

GPU computing at TeideHPC

The TeideHPC cluster has a number of nodes that have NVIDIA general purpose graphics processing units (GPGPU) attached to them. It is possible to use CUDA tools to run computational work on them and in some use cases see very significant speedups.

As we explained in the [how to run jobs section](#) we can use 3 different ways for sending a job to the job queue: using an interactive session, launching the application in real time or by means of an execution script.

GPUs on Slurm

To request a single GPU on slurm just add `#SBATCH --gres=gpu` to your submission script and it will give you access to a GPU. To request multiple GPUs add `#SBATCH --gres=gpu:n` where 'n' is the number of GPUs.

So if you want 1 CPU and 2 GPUs from our general use GPU nodes in the 'gpu' partition, you would specify:

```
#SBATCH -p batch
#SBATCH -n 1
#SBATCH --gres=gpu:2
```

If you prefer to use *interactive session* you can use:

```
salloc -p express --mem 8000 --gres=gpu:1
```

While on GPU node, you can run ***nvidia-smi*** to get information about the assigned GPU's.

Specifying GPU Type or MIG partition to use.

The GPU models currently available on our cluster can be found [here](#) but as we explain in [MIG section](#) we can specify gpu type or MIG partition to use. There are **two methods** that can be used.

Visit [request GPU and compute resources](#) page for a detailed explanation:

Select GPU using `--constraint=`

```
salloc -p express --mem 8000 --constrains=gpu,a100
```

Select GPU using `--gres=gpu:model:1` or `--gres=gpu:mig-partition:1`

Note that `--gres` specifies the resources on a per node basis, so for multinode work you only need to specify how many gpus you need per node.

List of model of GPUs and partitions

To find out what specific types of gpu's are available on a partition run `scontrol show partition` and look under the **TRES category**.

```
scontrol show partition express
```

```
PartitionName=express
...
MaxNodes=UNLIMITED MaxTime=03:00:00 MinNodes=0 LLN=NO
MaxCPUsPerNode=UNLIMITED
Nodes=node0303-2,node0304-[1-4],node1301-[1-4],node1302-[1-4],node1303-[1-4],
...
State=UP TotalCPUs=2424 TotalNodes=88 SelectTypeParameters=NONE
...
DefMemPerNode=UNLIMITED MaxMemPerNode=UNLIMITED
TRES=cpu=2424,mem=7565306M,node=88,billing=2424,gres/gpu=79,gres/gpu:1g.5gb=2,gres/gpu:2g.10gb=1,gres/gpu:3g.20gb=1,gres/gpu:a100=71,gres/gpu:t4=4
```

Nvidia A100

gpu:a100

a100-mig

- gpu:1g.5gb
- gpu:2g.10gb
- gpu:3g.20gb
- ...

Nvidia Tesla T4

- gpu:t4

Job script example for GPU job:

- Full GPU Nvidia A100

```
#!/bin/bash

#SBATCH --partition=batch
#SBATCH --nodes=1
#SBATCH --ntasks=1
#SBATCH --cpus-per-task=4
#SBATCH --gres=gpu:a100:1
#SBATCH --mem=8G
#SBATCH --time=1:00:00

module purge
module load CUDA/12.0.0

nvidia-smi
sleep 20
```

```
sbatch 01_gpu_basic_a100.sbatch
```

- MIG partition in Nvidia A100

```
#!/bin/bash

#SBATCH --partition=batch
#SBATCH --nodes=1
#SBATCH --ntasks=1
#SBATCH --cpus-per-task=4
#SBATCH --gres=gpu:2g.10gb:1
#SBATCH --mem=16G
#SBATCH --time=1:00:00

module purge
module load NVHPC/22.11-CUDA-11.8.0

nvidia-smi
sleep 20
```

```
sbatch 02_gpu_basic_mig_partition.sbatch
```

More examples.

Visit our repository in github https://github.com/hpciter/user_codes