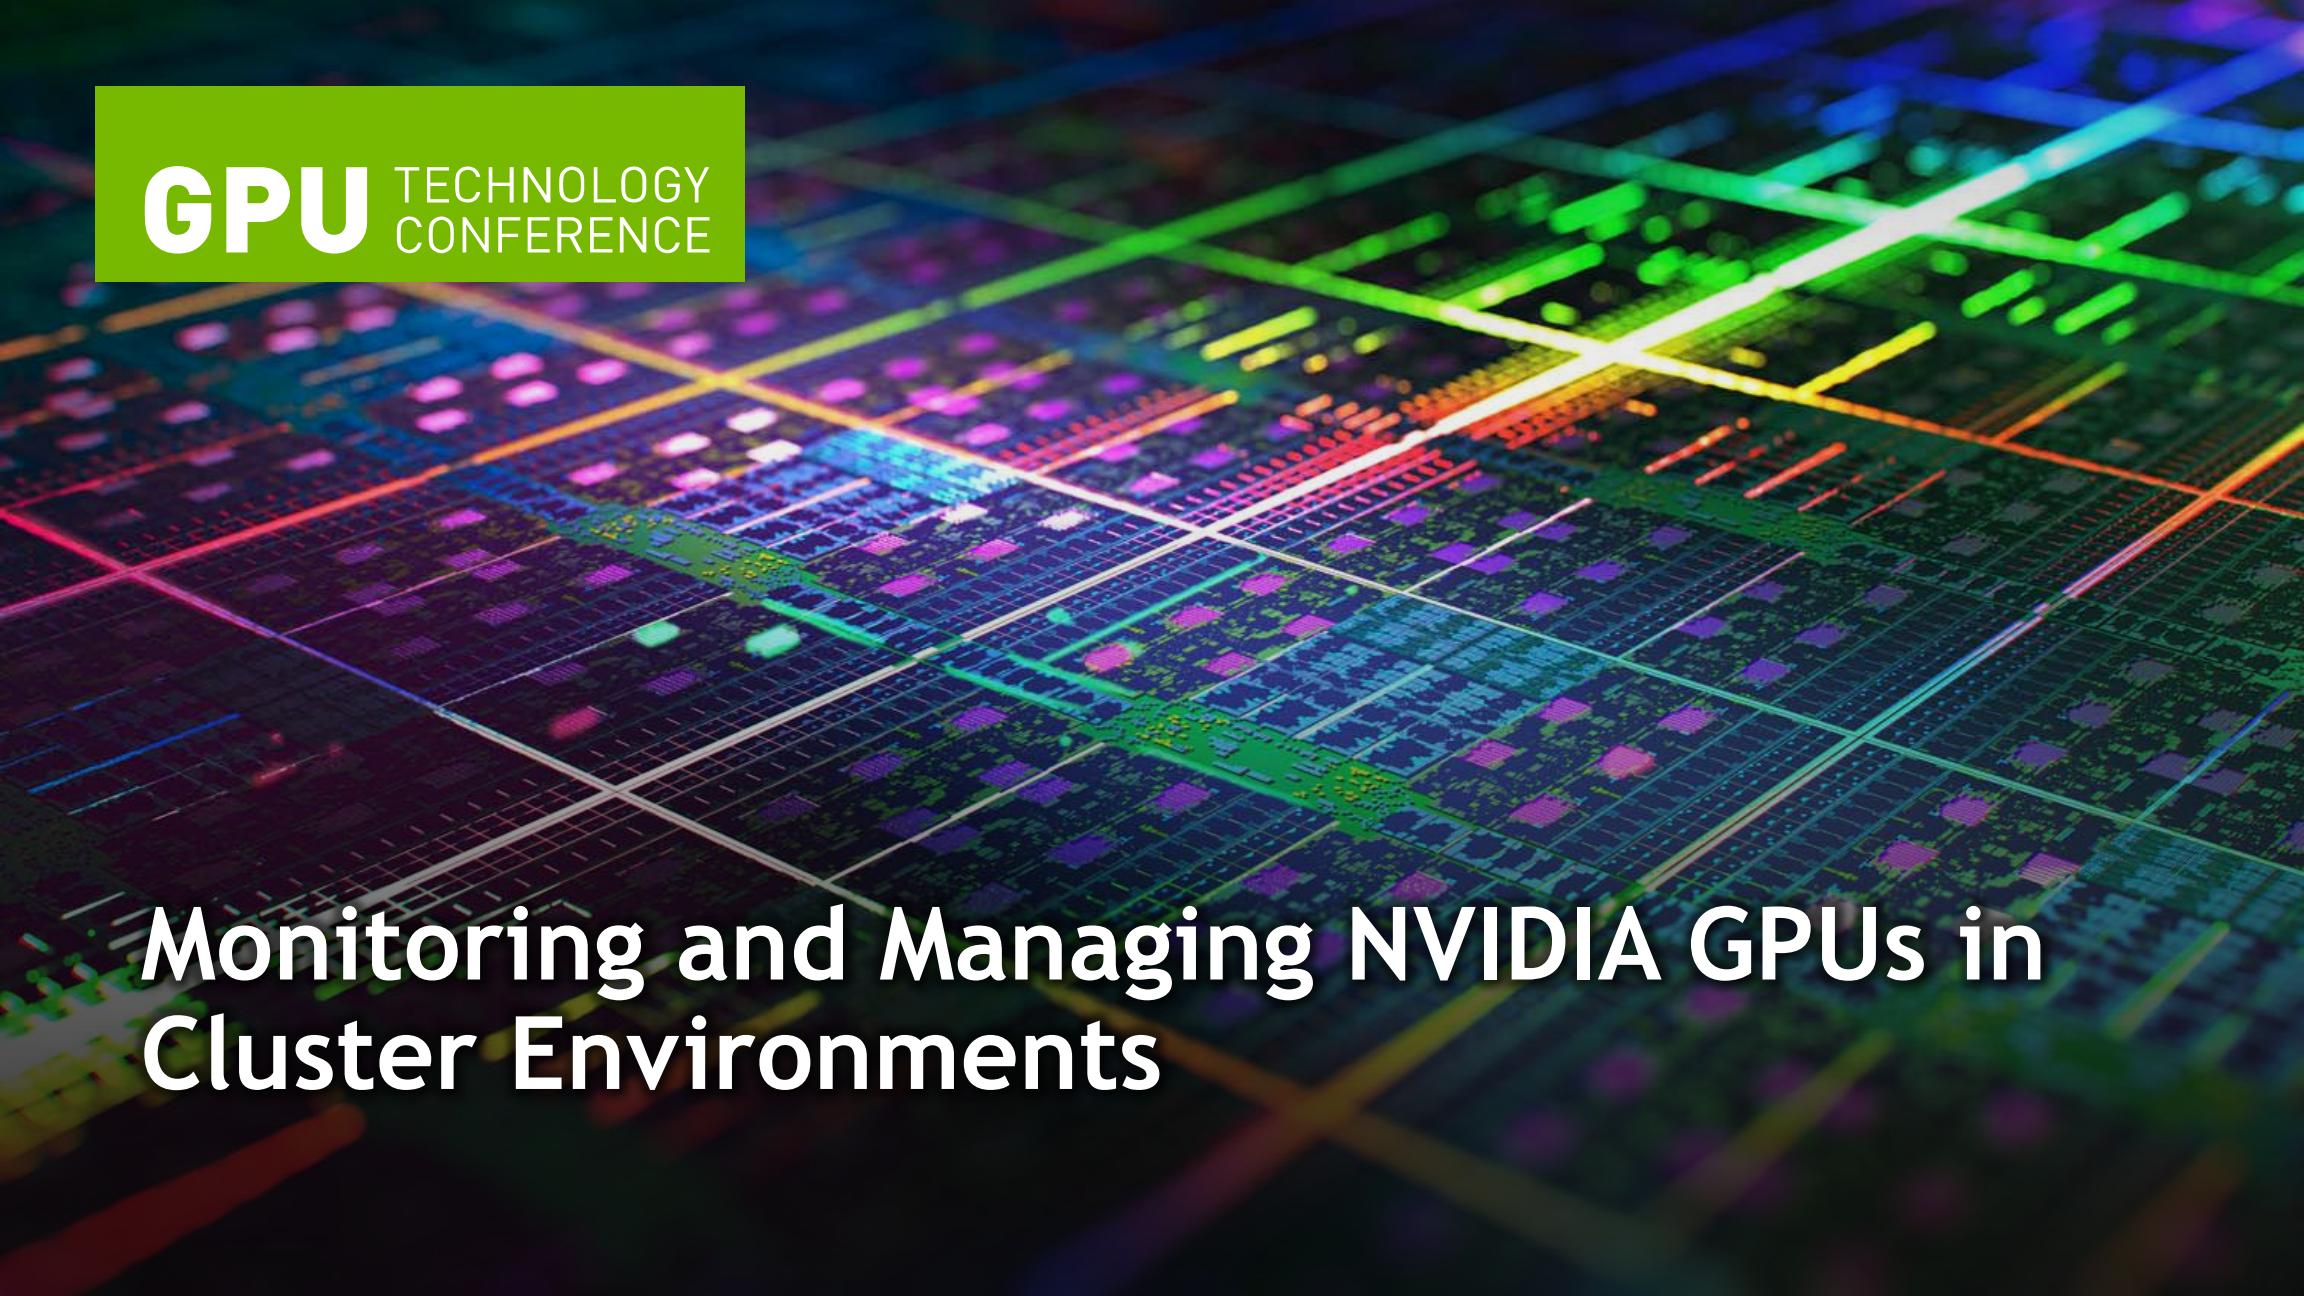
GPU TECHNOLOGY CONFERENCE
Monitoring and Managing NVIDIA GPUs in Cluster Environments

