SGI® Solutions for NVIDIA® Tesla®

2012 július 2

Agenda/Topics

- NVIDIA® Tesla® Solutions
- Intel Sandy Bridge
- SGI® Solutions
  - Optimized for cabinet-level flexibility
  - Optimized for HPC & Big Data

GPUs are Mainstream

Oil & Gas	Edu/Research	Government	Life Sciences	Finance	Manufacturing

Schlumberger

BR PETROBRAS

Total

Chevron

Paradigm

Chinese Academy of Sciences

Georgia Tech

HARVARD School of Engineering and Applied Sciences

Oak Ridge National Laboratory

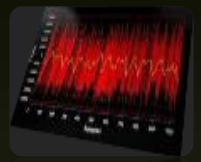

Air Force Research Laboratory

NASA

Naval Research Laboratory

BAE SYSTEMS

Boston Scientific

Max Planck Institute

Mass General Hospital

Beckman Coulter

life technologies

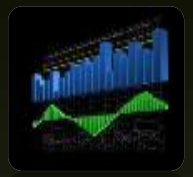

Bloomberg

BNP PARIBAS

STANDARD LIFE

J.P.Morgan

NumeriX

Agilent

ANSYS

Autodesk

SIMULIA

ACUSIM SOFTWARE

EMEA GPU users

- Poznan
  167szerver, 334db M2050
- Szeged, NIIFI
  2-4 szerver, 4-8 M2075 (?)
- Wigner SZFKI
  1 szerver, 2x M2075
- ...

NVIDIA® Current Gen Tesla® GPUs

Tesla Data Center & Workstation GPU Solutions

Tesla M-series GPUs
M2090 | M2075

Tesla C-series GPUs
C2075

		Servers & Blades		Workstations
		M2090	M2075	C2075
Cores		512	448	448
Memory		6 GB	6 GB	6 GB
Memory bandwidth (ECC off)		177.6 GB/s	150 GB/s	148.8 GB/s
Peak Perf Gflops	Single Precision	1331	1030	1030
	Double Precision	665	515	515

Max: 1mFt 700eFt 800eFt

NVIDIA® Current Gen **Quadros**

	Ultra High End		High End			Mid Range		Entry	
BOARD FEATURES	**Quadro 6000**	**Quadro FX 5800**	**Quadro 5000**	**Quadro 4000**	**Quadro 4000 for Mac**	**Quadro 2000**	**Quadro 2000D**	**Quadro 600**	**Quadro 400**
Memory Size	6GB[4] GDDR5	4GB GDDR3	2.5GB GDDR5	2GB GDDR5	2GB GDDR5	1GB GDDR5	1GB GDDR5	1GB DDR3	512MB DDR3
Memory Interface	384-bit	512-bit	320-bit	256-bit	256-bit	128-bit	128-bit	128-bit	64-bit
Memory Bandwidth	144 GB/s	102GB/s	120 GB/s	89.6 GB/s	89.6 GB/s	41.6 GB/s	41.6 GB/s	25.6 GB/s	12.3 GB/s
CUDA™ Parallel Processor Cores	448	240	352	256	256	192	192	96	48
	1,5mFt+								**100eFt-**
Max Power Consumption	204W	189W	152W	142W	142W	62W	62W	40W	32W
Number of slots	2	2	2	1	1	1	1	1	1
Display Connectors	DVI-I DP DP 3pin Stereo	DVI-I DVI-I DP 3pin Stereo	DVI-I DP DP 3pin Stereo	DVI-I DP DP 3pin Stereo[2]	DVI-I DP 3pin Stereo[2]	DVI-I DP DP	DVI-I DVI-I	DVI-I DP	DVI-I DP
Dual Link DVI	1	2	1	1	1	1	2	1	1
DisplayPort	2	1	2	2	1	2		1	1
ECC (Error Correcting Code)	✓		✓						
Fast Double Precision	✓		✓	✓	✓				
OpenGL	4.1	3.3	4.1	4.1	4.1[3]	4.1	4.1	4.1	4.1
Shader Model	5.0	4.0	5.0	5.0	5.0	5.0	5.0	5.0	4.0
DirectX	11	10.1	11	11	11	11.0	11.0	11.0	10.1
3D Vision Pro	✓	✓	✓	✓		✓	✓	✓	✓
Quadro® Mosaic Technology	✓	✓	✓	✓		✓	✓	✓	✓
NVIDIA® SLI® Multi OS	✓	✓	✓	✓		✓	✓		
NVIDIA® SLI Frame Rendering Support	✓	✓	✓						
NVIDIA® nView® Display Management Technology	✓	✓	✓	✓		✓	✓	✓	✓
Quadro® SDI option card	✓	✓	✓	✓					
Quadro® G-Sync option card	✓	✓	✓						
3D PRIMITIVE PERF									
Triangles per Second	1.3 Billion	300 Million	950 Million	890 Million	890 Million	410 Million	410 Million	210 Million	120 Million

