

GPU Cluster Administration Guide

The GPU cluster is managed by Docker. The docker containers are monitored by Zabbix.

Overview of the GPU Cluster Logic

- Principal Investigators (PIs) are assigned one or more Docker containers. The computational resources of the server, including CPU cores, memory, and storage, are shared among all PIs, with the exception of GPU resources.
- Each PI is allocated a specific number of GPUs. This allocation is determined based on the PI's requirements and contributions to the cluster. These GPU resources are exclusive to each PI and are not shared with others. However, a PI may choose to distribute these GPU resources across their own multiple containers as needed.
- PIs have the privilege to request user accounts on the GPU cluster. These user accounts are granted access to the Docker containers owned by the corresponding PI, facilitating collaboration and resource utilization within the PI's team.

Please refer to <https://docs.docker.com/engine/install/ubuntu/> for Docker Installation.

Please refer to <https://www.zabbix.com/download> for Zabbix Installation.

Please refer to <https://github.com/plambe/zabbix-nvidia-smi-multi-gpu> for adding GPU support to Zabbix.

Administration Operations

Create A New Docker Container

```
docker run -it --gpus '"device=0,1"' --shm-size=32gb
```

```
tensorflow/tensorflow:latest-gpu
```

`device` can be updated with the index of GPU assigned to PI.

`shm-size` increases the shared memory size, which is typically useful for memory-intensive AI tasks.

`tensorflow/tensorflow:latest-gpu` can be updated with other docker images, e.g., `nvcv.io/nvidia/pytorch:20.06-py3`

Attach a disk volume to the container

```
docker run -it --gpus '"device=0,1"' -v /data:/data/  
--shm-size=32gb tensorflow/tensorflow:latest-gpu
```

Create a new docker container with a GUI-enabled Docker image.

```
docker run -it --gpus device=0 -p 2080:80 -p 5900:5900  
-e RESOLUTION=1920x1080 -e VNC_PASSWORD=NpIfGCWBJMYr  
dorowu/ubuntu-desktop-lxde-vnc
```

After creating the docker container, please provide the container ID to the users.

Adding A New User

1. Create a user account on the server

```
sudo adduser [username]
```

2. Provide access to docker operations, e.g., attach, start, restart containers.

```
sudo vi /etc/sudoers  
[username]    ALL=(ALL:ALL) NOPASSWD: /usr/bin/docker  
attach *,/usr/bin/docker cp *,/usr/bin/docker
```

```
start *,/usr/bin/docker restart *
[username]    ALL=(ALL:ALL) NOPASSWD: /usr/bin/docker cp *
[username]    ALL=(ALL:ALL) NOPASSWD: /usr/bin/docker start *
[username]    ALL=(ALL:ALL) NOPASSWD: /usr/bin/docker restart *
```

Monitor Cluster Performance with Zabbix

The cluster is continuously monitored using Zabbix.

The Zabbix server is installed on node0. Zabbix agents are installed on all the nodes within the cluster.

Detailed monitoring information can be accessed through the following URL:

<http://node0.research.mtu.edu/zabbix>

The administrator's information has been added to the Zabbix system.