Designing for the Future –

GPU Solutions that Maximize Performance, Density and Energy Efficiency

Bill Chen

Director, Solutions Optimization Engineering

July, 2014 @ GTW Singapore

Supermicro Profile

Global Footprint:	**>80 Countries**
Years Profitable:	**21 Years (since day one, 1993)**
Production:	**Facilities in the US, Asia and EMEA**
Customers:	**Channel, SI/VAR, OEM direct**
Corporate Focus:	**Architecture Innovation, Energy Efficiency, Total Solution**

Performance & Efficiency – Industry Trends

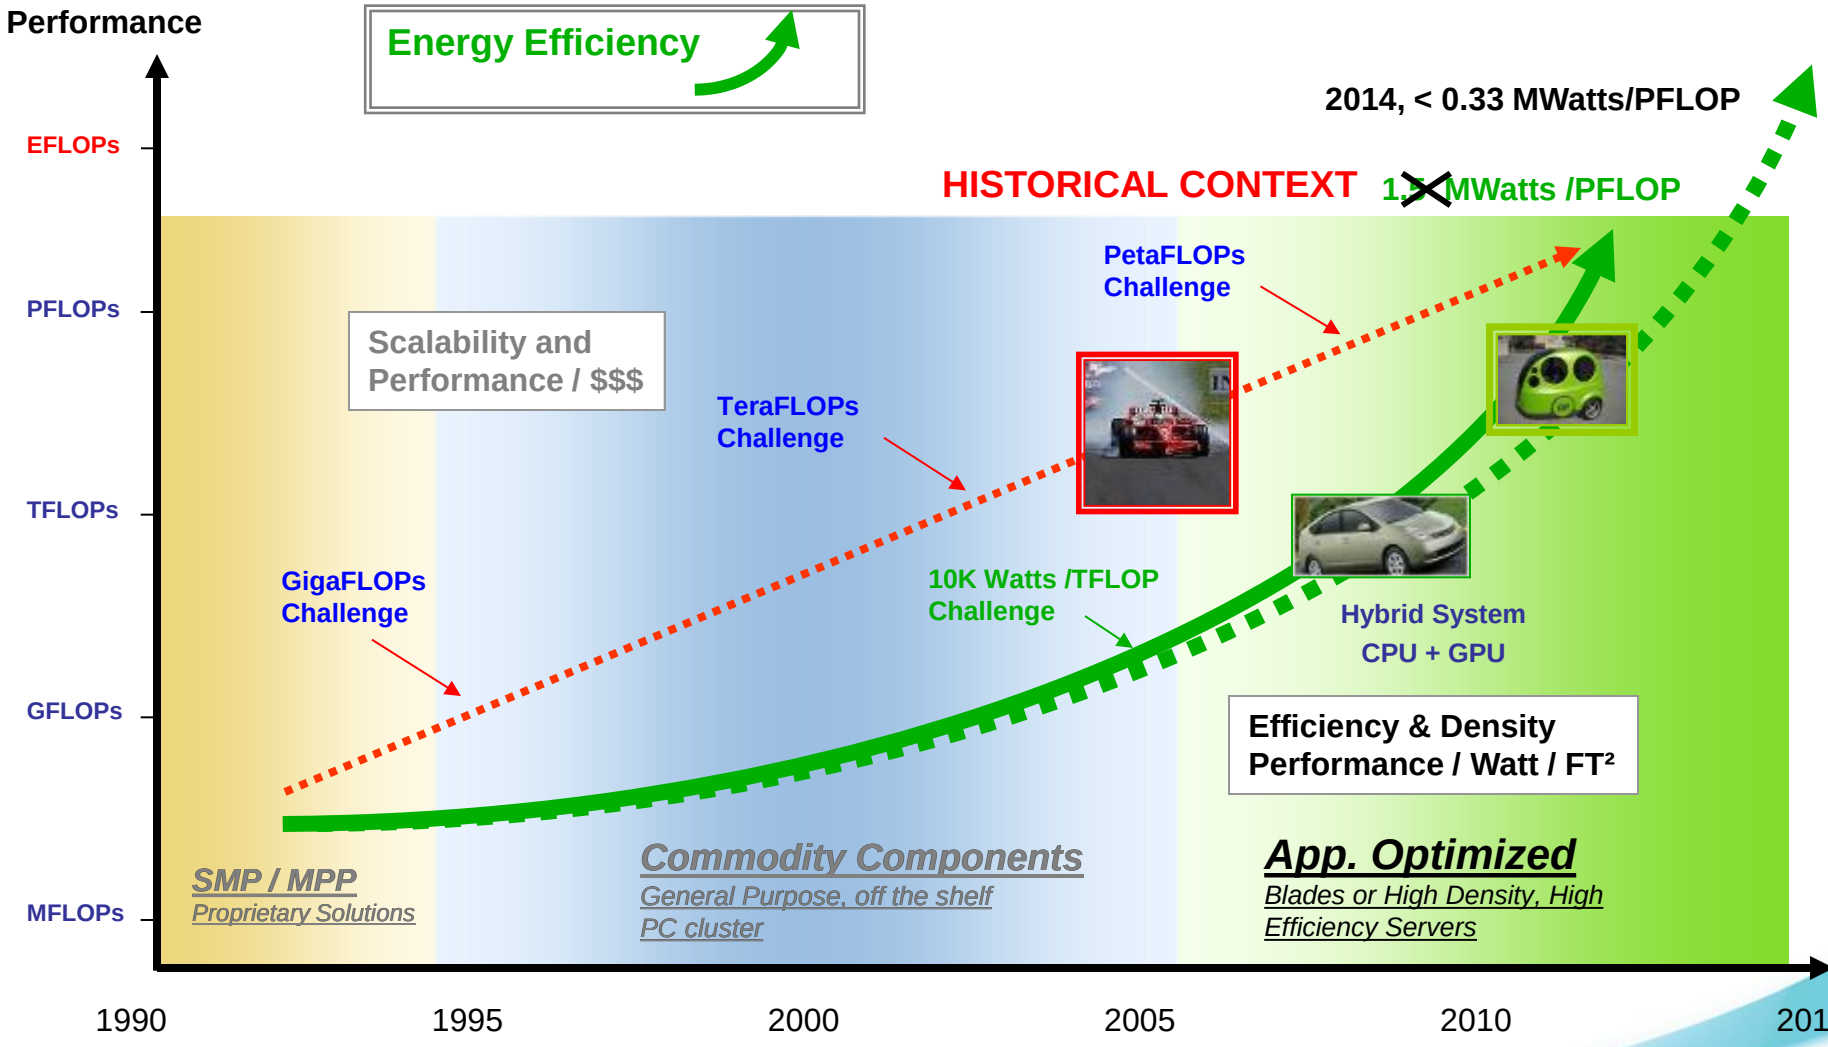

GPU Accelerated Computing on Top500 / Green500

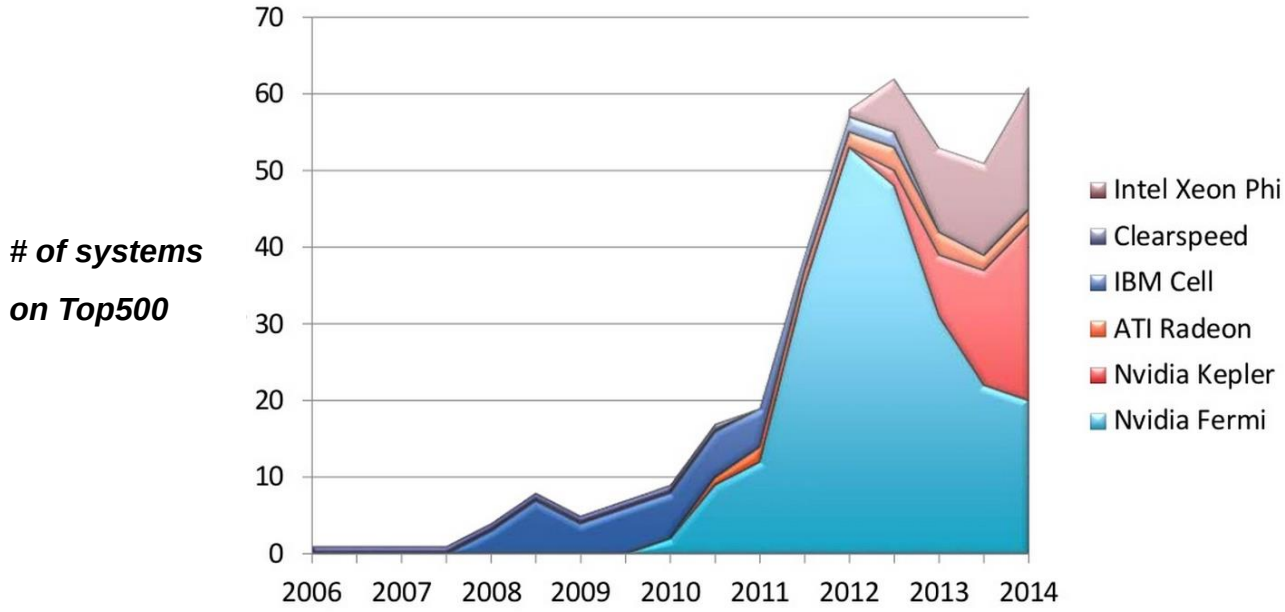

- **HPC Trend**: > 80% HPC sites have been use processor/co-processor/accelerator for either exploratory or production **(IDC: grows from 28.8% 2011 to 76.9% 2013)**
- **Performance and Efficiency** (performance per watt): 17 of the top Green500 list in June'14 are GPU/Co-processor accelerated HPC systems
- **Moving to the Top**: the greenest supercomputer (on Green500 list) is TSUBAME-KFC submerged GPU cluster: >4 MFLOPS/w
- **Wide Adoption**: GPU applications are beyond HPC: such as Finance, Gaming, Vitalization (E.g. VDI)…

GPU Computing Beyond HPC

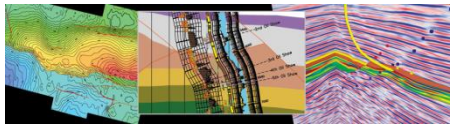

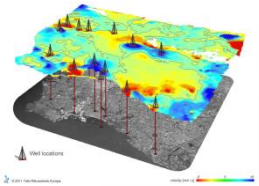

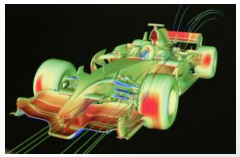

Oil and Gas/Seismic