sgi

NVIDIA® Next Gen Tesla® GPUs

Kepler GPU

NVIDIA

Fastest, Most Efficient HPC Architecture Ever

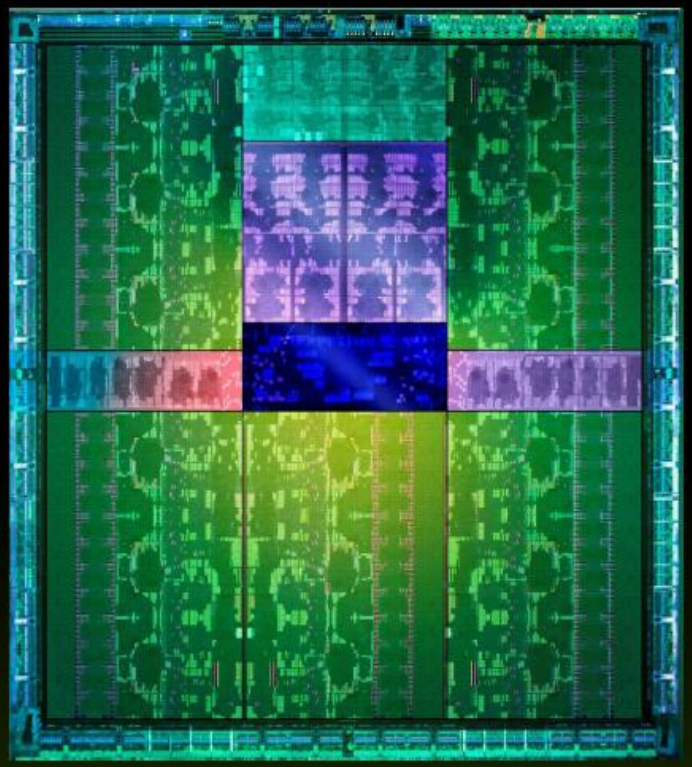

Performance : SMX

Efficiency : Hyper-Q

Programmability : Dynamic Parallelism

sgi

NVIDIA® Next Gen Tesla® GPUs

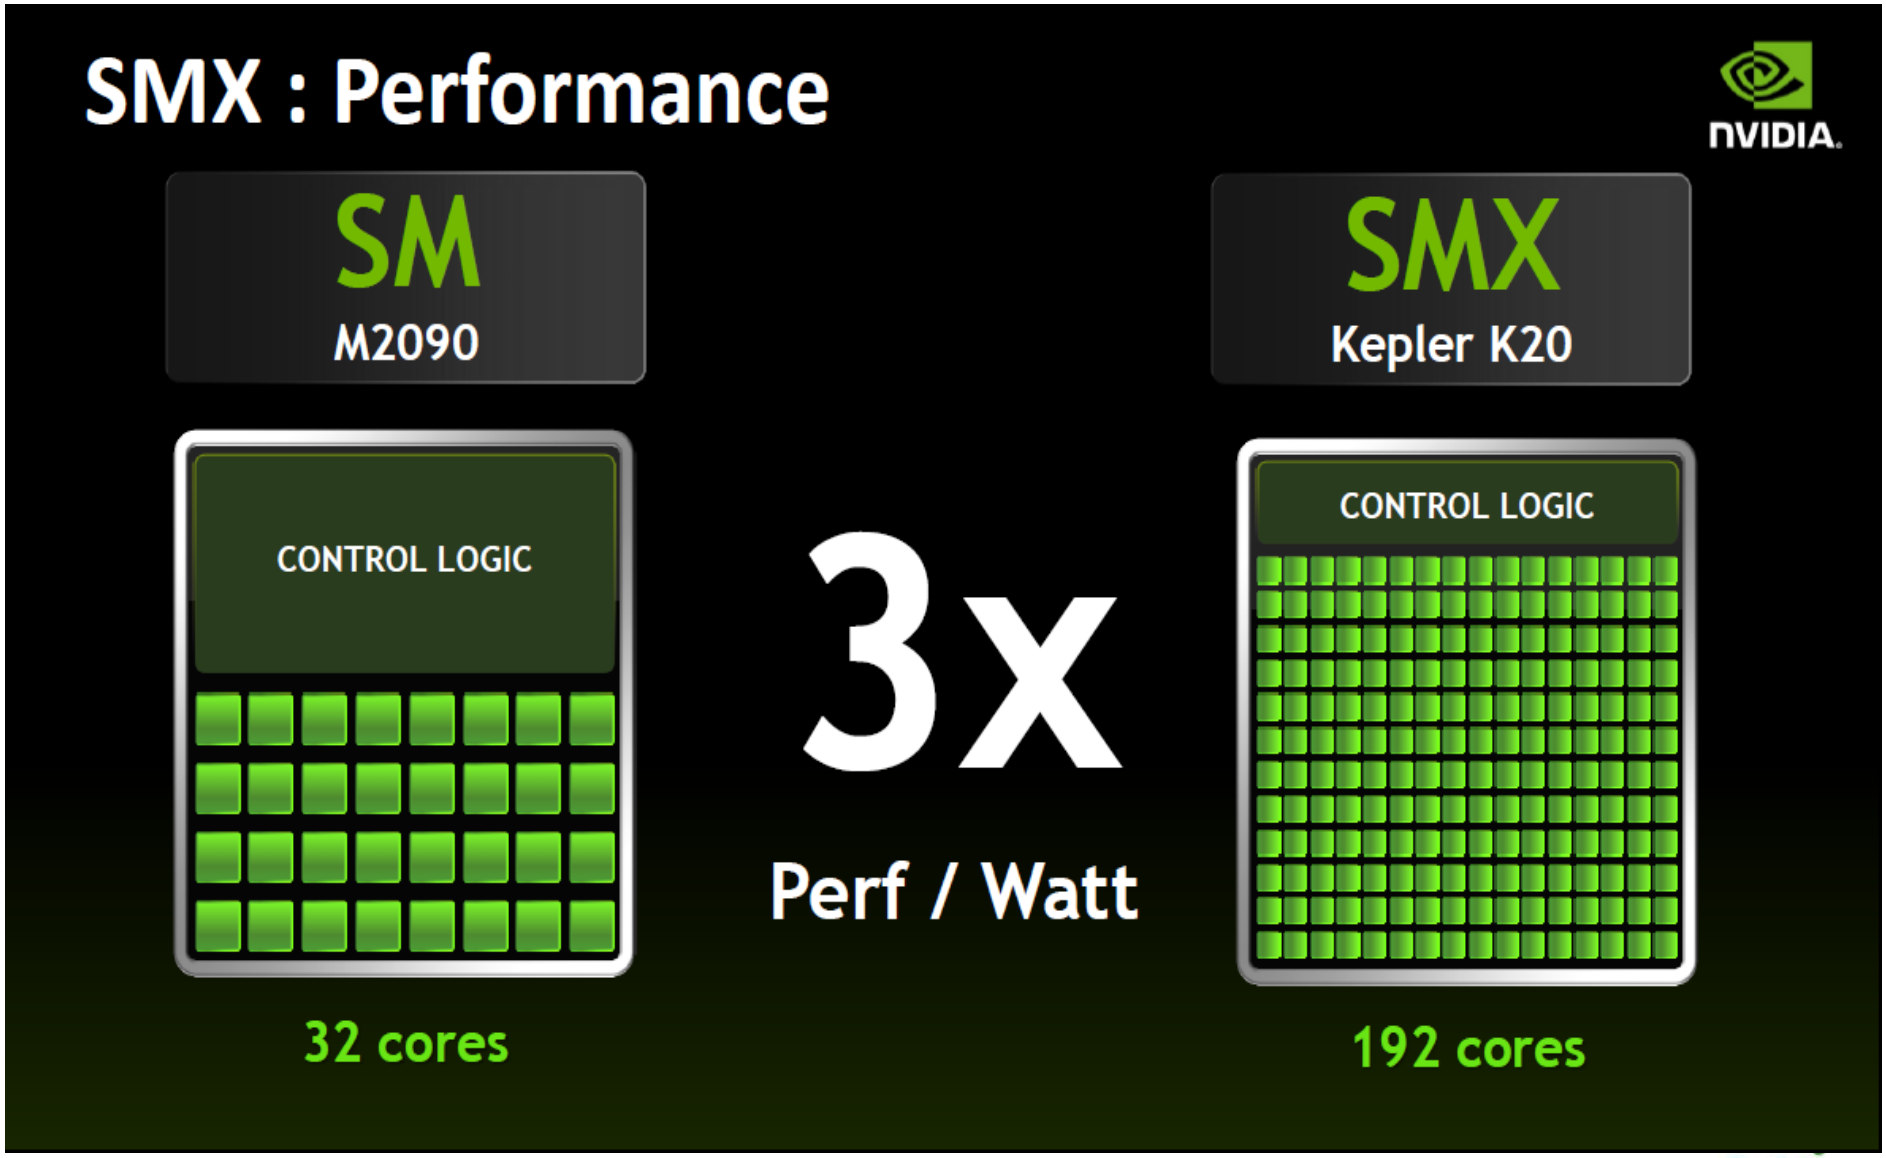

sgi

NVIDIA® Next Gen Tesla® GPUs

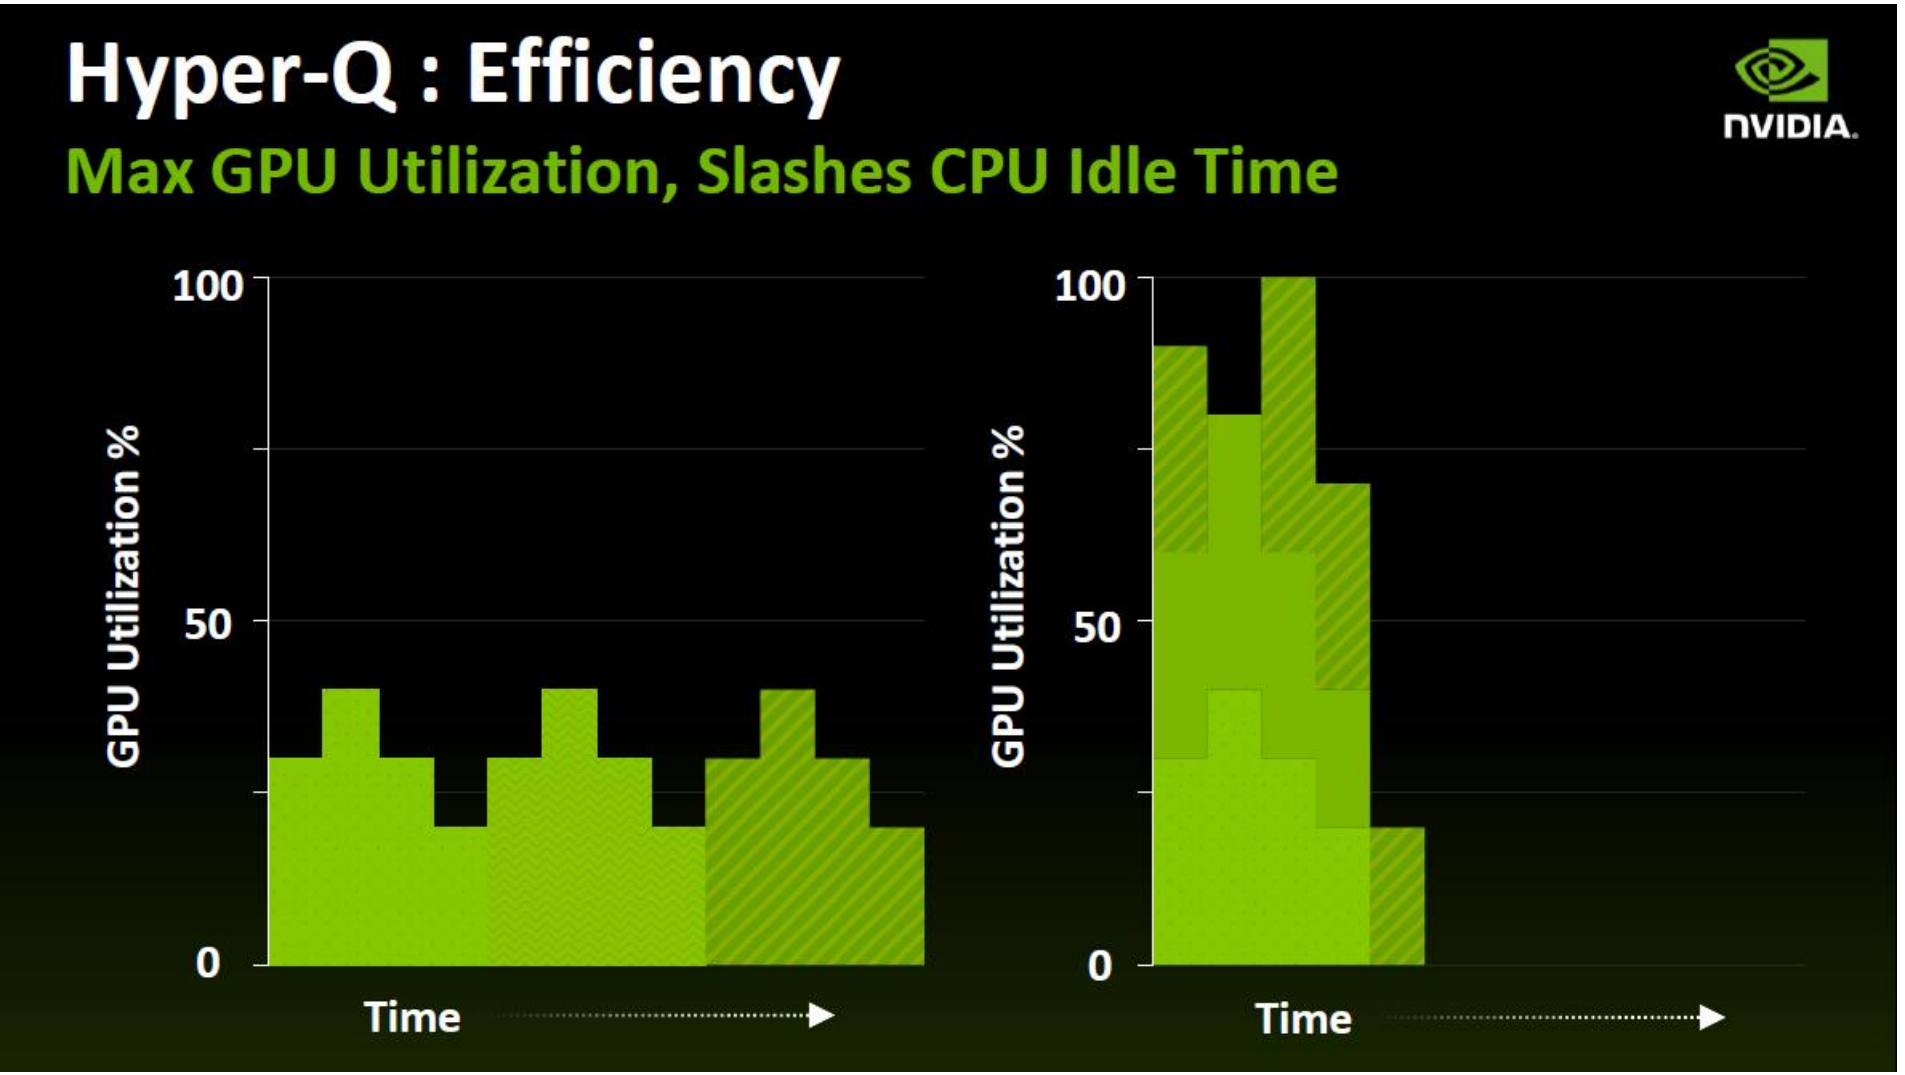

sgi

NVIDIA® Next Gen Tesla® GPUs

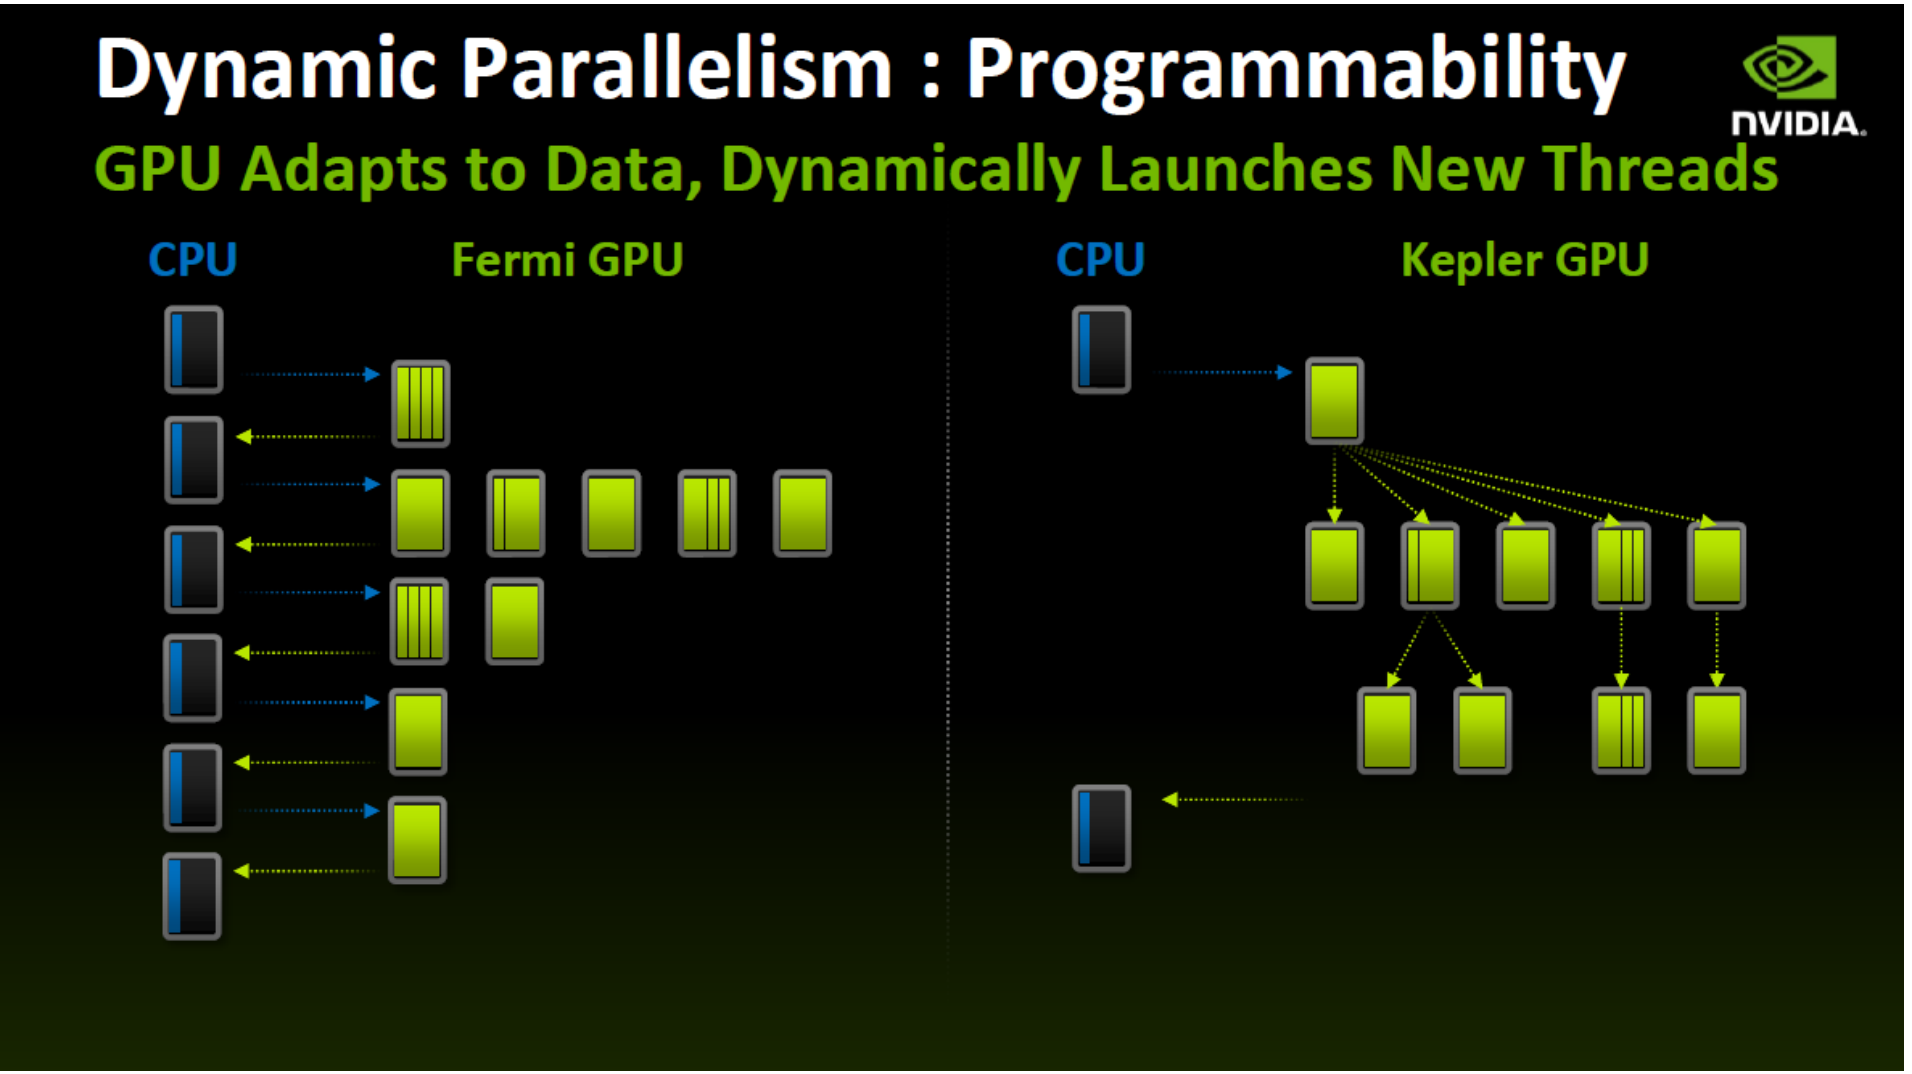

sgi

NVIDIA® Next Gen Tesla® GPUs

Kepler Product Family

NVIDIA

Tesla K10

3x Single Precision

1.8x Memory Bandwidth

Image, Signal, Seismic, Life Sci (MD)

Available Now

<1mFt

Tesla K20

3x Double Precision

Hyper-Q, Dynamic Parallelism

CFD, FEA, Finance, Physics

Available Q4 2012

sgi

SGI® Servers for Speed and Scale

HPC	Big Data	Cloud
Commercial	Hadoop	Public
