March, 2013

Boost Your Productivity with GPGPUs and IBM Platform Computing Software

NVIDIA GTC 2013
Chris Porter, IBM

Agenda

- IBM Platform Computing offerings
- GPGPU Adoption in the HPC Market
- GPGPU Scheduling & Management
  - IBM Platform Computing Solutions for GPGPUs
  - Benefits from Intelligent GPU Scheduling & Management
  - Use Case Examples
- Summary

IBM PLATFORM COMPUTING OFFERINGS

IBM Platform Computing

The leader in cluster, grid and HPC cloud management software

- Acquired by IBM in 2012 as part of mainstream Technical Computing strategy
- 20 year history delivering leading workload and resource management software for technical computing and big data/analytics environments
- 2000+ global customers including 23 of 30 largest enterprises
- Market leading scheduling engine with high performance, mission-critical reliability and extreme scalability
- Comprehensive capability from ready-to-deploy complete cluster systems to large global grids to HPC clouds
- Large ISV and global partner ecosystem
- Global services and support coverage

De facto Standard for Commercial HPC

60% of top Financial Services

Over 5 MM CPUs under management

IBM Platform Computing offerings

Workload Management	**Platform LSF Family**	Scalable, comprehensive workload management suite for heterogeneous compute environments	• Unmatched experience through market share • Powerful multi-policy scheduling engine • Unmatched scalability through high end accounts • Unmatched breadth of offering due to extent of add-ons
	Platform HPC for System x	Simplified, integrated, purpose-built HPC management software integrated with systems	• All-in-one integrated solution with leading web interface • Applicable to the smallest of clusters • Leverages Platform LSF technology base • Hardware bundled for turnkey purchasing and deployment
Analytics Infrastructure	**Platform Symphony Family**	High-throughput, low-latency compute and data intensive analytics applications	• Leading experience due to 50%+ major investment banks as customers (translates to other industries) • High scalability and better application performance due to fast, low latency processing (sub millisecond) • Proven business model for sharing grid infrastructure • Both compute and data intensive applications
Flexible Clusters	**Platform Cluster Manager**	Provisioning and management of HPC clusters	• Scalable offerings to simplify process of deploying and managing small clusters to global HPC clouds • Broad heterogeneous support enables managing broad technologies and multiple workload managers • Enables multi-tenant HPC clouds

The Application Accelerator Storm

- GPU adoption is increasing
  - 53 systems on the Top500 released in Nov, 2012 are using GPGPUs
  - GPGPUs are penetrating both high-end and mainstream HPC
- Nvidia is leading the accelerator race
  - 100's of K's of trained CUDA developers worldwide
  - 50 systems powered by Nvidia on the latest Top500 list
- Other accelerator technologies are emerging
  - Intel: Xeon Phi Coprocessor
  - AMD: FireStream

