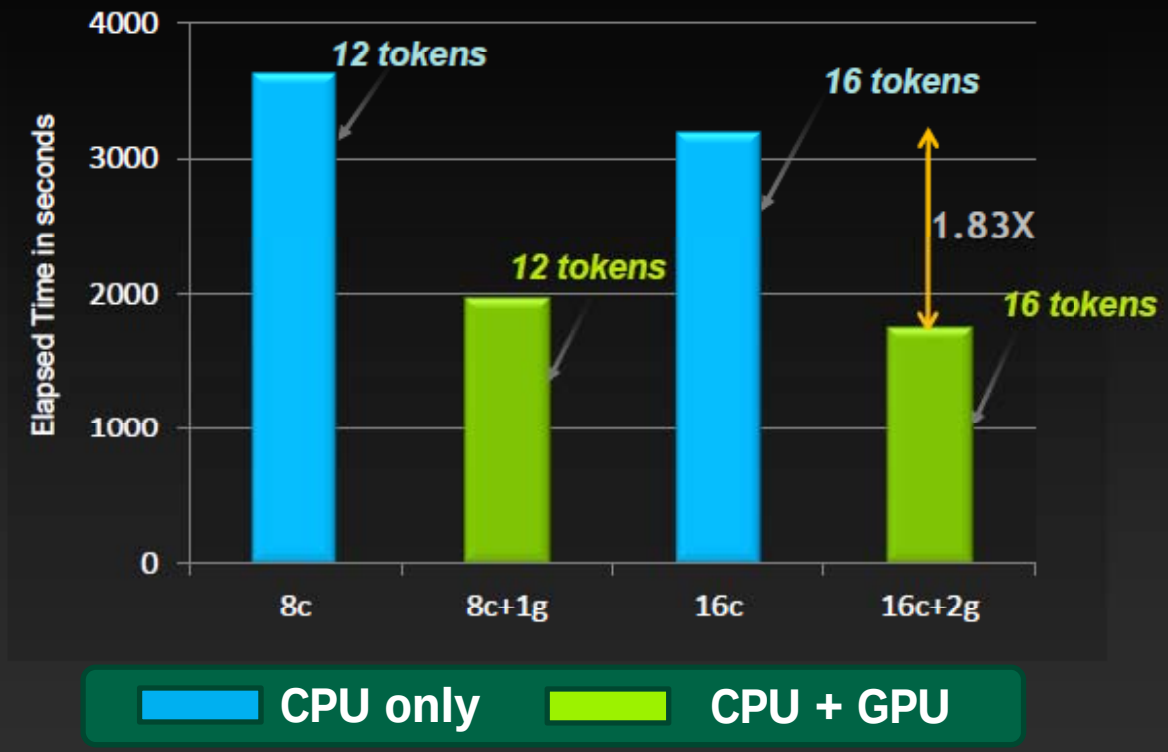

Dynamic (shock) analysis

- 1.1M equations; 8 TFLOPs
- Underwater explosion
- Sparse solver (un-symmetric)

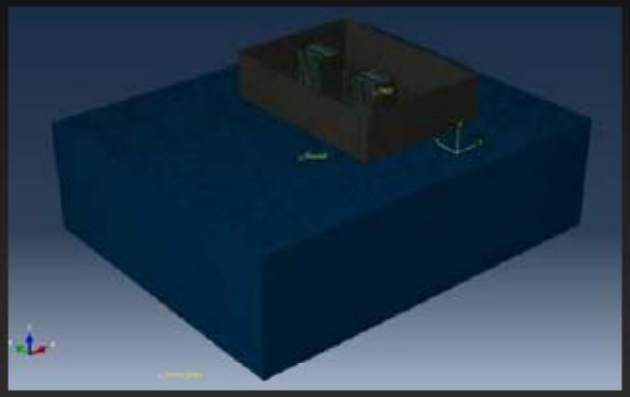

PSG Cluster: Sandy Bridge CPUs with (2x E5-2670 (8-core), 2.6 GHz), 128GB memory, 2x Tesla K20X, Linux RHEL 6.2, QDR IB, CUDA 5