- Seismic imaging
- Seismic Interpretation
- Reservoir Modeling
- Seismic Inversion

Weather and Climate

- Weather
- Atmospheric
- Ocean Modeling
- Space Sciences

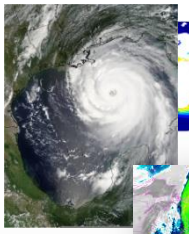

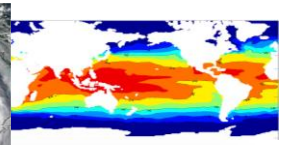

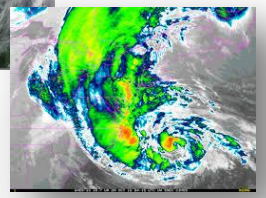

Scientific

- Computational fluid dynamics
- Materials science
- Molecular dynamics
- Quantum chemistry

Simulation & Creation Design

- Mechanical design & simulation
- Structural mechanics
- Electronic Design Automation

Data Mining

- Data parallel mathematics
- Extend Excel with OLAP for planning & analysis
- Database and data analysis acceleration

Entertainments

- Online gaming (Gaming Grid)
- Movie rendering / animation
- Video streaming / image processing

Massively parallel architecture accelerates scientific & engineering applications

Computational Finance & simulation

- Options pricing
- Risk analysis
- Algorithmic trading

Imaging and Computer Vision

- Medical imaging
- Visualization & docking
- Filmmaking & animation

FatTwin

TwinPro

SuperBlade

GPU / Xeon Phi

MicroCloud

MicroBlade

Storage

Embedded

Switch

Software

Complete Rack Solutions

GPU GRID for Virtualization, Gaming & Enterprise

- Industry's most comprehensive, power efficient and densest GPU solutions
- The first NVIDIA GRID-certified GPU-systems on the market

