GPU Cluster Monitoring and Management

November 15th, 2012

Introductions

- Robert Alexander
  - CUDA Tools Software Engineer at NVIDIA
  - Tesla Software Group

Overview

- NVIDIA Interfaces
- Third Party Interfaces
  - GPU Monitoring
  - GPU Management
  - GPU Scheduling
- nvidia-healthmon
- Links and contact info

NVIDIA Interfaces

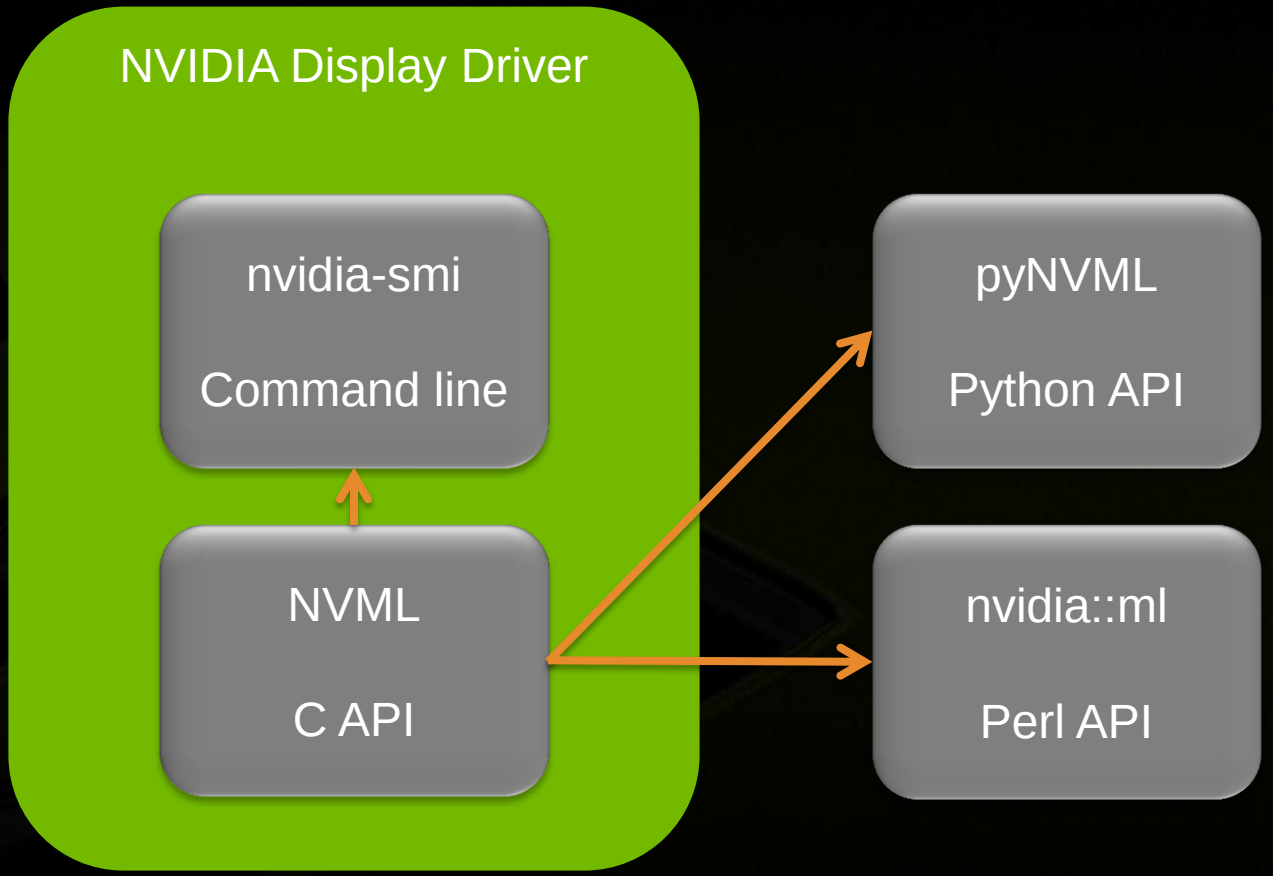

GPU Monitoring Ecosystem

Third Party Tools

NVIDIA APIs

nvidia-smi	NVML	pyNVML	nvidia::ml
Command line	C API	Python API	Perl API

sFlow + Ganglia GPU Monitoring