Multi-GPU Support for ANSYS Fluent 15.0

Available 2013.12

ANSYS Fluent 15.0 -111 M aerodynamic problem

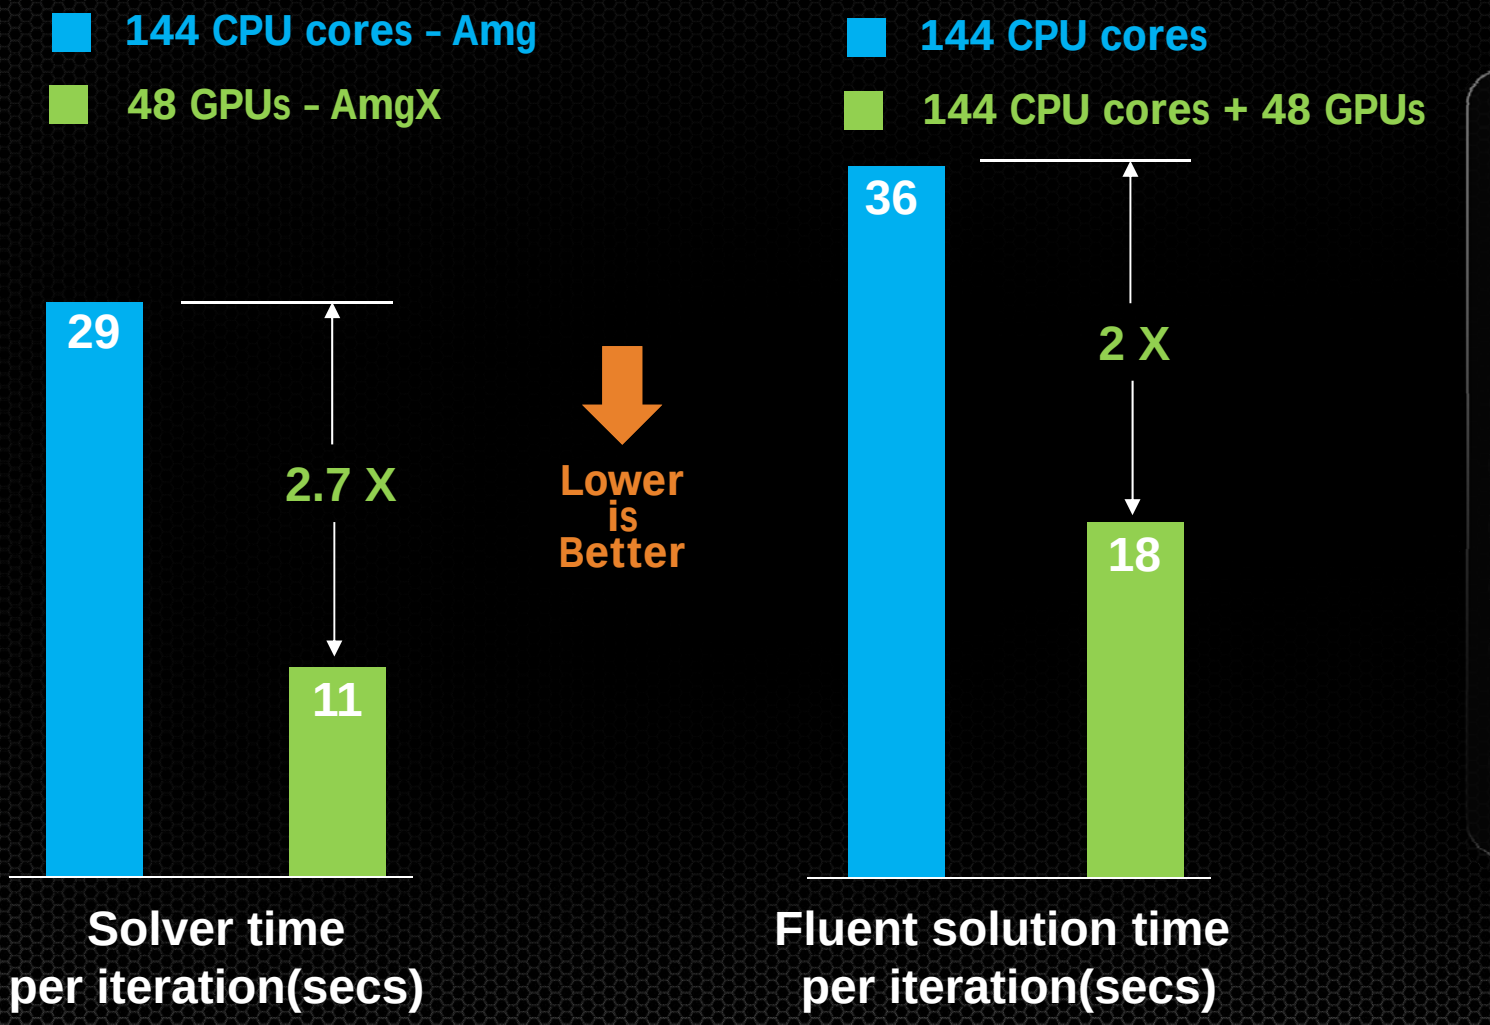

Truck Body Model

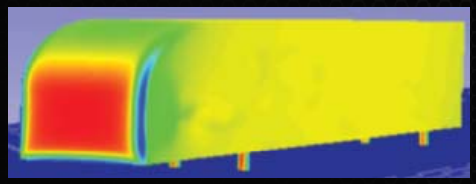

- 111 M Mixed cells
- External aerodynamics
- Steady, k-ε turbulence
- Double-precision solver
- CPU: Intel Xeon SB; 12 cores per node
- GPU: Tesla K40, 4 per node

