gpus1 – Ubuntu 10.04 Available via ssh

```
root@gpus1:[~]#lspci -v | grep VGA
01:04.0 VGA compatible controller: Matrox Graphics, Inc. MGA G200eW WPCM450 (rev 0a)
03:00.0 VGA compatible controller: nVidia Corporation GF100 [Tesla C2050 / C2070] (rev a3)

root@gpus1:[~]#cd /gpusrc/NVIDIA_GPU_Computing_SDK/C/bin/linux/release/
root@gpus1:[/gpusrc/NVIDIA_GPU_Computing_SDK/C/bin/linux/release]#./deviceQuery
./deviceQuery Starting...

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

There is 1 device supporting CUDA

Device 0: "Tesla C2070"
  CUDA Driver Version:                           3.20
  CUDA Runtime Version:                          3.20
  CUDA Capability Major/Minor version number:    2.0
  Total amount of global memory:                 5636554752 bytes
  Multiprocessors x Cores/MP = Cores:            14 (MP) x 32 (Cores/MP) = 448 (Cores)
  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 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 1
  Maximum memory pitch:                          2147483647 bytes
  Texture alignment:                             512 bytes
  Clock rate:                                    1.15 GHz
  Concurrent copy and execution:                 Yes
  Run time limit on kernels:                     No
  Integrated:                                    No
  Support host page-locked memory mapping:       Yes
  Compute mode:                                  Default (multiple host threads can use this device
simultaneously)
  Concurrent kernel execution:                   Yes
  Device has ECC support enabled:                Yes
  Device is using TCC driver mode:               No

deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 3.20, CUDA Runtime Version =
3.20, NumDevs = 1, Device = Tesla C2070


PASSED


========================================
```

gpus2 – Ubuntu 10.04 Available via ssh

```
root@gpus2:[~]#lspci -v | grep VGA
01:04.0 VGA compatible controller: Matrox Graphics, Inc. MGA G200eW WPCM450 (rev 0a)
06:00.0 VGA compatible controller: ATI Technologies Inc Antilles [AMD Radeon HD 6990]

root@gpus2:[~]#fglrxinfo
display: :0.0  screen: 0
OpenGL vendor string: ATI Technologies Inc.
OpenGL renderer string: AMD Radeon HD 6990
OpenGL version string: 4.1.10750 Compatibility Profile Context

dhe@gpus2:[/gpusrc/AMD-APP-SDK-v2.4-lnx64/bin/x86_64]$./clinfo
Number of platforms:                 1
  Platform Profile:                  FULL_PROFILE
  Platform Version:                  OpenCL 1.1 AMD-APP-SDK-v2.4 (595.10)
  Platform Name:                     AMD Accelerated Parallel Processing
  Platform Vendor:                   Advanced Micro Devices, Inc.
  Platform Extensions:               cl_khr_icd cl_amd_event_callback
cl_amd_offline_devices


  Platform Name:                     AMD Accelerated Parallel Processing
Number of devices:                   1
  Device Type:                       CL_DEVICE_TYPE_CPU
  Device ID:                         4098
  Max compute units:                 24
  Max work items dimensions:               3
    Max work items[0]:                     1024
    Max work items[1]:                     1024
    Max work items[2]:                     1024
  Max work group size:                     1024
  Preferred vector width char:       16
  Preferred vector width short:            8
  Preferred vector width int:        4
  Preferred vector width long:       2
  Preferred vector width float:            4
  Preferred vector width double:     0
  Native vector width char:          16
  Native vector width short:         8
  Native vector width int:           4
  Native vector width long:          2
  Native vector width float:         4
  Native vector width double:        0
  Max clock frequency:                     800Mhz
  Address bits:                      64
  Max memory allocation:             16923791360
  Image support:                     Yes
  Max number of images read arguments:     128
```