POWERED BY NVIDIA GRID

1U UP – Value

2 — 1017GR-TF

2 — 5017GR-TF

n — No. of GPUs per Node

1U/2U DP, Scalable, High Density

3 — 1027GR-72R2+, 1027GR-72RT2+, 1027GR-TR2+, 1027GR-TRT2+

6 — 2027GR-TRFH, 2027GR-TRFHT

2 — 1027GR-TRF+, 1027GR-TRFT+

4 — 2027GR-TRF, 2027GR-TRFT

3 — 1027GR-TRF, 1027GR-TSF, 1027GR-TRFT

4 — 1027GR-TQF, 1027GR-TQFT

3U & Above – Powerful

4 — 7047GR-TPRF

3 — GPU FatTwin

2 — 6037R-72RFT+

2 — GPU Blade

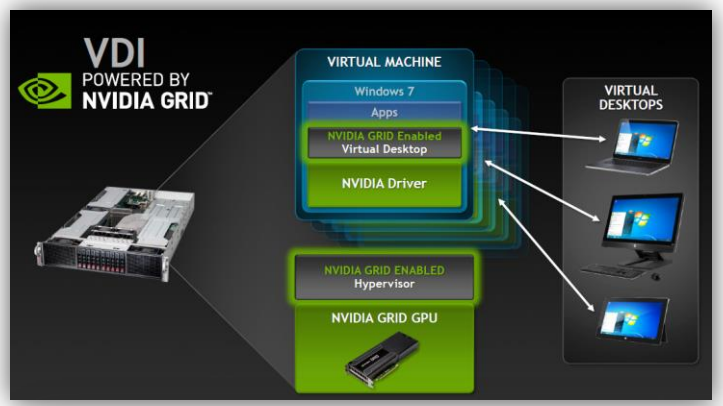

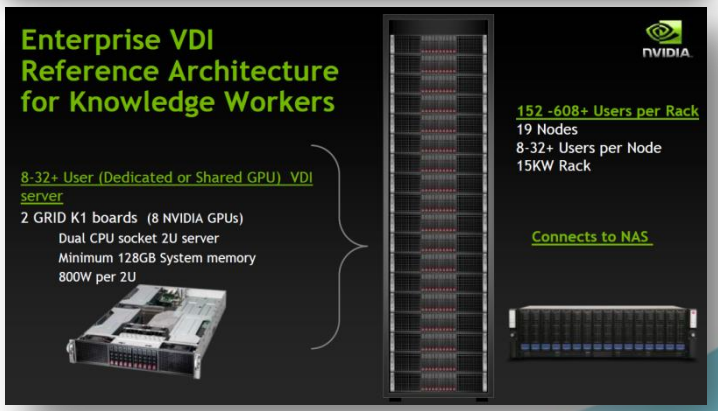

Supermicro GPU Solution Evolution

The most comprehensive product line in the Industry

NVIDIA

GPGPU Where it started…

Telsa S1070

1U 4-GPU Standalone box

PCI-E x16

1U Twin™

Hybrid Computing Pioneer

The fastest 1U server in the world

The most powerful PSC

Efficiency

Ultra High Efficiency

2U GPU w/ QDR IB onboard

2U Twin

Density

7U GPU Blades
20 CPUs + 20 GPUs

2U 4-GPU

1U 3-GPU

Mainstream

3 GPUs Blade

4 GPUs Workstation / 4U

X9 (UP) 1U 2-GPUs

1U 4-GPUs

4U 8-GPUs

ıtTwin™ 2-node
6 GPUs or MICs per node

FatTwin™ 4-node
3 GPUs per node

2U 6-GPUs

2008 | 2009 | 2010 | 2011 | 2012 | 2013 | 2014

FatTwin | TwinPro | SuperBlade | GPU / Xeon Phi | MicroCloud | MicroBlade | Storage | Embedded | Switch | Software | Complete Rack Solutions

GPU Optimized System Lineup

❖ SuperServer (Passive Cooling)

1U UP – Value

2

1017GR-TF

2

5017GR-TF

1U/2U DP – Scalable, High Density

4

1027GR-TQF
1027GR-TQFT

6

2027GR-TRFH
2027GR-TRFHT

3

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

4

2027GR-TRF
2027GR-TRFT

3U & Above – Performance

4

7047GR-TPRF

3

FatTwin

2

6037R-72RFT+

2

SuperBlade

❖ Workstation (Active Cooling)

1

5037A-iL

2

5037A-i

1

7037A-iL

2

7037A-i

2

7047A-T

4

7047GR-TRF

Designing GPU Optimized Systems

- **Performance**
  - PCI-e lanes arrangement, PCB placement, interconnection…
  - CPU, MEM, I/O, Networking, Storage…
- **Mechanical design**
  - mounting, location, space utilization
- **Thermal**
  - air flow, fan speed control, location, noise control
- **Power supply**
  - PSU efficiency, wattage options, power monitoring & management
  - Number of power connectors (& location)

