Using GPUs with --gres

Now that you have a handle on basic and more advanced jobs, it's time to get even more advanced. Some researchers require GPUs for their work, and luckily, that is really easy to do with Slurm. We currently have 16 Nvidia V100 cards and 96 Nvidia P4 cards jobs can be run on. CUDA is available through either Spack or module if it is needed.

What's Available?

If you are using GPUs through SPORC you will have **one V100** code or **two P4** cards per node.

If you are using sinteractive with theOcho you have access to **one V100** and **one rtx6000**.

Formatting the Command

The general form for --gres is as follows:

```
--gres=gpu:type:count
```

where type is the name of the kind of GPU you want and count is the number of GPUs of that type you want to request. For example:

```
--gres=gpu:v100:1
```

With this command, your job will run on one V100 GPU. The first available GPU will be used if you do not specify a type . Everything else about the SBATCH file is the same

Example

Coming Soon....