NVIDIA
REWRITING THE RULES WITH GPUS
Innovative Solutions for Energy Exploration and Data Processing, Nov 2013

Outline

- Leveraging GPUs for energy exploration
- Ultra-compact GPU configurations
- Best practices for GPU deployments

NVIDIA.

Outline

- Leveraging GPUs for energy exploration
- Ultra-compact GPU configurations
- Best practices for GPU deployments

NVIDIA.

Increasing need for energy exploration

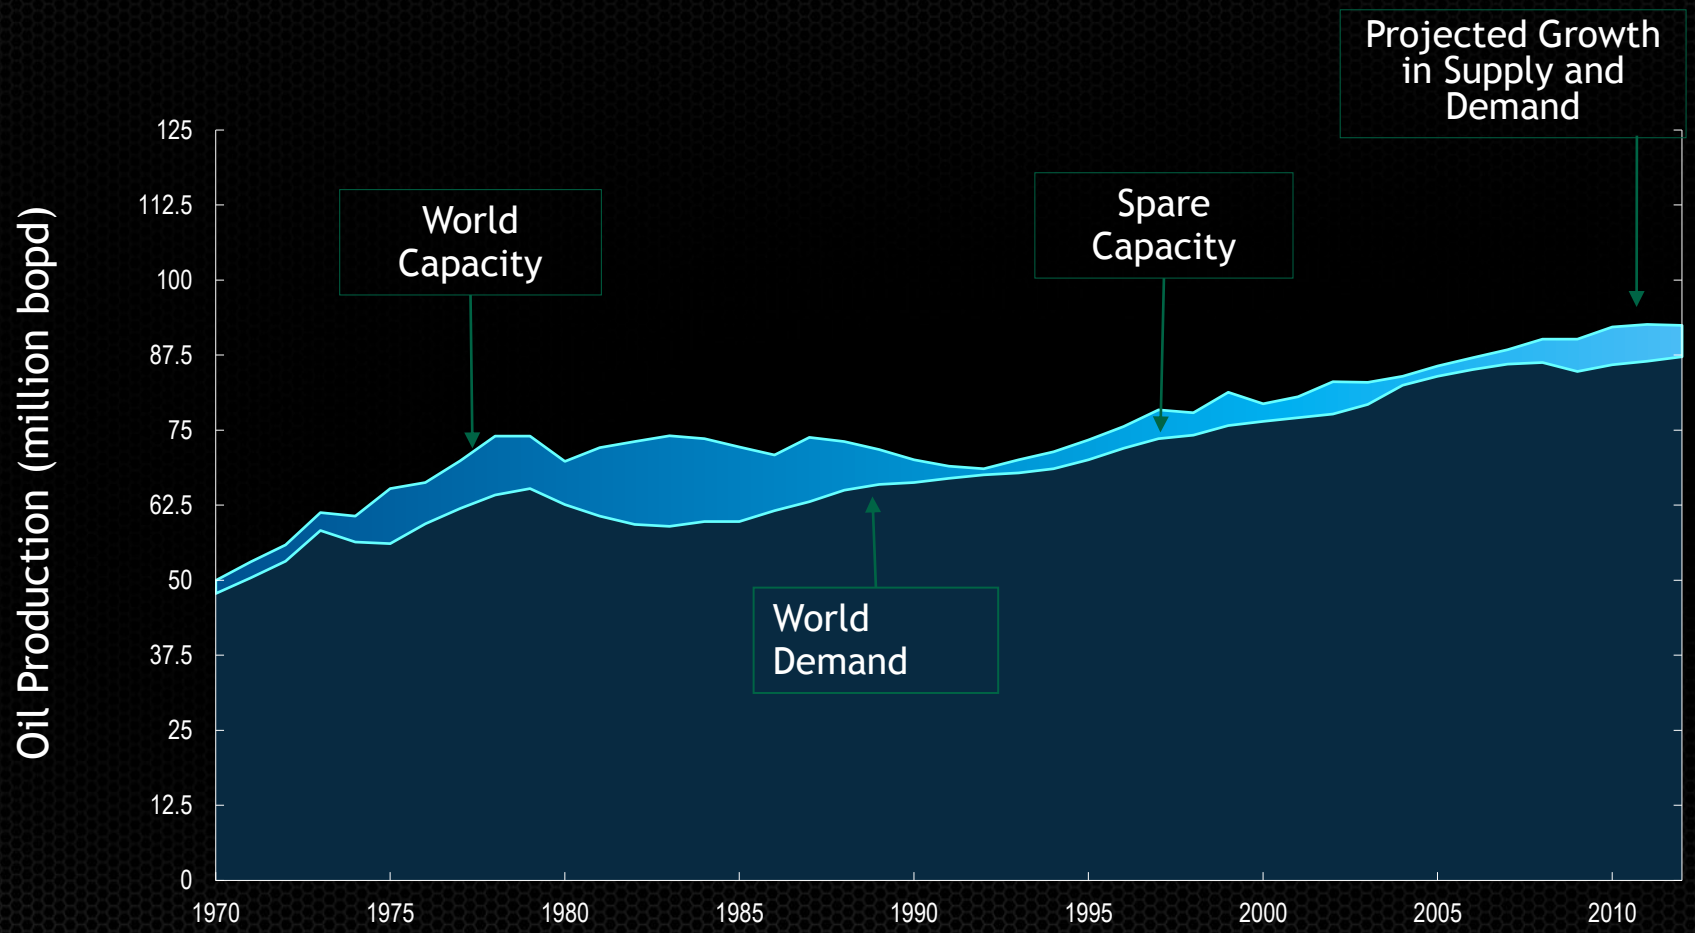

NVIDIA.

Increasing need for energy exploration

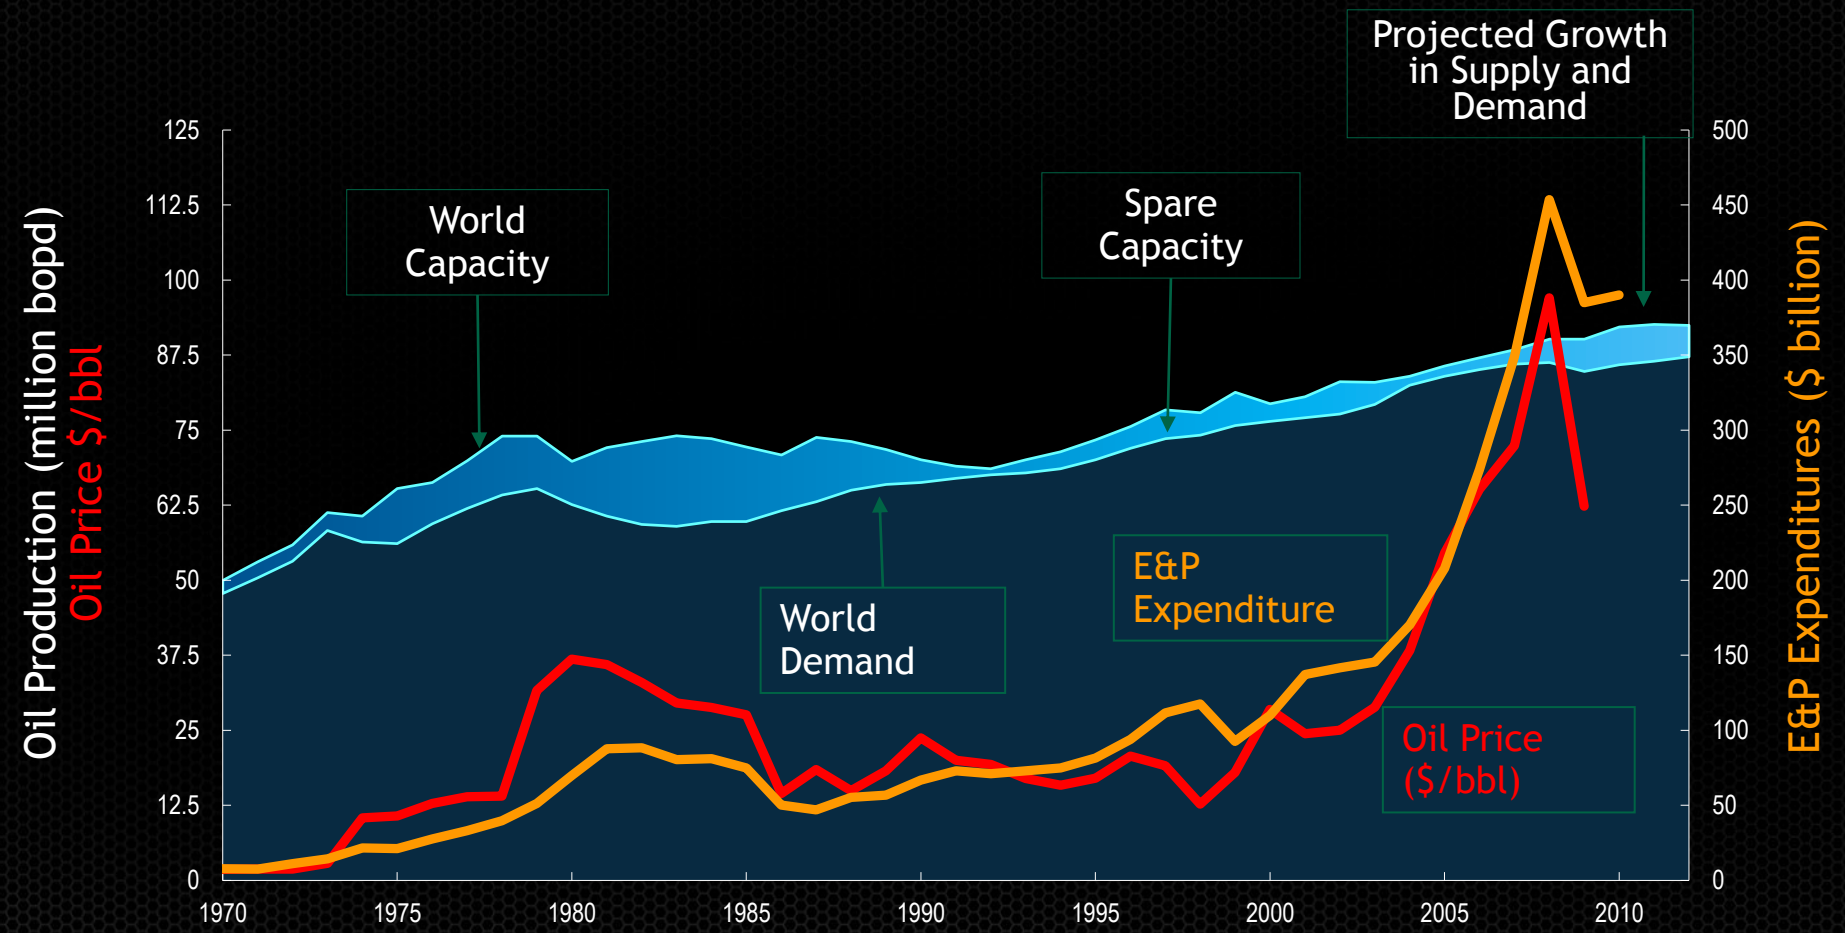

NVIDIA

Use simulations to reduce risk

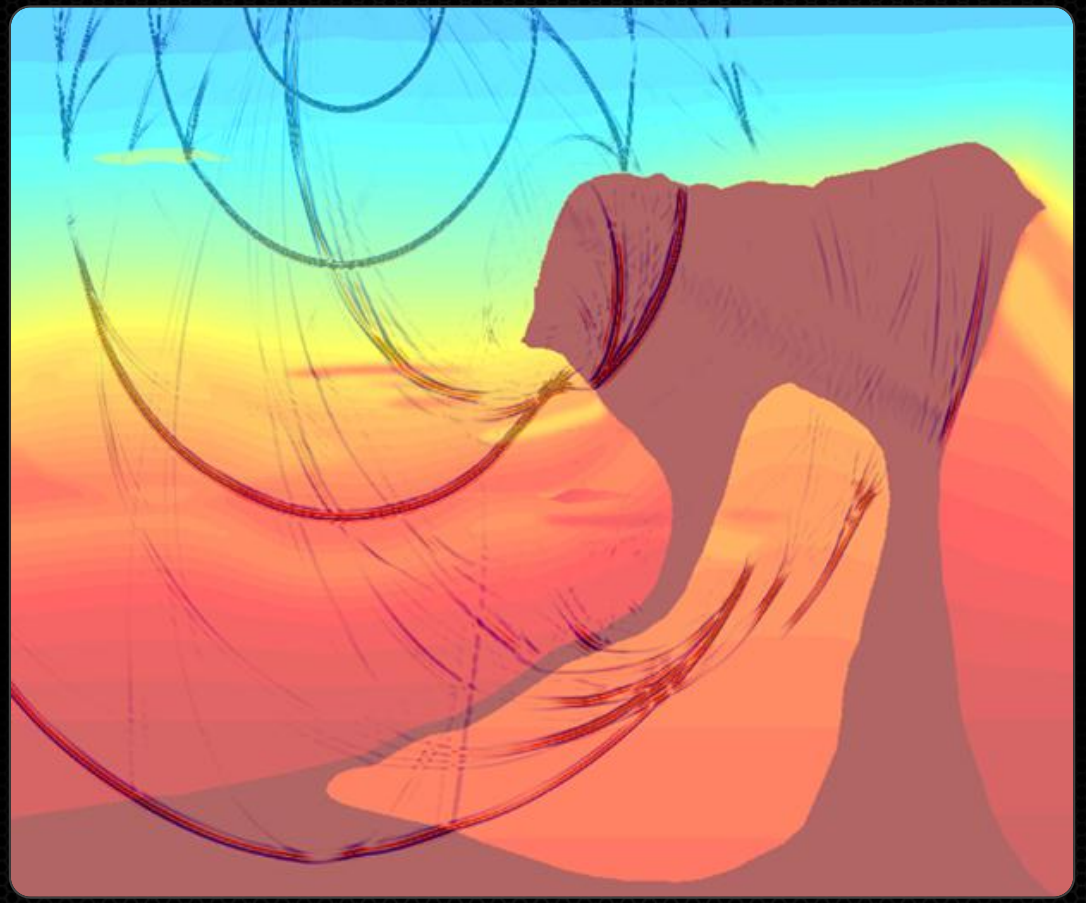

Image courtesy of Acceleware

NVIDIA.

