Engineering DGX AI Cluster: saidie.engr.colostate.edu

Head node

2 x AMD EPYC 9354 32-Core Processor – with multithreading

504 gb RAM

1 compute node

2 TB of RAM

2x56 core processors multithreaded

8 h200 GPUs

```
+-----------------------------------------------------------------------------------------+
| NVIDIA-SMI 570.172.08             Driver Version: 570.172.08     CUDA Version: 12.8     |
|-----------------------------------------+------------------------+----------------------+
| GPU  Name                 Persistence-M | Bus-Id          Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf          Pwr:Usage/Cap |           Memory-Usage | GPU-Util  Compute M. |
|                                         |                        |               MIG M. |
|=========================================+========================+======================|
|   0  NVIDIA H200                    On  |   00000000:1B:00.0 Off |                    0 |
| N/A   36C    P0             77W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   1  NVIDIA H200                    On  |   00000000:43:00.0 Off |                    0 |
| N/A   37C    P0             79W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   2  NVIDIA H200                    On  |   00000000:52:00.0 Off |                    0 |
| N/A   39C    P0             79W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   3  NVIDIA H200                    On  |   00000000:61:00.0 Off |                    0 |
| N/A   38C    P0             75W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   4  NVIDIA H200                    On  |   00000000:9D:00.0 Off |                    0 |
| N/A   36C    P0             77W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   5  NVIDIA H200                    On  |   00000000:C3:00.0 Off |                    0 |
| N/A   35C    P0             78W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   6  NVIDIA H200                    On  |   00000000:D1:00.0 Off |                    0 |
| N/A   37C    P0             78W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+
|   7  NVIDIA H200                    On  |   00000000:DF:00.0 Off |                    0 |
| N/A   38C    P0             76W /  700W |       0MiB / 143771MiB |      0%      Default |
|                                         |                        |             Disabled |
+-----------------------------------------+------------------------+----------------------+

+-----------------------------------------------------------------------------------------+
| Processes:                                                                              |
|  GPU   GI   CI              PID   Type   Process name                        GPU Memory |
|        ID   ID                                                               Usage      |
|=========================================================================================|
|  No running processes found                                                             |
+-----------------------------------------------------------------------------------------+
```

CLUSTER NOTES

- The head node is for submitting jobs, testing and working out your code. DO NOT run your job on the head node, this is what the rest of the cluster is for. Any jobs running on the head node will be killed if found running without warning.
- Shared libraries will be in /cm/shared. This is mounted on the DGX as well.
- Your /home directory is mounted on the DGX. This is capped at 2 TB. Please put in an ETS ticket if you need more space.
- Saidie runs a module system.
    - You can check the current modules with ‘module avail’.

- The Slurm module is loaded by default when you log in.
- When running on the DGX, you will need to load the necessary modules that you need. Slurm will not automatically load them for you.

SLURM BASICS

Slurm is the job scheduler on saidie. All jobs must go through the queueing system. Keep in mind the following when running on saidie.

- If any jobs or processes are found to be running on the DGX outside the Slurm queuing system they will be killed without warning.
- There are two queues:
  -
```
root@saidie-c2:/home# sinfo
PARTITION AVAIL  TIMELIMIT  NODES  STATE NODELIST
general*     up 2-00:00:00      1   idle saidie-dgxl
long-runs    up 7-00:00:00      1   idle saidie-dgxl
root@saidie-c2:/home#
```
  - The general queue is capped at a total run time of 2 days.
  - The long-runs queue is capped at 7 days.
  - Everyone who submits a job needs a Slurm account to run. You will want to include the **#SBATCH --account=<XXX>** in your sbatch file. The account will be based on your PIs last name. For example:

Research Group	slurm account	ldap group
Debasish Jana	jana-saidie	hpc-jan
Yinshuang Xiao	xiao-saidie	hpc-xiao
Tim Hansen	hansen-saidie	hpc-han
Shirin Panahi	panahi-saidie	hpc-pan
Thomas Bradley	bradley-saidie	hpc-brad
John Volkens	volkens-saidie	hpc-vol