Max number of images write arguments: 8
Max image 2D width: 8192
Max image 2D height: 8192
Max image 3D width: 2048
Max image 3D height: 2048
Max image 3D depth: 2048
Max samplers within kernel: 16
Max size of kernel argument: 4096
Alignment (bits) of base address: 1024
Minimum alignment (bytes) for any datatype: 128
Single precision floating point capability
Denorms: Yes
Quiet NaNs: Yes
Round to nearest even: Yes
Round to zero: Yes
Round to +ve and infinity: Yes
IEEE754-2008 fused multiply-add: No
Cache type: Read/Write
Cache line size: 64
Cache size: 65536
Global memory size: 67695165440
Constant buffer size: 65536
Max number of constant args: 8
Local memory type: Global
Local memory size: 32768
Kernel Preferred work group size multiple: 1
Error correction support: 0
Unified memory for Host and Device: 1
Profiling timer resolution: 1
Device endianess: Little
Available: Yes
Compiler available: Yes
Execution capabilities:
Execute OpenCL kernels: Yes
Execute native function: Yes
Queue properties:
Out-of-Order: No
Profiling : Yes
Platform ID: 0x7f23cec69800
Name: AMD Opteron(tm) Processor 6168
Vendor: AuthenticAMD
Driver version: 2.0
Profile: FULL_PROFILE
Version: OpenCL 1.1 AMD-APP-SDK-v2.4 (595.10)
Extensions: cl_khr_fp64 cl_amd_fp64 cl_khr_global_int32_base_atomics cl_khr_global_int32_extended_atomics cl_khr_local_int32_base_atomics cl_khr_local_int32_extended_atomics cl_khr_int64_base_atomics cl_khr_int64_extended_atomics cl_khr_byte_addressable_store cl_khr_gl_sharing cl_ext_device_fission cl_amd_device_attribute_query cl_amd_vec3 cl_amd_media_ops

cl_amd_popcnt cl_amd_printf

dhe@gpus2:[/gpusrc/AMD-APP-SDK-v2.4-lnx64/bin/x86_64]$
==========

gpus3 – Windows 2008r2
Windows 2008 R2 Server ready to test. You can use the standard rdp option to connect and test. Rdesktop to gpus3.cise.ufl.edu. If you need access to the Nvidia Control Panel applet, download and install Tightvnc 2.0.3 Viewer from the link below.

http://www.tightvnc.com/download.php

Enter the vncviewer password to connect.
Vncviewer password Vncv;ew

Packages installed:

Furmark GPU Stress Test
GPU Caps Viewer
Visual Studio 2008
Visual Studio 2010
Microsoft DirectX SDK
Nvidia 3D Vision Controller and Drivers
Nvidia Cuda Toolkit v.4
Nvidia Cuda Tools API
Nvidia GPU Computing SDK 4.0
Nvidia Parallel Nsight 2.0 for Visual Studio
Nvidia Physx System Software
Tightvnc Server

==========

gpus4 – Ubuntu 10.04 Available via ssh
root@gpus4:[~]#lspci -v | grep VGA
01:04.0 VGA compatible controller: Matrox Graphics, Inc. MGA G200eW WPCM450 (rev 0a)
05:00.0 VGA compatible controller: nVidia Corporation Device 1088 (rev a1)

root@gpus4:[/gpusrc/NVIDIA_GPU_Computing_SDK/C/bin/linux/release]#./deviceQuery
[deviceQuery] starting...
./deviceQuery Starting...

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

Found 2 CUDA Capable device(s)

Device 0: "GeForce GTX 590"
CUDA Driver Version / Runtime Version 4.0 / 4.0

CUDA Capability Major/Minor version number: 2.0
Total amount of global memory: 1536 MBytes (1610285056 bytes)
(16) Multiprocessors x (32) CUDA Cores/MP: 512 CUDA Cores
GPU Clock Speed: 1.26 GHz
Memory Clock rate: 1728.00 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 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 execution: Yes with 1 copy engine(s)
Run time limit on kernels: No
Integrated GPU sharing Host Memory: No
Support host page-locked memory mapping: Yes
Concurrent kernel execution: Yes
Alignment requirement for Surfaces: Yes
Device has ECC support enabled: No
Device is using TCC driver mode: No
Device supports Unified Addressing (UVA):   Yes
Device PCI Bus ID / PCI location ID: 6 / 0
Compute Mode:
   < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

Device 1: "GeForce GTX 590"
CUDA Driver Version / Runtime Version   4.0 / 4.0
CUDA Capability Major/Minor version number: 2.0
Total amount of global memory: 1536 MBytes (1610285056 bytes)
(16) Multiprocessors x (32) CUDA Cores/MP: 512 CUDA Cores
GPU Clock Speed: 1.26 GHz
Memory Clock rate: 1728.00 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 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 execution:                 Yes with 1 copy engine(s)
  Run time limit on kernels:                     No
  Integrated GPU sharing Host Memory:            No
  Support host page-locked memory mapping:       Yes
  Concurrent kernel execution:                   Yes
  Alignment requirement for Surfaces:            Yes
  Device has ECC support enabled:                No
  Device is using TCC driver mode:               No
  Device supports Unified Addressing (UVA):      Yes
  Device PCI Bus ID / PCI location ID:           5 / 0
  Compute Mode:
    < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 4.0, CUDA Runtime Version = 4.0,
NumDevs = 2, Device = GeForce GTX 590, Device = GeForce GTX 590
[deviceQuery] test results...
PASSED