GPGPU ADOPTION IN THE HPC MARKET

Market Landscape: Technical Applications are Exploding

Adoption Drivers

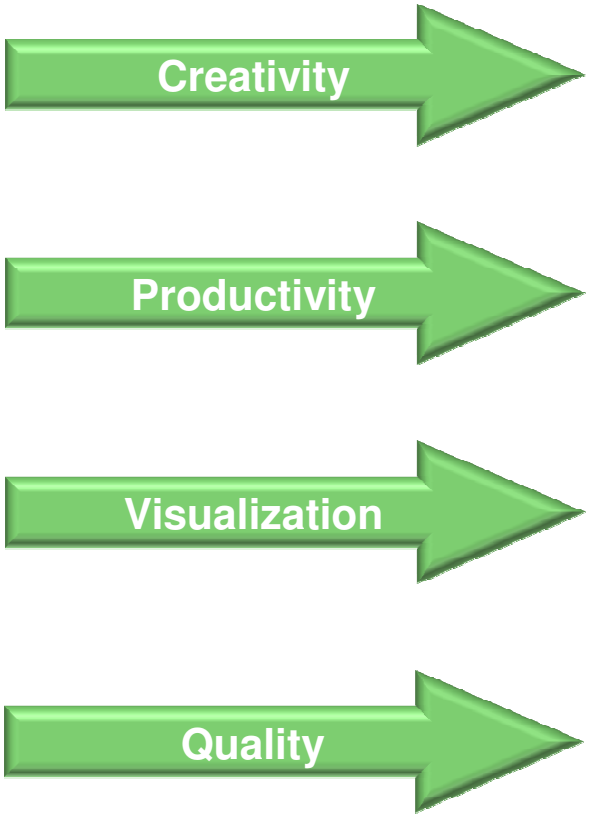

Technical Applications

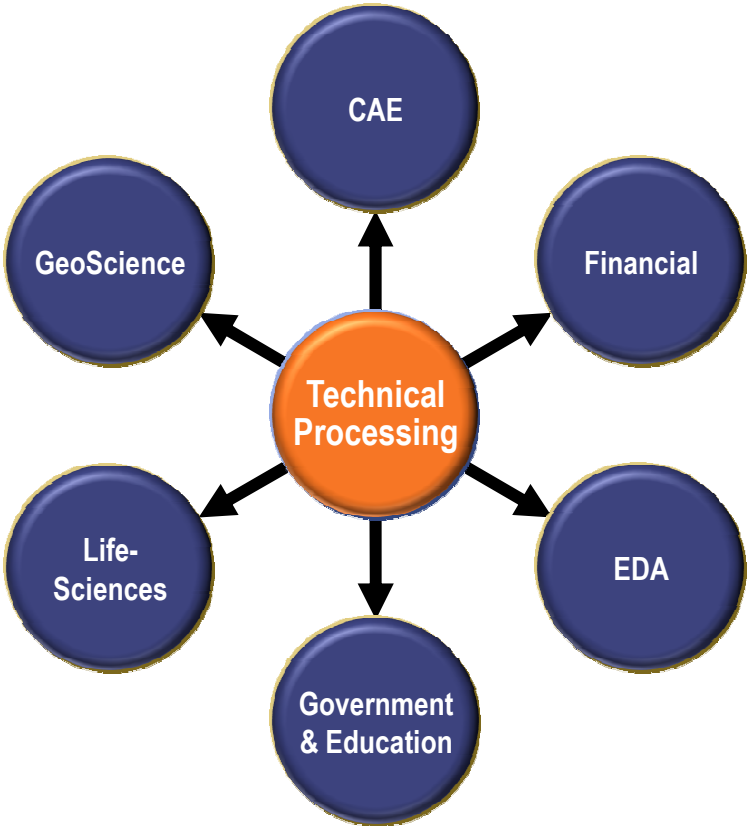

The Big Buzz in HPC: Hybrid Computing

- Hybrid Computing: CPUs and GPUs working together

When do I use them?
What is the ROI?
How do I schedule jobs to them?
How to maximize utilization, various published benchmarks showing dramatic performance increases?

- Applications Taking Advantage Of GPUs
  - Life Sciences
    - Unipro UGENE, Agile Molecule, many others
  - Financial Services
    - Volmaster FX, ClusterTech Financial Library, many others
  - Manufacturing
    - Fidesys, Ansys, 3ds, many others
  - Oil and Gas
    - Acceleware Seismic Solvers, many others

GPU SCHEDULING & MANAGEMENT

What do Intelligent GPU Scheduling and Management Bring to You?

- Improved application performance by allocating GPU suitable workloads on those resources and free up CPUs for other types of workloads.
- Reduced infrastructure cost by maximizing cluster utilization.
- Simplified system management via easy to use GUI and timely alerts.
- Increased productivity for administrators and application developers.

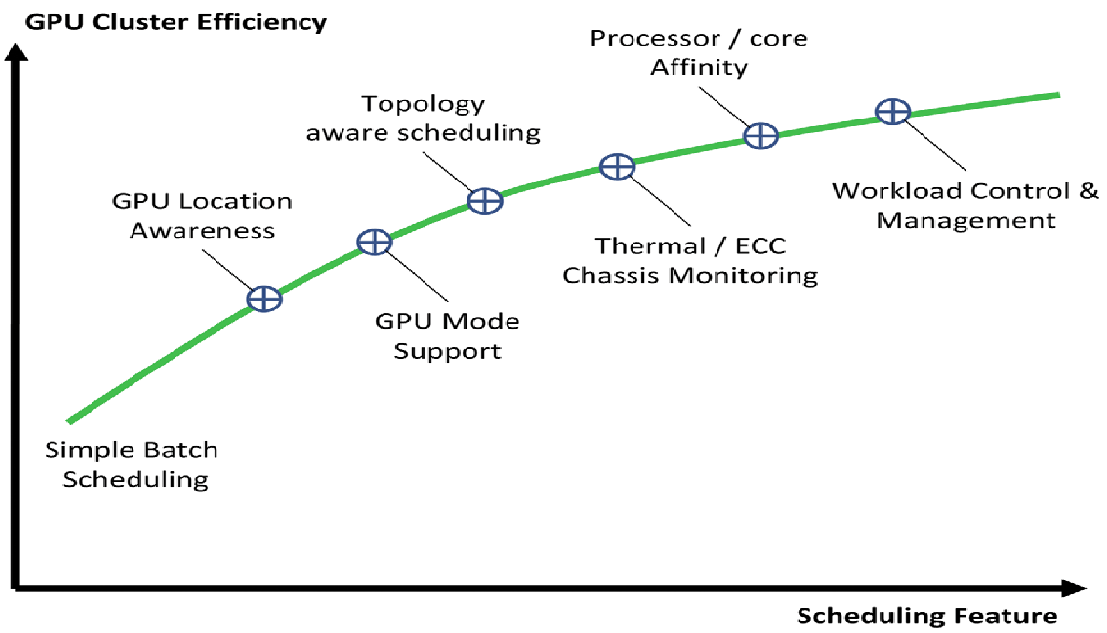

Intelligent scheduling improves cluster efficiency

GPUs: Schedule, Monitor & Manage

- DEPLOY: Quickly deploy workload to GPU resources
  - Easy job submission to GPUs in a cluster via CUDA job submission wrappers
  - Install CUDA across a cluster is a couple of clicks
- MANAGE: Easily manage heterogeneous clusters
  - Deploy & manage both CPU & GPU resources in the same cluster
  - Remotely manage & view the status of your jobs

