

GPU Tutorial Oct 31 - Nov 2 2023

Preparation for Using the NVIDIA DLI Training Cluster

please setup your account on the NVIDIA training cluster as follows

1. Create or log into your NVIDIA Developer Program account at <https://developer.nvidia.com/developer-program> . This account will provide you with access to all of the DLI training materials during and after the workshop.
2. Visit websocketstest.courses.nvidia.com and make sure all three test steps are checked "Yes." This will test the ability of your system to access and deliver the training contents. If you encounter issues, try updating your browser. Note: Only Chrome and Firefox are supported.
3. Go to URL <https://courses.nvidia.com/dli-event>
4. Enter event code that was distributed by email.

For the session on day 3, devoted to GPU programming with Python, follow [these](#) instructions to install the Python packages into your DLI Jupyter.

Preparation for Using Your Own System

You don't need access to the NVIDIA Training cluster if you have access to a machine with an NVIDIA GPU and the CUDA SDK installed. You can run the exercises yourself, either by hand issuing the appropriate **nvcc** commands, or by running a Jupyter notebook.

Jupyter Notebooks run from a web browser and connect to the server via port 8888, so if you want to try to run one locally, you will need for this port to be open. You can change the port number if needed.

To install the Jupyter notebook infrastructure simply do

```
pip install notebook
```

To test your installation, you can download the Jupyter notebook for the tutorial from [here](#)

Make sure **nvcc** is in your **\$PATH**. Untar the notebook you just downloaded and run it:

```
tar -xzf JupyterNotebook.tgz
cd CUDA
jupyter notebook START_HERE.ipynb
```

this should launch a web browser window with the start page of the notebook.

If you want to try running the exercises by hand, try the following after unpacking the above notebook and ensuring that **nvcc** is in your path:

```
cd CUDA-labs/CUDA/01_Introduction_to_CUDA/solutions
nvcc -o hello hello.cu
./hello
```

This should print **Hello World**.

For the session on day 3, devoted to GPU programming with Python, you can install the Python packages on your own computer by creating an environment in [your conda distribution](#) with this [environment.yml](#) or this [environment.lock.yml](#), followed by installing Jupyter in the environment.

```
mamba env create environment.yml  
mamba install jupyterlab
```

```
jupyter lab
```

(JupyterLab needs to be in the same environment with the Python packages; the installation via JupyterNotebook.tgz installs it outside of the conda environment.)

If you have any trouble installing the software on your computer, use the DLI system!