============================================================
```

gpus5 – Ubuntu 10.04 Available via ssh

```
root@gpus5:[~]#lspci | grep VGA
01:04.0 VGA compatible controller: Matrox Graphics, Inc. MGA G200eW WPCM450 (rev 0a)
03:00.0 VGA compatible controller: ATI Technologies Inc Cayman PRO [AMD Radeon 6900 Series]
41:00.0 VGA compatible controller: ATI Technologies Inc Cayman PRO [AMD Radeon 6900 Series]

root@gpus5:[~]#fglrxinfo
display: :0.0  screen: 0
OpenGL vendor string: ATI Technologies Inc.
OpenGL renderer string: AMD Radeon HD 6900 Series
OpenGL version string: 4.1.10750 Compatibility Profile Context

root@gpus5:[/gpusrc/AMD-APP-SDK-v2.4-lnx64/bin/x86_64]#./clinfo
Number of platforms:                           1
  Platform Profile:                            FULL_PROFILE
  Platform Version:                            OpenCL 1.1 AMD-APP-SDK-v2.4 (595.10)
  Platform Name:                               AMD Accelerated Parallel Processing
  Platform Vendor:                             Advanced Micro Devices, Inc.
  Platform Extensions:                                 cl_khr_icd cl_amd_event_callback
cl_amd_offline_devices


  Platform Name:                               AMD Accelerated Parallel Processing
Number of devices:                             2
  Device Type:                                 CL_DEVICE_TYPE_GPU
  Device ID:                                   4098
```

Max compute units:	22
Max work items dimensions:	3
Max work items[0]:	256
Max work items[1]:	256
Max work items[2]:	256
Max work group size:	256
Preferred vector width char:	16
Preferred vector width short:	8
Preferred vector width int:	4
Preferred vector width long:	2
Preferred vector width float:	4
Preferred vector width double:	0
Native vector width char:	16
Native vector width short:	8
Native vector width int:	4
Native vector width long:	2
Native vector width float:	4
Native vector width double:	0
Max clock frequency:	0Mhz
Address bits:	32
Max memory allocation:	134217728
Image support:	Yes
Max number of images read arguments:	128
Max number of images write arguments:	8
Max image 2D width:	8192
Max image 2D height:	8192
Max image 3D width:	2048
Max image 3D height:	2048
Max image 3D depth:	2048
Max samplers within kernel:	16
Max size of kernel argument:	1024
Alignment (bits) of base address:	32768
Minimum alignment (bytes) for any datatype:	128
Single precision floating point capability	
Denorms:	No
Quiet NaNs:	Yes
Round to nearest even:	Yes
Round to zero:	Yes
Round to +ve and infinity:	Yes
IEEE754-2008 fused multiply-add:	Yes
Cache type:	None
Cache line size:	0
Cache size:	0
Global memory size:	536870912
Constant buffer size:	65536
Max number of constant args:	8
Local memory type:	Scratchpad
Local memory size:	32768
Kernel Preferred work group size multiple:	64

Error correction support: 0
Unified memory for Host and Device: 0
Profiling timer resolution: 1
Device endianess: Little
Available: Yes
Compiler available: Yes
Execution capabilities:
Execute OpenCL kernels: Yes
Execute native function: No
Queue properties:
Out-of-Order: No
Profiling : Yes
Platform ID: 0x7f859d4dc800
Name: Cayman
Vendor: Advanced Micro Devices, Inc.
Driver version: CAL 1.4.1385
Profile: FULL_PROFILE
Version: OpenCL 1.1 AMD-APP-SDK-v2.4 (595.10)
Extensions: cl_amd_fp64 cl_khr_global_int32_base_atomics cl_khr_global_int32_extended_atomics cl_khr_local_int32_base_atomics cl_khr_local_int32_extended_atomics cl_khr_3d_image_writes cl_khr_byte_addressable_store cl_khr_gl_sharing cl_amd_device_attribute_query cl_amd_printf cl_amd_media_ops cl_amd_popcnt

Device Type: CL_DEVICE_TYPE_CPU
Device ID: 4098
Max compute units: 24
Max work items dimensions: 3
Max work items[0]: 1024
Max work items[1]: 1024
Max work items[2]: 1024
Max work group size: 1024
Preferred vector width char: 16
Preferred vector width short: 8
Preferred vector width int: 4
Preferred vector width long: 2
Preferred vector width float: 4
Preferred vector width double: 0
Native vector width char: 16
Native vector width short: 8
Native vector width int: 4
Native vector width long: 2
Native vector width float: 4
Native vector width double: 0
Max clock frequency: 800Mhz
Address bits: 64
Max memory allocation: 16923774976
Image support: Yes
Max number of images read arguments: 128