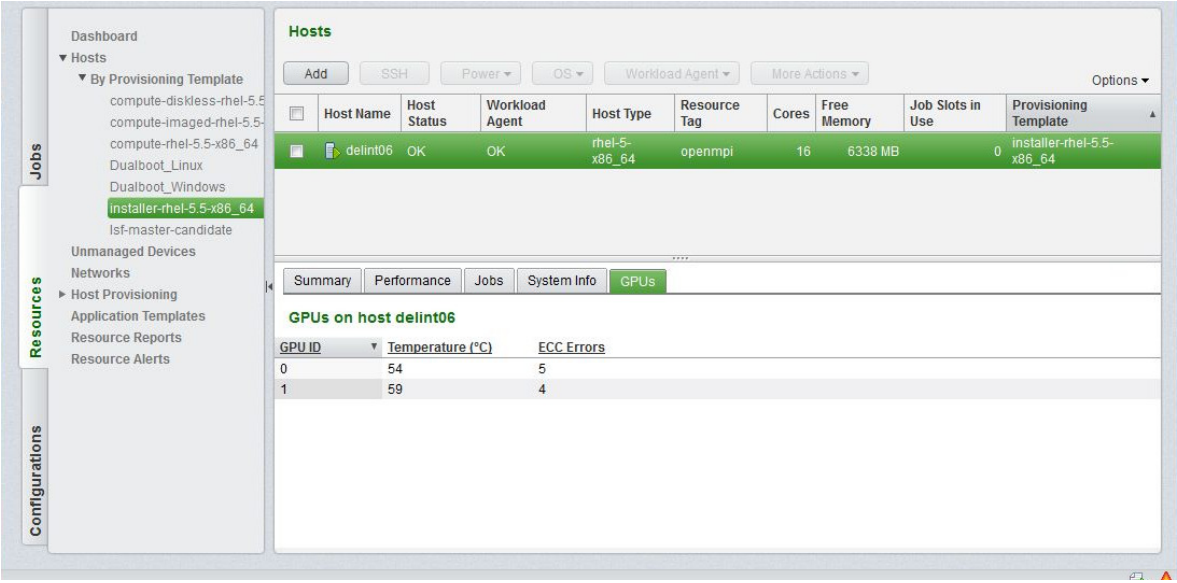

Take immediate advantage of the exceptional HPC performance provided by GPUs

GPUs: Schedule, Monitor & Manage

- MONITOR: Monitor GPU metrics
  - GPU slot utilization, temperature & status
  - Detect ECC error accumulation