Scientific	In-memory	Private
Mode		

NVIDIA® Tesla Solutions Available!

SGI® ICE

Rackable™

SGI® UV

Rackable™

CloudRack™ C2

Romley Platform

Next Stage on Intel Tick/Tock

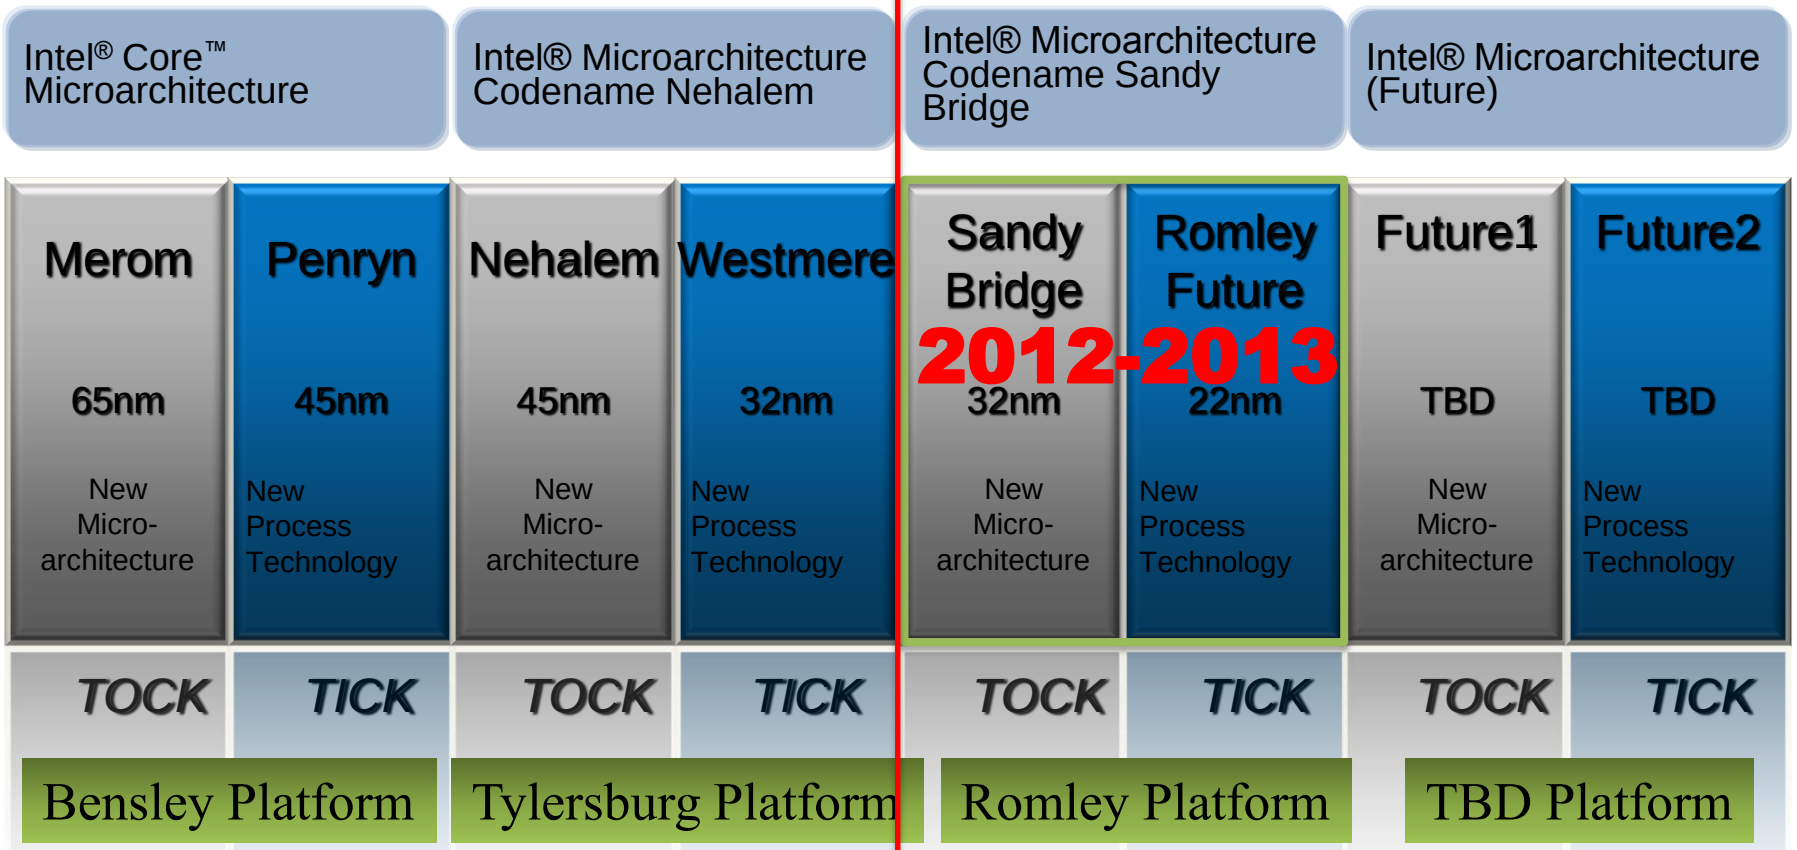

Gflps /socket

43 43 43 64 166.4

1X 1.5X ~3X!

*SGI® estimate based upon
Assumed 2.66 GHz sku, 8 cores/socket
8 flops/core

Romley: Integrated I/O

- First integrated PCIe into the server processor
- Delivers breakthrough performance:
  - I/O latency reduced > 30%
  - Up to 40 PCIe lanes per socket
- **2X the I/O bandwidth**!

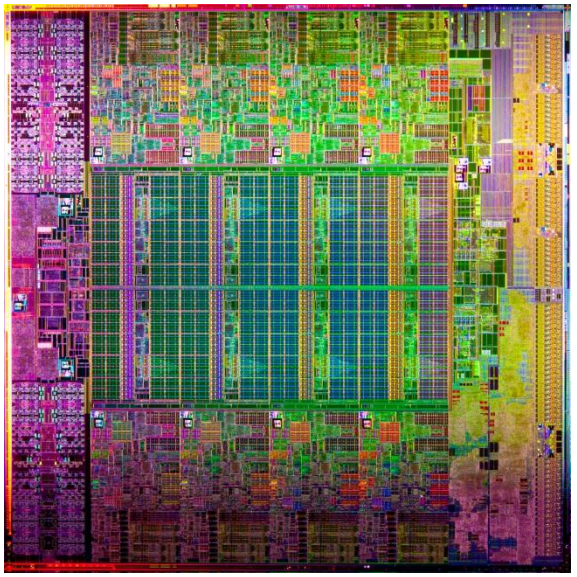

sgi

Key Romley Components

- *SandyBridge* Processor (later IvyBridge)
- Data Center Manager Interface (DC[illegible] and Intel Power Node Manager
- *Patsburg* Chipset
- PCIe gen 3.0 controller integrated into the processor

