Remote GPU Virtualization: Is It Useful?

F. Silla, J. Prades, S. Iserte, and C. Reaño

The 2nd IEEE International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era, HiPINEB 2016

What is "Remote GPU Virtualization"?

What is "Remote GPU Virtualization"?

What is "Remote GPU Virtualization"?

Basic behavior of CUDA

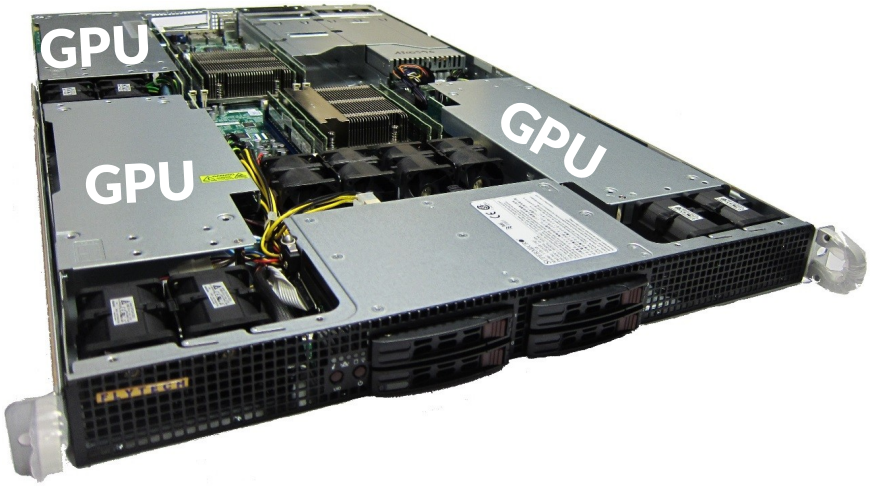

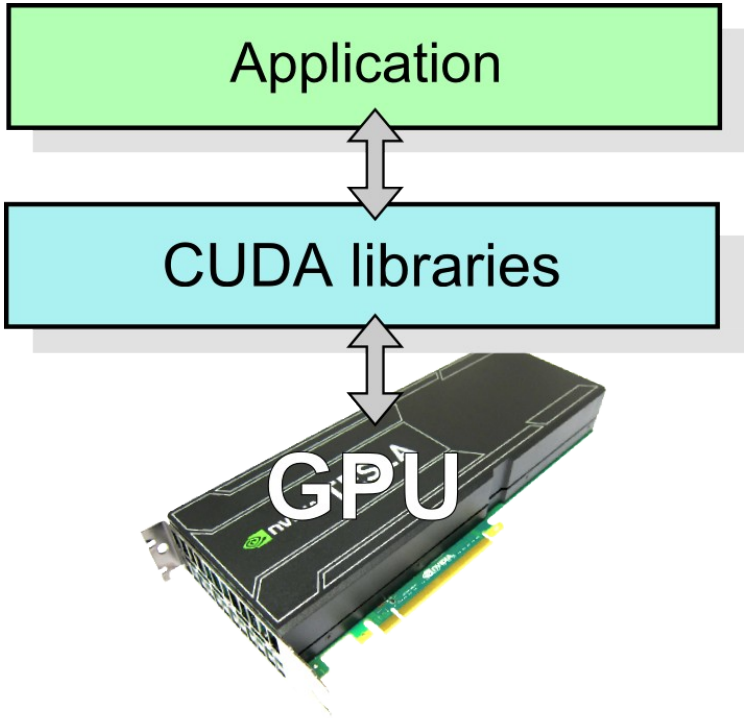

What is "Remote GPU Virtualization"?

Basic behavior of remote GPU virtualization solutions

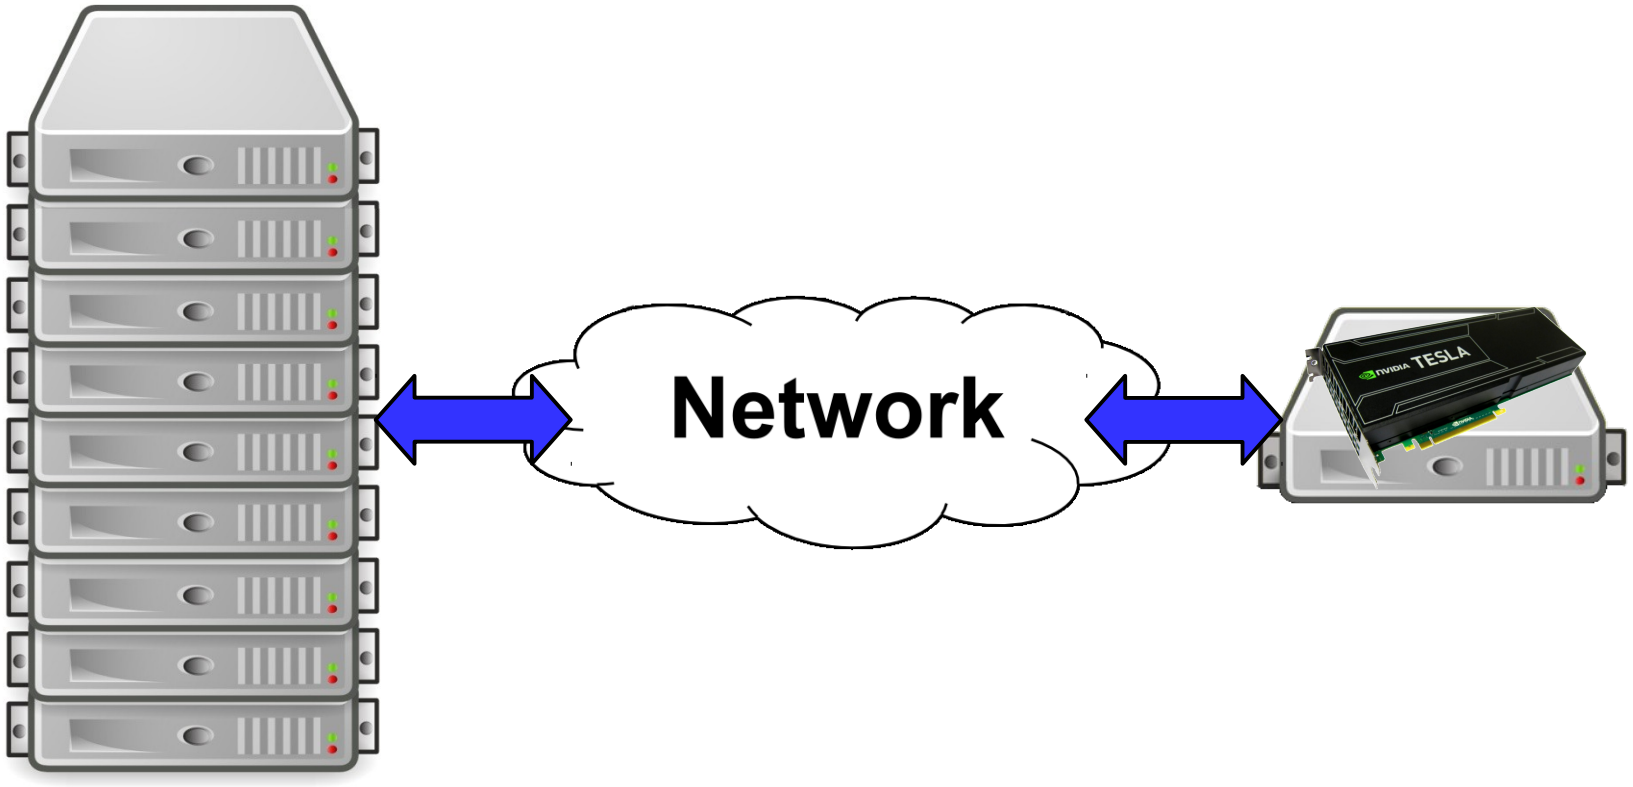

What is "Remote GPU Virtualization"?

Remote GPU virtualization solutions

- **rCUDA (CUDA 7.0)**
- GVirtuS (CUDA 3.2)
- DS-CUDA (CUDA 4.1)

Publicly available

- vCUDA (CUDA 1.1)
- GViM (CUDA 1.1)
- GridCUDA (CUDA 2.3)
- V-GPU (CUDA 4.0)

NOT publicly available

What is "Remote GPU Virtualization"?

Basic behavior of rCUDA

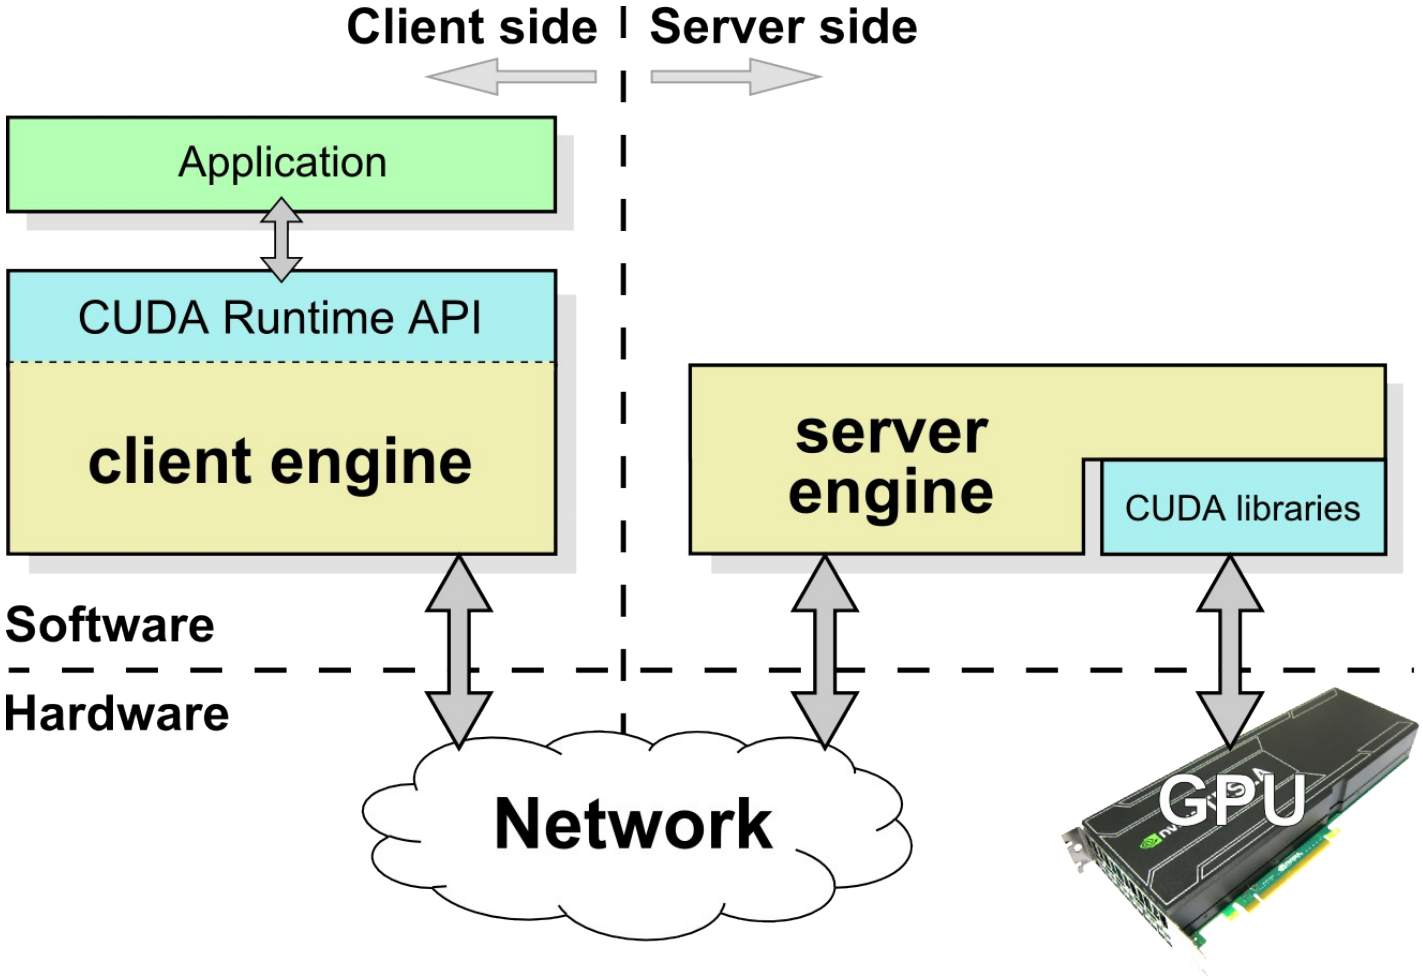

What is "Remote GPU Virtualization"?