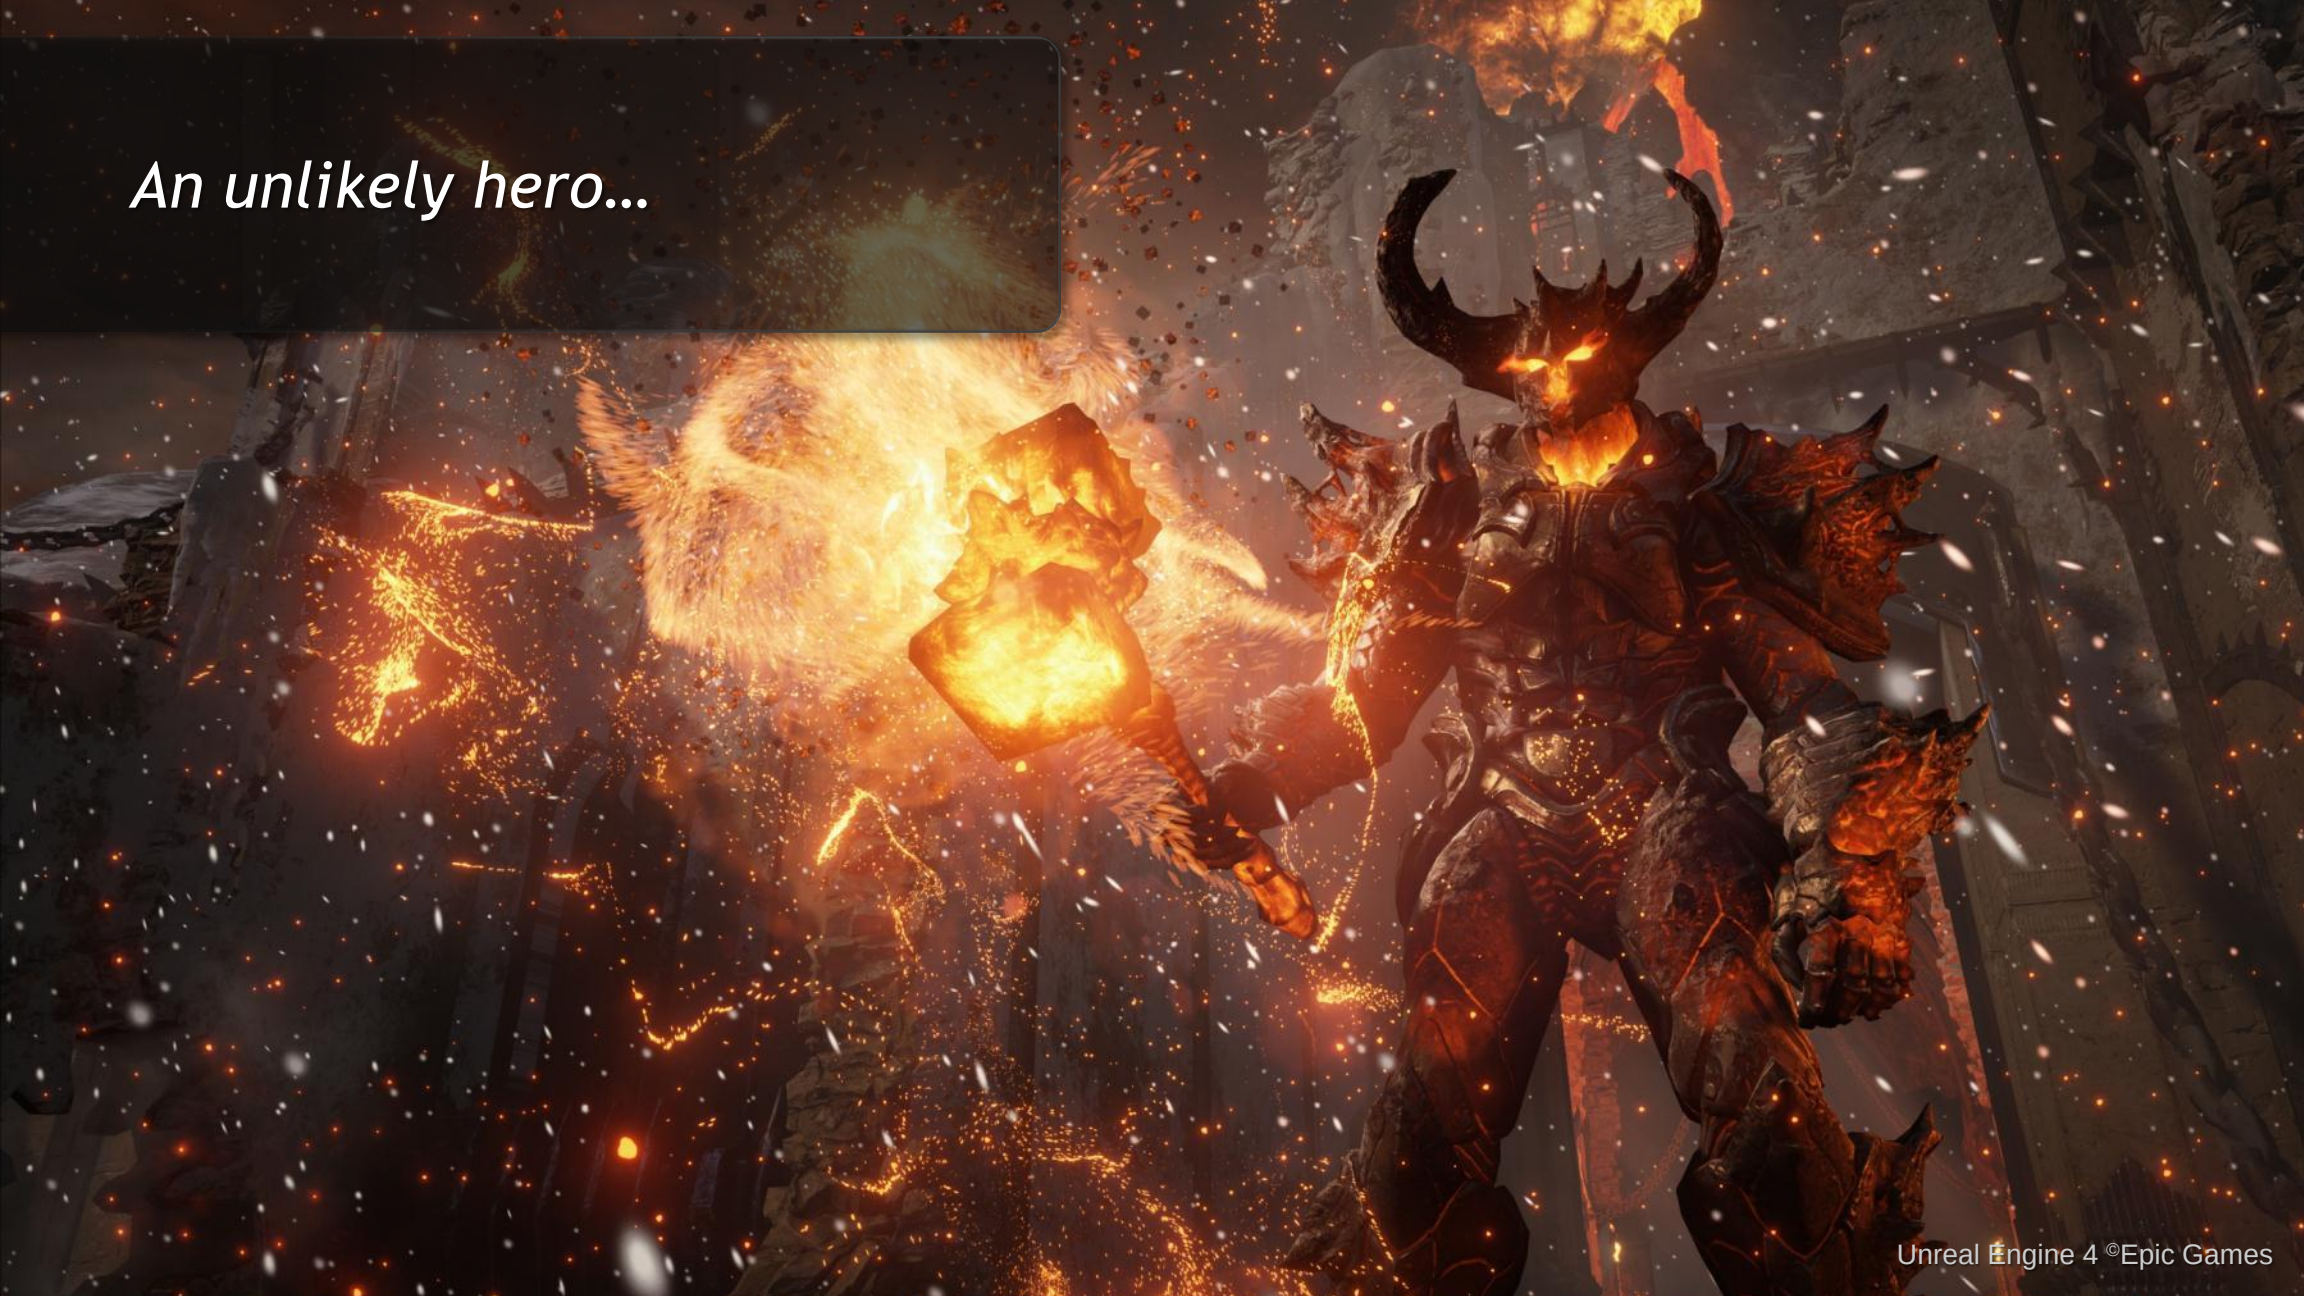
An unlikely hero…
Unreal Engine 4 ©Epic Games

Computer graphics require billions of parallel computations

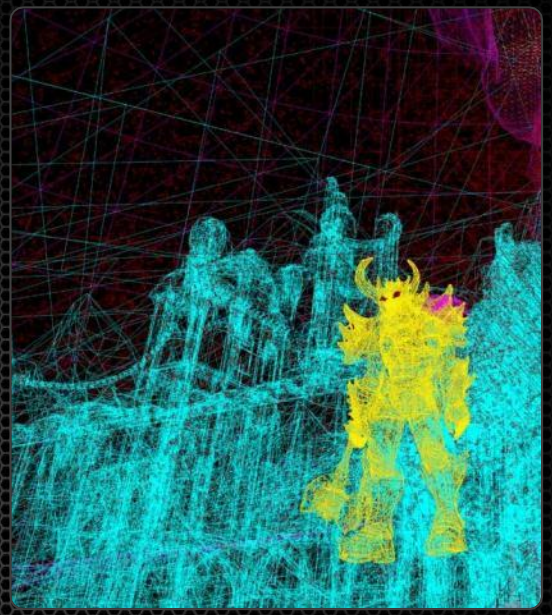

Millions of triangles

Millions of pixels

Why are so many parallel operations needed?

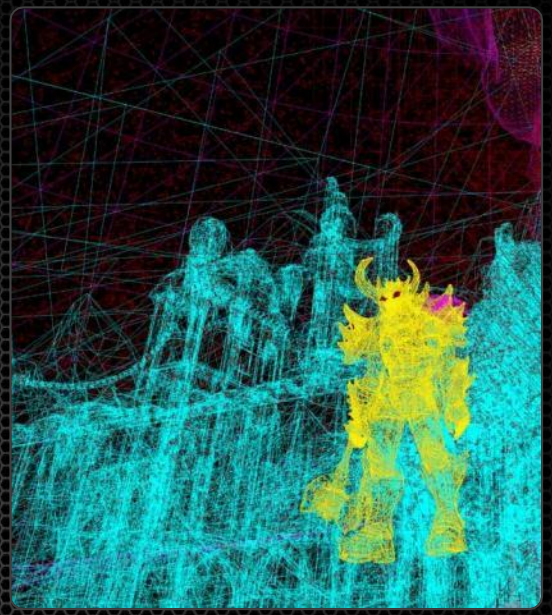

Millions of triangles

Millions of pixels

Why are so many parallel operations needed?

Input triangle

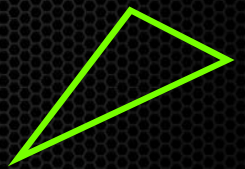

Transform vertices

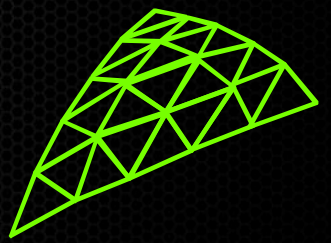

Tessellate

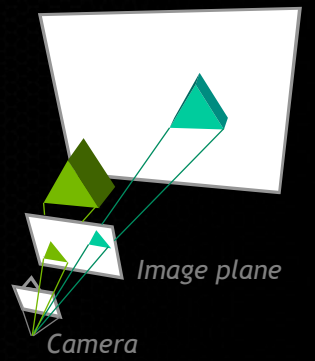

Projection

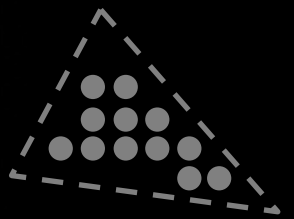

Rasterize

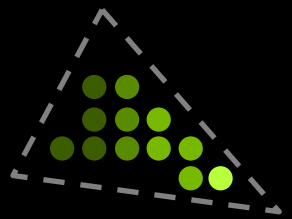

Shade

E&P workflow

Seismic Acquisition

Seismic Imaging

Interpretation

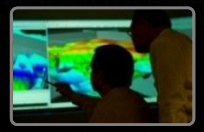

Reservoir Characterization

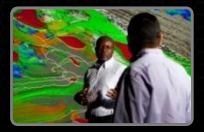

Petrophysics

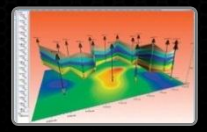

Well Planning

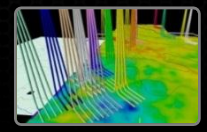

Drilling

Reservoir Simulation

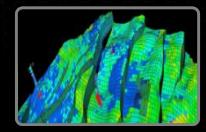

Economics

Images courtesy of Schlumberger

NVIDIA.

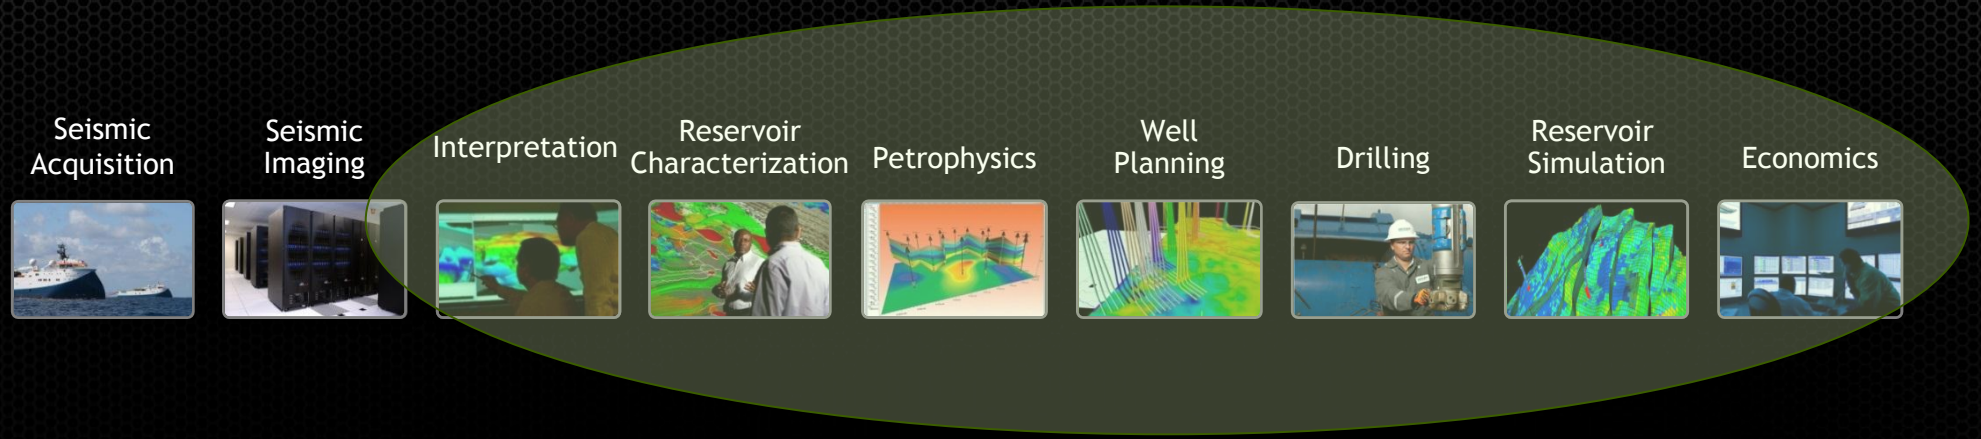

Quadro™ Visualization

Images courtesy of Schlumberger

NVIDIA.

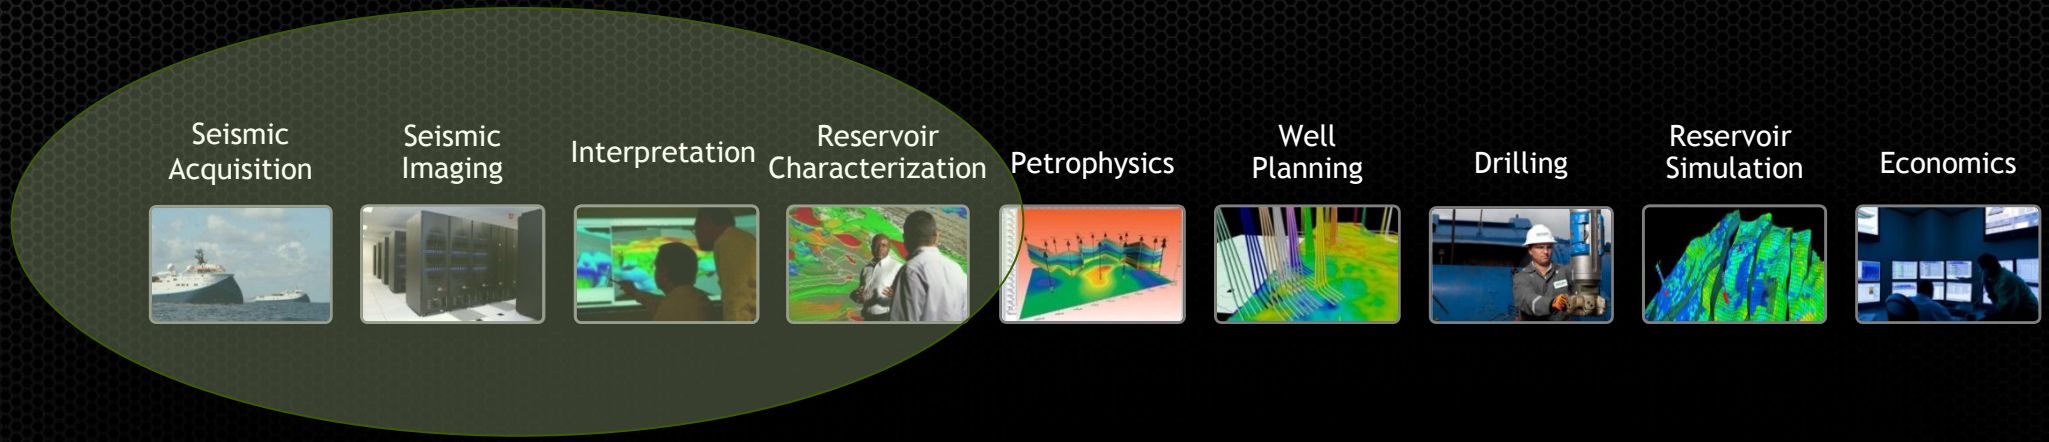

Tesla™ Computation

Images courtesy of Schlumberger

NVIDIA.

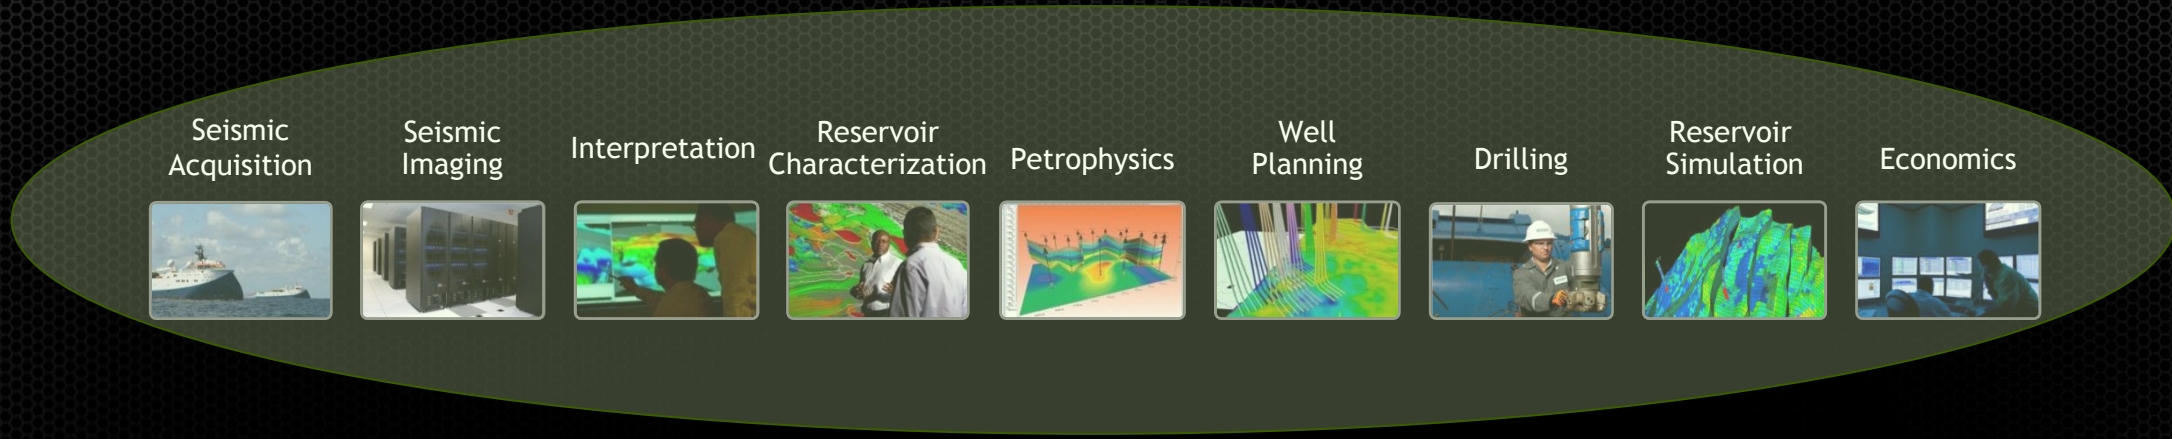

NVIDIA Visual Computing

Transforming E&P workflow

Images courtesy of Schlumberger

NVIDIA

E&P workflow

Seismic Imaging

Reservoir Characterization

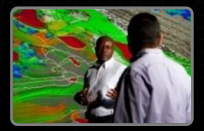

Petrophysics

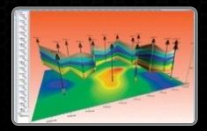

Well Planning

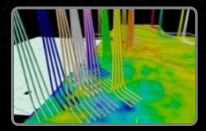

Drilling

Reservoir Simulation

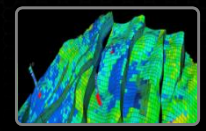

Economics

Images courtesy of Schlumberger

NVIDIA.

