



udocker - *be anywhere*

Part 4 - Hands On: submission to SLURM clusters

<https://github.com/indigo-dc/udocker>

Mario David david@lip.pt, Jorge Gomes jorge@lip.pt



Before the beginning (slide deck 2)

Access the ACNCA (former INCD) advanced computing facility at Lisbon using ssh:

```
ssh -l <username> cirrus.a.acnca.pt  
module load python
```

- The end user can download and execute `udocker` without system administrator intervention.
- Install from a released version: <https://github.com/indigo-dc/udocker/releases>:

```
wget https://github.com/indigo-dc/udocker/releases/download/1.3.17/udocker-1.3.17.tar.gz  
tar zxvf udocker-1.3.17.tar.gz  
export PATH=$HOME/udocker-1.3.17/udocker:$PATH
```

In the beginning - I

Make a directory for the tutorial and set the environment variable of `udocker` to that dir:

```
mkdir udocker-tutorial
cd udocker-tutorial
export UDOCKER_DIR=$HOME/udocker-tutorial/.udocker
udocker version
```

Check that the `UDOCKER_DIR=$HOME/udocker-tutorial/.udocker` was created

```
echo $UDOCKER_DIR
ls -al $UDOCKER_DIR
```

In the beginning - II

I assume that the compute/worker nodes mount your \$HOME directory or, you can do this in some directory mounted in the compute/worker nodes.

Git pull the repository to get the necessary input files, for the tensorflow/keras application:

```
git clone https://github.com/LIP-Computing/tutorials.git
```

In particular, you will need the files and scripts in `tutorials/udocker-files/`

```
cp -r tutorials/udocker-files .
```

Pull a nice image

```
udocker pull tensorflow/tensorflow:2.20.0-gpu
```

First we create and prepare the container, later we run the actual job, the creation of the container may take some time (a few minutes), thus we do it once initially. And we can use some fast/low resource queue.

Modify the script `udocker-files/prep-keras.sh` to suit your slurm options and partition settings:

Submit job to create the container

In general just submit this script to slurm, we assume using GPU partition:

```
cd udocker-files; chmod 755 prep-keras.sh # if needed  
sbatch prep-keras.sh
```

Check job status with `squeue`

Creates the container and setup exec mode

Creating a container:

```
udocker create --name=tf_gpu tensorflow/tensorflow:2.20.0-gpu
```

Set the nvidia mode:

```
udocker setup --nvidia --force tf_gpu
```

Check the output of the slurm job `cat slurm-NNNN.out`

Run the container

Check the script `udocker-files/run-keras.sh` and modify the slurm options and partition as you see fit:

```
sbatch run-keras.sh
```

The script executes:

```
udocker run -v $TUT_DIR/udocker-files/tensorflow:/home/user -w /home/user tf_gpu python3 keras_2_small.py
```


Job output of tensorflow run

And, if all goes well you should see in the keras-xxx.out something like this:

```
#####  
Downloading data from https://storage.googleapis.com/tensorflow/tf-keras-datasets/mnist.npz  
11490434/11490434 [=====] - 1s 0us/step  
Epoch 1/5  
1875/1875 [=====] - 6s 2ms/step - loss: 0.2912 - accuracy: 0.9153  
Epoch 2/5  
1875/1875 [=====] - 3s 2ms/step - loss: 0.1427 - accuracy: 0.9574  
Epoch 3/5  
1875/1875 [=====] - 3s 2ms/step - loss: 0.1063 - accuracy: 0.9678  
Epoch 4/5  
1875/1875 [=====] - 3s 2ms/step - loss: 0.0890 - accuracy: 0.9721  
Epoch 5/5  
1875/1875 [=====] - 3s 2ms/step - loss: 0.0762 - accuracy: 0.9765  
313/313 - 1s - loss: 0.0769 - accuracy: 0.9771 - 594ms/epoch - 2ms/step
```

And now Gromacs

- I have a tarball that I built with docker from a Dockerfile in part 3 of this tutorial: `gromacs.tar`.
- It was saved with:
 - `docker save -o gromacs.tar gromacs`
- Now we will load the tarball with `udocker`:
 - `udocker load -i gromacs.tar gromacs`

Gromacs image in **udocker**

```
udocker images
```

```
REPOSITORY
```

```
gromacs:latest      .
```

```
tensorflow/tensorflow:2.20.0-gpu    .
```

Check in the filesystem:

```
ls -al $HOME/udocker-tutorial/.udocker/repos
```

```
total 16
```

```
drwxr-x---+ 4 david csys 4096 jan 20 17:38 .
```

```
drwxr-x---+ 8 david csys 4096 jan 20 16:54 ..
```

```
drwxr-x---+ 3 david csys 4096 jan 20 17:38 gromacs
```

```
drwxr-x---+ 3 david csys 4096 jan 20 17:04 tensorflow
```

Submit job to create container

```
cd udocker-files; chmod 755 prep-gromacs.sh # if needed  
sbatch prep-gromacs.sh
```

Submit Gromacs job

Prepare input dir and file

```
mkdir -p $HOME/udocker-tutorial/gromacs/input $HOME/udocker-tutorial/gromacs/output  
cd $HOME/udocker-tutorial/gromacs/input/  
wget --no-check-certificate https://download.a.acnca.pt/webdav/gromacs-input/md.tpr
```

```
sbatch run-gromacs.sh
```

Job output of Gromacs run

The Gromacs output files can be found in `$HOME/udocker-tutorial/gromacs/output`, and the slurm job output in `$HOME/udocker-tutorial/udocker-files/gromacs-*.out/err`

End of Hands On part III