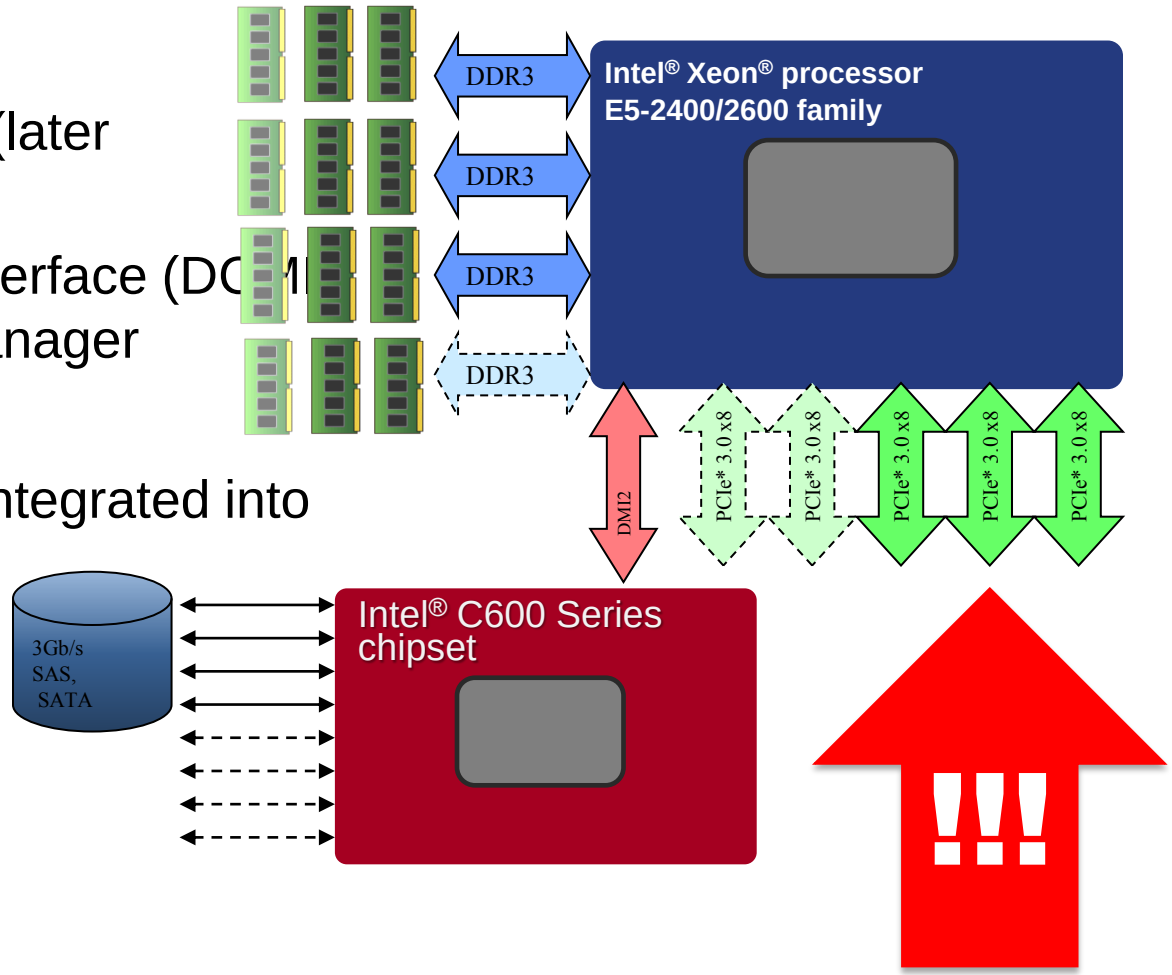

Intel® Xeon® Processor E5-2600 Product Family Integrated PCIe* Bandwidth

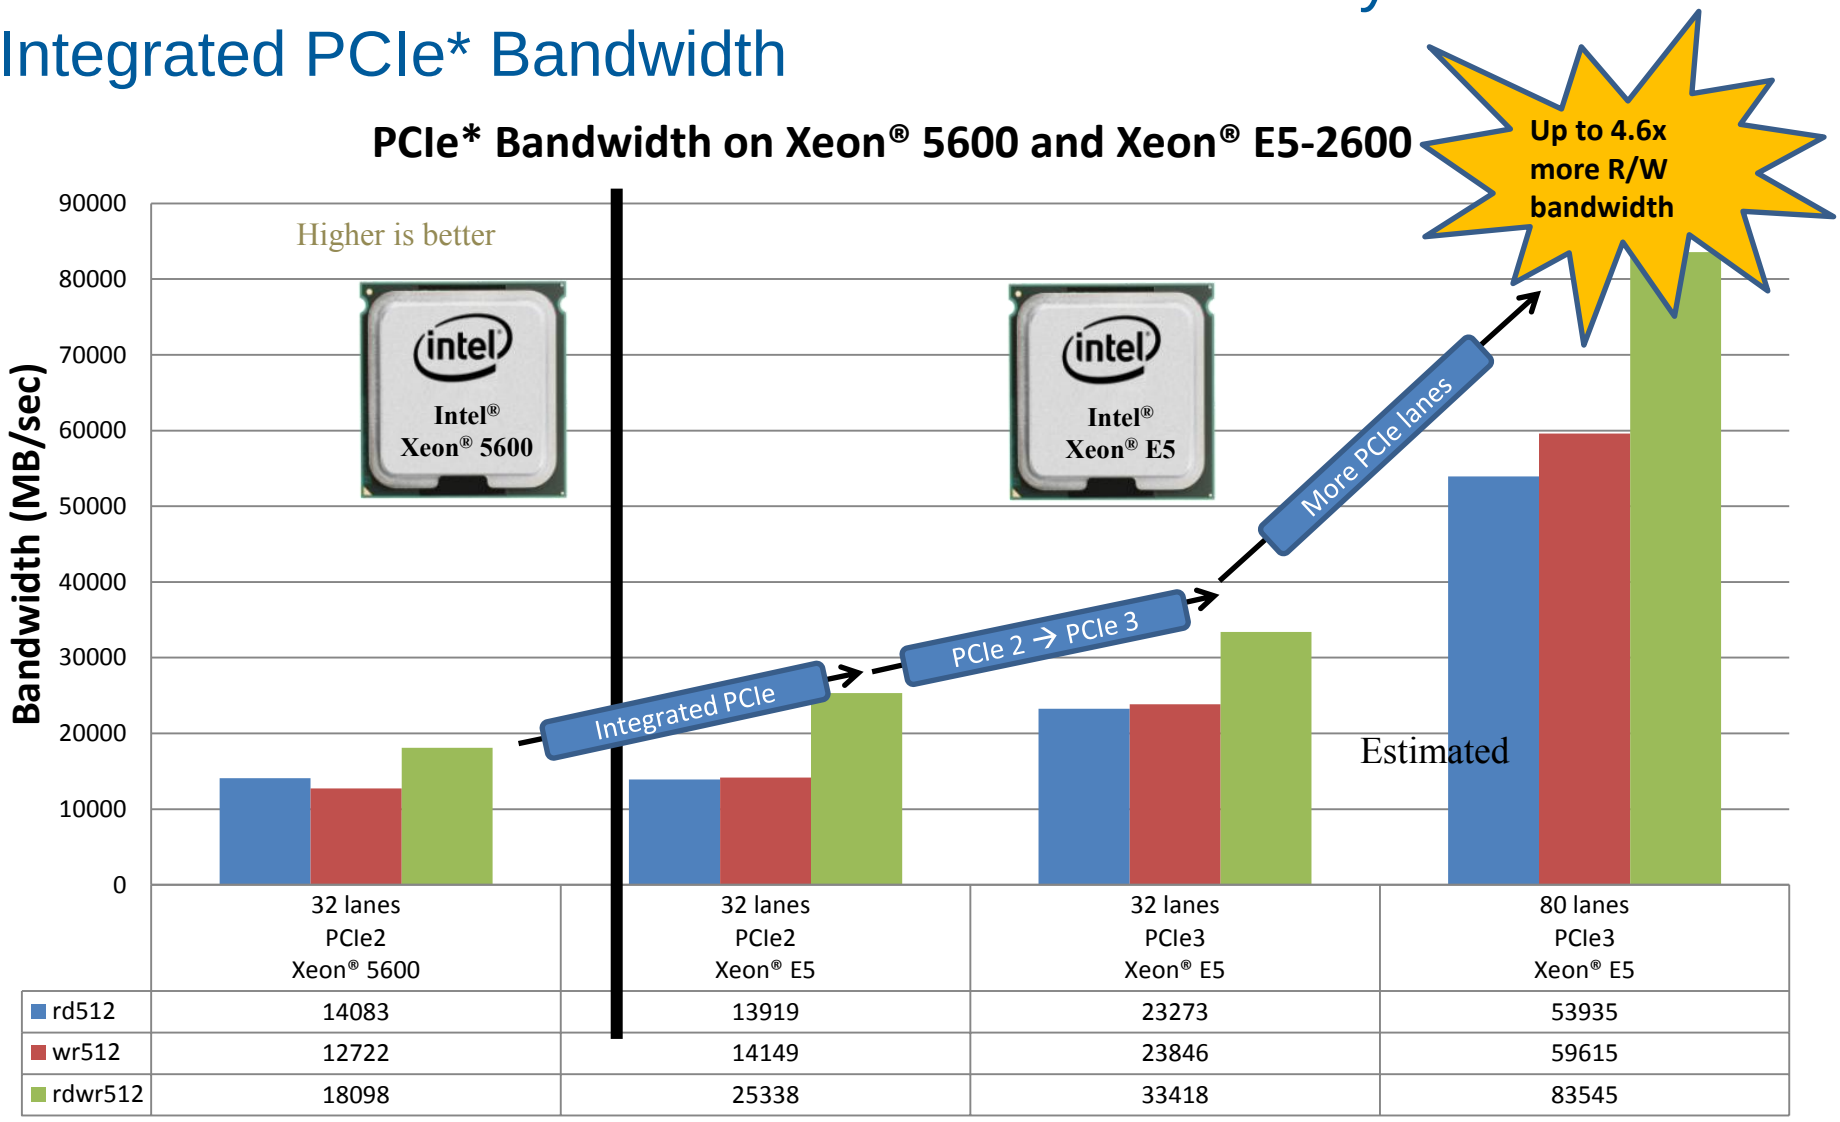

	32 lanes PCIe2 Xeon® 5600	32 lanes PCIe2 Xeon® E5	32 lanes PCIe3 Xeon® E5	80 lanes PCIe3 Xeon® E5
rd512	14083	13919	23273	53935
wr512	12722	14149	23846	59615
rdwr512	18098	25338	33418	83545