Compute requirements

Seismic Imaging

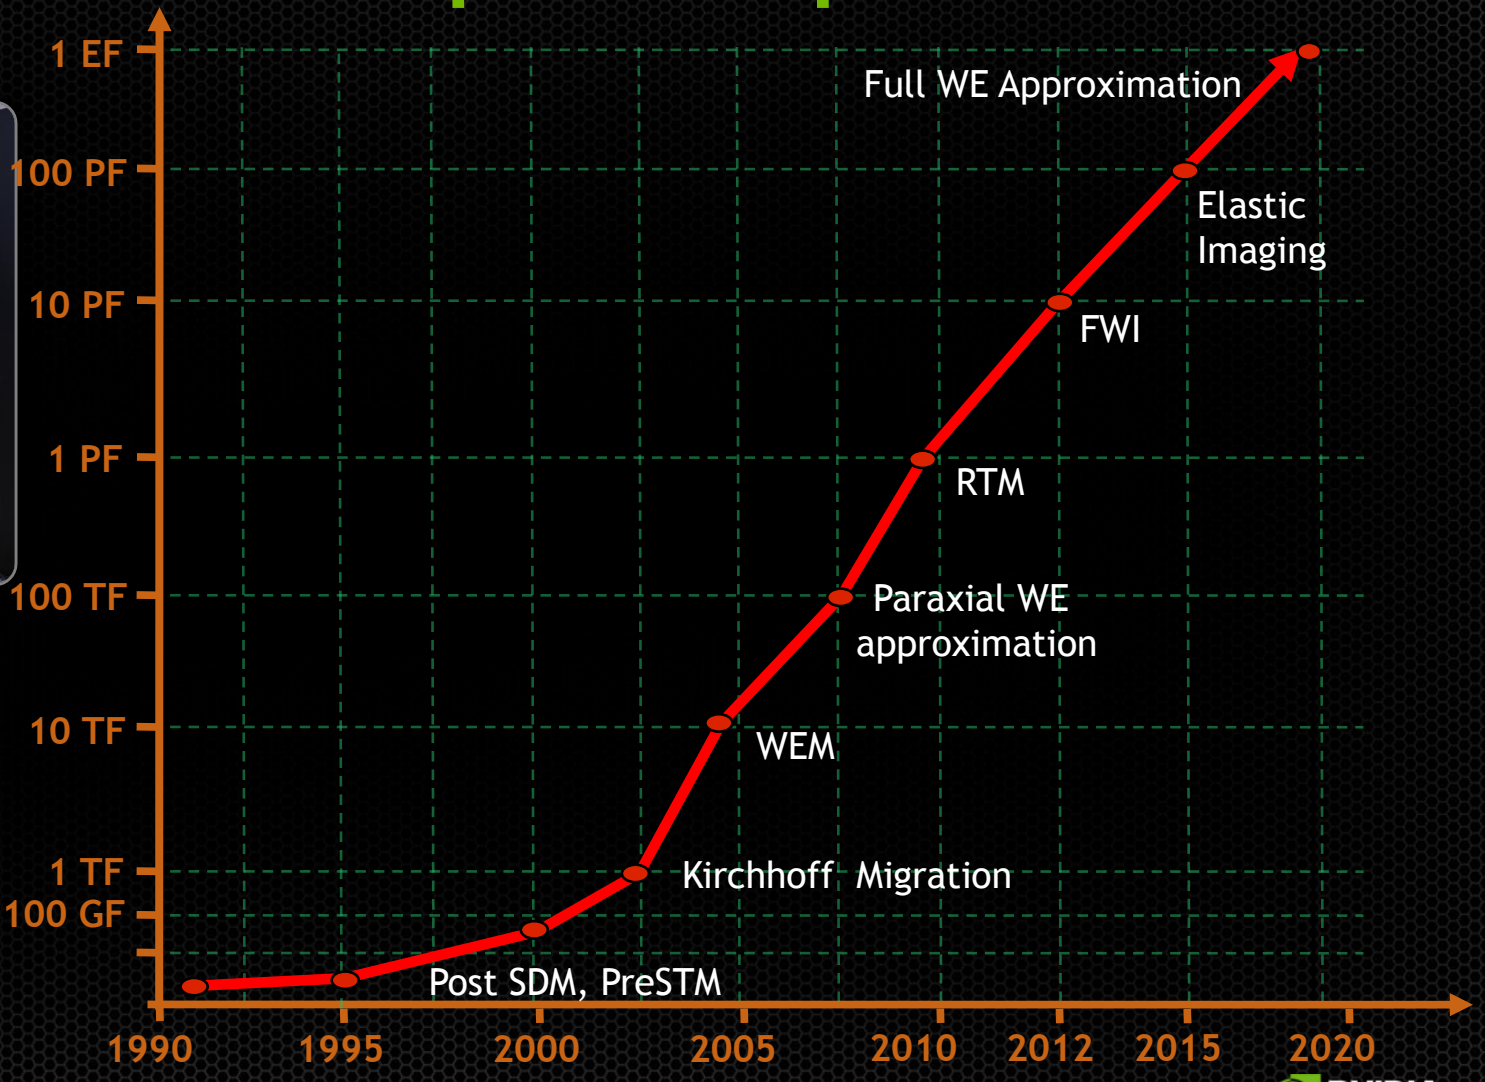

Chart courtesy of TOTAL

NVIDIA

Why augment infrastructure with GPUs?

4x to 6x more throughput

12% cost reduction

4x Perf/W advantage

More revenue
Same cost basis

Sharper image
Higher throughput
Better drilling decisions

NVIDIA.

Why augment infrastructure with GPUs?

4x to 6x more throughput

12% cost reduction

4x Perf/W advantage

28 Days to 7 Days

¼ Power & Cooling Costs

Improved Opex/Capex

More revenue

Same basis

Sharper image

Higher throughput

Better drilling decisions

NVIDIA.

Solution to challenges: just add GPUs

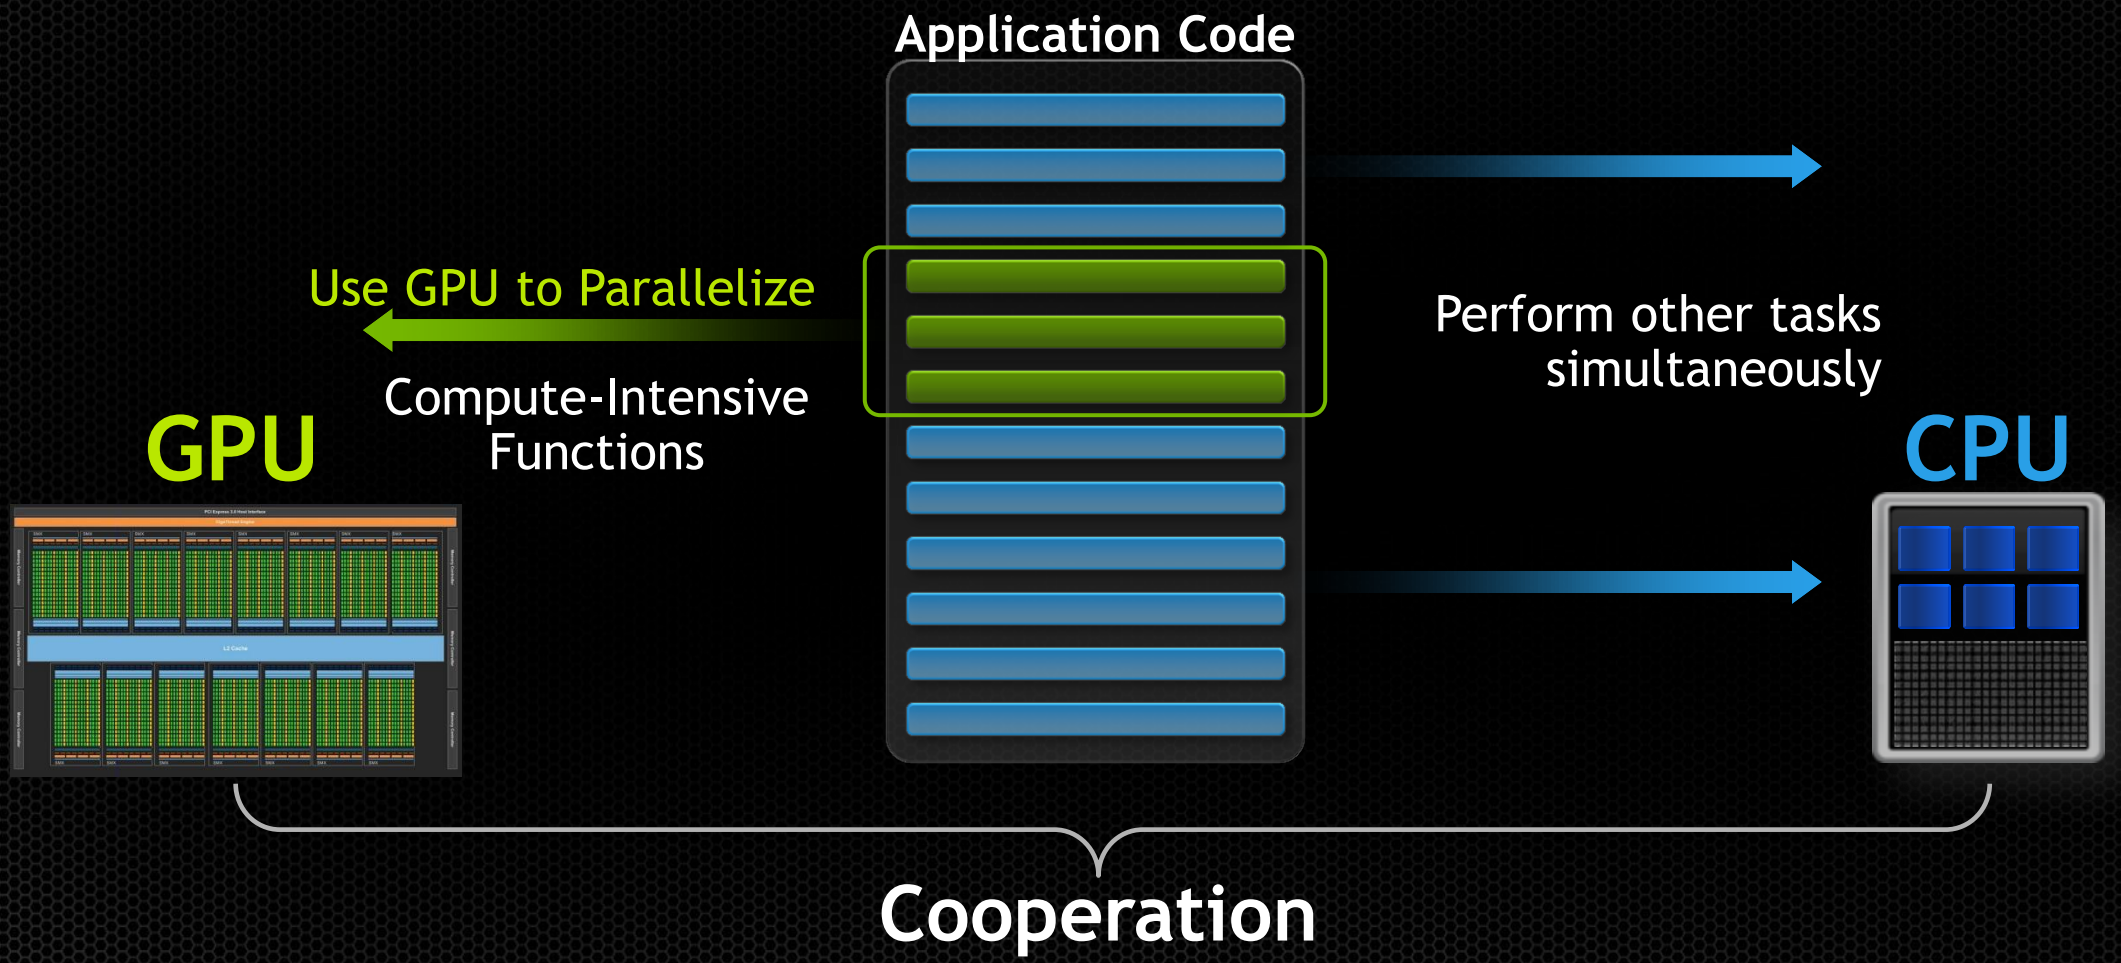

NVIDIA.

GPUs: Two year heartbeat

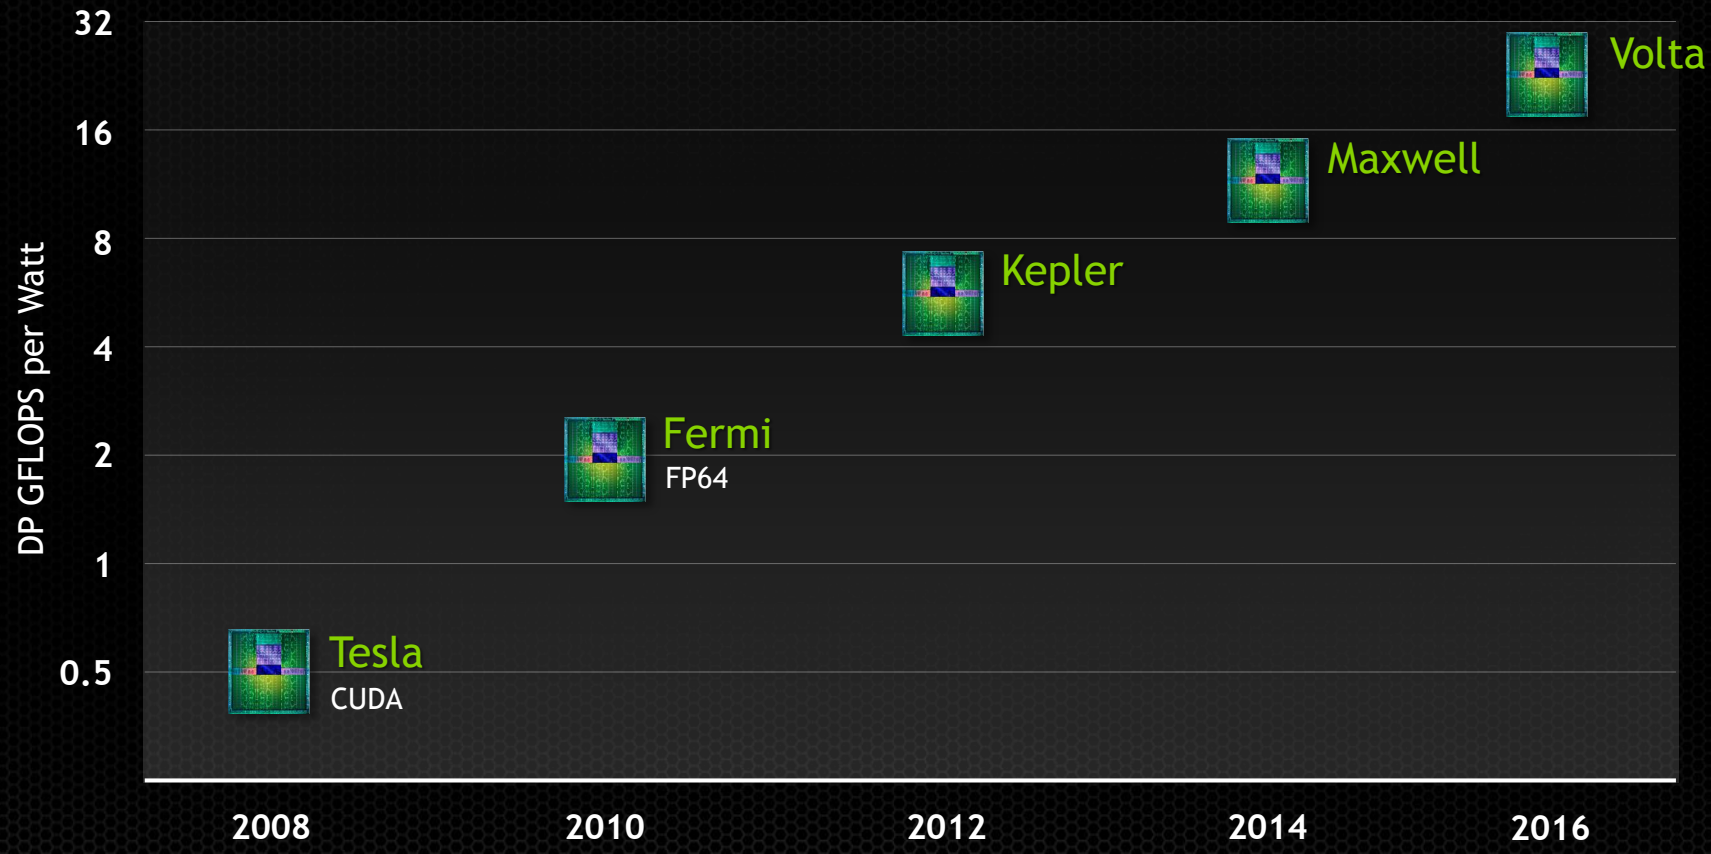

NVIDIA.