    - Jana's group will use jana-saidie as their account, for example.
    - This account is how we manage usage.
    - Everyone has a default account they will start with.
  - Both queues have the same priority, one is not better than the other.
  - The two queues are running on the same DGX. As such, they have the same resources to work with.
  - To get detailed information about the queues you can use the following command:
    - scontrol show partition long-runs

- ```
  aredman@saidie-c2:~/testSubmit$ scontrol  show partition long-runs
  PartitionName=long-runs
     AllowGroups=ALL AllowAccounts=testai,hansen-saidie,jana-saidie,panahi-saidie,wang-saidie,xiao-saidie AllowQos=ALL
     AllocNodes=ALL Alternate=general Default=NO QoS=N/A
     DefaultTime=NONE DisableRootJobs=NO ExclusiveUser=NO ExclusiveTopo=NO GraceTime=0 Hidden=NO
     MaxNodes=UNLIMITED MaxTime=7-00:00:00 MinNodes=1 LLN=NO MaxCPUsPerNode=UNLIMITED MaxCPUsPerSocket=UNLIMITED
     Nodes=saidie-dgx1
     PriorityJobFactor=10 PriorityTier=10 RootOnly=NO ReqResv=NO OverSubscribe=NO
     OverTimeLimit=NONE PreemptMode=OFF
     State=UP TotalCPUs=224 TotalNodes=1 SelectTypeParameters=NONE
     JobDefaults=DefCpuPerGPU=1
     DefMemPerNode=4000 MaxMemPerNode=UNLIMITED
     TRES=cpu=224,mem=1857527M,node=1,billing=224,gres/gpu=8,gres/gpu:h200=8
  ```
- Currently the DGX is set up to allow full use of all 8 GPUs. Please only request the minimum amount of GPUs that you need so others can run.
- Basic Slurm script
  - You will need to run **sbatch** with your batch script to submit your jobs.
    - sbatch <job script name>
  - This is the bare minimum you need. **Make sure the --account line is included.**
  - Ensure that you include the **#SBATCH --gpus-per-node=** in your submit script
  - If you do not include –mem, --cpus and –gpus-pre-node your job WILL take over the entire node. Even if you only need 1 of them. You will want to estimate how much memory your run takes. This might take a few tries.
  - Include a ‘module load’ command if using modules
  - There are many more SBATCH commands you can add to this: https://slurm.schedmd.com/sbatch.html

BASIC SLURM SUBMIT SCRIPT

```
==================================================
#!/bin/bash

#SBATCH --partition=general
#SBATCH --job-name=GPU hello world
#SBATCH --gpus-per-node=h200:1
#SBATCH --account=wang-saidie
#SBATCH --mem=100g
#SBATCH --cpus-per-task=10

module load <module name>
/home/aredman/testSubmit/helloWorld
sleep 120
/home/aredman/testSubmit/helloWorld
==================================================
```
```

Vim NOTE

*vim is installed on the cluster. For a better user experience, you can add the following to ~/.vimrc

==========

syntax on

set background=dark ##set dark or light depending on your terminal

set backspace=indent,eol,start

==========

*The middle line is the important one as you can change the background between 'light' and 'dark'.

GETTING HELP AND REPORTING ISSUES

If you encounter issues while running on saidie please submit an ETS freshservice ticket:
https://csu-ets.freshservice.com/support/home

SLURM – BASIC COMMANDS

Job scheduling commands

Commands	Function	Basic Usage	Example
sbatch	submit a slurm job	sbatch [script]	$ sbatch job.sub
scancel	delete slurm batch job	scancel [job_id]	$ scancel 123456
scontrol hold	hold slurm batch jobs	scontrol hold [job_id]	$ scontrol hold 123456
scontrol release	release hold on slurm batch jobs	scontrol release [job_id]	$ scontrol release 123456

Job management commands

Job Status	Commands
sinfo -a	list all queues
squeue	list all jobs
squeue -u userid	list jobs for userid
squeue -t R	list running jobs
smap	show jobs, partitions and nodes in a graphical network topology

*There are many more commands for Slurm. The Slurm documentation / quick start user guide may be found here: https://slurm.schedmd.com/quickstart.html .