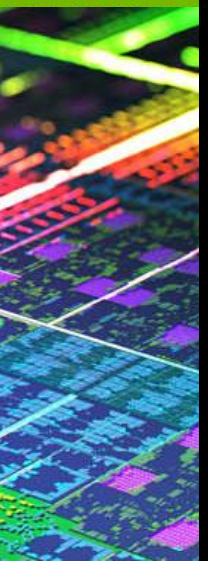

Introductions

- Przemyslaw Zych
  - Senior CUDA Tools Software Engineer at NVIDIA
  - Tesla Software Group

  - Feel free to contact my team at cudatools@nvidia.com

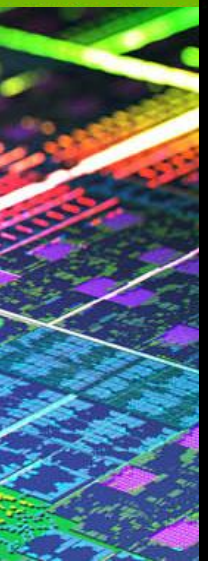

Agenda

- NVIDIA Monitoring & Management API
- Nvidia-healthmon
- 3rd Party Software

NVIDIA-SMI & NVML

- NVML
  - C library
  - SDK Part of TDK
  - Perl & Python Bindings
- NVIDIA-SMI
  - Command line application

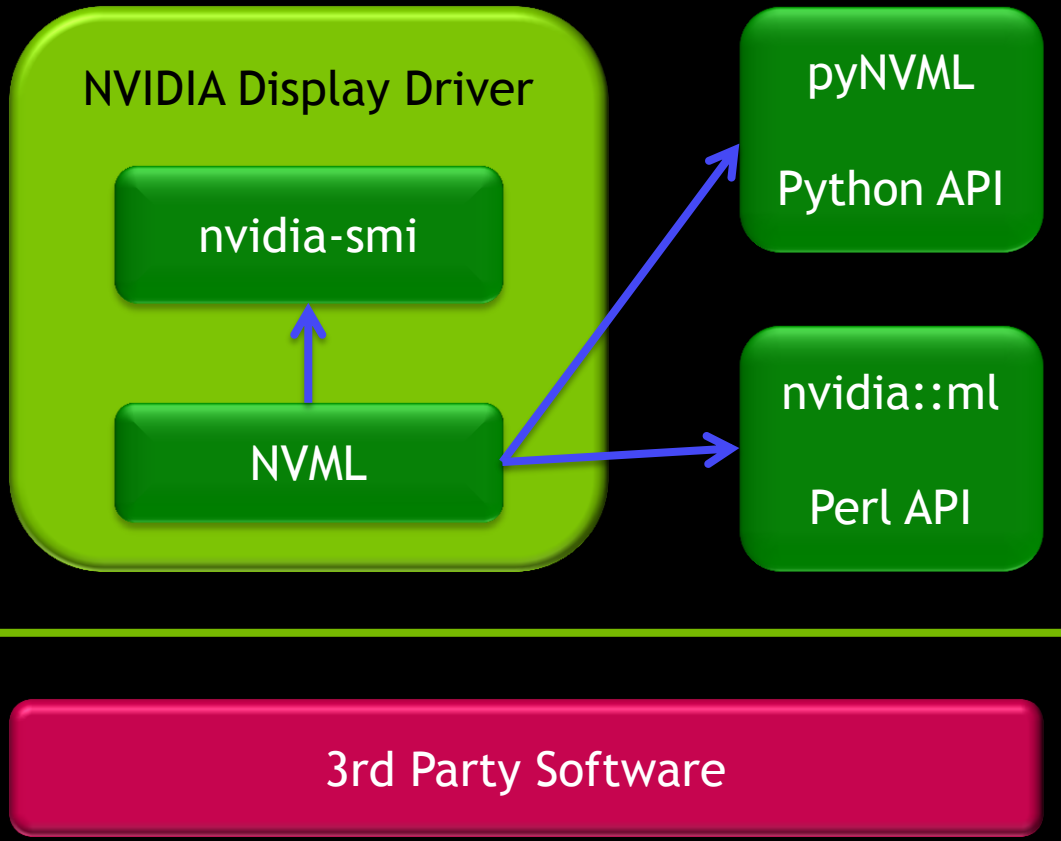

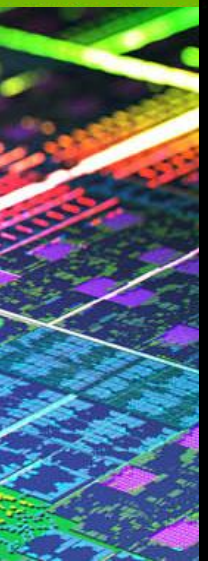

NVIDIA-SMI & NVML

- Supported OS platforms:
  - Linux: 32-bit and 64-bit
  - Windows: Windows Server 2008 R2 64bit, Windows 7 64bit
- Full support of Tesla and Quadro GPUs
- Limited functionality of GeForce
- Follows CUDA toolkit release cycle