Five years of NVIDIA GPUs in seismic

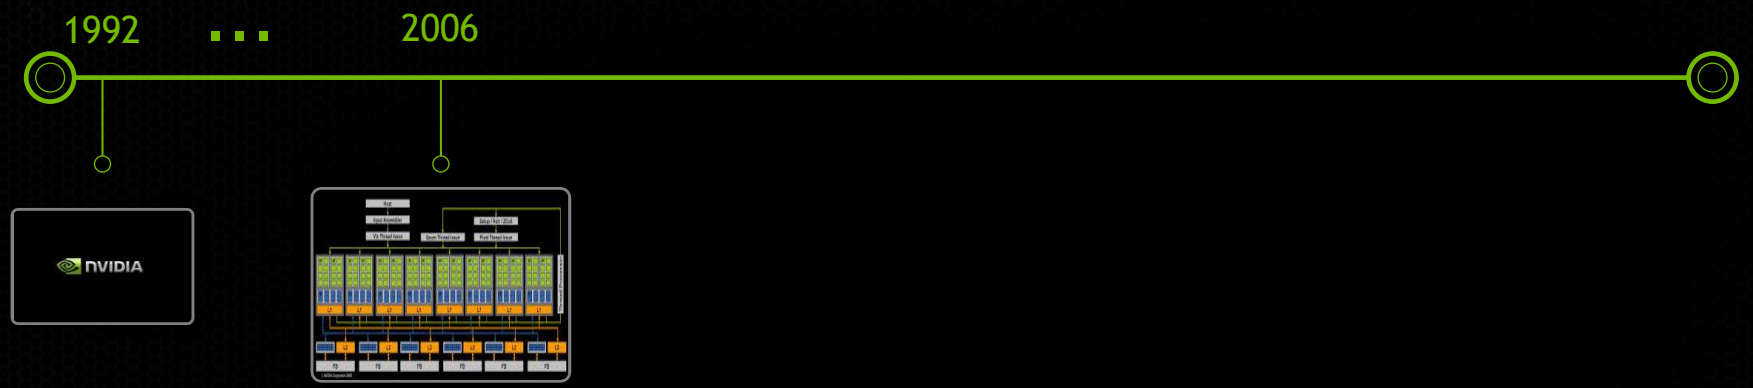

NVIDIA Founded

1st Generation
Tesla M870
CUDA Released

Five years of NVIDIA GPUs in seismic

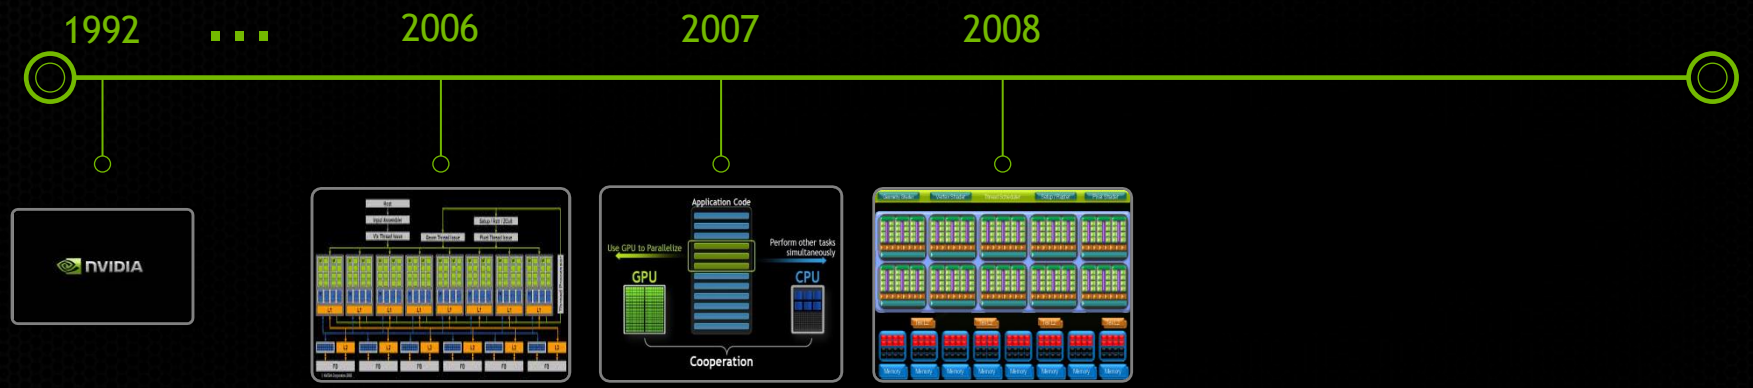

NVIDIA Founded

1st Generation
Tesla M870
CUDA Released

App Ports
KDM
KTM
WEM

2nd Generation
Tesla M1060
Port RTM

NVIDIA

Five years of NVIDIA GPUs in seismic

1992 . . . 2006 2007 2008 2010

NVIDIA Founded

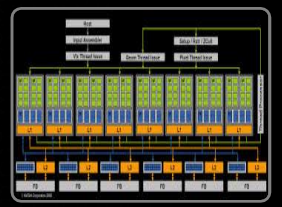

1st Generation
Tesla M870
CUDA Released

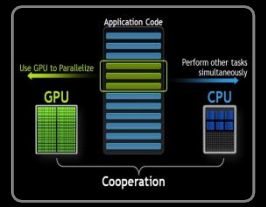

App Ports
KDM
KTM
WEM

2nd Generation
Tesla M1060
Port RTM

3rd Generation
Tesla M2090
“Fermi”
Port FWI

NVIDIA

Five years of NVIDIA GPUs in seismic

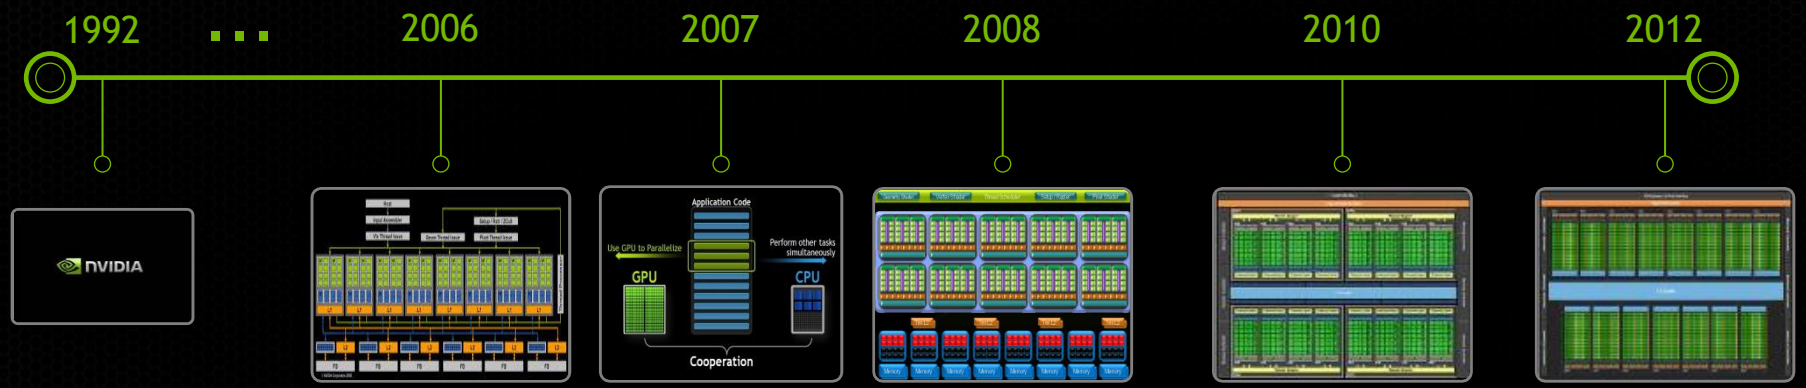

NVIDIA Founded

1st Generation
Tesla M870
CUDA Released

App Ports
KDM
KTM
WEM

2nd Generation
Tesla M1060
Port RTM

3rd Generation
Tesla M2090
“Fermi”
Port FWI

4th Generation
Tesla K10
“Kepler”
Port Elastic

NVIDIA

Industry trends and challenges

GPU-based servers

Consolidate workloads

Improve datacenter operations

Image courtesy of Schlumberger

NVIDIA.

Industry trends and challenges

GPU-based servers

Consolidate workloads
Improve datacenter operations

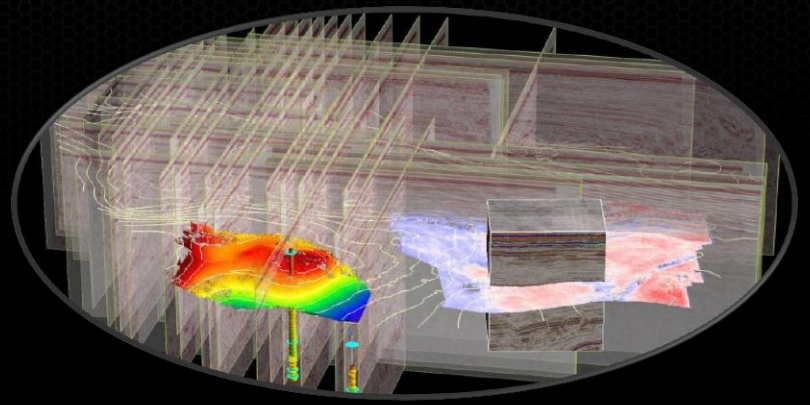

Regional field studies

Wide azimuth
Ocean bottom sensors
Time lapse surveys

Images courtesy of Schlumberger, Shell

NVIDIA.

Industry trends and challenges

GPU-based servers

Consolidate workloads
Improve datacenter operations

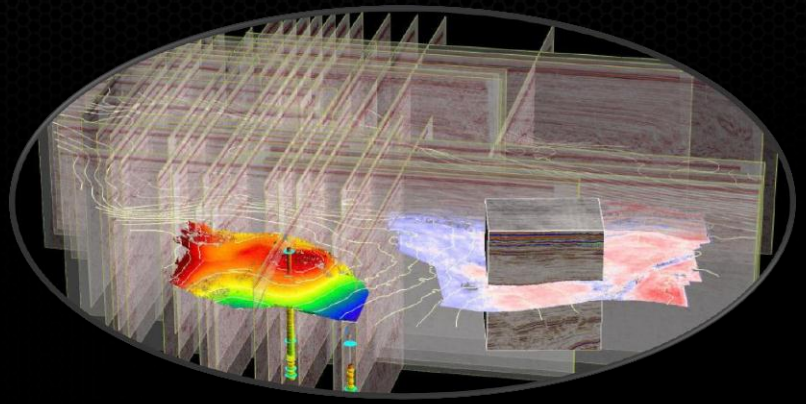

Regional field studies

Wide azimuth
Ocean bottom sensors
Time lapse surveys

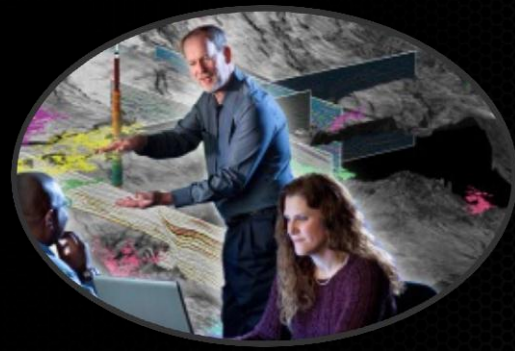

Collaborative workforce

Globally disperse
Ever-changing environments

Images courtesy of Schlumberger, Shell, and Paradigm

NVIDIA

Remote collaboration architecture

Notebook PCs running Client

Server platform

NIC

VT-d capable chipset

GRID K2

GRID K2

Multiple users on virtualized server platform

Public or Private Cloud

NVIDIA

Remote collaboration architecture

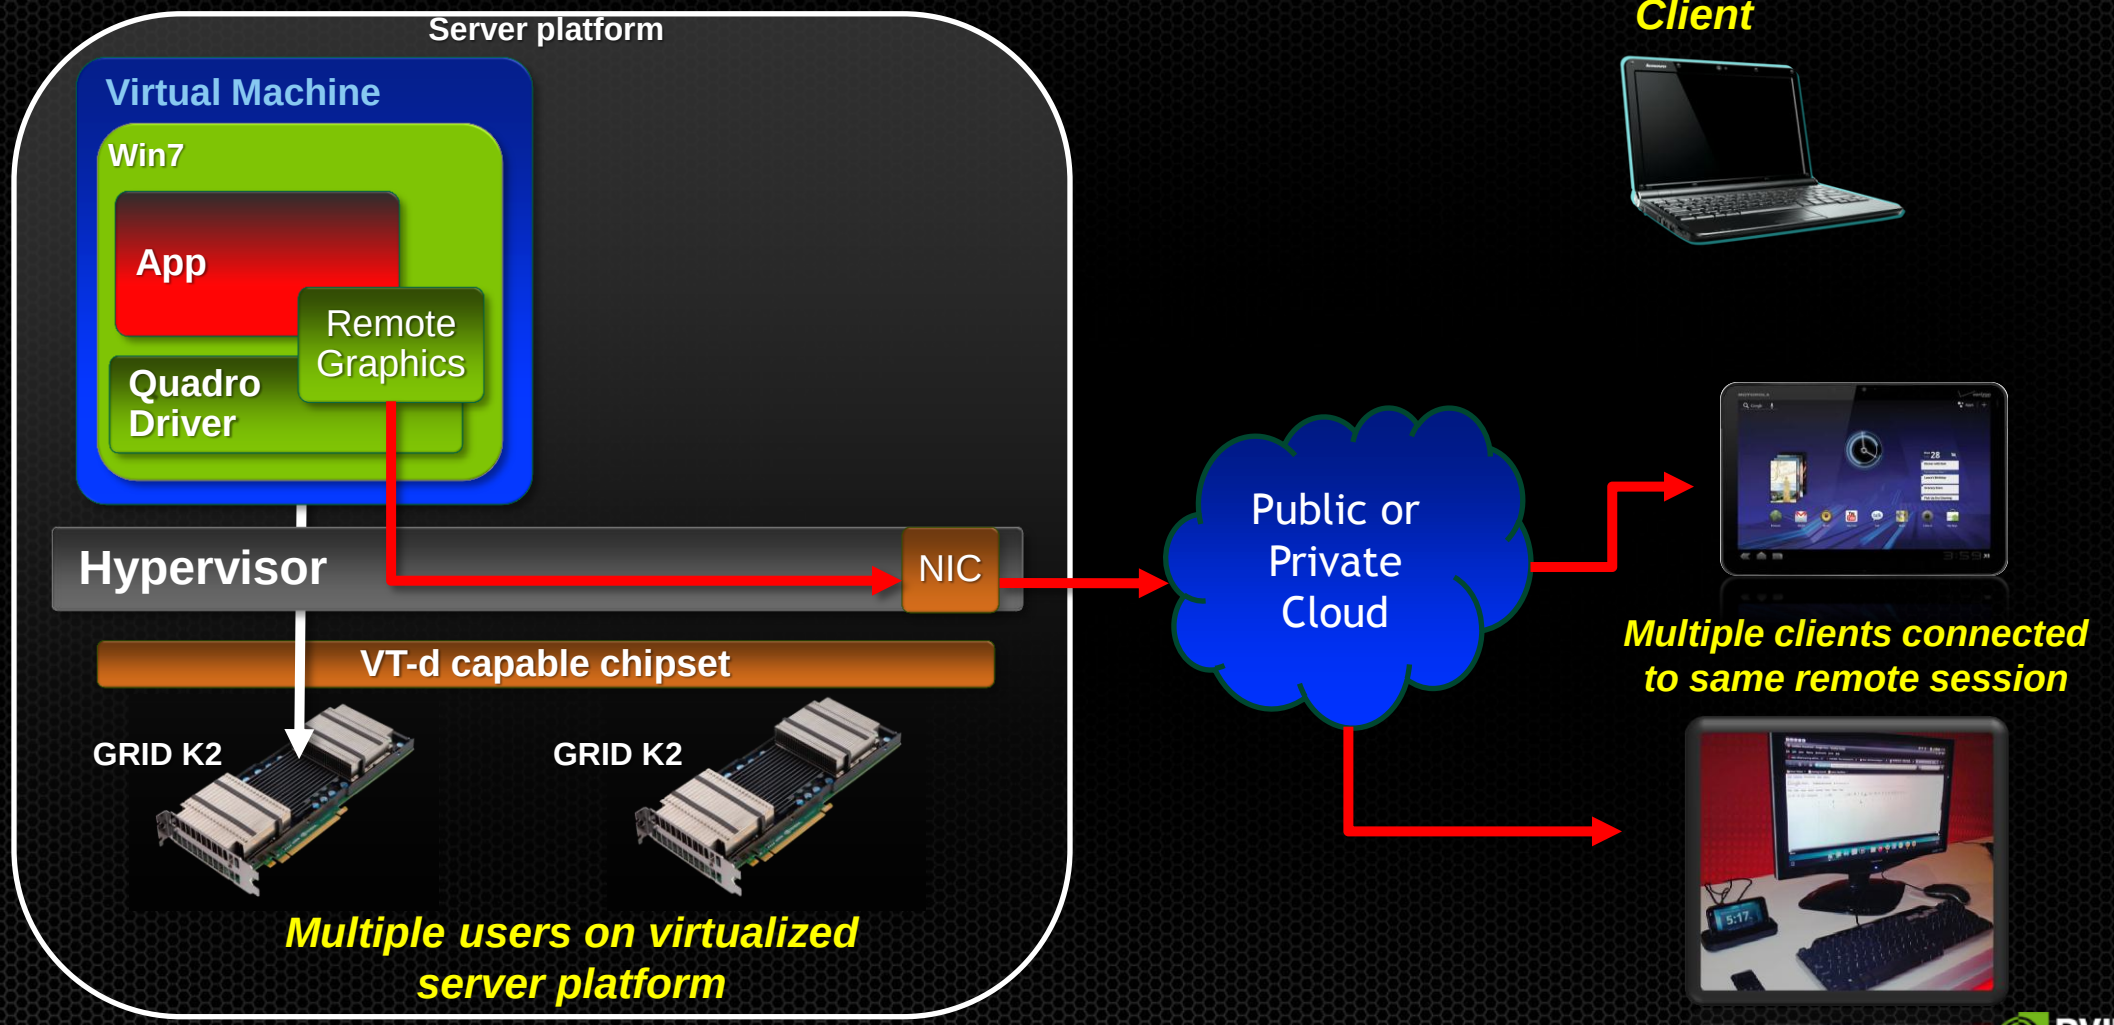

NVIDIA.