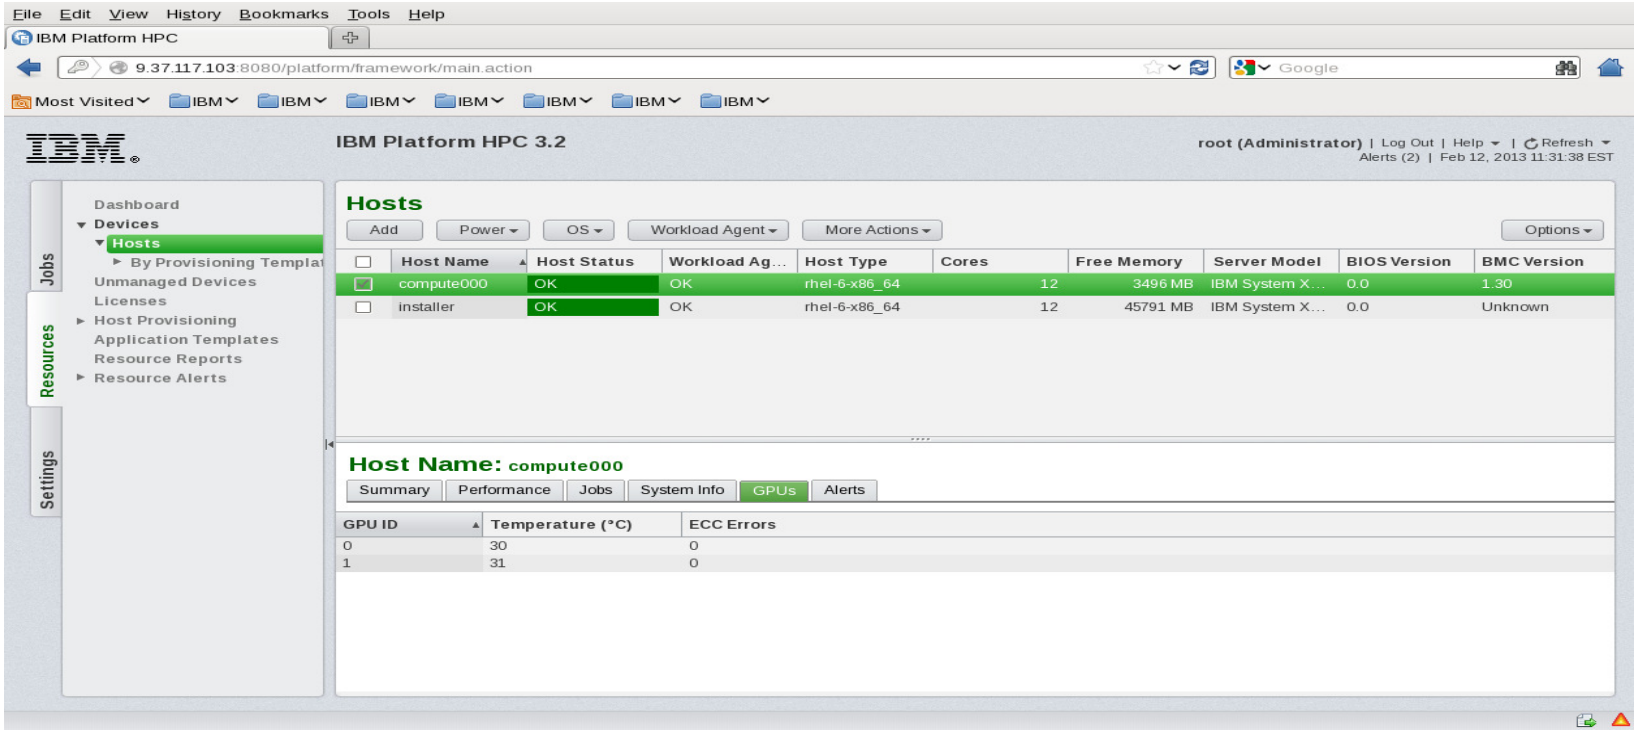

Scheduling to GPGPUs Today

- Managing latest GPGPUs and CUDA (V5.0) applications using:
  - IBM Platform LSF
  - IBM Platform HPC
  - IBM Platform Symphony
- GPU ELIM provides:
  - Monitoring and detection of GPUs
  - Group hosts with GPU(s) into a resource group
  - Compute slots on these hosts are user configured

GPU Management:
- # of GPU
- # of GPU in “normal”
- # of GPU in “exclusive”
- # of GPU in “prohibited”

GPU Monitoring:
- Mode (normal, exclusive, prohibited)
- Temperature
- ECC error count