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0

+-----------------------------------------------------+
| NVIDIA-SMI 4.319.12   Driver Version: 319.12        |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id               | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap| Memory-Usage         | GPU-Util  Compute M. |
|===============================+======================+======================|
|   1  Tesla K20c          On   | 0000:06:00.0         |                    0 |
| 53%   63C    P0   167W / 225W |       477MB / 4799MB |    100%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                               GPU Memory |
|  GPU       PID  Process name                                     Usage      |
|=============================================================================|
|    1     22302  ./sgemm_cublas                                      461MB  |
+-----------------------------------------------------------------------------+
```

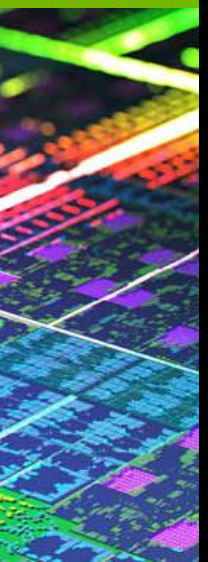

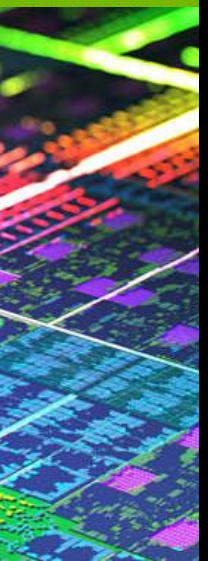

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0 --query

    ==============NVSMI LOG==============

    Timestamp                           : Fri Feb  1 15:38:55 2013
    Driver Version                      : 319.12

    Attached GPUs                       : 2
    GPU 0000:06:00.0
        Product Name                    : Tesla K20c
        Display Mode                    : Disabled
        Display Active                  : Disabled
        Persistence Mode                : Enabled
        Driver Model
            Current                     : N/A
            Pending                     : N/A
        Serial Number                   : 12345678
        GPU UUID                        : GPU-c76cf543-ec97-c570-6e76-a4f22059cf20
        VBIOS Version                   : 80.10.23.00.00
```

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0 --query

    ==============NVSMI LOG==============

    Timestamp                           : Fri Feb  1 15:38:55 2013
    Driver Version                      : 319.12

    Attached GPUs                       : 2
    GPU 0000:06:00.0
        Product Name                    : Tesla K20c
        Display Mode                    : Disabled
        Display Active                  : Disabled
        Persistence Mode                : Enabled
        Driver Model
            Current                     : N/A
            Pending                     : N/A
        Serial Number                   : 12345678
        GPU UUID                        : GPU-c76cf543-ec97-c570-6e76-a4f22059cf20
        VBIOS Version                   : 80.10.23.00.00
```

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0 --query

    PCI
            Bus                         : 0x06
            Device                      : 0x00
            Domain                      : 0x0000
            Device Id                   : 0x102210DE
            Bus Id                      : 0000:06:00.0
            Sub System Id               : 0x098210DE
            GPU Link Info
                PCIe Generation
                    Max                 : 2
                    Current             : 2
                Link Width
                    Max                 : 16x
                    Current             : 16x
```

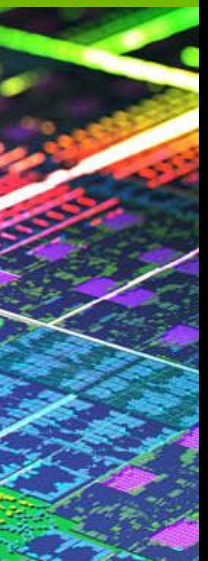

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0 --query

    PCI
            Bus                         : 0x
```

Gen 1 is 250MB/s per lane
Gen 2 is 500MB/s per lane
Gen 3 is 1000MB/s per lane

500MB/s/lane * 16 lanes = 8 GB/s of raw theoretical bandwidth in each direction

```
$ lspci –t –v

 +-[0000:80]-+-01.0-[81-82]--+-00.0  Intel Corporation Ethernet Controller 10 Gigabit X540-AT2
 |           |               \-00.1  Intel Corporation Ethernet Controller 10 Gigabit X540-AT2
 |           +-02.0-[83-86]----00.0-[84-86]--+-08.0-[85]----00.0  nVidia
 |           |                               \-10.0-[86]----00.0  nVidia
 |           +-03.0-[87-8a]----00.0-[88-8a]--+-08.0-[89]----00.0  nVidia
 |           |                               \-10.0-[8a]----00.0  nVidia
 \-[0000:00]-+-00.0  Intel Corporation Sandy Bridge DMI2
             +-01.0-[01]--
             +-02.0-[02-0b]----00.0-[03-0b]--+-08.0-[04-07]----00.0-[05-07]--+-08.0-[06]----00.0  nVidia
             |                               |                               \-10.0-[07]----00.0  nVidia
             |                               \-10.0-[08-0b]----00.0-[09-0b]--+-08.0-[0a]----00.0  nVidia
             |                                                               \-10.0-[0b]----00.0  nVidia
```

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0 --query

    Compute Mode                        : Default
        Utilization
            Gpu                         : 99 %
            Memory                      : 41 %
                                        : Enabled
                                        : Enabled
        Volatile
            Single Bit
                Device Memory           : N/A
                File                    : N/A
                                        : N/A
                                        : N/A
                emory                   : N/A
                                        : N/A
            Double Bit
```