Remote collaboration architecture

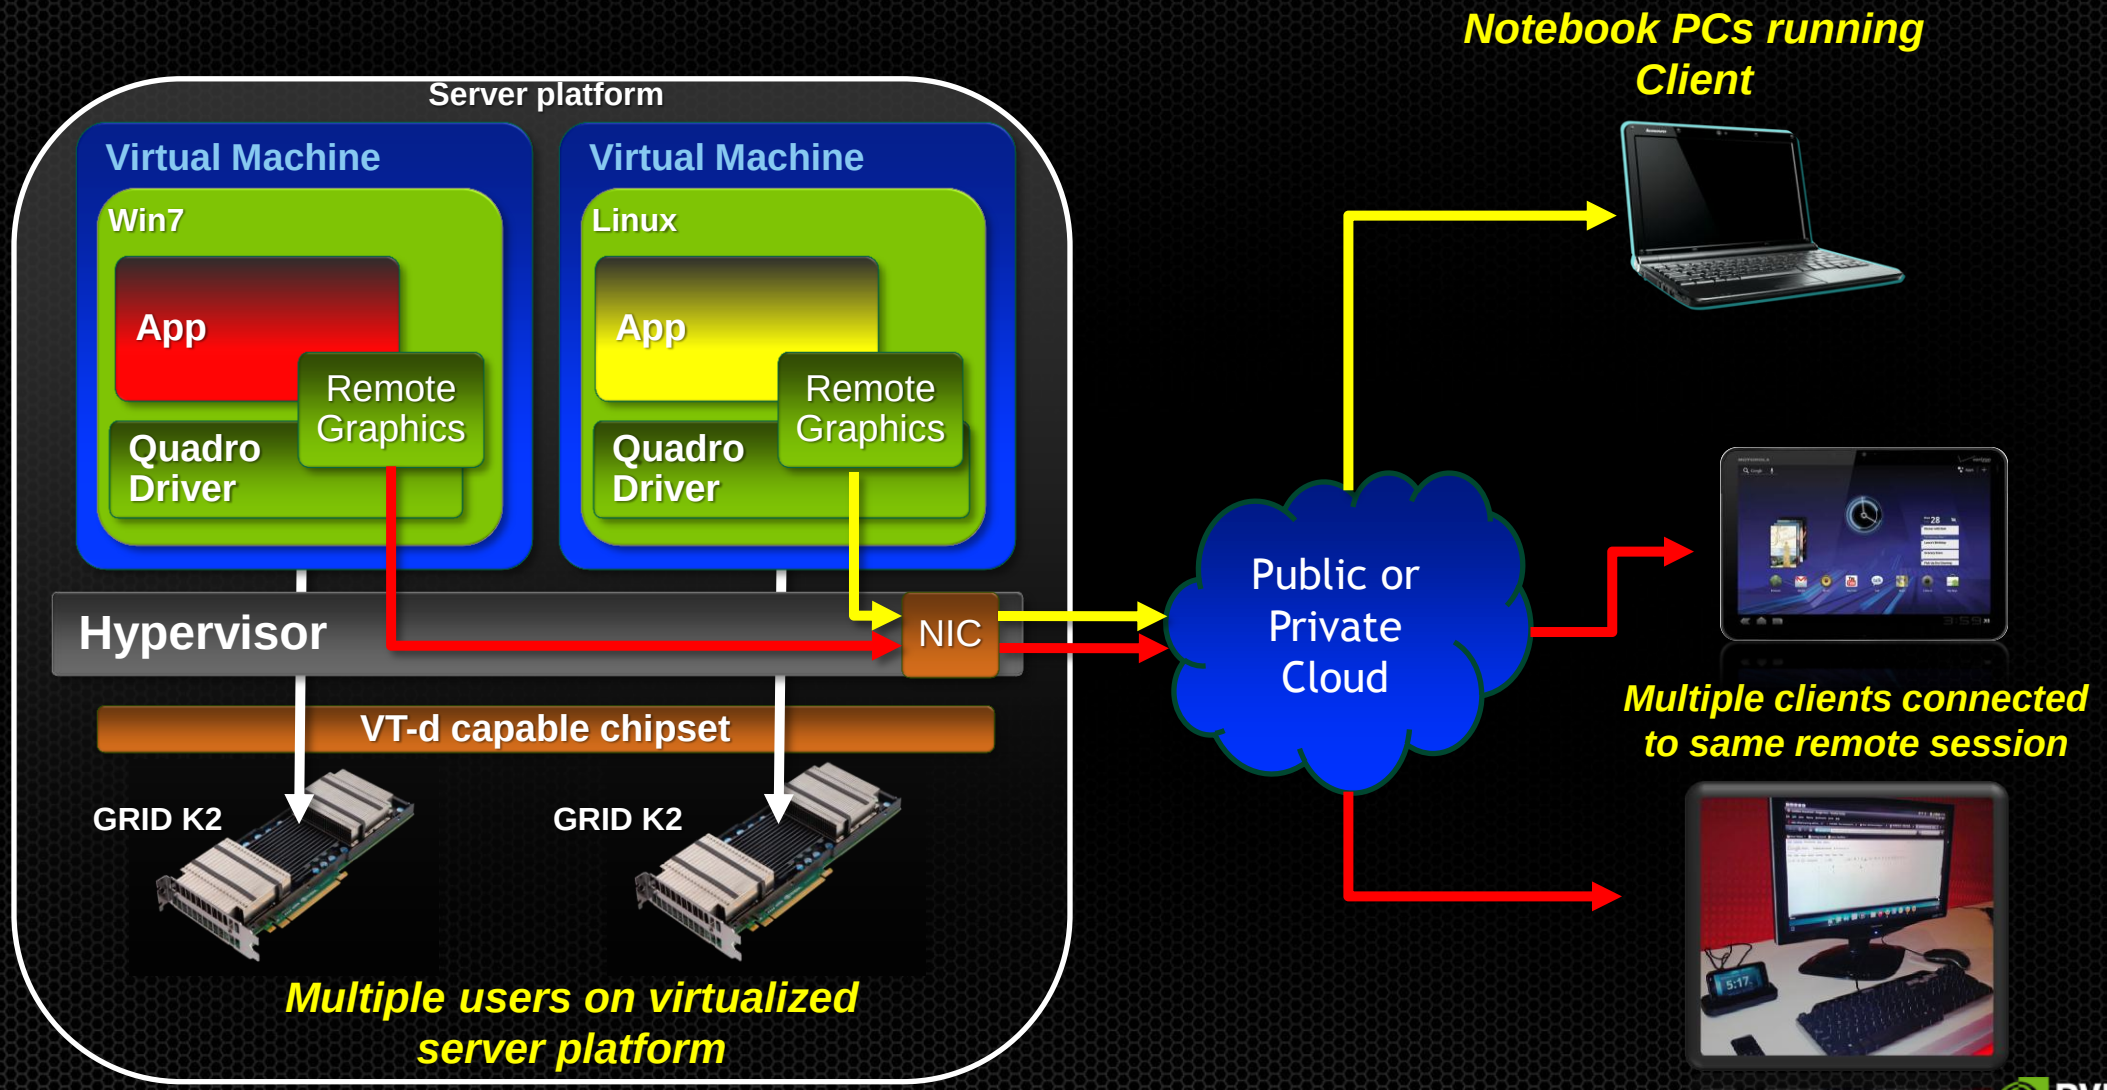

NVIDIA.

Transformational hybrid architecture

Compute Intensive Apps

GPU Cluster

Seismic Imaging
Electromagnetic Imaging
Reservoir Simulation

Image courtesy of Schlumberger

NVIDIA.

Transformational hybrid architecture

Compute Intensive Apps

GPU Cluster

Seismic Imaging
Electromagnetic Imaging
Reservoir Simulation

High-speed network

Data Storage

Interactive Apps

GPU Cluster

Adapt to client needs
Scalable rendering
Fast access to data
Automatic updates

Image courtesy of Schlumberger

NVIDIA.

Transformational hybrid architecture

Compute Intensive Apps

GPU Cluster

Seismic Imaging
Electromagnetic Imaging
Reservoir Simulation

High-speed network

Data Storage

Interactive Apps

GPU Cluster

Adapt to client needs
Scalable rendering
Fast access to data
Automatic updates

Datacenter

Client Environments

Workstations
Laptops
Quad-HD displays
Immersive rooms
Tablets
Smart phones

Image courtesy of Schlumberger

NVIDIA

Depth-correct transparency rendering

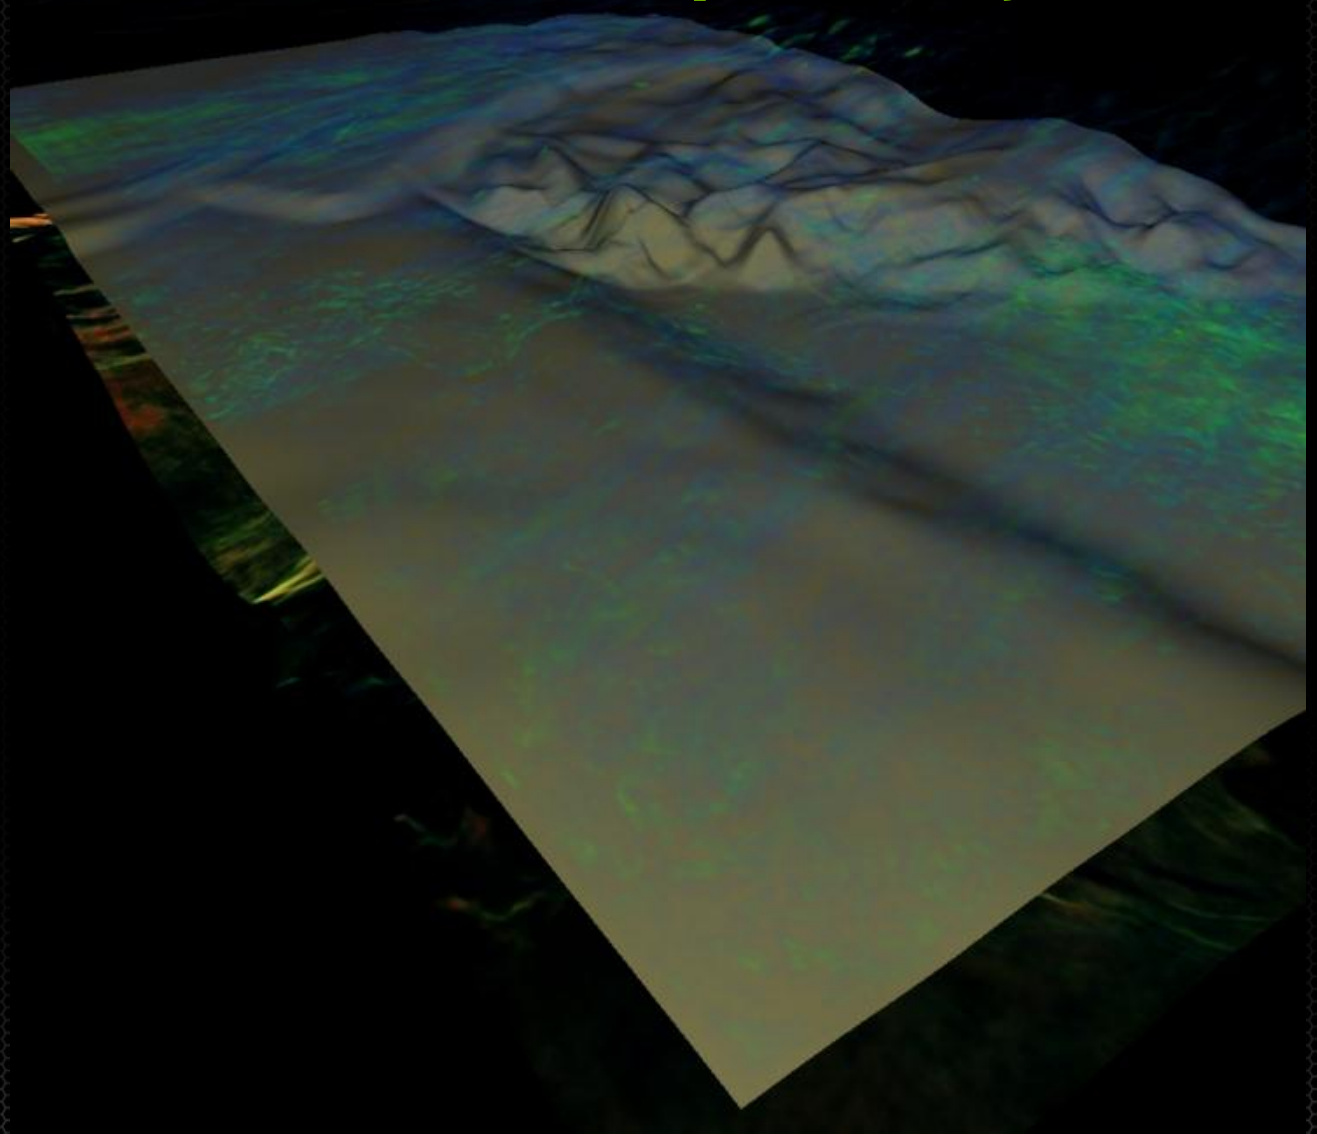

NVIDIA.

Depth-correct transparency rendering

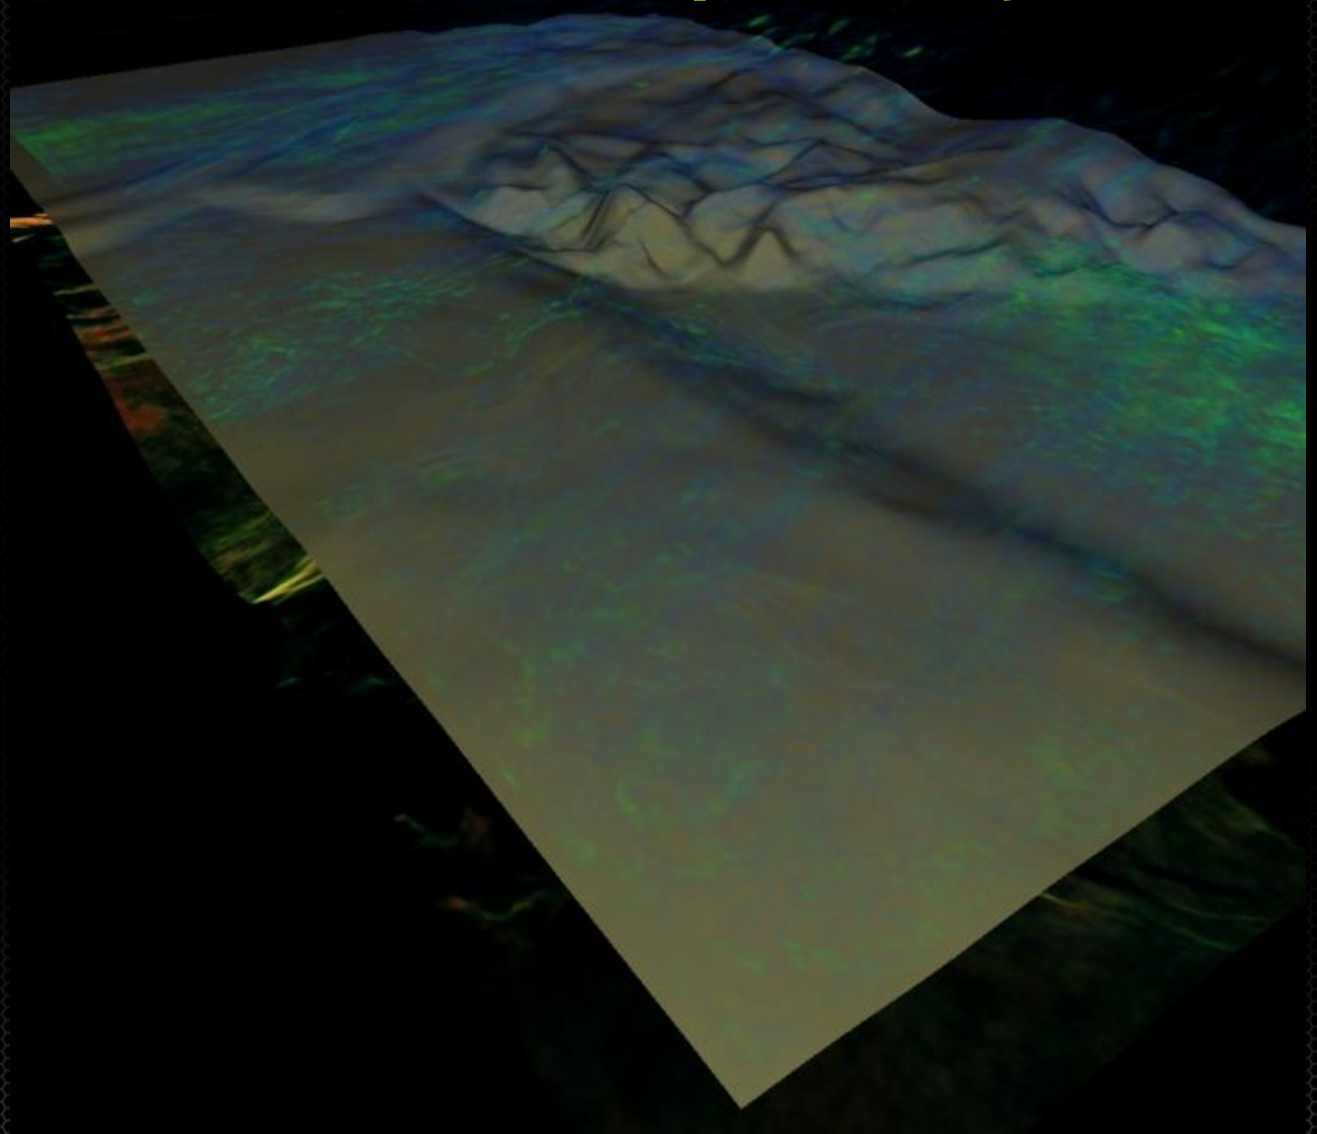

NVIDIA.