GPU Virtualization
with NVIDIA GRID

DESIGNER

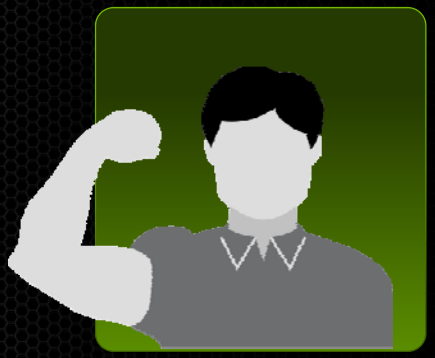

POWER USER

KNOWLEDGE WORKER

NVIDIA GRID K1

NVIDIA GRID K2

	NVIDIA GRID K1	NVIDIA GRID K2
GPU	4 Kepler GPUs	High End Kepler GPUs
CUDA Cores	768 (192/GPU)	072 (1536/GPU)
Memory Size	16GB DDR3 (4GB/GPU)	GB GDDR5 (4GB/GPU)
Max Power	130 W	225 W
Equivalent Quadro with Pass-through	Quadro K600 (entry)	Quadro K5000 (high end)

[1] Number of users depends on software solution, workload, and screen resolution

DELL | NVIDIA

Graphics Options in Virtualization

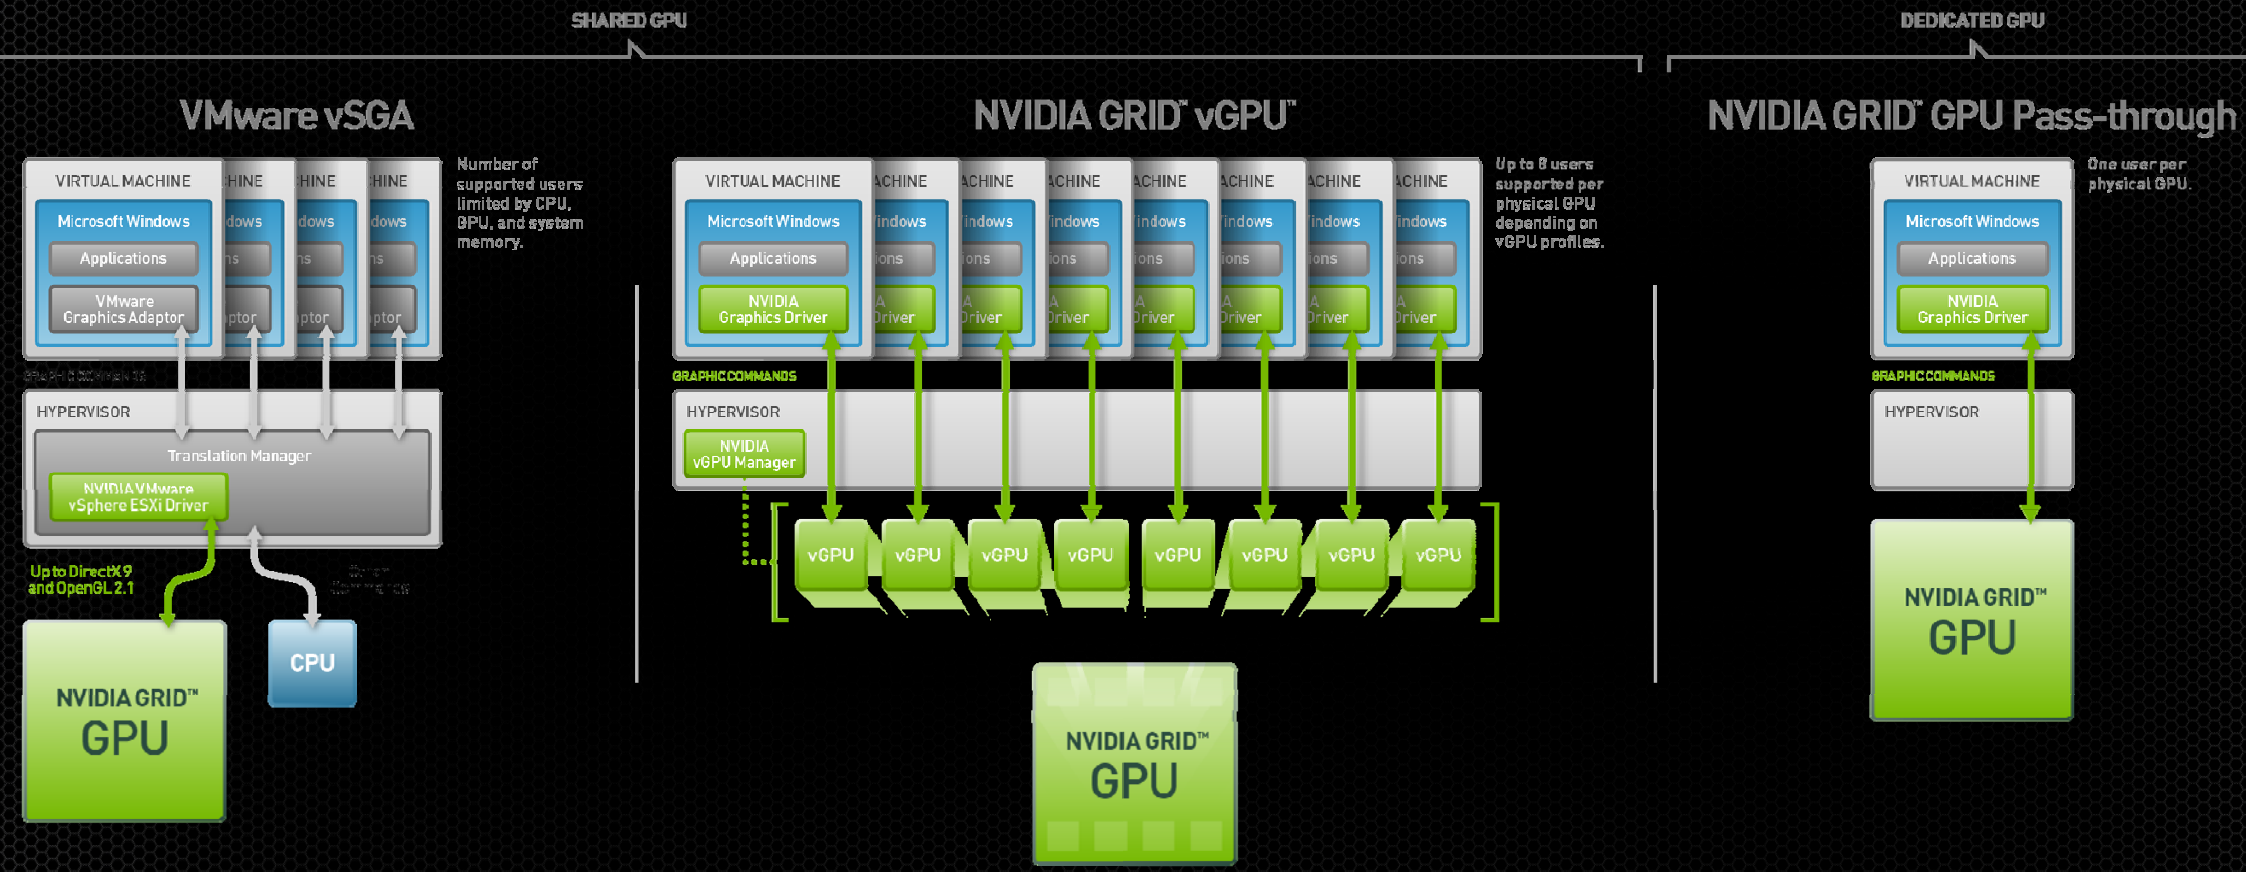

NVIDIA vGPU

GRID boards support several different logical GPU Types

- e.g. GRID K2 supports K200, K240Q, K260Q vGPUs
- vGPU types differ in amount of framebuffer, virtual display heads, max resolution, and number of simultaneous instances

Each physical GPU on a GRID board can host at most one vGPU type at a time

Card	Physical GPUs	Virtual GPU	Intended Use Case	Framebuffer (Megabytes)	Virtual Display Heads	Max Resolution per Monitor	Maximum vGPUs	
							Per GPU	Per Board
GRID K1	4	GRID K140Q	Power User / Designer	1024	2	2560x1600	4	16
		GRID K100	Knowledge Worker	256	2	1920x1200	8	32
GRID K2	2	GRID K260Q	Power User / Designer	2048	4	2560x1600	2	4
		GRID K240Q	Power User / Designer	1024	2	2560x1600	4	8
		GRID K200	Knowledge Worker / Power User	256	2	1920x1200	8	16

감사합니다