

NVIDIA DGX A100

In publications and presentations that use results obtained on this system, please include the following acknowledgement: "This work utilizes resources supported by the National Science Foundation's Major Research Instrumentation program, grant #1725729, as well as the University of Illinois at Urbana-Champaign".

[Main -> Systems](#)
-> DGX A100

System Description


Host name: **hal-dgx.ncsa.illinois.edu**

Hardware

- [nvidia-dgx-a100-datasheet.pdf](#)
- [dgxa100-user-guide.pdf](#)

Software

- CentOS 7
- CUDA 11
- NCCL
- cuDNN
- TensorRT



Contact us

Request access to ISL resources:
Application

Contact ISL staff: [Email Address](#)

Visit: [NCSA](#), room 3050E

User Guide

To request an account: email kindrtnk@illinois.edu



Use `/home/<username>` for basic stuff only, do not put any code/data here as the `/home` partition is very small

Use `/raid` partition for all your data/code. Create a subfolder in this partition for your project and keep your stuff there.

To access the system, submit jobs, etc.: [Access hal-dgx and overdrive with hal-login3 Node](#)

To compile CUDA applications, first enable devtoolset 9

```
scl enable devtoolset-9 bash
```

To work with Python, create and enable virtual environment first

```
python3 -m venv python3-virtualenv  
source python3-virtualenv/bin/activate
```

You can use `pip` to install python packages within this environment.

To start a Jupyter notebook on hal-dgx

- on hal-dgx:

```
source python3-virtualenv/bin/activate  
pip install jupyter # this needs to be done only once  
jupyter notebook # this will start the jupyter server
```

- on your own computer:

```
# This opens a connection to the hal-dgx.ncsa.illinois.edu Jupyter server,  
and  
# forwards any connection to port 8888 on the local machine to port 8888  
on hal-dgx.ncsa.illinois.edu.  
ssh -L 8888:localhost:8888 hal-dgx.ncsa.illinois.edu
```

- Finally, on your own computer, open web browser and point it to the address you see after running 'jupyter notebook' on hal-dgx, something like `http://localhost:8888/?token=...`