Depth-correct transparency rendering

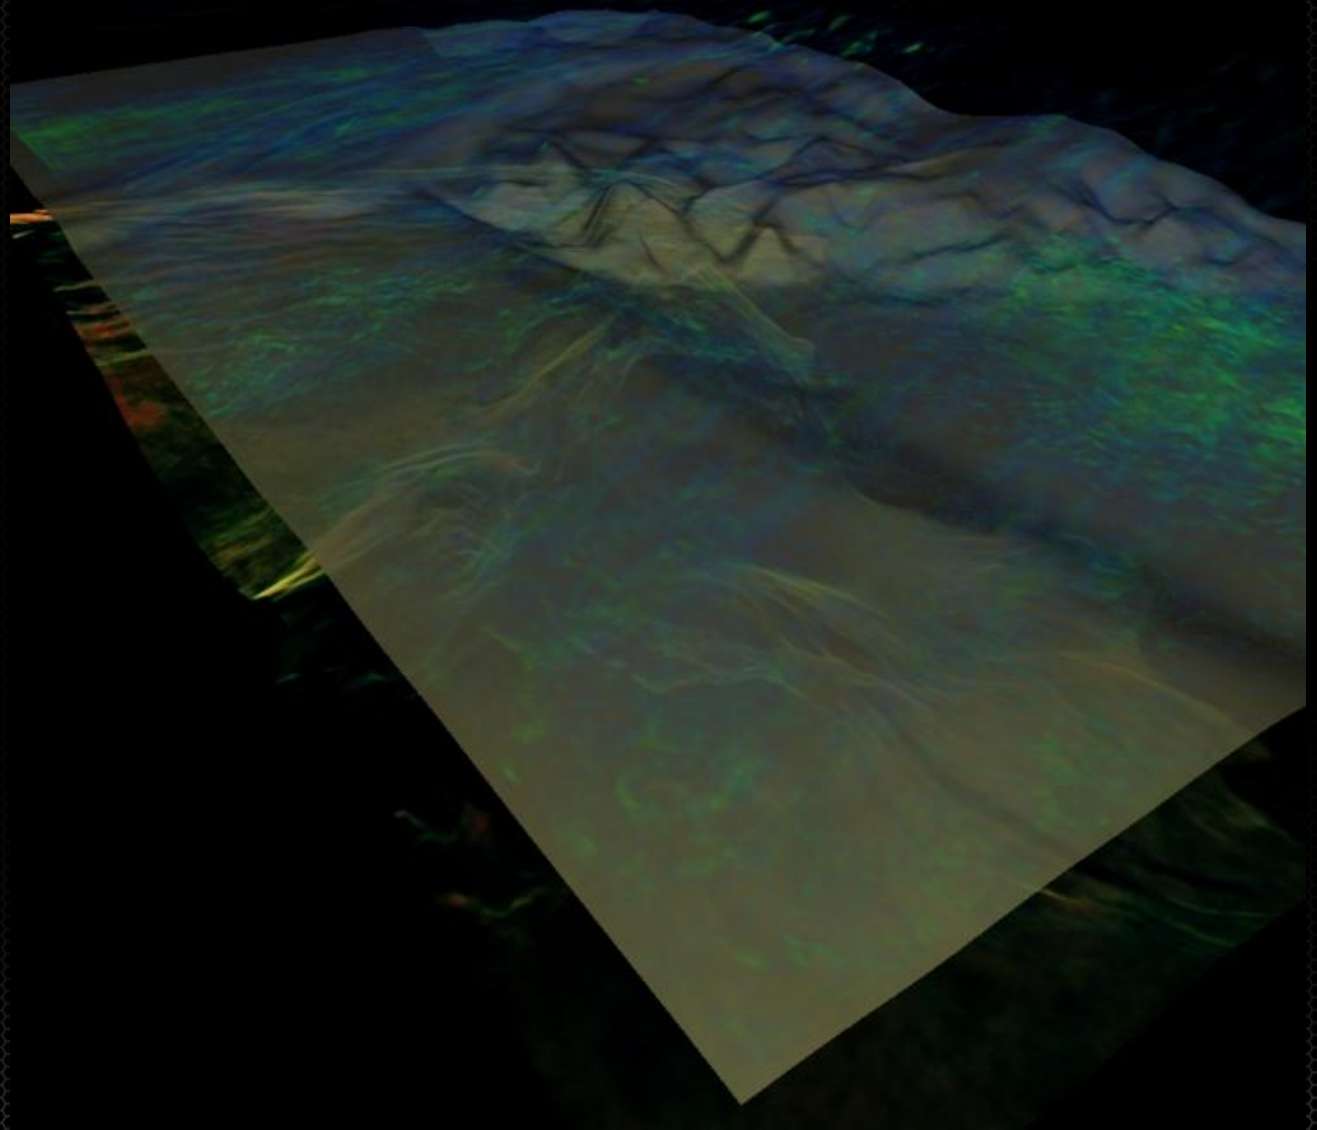

NVIDIA.

Depth-correct transparency rendering

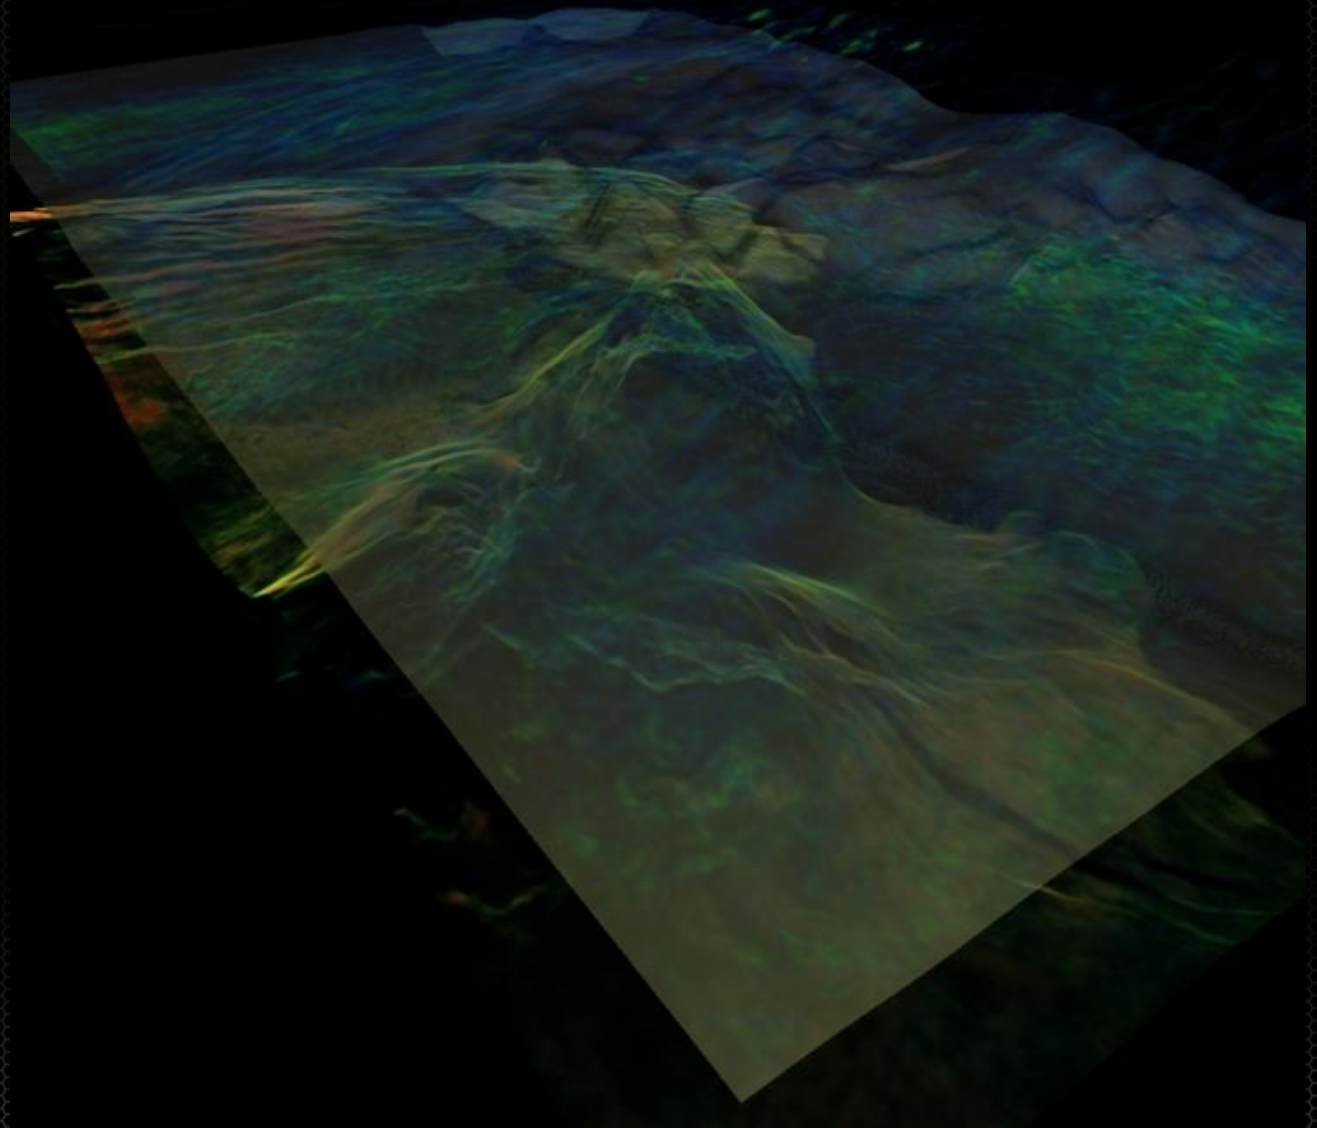

NVIDIA.

Depth-correct transparency rendering

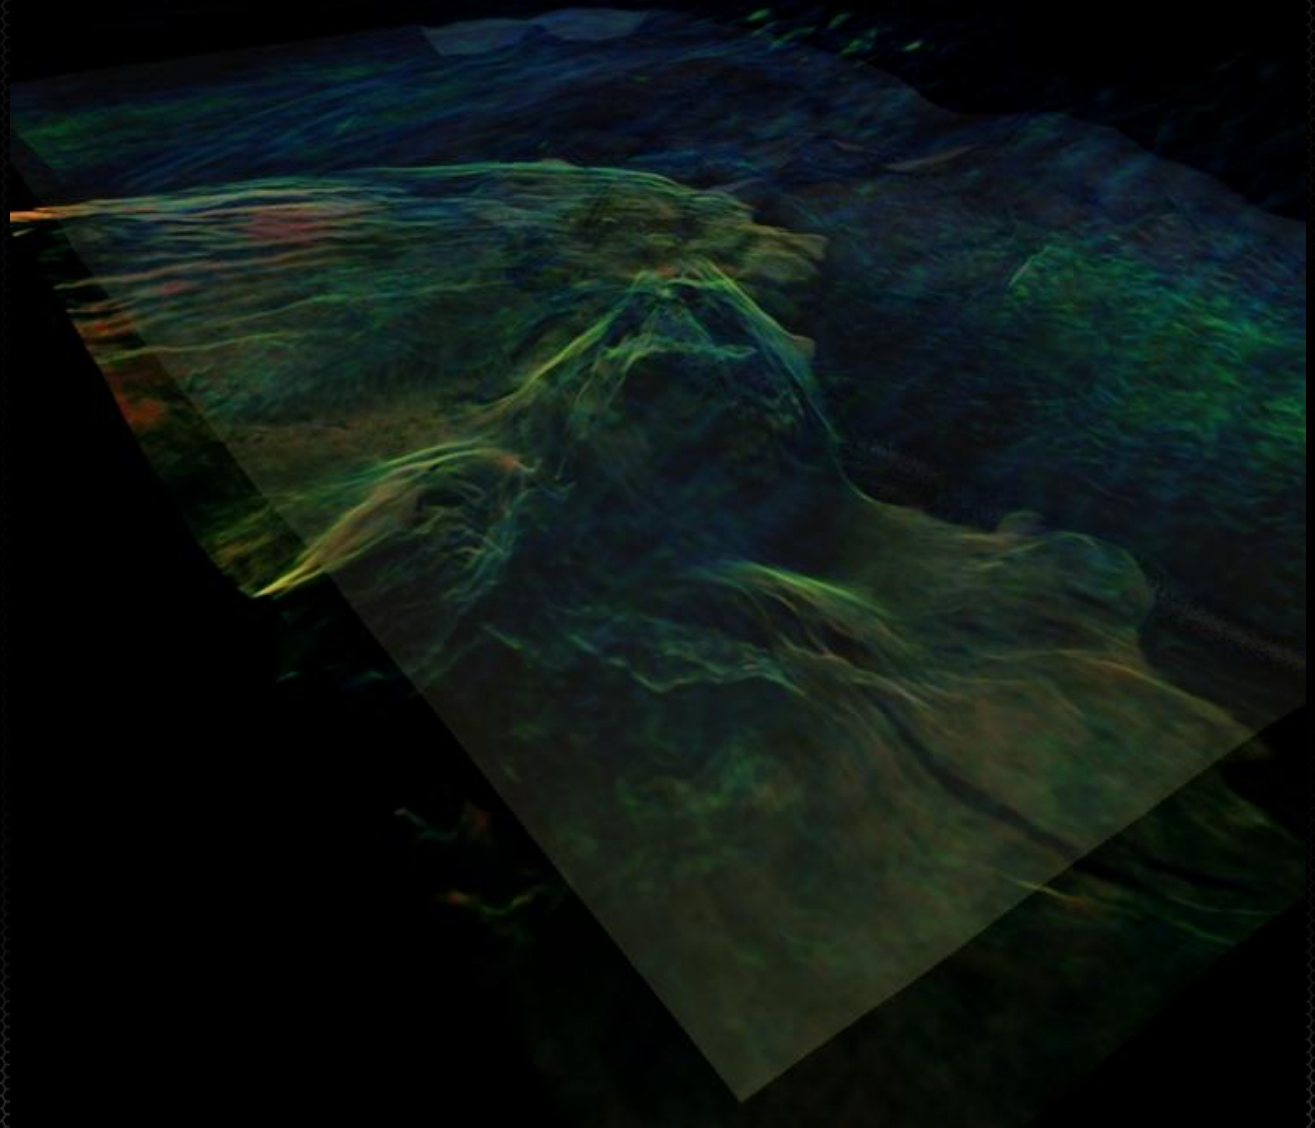

NVIDIA.

Depth-correct transparency rendering

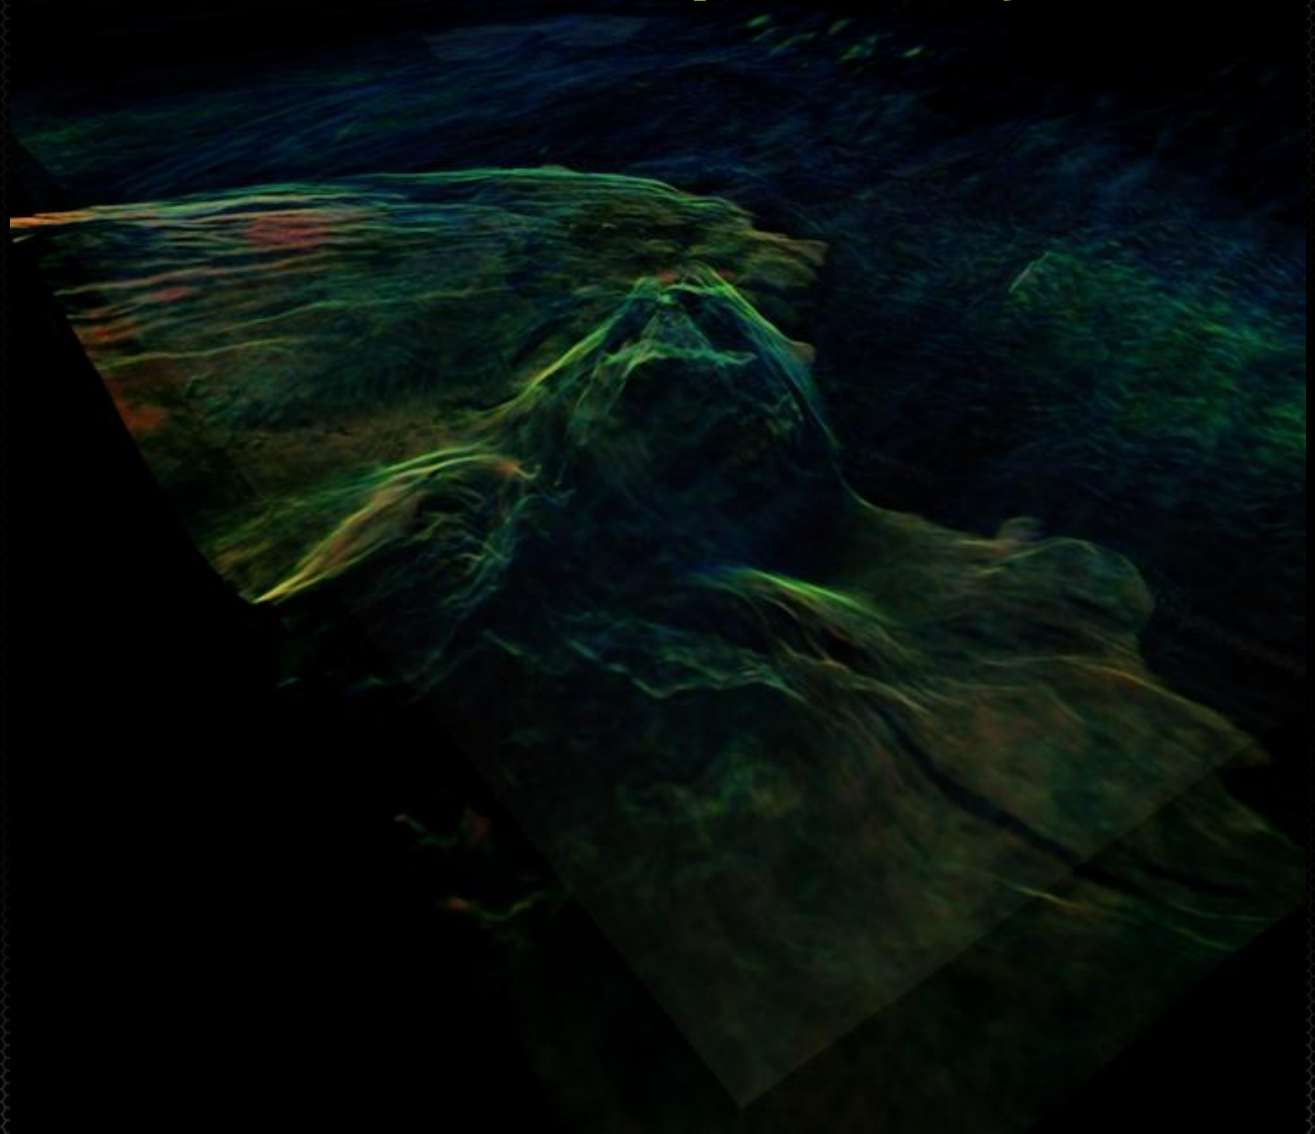

NVIDIA.