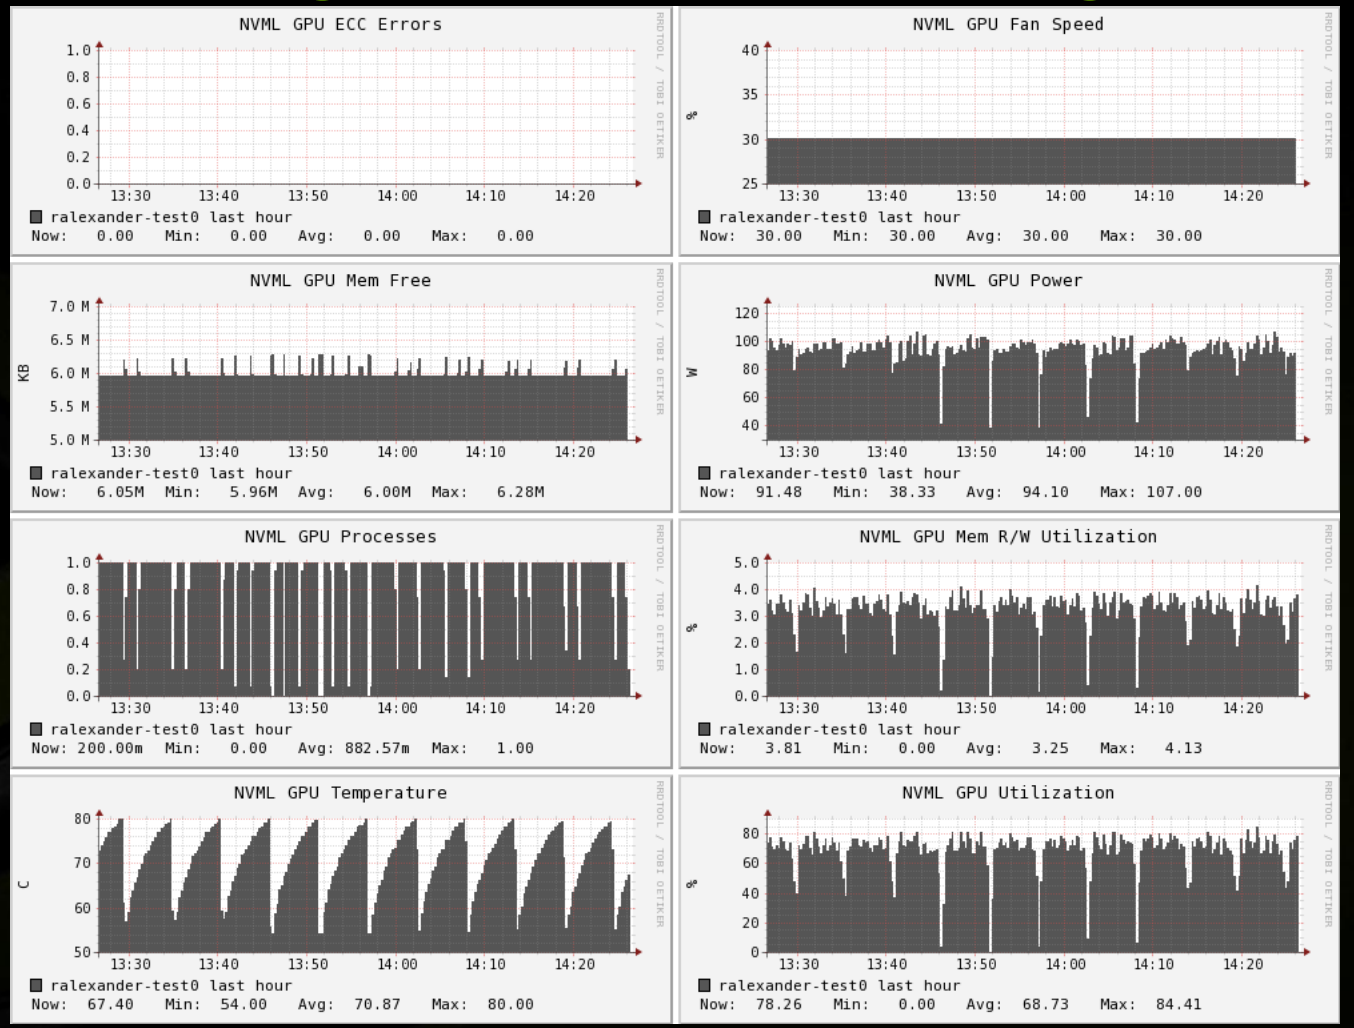

From http://blog.sflow.com

HP Insight Cluster Management Utility

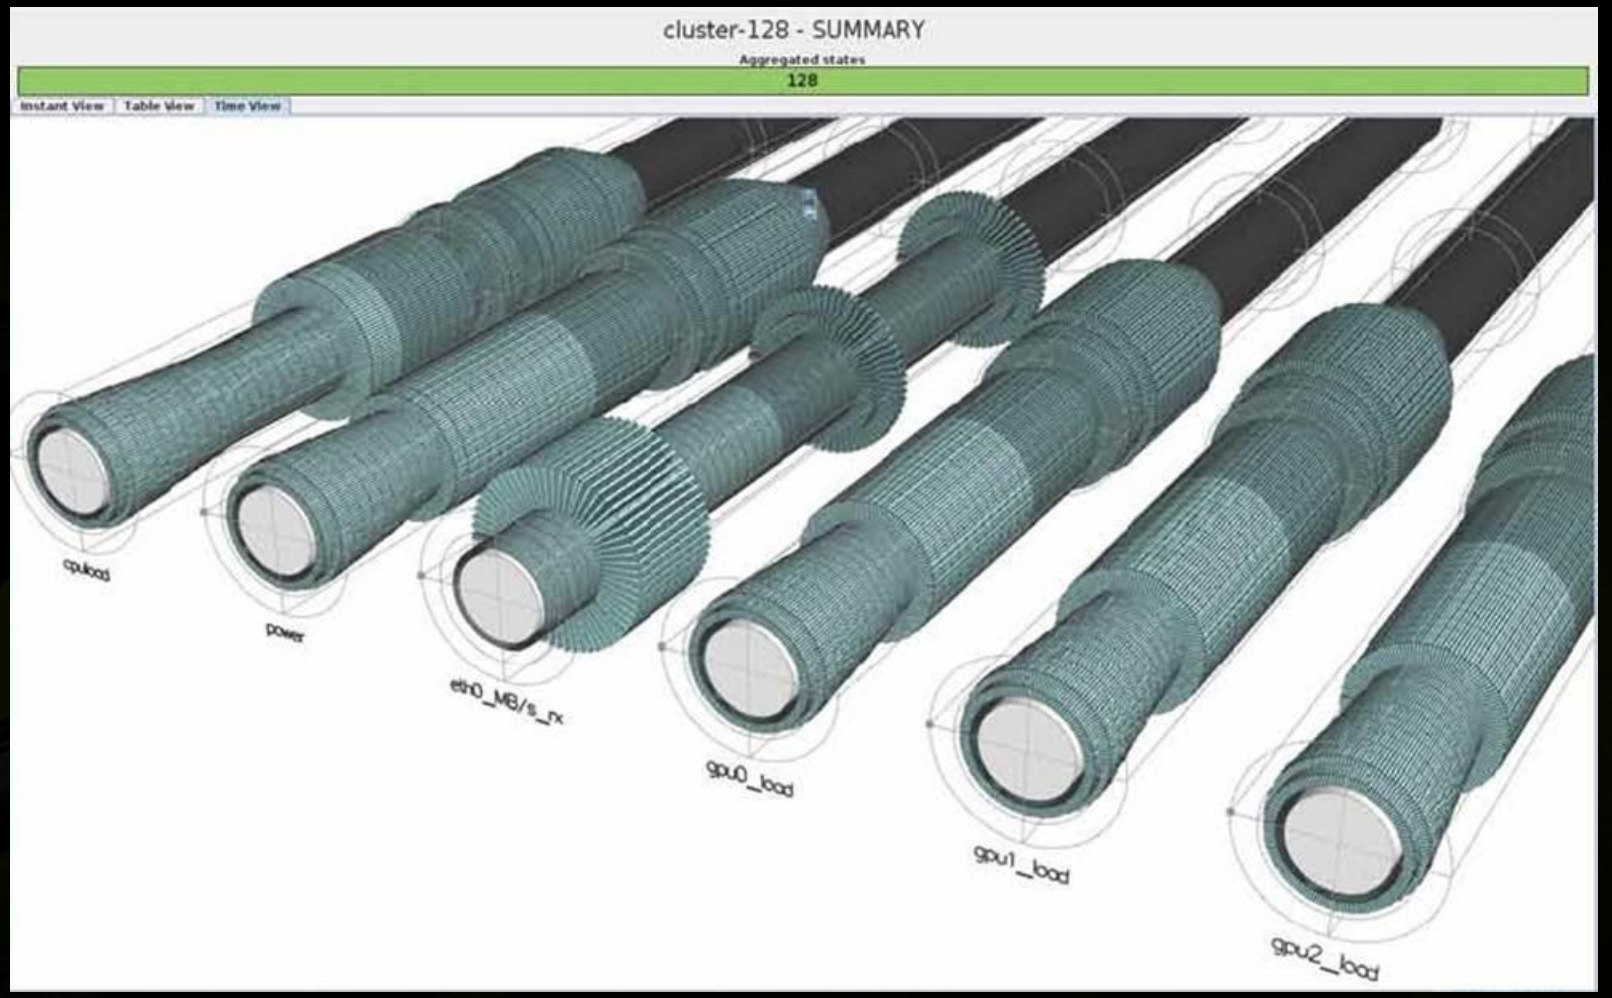

From http://hp.com

GPU Management Ecosystem

Third Party Tools

NVIDIA APIs

nvidia-smi

Command line