Default

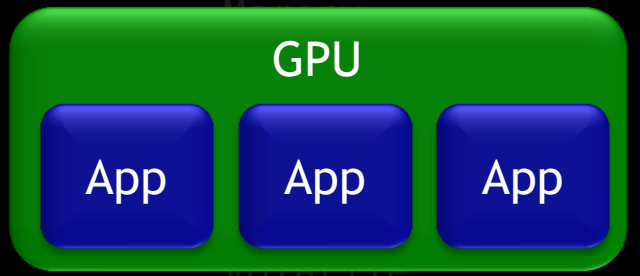

Exclusive Process

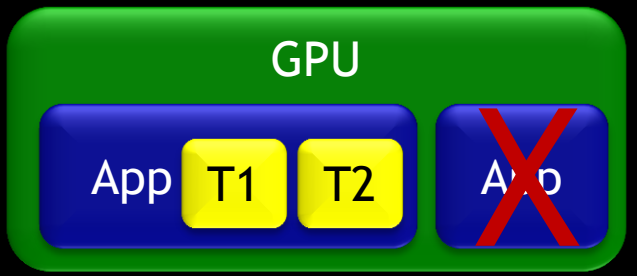

Exclusive Thread

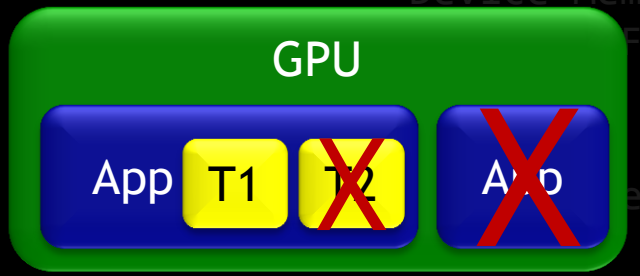

Prohibited

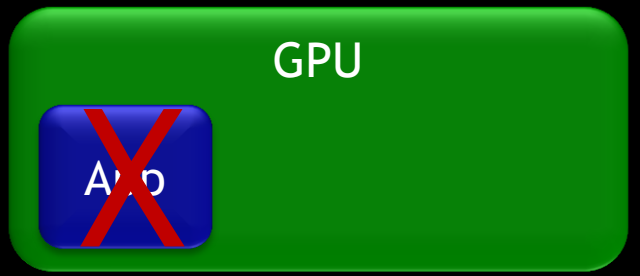

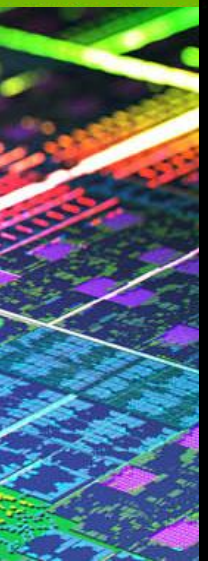

NVIDIA-SMI

```
$ nvidia-smi -i 0000:06:00.0 --query –-display=ECC

        Ecc Mode
            Current                     : Enabled
            Pending                     : Enabled
        ECC Errors
            Volatile                                    Aggregate
                Single Bit                                  Single Bit
                    Device Memory   : 0                         Device Memory   : 0
                    Register File   : 0                         Register File   : 0
                    L1 Cache        : 0                         L1 Cache        : 0
                    L2 Cache        : 0                         L2 Cache        : 0
                    Texture Memory  : 0                         Texture Memory  : 0
                    Total           : 0                         Total           : 0
                Double Bit                                  Double Bit
                    Device Memory   : 0                         Device Memory   : 0
                    Register File   : 0                         Register File   : 0
                    L1 Cache        : 0                         L1 Cache        : 0
                    L2 Cache        : 0                         L2 Cache        : 0
                    Texture Memory  : 0                         Texture Memory  : 0
                    Total           : 0                         Total           : 0
```

NVIDIA-SMI

```
$ nvidia-smi –i 1 --query --display=POWER,CLOCK,PERFORMANCE

Power Readings
    Power Management                : Supported
    Power Draw                      : 40.10 W
    Power Limit                     : 225.00 W
    Default Power Limit             : 225.00 W
    Min Power Limit                 : 150.00 W
    Max Power Limit                 : 225.00 W

Clocks Throttle Reasons
    Idle                            : Not Active
    Applications Clocks Setting     : Active
    SW Power Cap                    : Not Active
    HW Slowdown                     : Not Active
    Unknown                         : Not Active

Clocks
    Graphics                        : 705 MHz
    SM                              : 705 MHz
    Memory                          : 2600 MHz
Applications Clocks
    Graphics                        : 705 MHz
    Memory                          : 2600 MHz
Default Applications Clocks
    Graphics                        : 705 MHz
    Memory                          : 2600 MHz
Max Clocks
    Graphics                        : 758 MHz
    SM                              : 758 MHz
    Memory                          : 2600 MHz
```

pynvml

- Sample Python Code:

```
>>> from pynvml import *
>>> nvmlInit()
>>> print "Driver Version:", nvmlSystemGetDriverVersion()
Driver Version: 319.12
>>> deviceCount = nvmlDeviceGetCount()
>>> for i in range(deviceCount):
...     handle = nvmlDeviceGetHandleByIndex(i)
...     print "Device", i, ":", nvmlDeviceGetName(handle)
...
Device 0 : Tesla K20c
Device 1 : Tesla K20c

>>> nvmlShutdown()
```

Initialize NVML Library

Get Driver Version

Get GPU Count

Query for device handle to perform operations on a device

Get the Device Name

Shut down NVML by releasing all GPU resource

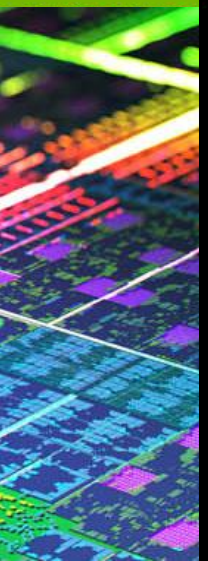

GPU Cluster Management Features

- Thermal and power monitoring and protection
- Controlling target performance
- ECC error management
- Version, Utilization and Configuration Info
- Accounting metrics (new in 5.5)
- Out of band access (improved in 5.5)

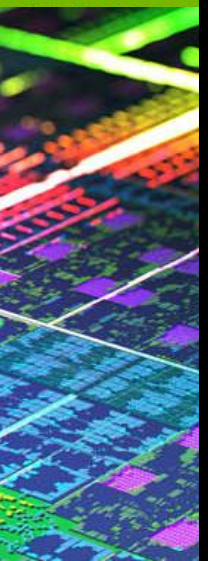

nvidia-healthmon

- Quick health check
- Not a full diagnostic
- Suggest remedies to SW and system configuration problems
- Expected to be run as an epilog/prolog script or periodically
- Distributed as part of TDK

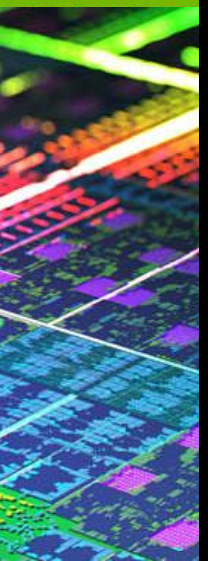

nvidia-healthmon - features

- Basic CUDA and NVML sanity check
- Diagnosis of GPU failure-to-initialize problems **(improved in 5.5)**
- Check for conflicting drivers (I.E. Nouveau)
- InfoROM validation
- Poorly seated GPU detection
- Check for disconnected power cables
- ECC error detection and reporting **(improved in 5.5)**
- Bandwidth test
- Note: Coordination with job schedulers is needed as nvidia-healthmon creates a CUDA context

nvidia-healthmon - configuration file

```
[global]
devices.tesla.count = 1
drivers.blacklist = nouveau

[Tesla C2070]
bandwidth.warn = 4000
bandwidth.min = 3000
pci.gen = 2
pci.width = 16
temperature.warn = 95
```

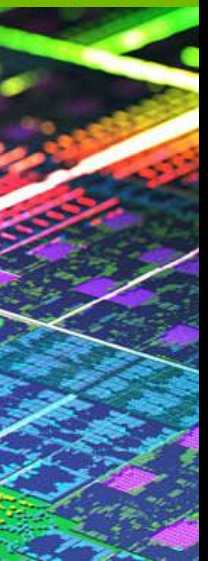

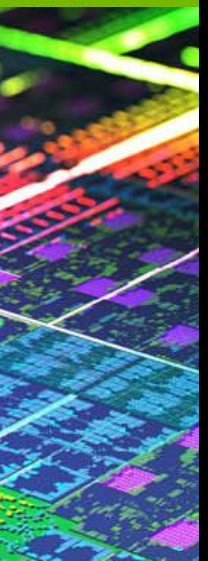

Downloads

- NVIDIA Driver (NVML, nvidia-smi)
  - http://developer.nvidia.com/cuda/cuda-toolkit
- TDK (NVML SDK + nvidia-healthmon)
  - http://developer.nvidia.com/tesla-deployment-kit
- Python NVML Bindings
  - http://pypi.python.org/pypi/nvidia-ml-py/
- Perl NVML Bindings
  - http://search.cpan.org/~nvbinding/nvidia-ml-pl/