Up to 4.6x more R/W Bandwidth with PCIe* 3; Up to 80 PCIe lanes

2S Sandy Bridge-EP HPC Application Performance

HPC Suite Workloads, Single Node

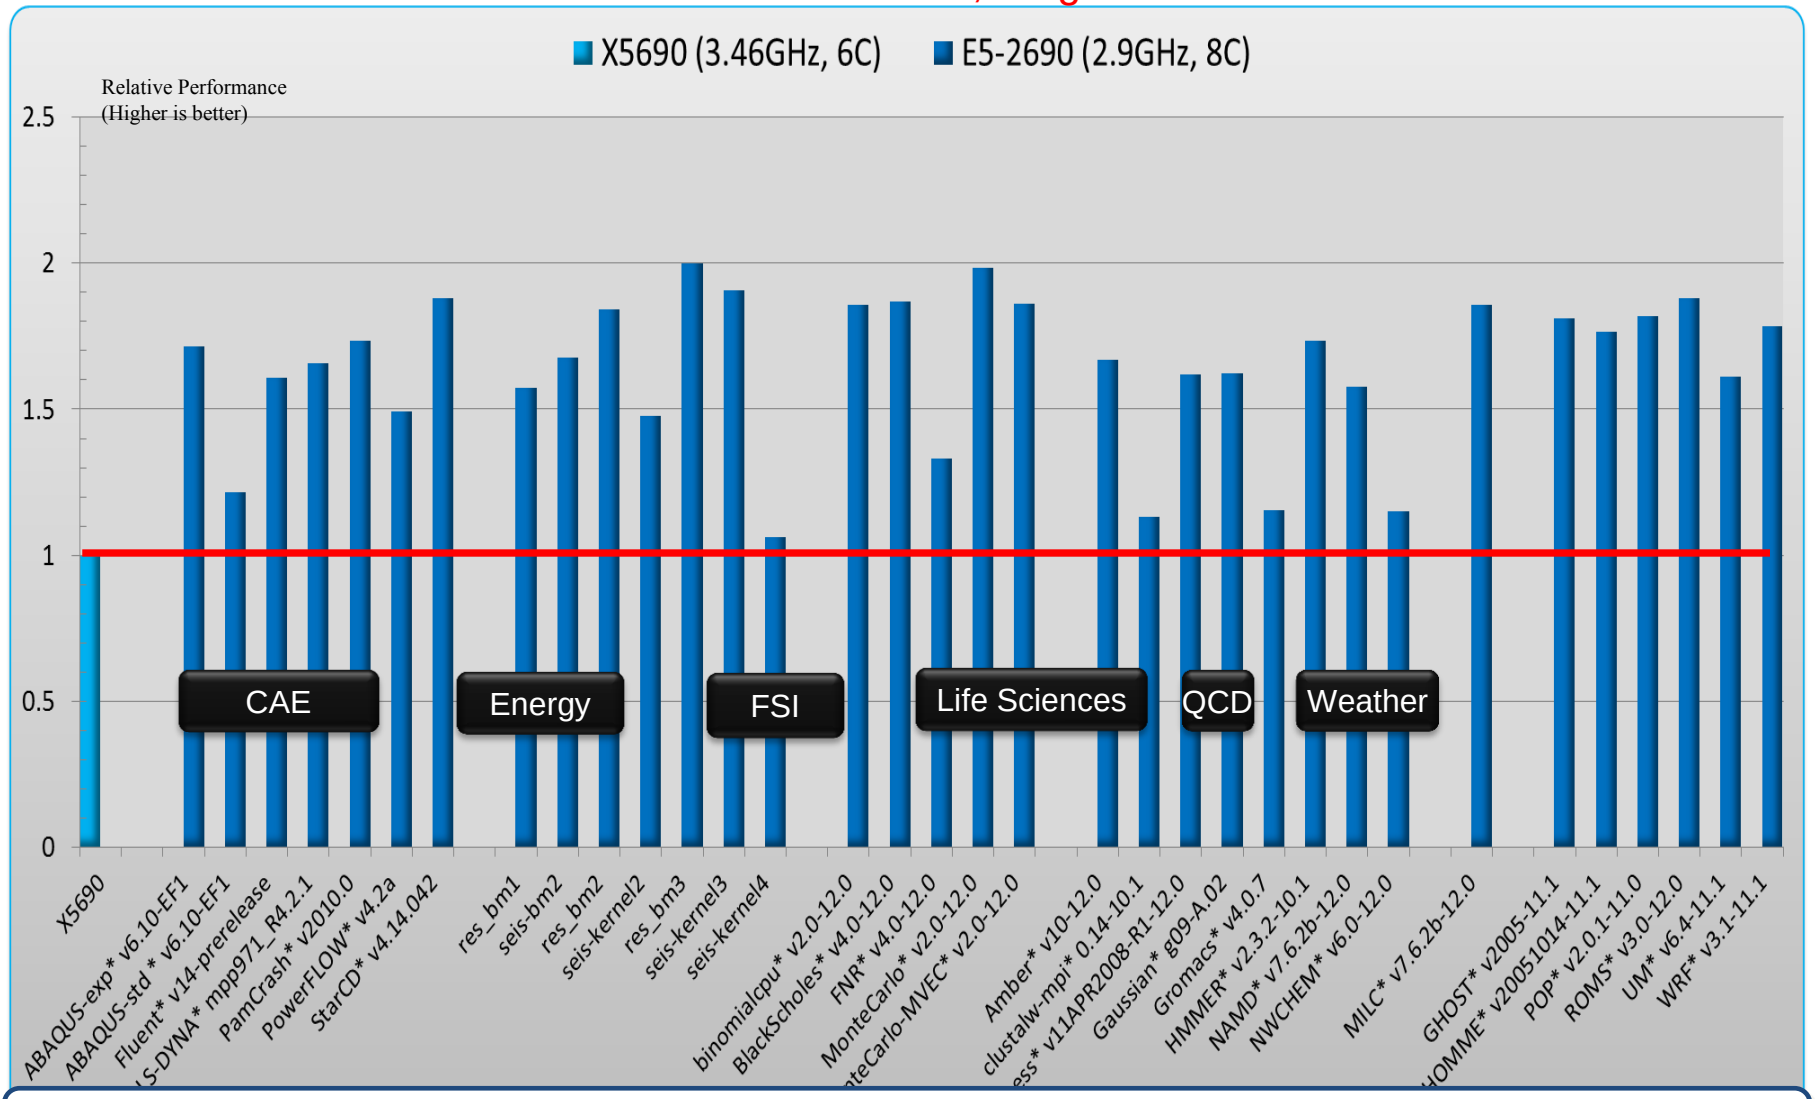

Xeon® E5-2690 delivering great generational performance increase

SGI® Products Overview

- It starts with software!
- Sandy Bridge evereywhere
- Optimized for cabinet-level flexibility
  - Rackable™
- SGI Management Center
- SGI GPU Clusters and Starter Kits
- Optimized for HPC & Big Data
  - SGI ICE™
  - SGI UV™ 2

Software Development Ecosystem for NVIDIA® Tesla®

	Open Source	Commercial
Compiler	OpenACC, OpenCL, OpenGL, NOAA F2C, FLAGON - Fortran-9X Library for GPU Numerics, PyCUDA Python for CUDA, jCuda – Java for CUDA, JOCL – Java for OpenCL	CUDA C/C++, Portland Group PGI Accelerator & CUDA Fortran, CAPS HMPP
Debugger		NVIDIA CUDA Toolkit (CUDA-GDB, CUDA-MEMCHECK), Rogue Wave TotalView, Allinea DDT
Libraries	Thrust – C++ Std Template Library	NVIDIA CUDA Toolkit (CUDA Math Library, cuFFT, cuBLAS, cuSPARSE, cuRAND, NPP, CUSP), Rogue Wave IMSL Fortran Numerical Library, Mathworks Matlab, Wolfram Mathematica, AccelerEyes ArrayFire & LibJacket, EmPhotonics CULA|Tools (LAPACK), Labview
Profiling & Analysis Tools	PAPI CUDA Component	NVIDIA CUDA Toolkit (Nvidia Visual Profiler), Tau, Vampir Trace
Workload Scheduler		Altair PBS Professional, Adaptive Computing Moab

sgi

SGI® Management Center

Premium Edition offers GPU Monitoring and Management!

- Single system management console with remote server monitoring and control
- Ease of use, full system management including GUI and CLI
- Policy driven, fine grained power load management for green computing
- Advanced fault, event, and alert management for improved reliability
- Advanced capabilities including GPU management, BIOS management, and high availability