Section re-cap

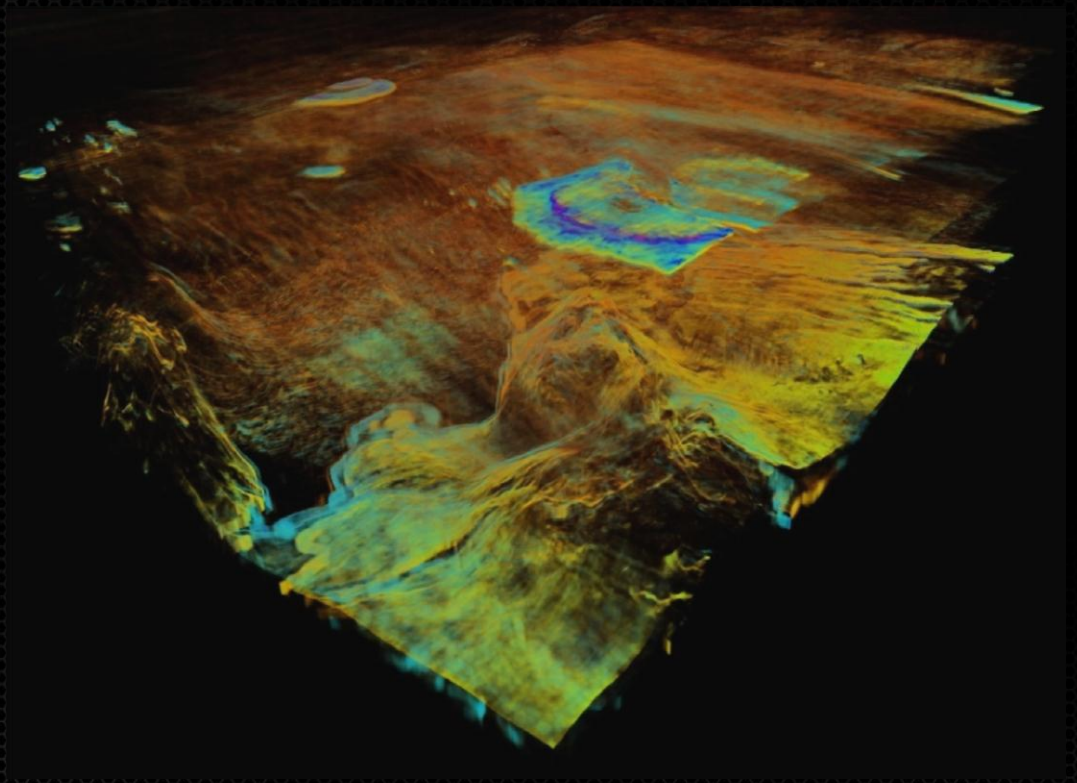

Energy Exploration challenges
Need to improve operational efficiency
Use simulations to reduce risk

GPUs deployed in datacenters
Better drilling decisions
Transforming E&P workflow

Leverage GPUs for collaboration
Need compact, reliable operation
Must monitor/manage GPUs

NVIDIA.

R Associates Inc.

RAI GPU Solutions

Bill Carter, President
November 2013

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

GPU Drivers in Seismic Exploration

- Kirchhoff Time Migration (KTM)
- Kirchhoff Depth Migration (KDM)
- Wave Equation Migration (WEM)
- Spectral Inversion
- Reverse Time Migration (RTM)
- Tesla GPUs have been deployed:
  - Schlumberger
  - Chevron
  - TOTAL
  - Petrobras
  - Repsol

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

RAI
R Associates Inc.

Decisions to Drive GPU Clusters

- Choose Correct Hardware
- Ensure Space, Power and Cooling
- Assemble and Physical Deployment
- Head Node Installation
- Compute Node Installation
- Management and Monitoring
- Run Benchmarks and Applications

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

RAI
R Associates Inc.

GPU Servers

1U UP – Value

1017GR-TF

5017GR-TF

n *# of GPU per node*

1U/2U DP, Scalable, High Density

1027GR-72R2
1027GR-72RT2
1027GR-TR2
1027GR-TRT2

2027GR-TRFH
2027GR-TRFHT

1027GR-TRF+
1027GR-TRFT+

2027GR-TR2
2027GR-TRT2

1027GR-TRF
1027GR-TSF
1027GR-TRFT

1027GR-TQF
1027GR-TQFT

3U & Above – Powerful

7047GR-TPRF

GPU FatTwin
F627G3-FT+
F627G3-F73+
F627G3-F73PT+
F627G2-FT+
F627G2-F73+
F627G2-F73PT+

6037R-72RFT+

GPU Blade
SBI-7127RG-E

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

DP GPU Server: SYS-1027GR-

- Motherboard: **X9DRG-HF+II/-HTF+II**
- 1U Chassis: **118GH-R1K65B**

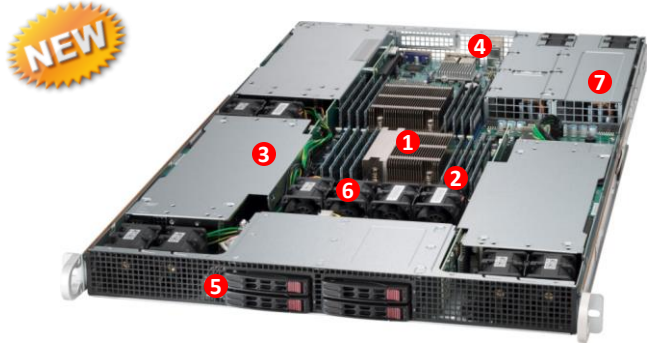

1	**Processor Support**	Dual Intel® Ivy Bridge EP (Socket R) series CPU
2	**Memory Capacity**	16 DIMM, Max of 512GB Reg. ECC DDR3 up to 1600MHz
3	**Expansion Slots**	3 PCI-e x16 Gen 3 for double width GPU cards 1 x8 (in x16 slot) LP card
4	**I/O ports**	1 VGA, 2 Gbit LAN, 2 USB 2.0, and 1 IPMI dedicated LAN port.
5	**Drive Bays**	4 hot-swap 2.5” drives bays
6	**System Cooling**	10 counter rotating fans with optimal fan speed control 1 air shroud
7	**Power Supply**	1600W Platinum level efficiency redundant power supply

80 PLUS PLATINUM

Key Features

- **Support up to 3 double width GPU cards (K10/K20M/K20X/K2/K1/Xeon Phi)**
- **16 DIMM, up to 512GB memories**
- **SAS 2208 controller**
- **Platinum level 1600W power supply**

RAI R Associates Inc.

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

DP GPU Server: SYS-2027GR-TR2/TRT2

- Motherboard: **X9DRG-HF+II/-HTF+II**
- 2U Chassis: **218GH-R1K65B**

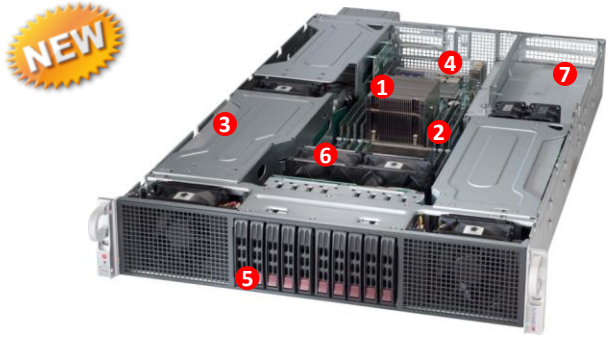

1	**Processor Support** Dual Intel® Ivy Bridge EP (Socket R) series CPU
2	**Memory Capacity** 16 DIMM, Max of 512GB Reg. ECC DDR3 up to 1600MHz
3	**Expansion Slots** 4 PCI-e x16 Gen 3 for double width GPU cards 1 x4 (in x16 slot) and1 x8 (in x16 slot) LP card
4	**I/O ports** 1 VGA, 2 Gbit LAN, 2 USB 2.0, and 1 IPMI dedicated LAN port.
5	**Drive Bays** 10 hot-swap 2.5” drives bays
6	**System Cooling** 5 counter rotating fans with optimal fan speed control 1 air shroud
7	**Power Supply** 1600W Platinum level efficiency redundant power supply

80 PLUS PLATINUM

Key Features

- **Support up to 4 or 6 double width GPU cards (K10/K20M/K20X/K2/K1/Xeon Phi)**
- **16 DIMM, up to 512GB memories**
- **Platinum level 1600W power supply**
- **Optional rear fans to support 300W GPU cards**
- **Optional two by 8 rear I/O add on card slots**

Key Application

- **Computational Finance**
- **HPC/Oil and gas**
- **Weather and Climate Analysis**

RAI
R Associates Inc.

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

GPU Blade: SBI-7127RG-E (20 GPUs)

Dual Port QDR/FDR IB or 10GbE Support

Integrated BMC – IPMI 2.0, KVM-over-IP, Remote Virtual Media

Highest GPU Blade Density
20 GPU + 20 CPU per 7U

1x SSD or
1x SATA-DOM

Greater TCO
No Cables
System Mgmt.
System Maintenance

Intel® Ivy Bridge
Qty:8 DDR3 DIMM Slots

Enhanced 3000W redundant power supply

2 Kepler

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

The RAI Advantage

- Experience
- Breadth of Partnerships
- Breadth of Solutions
- Dedicated Engineering
- Tested Segment-specific Solutions
- Small, Nimble, Flexible
- Innovation in our solutions
- Commitment to You!!

© Copyright 2013 R Associates Incorporated (RAI). The information contained herein is subject to change without notice, and is **RAI CONFIDENTIAL.** For use only under Confidentiality Agreement.

RAI
R Associates Inc.

Thank you

Rewriting the Rules with GPUs:
Solutions for Energy Exploration and Data Processing

Ian Lumb
Bright Evangelist

About Bright Cluster Manager

- Approach
  - Designed and written from the ground up
  - Single cluster management daemon/agent
    - Provides all functionality
  - Single, central database for configuration and monitoring data
  - Single CLI and GUI for ALL cluster management functionality
- Outcome: Bright Cluster Manager is ...
  - Extremely easy to use
  - Extremely scalable
  - Secure & reliable
  - Complete
  - Flexible
  - Maintainable

Bright Cluster Manager – Elements

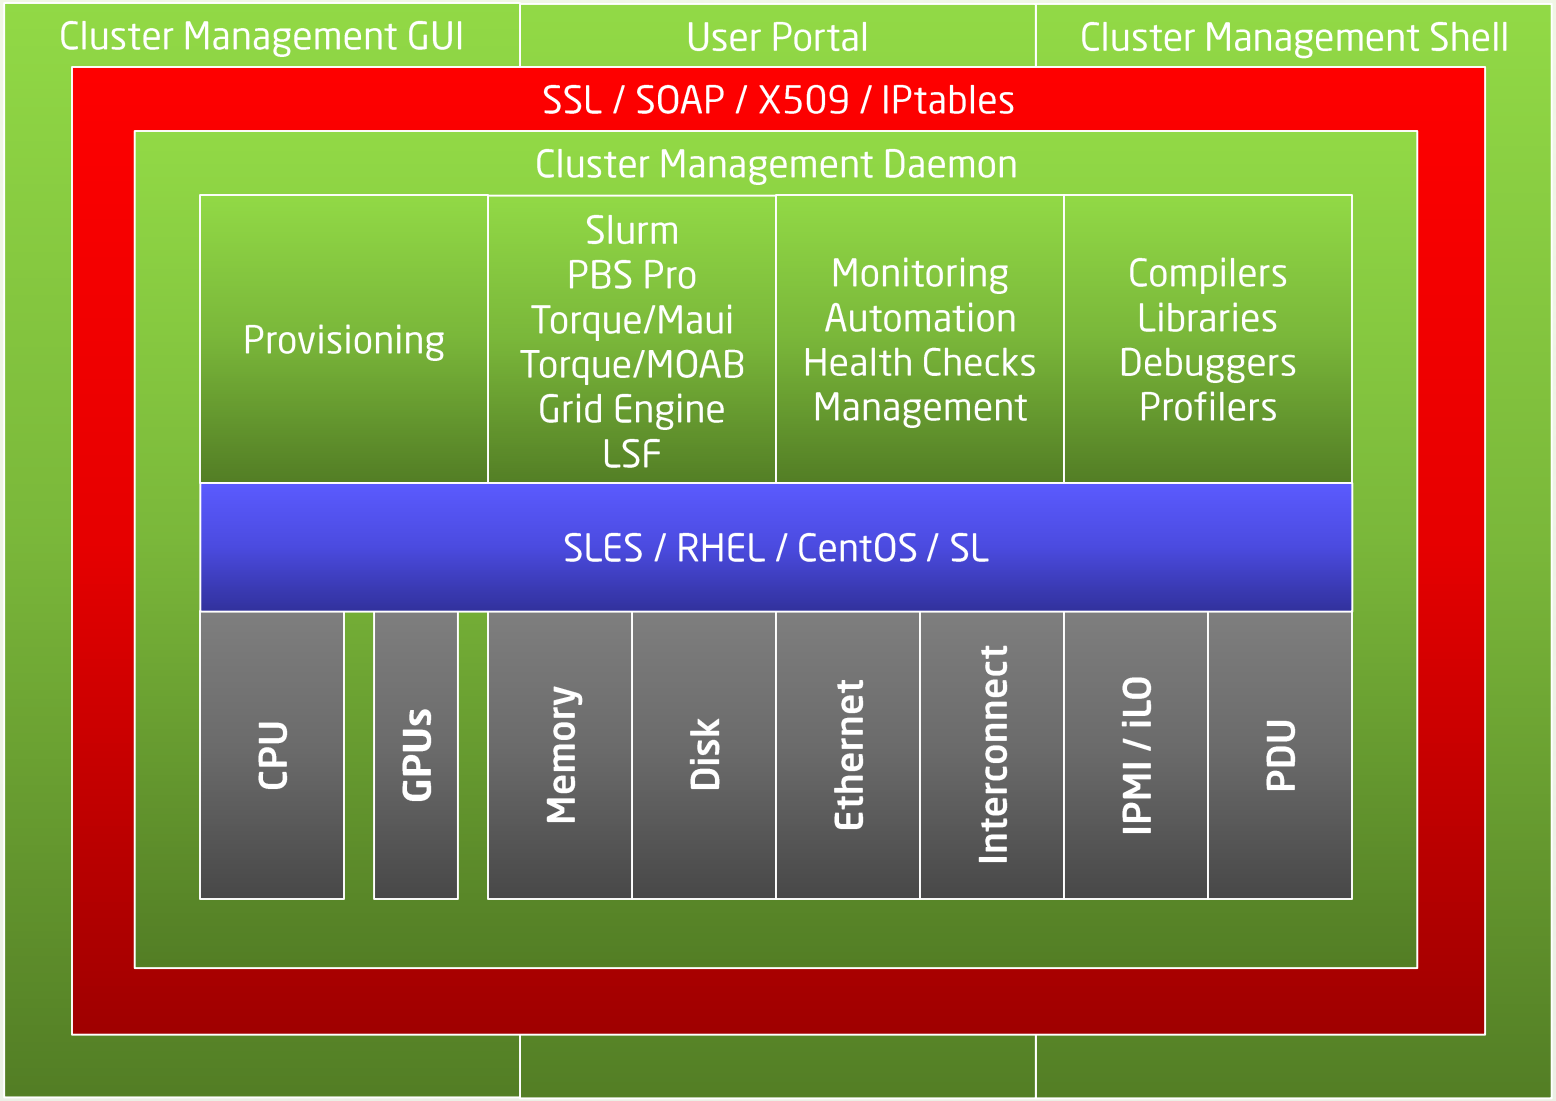

GPU Management in Bright (1)

Bright Computing

- CUDA and NVIDIA drivers distributed as Bright packages
  - Available through YUM/Zypper
- NVIDIA driver recompiled automatically against running kernel on node boot-up
  - Driver-kernel version mismatch eliminated
- Switch between CUDA versions using **modules** environment
- GPUs identified as **consumable resources**
  - Workload Management System (WLM)
    - Users can request GPUs
      - All popular WLMs supported
- NVIDIA GPUs easily integrated into cluster
  - Admins can manage and monitor
  - Users can request GPUs in their compute jobs
    - Use CUDA or OpenCL to offload computations