Monitoring & Management

hp
IBM Platform Computing
Be the Master of Your Clusters
Bright Computing
StackIQ
sFlow
Adaptive COMPUTING
slurm workload manager
Ganglia
PENGUIN COMPUTING

GPU Scheduling

Bright Computing
hp
UNIVA
GRID ENGINE
IBM Platform Computing
Be the Master of Your Clusters
StackIQ
Adaptive COMPUTING
PBS Works™
slurm workload manager
PENGUIN COMPUTING

sFlow + Ganglia GPU Monitoring

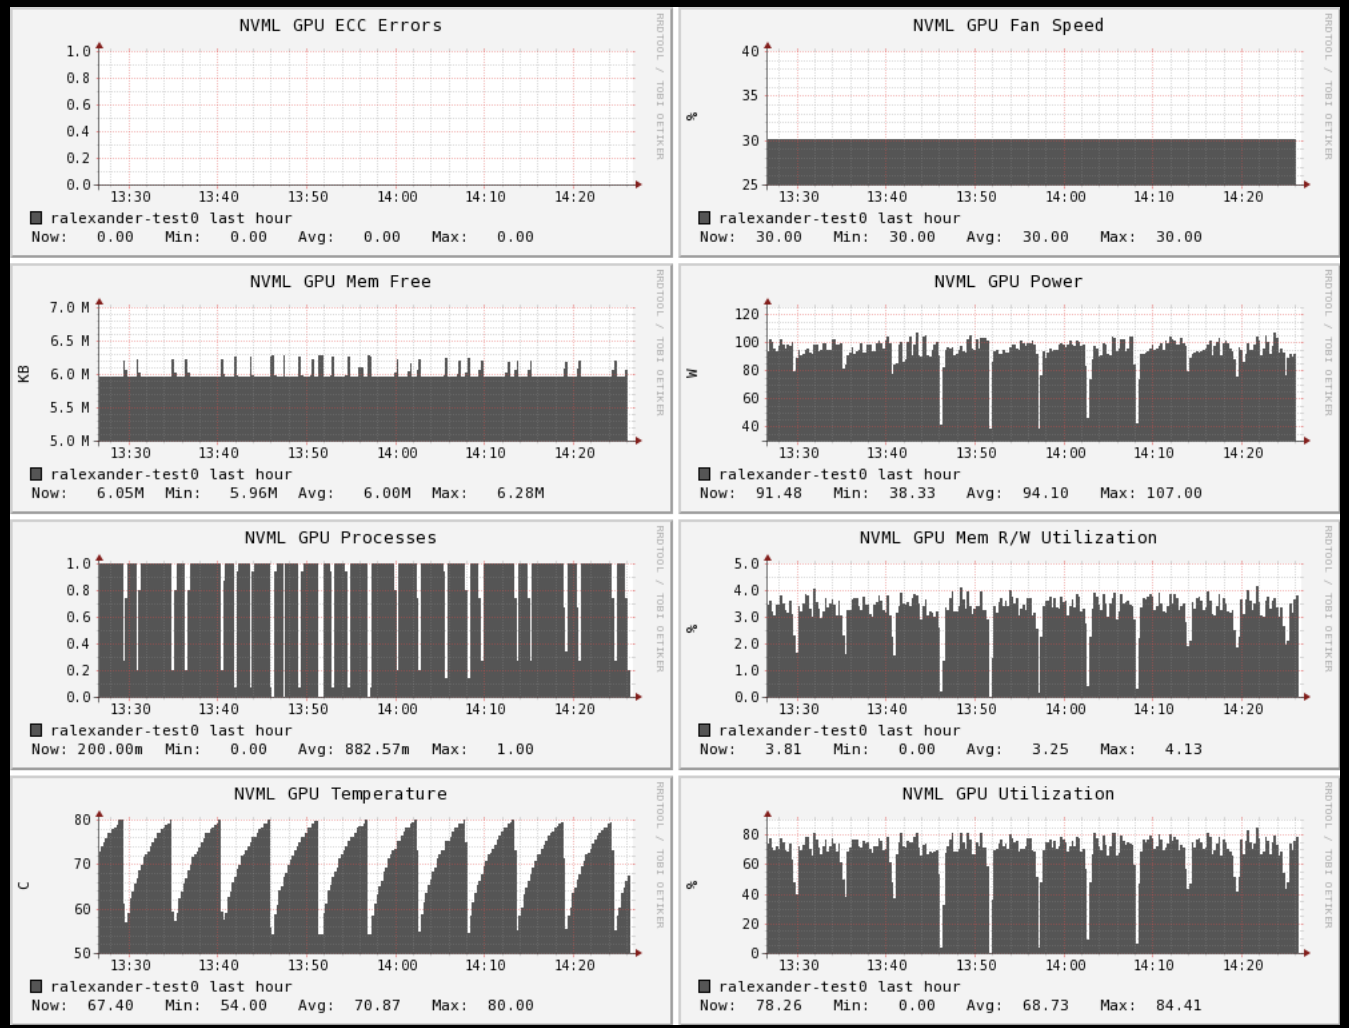

From http://blog.sflow.com

HP Insight Cluster Management Utility

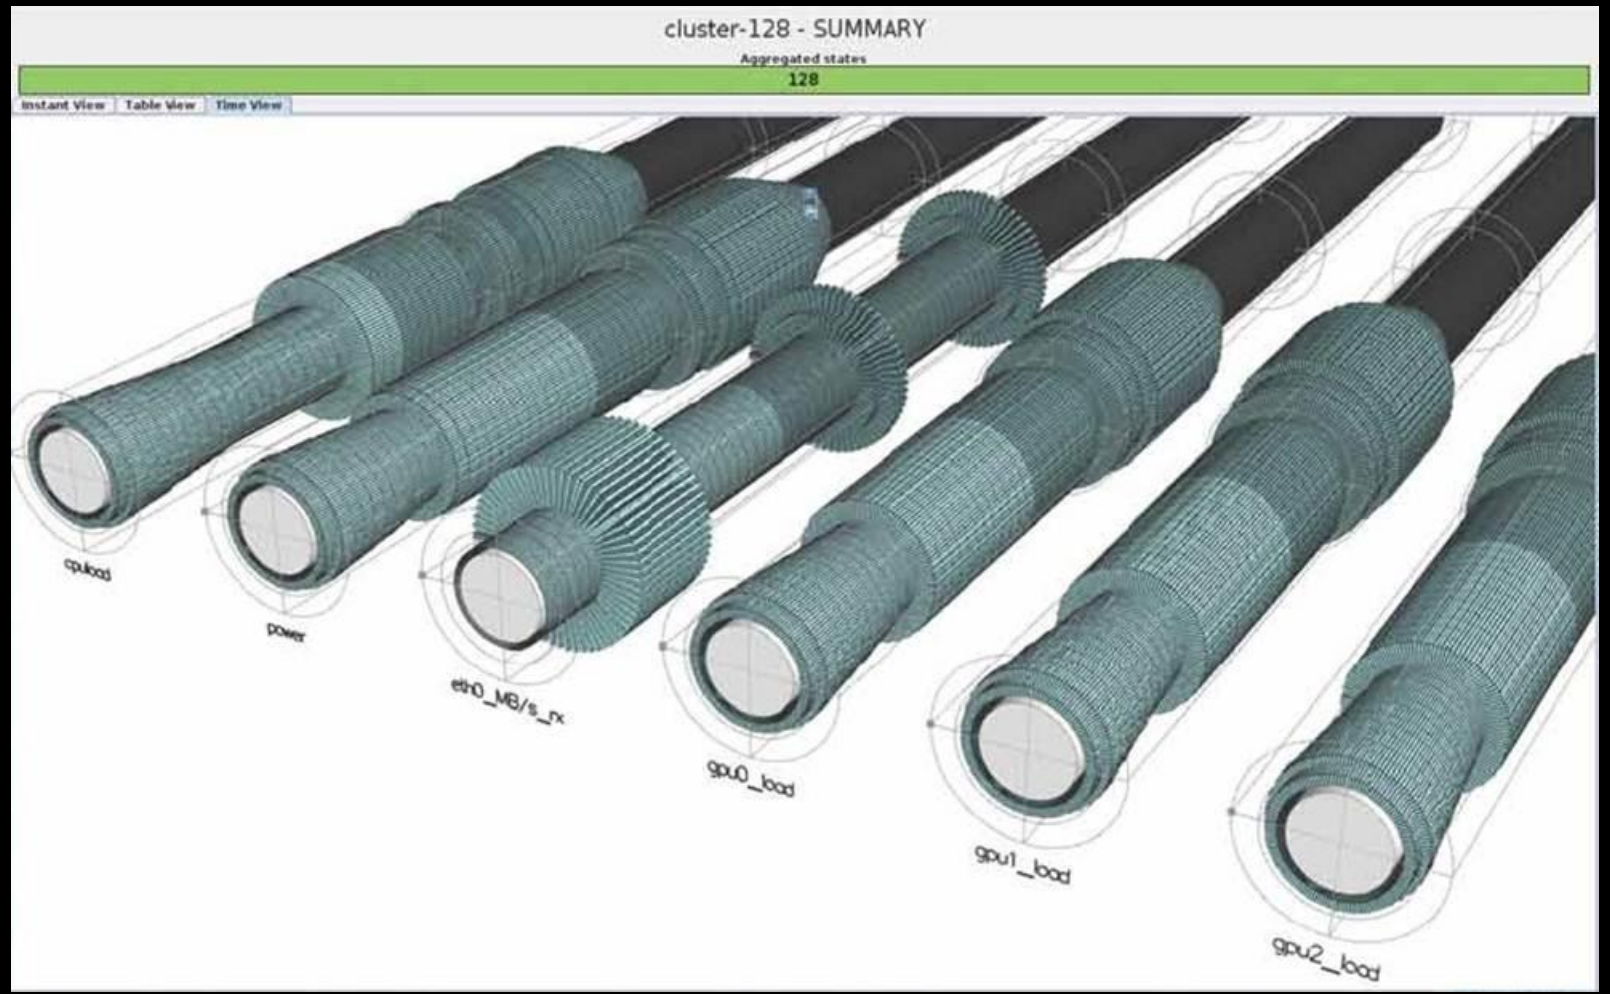

From http://hp.com

GPU TECHNOLOGY CONFERENCE

Managing GPUs with Bright Computing

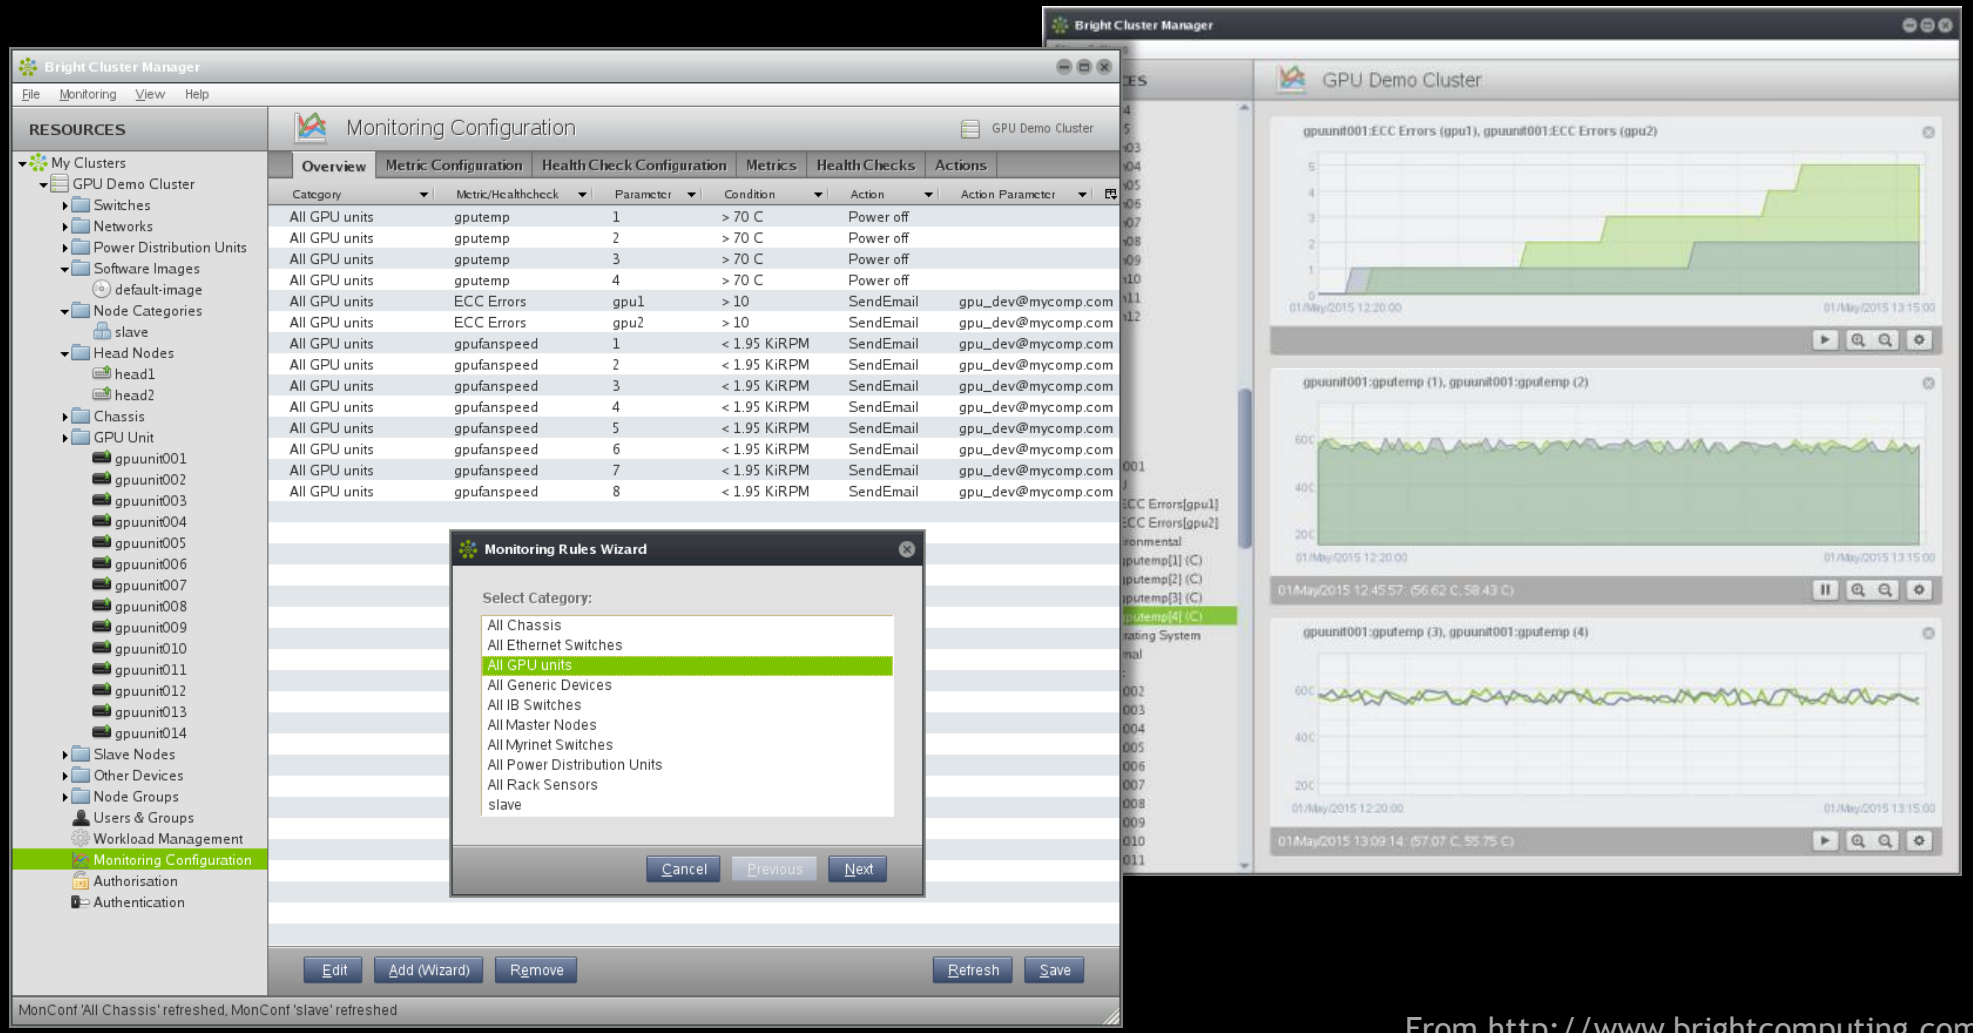

From http://www.brightcomputing.com/

Nagios / Icinga GPU Support

- Under developement
  - http://www.thomas-krenn.com/en/wiki/GPU_Sensor_Monitoring_Plugin