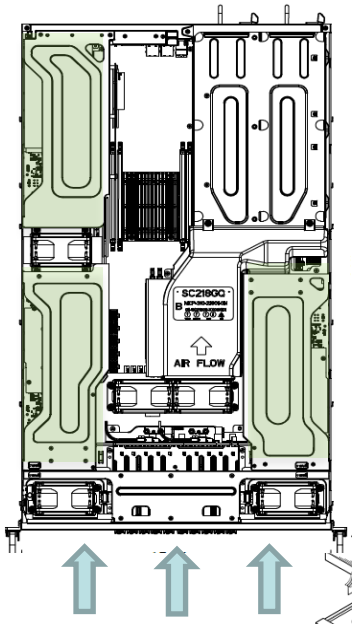

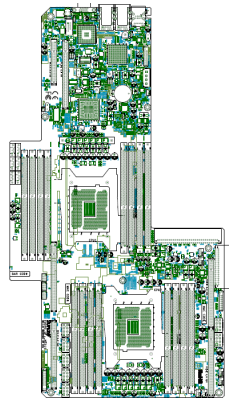

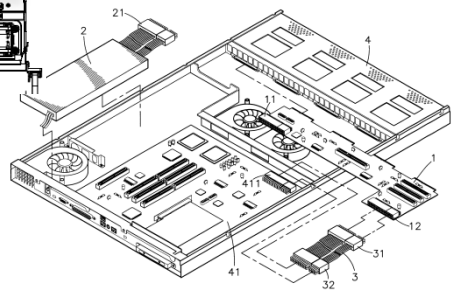

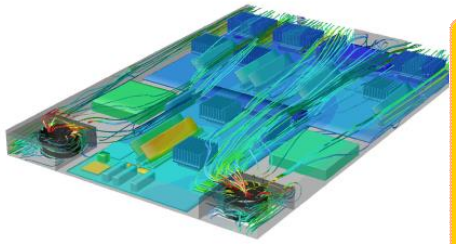

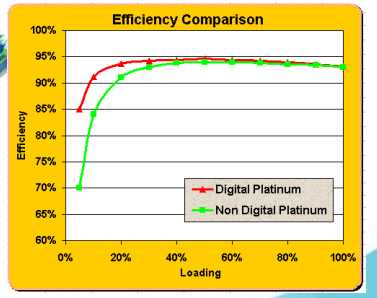

Design for Performance

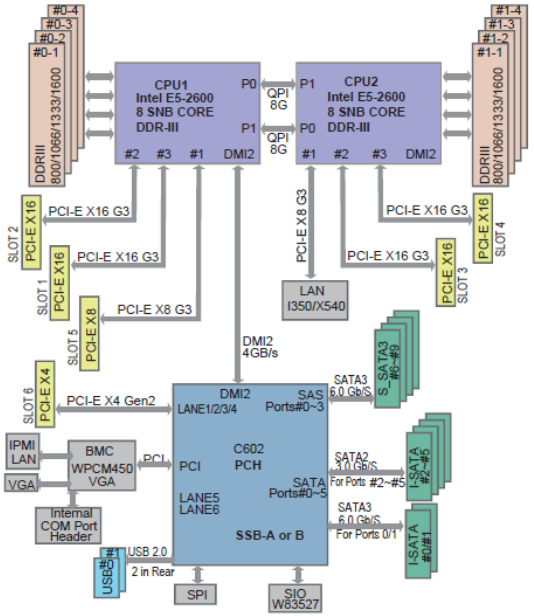

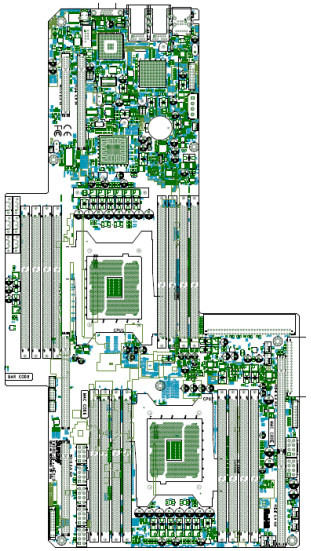

Platinum level high efficient **1800W** power supply

80 PLUS PLATINUM

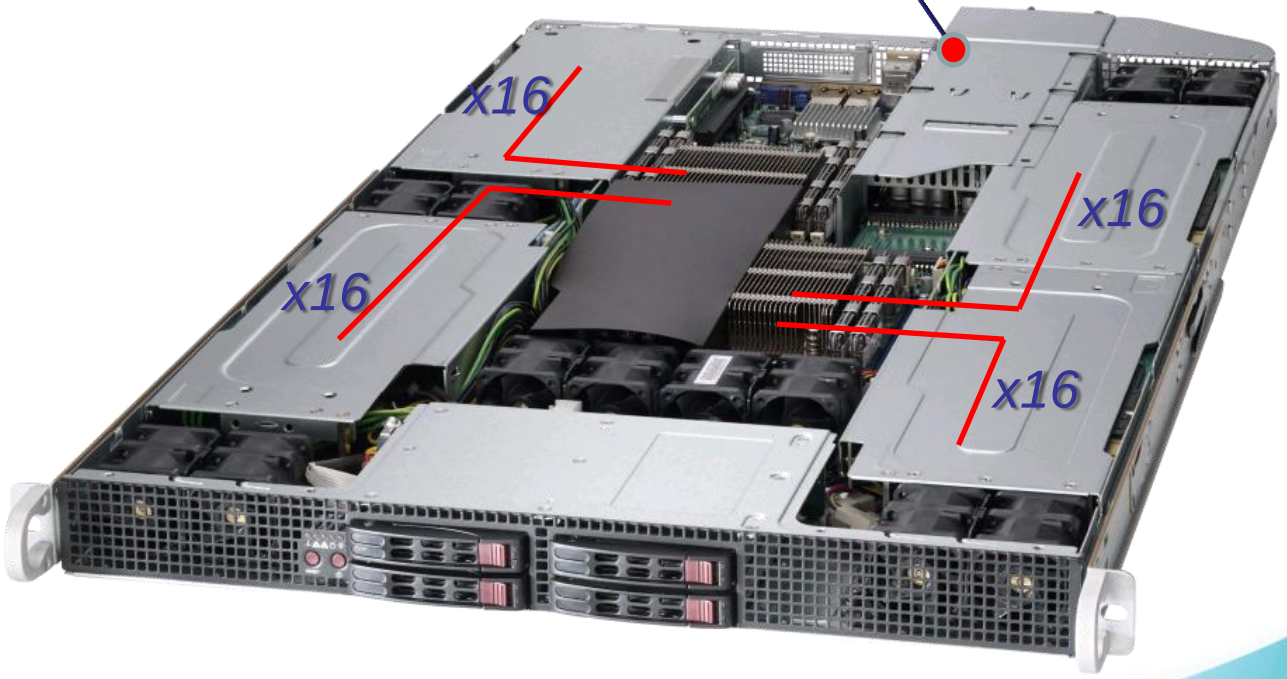

Communication Between GPUs

IB Switch

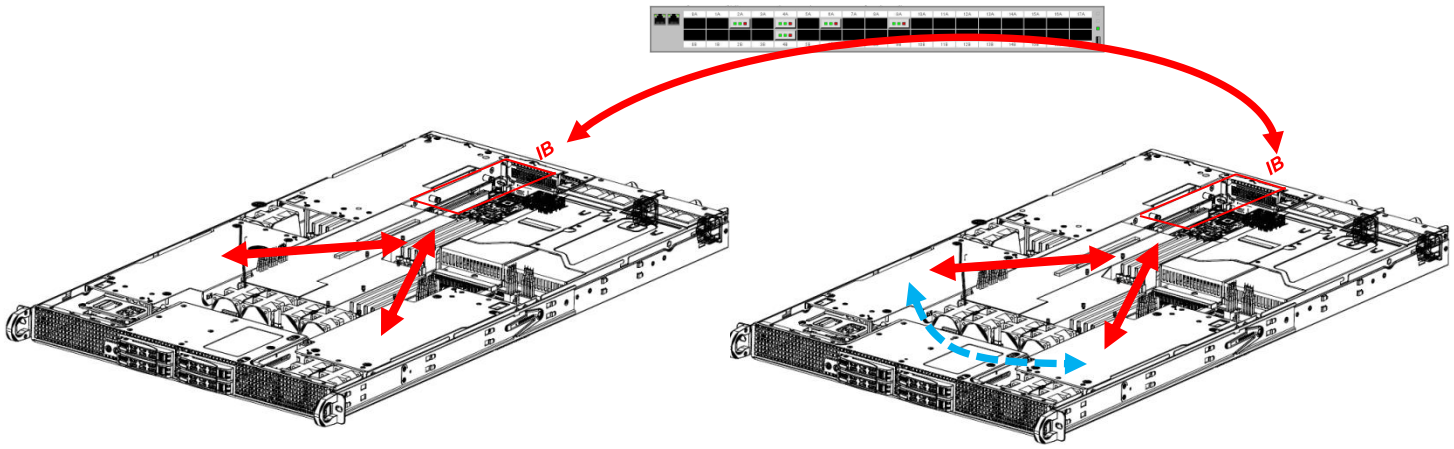

Implementation Example

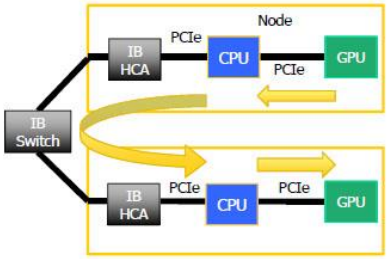