NVML

C API

pyNVML

Python API

nvidia::ml

Perl API

Managing GPUs with Bright Computing

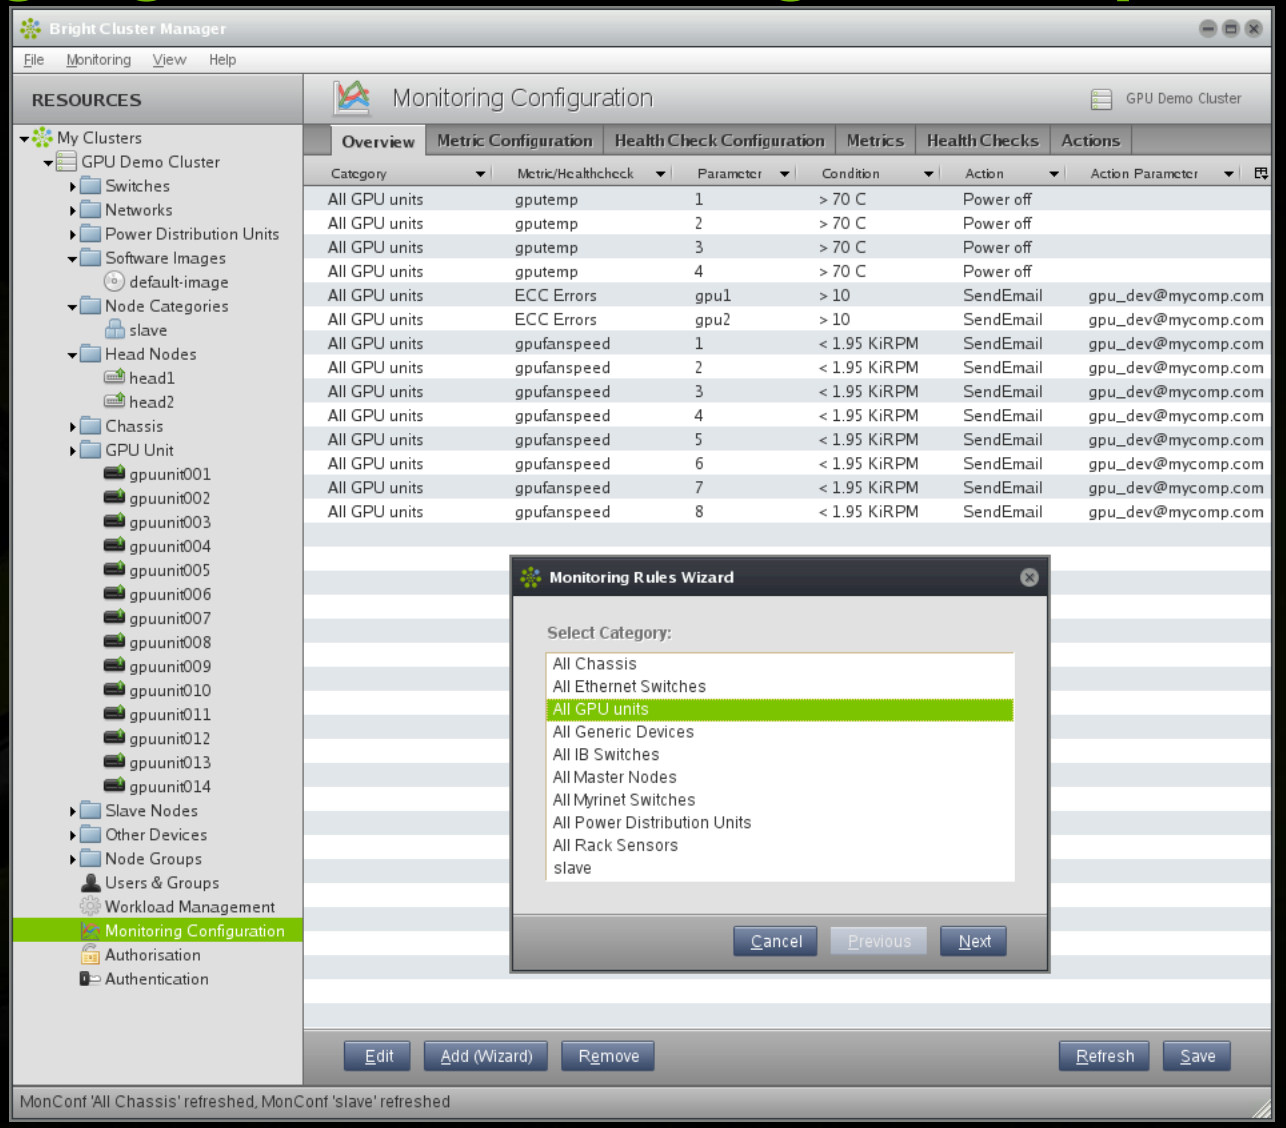

From http://www.brightcomputing.com/

GPU Scheduling Ecosystem

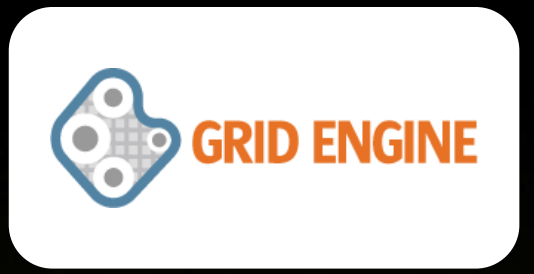

Third Party Tools

NVIDIA APIs

nvidia-smi	NVML	pyNVML	nvidia::ml
Command line	C API	Python API	Perl API

GPU Job Scheduling with TORQUE


```
$ ./configure –with-nvml-lib=<…> --with-nvml-include=<…>

$ qsub -l nodes=1:gpus=1 myjob.sh
$ qsub -l nodes=1:gpus=2:exclusive_process myjob.sh
```