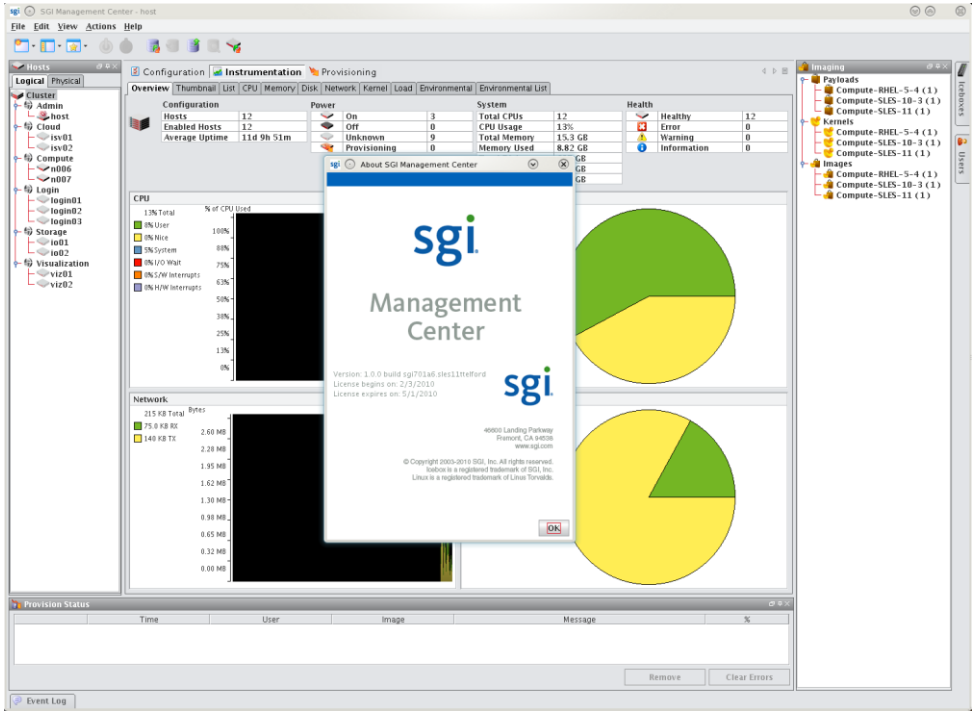

SMC Premium Edition– GPU Monitoring

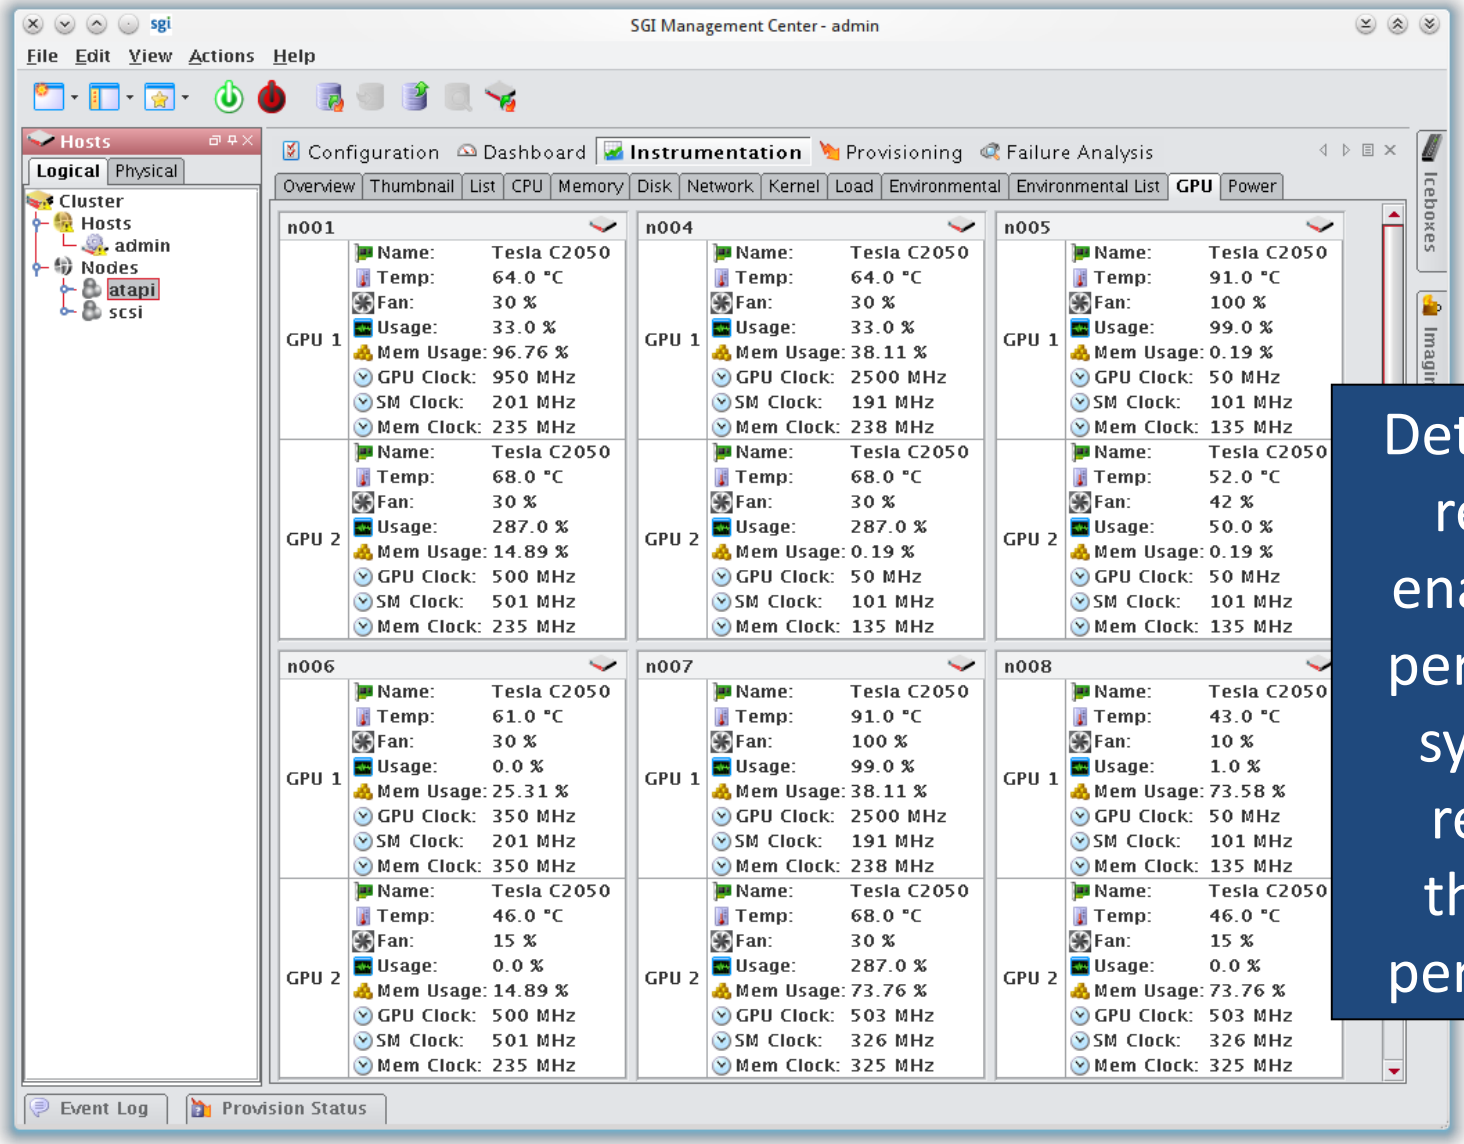

Detailed GPU reporting enables high performance systems to remain at their peak performance

SGI Rackable™ C1104G

NVIDIA GPU accelerator support:
•Tesla M2075, M2090**, K10 (Kepler I)
NVIDIA graphics GPU support:
•Quadroplex 7000

3x

**130W standard Max. 135W with ambient temp. restrictions

Model	C1104G-RP5
Chassis Profile	1U standard-depth
Servers/System	One dual-socket
Chipset	Intel® C600
Max. Processors	Two Intel® **Xeon® E5-2600**
Max. CPU TDP	115W
Memory Slots	8 DIMM slots
Memory Type	1600/1333/1066/800 MHz DDR3 ECC Reg
Max. Hard Disk Drives	4 x 2.5" drives
Expansion Slot	Three PCI-E 3.0 x16 (two internal and one external) double-width slots One external PCI-E 3.0 x8 low-profile slot
Networking (Onboard)	Dual-Port GigE controller (Intel® I350)
IPMI Remote Management	Integrated IPMI 2.0
Power Supply	1800W Redundant* Platinum Level

*Redundant per configuration

SGI Rackable™ C2110G

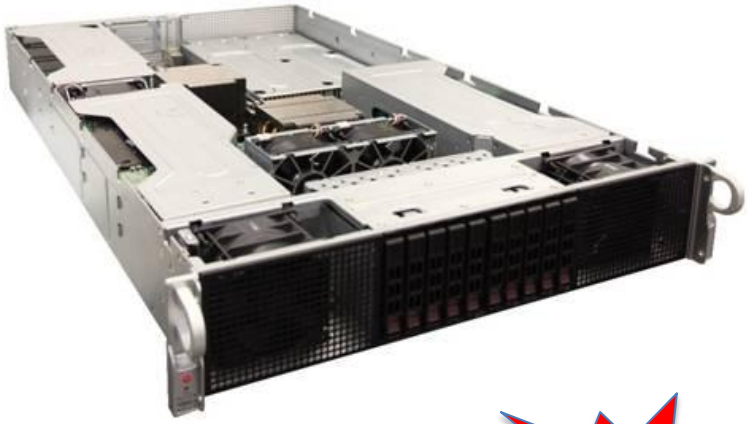

NVIDIA GPU accelerator support:
•Tesla M2075, M2090, K10 (Kepler I)
NVIDIA graphics GPU support:
•Quadroplex 7000

6x

Model	C2110G-RP5
Chassis Profile	2U standard-depth
Servers/System	One dual-socket
Chipset	Intel® C600
Max. Processors	Two Intel® **Xeon® E5-2600**
Max. CPU TDP	130W
Memory Slots	8 DIMM slots
Memory Type	1600/1333/1066/800 MHz DDR3 ECC Reg
Max. Hard Disk Drives	10 x 2.5” drives
Expansion Slot	Four internal PCI-E 3.0 x16 double-width (optional riser card for 2 additional external PCI-E 3.0 x16 double width slots*), One external PCI-E 3.0 x8 low profile and one PCI-E 2.0 x4 full height
Networking (Onboard)	Dual-Port GigE controller (Intel® I350)
IPMI Remote Management	Integrated IPMI 2.0
Power Supply	1800W Redundant** Platinum Level

*Not available at launch
**Redundant per configuration

SGI® Integrated Clusters for Tesla®

- A complete, managed GPU solution of software and hardware, all you need for a powerful deployment
- Hardware stack
  - Rackable™ C2108 head node
  - Either/or Rackable C1104G or C2110G compute nodes
  - Infiniband NICs and Switch
  - Ethernet switch for management
  - Add SGI InfiniteStorage solutions, either direct-attached to the head node or SGI NAS
- Software stack
  - SGI Management Center software
  - Altair PBSpro Load Management software
- Complete Factory Test & Integration

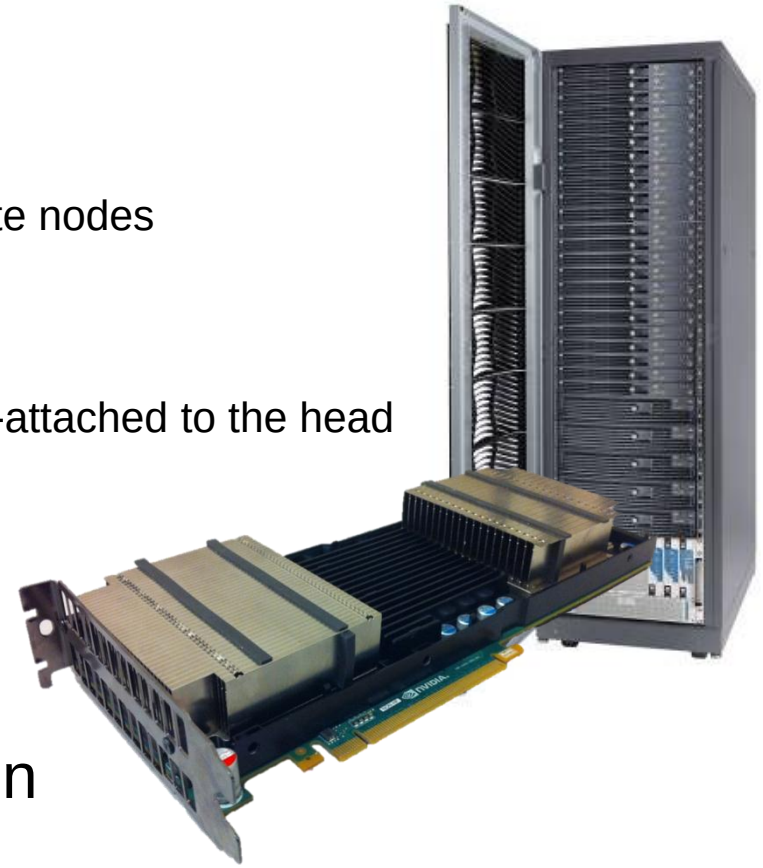

Optimized for HPC & Big Data Solutions

- SGI® ICE™ X
- SGI UV™ 20/2000

Introducing **SGI ICE X**…
Fifth Generation ICE (IB cluster)

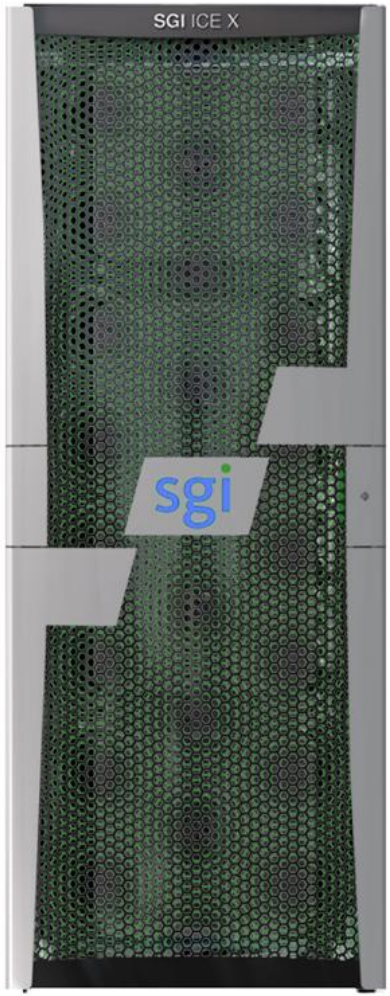

- **The world's fastest supercomputer just got faster!**
- **The world-renowned SGI quality and performance you love**
- **Flexible to fit your workload**