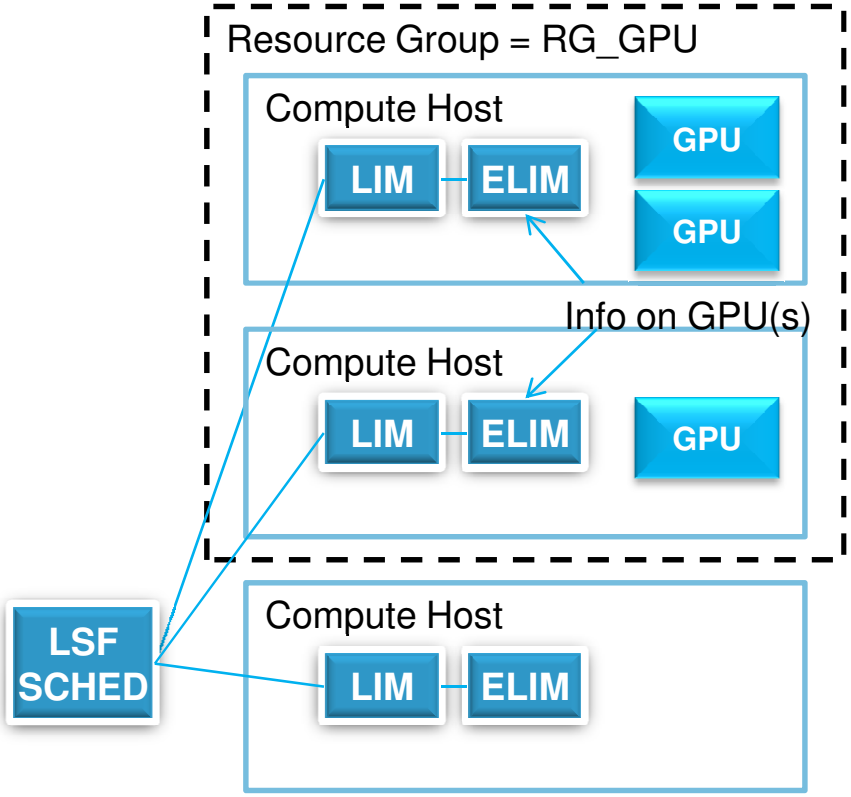

- *GPU-enablement is the responsibility of the application developer*

USE CASES

Use Case #1: Simple Use Case

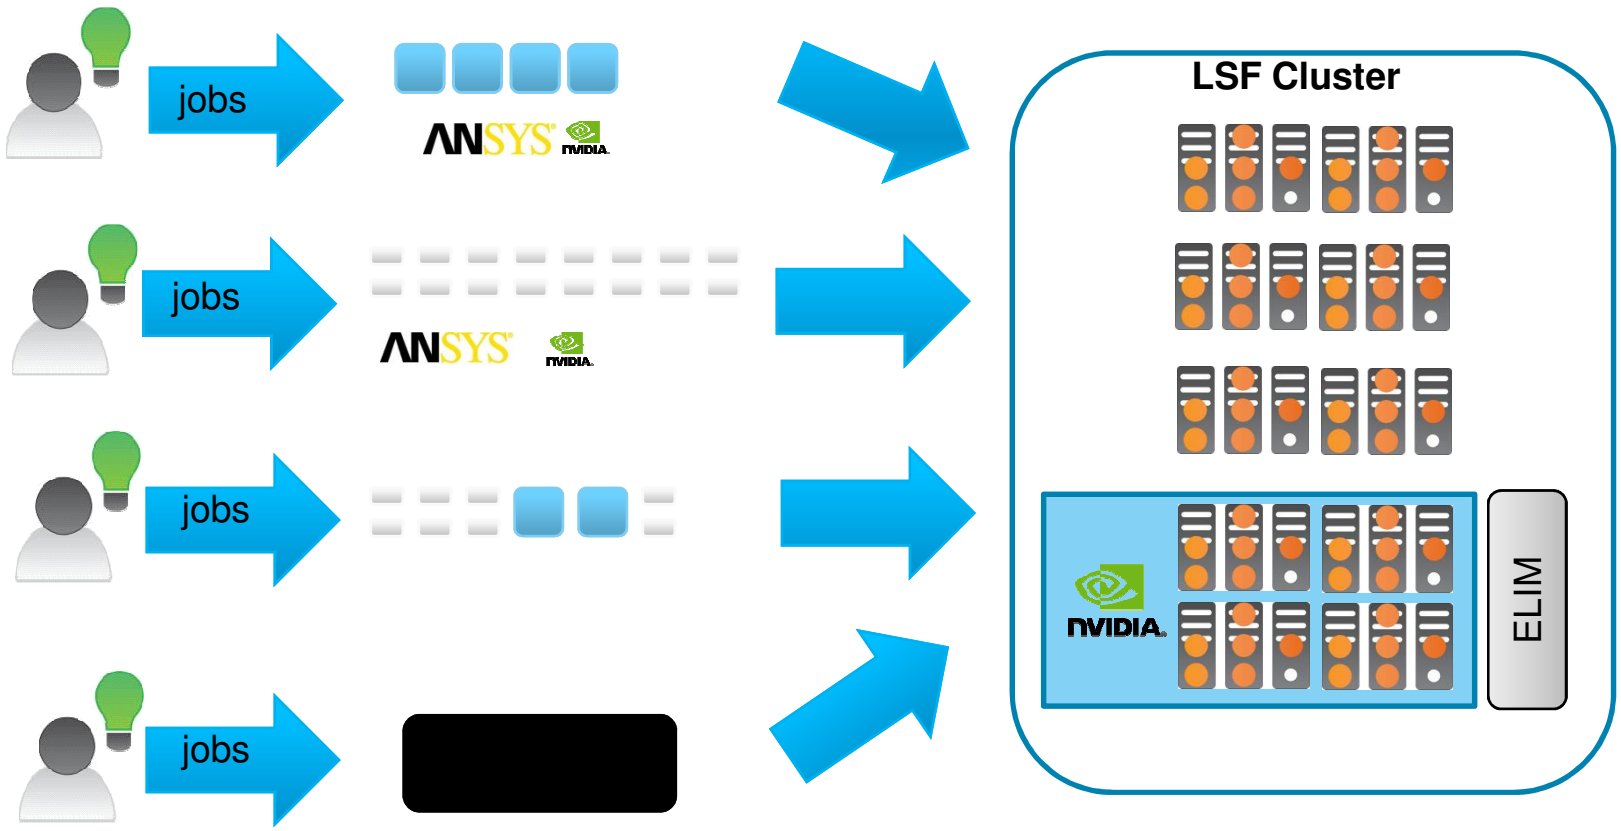

- Nvidia GPGPU only
- CUDA 5.0 and older