The model used by existing CPU-GPU Heterogeneous architectures for GPU-GPU communication. Data travels via CPU & Infiniband (IB) Host Channel Adapter (HCA) and Switch or other proprietary interconnect

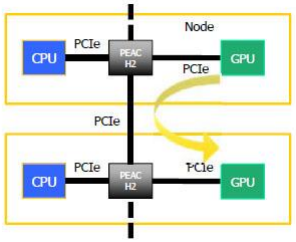

Data transfer between cooperating GPUs in separate nodes in a TCA cluster enabled by the PEACH2 chip.

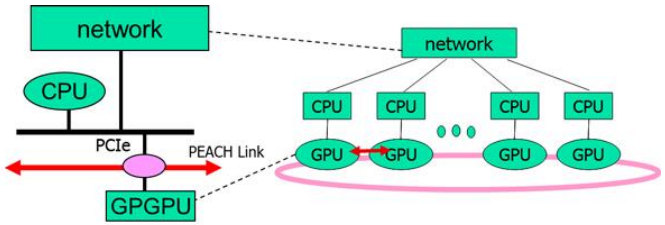

Schematic of the PEARL network within a CPU/GPU cluster

Source: Tsukuba University

Power Supply

- **High efficiency power supplies**
  - 95% platinum
  - Wattage choices & configurability
  - Redundancy & BBP support
- **Power management software**
  - Power capping
  - Core speed control for power management