GPU Job Scheduling with IBM PCM

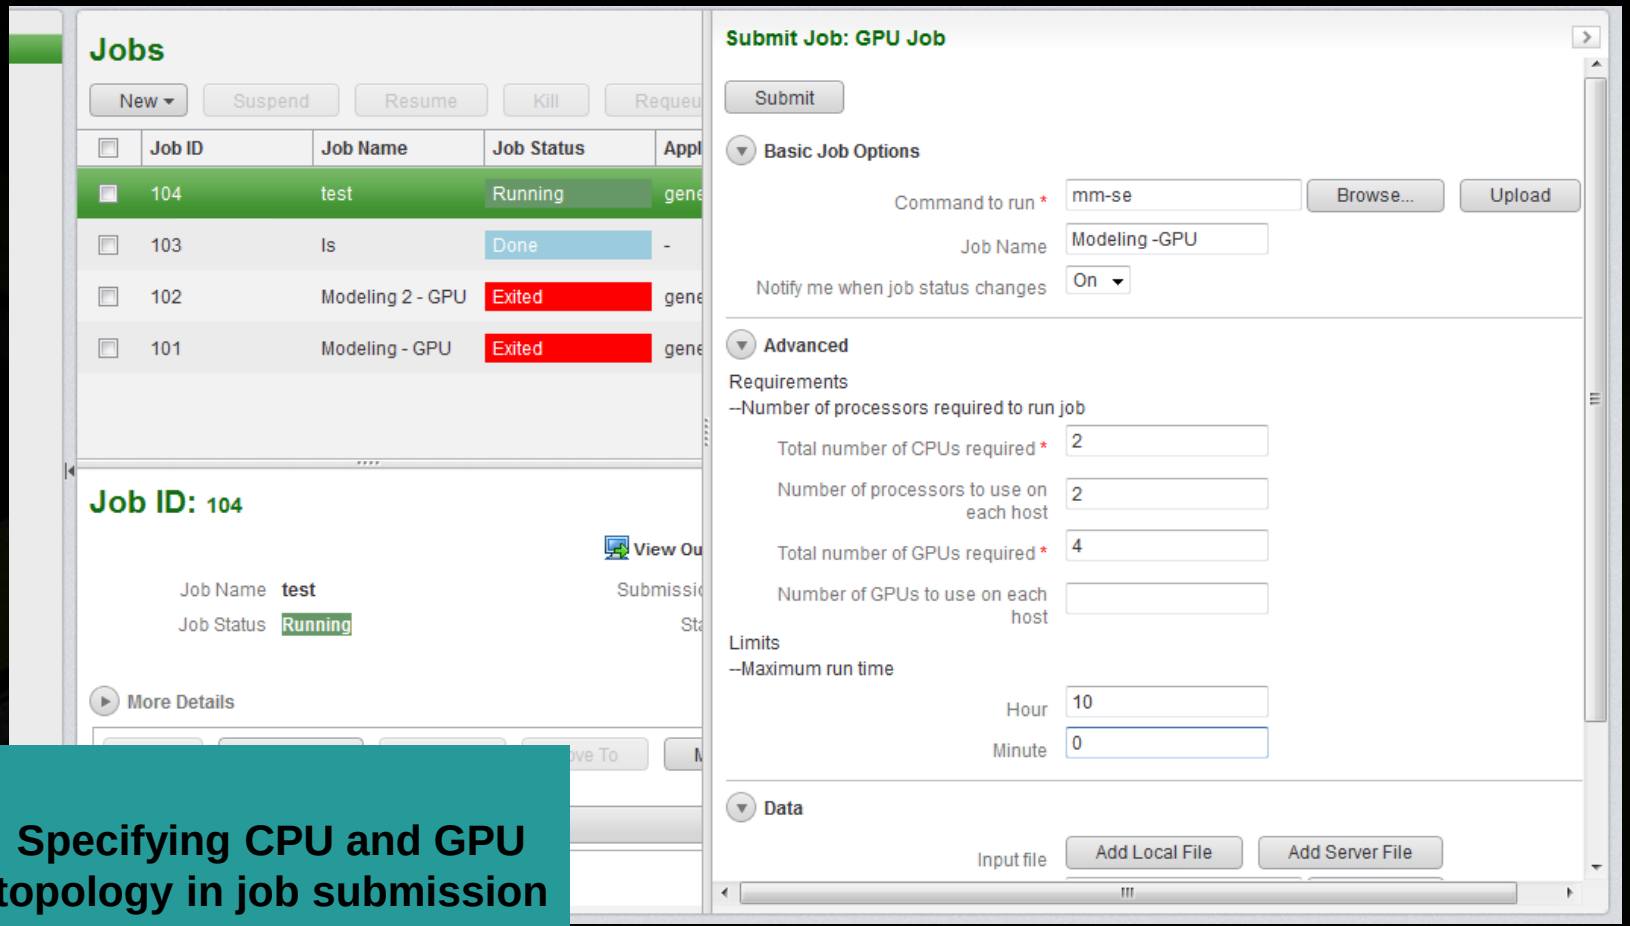

Specifying CPU and GPU topology in job submission

nvidia-healthmon

- Quick health check
- Not a full diagnostic
- Suggest remedies to SW and system configuration problems

nvidia-healthmon – Use Cases

- Cluster scheduler's prologue / epilogue script
- Heath and diagnostic suites
  - Designed to integrate into third party tools
  - Support in Warewulf NHC tool
  - http://warewulf.lbl.gov/trac/wiki/Node%20Health%20Check
- After provisioning cluster node
- Run directly, manually

nvidia-healthmon – Feature Set

- Basic CUDA and NVML sanity check
- Diagnosis of GPU failure-to-initialize problems
- Check for conflicting drivers (I.E. VESA)
- InfoROM validation
- Poorly seated GPU detection
- Check for disconnected power cables
- ECC error detection and reporting
- Bandwidth test

Downloads

- CUDA 5.0 Driver
  - http://developer.nvidia.com/cuda/cuda-toolkit
- 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/

Contact Info and Links

- Robert Alexander
  - ralexander@nvidia.com
- Developer Information
  - http://www.developer.nvidia.com/nvidia-management-library-nvml

Third Party Links

- http://ganglia.sourceforge.net/
- http://sflow.org/
- http://www.hp.com
- http://www.ibm.com
- http://www.adaptivecomputing.com/
- http://www.brightcomputing.com/
- http://www.penguincomputing.com/
- http://www.stackiq.com/
- http://www.univa.com/
- http://gridscheduler.sourceforge.net/
- https://computing.llnl.gov/linux/slurm/
- http://www.pbsworks.com/

Thanks!

