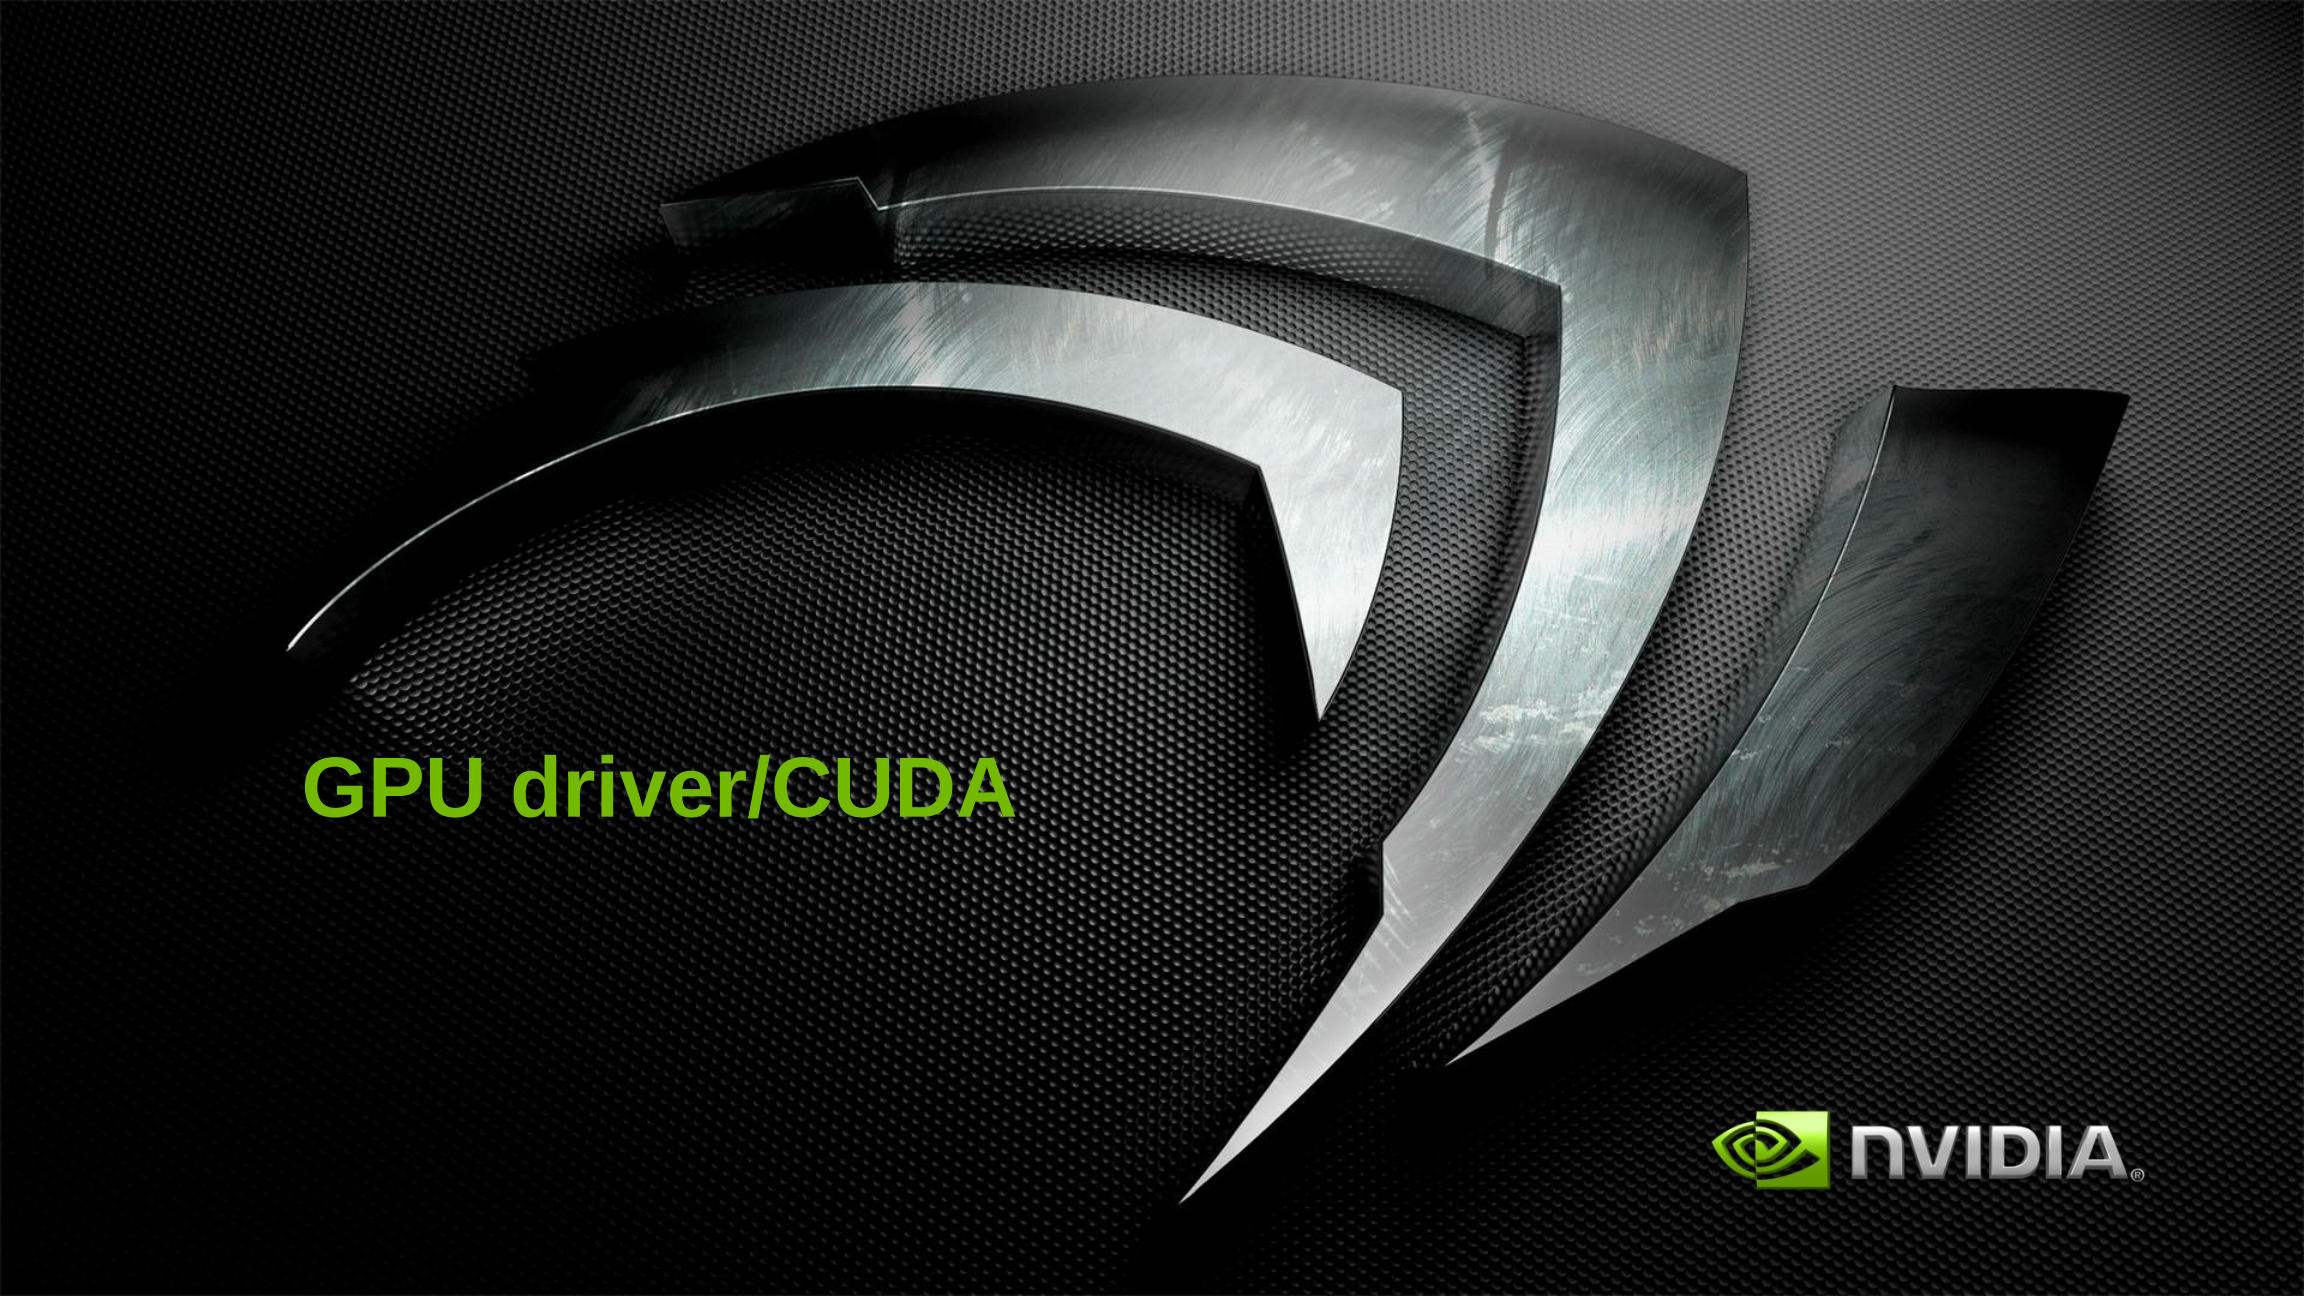
GPU driver/CUDA
NVIDIA

Software

- GPU Driver
- CUDA toolkit
  - Includes all the software necessary for developers to write applications
    - Compiler (nvcc), Libraries, Profiler, Documentation
- SDK
  - Not strictly required but a good idea for ensuring your system is running properly.
  - Many examples with code samples illustrating lots of the important programming constructs and techniques.
- www.nvidia.com/getcuda Above software from NVIDIA is free