1000W (BBP)
Battery Backup Power

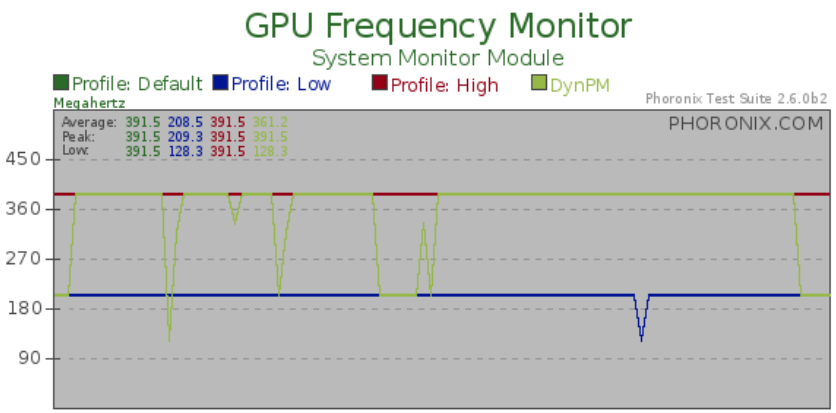

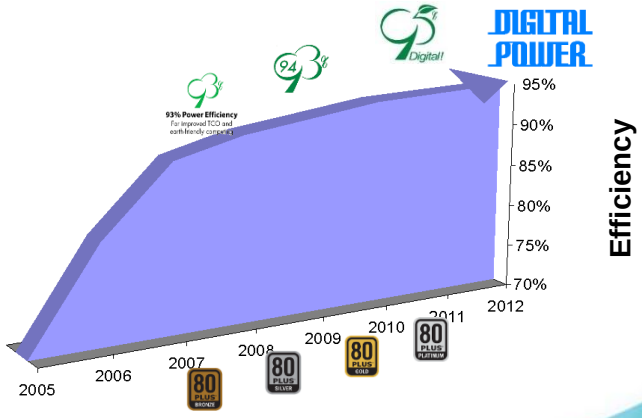

Platinum/Titanium (95%+) Digital Power Supply

Max. Power Requirements

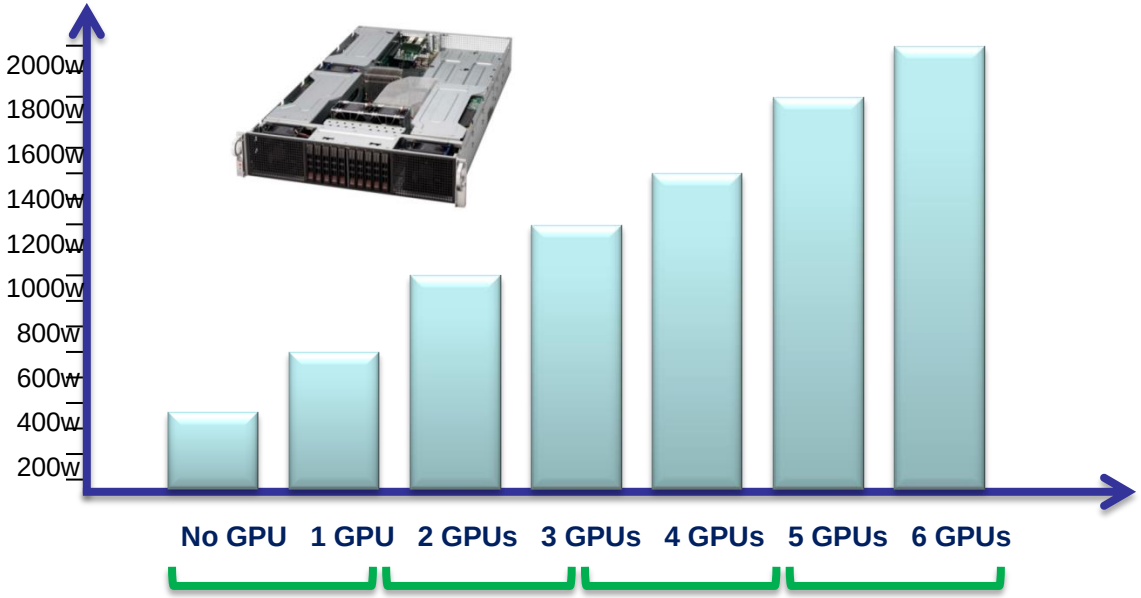

PSU Loading ***20%*** ***40%*** ***60%*** ***80%***

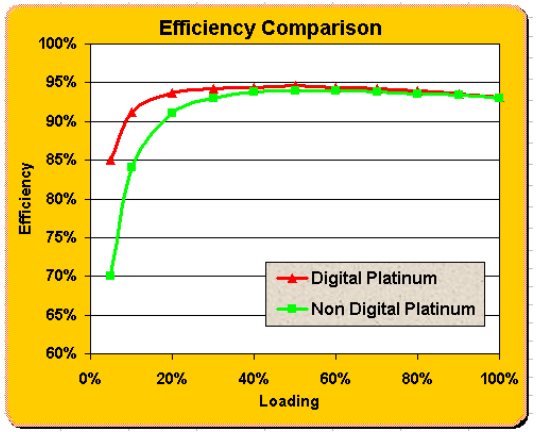

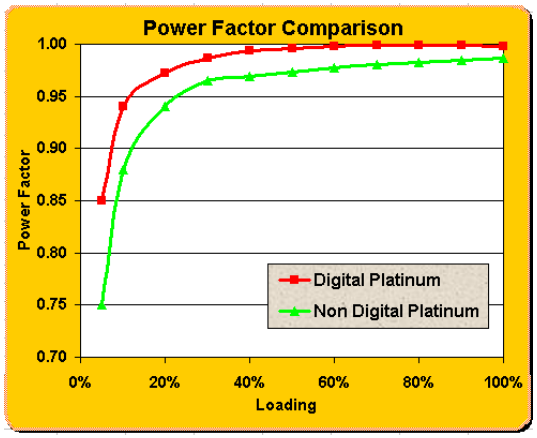

Digital Switching Power Supplier (95+%)

**** maintain high efficiency even at low loading***

Configurable Power Supplies

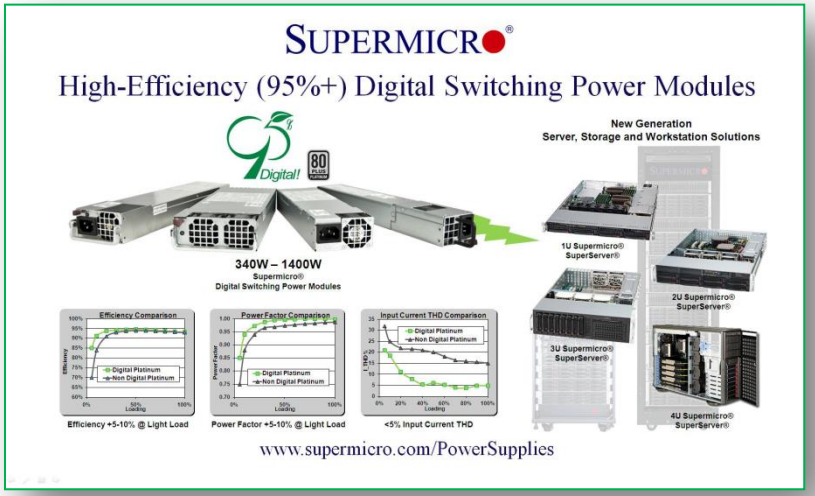

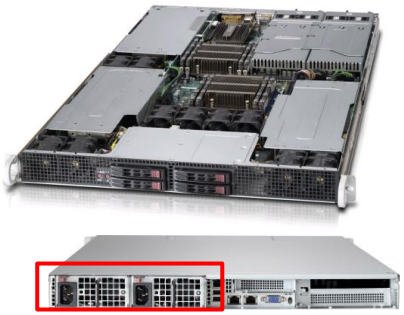

1+1 or 2+1

- ✓ *Standardize power supply module*
- ✓ *Design multiple capacity options (240W ~ 2000W)*
- ✓ *Provide application-optimized & energy-efficient configurations*
- ✓ *Feature power management / control*

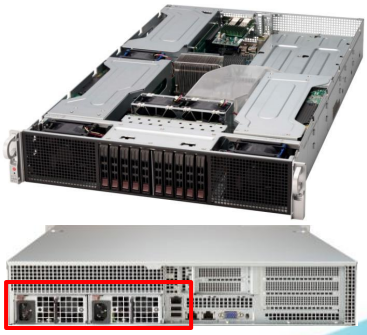

Thermal & Cooling Design

- Heatsink performance
- Passive & active
- High-performance Fan
- Fan speed control
- Multiple zones sensors
- Air shroud design
- Liquid cool