- Simple monitoring statistics

Use Case #2: Complex Use Case

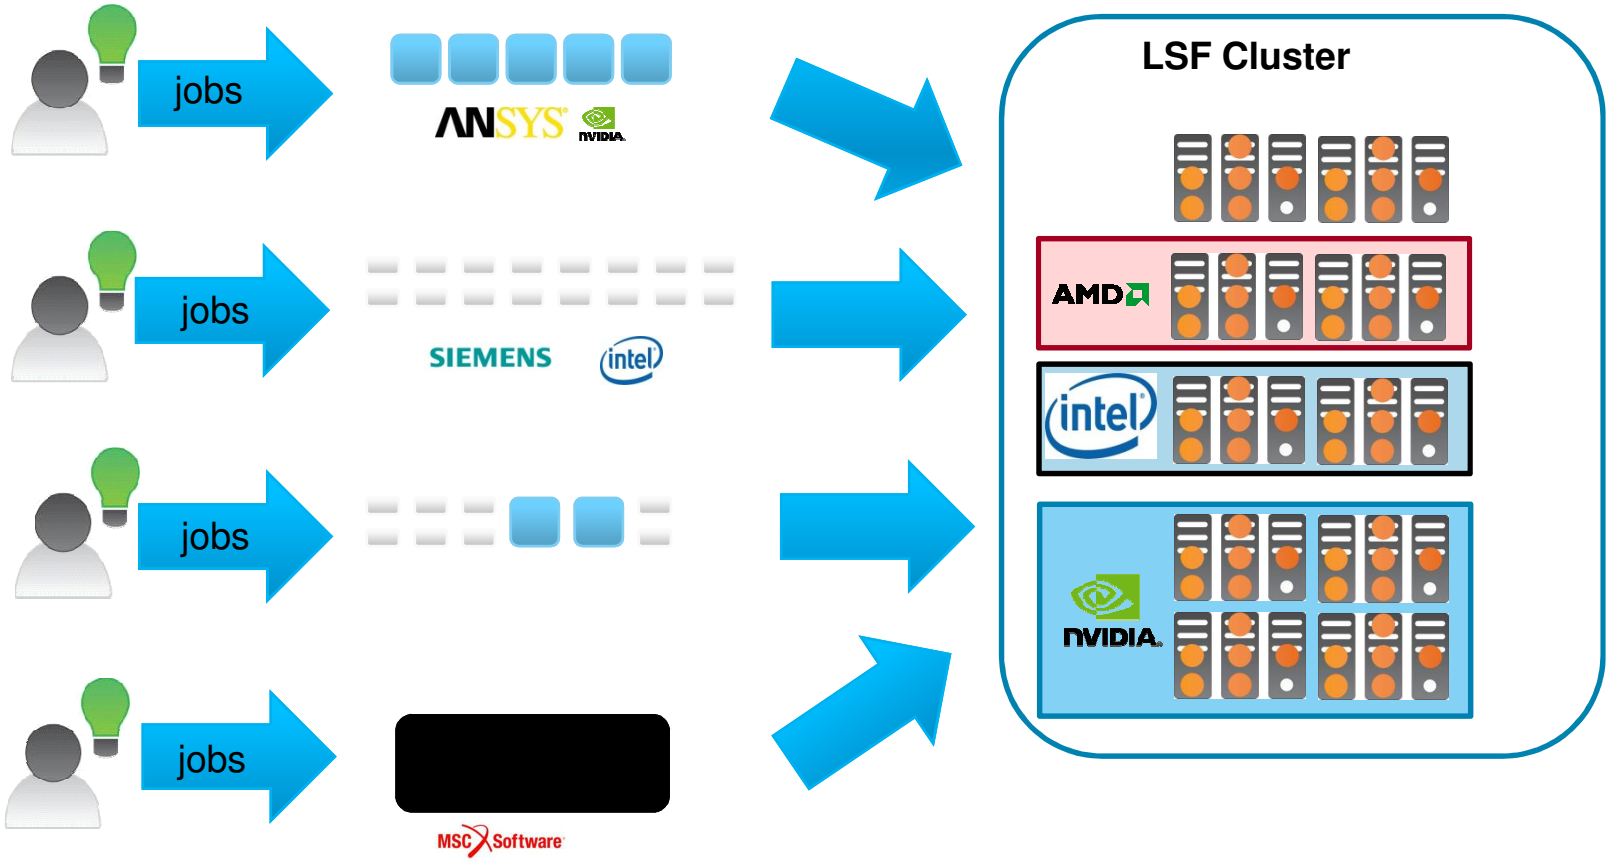

- Multiple GPGPU / accelerators **OR**
- Use of newer CUDA features > 3.2 **OR**
- Monitoring of memory and GPU core utilization