Examine GPU h/w and driver

- `nvidia-smi`
- -h for help
- -q for long query of all GPUs
  - PCIe Bus ID
  - Driver version
  - ECC state
  - Power state
  - Fans/Temps/Clockspeeds

nvidia-smi

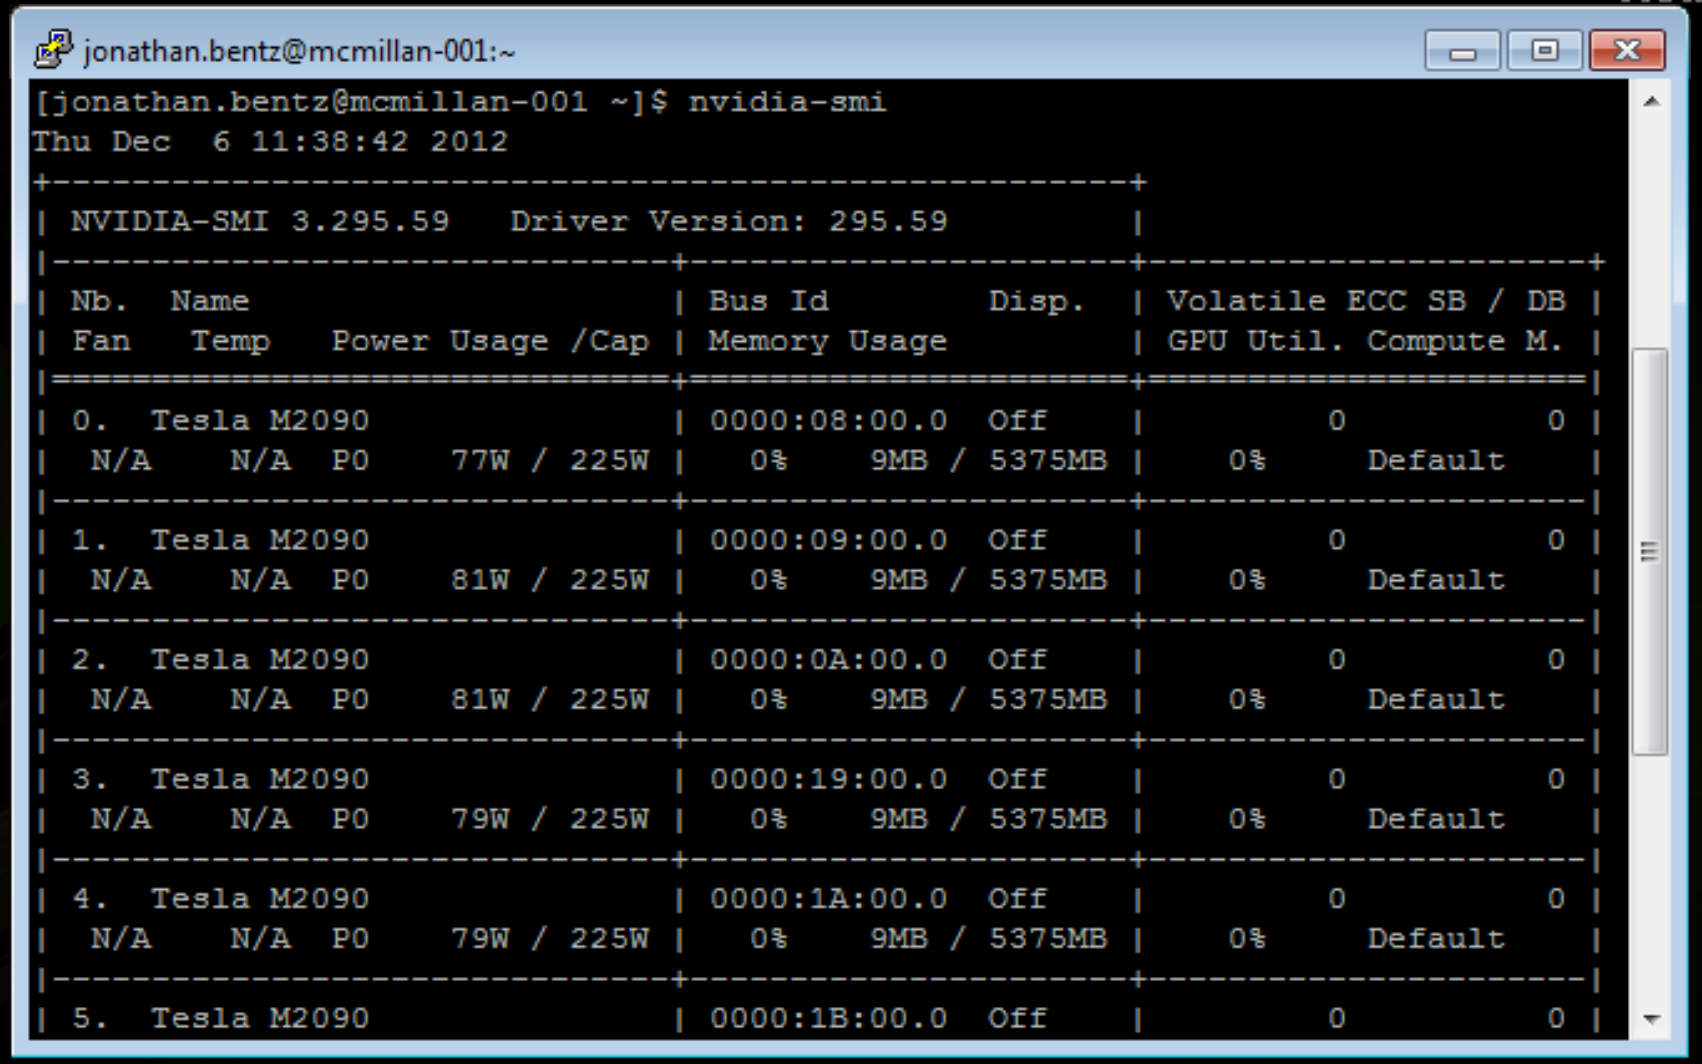
jonathan.bentz@mcmillan-001:~
[jonathan.bentz@mcmillan-001 ~]$ nvidia-smi
Thu Dec 6 11:38:42 2012
+--+
| NVIDIA-SMI 3.295.59 Driver Version: 295.59 |
|-------------------------------+----------------------+----------------------+
| Nb. Name | Bus Id Disp. | Volatile ECC SB / DB |
| Fan Temp Power Usage /Cap | Memory Usage | GPU Util. Compute M. |
|===============================+======================+======================|
| 0. Tesla M2090 | 0000:08:00.0 Off | 0 0 |
| N/A N/A P0 77W / 225W | 0% 9MB / 5375MB | 0% Default |
|-------------------------------+----------------------+----------------------|
| 1. Tesla M2090 | 0000:09:00.0 Off | 0 0 |
| N/A N/A P0 81W / 225W | 0% 9MB / 5375MB | 0% Default |
|-------------------------------+----------------------+----------------------|
| 2. Tesla M2090 | 0000:0A:00.0 Off | 0 0 |
| N/A N/A P0 81W / 225W | 0% 9MB / 5375MB | 0% Default |
|-------------------------------+----------------------+----------------------|
| 3. Tesla M2090 | 0000:19:00.0 Off | 0 0 |
| N/A N/A P0 79W / 225W | 0% 9MB / 5375MB | 0% Default |
|-------------------------------+----------------------+----------------------|
| 4. Tesla M2090 | 0000:1A:00.0 Off | 0 0 |
| N/A N/A P0 79W / 225W | 0% 9MB / 5375MB | 0% Default |
|-------------------------------+----------------------+----------------------|
| 5. Tesla M2090 | 0000:1B:00.0 Off | 0 0 |