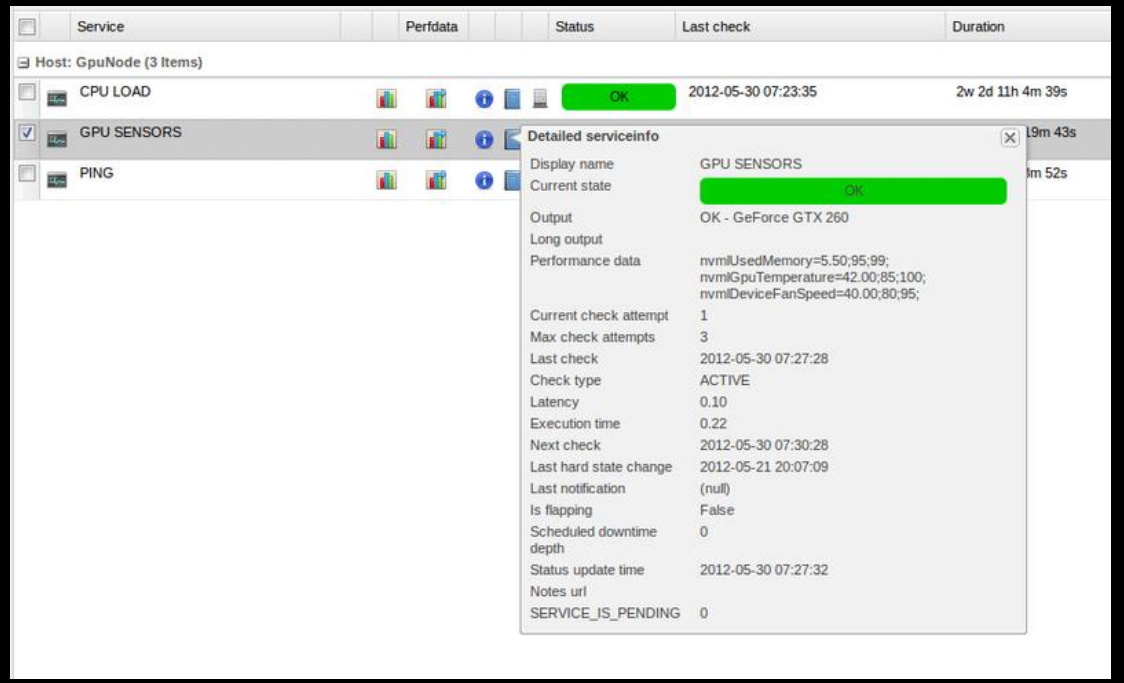

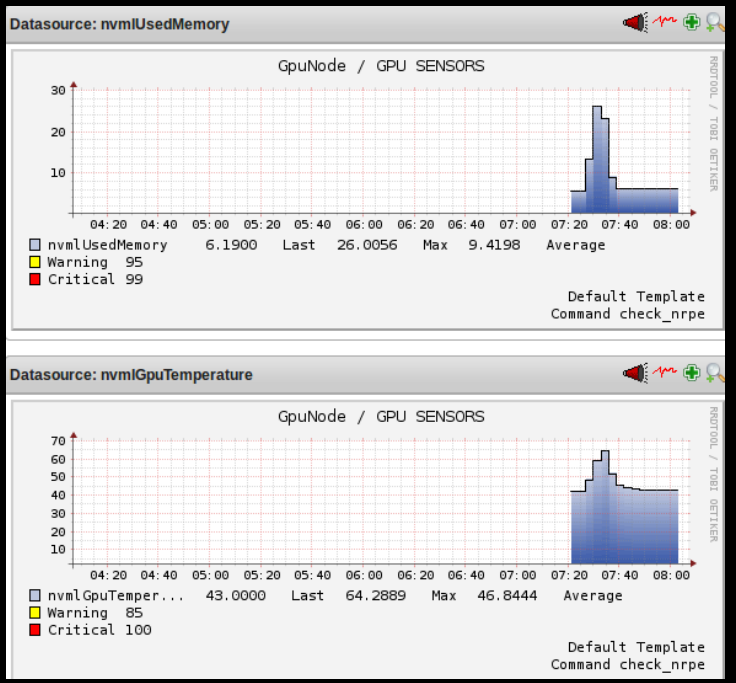

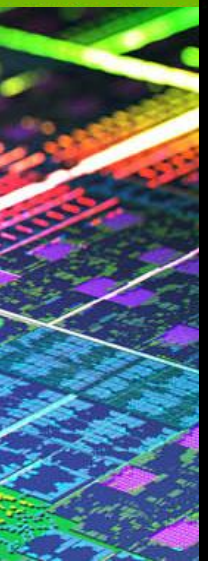

Learn more at GTC

- S3045 - Tips & Tricks for Getting the Most Out of GPU-accelerated Clusters by Rolf VandeVaart (NVIDIA)
  - Today 15:00 - 15:25 - Room 230C
- S3214 - CUDA in the Cloud: Enabling HPC Workloads in OpenStack by John Paul Walters (USC)
  - Today 16:00 - 16:25 - Room 111
- S3516 - Building Your Own GPU Research Cluster Using Open Source Software Stack by Pradeep Kumar Gupta (NVIDIA)
  - Tomorrow 10:00 - 10:25 - Room 210A

GPU TECHNOLOGY CONFERENCE

Learn more at GTC

- HP - Booth #107
- IBM - Booth #607
- Bright Computing
  - Booth #412
- Cray – Booth #301
- Dell – Booth #506

