

CUDA 10 – Jetson

L'image disque fournie par NVIDIA comprend CUDA 10.

Compiler les exemples :

```
$ cd /usr/local/cuda-10.0/samples
```

```
$ sudo GLPATH=/usr/lib make -j4
```

Patienter 20 minutes.

Vérification :

```
$ cd l*/deviceQuery
```

```
$ ./deviceQuery
```

```
Fichier Édition Affichage Rechercher Terminal Aide
ms@nano:/usr/local/cuda-10.0/samples/1_Uutilities/deviceQuery$ ./deviceQuery
./deviceQuery Starting...

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

Detected 1 CUDA Capable device(s)

Device 0: "NVIDIA Tegra X1"
  CUDA Driver Version / Runtime Version      10.0 / 10.0
  CUDA Capability Major/Minor version number: 5.3
  Total amount of global memory:             3956 MBytes (4148543488 bytes)
  ( 1) Multiprocessors, (128) CUDA Cores/MP: 128 CUDA Cores
  GPU Max Clock rate:                       922 MHz (0.92 GHz)
  Memory Clock rate:                        13 Mhz
  Memory Bus Width:                         64-bit
  L2 Cache Size:                            262144 bytes
  Maximum Texture Dimension Size (x,y,z)    1D=(65536), 2D=(65536, 65536), 3D=(4096, 4096, 4096)
  Maximum Layered 1D Texture Size, (num) layers 1D=(16384), 2048 layers
  Maximum Layered 2D Texture Size, (num) layers 2D=(16384, 16384), 2048 layers
  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: 2048
  Maximum number of threads per block:      1024
  Max dimension size of a thread block (x,y,z): (1024, 1024, 64)
  Max dimension size of a grid size (x,y,z): (2147483647, 65535, 65535)
  Maximum memory pitch:                    2147483647 bytes
  Texture alignment:                       512 bytes
  Concurrent copy and kernel execution:     Yes with 1 copy engine(s)
  Run time limit on kernels:                Yes
  Integrated GPU sharing Host Memory:       Yes
  Support host page-locked memory mapping:  Yes
  Alignment requirement for Surfaces:       Yes
  Device has ECC support:                   Disabled
  Device supports Unified Addressing (UVA):  Yes
  Device supports Compute Preemption:       No
  Supports Cooperative Kernel Launch:       No
  Supports MultiDevice Co-op Kernel Launch: No
  Device PCI Domain ID / Bus ID / location ID: 0 / 0 / 0
  Compute Mode:
    < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 10.0, CUDA Runtime Version = 10.0, NumDevs = 1
Result = PASS
ms@nano:/usr/local/cuda-10.0/samples/1_Uutilities/deviceQuery$ _
```