IBEX CHEAT SHEET

Ibex is a heterogeneous cluster with a mix of AMD, INTEL and NVIDIA GPUs.

To Login:

Intel nodes:

```
ssh -X <UserName>@ilogin.ibex.kaust.edu.sa
```

AMD nodes:

```
ssh -X <UserName>@alogin.ibex.kaust.edu.sa
```

GPU nodes (non-Volta):

```
ssh -X <UserName>@glogin.ibex.kaust.edu.sa
```

Intel Skylake nodes:

```
ssh -X <UserName>@slogin.ibex.kaust.edu.sa
```

GPU nodes (Volta):

```
ssh -X <UserName>@vlogin.ibex.kaust.edu.sa
```

Application installation :

All compilers, libraries and applications are installed on each login node due to variation in the system architecture. Intel, AMD and GPU based architecture specific applications are available through modules.

Application availability:

```
$module avail
$module avail <ApplicationName>
```

Application loading:

```
$module load <ApplicationName>
$module load <ApplicationName>/<version>
```

Job Submission (batch mode):

To set memory requirement: `--mem=<in MB>`

To select architecture specific node type:
`--constraint=intel|amd`
`--gres=gpu:<$$$>:<#>”`, where: <$$$> is the GPU architecture and <#> is for number of GPUs. For example, “--gres=gpu:gtx1080ti:4” is for 4 GTX GPUs

To set number of nodes: `--nodes`

To set number of tasks (for parallel processing): `--ntasks`

To set the number of core per tasks : `--cpus-per-task`

To set wall clock time: `--time`

To set the node as dedicated for the job: `--exclusive`

To set the file name for standard err: `--error`

To set the file name for standard out: `--output`

Tunable job script generator for IBEX is available in: https://www.hpc.kaust.edu.sa/ibex/job

Example Job Script:

```
#!/bin/bash
## SLURM Resource requirement:
#SBATCH --nodes=1
#SBATCH --cpus-per-task=8
#SBATCH --job-name=spades
#SBATCH -output=myjob.%J.out
#SBATCH --error=myjob.%J.err
#SBATCH --time=8:00:00

## Required software list:
  module load gaussian09/d.01/precompiled
## Run the application:
echo "This job ran on $SLURM_NODELIST dated
`date`";
  mpirun g09 < testgau.inp > testgau.out
```

Job Submission queues:

There are 2 queues, the default batch is for production runs and the debug is for interactive debugging the jobs.

To use debug queue (for example):

```
salloc --time=5:00 --nodes=1 \
--partition=debug
```

Other Slurm Commands:

`sbatch`: to run jobs
`sinfo`: to check node availability
`squeue`: to check job status
`scancel job#` : to cancel jobs

General Tips:

- Do to run on the logins nodes, always submit your jobs through scripts.
- Logins are designed for compilations and edits.
- Always run your jobs from the scratch.
- Remember to clean up your scratch.

Filesystem:

- `/home/<UserName>`: *Home directory for important data backup.*
- Always use the `/scratch` filesystem to *submit jobs from amd/intel/gpu* nodes.
- Use `/fscratch` if your jobs require a high number of IOPS.

Contact for Help/Support:

Application installation/failure/support:
cluster-apps@hpc.kaust.edu.sa
System issues/failure/support:
cluster-systems@hpc.kaust.edu.sa
Our website:
https://www.hpc.kaust.edu.sa/ibex