

SLURM Demo

1 Introduction

This is a simple tutorial on the workload manager SLURM on Linux platform. A complete manual can be referred to <https://slurm.schedmd.com>. In case of any problems, you may contact me via email: linkailuo@hsmc.edu.hk or Ext: 632.

The demo is based on Google's deep learning framework TENSORFLOW. The SLURM usage can be directly applied to other framework.

2 Procedure

1. Write your code. Our code simple sum up two variable and return a value.

```
# Import tensorflow
import tensorflow as tf

# Define two constant variable, a and b
a = tf.constant(1)
b = tf.constant(2)

# Create a tensorflow session to do the computation
sess = tf.Session() :

# Do the summation
c = sess.run(a+b)

# Print out message
print('{} + {} = {}'.format(sess.run(a),sess.run(b),c))
print('Test terminated')

# Close session
sess.close()
```

2. Construct shell script.

```
#!/bin/bash # use bash to interpret
```

```
#SBATCH --nodes=1      # number of nodes
#SBATCH --ntasks=1     # number of CPU (i.e., thread)
#SBATCH --workdir=/data/linkailuo/test_slurm # working directory
#SBATCH --time=12:00:00 # allocate time
#SBATCH --begin=now # when to begin the task
#SBATCH --mem=100000M  # how many memory to allocate to each node
#SBATCH --gres=gpu:tesla:1 # how many GPU to allocate to each node
#SBATCH --output=job-series.%j.out # job standard output, %j means job ID

source /usr/local/tensorflow/bin/activate # activate tensorflow

srun -N 1 -n 1 --gres=gpu:tesla:1 python3 -m test # run the task, "python3 -m test" is t
```

Then save your script as .sh file, e.g., **job.sh**.

3. Run your save task: **sbatch job.sh** A message will show “Submitted batch job 352”, and a file “job-series.352.out” will be generated once the task is completed, in which information is stored. The job ID may vary.