Use Case #3: NUMA optimization within a single server

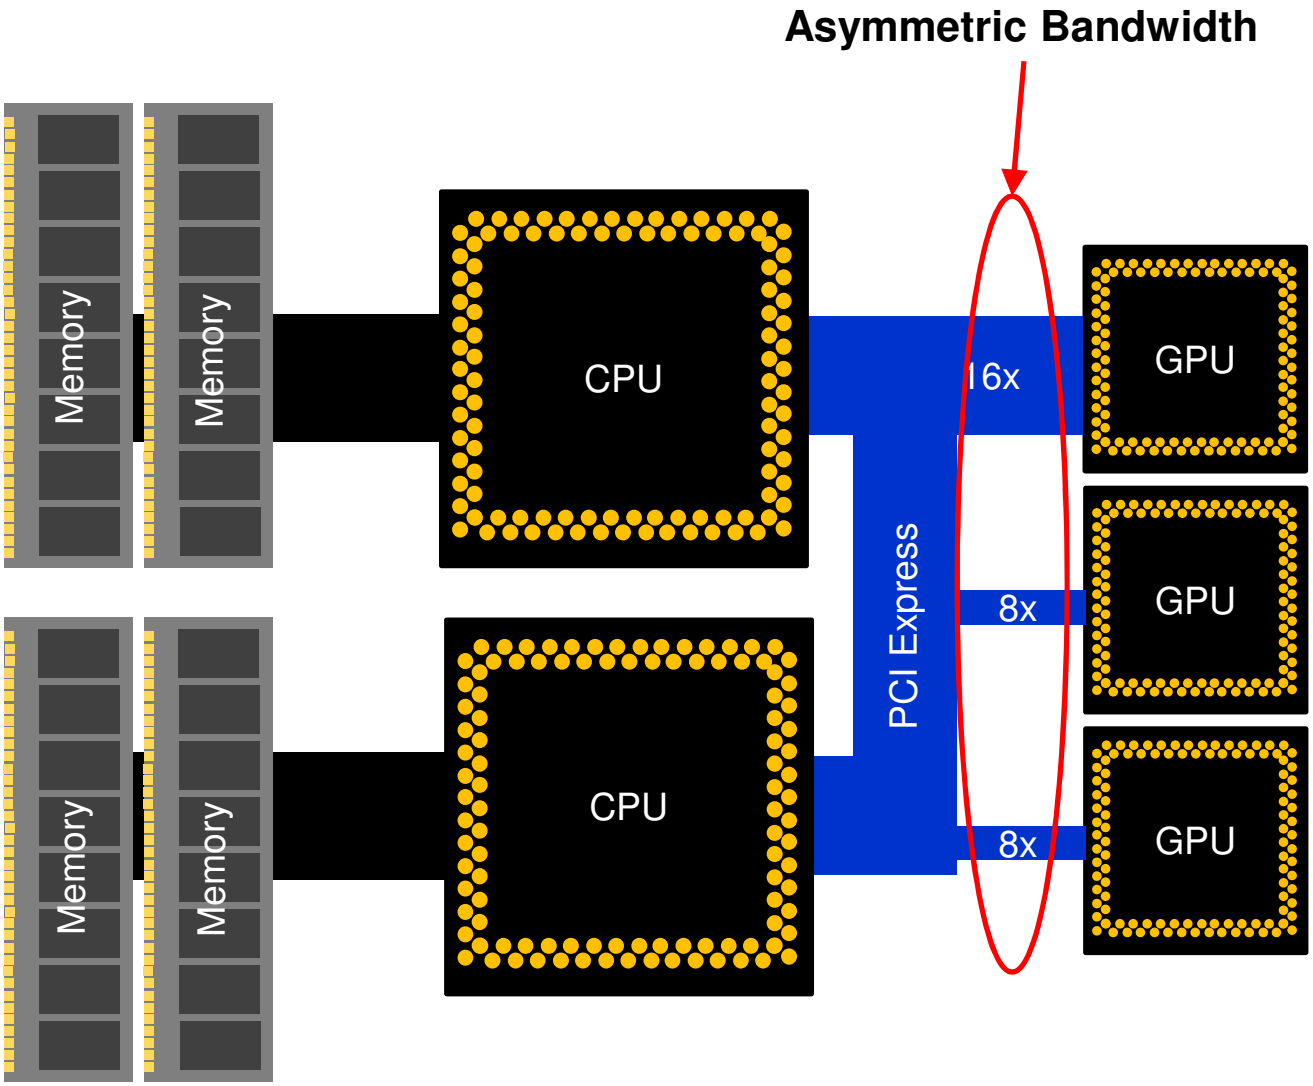

Use Case #3: NUMA Optimization within a single server

Asymmetric bandwidth requires:

- LSF: Non-GPU jobs to be scheduled to hosts without GPUs first
- LSF: Non-GPU jobs be scheduled to cores with low GPU bandwidth
- LSF: GPU jobs schedule to cores with maximum GPU bandwidth