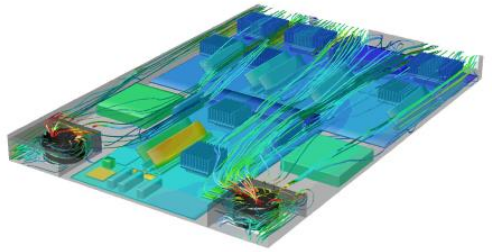

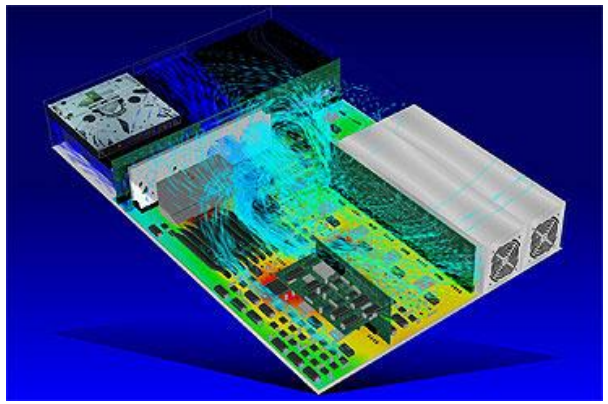

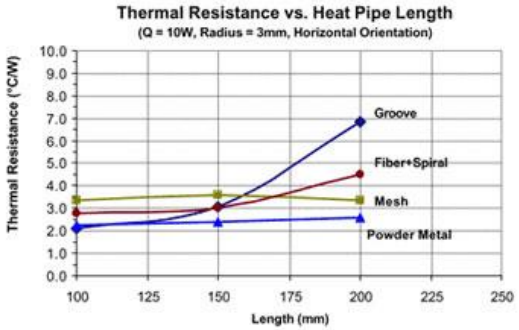

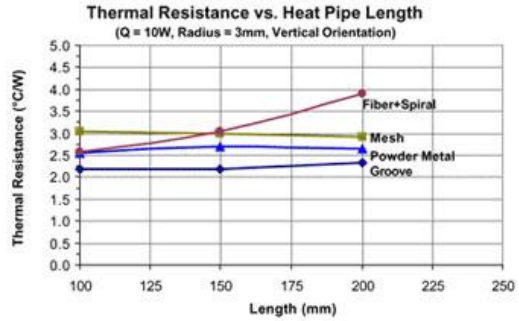

Optimized Airflow and Configurable Cooling

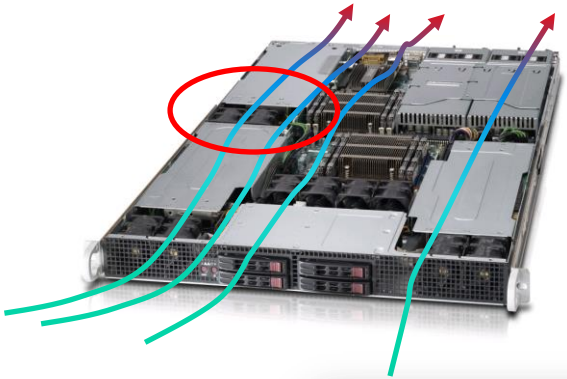

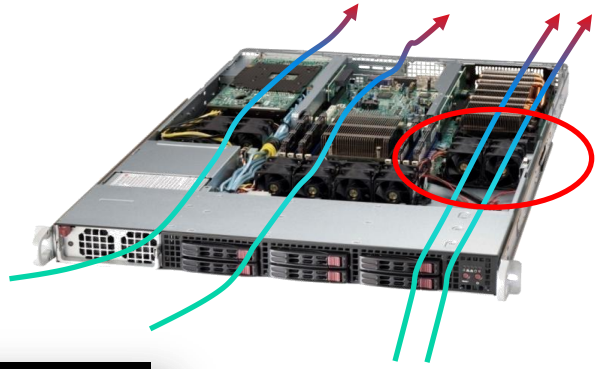

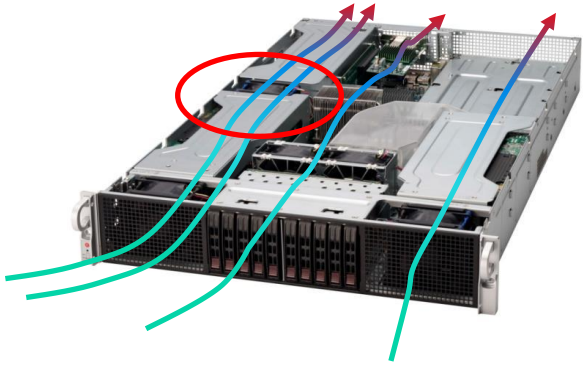

Workstation / 4U Server

(Accommodates both Active and Passive)

- ✓ *Consider total system-cooling design*
- ✓ *Remove unnecessary cooling component*
- ✓ *Enforce the hot zone airflow*
- ✓ *Provide application-optimized & energy-efficient configurations*

Water Cooling Example

Rack DCLC AHx™ - components used in the self-contained rack (CoolIT®)

Multiple Rack CHx

AHx Module
- Dual redundant fans
- Centralize pumping architecture
- CoolIT Command Center monitors/alerts on health of liquid system

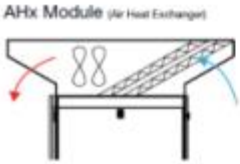

Manifold Module
- Steel body attaches similar to PDU
- All-metal dry-break quick connects

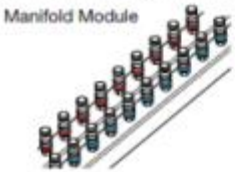

Server Module
- Passive cold plate technology
- All-metal dry-break quick connects

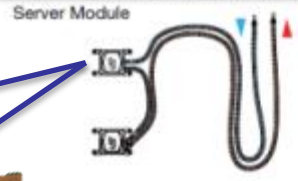

FatTwin™ GPU node

Cooling the both the CPUs and GPUs

- ***Example configuration: FatTwin GPU, 4U 4-node – 3 GPUs per node***
- ***Can be used in any form factors – 1U, 2U, 4U… GPU systems***
- ***Cold plates for CPUs and GPUs***
- ***Very low system fan speed for cooing other components***

Case Study – Submerged Liquid Cooling

- ✓ *Removed Fans and Heat Sinks*
- ✓ *Use SSD & Updated BIOS*
- ✓ *Reverse the handlers*

"Submerged Supermicro Servers Accelerated by GPUs"

- ❑ Supermicro 1U with
- ❑ No requirement for room-level cooling
- ❑ Operates at PUE ~ 1.12
- ❑ 25 kilowatts per rack – the breakpoint per rack (between regular air-cool and submerged cool)