Basic behavior of rCUDA

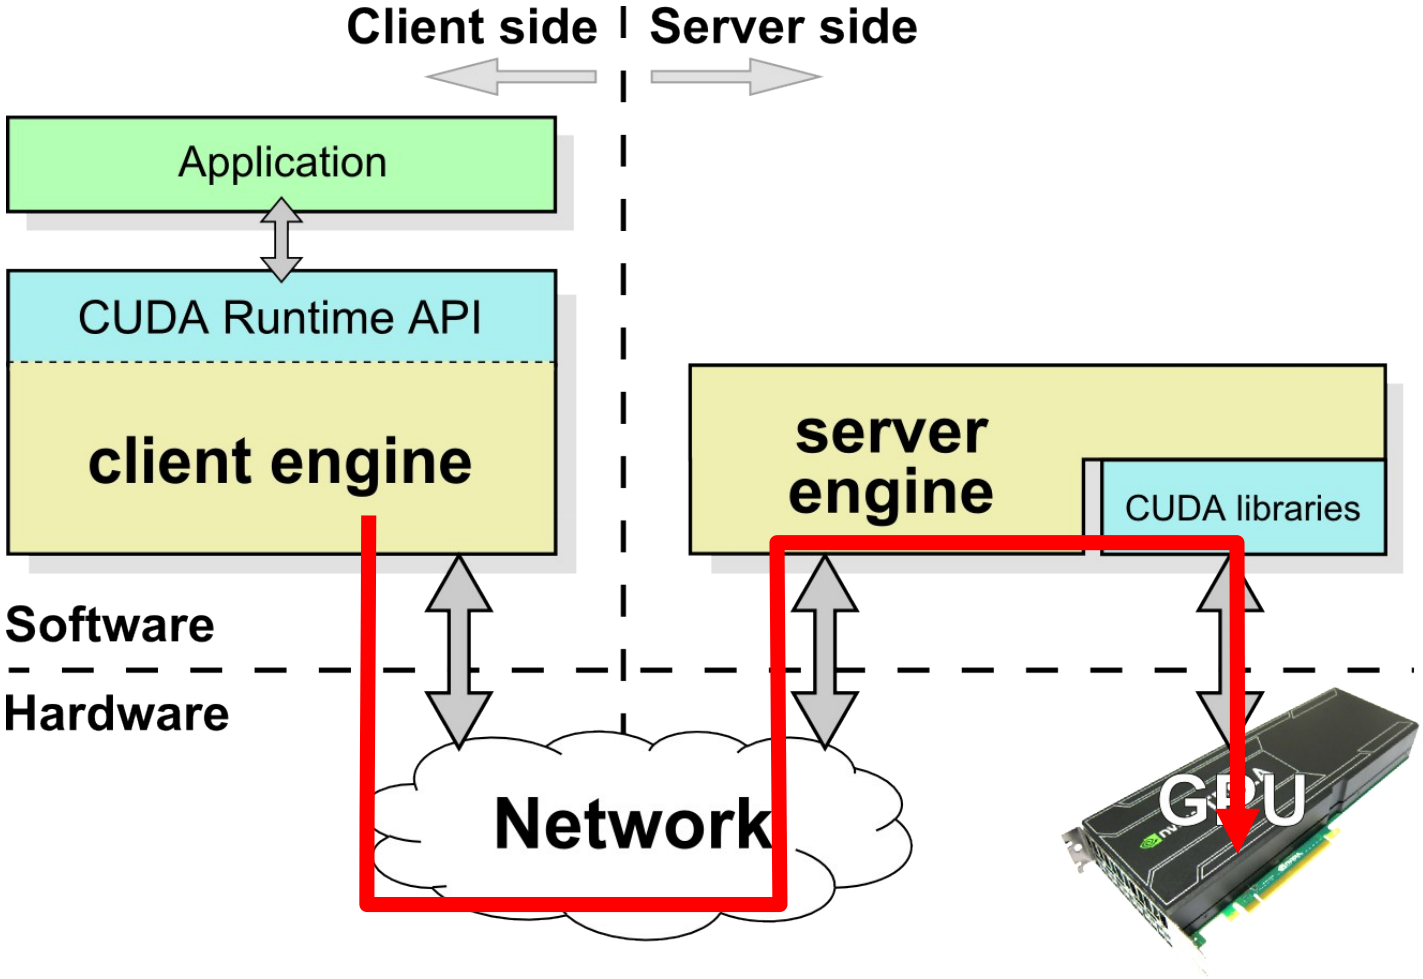

What is "Remote GPU Virtualization"?

Basic behavior of rCUDA

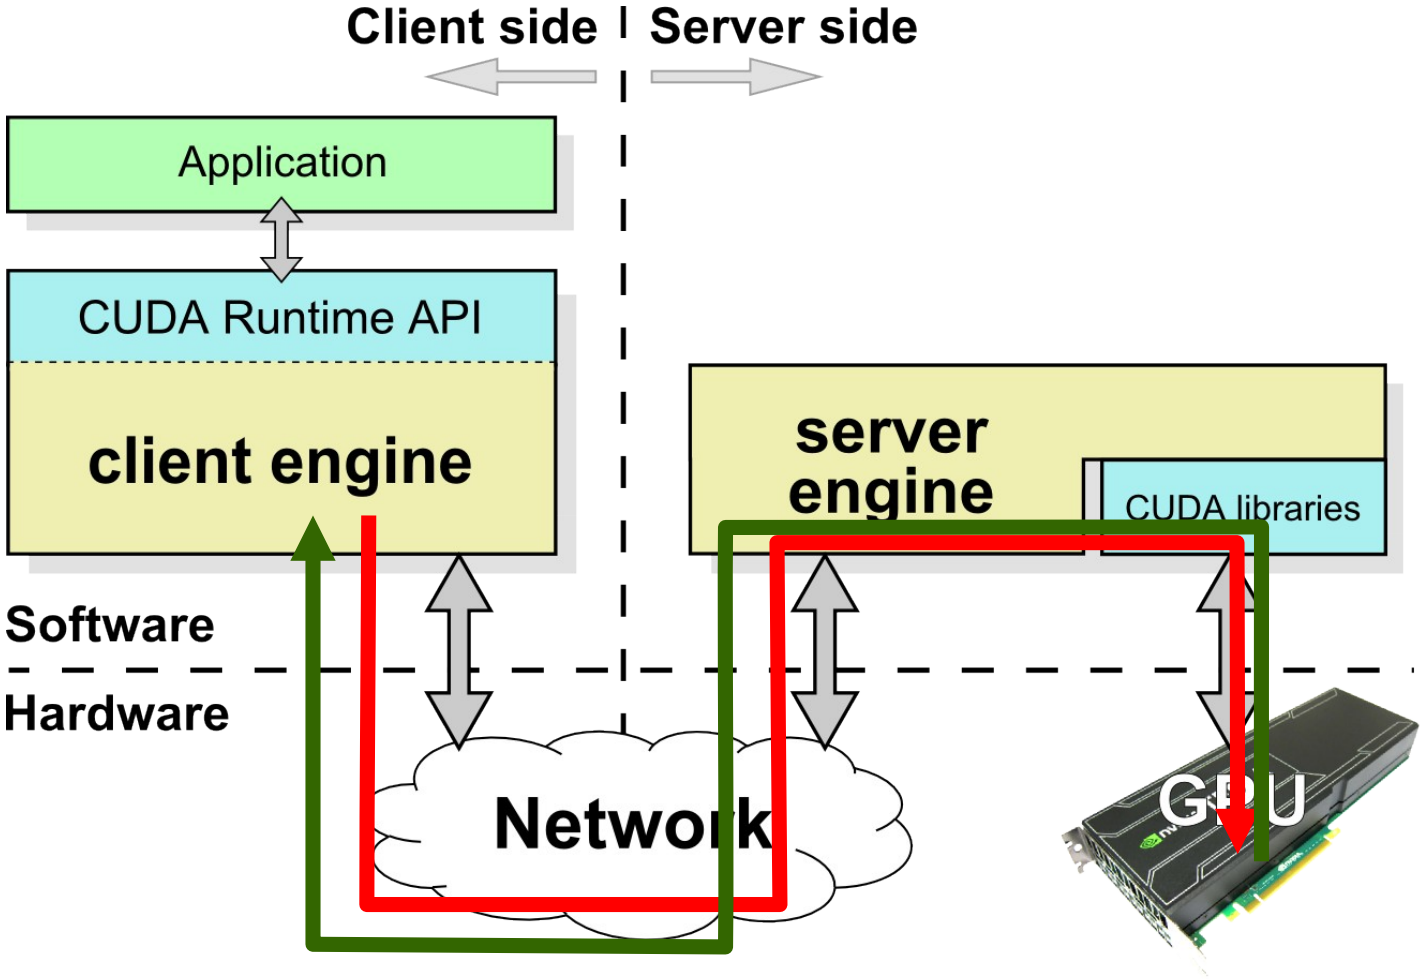

What is "Remote GPU Virtualization"?

Remote GPU virtualization envision

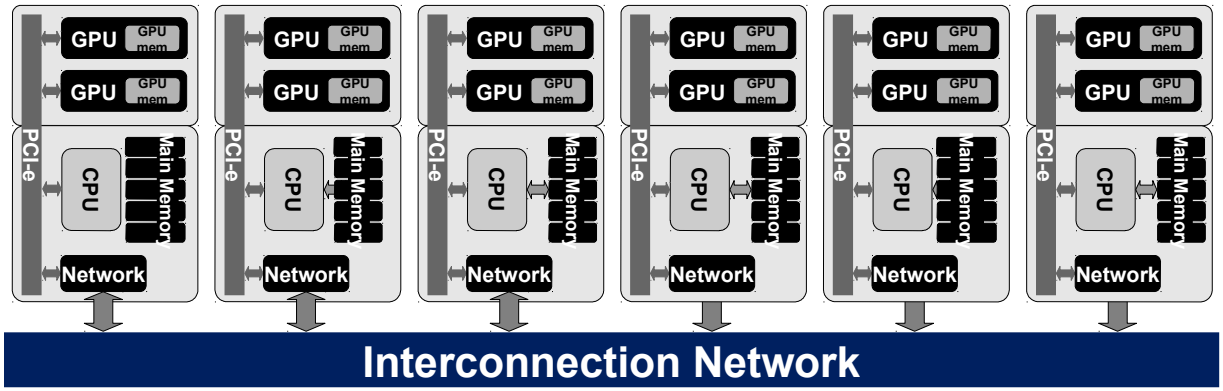

What is "Remote GPU Virtualization"?