SGI® ICE X twin blade

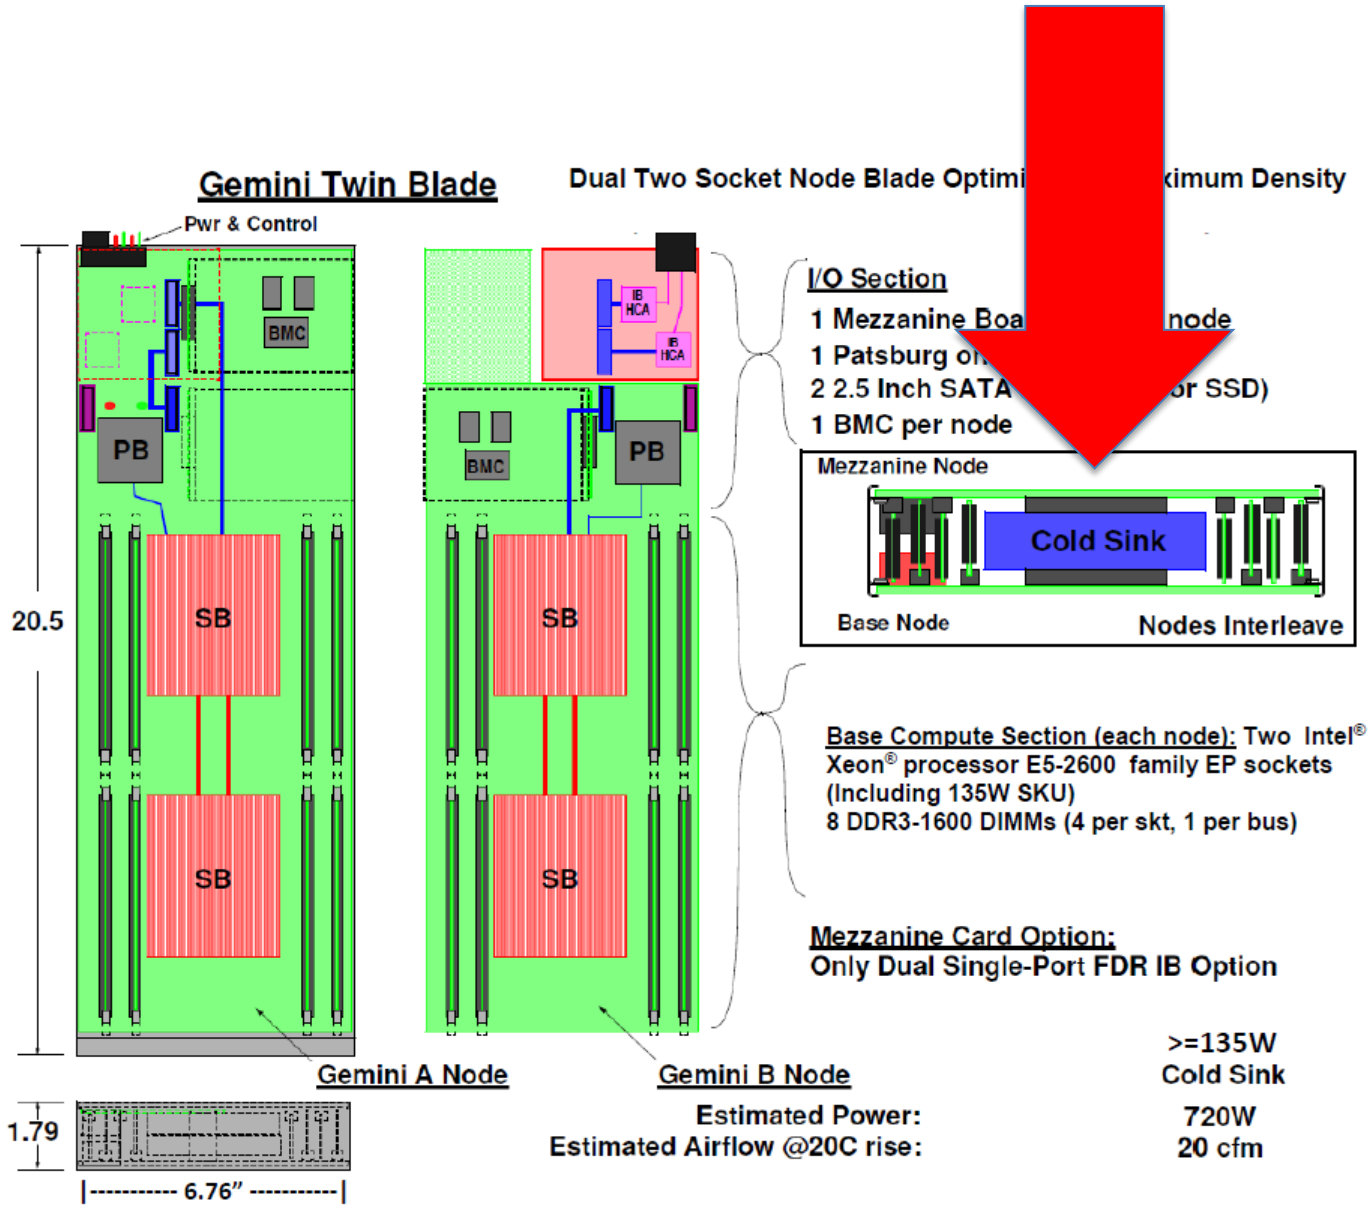

Main Features:

- Supports single plane FDR InfiniBand
- Supports four Intel® Xeon® processor E5-2600 family CPUs
  - Two dual socket nodes
- Supports four DDR3 DIMMs per socket @ 1600 MT/s
- Houses up to two 2.5" SATA drives for local swap/scratch usage
  - One per node
- Utilizes cold sinks (liquid)

SGI ICE blade with GPU

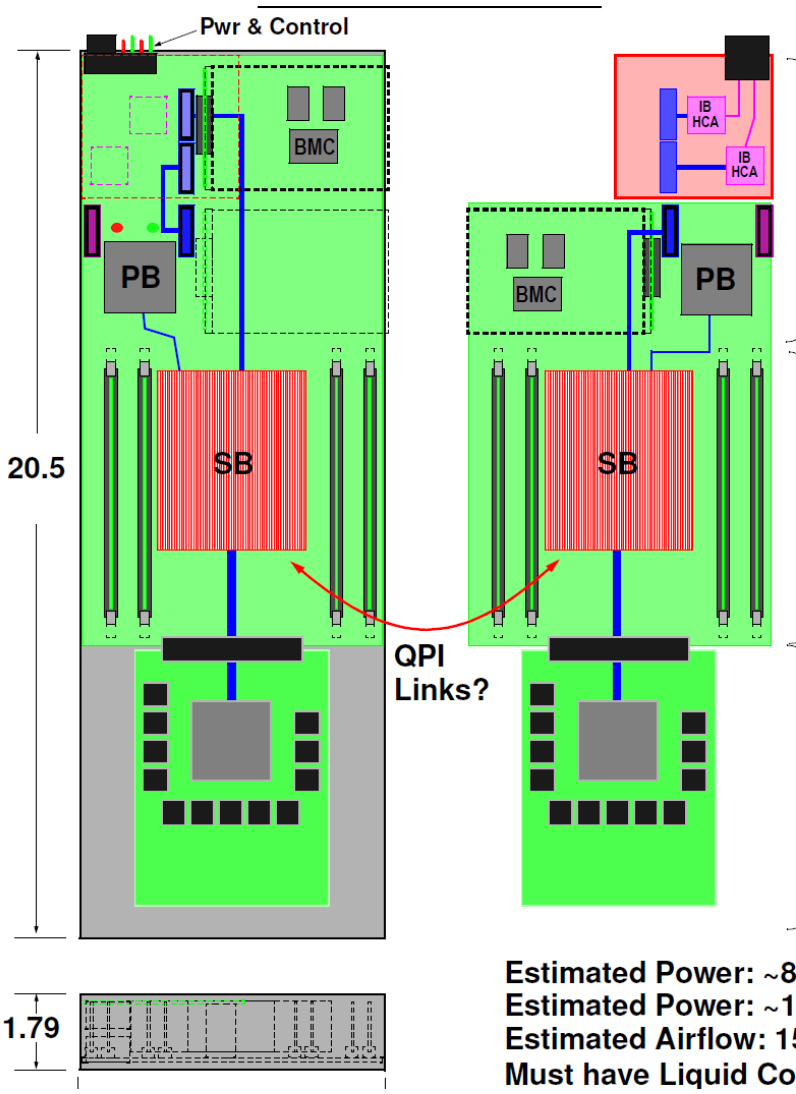

With 2 NVidia SXM GPUs

I/O Section
- 1 Mezzanine Board on base node
- 1 Patsburg on each node
- 2 2.5 Inch SATA Drive (HDD or SSD)
- 1 BMC per node

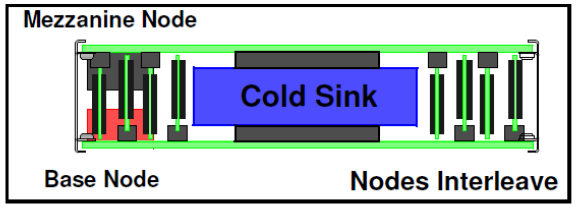

Base Compute Section (each node)
- 1 SandyBridge EP Sockets (130W SKU)
  - w/ Liquid Cooled Cold Sink
- 4 DDR3-1600 DIMMs (4 per skt, 1 per bus)

GPU
up to 300W each,

Mezzanine Card Option

Dual Single Port FDR IB

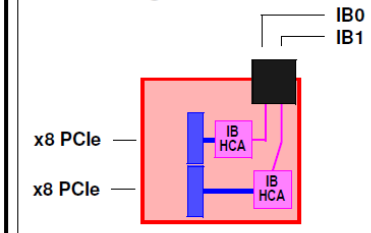

Estimated Power: ~870 Watts (225W GPU)
Estimated Power: ~1020 Watts (300W GPU)
Estimated Airflow: 15 cfm (@20C rise)
Must have Liquid Cooled Sinks for CPUs & GPUs

ICE Direct Water cooling (opt.)

