جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

KAUST SUPERCOMPUTING CORE LAB

Shaheen III GPU Partition

Getting Started

System Specifications

System	HPE Cray EX4000 with 7 cabinets
Nodes	700, each node with 4 GH200 Superchip
Total GPUs	2800 NVIDIA GH200 Superchips
Total ARM cores	201,600
Memory	120 GB/superchip, 480 GB/node
Interconnect	Slingshot-11
NIC	4 x Cassini 200 Gbps, 1 per superchip
Job scheduler	SLURM
Storage	Lustre parallel file-system Encrypted Scratch 20 PB

Login

Your login name starts with `s-`. You need to login `hpc-vpn` if you are trying to connect out side of campus area. Then you need to login Virtual Desktop Infrastructure (VDI): https://s3vdi.kaust.edu.sa Once you are in VDI, open a terminal and connect to `ulogin` and then `gh` as follows:

```
ssh -XY ulogin.hpc.kaust.edu.sa
ssh -XY gh
```

For any issues, please contact gpuhelp@hpc.kaust.edu.sa.

Compile Source Code

The following lines show how to compile CUDA code:

```
module swap PrgEnv-cray PrgEnv-nvidia
nvcc -o executable cudacode.cu
```

Scheduler and Queues

Run

SLURM is the batch scheduler. Here is a basic script for running a CUDA program on 1 GPU:

```
#!/bin/bash
#SBATCH --time=00:10:00
#SBATCH --gres=gpu:1
module switch PrgEnv-cray PrgEnv-nvidia
srun --ntasks=1 ./executable
```

The job script is submitted as follows:

```
sbatch job_script.sh
```

The following job script is used for MPI+CUDA workloads:

```
#!/bin/bash
#SBATCH --partition=workq
#SBATCH --nodes=2
#SBATCH --account=k2####
#SBATCH --time=01:00:00
#SBATCH --mail-type=ALL
#SBATCH --mail-user=emailAddress
export MPICH_GPU_SUPPORT_ENABLED=1
export OMP_NUM_THREADS=72
srun --ntasks=8 --cpus-per-task=${
    ↪ OMP_NUM_THREADS} --gpu-bind=verbose,map
    ↪ :0,1,2,3 --gres-flags=allow-task-sharing
    ↪ ./executable
```

Hyperthreading is enabled by default and might improve the performance of some codes. Use `--hint=nomultithread` option in the execution line to disable it.
Submit a job:

```
sbatch --gres=gpu:4 job_script.sh
```

List jobs:

```
squeue --me
```

Cancel a job:

```
scancel id_of_my_job
```

Queues

Use `sinfo` for the queue status.
`workq`: Default queue with a maximum 24 hours wall clock time.

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

KAUST SUPERCOMPUTING CORE LAB

Storage, Quotas, Allocations

Storage

	/scratch
Login nodes	RW
Compute nodes	RW

R:Read W:Write

- `/scratch/$USER`: Individual storage for data needed for execution.
- `/scratch/project/k2####`: Storage for the project.

Default quotas, capacity and performance:

	#files	Size
`/scratch/$USER`	1 M	10 TB
`/scratch/project/k2####`		80 TB*

* Per Principal Investigator (PI)

Find out your project groups:

```
groups
```

Check remaining core hours for a project:

```
sb k2####
```

Check project storage quota:

```
kpq k2####
```

Software & Libraries

Before requesting installation of software, please check if the desired software is already installed on the system:

```
module av
```

To find a specific software:

```
module av -S gro
```

To get information about the software usage:

```
module help gromacs
```

To get information about the software:

```
module show gromacs
```

Many modules provide example job scripts:

```
module show gromacs
/sw/ex113genoa/modulefiles/espresso/7.5:

module-whatis
-------------------------------
GROMACS is a versatile package to perform
    ↪ molecular dynamics, i.e. simulate ...
-------------------------------
...
setenv  MY_EXAMPLE /sw/ex113gracehopper/gromacs
    ↪ /2025.2/gcc13.2/example
-----------------------------------
```

The examples folder includes all necessary files, along with a SLURM job script to run a sample job:

```
ls -1 /sw/ex113gracehopper/gromacs/2025.2/gcc13
    ↪ .2/example/02
md_0_1.tpr
z_jobs_shaheen
```

A sample job can be launched as follows:

```
cd $SCRATCH
cp /sw/ex113gracehopper/gromacs/2025.2/gcc13.2/
    ↪ example/02/* .
sbatch z_jobs_shaheen
```

Contact Us

gpuhelp@hpc.kaust.edu.sa