CUDA toolkit

- `module load cudatoolkit/4.2.9`
  - `/usr/local/cudatoolkit/4.2.9`
- Compiler (nvcc)
- Libraries
  - BLAS, FFT, sparse, RNG, NPP, OpenCL
- Profiler
  - Visual or command-line profiling available.

SDK (free download from nvidia.com)

- Sample programs to illustrate CUDA and OpenGL programming constructs and algorithms.
- Useful diagnostic tests as well to query the GPU and its performance

SDK bandwidthTest

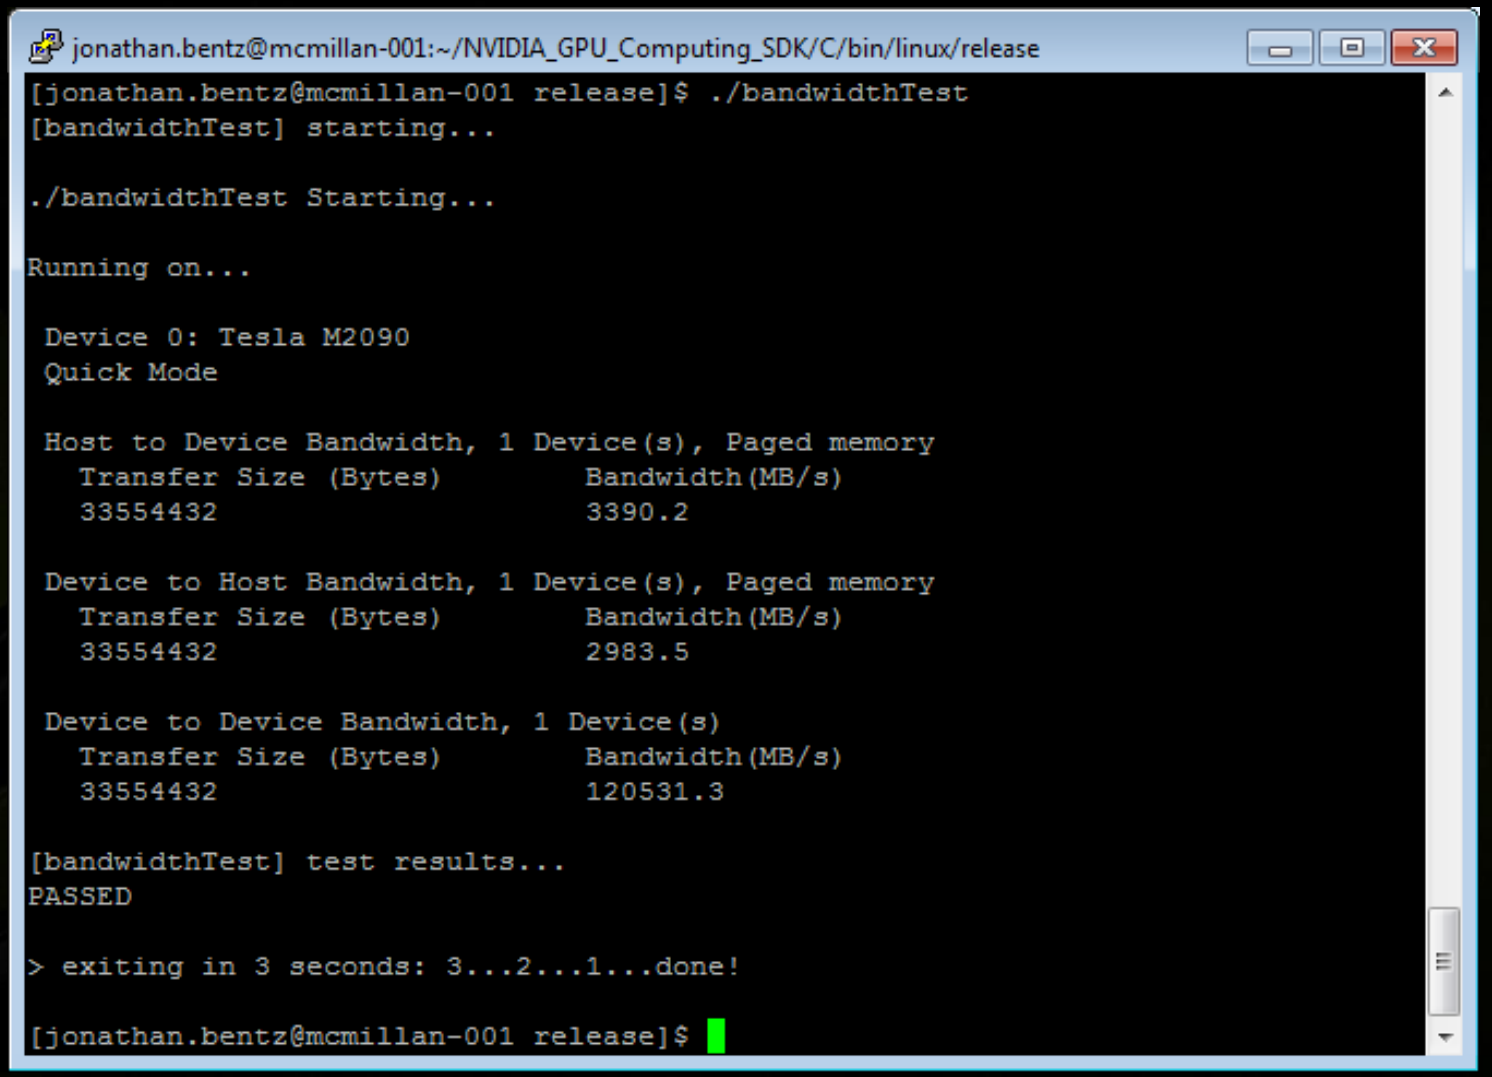

bandwidthTest --memory=pinned

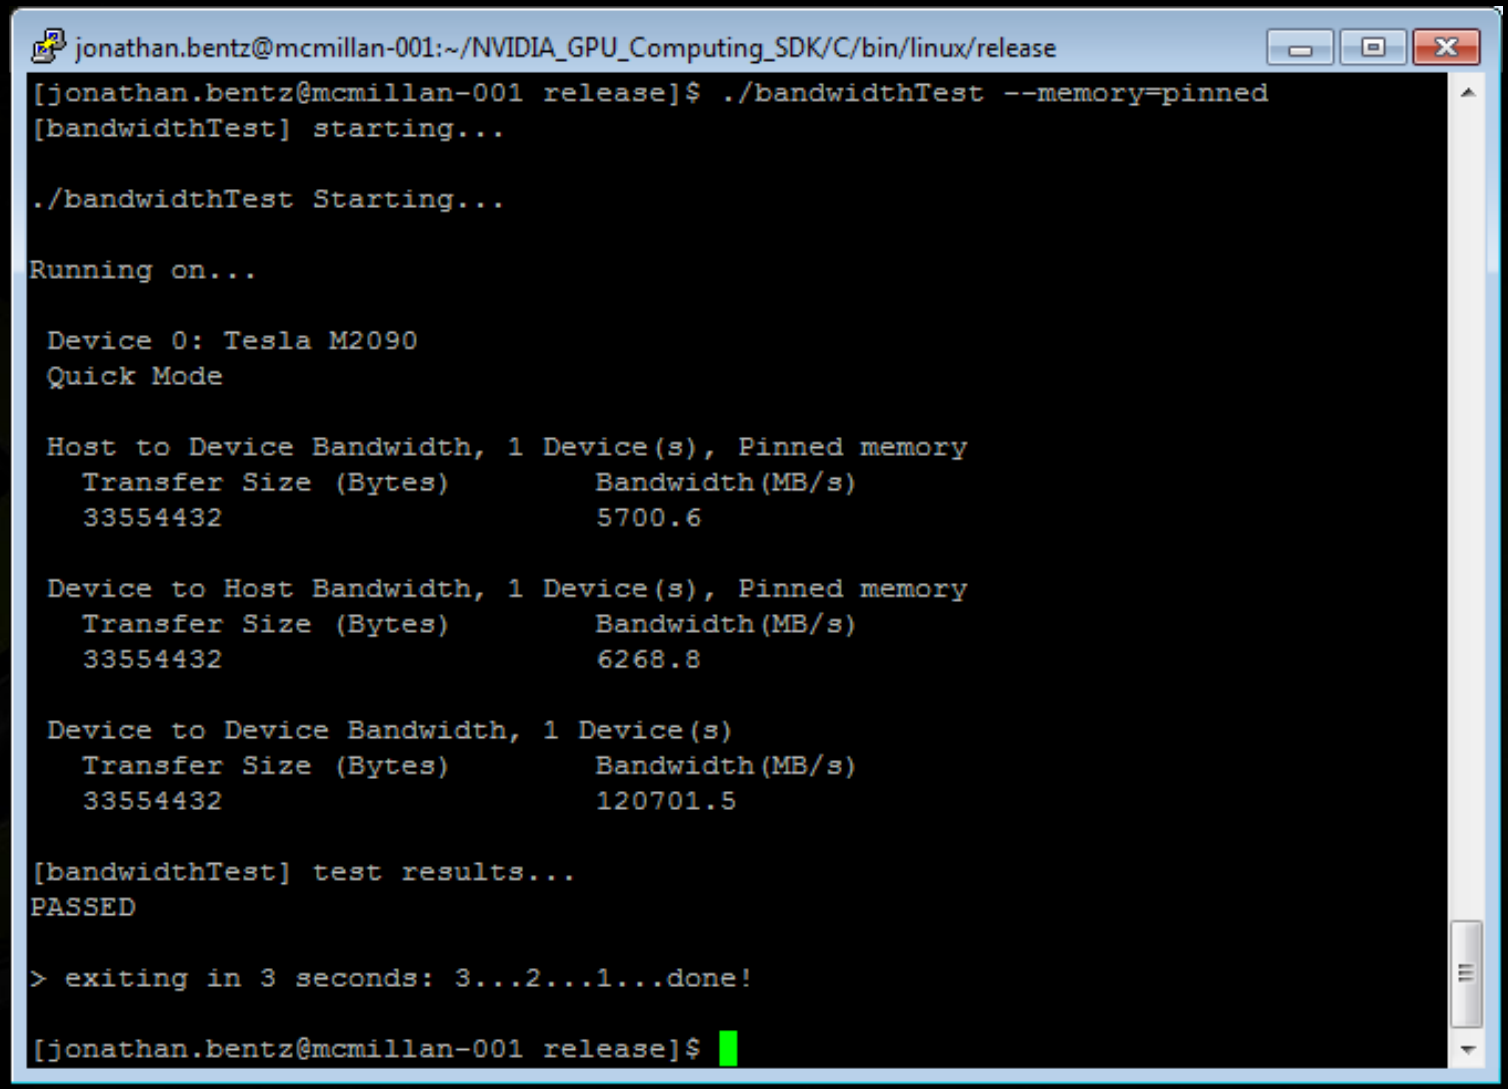

SDK deviceQuery

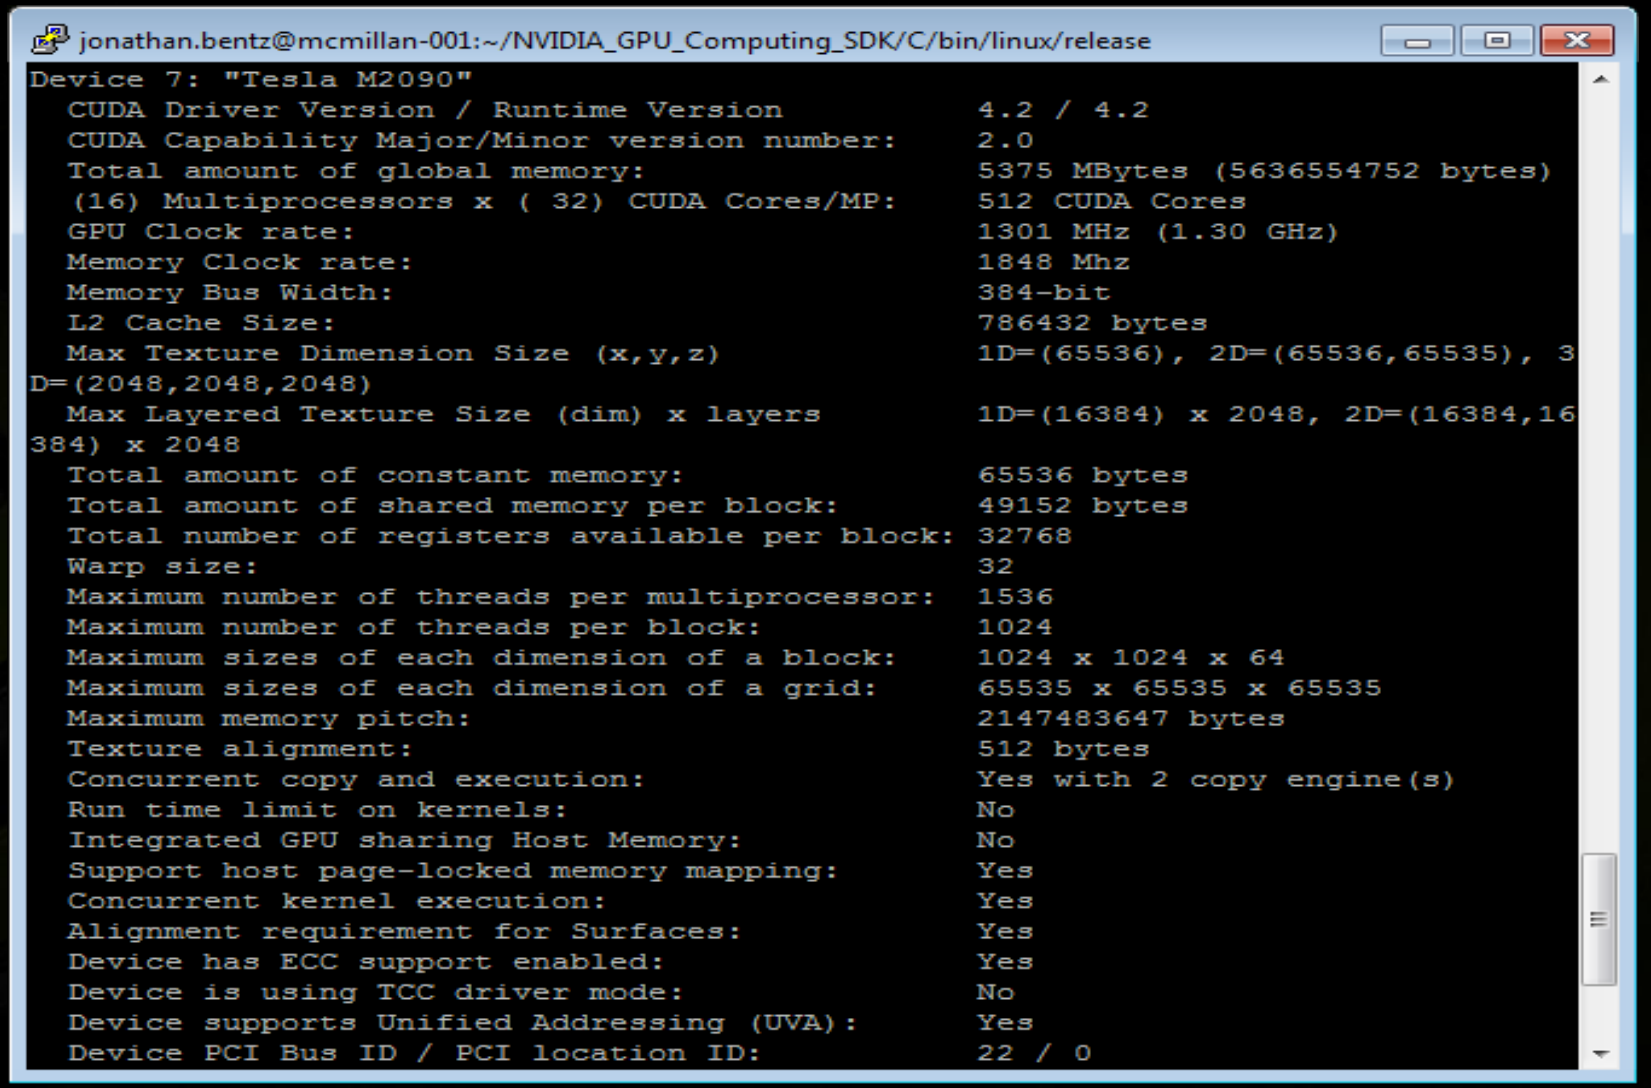

matrixMul

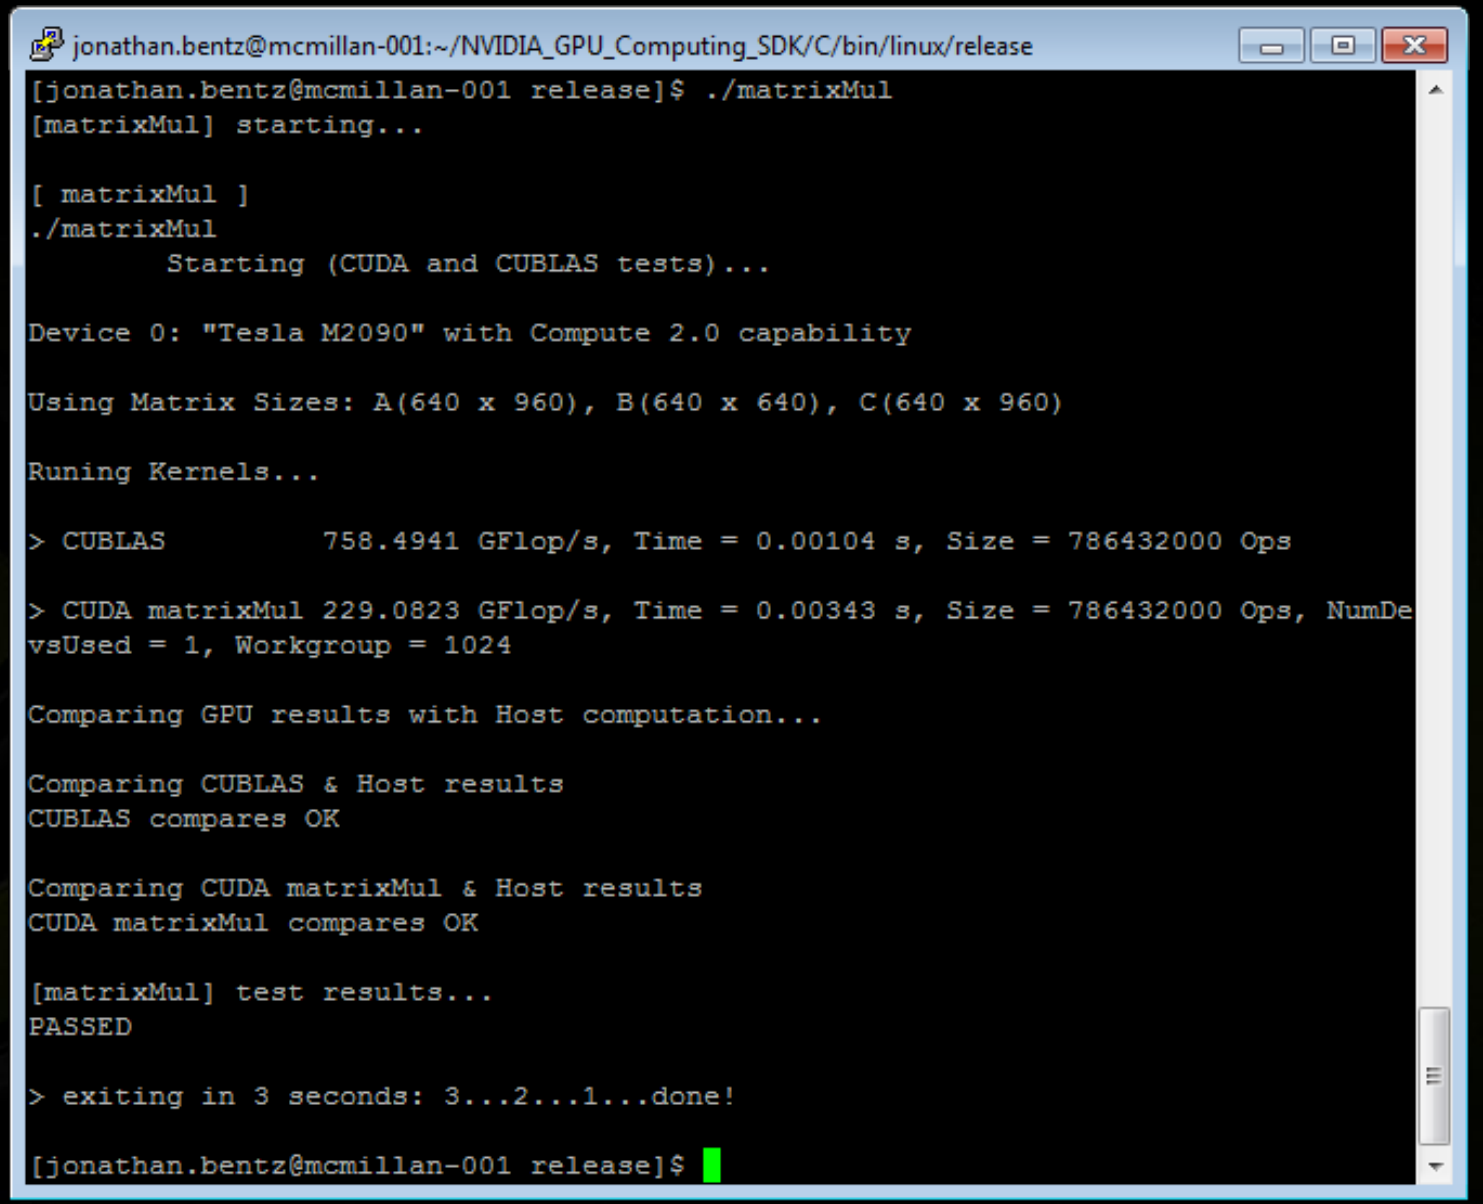