Remote GPU virtualization envision

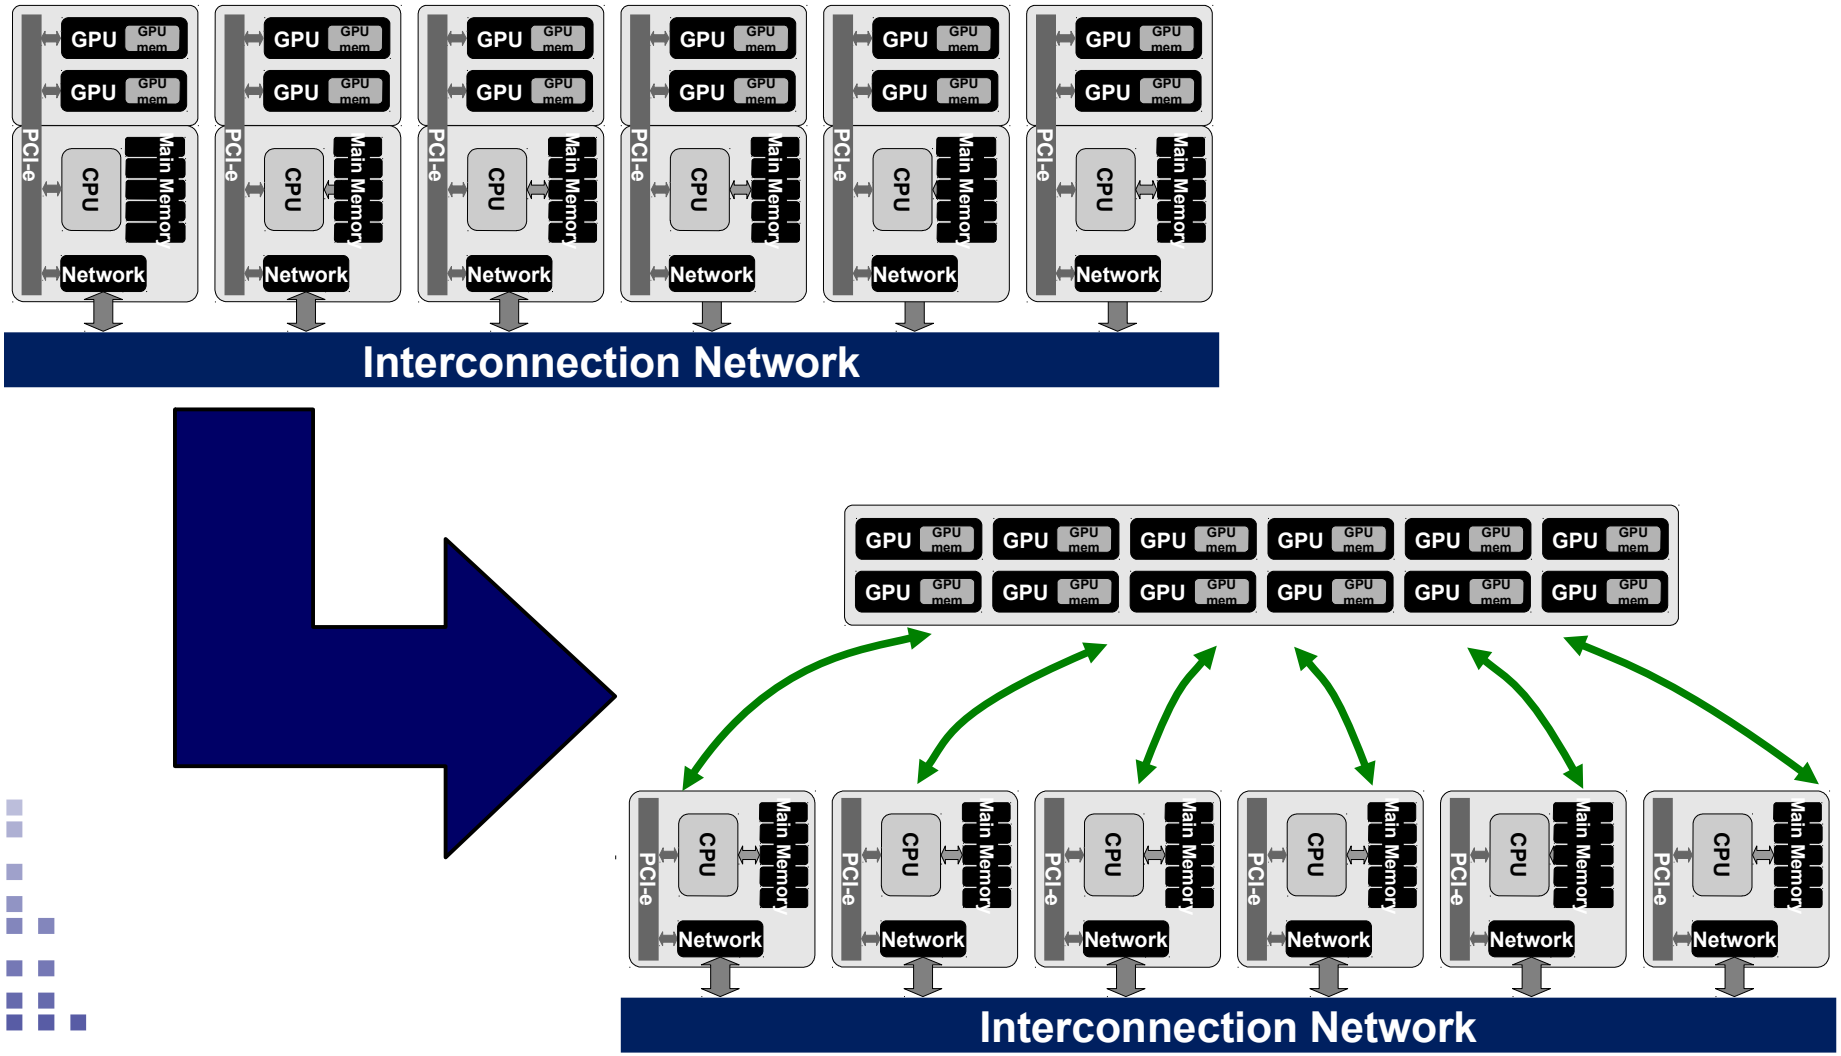

Benefits of Using Remote GPU Virtualization

Benefits of Using Remote GPU Virtualization

1st → More GPUs are available for a single application

Only the GPUs in the node can be provided to the application

Without GPU virtualization

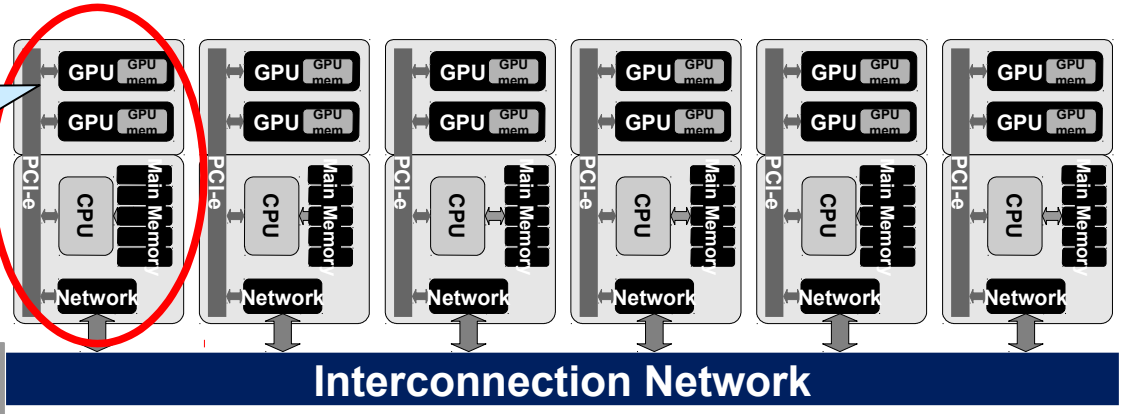

With GPU virtualization

Many GPUs in the cluster can be provided to the application

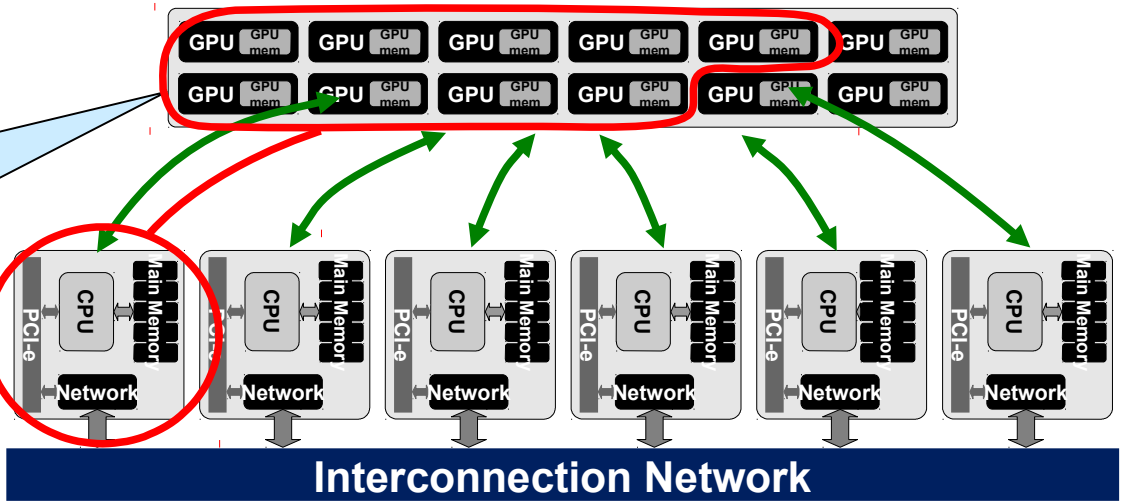

Benefits of Using Remote GPU Virtualization

1st → More GPUs are available for a single application

64 GPUs!

```
bsc19421@nvb127:~
./deviceQuery Starting...

 CUDA Device Query (Runtime API) version (CUDART static linking)

Detected 64 CUDA Capable device(s)

Device 0: "Tesla M2090"
  CUDA Driver Version / Runtime Version          5.0 / 5.0
  CUDA Capability Major/Minor version number:    2.0
  Total amount of global memory:                 6144 MBytes (6442123264 bytes)
  (16) Multiprocessors x ( 32) CUDA Cores/MP:    512 CUDA Cores
  GPU Clock rate:                                1301 MHz (1.30 GHz)
  Memory Clock rate:                             1848 Mhz
  Memory Bus Width:                              384-bit
  L2 Cache Size:                                 786432 bytes
  Max Texture Dimension Size (x,y,z)             1D=(65536), 2D=(65536,65535), 3D=(2048,2048,2048)
  Max Layered Texture Size (dim) x layers        1D=(16384) x 2048, 2D=(16384,16384) x 2048
  Total amount of constant memory:               65536 bytes
  Total amount of shared memory per block:       49152 bytes
  Total number of registers available per block: 32768
  Warp size:                                     32
  Maximum number of threads per multiprocessor:  1536
  Maximum number of threads per block:           1024
  Maximum sizes of each dimension of a block:    1024 x 1024 x 64
  Maximum sizes of each dimension of a grid:     65535 x 65535 x 65535
  Maximum memory pitch:                          2147483647 bytes
  Texture alignment:                             512 bytes
  Concurrent copy and kernel execution:          Yes with 2 copy engine(s)
  Run time limit on kernels:                     No
  Integrated GPU sharing Host Memory:            No
  Support host page-locked memory mapping:       No
  Alignment requirement for Surfaces:            Yes
  Device has ECC support:                        Disabled
  Device supports Unified Addressing (UVA):      Yes
  Device PCI Bus ID / PCI location ID:           2 / 0
  Compute Mode:
     < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

Device 1: "Tesla M2090"
  CUDA Driver Version / Runtime Version          5.0 / 5.0
```

Benefits of Using Remote GPU Virtualization

1st → More GPUs are available for a single application

MonteCarlo-MultiGPU (NVIDIA Samp.)

Supermicro SYS7047GR-TRF server with two Intel Xeon E5-2620 v2 processors.