Provisioning

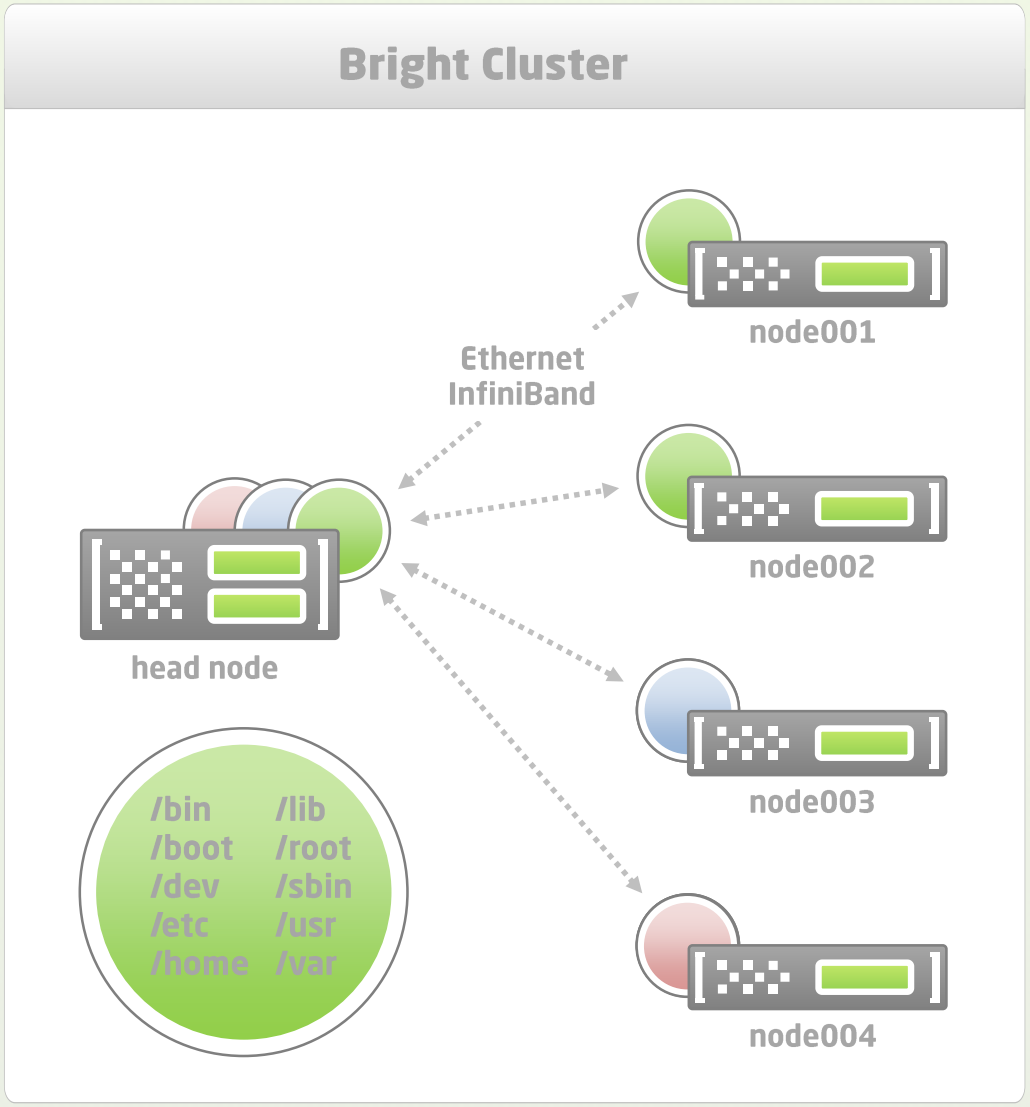

Features

- Create: clone, from distri, from server
- Modify: add, remove, edit file; rpm --root; yum; chroot
- Sync: to and from
- Where: to HDD, RAM disk or VM
- How: full install, incremental, live
- Scalability: off-loadable, linear scaling

GPU Management in Bright (2)

Bright Computing

- All **metrics** exposed by GPU are available
  - Fan speeds, temperatures, utilization, memory usage, etc. can be graphed
    - Different GPUs support different metrics
      - … cumbersome to keep overview up-to-date
  - Can be used in monitoring rules that take action automatically
- Bright also supports **GPU Units**
  - GPUs are exported to nodes
  - Monitoring data for the GPU Unit combines
    - Data collected from the GPU Unit chassis
    - GPU data from the nodes that GPUs are exported to

Architecture – Monitoring

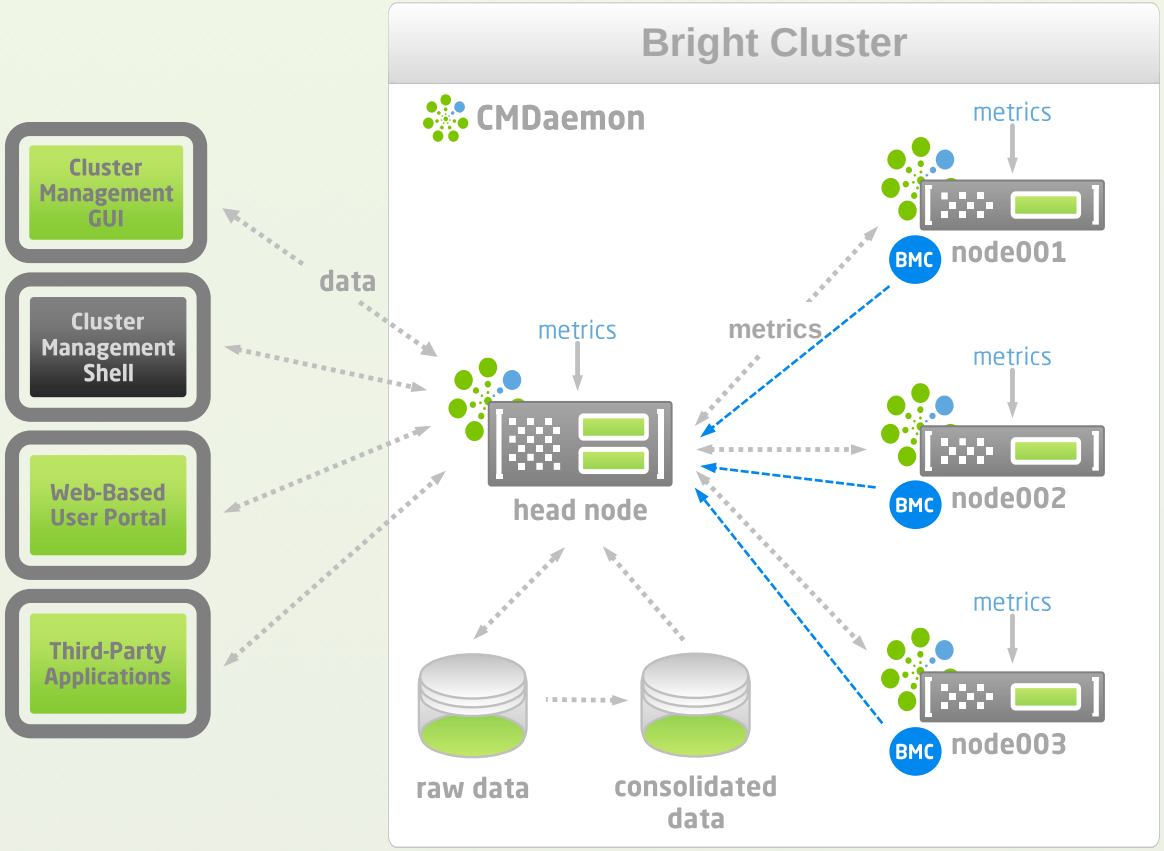

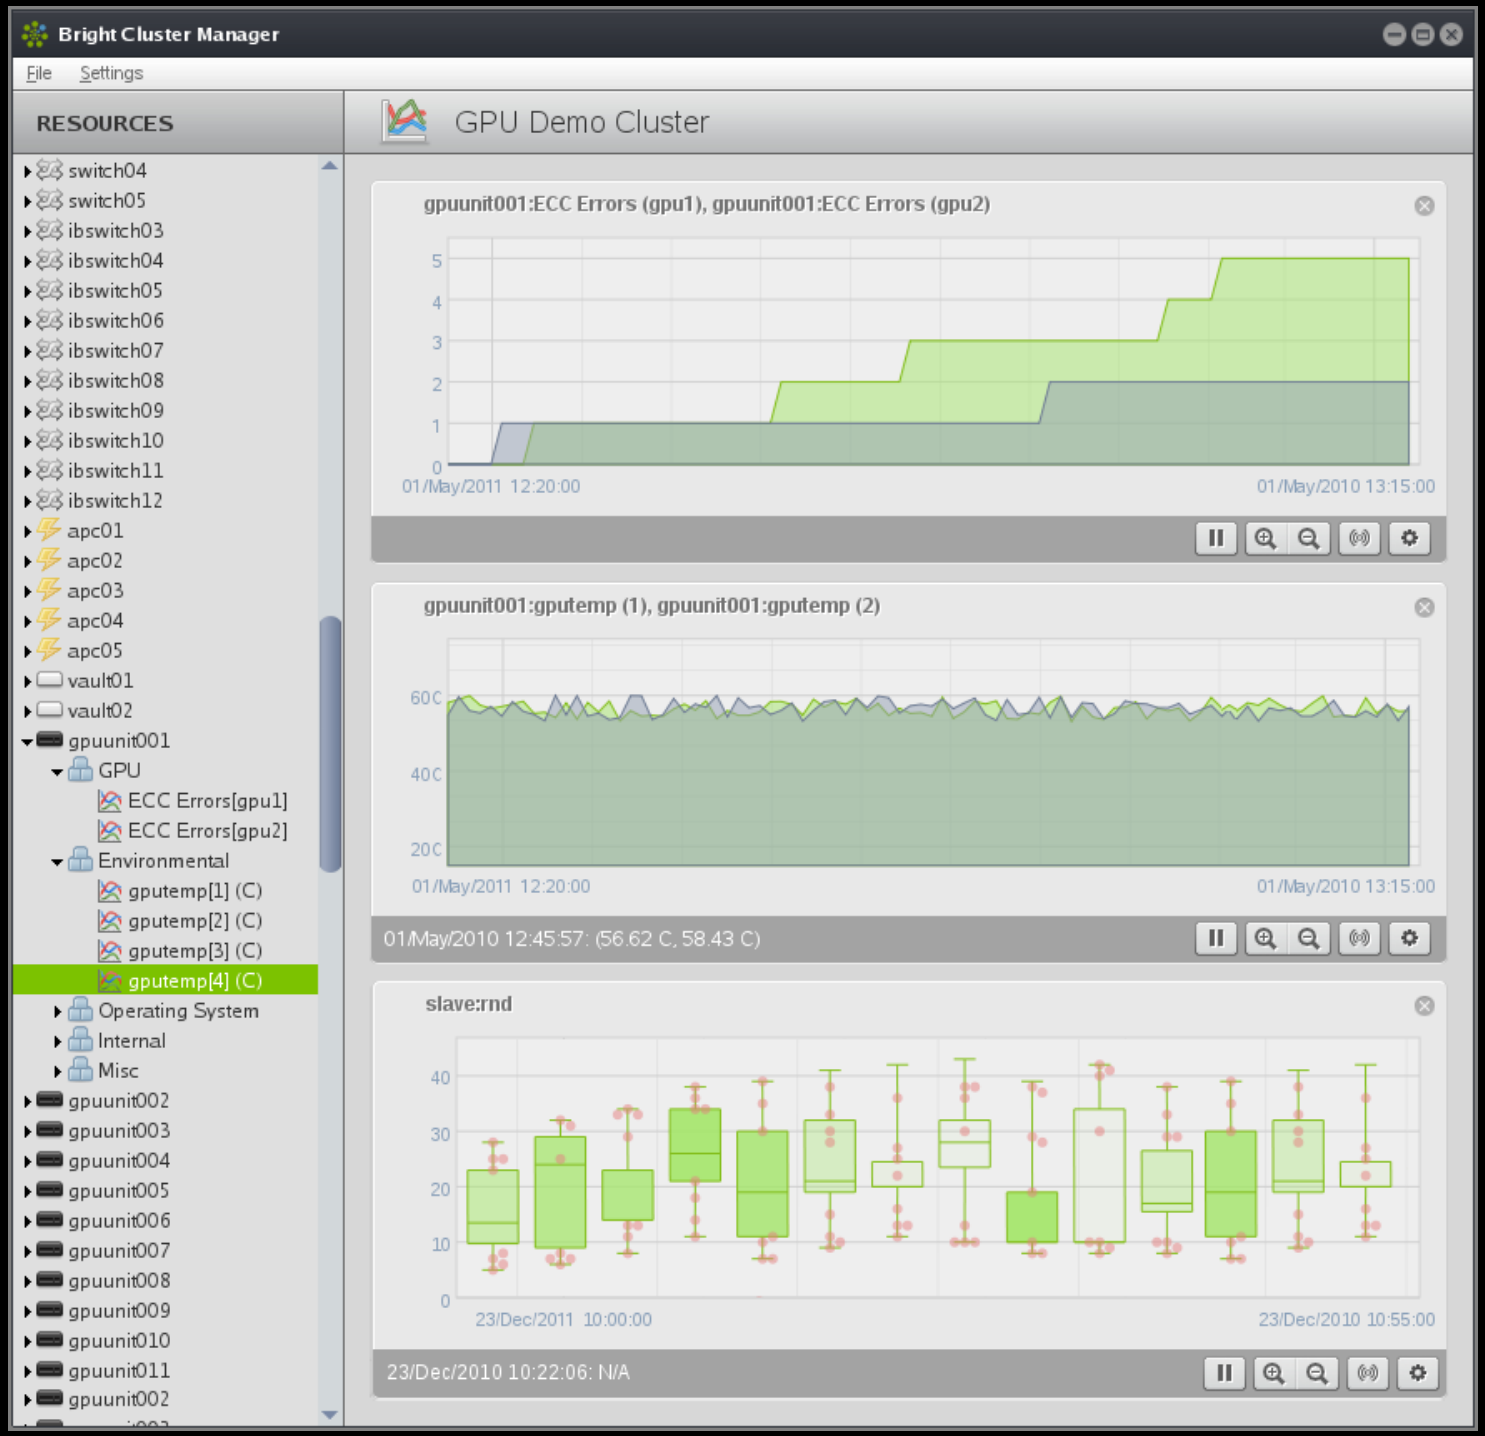
Bright Cluster Manager
File
Settings
RESOURCES
GPU Demo Cluster
switch04
switch05
ibswitch03
ibswitch04
ibswitch05
ibswitch06
ibswitch07
ibswitch08
ibswitch09
ibswitch10
ibswitch11
ibswitch12
apc01
apc02
apc03
apc04
apc05
vault01
vault02
gpuunit001
GPU
ECC Errors[gpu1]
ECC Errors[gpu2]
Environmental
gputemp[1] (C)
gputemp[2] (C)
gputemp[3] (C)
gputemp[4] (C)
Operating System
Internal
Misc
gpuunit002
gpuunit003
gpuunit004
gpuunit005
gpuunit006
gpuunit007
gpuunit008
gpuunit009
gpuunit010
gpuunit011
gpuunit002
gpuunit001:ECC Errors (gpu1), gpuunit001:ECC Errors (gpu2)
5
4
3
2
1
0
01/May/2011 12:20:00
01/May/2010 13:15:00
gpuunit001:gputemp (1), gpuunit001:gputemp (2)
60C
40C
20C
01/May/2011 12:20:00
01/May/2010 13:15:00
01/May/2010 12:45:57: (56.62 C, 58.43 C)
slave:rnd
40
30
20
10
0
23/Dec/2011 10:00:00
23/Dec/2010 10:55:00
23/Dec/2010 10:22:06: N/A

Additional Support for GPUs in Bright 6.1

- Additional configuration options
  - ECC mode
  - Compute mode
  - Operation mode
  - Persistence mode
  - Power limit
  - Clock speeds

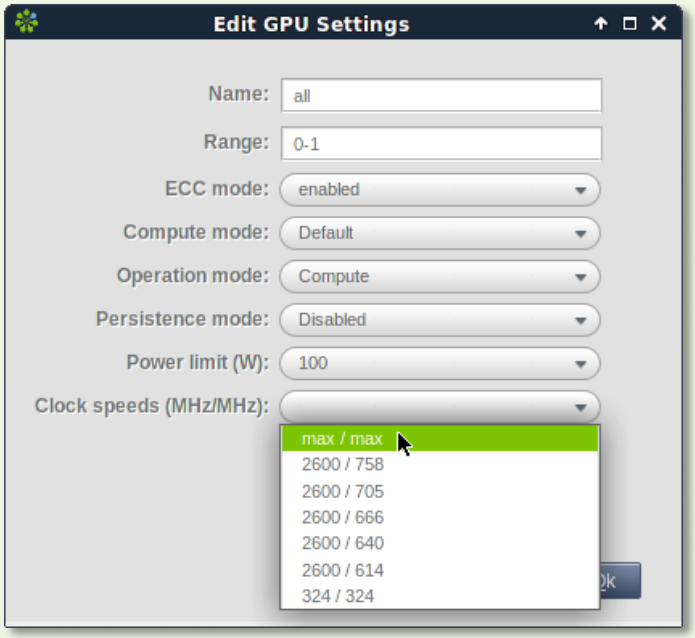

- Two additional nvidia-healthmon based health checks to verify that GPU is working properly:
  - Quick check which can be run periodically
  - Extensive check which should not be run concurrently with jobs

Bright in Lockstep with CUDA …

- *Innovation* characterizes the entire history and evolution of GPU programmability through CUDA
  - BUT … introduces challenges and opportunities …
- Bright Computing's approach leverages
  - People
    - Proactively maintaining business and technical relationships
  - Process
    - `Hands-on engineering' begins with CUDA release candidates
      - Preliminary to fully productized implementations (e.g., MPS)
  - Product
    - Bright Cluster Manager released once per year
      - Updates flow *continuously* …
- Key takeaways
  - CUDA is a *platform* for GPU computing
  - Bright Computing is an *agile* ISV

Bright Computing

Bright Cluster Manager makes it *easy* to install, manage and use clusters and clouds with NVIDIA GPUs.

Booth #1725

Questions?

Ian Lumb

ian.lumb@brightcomputing.com

Bright Computing

Bright Cluster Manager:

A CUDA-Ready Management Solution for GPU-Based HPC

Ian Lumb
Bright Evangelist

Last week via GTC Express!