

How to run parallel job's with srun

srun : Used to submit/initiate job or job step

Typically, srun is invoked from a SLURM job script but alternatively, srun can be run directly from the command, in which case srun will first create a resource allocation for running the parallel job (the salloc is implicit)

```
srun -N 1 -c 16 -p HPC_4_Days --time=1:00:00 --pty /bin/bash
```

This command will request 16 cores (-c) of one Node (-N) for 1h00 in the partition (-p) HPC_4_Days. Please note that this is subject to Nodes availability, if no Nodes are available your request will be put in the queue waiting for resources.

The srun may also be executed inside a shell script.

```
#!/bin/bash

#SBATCH -N 3
#SBATCH -p HPC_4_Days

echo Starting job $SLURM_JOB_ID
echo SLURM assigned me these nodes
srun -l hostname
```

This batch job will result in the following output:

```
Starting job 51057
SLURM assigned me these nodes
0: wn054.b.incd.pt
1: wn055.b.incd.pt
2: wn057.b.incd.pt
```

The 3 allocated nodes are released after the `srun` finish.

By default `srun` will use the `pmi2`, but you may consult the full list of the available mpi types.

```
$ srun --mpi=list
```

```
srun: MPI types are...
```

```
srun: pmi2
```

```
srun: openmpi
```

```
srun: none
```

To use a different mpi type e.g. `srun --mpi=openmpi`

“ For more detailed information, please see `man srun`

Revision #7

Created 29 November 2019 17:58:21 by Joao Machado

Updated 16 March 2023 15:23:38 by João Pina