FDR InfiniBand Interconnection Netw. and K20 GPUs

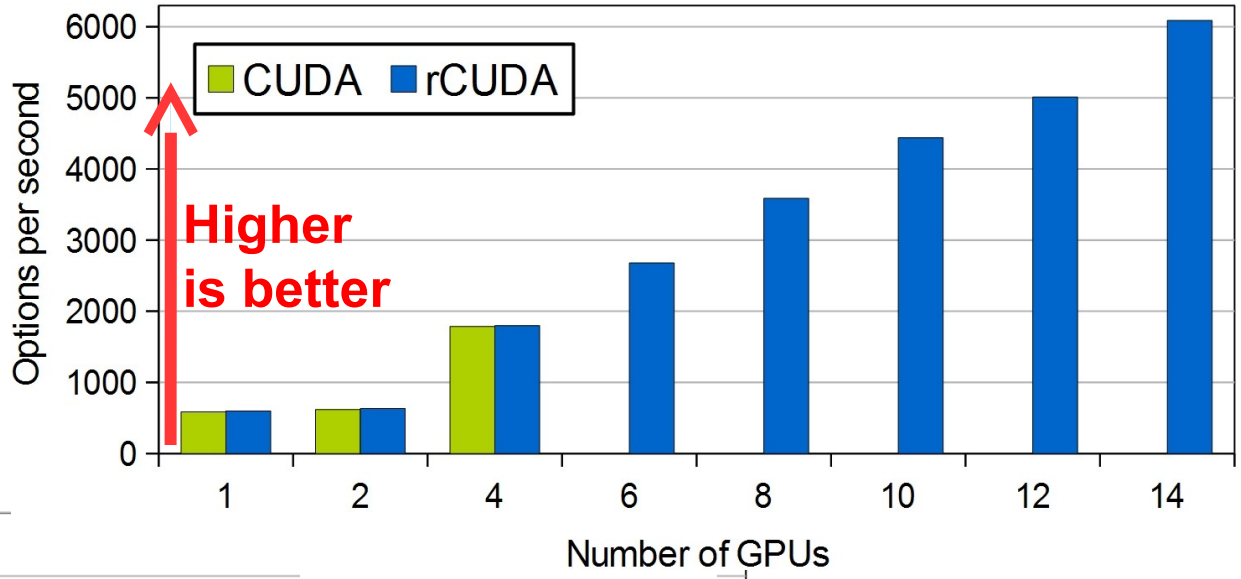

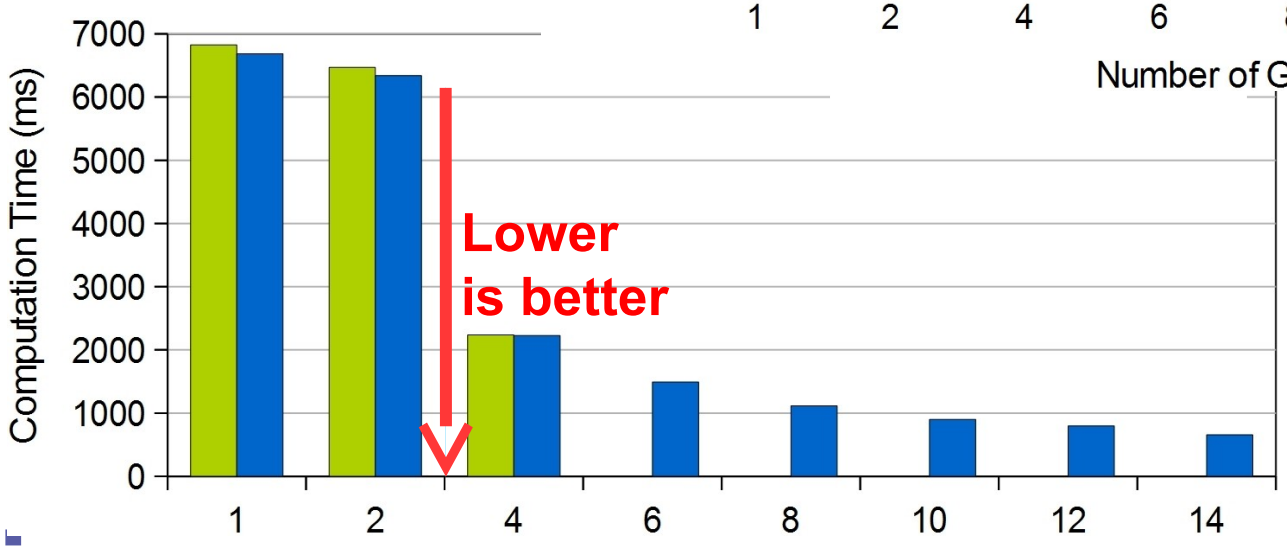

Benefits of Using Remote GPU Virtualization

2nd → Throughput is increased & energy consumption is reduced

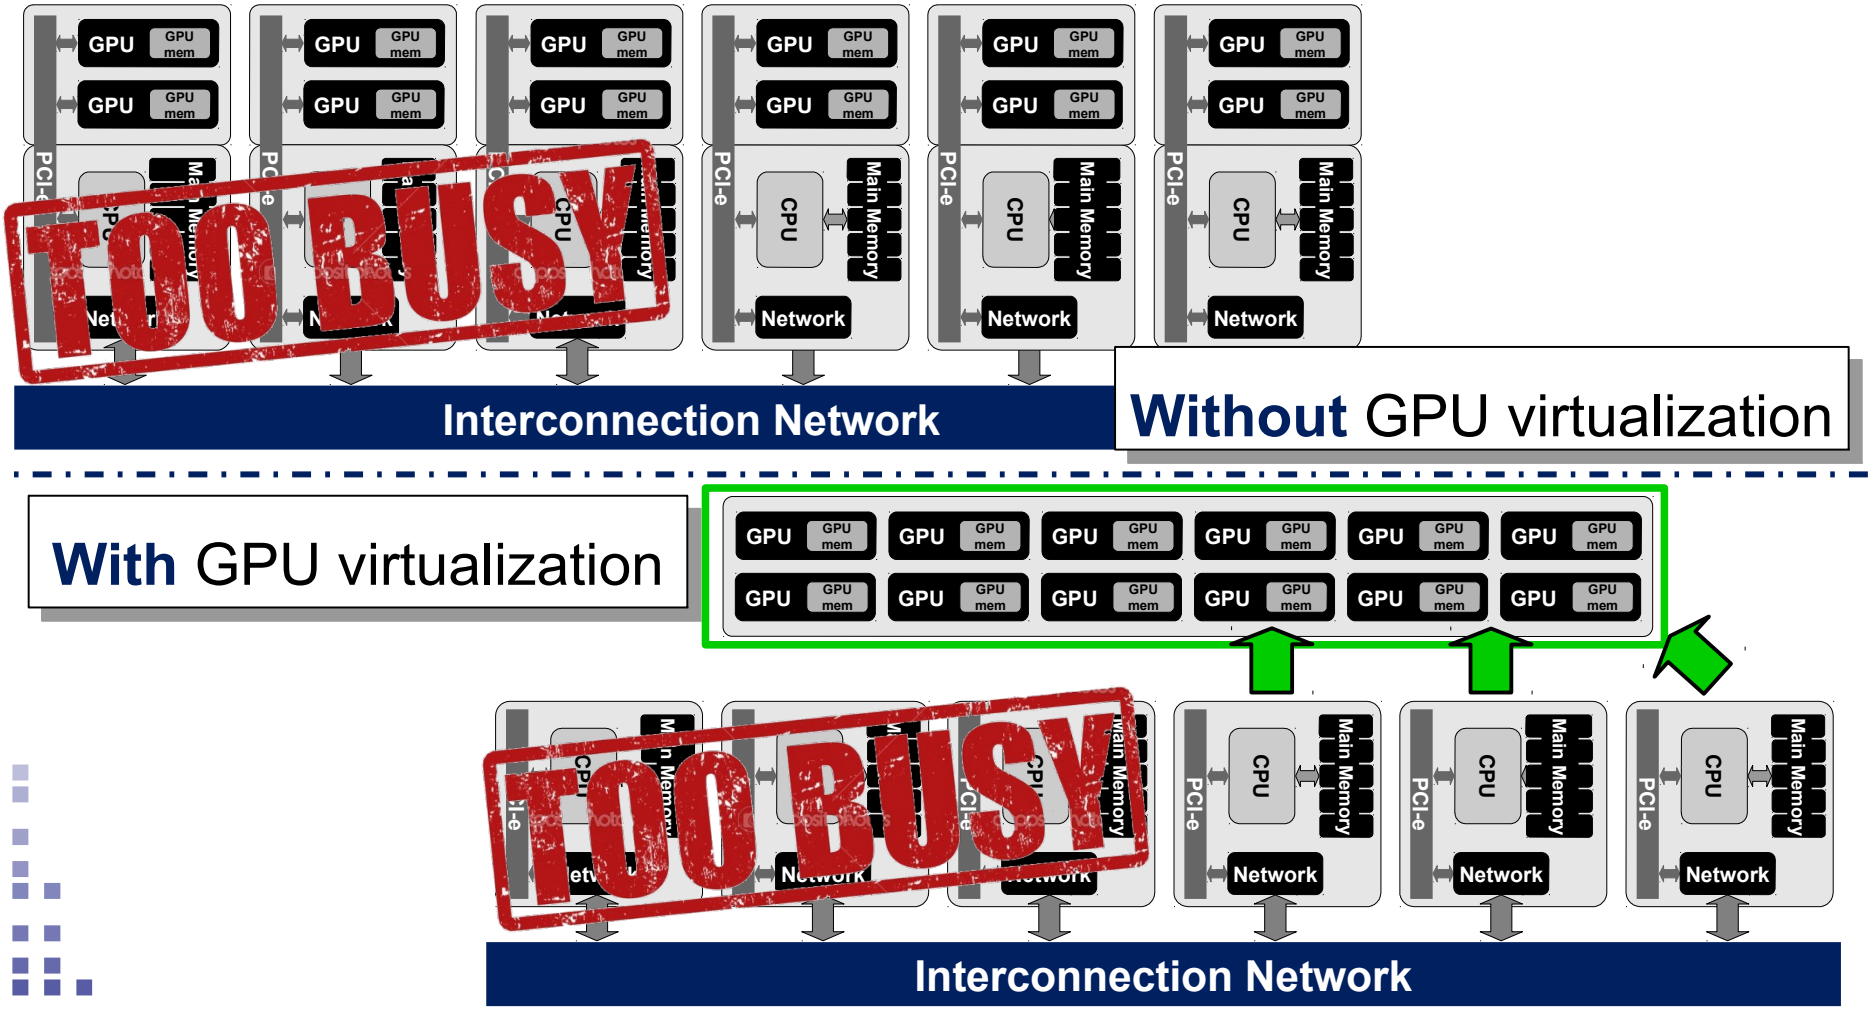

Benefits of Using Remote GPU Virtualization

2nd → Throughput is increased & energy consumption is reduced

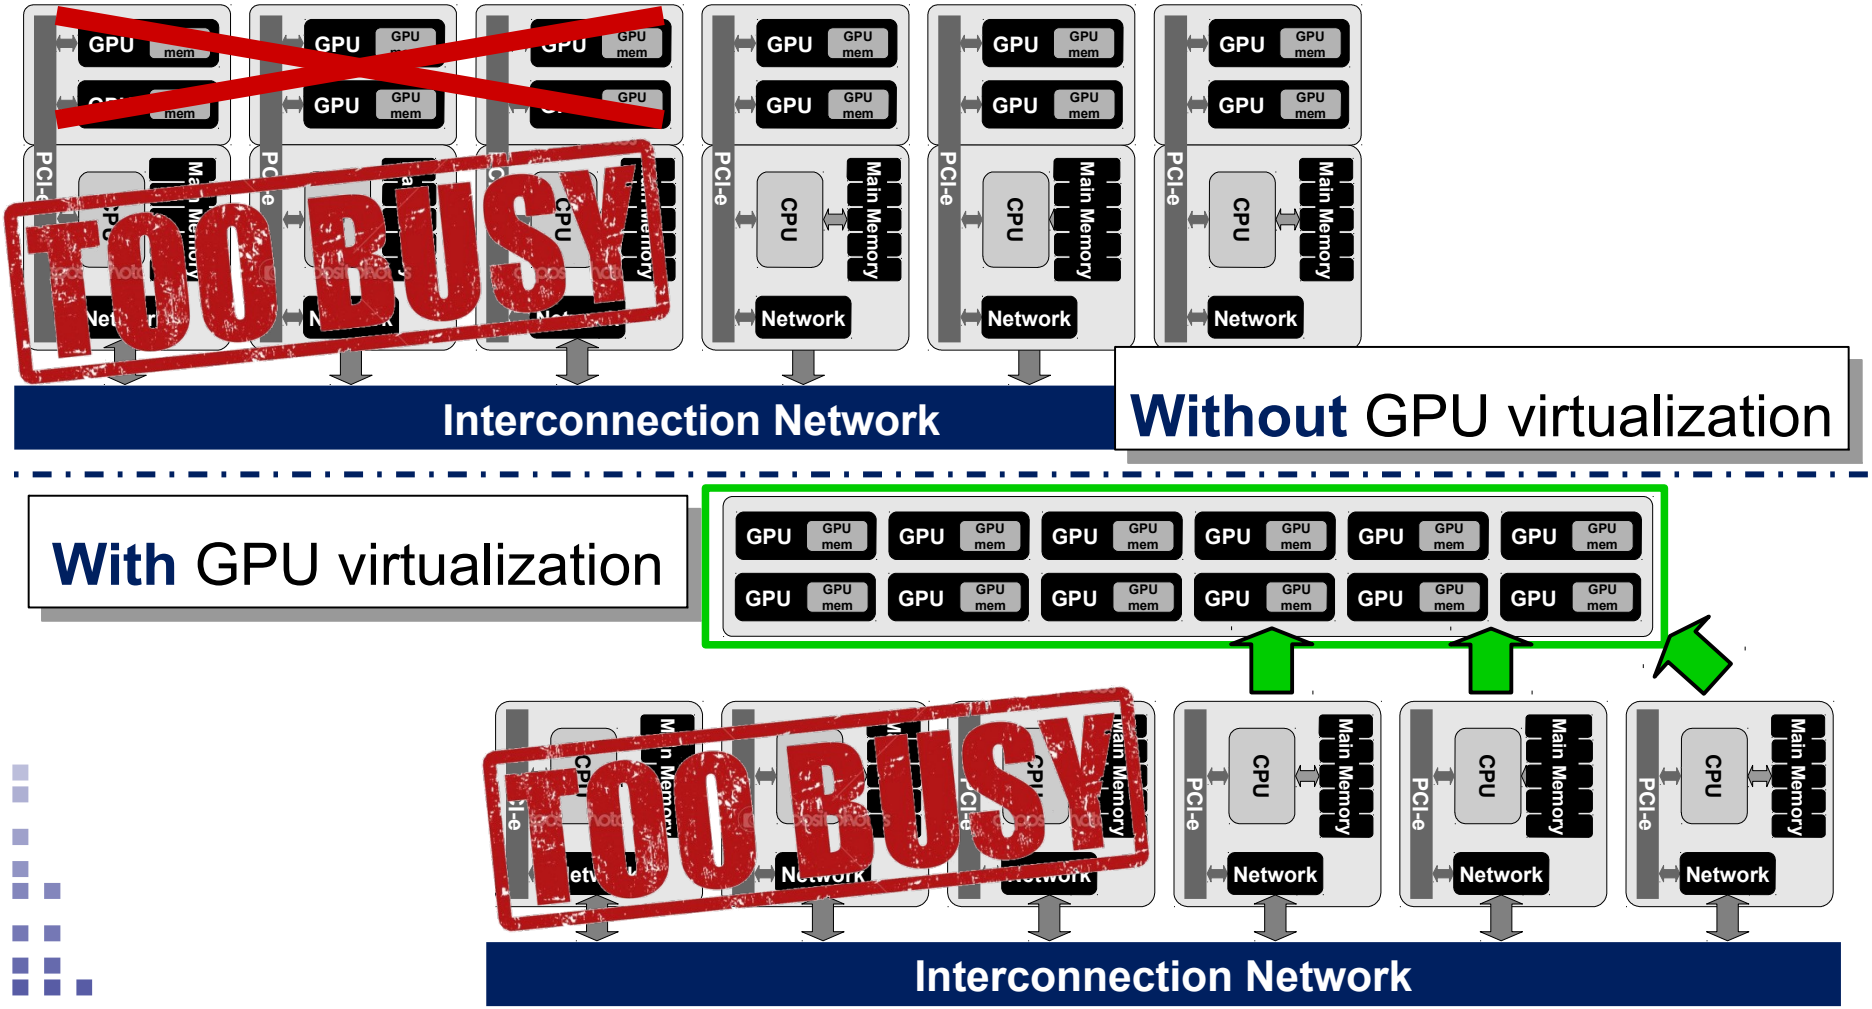

Benefits of Using Remote GPU Virtualization

