SLURM JOB SCRIPT CHEAT SHEET
GENERIC/ADVANCED

```
#!/bin/bash

#SBATCH --job-name=myjobname                    #Your Job Name
#SBATCH --nodes=1                               #Number of Nodes desired e.g 1 node
#SBATCH --time=00:10:00                         #Walltime: Duration for the Job to run HH:MM:SS
#SBATCH --mail-user=useremail@kaust.edu.sa      #Your Email address assigned for your job
#SBATCH --mail-type=ALL                         #Receive an email for ALL Job Statuses
#SBATCH --error=JobName.%J.err                  #The .error file name
#SBATCH --output=JobName.%J.out                 #The .output file name

#Go to your working directory
cd /my_working_dir/

#Module load the desired application if necessary
module load module_name                         #Always check the module needed on the login node "module avail"

#Edit below with the launching command:
your_commands_goes_here
```

Additional Options that can be added to your job script:

```
#SBATCH --exclusive                             #Use only for runs needing the whole node exclusively
#SBATCH --constraint={constraint}               #Run on a specific type of nodes
#SBATCH --ntasks-per-node=4                     #Number of tasks to run per node
#SBATCH --cpus-per-task=4                       #Number of cores assigned per task
#SBATCH --partition=batch                       #Partition name default for Ibex
#SBATCH --mem=2GB                               #Memory requested for e.g 2GB
#SBATCH --reservation=RESNAME                   #Specify your reservation node if any
#SBATCH --gres=gpu:1                            #Run on 1 GPU of any type
#SBATCH --gres=gpu:<GPU_type>:<number>          #Check for advanced tips on GPU constraints
#SBATCH --array=1-10                            #For Job Arrays
```

To submit a Job: *sbatch myjobscript*

To cancel a job: *scancel jobid*

To check the status of your jobs: *squeue –u username*

TIPS:

- Best practice is to use the Ibex Job generator : https://www.hpc.kaust.edu.sa/ibex/job
- Check with the system's team if you need to extend your job wall time.
- For more info on SLURM check their website: https://slurm.schedmd.com/

SLURM JOB SCRIPT CHEAT SHEET
GENERIC/ADVANCED

Here are some advanced tips for Ibex Users:

Intel & AMD Specific Constraints for different node types:			
CPU Family	#Nodes	#Cores per node	Constraint
AMD	23	64	#SBATCH –constraint=amd
Intel Skylake	105	40	#SBATCH –constraint=cpu_intel_gold_6148
Intel Ivybridge	174	20	#SBATCH --constraint=[cpu_intel_e5_2680_v2|cpu_intel_e5_2670_v2]
Intel SandyBridge	96	16	#SBATCH --constraint=[cpu_intel_e5_2670]

Large memory info for different node sizes **:** Just add the memory desired and automatic allocation will occur #SBATCH –mem=##			
CPU Family	Cores per node	Available number of nodes	Recommended max memory per node
AMD Abu Dhabi	64	3	755 GB
		6	995 GB
		1	1.40 TB
		3	1.50 TB
Intel Westmere	64	1	1.84 TB
		1	1.9 TB
Intel Westmere	80	1	1.9 TB
Intel Skylake	32	4	2.93 TB

GPU Specific Constraints for different node types:			
Description	Available GPU cards per node	Available number of nodes	Constraint
Kepler: K40m	8	3	#SBATCH --gres=gpu:tesla_k40m:1
Fermi:gtx1080ti	4	8	#SBATCH --gres=gpu:gtx1080ti:1
Pascal:p100	4	6	#SBATCH --gres=gpu:p100:1
Pascal:p6000	2	2	#SBATCH --gres=gpu:p6000:1
Volta:v100	4	8	#SBATCH --gres=gpu:v100:1
Fermi:gtx1080ti	8	4	#SBATCH --gres=gpu:gtx1080ti:1

Contact Us: https://kaust-ibex.slack.com/ - Use #general for simple queries

Open a ticket by sending an email to: ibex@hpc.kaust.edu.sa