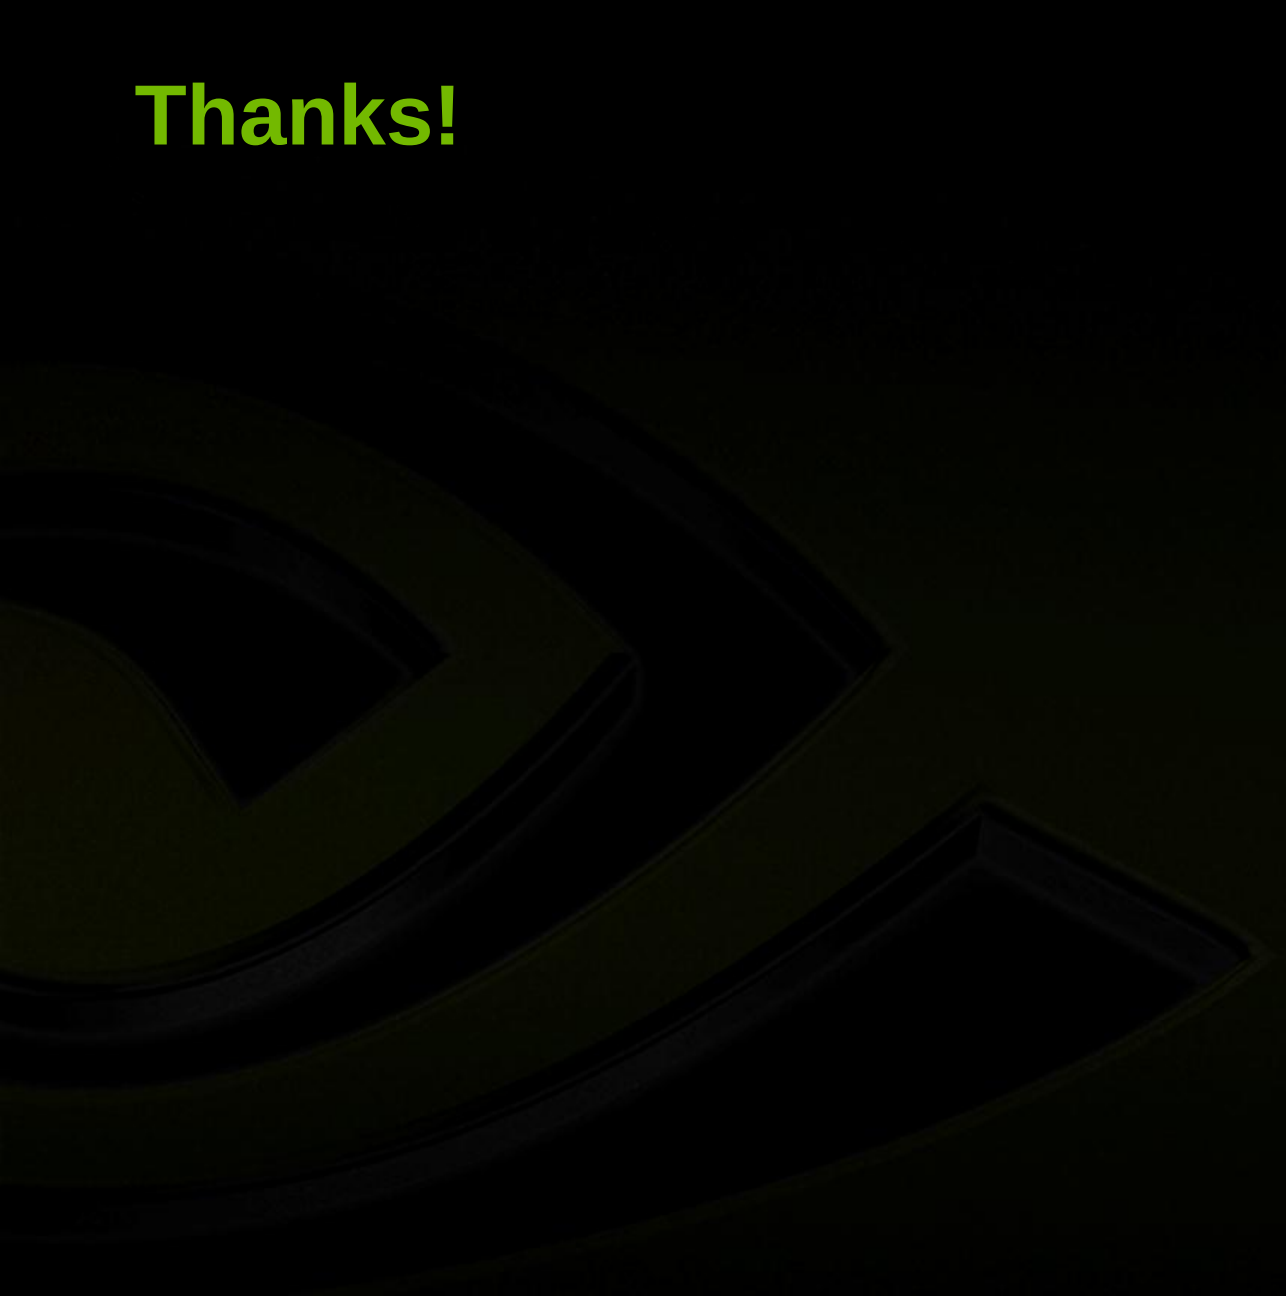

Questions?