2nd → Throughput is increased & energy consumption is reduced

Job scheduler uses a per-GPU granularity

Without GPU virtualization

With GPU virtualization

Job scheduler can use a more fine granularity

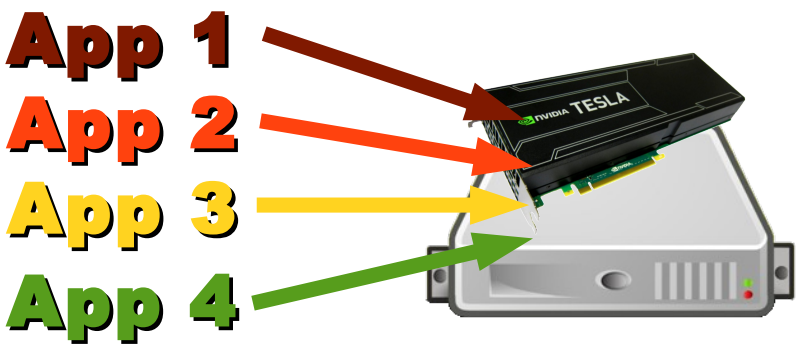

Benefits of Using Remote GPU Virtualization

2nd → Throughput is increased & energy consumption is reduced

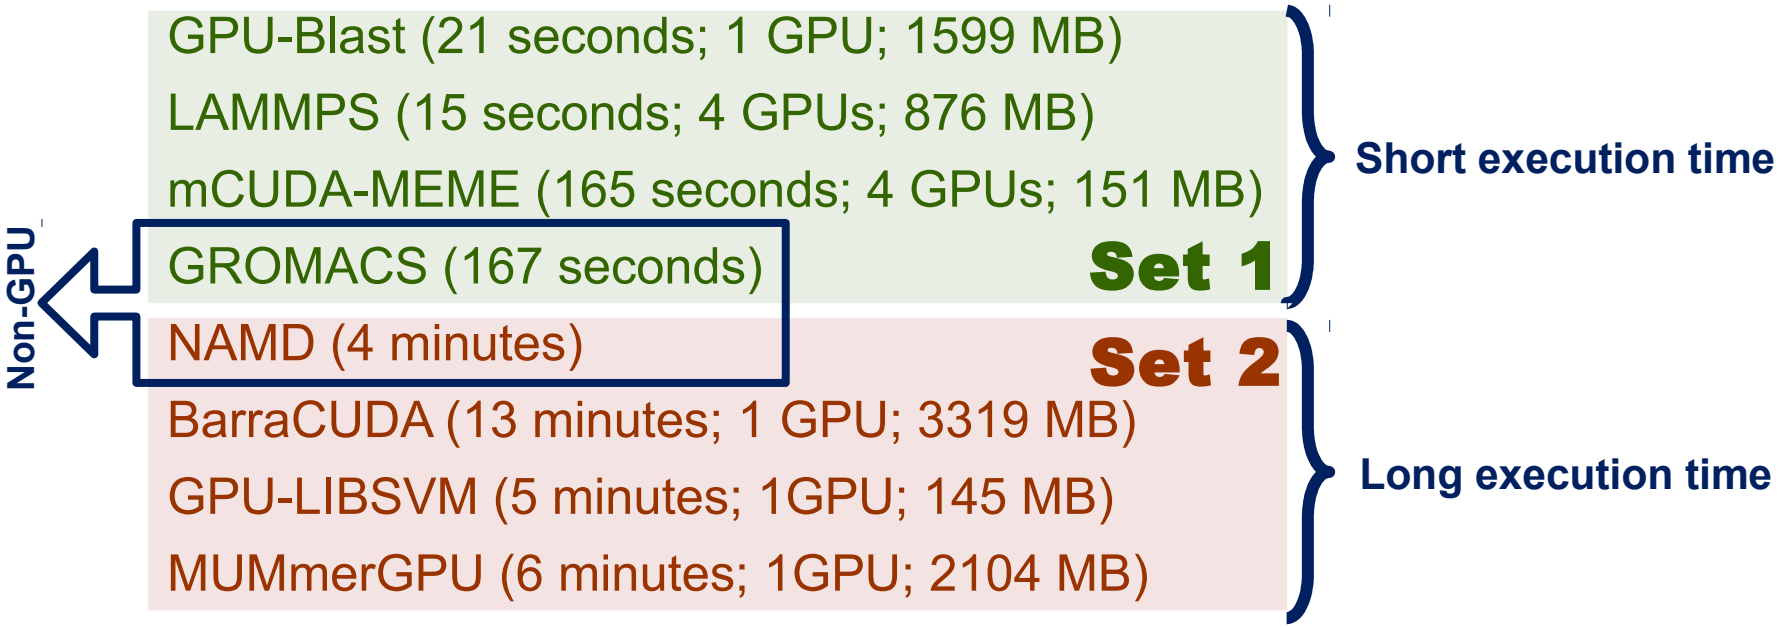

- Three workloads composed of 400 instances:
  - Set 1
  - Set 2
  - Set 1 + Set 2