For GPUs, **MIC**, twin-blades
SGI COLDSINK Technology

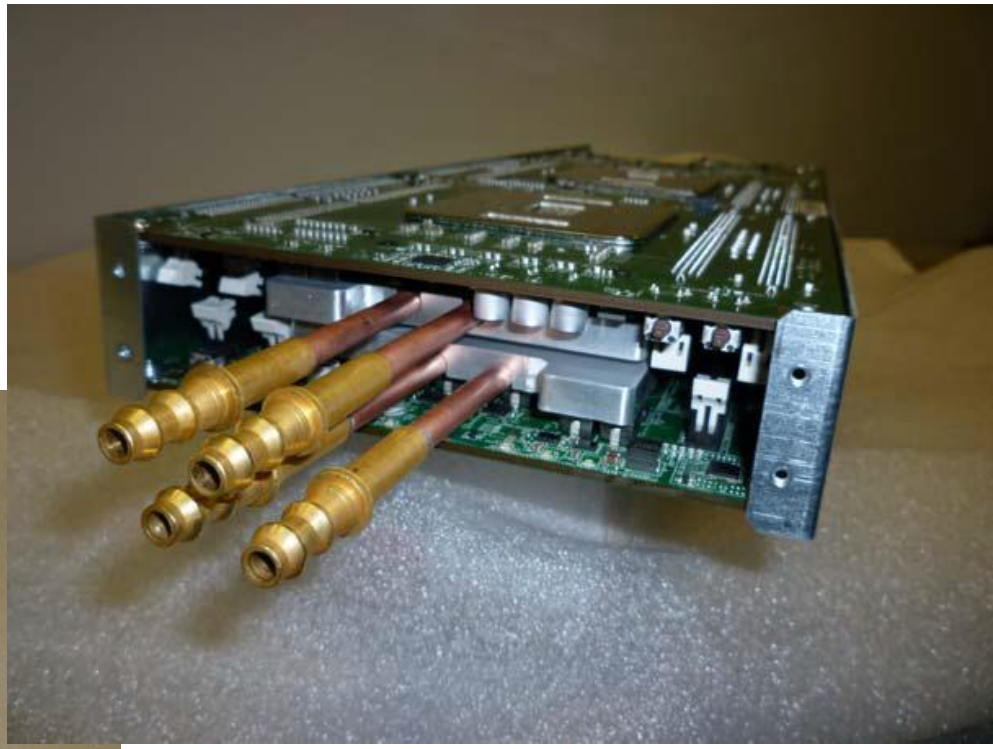

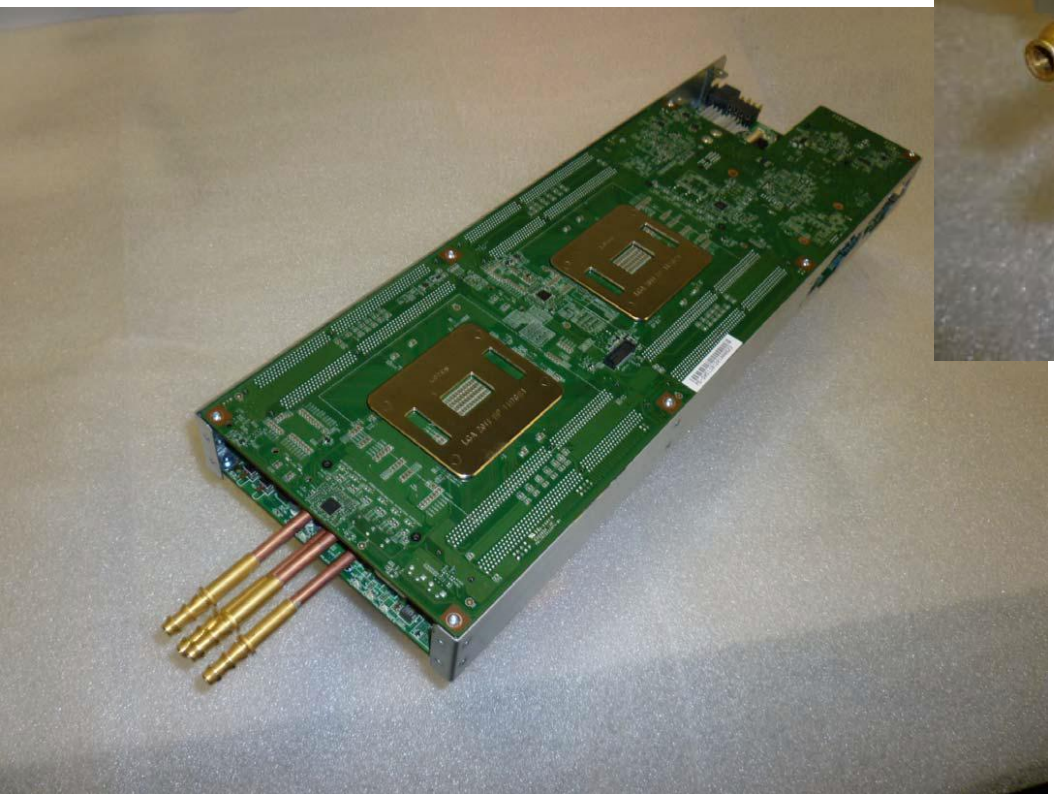

http://www.intel.com/content/www/us/en/architecture-and-technology/many-integrated-core/intel-many-integrated-core-architecture.html

sgi

SGI UV 2: The Big Brain Computer

World-leading Capability for Data Intensive Work, now with Intel® Xeon® Processor E5-4600 Product Family.

- ***Focus on Solving Your Problems, Not IT Problems***
- ***No-Limit Computing, Built on Industry Standards***
- ***World's Largest In-Memory System for Data-Intensive Applications***

- *Large scalable Shared Memory System*
- *Up to **4096 Cores and 64TB***
- *Károly Róbert Főiskola, **Gyöngyös***

SGI UV 2

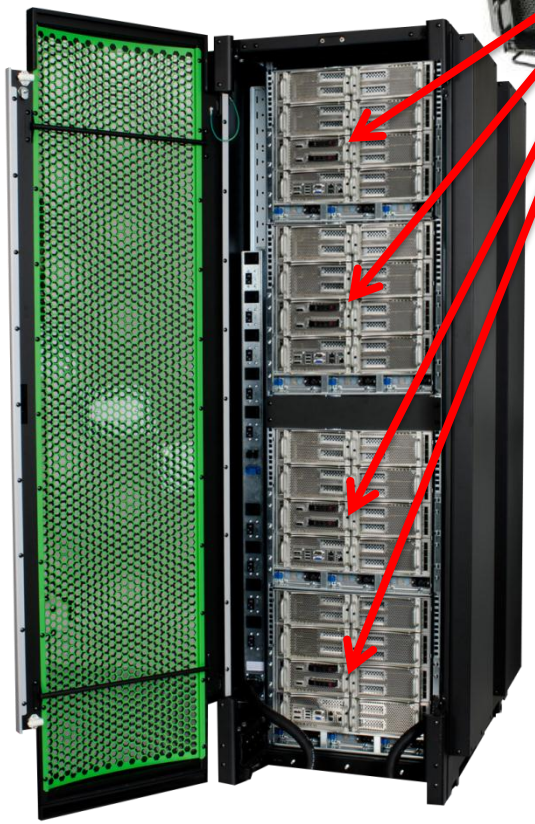

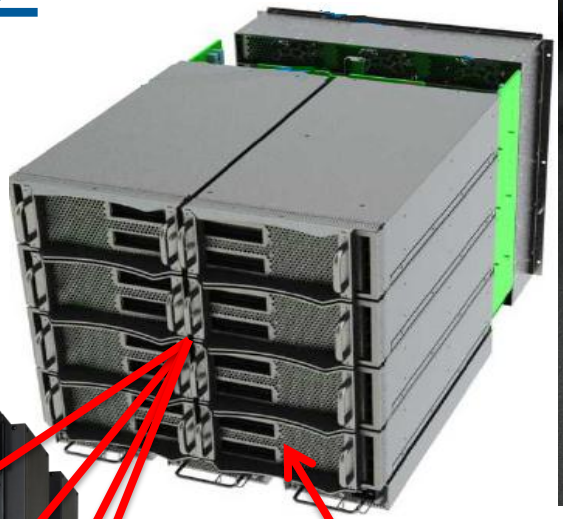

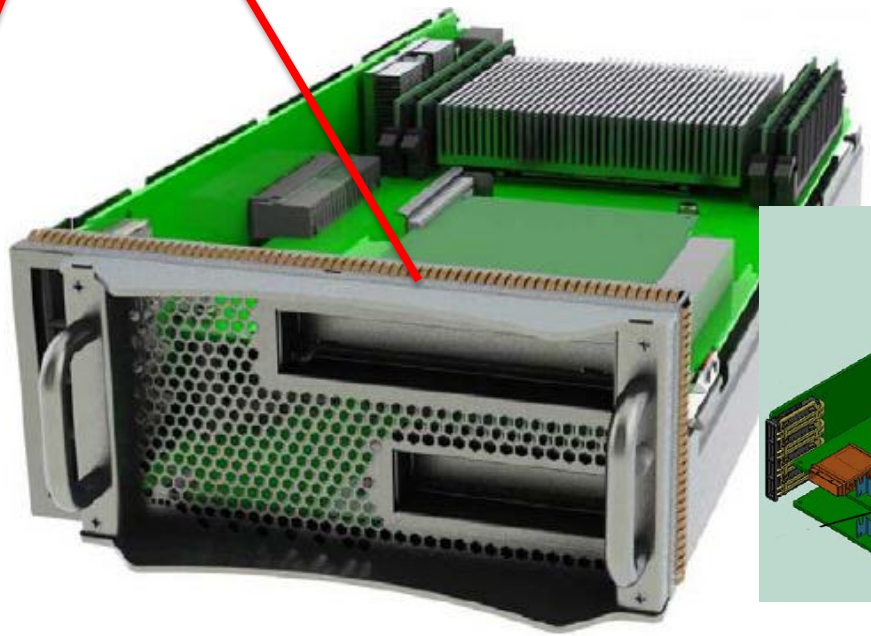

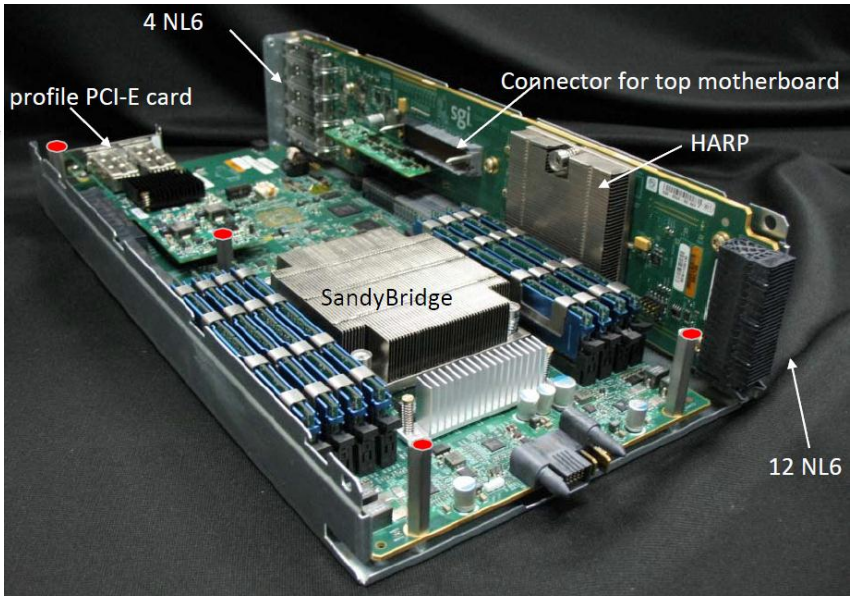

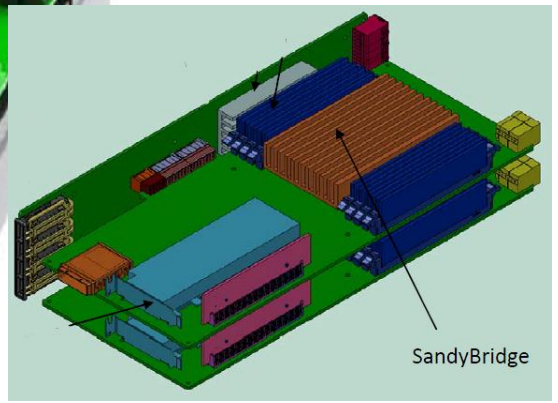

SGI® UV™ 20 GPU Solutions

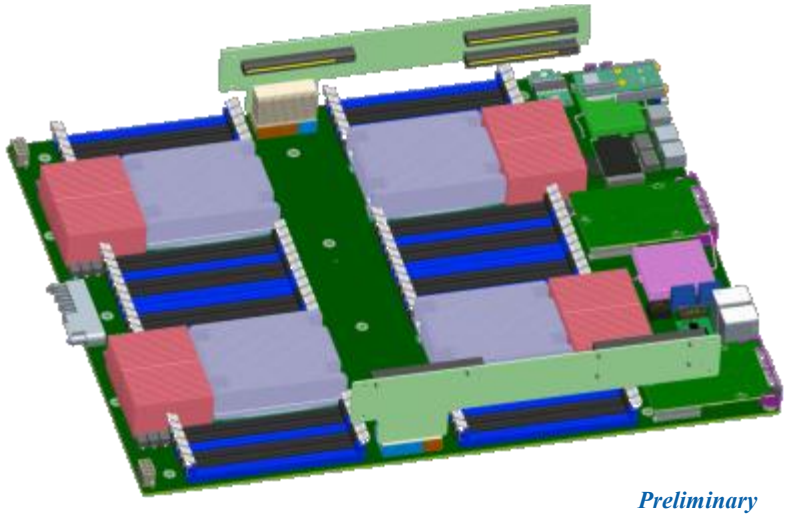

Preliminary

Jun-Sept '12

- **Up to two NVIDIA® Tesla® GPUs**
- **Intel® Xeon® processor EP E5-4600**
  - Up to 8 Cores and130W TDP
- **48 memory sockets support LR/U/R-DIMMs**
  - Sixteen 3DPC channels native DDR-3 (800/1066/1333/1600*)
- **Integrated BMC with remote management & KVM option**
  - Optional dedicated management LAN port & Activation key for advanced features
- **SNB-EP Integrated PCI-E Gen 3 I/O: Configuration:**
  - Two PCIe3 x48 Risers
    - Four x16 external slots (FLFH or HLFH)
    - Two x16 internal slots (HLFH)
  - Two IO x8 modules

SGI® Solutions for NVIDIA® Tesla®

QUESTIONS?

©2012 SGI