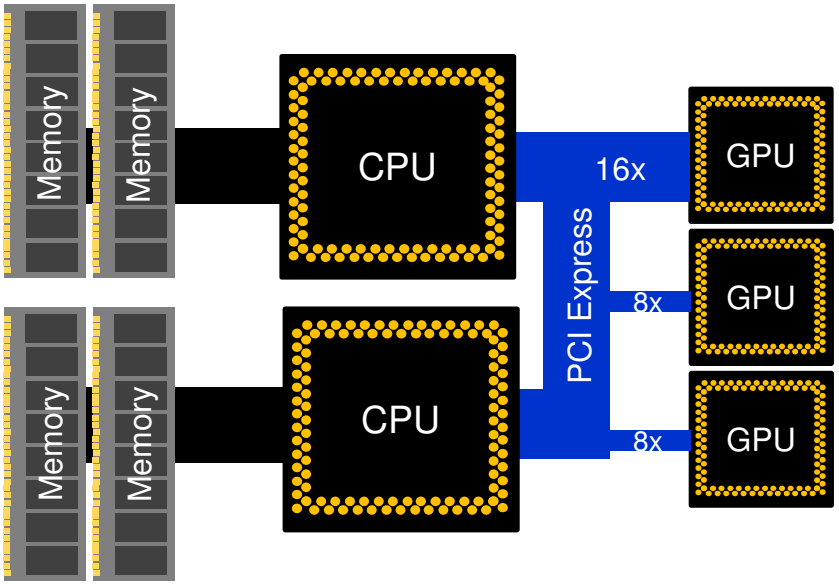

Use Case #4: NUMA optimization for multi-server MPI jobs

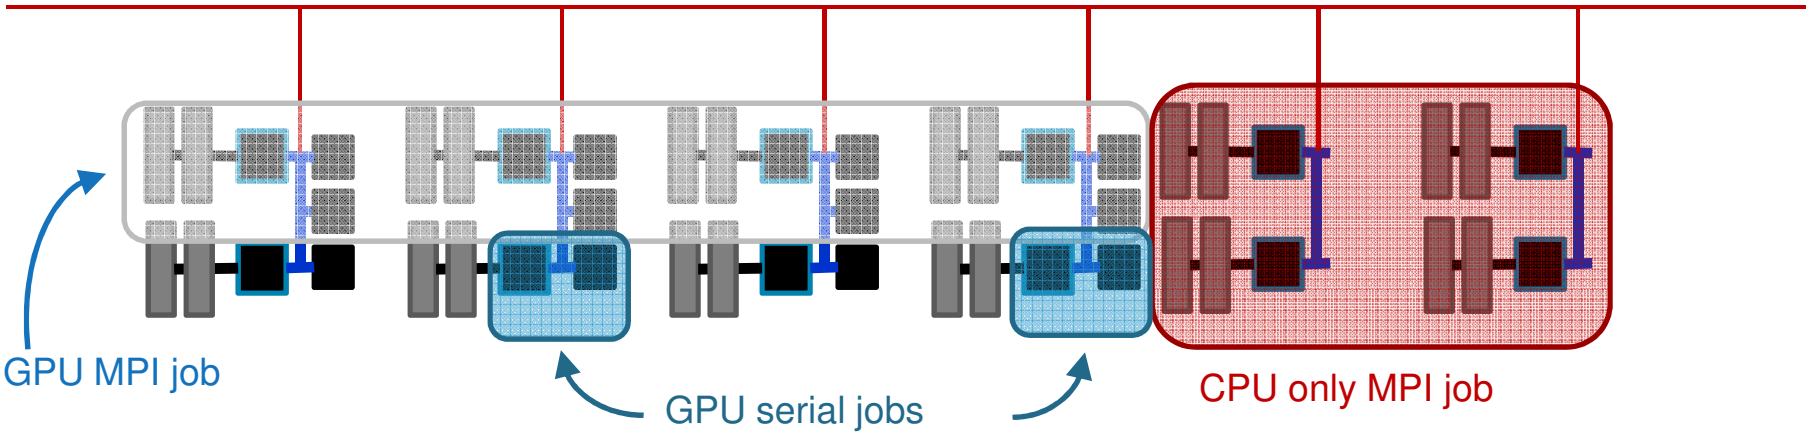

MPI job optimization

– MPI selects optimal cores for multi-host job MPI processes

Use Case #4: NUMA Optimization for multi-server MPI jobs

MPI based multi-server GPU and non-GPU jobs

– LSF: Single servers – LSF scheduling plugin controls core placement

– LSF: Multiple servers – LSF scheduling plugin does nothing

– MPI: Single servers – MPI scheduler does nothing

– MPI: Multiple servers – LSF scheduling plugin controls core placement

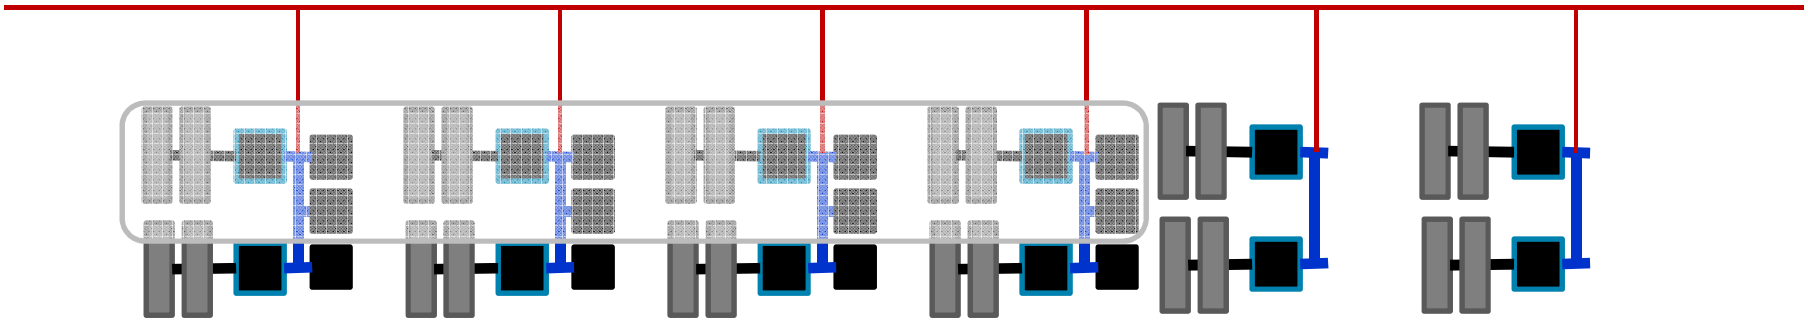

SUMMARY

Reality and Conclusions

Don’t have application developed for GPUs?

- Many ISVs are working hard to adopt CUDA and/or openCL for their applications

IBM Platform Computing has available solutions for

- Managing both GPU and CPU resources in a cluster
- Monitoring & visualizing important parameters for GPUs
- Scheduling serial jobs to available and functional GPUs
- Scheduling parallel jobs to available and functional GPUs
- Scheduling & optimizing mixed mode serial and parallel workload

QUESTIONS?