Benefits of Using Remote GPU Virtualization

2nd → Throughput is increased & energy consumption is reduced

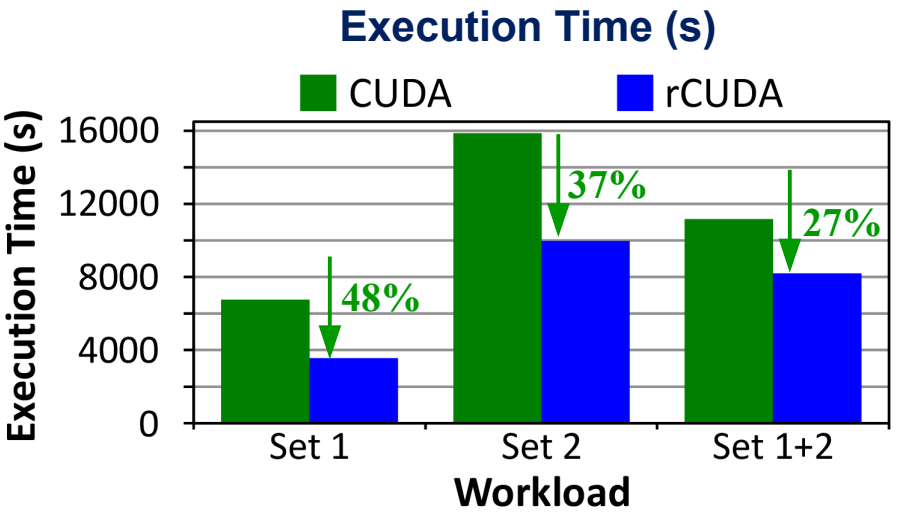

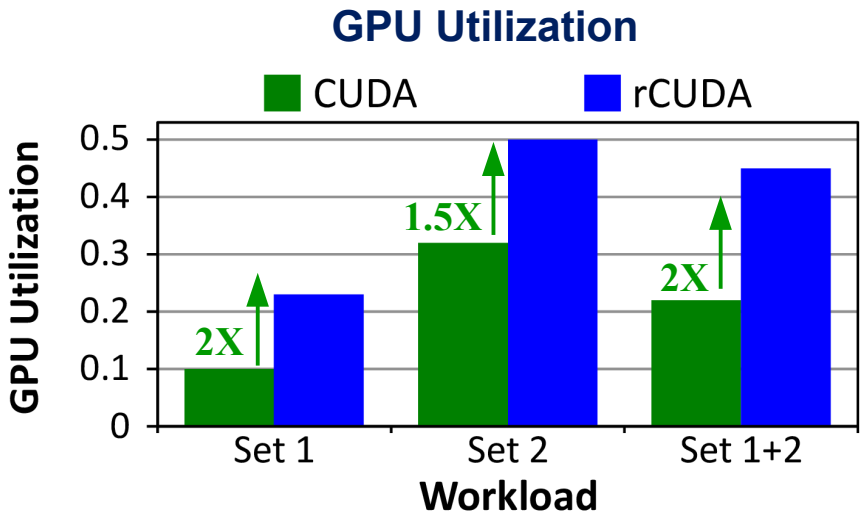

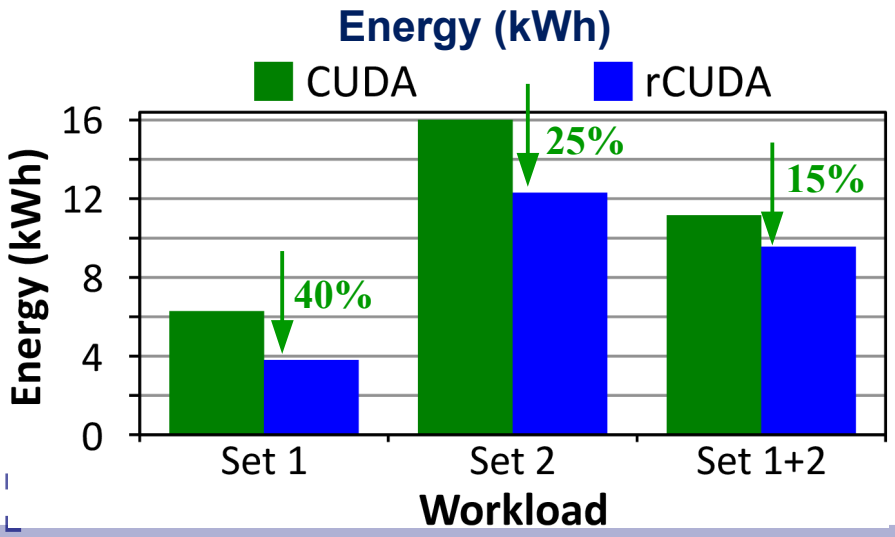

16-node cluster with

Supermicro 1027GR-TRF servers with two Intel Xeon E5-2620 v2 processors

FDR InfiniBand Interconnection Netw. and one K20 GPU per node

Benefits of Using Remote GPU Virtualization

3rd → Cheaper cluster upgrade

- A cluster without GPUs may be easily upgraded to use GPUs with rCUDA

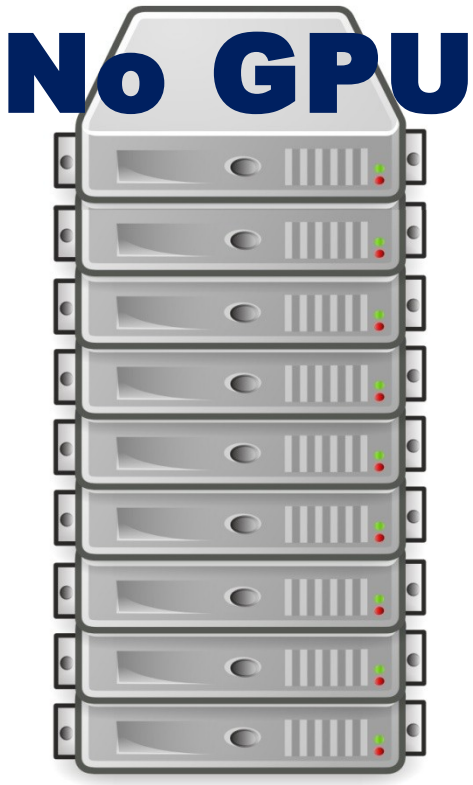

Benefits of Using Remote GPU Virtualization

3rd → Cheaper cluster upgrade

- A cluster without GPUs may be easily upgraded to use GPUs with rCUDA

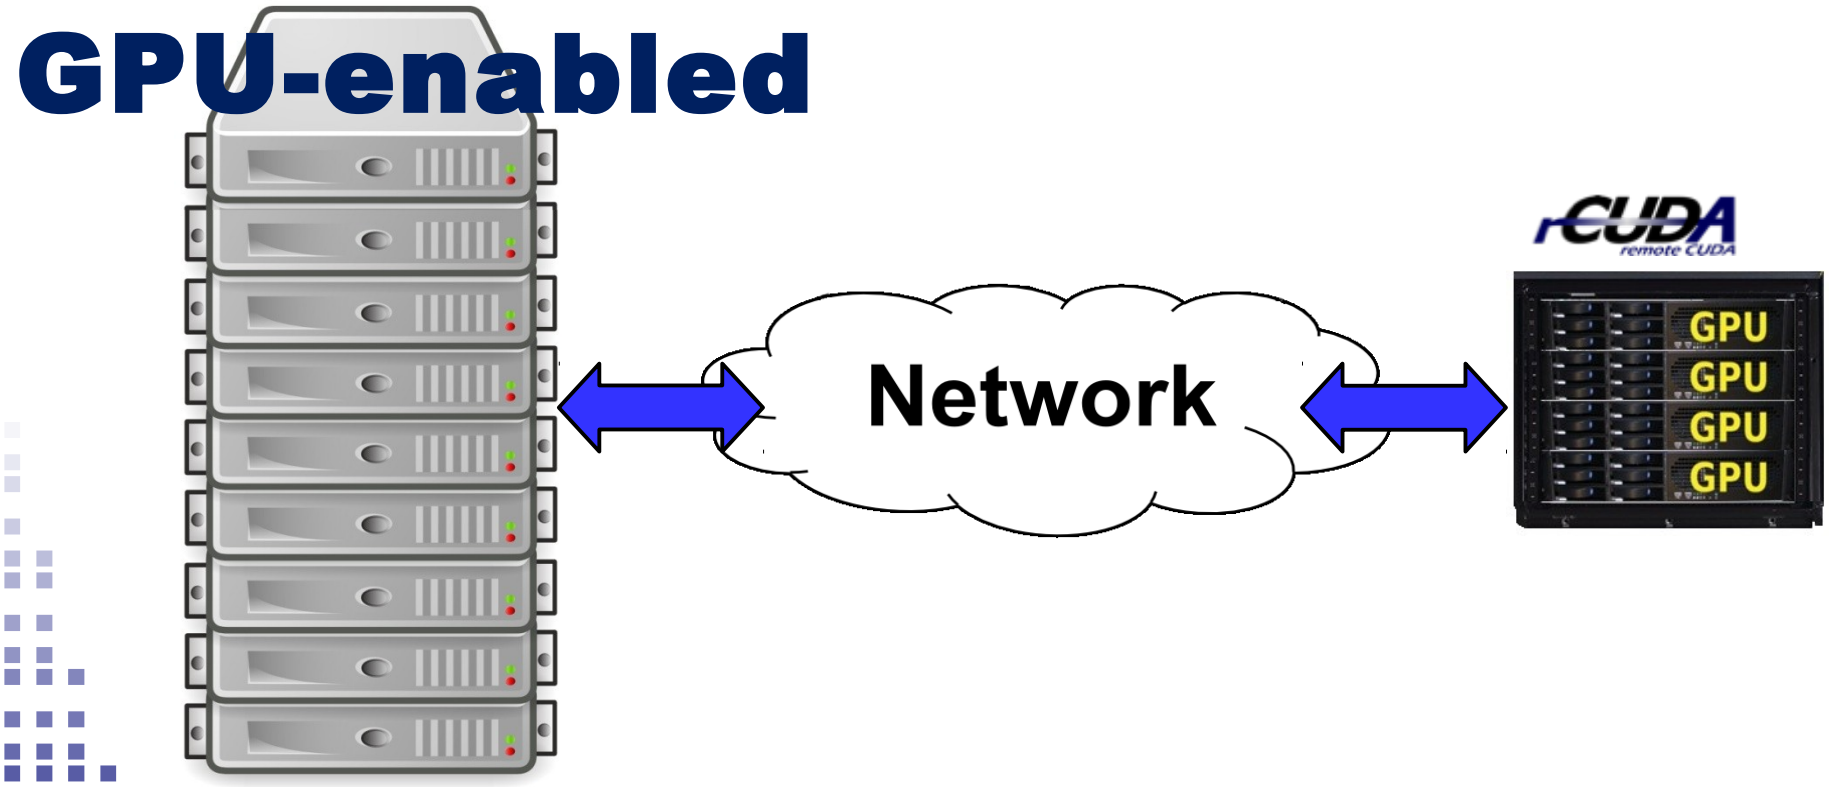

Benefits of Using Remote GPU Virtualization

3rd → Cheaper cluster upgrade

Application	Workload 1	Workload 2
GPU-BLAST	41	48
GROMACS	39	46
NAMD	20	10
BarraCUDA	20	10
GPU-LIBSVM	39	46
MUMmer-GPU	20	10
LAMMPS short	20	10
mCUDA-MEME short	40	40
LAMMPS long 2p	40	47
LAMMPS long 4p	41	47
mCUDA-MEME long 2p	40	46
mCUDA-MEME long 4p	40	40
Total Instances	400	400

Benefits of Using Remote GPU Virtualization

3rd → Cheaper cluster upgrade

Application	Workload 1	Workload 2
GPU-BLAST	41	48
GROMACS	39	46
NAMD	20	10
BarraCUDA	20	10
GPU-LIBSVM	39	46
MUMmer-GPU	20	10
LAMMPS short	20	10
mCUDA-MEME short	40	40
LAMMPS long 2p	40	47
LAMMPS long 4p	41	47
mCUDA-MEME long 2p	40	46
mCUDA-MEME long 4p	40	40
Total Instances	400	400