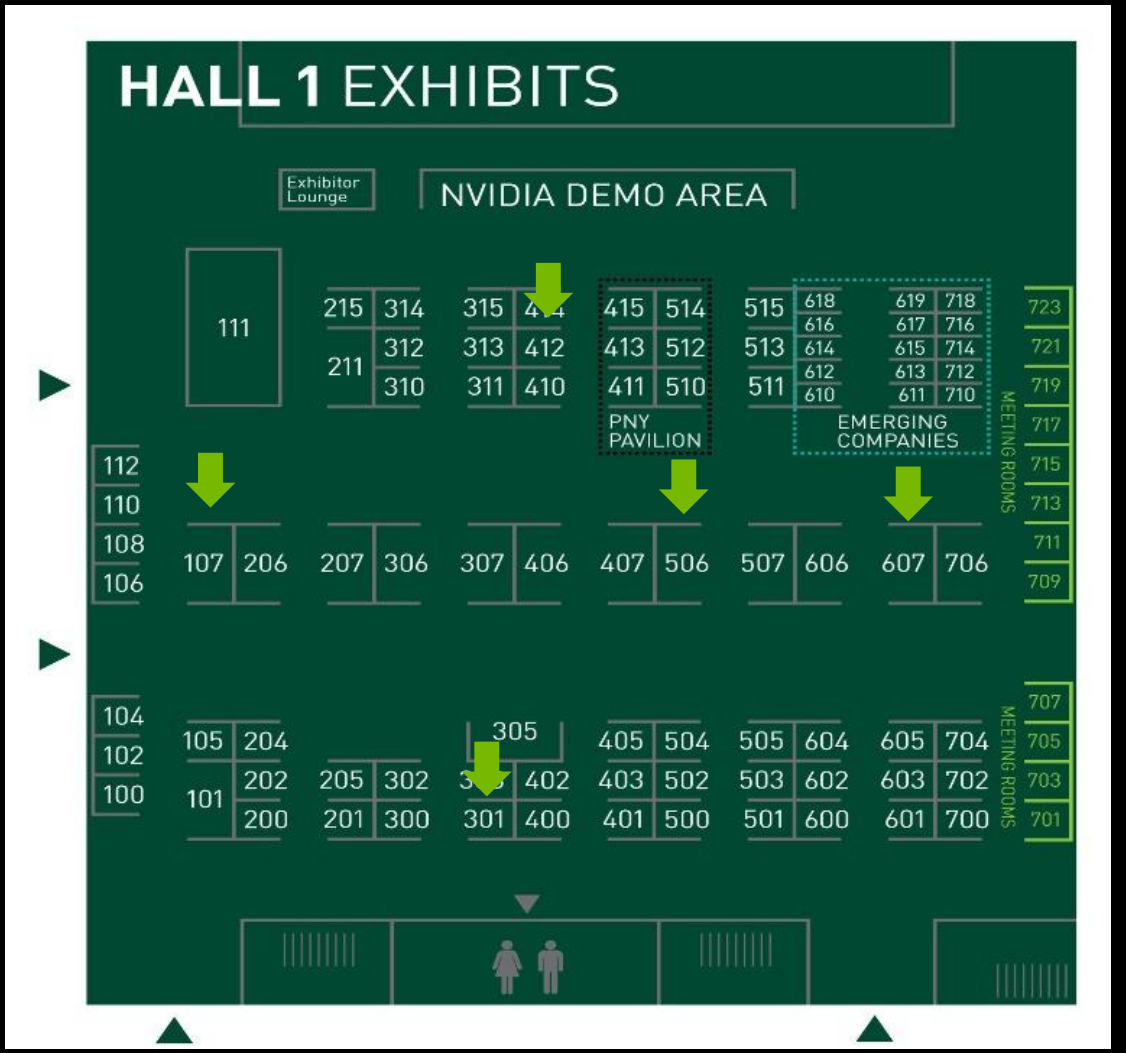

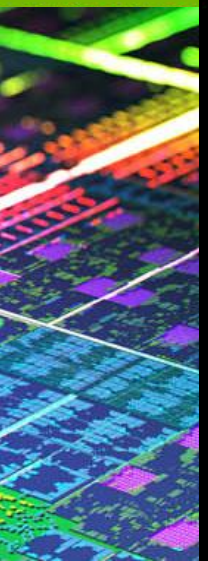

Learn more at GTC

- S3248 - Acceptance Testing a GPU Based Cluster (LANL)
- S3034 - Efficient Utilization of a CPU-GPU Cluster (NRL)
- S3556A - System Design of Kepler Based HPC Solutions (Presented by Dell Inc.)
- S3249 - Introduction to Deploying, Managing, and Using GPU Clusters (NVIDIA)
- S3536 - Accelerate GPU Innovation with HP Gen8 Servers (Presented by HP)
- S3578 - Optimizing GPU Utilization and Application Throughput in HPC Clusters (Presented by IBM)

Thank you

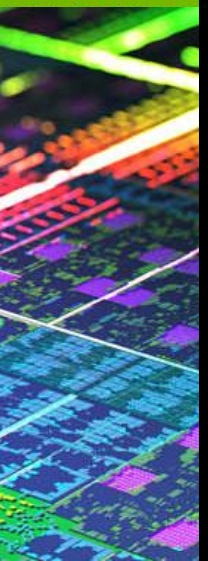

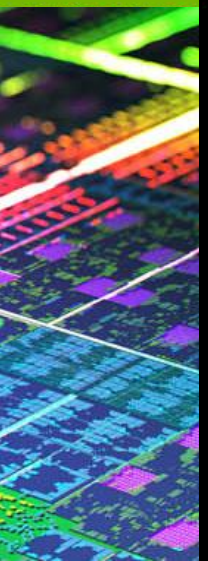

Contact information

- Contact my team at cudatools@nvidia.com

Questions?