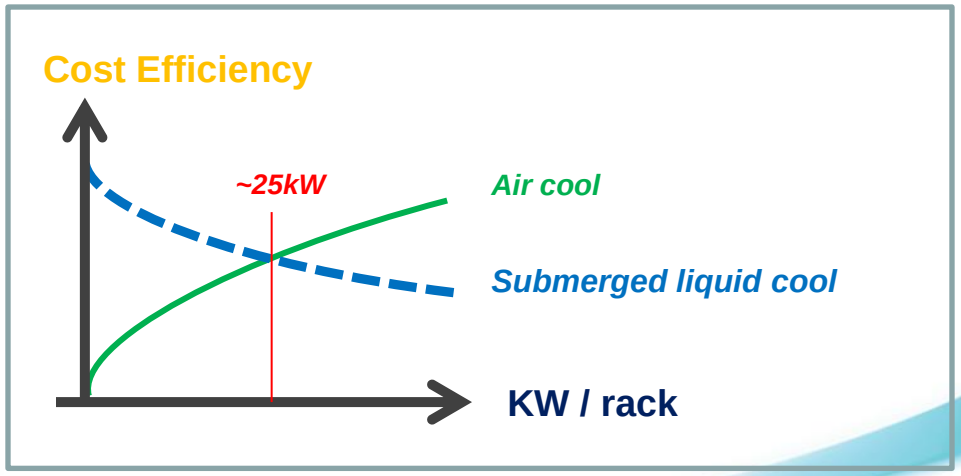

Green500 #1

Green Top-17 all employ Heterogeneous GPU Architectures

Tokyo Institute of Technology

http://www.supermicro.com/products/nfo/Green500.cfm

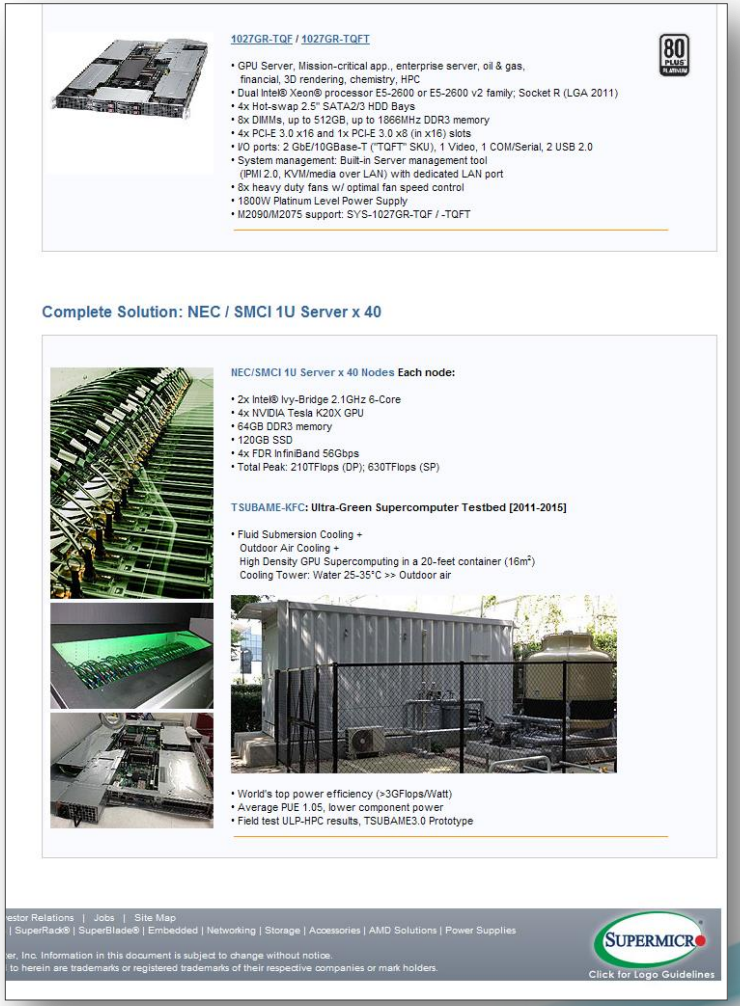

Top500 #311 (~4.5GFLOPS per Watt)

Summary

- A new era of hybrid computing – heterogeneous architecture with GPU / coprocessor acceleration
- There are more to come in the industry roadmap with new technologies, power management features and system architectures
- The trend towards heterogeneous architecture poses many challenges for system builders and software developers in making efficient use of the computing resources
- Configurable cooling & power for energy efficiency and performance are the key to optimized the GPU systems
- Specialized (or application-optimized) design is required for GPU Applications efficiency and scalability
- Supermicro offers the most comprehensive line of solutions supporting the full spectrum of GPU computing applications

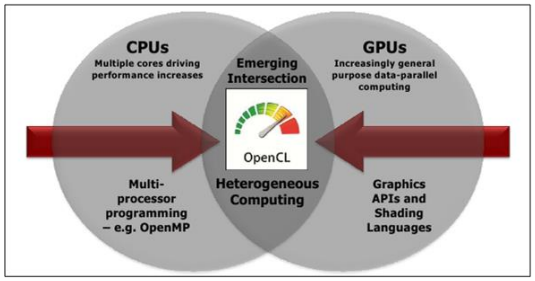

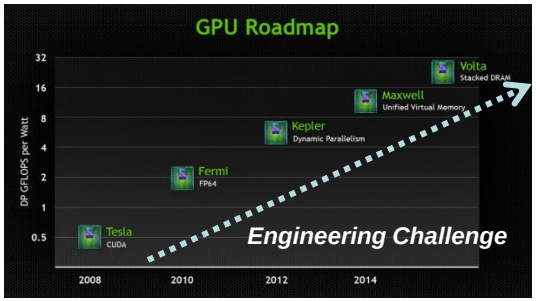

(http://www.supermicro.com/GPU/)

Thank you!

SUPERMICRO® SuperServer® FatTwin™ and SuperBlade® GPU Platforms

HPC, Enterprise Virtualization, Visual Computing and Cloud Gaming at NVIDIA GTC 2014

4U SuperServer®

8x GPU

NVIDIA® Iray VCA Platform

SYS-4027GR-TR/-TRT

4U/Tower

5x GPU NVIDIA Maximus™ Certified SYS-7047GR-TRF

4U Hyper-Speed

1x GPU Low-Latency Optimized SYS-7047AX-TRF

7U SuperBlade®

30x GPU

7U SuperBlade®

3x GPU

SBI-7127RG3

20x GPU

7U SuperBlade®

2x GPU

SBI-7127RG-E

4U 2-Node FatTwin™

12x GPU

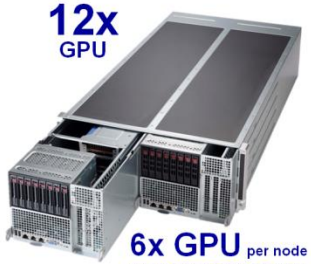

6x GPU per node

SYS-F647G2-FT+/FTPT+/-F73+/F73PT+

NVIDIA

K40 K40X | K20 K20X | K10 | GRID™ K1/K2

SuperServer®

2x GPU

3U SYS-6037R-72RFT+

6x GPU

2U SYS-2027GR-TRFH

3x GPU

1U SYS-1027GR-TRF

3x GPU

1U SYS-1027GR-72R2

2x GPU

1U SYS-1017GR-TF

www.supermicro.com/GPU