Benefits of Using Remote GPU Virtualization

3rd → Cheaper cluster upgrade

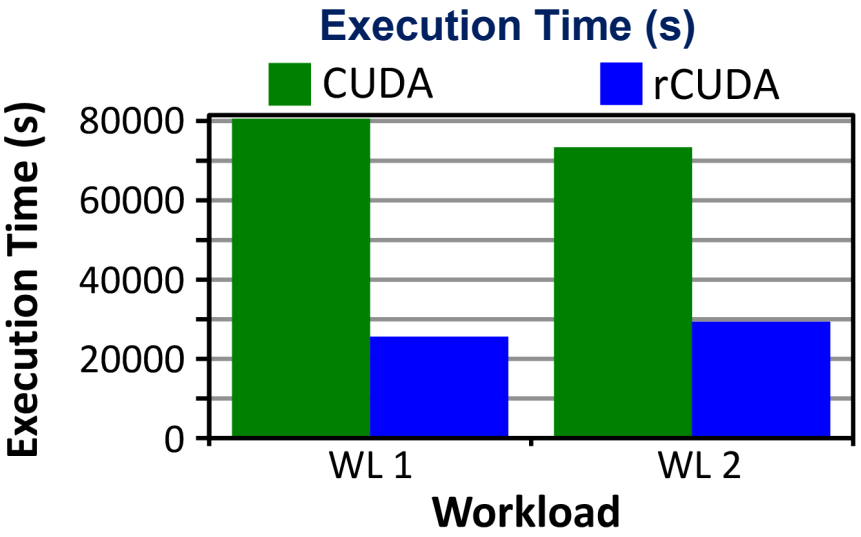

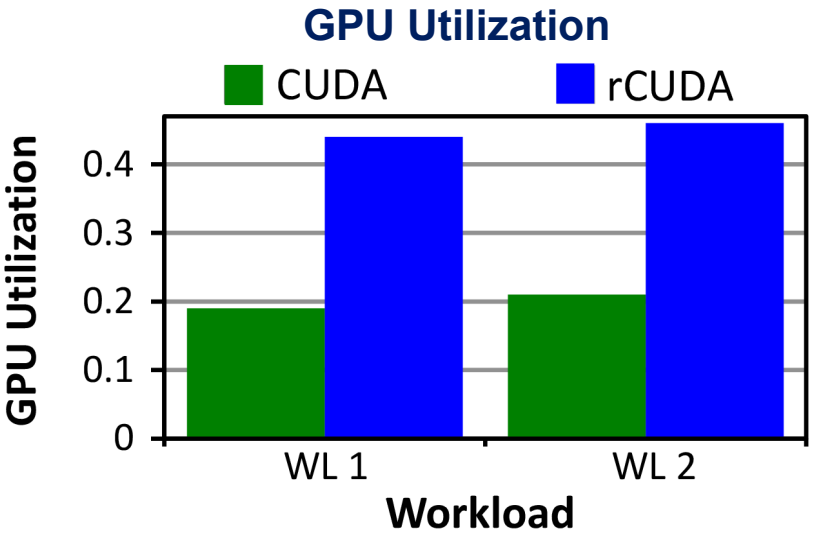

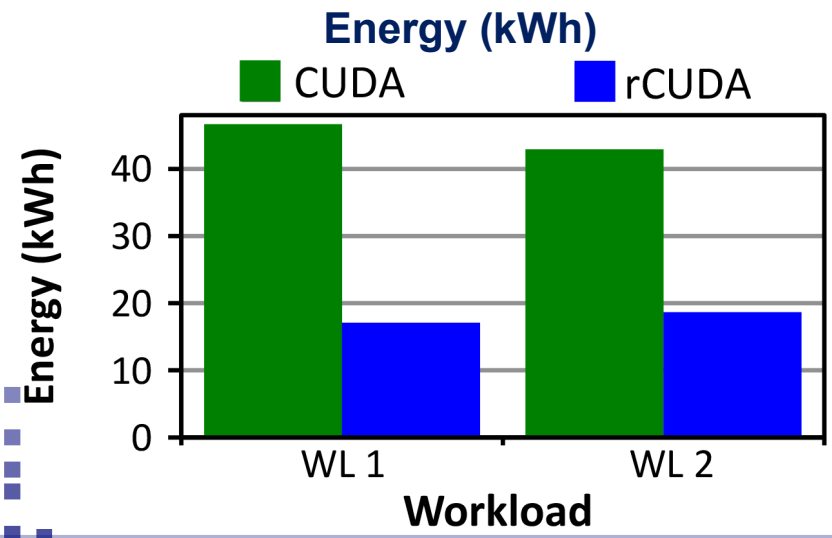

16-node cluster with 15 Supermicro 1027GR-TRF servers with two Intel Xeon E5-2620 v2 processors FDR InfiniBand Interconnection Netw. and one SYS7047GR-TRF server, populated with four NVIDIA Tesla K20

Benefits of Using Remote GPU Virtualization

4th → Virtual machines can easily share GPUs

Without GPU virtualization

- The GPU is assigned by using PCI passthrough **exclusively to a single virtual machine**
- Concurrent usage of the GPU is not possible

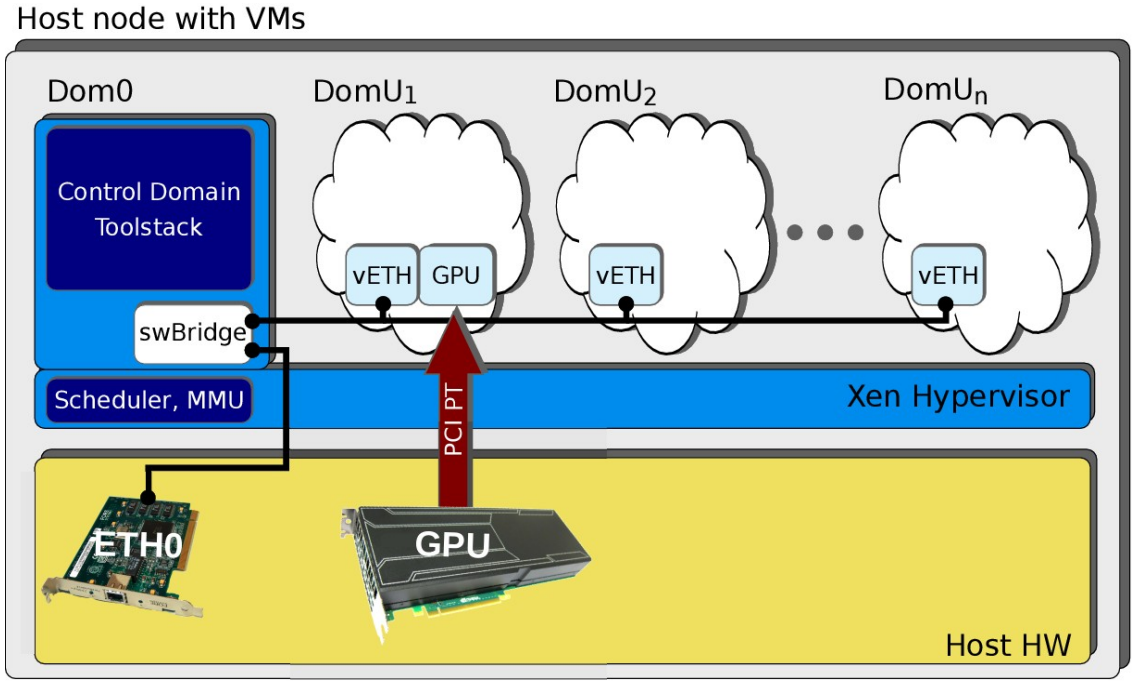

Benefits of Using Remote GPU Virtualization

4th → Virtual machines can easily share GPUs

Without GPU virtualization

- The GPU is assigned by using PCI passthrough **exclusively to a single virtual machine**
- Concurrent usage of the GPU is not possible