jonathan.bentz@mcmillan-001:~/NVIDIA_GPU_Computing_SDK/C/bin/linux/release
[jonathan.bentz@mcmillan-001 release]$ ./matrixMul
[matrixMul] starting...
[matrixMul]
./matrixMul
Starting (CUDA and CUBLAS tests)...
Device 0: "Tesla M2090" with Compute 2.0 capability
Using Matrix Sizes: A(640 x 960), B(640 x 640), C(640 x 960)
Runing Kernels...
> CUBLAS 758.4941 GFlop/s, Time = 0.00104 s, Size = 786432000 Ops
> CUDA matrixMul 229.0823 GFlop/s, Time = 0.00343 s, Size = 786432000 Ops, NumDe
vsUsed = 1, Workgroup = 1024
Comparing GPU results with Host computation...
Comparing CUBLAS & Host results
CUBLAS compares OK
Comparing CUDA matrixMul & Host results
CUDA matrixMul compares OK
[matrixMul] test results...
PASSED
> exiting in 3 seconds: 3...2...1...done!
[jonathan.bentz@mcmillan-001 release]$

Recap

- Driver
  - nvidia-smi to query the GPU hardware and state
- Toolkit
  - Development tools for GPU programming
- SDK
  - Sample code as well as diagnostic tests