```
  Max number of images write arguments:                8
  Max image 2D width:                                  8192
  Max image 2D height:                                 8192
  Max image 3D width:                                  2048
  Max image 3D height:                                 2048
  Max image 3D depth:                                  2048
  Max samplers within kernel:                          16
  Max size of kernel argument:                         4096
  Alignment (bits) of base address:                    1024
  Minimum alignment (bytes) for any datatype:          128
  Single precision floating point capability
    Denorms:                                           Yes
    Quiet NaNs:                                        Yes
    Round to nearest even:                             Yes
    Round to zero:                                     Yes
    Round to +ve and infinity:                         Yes
    IEEE754-2008 fused multiply-add:                   No
  Cache type:                                          Read/Write
  Cache line size:                                     64
  Cache size:                                          65536
  Global memory size:                                  67695099904
  Constant buffer size:                                65536
  Max number of constant args:                         8
  Local memory type:                                   Global
  Local memory size:                                   32768
  Kernel Preferred work group size multiple:           1
  Error correction support:                            0
  Unified memory for Host and Device:                  1
  Profiling timer resolution:                          1
  Device endianess:                                    Little
  Available:                                           Yes
  Compiler available:                                  Yes
  Execution capabilities:
    Execute OpenCL kernels:                            Yes
    Execute native function:                           Yes
  Queue properties:
    Out-of-Order:                                      No
    Profiling :                                        Yes
  Platform ID:                                         0x7f859d4dc800
  Name:                                                AMD Opteron(tm) Processor 6168
  Vendor:                                              AuthenticAMD
  Driver version:                                      2.0
  Profile:                                             FULL_PROFILE
  Version:                                             OpenCL 1.1 AMD-APP-SDK-v2.4 (595.10)
  Extensions:                                          cl_khr_fp64 cl_amd_fp64
cl_khr_global_int32_base_atomics cl_khr_global_int32_extended_atomics
cl_khr_local_int32_base_atomics cl_khr_local_int32_extended_atomics cl_khr_int64_base_atomics
cl_khr_int64_extended_atomics cl_khr_byte_addressable_store cl_khr_gl_sharing
cl_ext_device_fission cl_amd_device_attribute_query cl_amd_vec3 cl_amd_media_ops
```

cl_amd_popcnt cl_amd_printf

root@gpus5:[/gpusrc/AMD-APP-SDK-v2.4-lnx64/bin/x86_64]#
==========
Important Directory Notes:
/gpusrc
This is a local place that you can copy your code and store your own personal copy of the sdk.
Warning: This space is not backed up!

There is a modified cuda SDK in /gpusrc.
NVIDIA_GPU_Computing_SDK

The ATI Dev Kit can also be found in /gpusrc under:
AMD-APP-SDK-v2.4-lnx64

There are some path-ing issues with the default download which have been corrected in the sdk copied there.

==========
CUDA Programming Notes:
* Please make sure your PATH includes /usr/local/cuda/bin
* Please make sure your LD_LIBRARY_PATH
* for 32-bit Linux distributions includes /usr/local/cuda/lib
* for 64-bit Linux distributions includes /usr/local/cuda/lib64:/usr/local/cuda/lib
* OR
* for 32-bit Linux distributions add /usr/local/cuda/lib
* for 64-bit Linux distributions add /usr/local/cuda/lib64 and /usr/local/cuda/lib
* to /etc/ld.so.conf and run ldconfig as root

* Please read the release notes in /usr/local/cuda/doc/

These paths can automatically be setup for you, by dot'ting your shell with either of these two files, depending on the shell you are using:

root@gpus4:[/gpusrc]#ls
cuda_bash_setup cuda_tcsh_setup

example:
$. cuda_bash_setup

==========
ATI NOTES:

You must have
amd/ati libraries in your path for development, to make this easier for you, there is an example setup under /gpusrc you can simply dot your shell with this file.

opencl_bash_setup

```
$. opencl_bash_setup
```