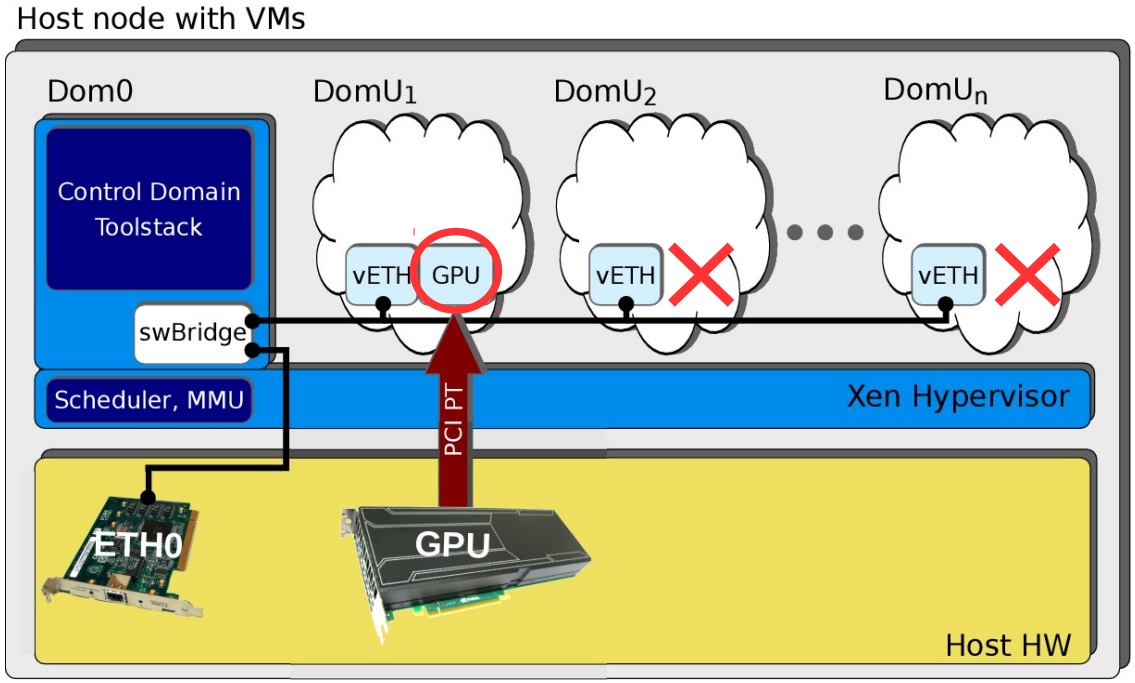

Benefits of Using Remote GPU Virtualization

4th → Virtual machines can easily share GPUs

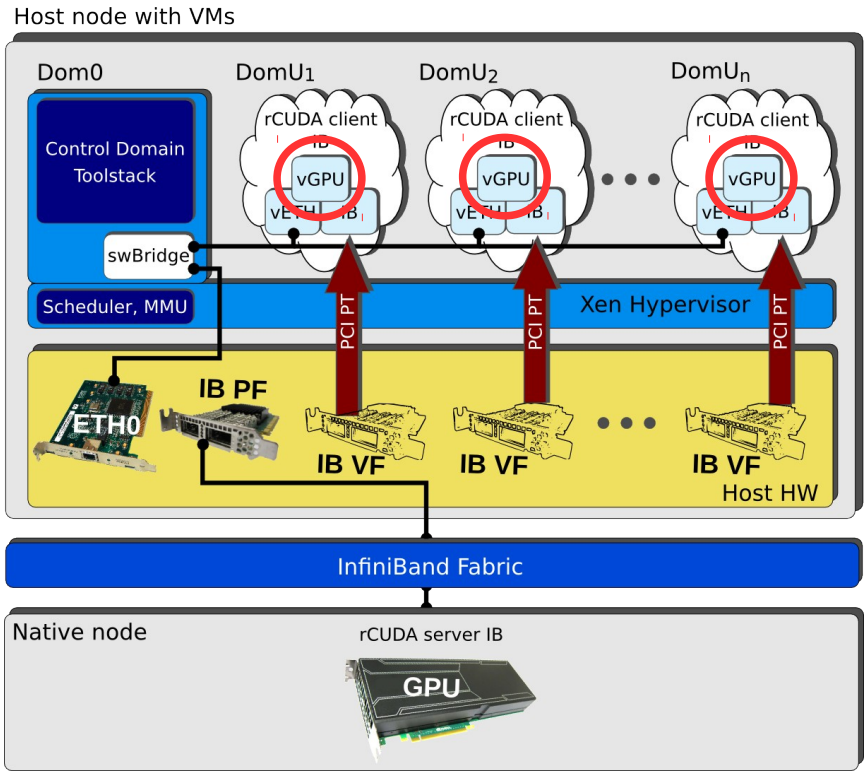

With GPU virtualization

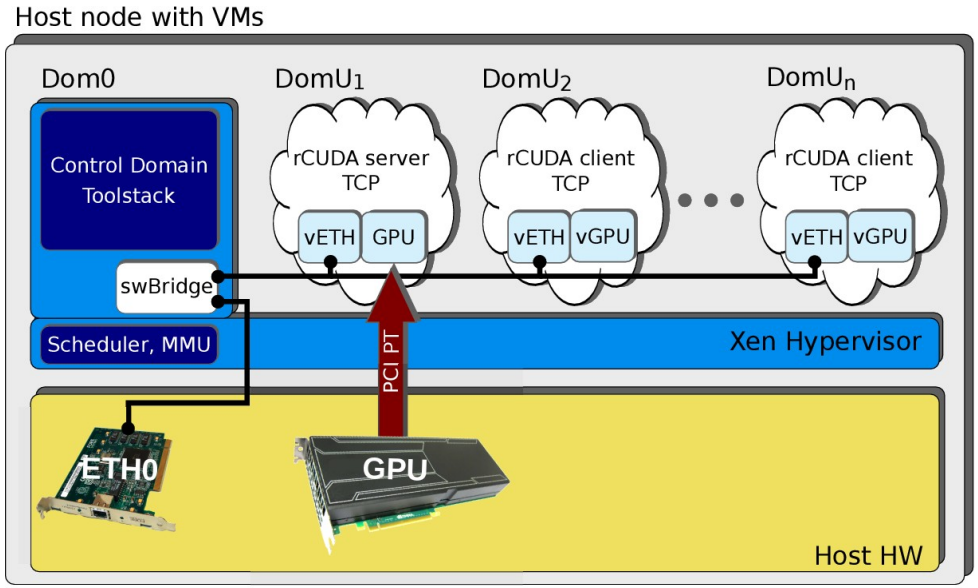

Benefits of Using Remote GPU Virtualization

4th → Virtual machines can easily share GPUs

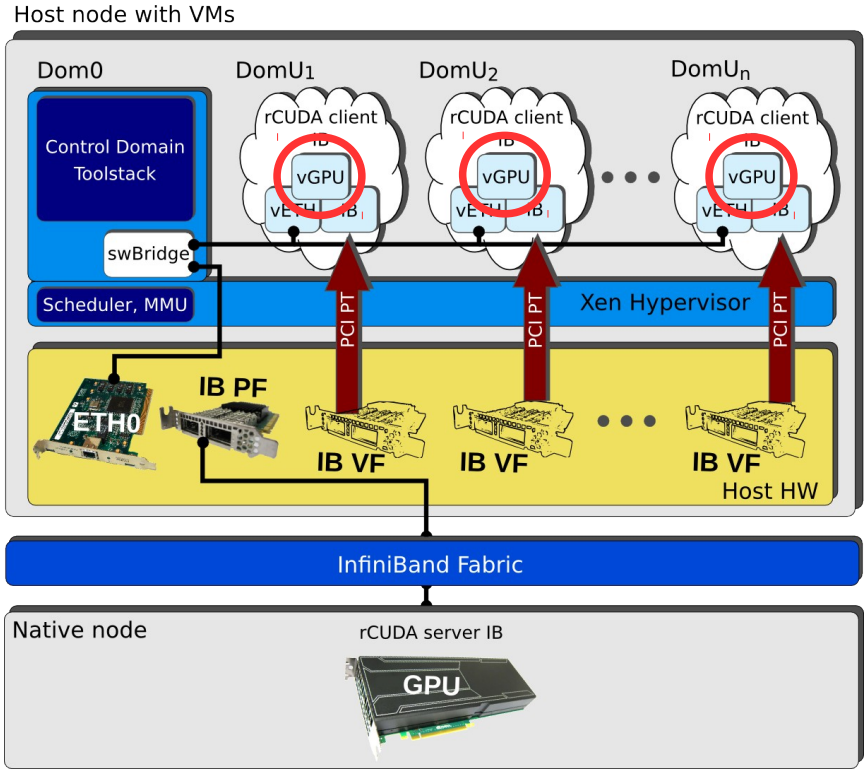

With GPU virtualization

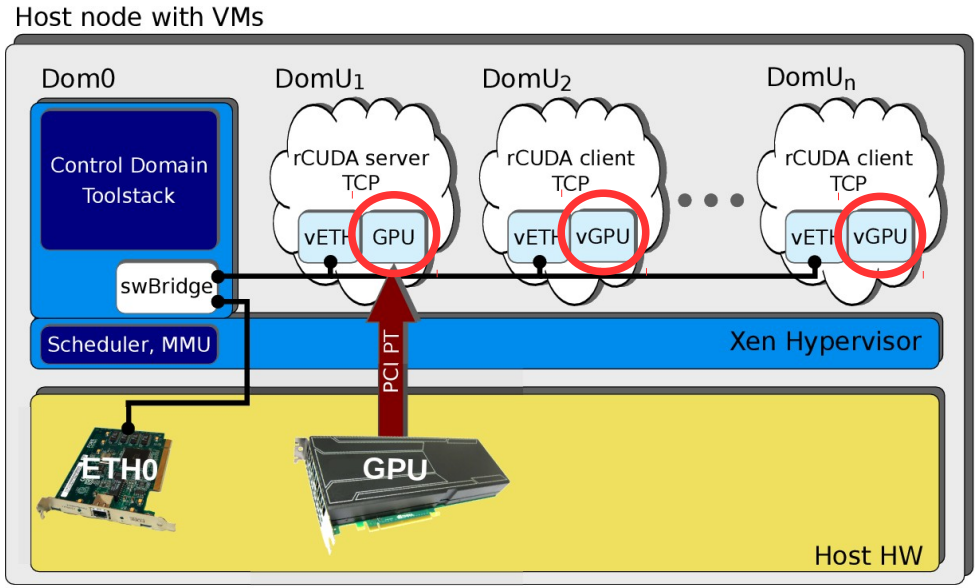

Benefits of Using Remote GPU Virtualization

4th → Virtual machines can easily share GPUs

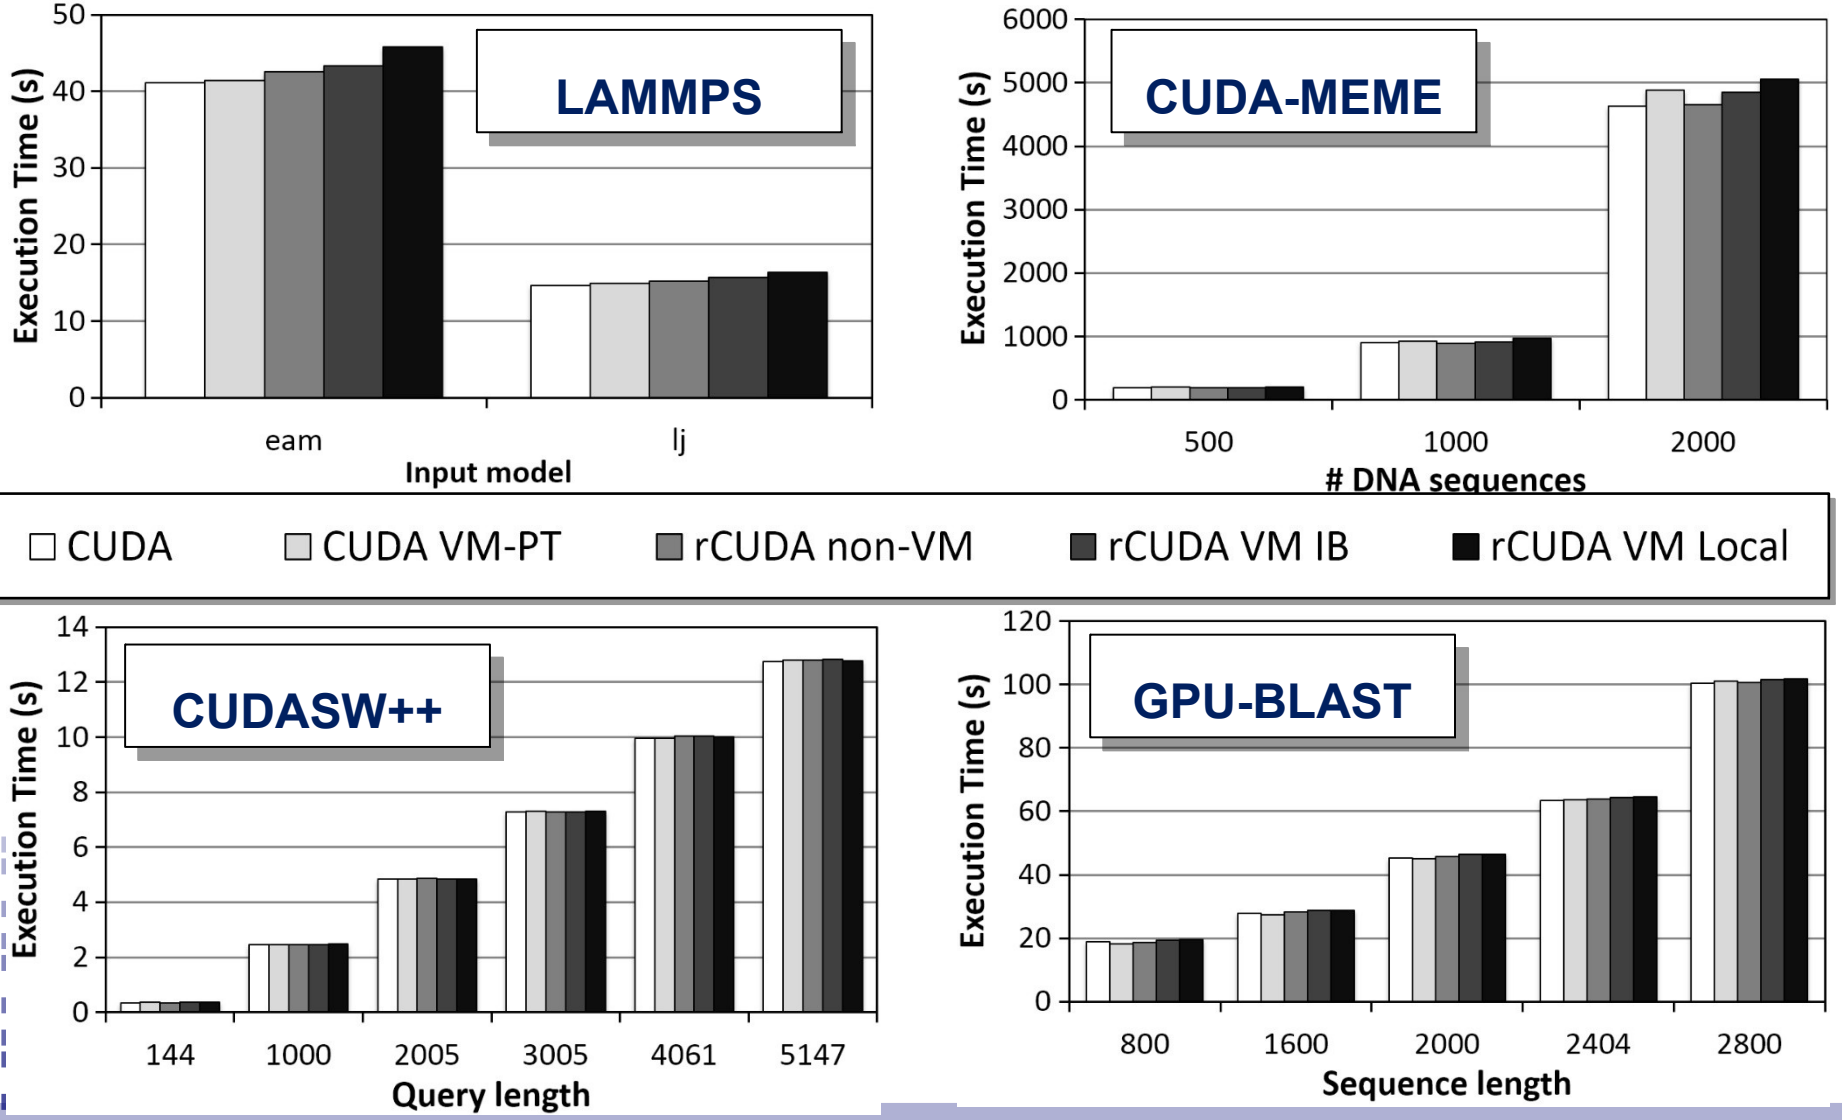

Main Conclusions

Main Conclusions

Remote gpu virtualization provides several benefits

- ✔ Similar performance with smaller investment
- ✔ Increased cluster throughput
- ✔ Many GPUs for a single application
- ✔ Concurrent GPU access to virtual machines
- ✔ Easier and cheaper cluster upgrade
- ✔ Reduced energy consumption
- ✔ Increased GPU utilization

Get a free copy of rCUDA at

http://www.rcuda.net

More than 650 requests world wide

@rcuda_

rCUDA is a development by Technical University of Valencia

Questions?